

IAMM[®]

INNOVATION ACCESS METHOD

USER DOCUMENTATION

**INNOVATION DATA PROCESSING
IAM USER MANUAL**

VERSION V6.3

PURPOSE OF THE GUIDE This Users Guide outlines capabilities, describes the basic features of these capabilities and provides direction on their use. It was written primarily for technical personnel who have responsibilities related to the design and implementation of systems which will use IAM. This manual is intended to give those personnel a reference that will help them use IAM as a simple, efficient, reliable alternative to VSAM KSDS and ESDS processing in batch and online applications.

IAM®, FDR®, DSF®, CPK® and ABR® are proprietary program products of INNOVATION DATA PROCESSING, Innovation Plaza, 275 Paterson Avenue, Little Falls, New Jersey 07424-1658.

They are available for license exclusively from INNOVATION DATA PROCESSING, INC.

EUROPEAN OFFICES:

France

Innovation Data Processing S.A.R.L.
66 Avenue François Arago
92000 Nanterre
Tel: 01-47-69-15-15 Fax: 01-47-69-15-16
E-mail: 76543.1377@compuserve.com

Netherlands

Innovation Data Processing
Brouwerstraat 8
1315 BP Almere
Tel: 036-534 1660 Fax: 036-533 7308
E-mail: 76543.1376@compuserve.com

Germany

Innovation Data Processing
International Ltd.
Orleansstraße 4a
D-81669 München
Tel: 089-489 0210 Fax: 089-489 1355
E-mail: 76543.1401@compuserve.com

United Kingdom

Innovation Data Processing Ltd.
Clarendon House
125 Shenley Road
Borehamwood, Herts, WD6 1AG
Tel: 0208-905 1266 Fax: 0208-905 1428
E-mail: 76543.1374@compuserve.com

Foreign Representatives:

Computer Supervisory
Services
4 Brandon Street
Clovelly NSW
Australia 2031
Tel: 61-2-9665-1104
Fax: 61-2-9665-7382

Eckert Research International
Corporation
2-7-10 Nihonbashi Kayabacho
Chuo-ku, Tokyo 103
Japan
Tel: 03-668-8866
Fax: 03-668-8874

Stratos-Strategic Tools &
Services
Via Assarotti, 9
10122 Torino, Italy
Tel: 11-5620 354
Fax: 11-5621 291

© Copyright 1972, 1976, 1987, 1988, 1990, 1994, 1996

© INNOVATION DATA PROCESSING, INC.

Reproduction of this User Manual is prohibited except by licensed users for their internal use.

Main Phone Number (US)
201-890-7300

Fax Number (US)
201-890-7147

E-mail (US): 76322.2076@compuserve.com
<http://www.innovationdp.fdr.com>

THIS PAGE WAS INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

SECTION 00: INTRODUCTION TO MANUAL

00.01	ORGANIZATION OF THE MANUAL	7
00.02	SUMMARY OF MODIFICATIONS	9
00.03	CONTROL STATEMENT FORMAT	19
00.04	DOCUMENTATION NOTATION FORMAT	21
00.05	CONSIDERATIONS FOR PRIOR USERS	22
00.06	SAMPLE VENDOR PRODUCTS THAT HAVE CONVERTED VSAM FILES TO USE IAM	23
00.07	SAMPLE VENDOR PRODUCTS THAT WORK WITH IAM	24

OVERVIEW AND INTRODUCTION

SECTION 01: IAM VERSION 6.3 ENHANCEMENTS

01.00	IAM VERSION 6.3 ENHANCEMENTS	1
01.10	HIGHLIGHTS OF IAM VERSION 6.3 ENHANCEMENTS	1
01.20	ENHANCED FILE STRUCTURE	3
01.30	DYNAMIC FILE EXPANSION	8
01.40	ONLINE (CICS) ENHANCEMENTS	13
01.50	ADDITIONAL ENHANCEMENTS	20
01.60	CONVERSION CONSIDERATIONS	27

SECTION 02: IAM TECHNICAL DESCRIPTION

02.00	IAM TECHNICAL DESCRIPTION	47
02.20	IAM FILE ALLOCATION AND SPACE SAVINGS	48
02.30	DATA COMPRESSION	50
02.40	IAM AUTOMATIC RELEASE	51
02.50	IAM's REAL TIME TUNING	52
02.60	IAM's DYNAMIC TABLING 'DATA RECORDS-IN-VIRTUAL'	55
02.70	SMS SUPPORT IN IAM	56
02.80	IAM STORAGE REQUIREMENTS	59
02.90	ADDING RECORDS TO AN IAM FILE	60
02.100	IAM MISCELLANEOUS FUNCTIONS	61

SECTION 03: IAM DATA SET STRUCTURE

03.00	IAM DATA SET STRUCTURE	63
03.20	IAM FILE ALLOCATION	64
03.30	IAM COMPATIBLE FORMAT FILE AREAS	65

IDCAMS SUPPORT

SECTION 10: IAM AND ACCESS METHOD SERVICES (IDCAMS)

10.01	GETTING STARTED WITH IAM DEFINING IAM FILES	1
10.02	IDCAMS OVERVIEW	3
10.03	IDCAMS JCL REQUIREMENTS	5
10.04	IDCAMS DEFINE CLUSTER COMMAND	13
10.05	IDCAMS DEFINE EXAMPLES	16
10.06	IDCAMS DELETE COMMAND AND EXAMPLES	17
10.07	IDCAMS LISTCAT COMMAND AND EXAMPLES	17
	IAMPRINT REPORT	17
	IAMPRINT REPORT FIELDS	20
10.08	IDCAMS PRINT COMMAND AND EXAMPLES	23
10.09	IDCAMS REPRO COMMAND AND EXAMPLES	24
10.10	REORGANIZING AN IAM FILE	26
10.11	IMPORT AND EXPORT OF AN IAM FILE	27

IAM OVERRIDES

SECTION 11: IAM OVERRIDE STATEMENTS

11.01	IAM OVERRIDE STATEMENTS OVERVIEW	1
11.02	CREATE OVERRIDE STATEMENT – FORMAT	2
11.03	ACCESS OVERRIDE STATEMENT – FORMAT	9
11.04	IAM OVERRIDE STATEMENT EXAMPLES	15

SECTION 12: IAM RUN TIME STATISTICS REPORT AND SMF RECORDING

12.01	IAM RUN STATISTICS (INFO) OVERVIEW	17
12.02	REQUIREMENTS TO ACTIVATE IAM'S RUN TIME STATISTICS	18
12.03	EXAMPLE OF RUN TIME STATISTICS REPORTING	19
12.04	KEY INFORMATION IN RUN TIME STATISTICS (INFO) REPORT	21

SYSTEM ANALYSIS UTILITIES

SECTION 40: VSAM USAGE SMF ANALYSIS PROGRAM (IAMSMFVS)

40.01	VSAM USAGE SMF ANALYSIS OVERVIEW	1
40.02	JCL REQUIREMENTS	3
40.03	DEFAULT COMMAND	4
40.04	END COMMAND	6
40.05	EXTRACT COMMAND	7
40.06	HELP COMMAND	11
40.07	REPORT COMMAND	12
40.08	SAMPLE REPORTS	16
40.09	EXAMPLES	20

<u>SECTION 41:</u>	<u>SMF JOB RESOURCE USAGE ANALYSIS PROGRAM</u>	
41.01	JOB RESOURCE USAGE ANALYSIS OVERVIEW	1
41.02	JCL REQUIREMENTS	2
41.03	COPY COMMAND	3
41.04	DEFAULT COMMAND	5
41.05	END COMMAND	9
41.06	HELP COMMAND	10
41.07	PRINT COMMAND	11
41.08	QUERY COMMAND	13
41.09	REPORT COMMAND	15
41.10	IAMINFO COMMAND	17
41.11	EXAMPLES AND SAMPLE REPORTS.....	22
<u>SECTION 42:</u>	<u>SPACE SAVINGS ANALYSIS (IAMSIMVS)</u>	
42.01	SPACE SAVINGS ANALYSIS PROGRAM OVERVIEW.....	1
42.02	JCL REQUIREMENTS	3
42.03	REPORT COMMAND	4
42.04	SELECT COMMAND.....	5
42.05	SAMPLE REPORTS	7
42.06	EXAMPLES	8
<u>SECTION 43:</u>	<u>TRACE/DEBUG FACILITY</u>	
43.01	IAM TRACE/DEBUG FACILITY OVERVIEW	1
43.02	EXECUTION TIME JCL REQUIREMENTS.....	2
43.03	REPORT GENERATOR JCL REQUIREMENTS.....	3
43.04	TRACE REPORT COMMAND.....	4
43.05	TRACE/DEBUG EXAMPLES	7
<u>SECTION 44:</u>	<u>EVENT RECORDING MONITOR</u>	
44.01	IAM EVENT RECORDING MONITOR OVERVIEW	1
44.02	EXECUTION TIME JCL REQUIREMENTS.....	2
44.03	REPORT GENERATOR JCL REQUIREMENTS.....	3
44.04	MONITOR REPORT COMMAND	4
44.05	EXAMPLES	7
44.06	SAMPLE REPORTS	8
<u>SECTION 45:</u>	<u>FILE RECOVERY (IAMRECVR)</u>	
45.01	FILE RECOVERY OVERVIEW	1
45.02	JCL REQUIREMENTS	2
45.03	APPLY COMMAND	3
45.04	DEFAULT COMMAND	4
45.05	DIAGNOSE COMMAND	5
45.06	HELP COMMAND	6
45.07	LIST COMMAND	7
45.08	PRINT COMMAND	8
45.09	RECOVER COMMAND.....	10
45.10	JCL EXAMPLES	13
<u>SECTION 50:</u>	<u>IAM'S ISPF INTERFACE</u>	
50.01	IAM'S ISPF INTERFACE.....	1

ESDS SUPPORT

SECTION 60: IAM ESDS SUPPORT

60.10	IAM ESDS SUPPORT OVERVIEW	1
60.20	IAM ESDS FEATURES	2

SECTION 61: USING IAM ESDS FILES

61.10	DEFINING IAM ESDS FILES WITH IDCAMS	1
61.20	DELETING IAM ESDS FILES WITH IDCAMS	3
61.30	LISTCAT OF IAM ESDS FILES WITH IDCAMS	4
61.35	EXPLANATION OF LISTCAT OUTPUT	5
61.40	COPYING AND REORGANIZING IAM ESDS FILES WITH IDCAMS	7
61.50	JCL CONSIDERATIONS WITH IAM	9
61.60	IAM OVERRIDES FOR ESDS FILES	10

SECTION 62: SPECIAL CONSIDERATIONS

62.10	FREESPACE AND IAM ESDS FILE REORGANIZATIONS	1
62.20	STORAGE USE	1
62.30	VSAM COMPATIBILITY	2

SECTION 63: IAM ESDS FILE RUN STATISTICS (INFO) REPORT

63.10	EXAMPLES OF RUN TIME STATISTICS REPORTING	1
63.20	KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT	3

SECTION 64: IAM ESDS TRACE FACILITY

64.10	IAM 6.2 TRACE FACILITY FOR ESDS FILES	1
-------	---	---

MESSAGES AND CODES

SECTION 80: IAM V6.0 IAM MESSAGES AND CODES

80.01	IAM MESSAGES AND CODES INTRODUCTION	1
80.02	WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES	2
80.03	IAM SYSTEM SYSPRINT MESSAGES	16
80.10	IAM USER ABEND, ERROR AND COMPLETION CODES	39
80.20	IAM's VSAM INTERFACE (VIF) ERROR CODES ISSUED BY IAM:	50
80.21	IAM ERROR CODES PASSED IN THE RPL	51
80.22	IAM ERROR CODES PASSED IN THE ACB	52

INSTALLATION

SECTION 90: INSTALLATION INSTRUCTIONS IAM VERSION 6.3

90.01	IAM V6.3 INSTALLATION INSTRUCTIONS	1
90.01	IAM V6.3 INSTALLATION INSTRUCTIONS IAM VERSION 6.3	3
90.02	INSTALLING THE IAM INSTALLATION CONTROL LIBRARY	5
90.03	IAM PROGRAM LOAD LIBRARY	7
90.04	ACTIVATING THE IAM VSAM INTERFACE (VIF)	9
90.05	TESTING IAM	13
90.06	IAM INSTALLATION QUESTIONS	15
90.07	INSTALLING THE USER DOCUMENTATION	17
90.08	DFSMS SUPPORTED BY NEW FEATURES IN IAM V6.2	20
90.10	SPECIAL CONSIDERATIONS FOR CICS WITH COMPATIBLE FORMAT FILES ..	23
90.30	TECHNICAL OVERVIEW AND EXCEPTION CONSIDERATIONS	26
90.40	INSTALLING THE IAM ISPF DIALOG (IAM 6.3)	28
90.41	APF AUTHORIZATION FOR THE IAM ISPF DIALOG	29

SECTION 91: GLOBAL OPTION CHANGE FACILITY 155

91.01	GLOBAL OPTION CHANGE FACILITY - Overview and JCL	31
91.02	GLOBAL OPTION CHANGE FACILITY - Format and options	32
91.03	GLOBAL OPTION CHANGE FACILITY - ZAP command	33
91.04	IAMZAPOP JCL EXAMPLES	41

00.01 ORGANIZATION

00.01 ORGANIZATION

APPLICABILITY This IAM Users Guide applies to Version 6.3 Level 0.2 or higher of IAM.

FORMAT OF THE GUIDE The User's Guide is divided into sections. Within sections, searchwords appear near the left margins. Each page displays its section number in the upper outer corner. Page numbering is sequential within each section.

CHANGES IAM Version 6.3 is a major new release. Changes from a prior version are indicated by a single vertical line to the left of the change.

LIST OF SECTIONS This manual contains the following major sections.

[SECTION 00](#) – INTRODUCTION AND TABLE OF CONTENTS
[SECTION 01](#) – VERSION 6.3 ENHANCEMENTS
[SECTION 02](#) – IAM TECHNICAL DESCRIPTION
[SECTION 03](#) – IAM DATA SET STRUCTURE
[SECTION 10](#) – ACCESS METHOD SERVICES (IDCAMS)
[SECTION 11](#) – IAM OVERRIDE STATEMENTS
[SECTION 12](#) – IAM RUN TIME STATISTICS AND SMF RECORDING
[SECTION 40](#) – IAM VSAM USAGE ANALYSIS PROGRAM
[SECTION 41](#) – IAM SMF ANALYSIS PROGRAM
[SECTION 42](#) – IAM SPACE SAVINGS SIMULATION
[SECTION 43](#) – IAM TRACE/DEBUG
[SECTION 44](#) – IAM MONITOR
[SECTION 45](#) – IAM RECOVERY PROGRAM
[SECTION 50](#) – IAM ISPF INTERFACE
[SECTION 60](#) – ESDS SUPPORT
[SECTION 80](#) – MESSAGES AND CODES
[SECTION 90](#) – INSTALLATION INSTRUCTIONS
[SECTION 91](#) – IAM GLOBAL CHANGE FACILITY

IMPROVING THE GUIDE We have tried to make this manual as complete, precise and error free as possible. However, in spite of our best efforts, errors and vague explanations may have crept in. Should you encounter any of these, we would appreciate your corrective criticism. It is primarily through your feedback that we can improve this manual.

EXAMPLES All examples and Job Control Language statements shown in this manual are for ILLUSTRATIVE PURPOSES ONLY. The user is expected to modify them as required by his/her environment.

THIS PAGE WAS INTENTIONALLY LEFT BLANK

00.02 SUMMARY OF MODIFICATIONS: IAM VERSION 6.3

00.02 SUMMARY OF MODIFICATIONS: IAM VERSION 6.3

IAM Version 6.3 includes many new features and enhancements that are anticipated to make IAM even easier to use than before. A new, enhanced IAM file structure is now available for KSDS files, and is quite similar to the file structure introduced with IAM Version 6.2 for ESDS type of files. IAM files in the new enhanced format can only be used by IAM Version 6.3, and subsequent versions. Prior versions of IAM can not process the enhanced KSDS file structure. IAM Version 6.3 does still include the capability to process and create KSDS type of files in the original IAM file format, which will be referred to as the Compatible IAM file format.

Full details on all of the new features and capabilities are in the IAM Version 6.3 Enhancements Section. Below is a summary of the major features:

DYNAMIC FILE EXPANSION

IAM Version 6.3 has the ability to dynamically expand the size of an IAM file, as needed for file growth. This is commonly referred to as the ability to take secondary extents. This capability is only available with the IAM *Enhanced File Structure*. The IAM Enhanced File Structure eliminates the preformatted Independent Overflow, along with the size limitation of 64,000 blocks, and the Prime Extension area (PE). As a file expands in size, additional blocks are formatted in the allocated but unused areas of the file, up to one cylinder at a time. These blocks are referred to as Extended blocks, and will be assigned to either the Extended Overflow area, which replaces the Independent Overflow area, or as Extended PE blocks, which replaces the Prime Extension area. When all of the allocated space is used, additional DASD space is requested through the MVS EOVS service.

IAM has incorporated some additional new features that enhance the file expansion feature. The Automatic Space Release Feature has been enhanced to incorporate a *Space Reserve* function. Using the CA% Freespace, or the Overflow override, IAM will retain some of the DASD space allocated after the file load is completed for future expansion. IAM will set a target for the amount of space to be reserved by using 1/2 of the CA Freespace % applied to the amount of data loaded into the file. DASD space allocated beyond the target reserve will be released. The reserved quantity is intended to prevent a file from immediately taking secondary extents when the file expands in size. Note that additional DASD space will NOT be allocated to guarantee the size of the reserved area. IAM Automatic Space Release is still present to provide automatic space savings for over allocated files.

IAM Version 6.3 also has a *Dynamic Secondary Space Adjustment* feature, to react to circumstances when a file has been allocated with an insufficient quantity of space. After a file has acquired five extents of DASD space, IAM will automatically increase the secondary space quantity based on a factor from the IAM Global Options Table. This feature is present for both file loads, as well as file expansion. Coupled with that capability is the feature that will allow customers to specify whether IAM should attempt to obtain the primary space quantity, or the secondary space quantity, when space is needed on another DASD volume.

MVS/ESA SP5.2 AND DFSMS 1.3 SUPPORT

IAM Version 6.3 has been enhanced for both Compatible and Enhanced format files to run under MVS/ESA/SP 5.2 with DFSMS 1.3. This includes changes for UCB's above the line and changes in the VSAM control blocks and services that affect the SHOWCB/TESTCB/MODCB and the OPEN/CLOSE VSAM macros. Because of some significant changes in VSAM with DFSMS 1.3, IAM Version 6.3 is necessary with for installations with that level of VSAM. Using prior versions of IAM with DFSMS 1.3 may cause problems with COBOL programs, and others that utilize the above mentioned VSAM macros.

00.02 SUMMARY OF MODIFICATIONS: IAM VERSION 6.3

ENHANCEMENTS FOR ONLINE SYSTEMS (CICS)

With the IAM Enhanced File Structure, there are numerous other enhancements that are anticipated to provide *improved response time* for online (CICS) systems using IAM files. The maximum number of buffers (*MAXBUFNO*) that can be used is now 255. This will provide improved response time for files with extremely high file activity. IAM can now perform *concurrent logical I/O* requests, which eliminates unnecessary waits on requests doing physical I/O by requests that can be satisfied from data already in storage. IAM now has the ability to perform *concurrent physical I/O* requests, which will benefit large multivolume files. This is because there can be an I/O active to each volume that the IAM file resides on concurrently, rather than only one physical I/O active at any given time. The IAM internal processing for *asynchronous requests* has been significantly enhanced, which will provide for improved response times and alleviate problems encountered due to the CICS AMAXTASK parameter and with cancelling transactions.

IAM Version 6.3 enhanced file structure also brings *usability enhancements* to CICS processing. IAM files no longer have to be removed from the LSR pools and in fact due to CICS usage of the UPAD exit, there may be even more *CPU time savings* by leaving the IAM files in the LSR pools. IAM does not use the LSR buffers, so the size of the LSR buffer pools can be reduced as files are converted to IAM. IAM files if left in the LSR pools, can now be used by the *CICS VSAM Data Table* facility. IAM has moved most of the data areas and modules it requires into extended private storage, *reducing the need for OSCORE*. New with IAM 6.3 is an optional *IAM file information and statistics* transaction. If installed into your CICS system, this transaction will find all of the enhanced format IAM files that are open, and provide additional detailed information about the file(s) that you request. IAM Version 6.3 has a *Dynamic Region Adjustment* feature which will dynamically adjust the size of the extended private region. This will reduce or eliminate the need to change MVS IEFUSI exit as storage requirements increase.

IAM Version 6.3 provides a response for *BWO (Backup While Open)* calls issued by various CICS VSAM recovery products. This enhancement should eliminate the warning and error messages issued by those products when being used for IAM files. The *Journaling Exit (JRNAD)* support has been enhanced add support for the call issued prior to a buffer write, to allow proper synchronization with packages that provide file recovery.

ENHANCEMENTS FOR BATCH JOBS

There are a number of additional enhancements that are more related to batch jobs. IAM has *improved the sequential I/O* capability. While prior versions of IAM did provide a sequential read ahead capability, this has been enhanced in Version 6.3 to allow for processing of data while subsequent data blocks are still being read into virtual storage. This feature may reduce elapsed time for long running batch jobs that access the files sequentially. IAM can now write multiple blocks per I/O on mass sequential update jobs, which will also reduce elapsed time and EXCP's.

IAM now offers a *backup of compressed data* option. This capability is activated through the IAM overrides, by the keyword BACKUPCOMPRESSED. This will prevent the decompressing data as it is read, and the compressing data as it is being loaded into a file. This will help improve backup and file reorganization time. FDRREORG Version 5.2/50 will automatically make use of that feature, plus will provide a *Parallel Sequential Read* ability for multivolume files to further reduce the time to reorganize IAM files. The compressed data on the sequential media can be decompressed through the IAMRECV utility, if ever needed.

IAM Version 6.3 files can be used as the object of the IDCAMS *MODEL* and the JCL *LIKE* parameters. Included in this capability is the saving of additional file definition characteristics, including CA Freespace, primary and secondary space allocations, and the control interval size. The values of these fields are available in the IAM ListCat output. Either IAM or VSAM files can now be defined based on the attributes of an IAM file.

Included in Version 6.3 is optional support for the software products from BIM. This support is enabled through the use of the IAM Global Options Change Facility, and does require a level of BIM which provides this additional IAM support.

00.02.00 SUMMARY OF MODIFICATIONS AND IAM V6.2 ENHANCEMENTS

00.02.00 SUMMARY OF MODIFICATIONS AND IAM V6.2 ENHANCEMENTS

ESDS The major enhancement to IAM in V6.2 is support for VSAM ESDS type files. The benefits IAM offers ESDS files include an average file size reduction of 30 to 50%, elapsed time reduction of 30 to 50%, EXCP's reduced by 30% to 60. The support for ESDS files includes enhancements to the product that presently are only available for ESDS files. Included in these enhancements are support to dynamically take additional extents when records are being added to the file and more DASD space is required, enhancements to the Real Time Tuning, including the capability to utilize a larger number of buffers, improvements in sequential I/O processing to improve I/O and processing overlap, and enhancements for concurrent I/O request processing to improve response time for online systems. Additionally, IAM ESDS files can be left in the LSR buffer pools, although IAM will not utilize any of the LSR buffers. Please review the ESDS section of the manual for complete information on using IAM ESDS files.

GLOBAL OPTIONS Several new global options have been added for IAM V6.2. When installing the new version, please be sure to review the new options available, as well as reviewing what options may have been changed when the prior version of IAM was installed. The AUDIT function of IAMZAPOP will identify which options were changed from their defaults in V6.1, at level 08 or above. The new global options include the following:

ENABLE=NOREUSE will indicate that IAM is to honor the REUSE or NOREUSE specification from the IDCAMS DEFINE of the IAM file. The default is that IAM files are always reusable, and the NOREUSE is ignored.

MAXOVERFLOW= can be used to revise the largest size of INDEPENDENT OVERFLOW that will be calculated for IAM KSDS files based on CA% FREESPACE. The default is 50,000 records.

DYNCAT=YES can be used to indicate that on an IDCAMS DEFINE of an IAM file, that IAM is to let Dynamic Allocation catalog the dataset. Please note that this is automatically done for SMS managed data sets already. This option should be set for those customers that have the POOLDASD product from EMPACT.

DSORG=PS can be used to cause IAM files to be allocated with a DSORG=PS instead of DSORG=DA. This is primarily of use for installations using DFHSM to manage IAM files, when DFHSM might move IAM files between different device types.

WORKUNIT= Global option, in addition to being used for the default unit for the work file that may be allocated during a file load, will now also be used to specify the default unitname used for IAM ANYVOL allocations.

Due to various problems, the LOADJCLALLOC global option has been removed. For those customers which are using the JCL VSAM file definition capability with IAM, IAM is required to be in the link list.

ESDSINTEGRATED= specifies as a percentage, the amount of freespace to be left in each block of an ESDS file, to accommodate possible increases in the length of records due to updated records not being as compressible as they had been originally. The default is 0. Installations with heavily updated ESDS files may want to specify a nonzero value, in the range of 5 to 10%.

00.02.00 SUMMARY OF MODIFICATIONS AND IAM V6.2 ENHANCEMENTS

- IAMOVRID** There have been some changes to the overrides for IAM Version 6.2. These include the following:
- DYNCORE can now have values up to 16000K.
 - MAXBUFNO can now take values up to 255 in support of the IAM ESDS files. This does not change the 32 buffer maximum which remains unchanged for the KSDS file type.
 - OCOREX% can now have values in the range of 1 to 100. The smaller values may provide a reduction in CPU time for jobs that do massive inserts into the IAM Independent Overflow area.
 - OVERFLOW (O=) can now have values up to 3,000,000 records. Please note that the 64,000 block limitation on the size of Independent Overflow still remains.
 - PSEUDORBA is a new CREATE override keyword for ESDS file types. By specifying this override, IAM ESDS files may exceed the VSAM 4 gigabyte size limitation. However, the "RBA" values which are fed back to the application will be different than the official VSAM RBA, which may result in some application processing errors. Applications which have dependencies on the true VSAM RBA values will not be able to use this capability.
 - REREAD is a new keyword, which when specified will result in IAM rereading the overrides every time an IAM file is opened. This is primarily intended to aid online system users, such as CICS, so that they can revise their IAM overrides, and have them take effect whenever the appropriate IAM files are reopened. For this to work, the overrides need to be in a dataset, not a part of a sysin input stream. By default, IAM for any given job step, will only read the IAMOVRID file once, the first time an IAM file is opened, and retain the values in storage.
 - UPDATENQ=NONE will now also eliminate the ENQ issued to prevent a file from being loaded. This is intended to help jobs using the MVS CHKPT/RESTART services from receiving error messages due to the ENQ.
- MULTIVOL** The limitation of 20 volumes on defining an IAM file has been eliminated. For most defines, the limitation is now 255 volumes. Some defines that end up using a DADSM allocation are limited to 10 volumes.
- TRUE PE=0** To help conserve DASD space on small files, IAM will no longer automatically add a PE block to the file when PE=0 is specified as an override and the index to the file is smaller than the size required for a compressed index, which defaults to 8000 bytes. For compatibility with the prior IAM releases, there will always be at least 1 PE block when the file has a compressed index structure.
- TESTING VIF** The testing support for new levels of IAM through VIF (VSAM INTERFACE) has been enhanced in Version 6.2. Review [Section 90.01](#) for further information.
- MVS/ESA 5.1** IAM 6.2 requires no changes to operate with MVS/ESA 5.1 and DFSMS 1.2.
- BOOK-MANAGER FORMAT AVAILABLE** Your IAM installation tape now contains a SOFTCOPY version of the documentation that can be downloaded and used with BOOKMANAGER to view the manual online. The SOFTLOAD member in the ICL library will download the PDS from the installation tape containing the SOFTCOPY doc. The BLDBOOKS member in the ICL library can then be used to create the sequential files needed to use the books with BOOKMANAGER. Follow the directions in the BLDBOOKS member to change the data set names to your installations specifications, BLDBOOKS will also create a file that contains the bookshelf index. We have also included a bookshelf in the ICL library, member BKSHELF, that can be copied into a sequential file with the appropriate file attributes for BOOKMANAGER. The bookshelf must be DSORG=PS,RECFM=VB,LRECL=255,BLKSIZE=8000 and can be allocated as just 1 track.

00.02.01 SUMMARY OF MODIFICATIONS AND IAM V6.1 LEVEL 12 ENHANCEMENTS

00.02.01 SUMMARY OF MODIFICATIONS AND IAM V6.1 LEVEL 12 ENHANCEMENTS

Changes added to IAM V6.1 Level 12

- LISTCAT** The IAM LISTCAT IAMPRINT report has been enhanced to include the actual space allocated on each volume, and a total space allocated figure. Each volume to which the IAM data set is cataloged to will be listed, not just the first five volumes as had been done in prior versions.
- MVS/ESA** The IAM file load processor has been enhanced to utilize an MVS/ESA Data Space to hold the index structure while the file is being created, instead of utilizing a dynamically allocated temporary data set. In support of this capability, a new CREATE override and a new global option of DATASPACE indicates the size, in megabytes, of the Data Space to be requested. The default Data Space size requested is 128 megabytes, which should be sufficient for files up to 4 gigabytes in size. This feature can be disabled by specifying a Data Space size of zero, i.e., DATASPACE=0. Installation which previously had to change the work file allocation parameters in the Global Options Table to handle larger files should review the total amount of space requested by those values, including the maximum secondary allocations, and change the default Data Space size if necessary to accommodate a similar total amount of storage.
- IAMINFO** IAMINFO has been enhanced to include statistics on Data Space usage. Additionally, to aid in recovering an IAM file with IAMRECVR, the number of "Data Blocks" in the file has been added to the IAMINFO report. This figure is needed to recover the IAM file if either the first block and/or the first track of the IAM file has been damaged. This figure previously was only available from the LISTCAT IAMPRINT report with IAM Version 6.1
- IAMZAPOP** IAMZAPOP has been enhanced to provide a COPY and AUDIT function. Also the IAM panels were updated to include these new options in the IAMZAPOP utility panel. The COPY function will copy the changed options in the Global Options Table from your old IAM 6.1 level to the new IAM 6.1 level. This function will not work for copying prior levels of the options table to 6.1. The AUDIT function will produce a list of all the Global Options that have been changed, and their original distributed values.
- IAMISPF**
- Panel IAMU0200 provides the options for you to select the new functions for IAMZAPOP. If you do not want your users to have this capability in their panels you can rename IAMU0201 to IAMU0200 for the production panel library and the new options will be disabled via the panels.
- FIXES** Customers with custom zaps should contact INNOVATION for an updated zap(s).

Changes added to IAM Level V6.1 Level 9

- COPY OPTION** An extended Copy Option has been added to the Copy/Move Panel. If yes is specified, an additional panel will be displayed where the user can limit the copy to a maximum number of records, or request a specific range of records to be copied.

00.02.02 SUMMARY OF MODIFICATIONS AND IAM V6.1 LEVEL 08 ENHANCEMENTS

00.02.02 SUMMARY OF MODIFICATIONS AND IAM V6.1 LEVEL 08 ENHANCEMENTS

Changes added to IAM V6.1 Level 08

- VSAM SUBTASKING** VSAM subtasking no longer has to be turned off for MVS/XA or MVS/ESA when converting VSAM files to IAM in a CICS region.
- ISPF PANELS** ISPF panels can be used to perform many utility functions against IAM and VSAM files.
- Some of the key features included in the IAM ISPF panels are:
- Allocation of new IAM data sets with full Override support.
 - Allocation of new VSAM clusters.
 - Allocation with an IAM or VSAM model data set.
 - Ability to specify SMS classes through the ISPF panels.
 - Deletion of Data sets, Clusters, Paths, and Alternate Indexes.
 - Renaming of Data sets, Clusters, Paths, and Alternate Indexes.
 - Copy/Move support of IAM data sets, VSAM ESDS's, and VSAM KSDS's.
 - Full IAM data set information.
 - VSAM cluster information.
 - Interactive execution of selected IAM utility functions.
- SMS** A new global option, STORCLASS, has been added to enhance IAM SMS support. A customer can now specify a default SMS Storage Class for an IAM file, that will be used for the IAM file allocation when SMS is active and there was no Storage Class passed on the define request. This can occur when either the user did not specify a Storage Class and a Storage Class was not selected by the ACS routines, or when the ACS routine nullified the Storage Class, which can occur with the NONSMS storage class from the starter set. This option can be used to help identify a file as an IAM file in the ACS routines, or can be used to set a storage class of NONSMS for IAM file allocations to volumes that are not SMS managed. The default action without specifying this global option is that no Storage Class will be specified on the IAM allocation, unless a Storage Class is passed along on the DEFINE request.
- ERASE** The IDCAMS DELETE CLUSTER ERASE function is now supported for IAM files. Previously, ERASE had no effect for deletions of IAM files.
- MVS/ESA SP 4.2** MVS/ESA SP 4.2 support for dynamic I/O reconfiguration has been included in this level of IAM. IAM files can now be defined and accessed on devices dynamically added to the system.
- SMFVS** The IAMSMFVS REPORT and EXTRACT commands have been enhanced to include new parameters of FROMTIME and TOTIME to aid in record selection criteria. When specified with FROMDATE and/or TODATE, they indicate a precise starting and ending interval. When specified without FROMDATE or TODATE, they are used as time criteria for each day that is processed. Additionally, a keyword of CURRENT has been added to the REPORT command. When specified, this indicates that the most recent IAM overflow statistics are printed, instead of the maximum. This is useful for determining when an IAM file needs to be reorganized. For further information, run the IAMSMFVS HELP S(REPORT).
- LISTCAT** The IAM LISTCAT processing can now be turned off by specifying a DD card of IAMNOLIC to the JCL for jobs which do listcats. This is provided primarily for jobs which list off entire catalogs, where the IAM listcat processing was resulting in elongated run times. The IAMNOLIC can be a DD DUMMY.
- MULTI-VOLUME** For multi-volume non-SMS managed IAM file definitions, IAM will now verify that the file being defined does not have an extent on any of the volumes. If there are any extent(s), the define will now fail with an error code indicating duplicate data set.
- VAM SUPPORT** Support for customers using VAM has been enhanced to correct problems encountered with the latest release of VAM. The zaps previously issued for VAM are not required. VAM support must still be specified in the IAM Global Option Table. Also, the problem with multiple IAMW57 messages indicating not enough space on volume have been eliminated.
- ZAPS** The following custom zaps are no longer necessary – C-61.0001, C61.0002, C-61.0004, C-61.0041, C-61.0100, C-61.0102, C-61.0104, C-61.0107, C-61.0108, C-61.0109, C-61.0112, C-61.0114, C-61.0116. Customers with other custom zaps should contact INNOVATION for an updated zap(s). Problems with allocating space for a duplicate dataset, and then not scratching it because it is allocated to another job have also been corrected.

00.02.03 SUMMARY OF MODIFICATIONS AND IAM V6.0 LEVEL 02 ENHANCEMENTS

00.02.03 SUMMARY OF MODIFICATIONS

INSTALLATION CONSIDERATIONS FOR IAM 6.0 LEVEL 02 USERS:

The fixes on the tape affect both the base IAM access method code as well as the VSAM Interface, commonly referred to as VIF. While the VIF modules are compatible between the base level of IAM Version 6.0, and this updated level, to have the full benefit VIF must be started from the new level, and all jobs must use the new level code. This will generally require a re-IPL of all systems running VIF. While some initial testing can be done in a mixed environment, that is running the base VIF with the new library, full acceptance testing and all production work should be done with the new VIF active and the new library.

Any users who have installed custom zaps, (IDRDATA of C60.xxxx) please contact INNOVATION for information on the corresponding fix for this level.

ENHANCEMENTS TO IAM:

The following enhancements are on Version 6.0 Level 02:

1. The file load strategy for single record load operations has been enhanced. This enhancement should reduce the problems encountered with running out of space for software that initializes files with a single record, such as most of the McCormack & Dodge packages, and may also result in some performance improvements for that software. With the new enhancements, when only 1 record is loaded, instead of formatting up to 90% of the primary allocation completely as PE, some amount of that area will be used to build a larger independent overflow, which may eliminate the need to provide an override for independent overflow for single record load applications. Also, in most circumstances of single record loads, the record will be put into independent overflow, which for many applications will result in more effective use of the PE area, reducing the need for an extraordinarily large independent overflow area. The use of this automatic load strategy can be prevented by providing an override value for PE on either the DEFINE or initial load of the file.
2. Changes to the default block size calculation have been made for the larger record sizes to improve dasd space utilization.
3. IAMSMF has been enhanced to produce IAMINFO reports from the IAM SMF records. This feature provides a range of selection criteria, and the capability to merge the data into a single report by job or jobstep. For information on how to use this new facility, use the HELP S(IAMINFO) under IAMSMF. Note that the IAM SMF recording must be enabled by the use of IAMZAPOP, as described in the installation section of the IAM manual.
4. IAMSMFVS has also been enhanced to produce reports from the IAM SMF records. This includes an EXCP report, and a SIZE report, that are similar in content and appearance to the VSAM EXCP and SIZE reports. These reports will be produced automatically by the REPORT command of IAMSMFVS if the IAM SMF records are in the input SMF data, unless DSORG has been specified. To get just IAM data sets in the reports, specify REPORT DSORG=IAM. See the user documentation updates for definition of the fields for IAM files on the reports produced by IAMSMFVS.

00.02.03 SUMMARY OF MODIFICATIONS AND IAM V6.0 LEVEL 02 ENHANCEMENTS

5. Statistics about the IAM file processing are now available real time under VIF, through the use of a special VSAM GET macro. To obtain this data, the application program issues a VSAM GET macro, with RPL options set for random control interval processing, OPTCD=(DIR,CNV,MVE), with an area size of at least 512 bytes, and requesting an RBA of x 'FFFFFFF'. IAM will still return a logical error, with an RPL error code of x '68'.

The format of the data moved to the specified area is:

- 0-3 Length of IAM INFO area plus 4 bytes for the length area.
- 4-LEN IAM INFO area as documented by the IAMUSMF macro in the IAM installation control library. This area starts with the characters 'IDP1'

Following the above area is the following data:

- 0-3 Length of IAM STATS area plus 4 bytes for the length area.
- 4-LEN IAM STATS area as documented by the IAMUSMF macro in the IAM installation control library. This area starts with the characters 'IDP2'

00.02.04 SUMMARY OF MODIFICATIONS AND IAM V6.0 LEVEL 01 ENHANCEMENTS

00.02.04 SUMMARY OF MODIFICATIONS

MAJOR NEW RELEASE IAM Version 6.0 is a major new release and, with IAM as with other Innovation products, great care has been taken to protect the current users investments.

Applications presently executing with IAM Version 4.0 or 5.0, in most cases will run as is, under Version 6.0 without change.

SYSTEM LEVEL VSAM INTERFACE IAM V 6.0 contains an enhanced system level VSAM interface VIF. IAM files now can be used in place of single indexed KSDS VSAM clusters by any application no matter what language it is coded in. An application can be coded in Assembler, a higher level language (ex: COBOL, PL1) or one of the new 4GL languages. VIF allows an unaltered application program executing under MVS, MVS/XA, or MVS/ESA to access IAM files in place of single index KSDS VSAM clusters. With VIF, IAM's COBOL VSAM Interface and the SMP modifications for CICS are no longer needed to process IAM files.

IAM DATA COMPRESSION IAM can optionally compress the data in its files and it is completely transparent to the programs that create and use those files. IAM's proprietary Data Compression option typically gains an additional 20 to 50% reduction in file size over the 20% to 40% savings that results from simply converting to IAM.

Data Compression can be used on all your IAM files. If a particular file is found by IAM to be uncompressable, there is no penalty in CPU time to process that file after the load. It is as if compression had never been requested for that file.

ADDITIONAL NEW FEATURES There are many additional new features in Version 6.0.

- ..."Real Time Tuning" buffer management
- ...All storage for buffers above the 16MB line
- ...Automatic release of unused space

- ...Support for programs coded for AMODE(31) such as COBOL II
- ...Support for IDCAMS REPRO, PRINT, LISTCAT and RECATALOG
- ...Support for SMP/E CSI files
- ...Reductions in CPU processing and I/O
- ...Deferred rewrite of blocks in buffer pool
- ...Etc.

PRODUCT OVERVIEW [Sections 01 and 02](#) of this manual contain detailed information on the new features distributed with IAM V6.0. Reading these sections will familiarize new users with IAM's capabilities and will highlight IAM's new and expanded capabilities for existing users.

THIS PAGE WAS INTENTIONALLY LEFT BLANK

00.03.00 CONTROL STATEMENT FORMAT

00.03.00 CONTROL STATEMENT FORMAT

GENERAL The IAM control statements consist of 80-character logical records. The general format of these records is:

COLUMN 1				COLUMN 72	COLUMN 80

	COMMAND	OPERANDS	COMMENTS		

WHERE:

- Columns 1 to 71 – contain the command, operands, and comments fields, except when continued to subsequent logical records.
- Column 72 – must be blank if the command or operands extend to column 71.
- Columns 73 to 80 – not used by IAM. We suggest you use them as an identification or sequence field.

COMMAND FIELD The command field identifies the control statement and consists of a one or more character command word. It may appear anywhere within columns 1 to 71 and may only be preceded by blanks. The command word must appear in its entirety within columns 1 to 71; it may not be continued.

OPERAND FIELD The operand field, if present, follows the command field and is separated from it by at least one blank. The operand field consists of one or more keywords and/or positional parameters, separated by commas. It may not contain embedded blanks except within quoted strings. Operand fields may be continued onto subsequent logical records. If operands are to be used with a command, at least one operand must be on the logical record containing the command.

COMMENTS FIELD The comments field, if present, follows the operand field(s) and is separated by one or more blanks. It may contain any information deemed helpful by the person who codes the control statement. Comments fields may not be continued, i.e., they must end at or before column 71. Comments are not permitted on a control statement that allows operands but on which no operands have been specified.

CODING AN OPERAND FIELD An operand field consists of one or more positional or keyword parameters separated by commas.

POSITIONAL PARAMETERS Positional parameters must be coded in a specific order relative to one another. This means that the variable data you substitute for positional parameter 1 must precede the variable data for positional parameter 2, and so on. The absence of variable data to be substituted for a positional parameter is indicated by coding a comma in its place. However, you may omit the commas when the absent parameter is the last one, or if all following positional parameters are absent.

00.05.00 CONSIDERATIONS FOR PRIOR USERS

00.05.00 CONSIDERATIONS FOR PRIOR USERS

PROCEDURE FOR TESTING NEW IAM LEVELS

Beginning with IAM Version 6.2, it is possible to test different levels of the same version of IAM concurrently. This ability is required when different levels of the same version utilize the same VIF key. Please follow these instructions which will limit the test level of IAM to interact with only certain jobs for testing purposes.

Dynamically install the test version of IAM by executing a batch job using the following JCL:

```
//IAMSTRTT JOB 'DYNAMICALLY INSTALL TEST VERSION  
// OF IAM',IAM  
//IAMSTRTT EXEC PGM=IAMSTART,  
// PARM='TEST, JOBNAME=jobname  
// =jobgroup* See Note 1  
//  
//STEPLIB DD DSN=iam.test.library,DISP=SHR See Note 2  
//SYSLIB DD DSN=iam.test.library,DISP=SHR See Note 2  
//ABNLDUMP DD DUMMY FOR ABEND-AID  
//SYSUDUMP DD SYSOUT=A
```

This JCL is supplied in member IAMSTRTT in the Installation Control Library (loaded as described in [Section 90.02](#)).

Note 1: On the EXEC statement, JOBNAME= specifies the job(s) that are to be processed by the test version of the exits. JOBNAME= may be abbreviated as JOB=. The value specified may be either a single jobname, or a group name specified as a jobname prefix followed by an asterisk (*). TSO userids and started task names will also be checked for a match and will be eligible for processing by the test version of IAM. Any job that does not match will be processed by the production version of IAM. If JOBNAME=* or JOB=* is specified, then all jobs will be processed by the test version; in effect, the production version will be temporarily replaced.

Note 2: The SYSLIB DD statement, and, if used, the STEPLIB DD statement must specify the dsname of the test library, created as discussed above in Step A.

00.06.00 SAMPLE VENDOR PRODUCTS THAT HAVE CONVERTED VSAM FILES TO USE IAM

00.06.00 SAMPLE VENDOR PRODUCTS THAT HAVE CONVERTED VSAM FILES TO USE IAM

**This is a sample of some of the Vendors' products that Customers
have converted VSAM files to use IAM**

American Software – Materials Management

American Management System – CUFS, ACAPS

Arthur Andersen – DCS

Credit Card Software – Banking Applications

Computer Applications Services – Abend Catcher

Computer Associates – CA-7, C-11, Activator, UFO,
Intertest, Netman, Easytrieve+

Cyborg – Payroll System

Datalink Systems – Fast Teller System

Dun & Bradstreet Software – Accounts Payable,
General Ledger, Millennium Series Products

Dyatron – Employee Benefits

Erisco – Claim Facts

First Bankcard System – TBS, TCS

Global – CARMS

Group 1 – Code 1/PLUS, Mailsort

H&W Computer – Wizard, Mail, SYSM, SYSB2

Healthquest (HBO) – Patient Information File

Hogan Software – Demand Deposit

IBM – SMP/E, RMDS

ISS America – SICS

Impel Pacific – Power Plant Maintenance

Information Builders – Focus

Information System of America – Prism

Integral Systems – Human Resources/Payroll,
HRMS

Kirchman (Florida Software) – Installment Loan,
CIF, TESAREC, Dealer Floor Plan

LPC – Mailers, Finalist, Choice

Legent – Bundl

MacKinney Systems – Job and Syslog Facility

Mobius – INFOPAC

Newtrend – Infopoint

Policy Management Systems – APS

SAP – RF (Financial Accounting), RV (Sales and
Distribution), RM-PPS, RM-MAT

Shared Medical Systems (SMS) – Invision
(Active Patient Database file)

Shaw – Installment Loan, Commercial Loan

Sterling – NetMaster

Stockholders Systems – PEP+, CSS II

Systematics –

VIPS – Medical Part B

Vantage – Vantage 1 Annuities

Walker – General Ledger, CARMS, Materials
Management

Xycor – ADAM

00.07.00 SAMPLE VENDOR PRODUCTS THAT WORK WITH IAM

00.07.00 SAMPLE VENDOR PRODUCTS THAT WORK WITH IAM

IAM works with many system management software products that are used with VSAM today. Below is a list of products known to work with IAM. In many circumstances, these vendors have made enhancements to their products so that they can be used with IAM files. This is by no means a complete list of all the products that do work with IAM files.

JOURNALING AND RECOVERY:

FILE SAVE	(CA)
DRS	(BMC)
CICS/VR	(IBM)

SHARING PACKAGES:

SHARE OPTION 5	(CA)
SYSB	(H & W)
VSAM ACCESS/CICS	(BMC)

MISCELLANEOUS PRODUCTS:

FILE-AID	(COMPUWARE)
ABEND-AID	(COMPUWARE)
CICS	(IBM)
ISPF	(IBM)
NETWORK DATA MOVER	(STERLING)
(NDM)	
SHRINK	(STERLING)

DASD MANAGEMENT SOFTWARE:

FDR/ABR	(INNOVATION DATA PROCESSING)
FDRREPORT	(INNOVATION DATA PROCESSING)
FDRREORG	(INNOVATION DATA PROCESSING)

DF/SMS	(IBM)
DF/HSM	(IBM)
DF/DSS	(IBM)
VAM	(STERLING)
DMS/DS	(STERLING)
POOLDASD	(EMPACT)
STOPX37	(EMPACT)

SORT PRODUCTS:

SYNCSORT	(SYNCSORT)
DF/SORT	(IBM)
CA/SORT	(CA)

SECURITY PRODUCTS:

RACF	(IBM)
ACF/2	(CA)
TOPSECRET	(CA)

PERFORMANCE MONITORS:

OMEGAMON	(CANDLE)
THE MONITOR TMON	(LANDMARK)
STROBE	(PROGRAM ART)

PROGRAMMING LANGUAGES:

VS/COBOL	(IBM)
COBOL II	(IBM)
FORTRAN	(IBM)
PL/1	(IBM)
BAL	(IBM)
CA/OPTIMIZER	(CA)
SAS	(SAS INSTITUTE)

IAMM

INNOVATION ACCESS METHOD

INTRODUCTION & V6.3 ENHANCEMENTS

THIS PAGE WAS INTENTIONALLY LEFT BLANK

SECTION 01: IAM VERSION 6.3 ENHANCEMENTS

01.00 IAM VERSION 6.3 ENHANCEMENTS

01.00 IAM VERSION 6.3 ENHANCEMENTS

This section is intended to present to you, in a centralized location, the new features and enhancements of IAM Version 6.3. This section includes instructions and examples for using the new features. The IAM User Manual is in the process of being completely rewritten, and a replacement for this manual will be available by 3rd quarter, 1996.

01.10 HIGHLIGHTS OF IAM VERSION 6.3 ENHANCEMENTS

01.10 HIGHLIGHTS OF IAM VERSION 6.3 ENHANCEMENTS

IAM Version 6.3 includes many new features and enhancements that are anticipated to make IAM even easier to use than before. A new, enhanced IAM file structure is now available for KSDS files, and is quite similar to the file structure introduced with IAM Version 6.2 for ESDS type of files. Unless otherwise stated, the enhancements described in this section are only applicable to the enhanced file structure KSDS and ESDS type of IAM files. IAM files in the new enhanced format can only be used by IAM Version 6.3, and subsequent versions. Prior versions of IAM can not process the enhanced KSDS file structure. IAM Version 6.3 does still include the capability to process and create KSDS type of files in the original IAM file format, which will be referred to as the Compatible IAM file format.

While much has been changed with the new enhanced file format, there is much that remains the same. The intent with Version 6.3 was to address some of the limitations with IAM, without doing so much that IAM file integrity or reliability would be put at risk. This was done so that our customers can quickly and confidently move to the new file structure, and be positioned for future IAM product enhancements.

IAM Version 6.3 major enhancements include the ability to dynamically expand the size of an IAM file, as needed for file growth. This is commonly referred to as the ability to take secondary extents. With this capability, which is only available with the enhanced file structure, the preformatted Independent Overflow and PE (Prime Extension) areas no longer exist. As a file expands in size, additional blocks are formatted in the allocated but unused areas of the file, up to one cylinder at a time. These blocks are referred to as Extended blocks, and will be assigned to either the Extended Overflow area, which replaces the Independent Overflow area, or as Extended PE blocks, which replaces the old PE, on an as needed basis. When all of the allocated space is used, additional space is requested through the MVS EOVS service.

With the IAM Enhanced File Structure, there are numerous other enhancements that are anticipated to provide improved response time for online (CICS) systems using IAM files. IAM is no longer limited to a maximum of 32 buffers, which will provide relief for files with extremely high file activity. IAM also now has the ability to perform concurrent physical I/O requests, which will benefit large multivolume files. This is because there can be an I/O active to each volume that the IAM file resides on concurrently, rather than only one physical I/O active at any given time. The IAM internal processing for asynchronous type requests has been significantly enhanced, which will provide for improved response times and eliminate the problems encountered due to the CICS AMAXTASK parameter.

01.10 HIGHLIGHTS OF IAM VERSION 6.3 ENHANCEMENTS

IAM Version 6.3 enhanced file structure also brings some usability enhancements to CICS processing. IAM files no longer have to be removed from the LSR pools and in fact due to CICS usage of the UPAD exit, there may be even more CPU time savings by leaving the IAM files in the LSR pools. IAM does not use the LSR buffers, so the size of the LSR buffer pools can be reduced as files are converted to IAM. Also, because IAM files can now be left in the LSR pools, IAM files can now be used by the CICS VSAM Data Table facility. IAM has moved a lot of the data areas and modules it requires into 31-bit addressable storage, reducing the need for OSCORE. With IAM 6.3 is an optional IAM file information and statistics transaction. If installed into your CICS system, this transaction will find all of the enhanced format IAM files that are open, and provide additional detailed information about the file(s) that you request. IAM can also dynamically adjust the size of the extended private region, which will minimize or eliminate the need to change MVS IEFUSI exit as storage requirements increase.

There are a number of additional enhancements that are more related to batch jobs. IAM has improved the sequential I/O capability to provide an I/O read ahead capability, which may improve I/O and processor overlap. This feature may reduce elapsed time for long running batch jobs that access the files sequentially. IAM can now write multiple blocks per I/O on mass sequential update jobs, which will also reduce elapsed time and EXCP's. IAM now offers an option to not decompress the data in a file on backup, improving backup and file reorganization time. Also, with the FDRREORG product, a Parallel Sequential Read ability for multivolume files is available to further reduce the time to reorganize IAM files.

IAM Version 6.3 has been enhanced for both Compatible and Enhanced format files to run under MVS/ESA/SP 5.2 with DFSMS 1.3. This includes changes for UCB's above the line and changes in the VSAM control blocks that affect the SHOWCB/TESTCB/MODCB VSAM macros. IAM will also now provide a response for BWO (Backup While Open) calls issued by various CICS VSAM recovery products.

Details on all the above enhancements, plus other new capabilities will be presented in this guide. A summary of the new IAM Global Options and Overrides will also be provided, along with some examples and explanations of the revised IAMPRINT and IAMINFO reports for the enhanced IAM file structure.

01.20 ENHANCED FILE STRUCTURE

01.20 ENHANCED FILE STRUCTURE

FILE STRUCTURE OVERVIEW

The IAM Version 6.3 Enhanced File Structure eliminates the predefined and preformatted Independent Overflow area and Prime Extension (PE) area that existed in the original file structure. As will be seen below, the concept of an Overflow area and PE area still exist for dealing with the need to expand the file size. These areas are now dynamically acquired and formatted as needed. The size restriction of 64,000 blocks for the Overflow area has been eliminated. There is no longer any need to specify the size of the Overflow or PE areas of the file. If specified, the PE override will be ignored. The Overflow override will influence the amount of DASD space reserved, as discussed later in DASD Space Reserve and Space Release, on page 8. Additionally, the method of handling the single record load followed by mass insertion of records has been improved, eliminating the need to over allocate the file.

All records contained within the IAM Enhanced File Structure are treated internally as variable length records. This eliminates the usability problem encountered when VSAM definitions indicated that records were fixed in length, as indicated by the average and maximum record lengths being equal. Also, if specified or defaulted to through the IAM Global Options, the records are stored in a proprietary data compressed format.

IAM files are self-defined, non-VSAM files. This means that the information about the file structure is contained within the file itself, rather than in the system catalog and VVDS, as VSAM files are. The information kept within these system control areas, the system catalog, the VTOC, and the VVDS for DFSMS managed volumes is the same as for any other non-VSAM file. IAM files are treated as non-VSAM files by most DASD management products, including the FDR family of products.

An IAM file, as it appears on DASD, consists of a set of unique areas. Each area will be explained in further detail below. The basic structure is:

1. File Definition and Structure Data (Two blocks)
2. Prime Data Area
3. Index to Prime Data Area
4. Extended File Area
 - Extended area index and contents description
 - Extended Overflow and/or Extended PE blocks

FILE DEFINITION AND STRUCTURE DATA

The first block is initialized when an IAM file is defined with the basic file definition information, such as record length, key size, key offset, and so forth. After a file has been successfully loaded, the last I/O done to the file is to store the file structure data in the first block. Included is an indication that the file has been successfully loaded, and of course information about the where the index begins, size of index, plus additional required data. The only updates done to this block after a file load are for statistical information that is written during file close processing, including information on number of inserts, deletes, and updates. The bulk of the information that is presented on the IAM Listcat report is taken from this control block. The statistical information, and information about the extended file area that are kept in the first block may not be accurate, and do not have to be accurate as they are for informational purposes only. A listcat that is done while the file is open to any application program will not reflect the exact status of the file, in much the same way as a listcat on an open VSAM file will also not reflect the exact status and statistics for the file. The statistical information can become inaccurate if a system or address space fails while the file is opened for update processing, and has not been successfully closed.

The second block contains information about the extended file area, including maximum size, and the location, by relative block, of where the information about the extended area is stored. This block is typically read when during open processing, and is updated immediately after a new DASD extent has been acquired, and during close processing. If the file has not been successfully closed, then open processing will detect any inconsistencies, and update this data with the correct information.

01.20 ENHANCED FILE STRUCTURE

PRIME DATA AREA The Prime Data Area is built as the file is being loaded. This area contains the data records that were passed to IAM during the file load process. There is included some imbedded freespace within each block, called the Integrated Overflow area. This area is similar in concept to the VSAM CI freespace, and the size is indicated by the CI% Freespace parameter on the IDCAMS file DEFINE. Every loaded IAM file has a Prime Data Area, except for files that have been loaded with a single record.

INDEX TO PRIME DATA AREA This index is built during close processing for a file load. This index consists of the high key in each prime data block, and may be in a compressed format. Once a file is loaded, this index structure is never changed, until the file is reorganized or reloaded.

EXTENDED FILE AREA The Extended File Area consists of the data that is added to the file, either as inserted records or from updates that increased the size of data records, which could not fit into the Prime Data Area of the file. This area consists of blocks containing control information about the Extended file area, and extended data blocks, which can be either Extended Overflow, or Prime Extension (PE). The control information for the Extended File area is based on the internal logical structure of the file, and is not necessarily tied to actual DASD extents. This way, IAM files can have their extents merged and eliminated by products such as FDR/CPK without impacting the integrity of the IAM file structure or the data it contains.

The Extended File Area is acquired and formatted as needed. When there is a need to acquire an additional data block, for either Overflow or PE, a segment of the allocated and unused space is formatted. Normally, for batch processing, up to one cylinder will be formatted with empty blocks. For online processing, normally only one track will be preformatted at a time, to reduce the impact on response time. Blocks are then assigned as needed to either Extended Overflow, or PE. Once all of the allocated space is used, additional DASD space will be requested through the normal MVS EOVS service.

If an error condition occurs during the EOVS processing, such as an X37 abend condition, it is captured by the IAM DCB ABEND exit to prevent the job from actually abending. The request requiring the additional DASD space is failed, with a file full logical error. The avoidance of the abend is done to be compatible with VSAM, which will not abend either. The user will see the IBM error messages relating to the error condition encountered.

When a new extent is acquired, additional Extended Area control blocks are formatted and written as necessary, and up to one track of empty blocks will also be formatted. The control information for the Extended File area consists of identifying the blocks that are assigned to Extended Overflow, and the blocks that are assigned to PE. Additionally, the high key for each assigned PE block is retained as the index for the PE area.

EXTENDED PE As records are being added to the logical end of the file, PE blocks are assigned from the Extended File Area. Once a PE block is considered full, then the high key in that block is used as the index entry. Just as with the Prime Data blocks, Integrated Freespace, or CI% Freespace, is left in each block. This will allow for records increasing in size, as well as for later record insertions. Once a block is assigned as PE, it will remain as a PE block until the file is reorganized or reloaded. A PE block is only be able to hold records that fall into the established index key range. If records are deleted from the PE block, the freespace is available for expansion of existing records within that block, or for new records added to the file within the established key range.

01.20 ENHANCED FILE STRUCTURE

EXTENDED OVERFLOW

The IAM Extended Overflow area is a record based overflow area, similar to the Independent Overflow area of the Compatible File Format. Extended Overflow blocks are used to handle records that are being inserted within the file when there is insufficient space within the Prime Data or Prime Extension block that the record would have been assigned to, based on the established index. Extended Overflow space will also be used when an updated data record increases in size, and there is insufficient room within the Prime Data or Prime Extension block for the larger record. As with prior versions of IAM, the philosophy that once a record has been placed in a particular Overflow block it will not be moved to a different Overflow block is retained. The format of the data in the Extended Overflow block is identical to the data in the prime blocks, only the number of records actually stored in the block is presently limited by the number of maximum size records that will fit. This philosophy has proven itself to provide a very high level of file integrity, even though some DASD space may be underutilized for Overflow. When a record is deleted from Extended Overflow, the space it occupied is immediately available for reuse by ANY inserted or updated record, regardless of the key value. This eliminates the unusable lost space condition that can occur within VSAM files taking CI/CA splits, as records are deleted from certain key ranges, and new records are added in different key ranges.

The index to the Extended Overflow area is record based, that is each record in Extended Overflow has an entry in the index, consisting of the key and the block number of the Extended Overflow block containing the record. The index is built when the file is opened, by reading all of the used Extended Overflow blocks, as indicated by the control information. The limitation of 64,000 blocks in Overflow has been eliminated for the Extended Overflow area by increasing the size of the block identifier in the Overflow index from two to four bytes.

IAM FILE INTEGRITY

IAM files have been designed to provide the highest level of both file and data integrity. Some of the design factors that enhance file integrity include a non-updatable prime index structure, minimizing as much as possible the need to update other control information about the file structure, designing a structure of adding to the existing structure, as opposed to modifying the existing file structure. The dynamically extended portion of the file structure is validated during open, with efforts made to recover from most anticipated type of structure errors that may occur. Every effort has been made to provide an expandable, non-destructible file structure that provides an outstanding level of performance and responsiveness.

As with all prior versions of IAM, the access method will never pass back to an application program the address of an actual buffer. This is done to prevent inadvertent corruption of the data contained within the buffers, as much as possible. Even for requests indicating OPTCD=LOC (locate) mode processing, IAM will always make a copy of the record in another storage area, rather than pass back the address of the data record within the buffer. Also, with the Version 6.3 Enhanced File Format code, all of the critical data areas are acquired within their own 4K pages, to prevent the inadvertent storage overlays that sometimes occur when there are multiple programs using the same page of virtual storage.

In spite of all these efforts, should a problem occur within the IAM file structure, IAM provides a recovery program, IAMRECVR, which will recover all the data that it is possible to recover for the file. This program reads the data blocks on it's own, without going through the index. This program can be used for recovering from physical I/O errors, or structural errors that may have occurred.

01.20 ENHANCED FILE STRUCTURE

SPECIFYING FILE STRUCTURE TYPE WITH OVERRIDES

For KSDS type of files, the user has the option to specify either the Enhanced file structure, or the Compatible file structure. File structure is determined when the file is defined, either with the IAM CREATE override, or defaulted based on the IAM Global Options Table. IAM Version 6.3 is being shipped with the IAM Global Options Table indicating a default of the Compatible File Structure. Existing files that are reorganized without being deleted and redefined will retain their previously existing type of file structure. Files that are deleted and defined will default to the file structure specified in the IAM Global Option table. It is recommended that initially customers use the IAM Create Override of ENHANCED to define Enhanced Format Files. This will prevent files from being inadvertently converted to the new file format. Examples of using the override are presented below.

ESDS type of IAM files are always in the Enhanced Format when defined under Version 6.2 or Version 6.3 of IAM. The ESDS file formats are completely compatible between the two versions, and can be used interchangeably.

The file structure can be specified on the IAM CREATE Override control statements using one of the following keywords:

1. **ENHANCED** For Enhanced format files, or
2. **COMPATIBLE** For a format compatible with prior versions of IAM.

While the file structure type override can be specified on the file load, it is highly recommended that it be specified for the file definition. Some decisions are made during file definition that may result in certain capabilities not being fully functional. For example, if a file is defined with the space allocation specified in records defaulting to the Compatible format, then if changed to Enhanced format on the file load, the Dynamic Secondary Space Adjustment feature may not function, due to how the space allocation is saved.

The following example demonstrates how to define an IAM file with the new Enhanced File Format being specified on an IAM Override. Due to the use of the DD=&ALLDD, all IAM files defined in this job step will use the Enhanced File Format.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD     SYSOUT=*
//IAMOVRID   DD     *
              CREATE DD=&ALLDD, ENHANCED
/*
//SYSIN      DD     *
              DEFINE CLUSTER ( . . . . . )
/*
```

Figure 1: Defining an Enhanced Format IAM File

In case your Global Option default has been changed, it is still possible to define and create a Compatible Format file through the use of an IAM Override. Files created with the COMPATIBLE file format can be used with the prior versions of IAM without any concerns or special considerations. The benefits of the Enhanced File Structure will not be available to such files.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD     SYSOUT=*
//IAMOVRID   DD     *
              CREATE DD=&ALLDD, COMPATIBLE
/*
//SYSIN      DD     *
              DEFINE CLUSTER ( . . . . . )
/*
```

Figure 2: Defining a Compatible Format File

01.20 ENHANCED FILE STRUCTURE

CHANGING THE DEFAULT FILE STRUCTURE

The Global Options can also be easily changed, if so desired. The following example demonstrates how to change the IAM file structure default to the Enhanced File Format.

```
//ZAPOP      EXEC  PGM=IAMZAPOP
//STEPLIB    DD    DSN=my.i am. loadlib,DISP=SHR
//SYSLIB     DD    DSN=my.i am. loadlib,DISP=SHR
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
              ZAP  ENABLE=ENHANCED
              PRINT
/*
```

Figure 3: Setting Enhanced File Format as the IAM Default

The example below demonstrates how to change the default file format back to Compatible:

```
//ZAPOP      EXEC  PGM=IAMZAPOP
//STEPLIB    DD    DSN=my.i am. loadlib,DISP=SHR
//SYSLIB     DD    DSN=my.i am. loadlib,DISP=SHR
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
              ZAP  DISABLE=ENHANCED
              PRINT
/*
```

Figure 4: Setting the Default File Format to Compatible

01.30 DYNAMIC FILE EXPANSION

01.30 DYNAMIC FILE EXPANSION

With the IAM Version 6.3 Enhanced File Structure, IAM files can now acquire additional DASD extents when the file is being updated, in addition to the prior capability to take extents during file loads. Because IAM files are stored on DASD as a non-VSAM file type, IAM files can only take 16 extents per volume, in contrast to the 123 extents that VSAM can take. IAM files are limited to a total of 255 extents, across multiple volumes, due to the structure of the MVS DEB. Because of these limitations imposed by the operating system, IAM has revised how the automatic space release function works. For Enhanced Format files IAM will dynamically modify the secondary space quantity in response to unanticipated file growth. These new capabilities are described below.

DASD SPACE RESERVE AND SPACE RELEASE

With IAM Version 6.3, the automatic space release capability has been enhanced for the new Enhanced Format files. IAM will still automatically release DASD space on the first load of a file, providing that there is a non-zero secondary allocation value specified, and the RELEASE=NO IAM override or Global Option have not been set. Because of the elimination of the preformatted and preallocated overflow areas, rather than eliminate all of the unused overallocated space, an automatic space reserve feature was developed. With this feature, some of the overallocated space will be retained rather than released at the end of the first load to allow for file growth without having to immediately go into secondary extents. The space reserved will be contained within the allocated extents after the file has been loaded. No additional extents will be obtained for the space reserve.

The desired amount of space to reserve will be calculated by using either the value specified for Overflow records on a CREATE IAM Override card, or by using the CA% Freespace value specified on the Define of the file. The calculation when CA% Freespace is used is similar to the calculation of the size of Overflow for the Compatible file format. Basically, the CA% Freespace value is cut in half, and then that percentage of the total DASD space required for the data set as loaded will be set as the desired reserve quantity. For example, consider a file where CA% Freespace is defined as 20%, and the file ends up using 300 tracks of space. Then, using half of the CA% Freespace, or 10% of the 300 tracks, which is 30 tracks, will be set as the desired reserve quantity. When the Overflow override is used, IAM calculates the amount of DASD space required to hold the specified number of Overflow records, and that is set as the desired reserve quantity.

If CA% Freespace is 0 and there is no Overflow override, then no space reserve will be done, and all excess allocated space will be released. Also, if the desired space reserve quantity is less than one cylinder, then no space reserve will be done, because most files are cylinder allocated, and the space release is done at a cylinder boundary.

When the program loading the file issues the close, IAM looks at how much DASD space is allocated to the data set, but is not yet used. If the allocated but unused space is less than the desired reserve quantity, then no space will be released. However, if the allocated but unused space is greater than the desired reserve quantity, then space will be released down to the desired reserve quantity. In all cases, the end of file is set at the end of the used area, and is indicated so in the VTOC LSTAR field. This will allow a DASD space management utility, such as Innovation's FDR/CPK product, to release the unused space that was reserved if it has not yet been used.

01.30 DYNAMIC FILE EXPANSION

DYNAMIC SECONDARY SPACE ADJUSTMENT

To compensate for IAM being able to acquire only sixteen extents per volume, with the new Enhanced File Format IAM will dynamically adjust the secondary space quantity. This algorithm will take effect once a file has used five extents on a volume. At that point in time, IAM will increase the secondary space allocation by a multiplication factor specified in the Global Options Table, or from IAM Overrides. The default secondary space multiplication factor as IAM is shipped is ten for file loads, and five for file updates. The Secondary Space Adjustment feature is subject to the following rules:

1. The secondary space quantity will not be increased to value that exceeds the original primary space quantity.
2. The secondary space quantity will not be increased to a value that exceeds the size of the largest contiguous available extent on the volume. Note this is only effective once the file has obtained five extents.
3. If the original secondary space quantity is higher than the original primary space quantity, the secondary space quantity will not be adjusted.
4. For single volume files, the original secondary may be decreased from the original value to the largest extent available on the volume, just to try to keep it running as long as possible.

For example, if a file is Defined with a primary quantity of 500 cylinders, and a secondary of 10 cylinders, after five extents have been acquired, the secondary during a file load will be adjusted up to 100 cylinders. During a file load, the maximum space that will be used for this file is:

- 1 extent of 500 cylinders
- 4 extents of 10 cylinders each
- 11 extents of 100 cylinders

This will bring the total maximum space up to 1,640 cylinders, which is slightly less than the VSAM maximum of 1,720 cylinders. Although IAM will use less DASD space than VSAM, the effective size of the IAM file will generally be larger in terms of actual data storage capability than VSAM. This is due to IAM's efficient use of DASD devices. So, by providing the Dynamic Secondary Space Adjustment feature, IAM files have the potential to grow in size as large as VSAM files will, although it will be done in fewer extents.

For this same allocation, the secondary will be adjusted up to 50 cylinders during a file update run. However, if the primary space is 20 cylinders, and the secondary is 10 cylinders, then the maximum value that will be used for the secondary is limited to 20 cylinders. The maximum secondary space quantity that will be requested for any particular file is included on the IAMPRINT LISTCAT output.

01.30 DYNAMIC FILE EXPANSION

DYNAMIC SECONDARY SPACE ADJUSTMENT OVERRIDES

The secondary extent multiplication factor can be changed through the use of the IAM Overrides on a file by file basis, using the MAXSECONDARY keyword. Regardless of the value set, the basic rules for modifying the secondary space quantity remain as explained above. The factor can be specified on the IAM CREATE override during the file definition or file load, and on file updates with an IAM ACCESS override. The values permitted are from zero to ten. Values of zero or one will prevent IAM from increasing secondary allocation. The secondary allocation value may be reduced for single volume files when there is insufficient space for the secondary. When the secondary factor is specified on a CREATE override during file definition, the value is stored with the file. That value will be used for subsequent file loading and file updates, unless overridden at run time. Following are some examples of using the IAM overrides to modify the secondary space factor.

The following JCL example demonstrates how to set the secondary space adjustment factor to eight. This override will be effective for all IAM files defined in this job step. The secondary factor will be saved with the file control information, and be used both during file load and during file updates, unless those job steps are overridden with a different value.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//IAMOVRID   DD    *
           CREATE DD=&ALLDD,MAXSECONDARY=8
/*
//SYSIN      DD    *
           DEFINE CLUSTER .....
/*
```

Figure 5: Specifying a Secondary Space Factor on File Definition

This next JCL example demonstrates how to override the secondary space adjustment factor on a file load. In this example, an IDCAMS REPRO is being used to load the file. However any application program that does a valid file load would use a similar override.

```
//LOAD       EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//INFILE     DD    DSN=backup.file,DISP=OLD
//IAMFILE    DD    DSN=my.i.am.file,DISP=OLD
//IAMOVRID   DD    *
           CREATE DD=IAMFILE,MAXSECONDARY=5
/*
//SYSIN      DD    *
           REPRO INFILE(INFILE) OUTFILE(IAMFILE) REUSE
/*
```

Figure 6: Specifying secondary space factor on file load

The following example demonstrates how to override the secondary space factor during a file update run, using an IAM ACCESS override:

```
//UPDATE     EXEC  PGM=mysgm
//IAMFILE    DD    DSN=my.i.am.file,DISP=OLD
//IAMOVRID   DD    *
           ACCESS DD=IAMFILE,MAXSECONDARY=10
/*
```

Figure 7: Specifying a Secondary Space Factor on File Update

01.30 DYNAMIC FILE EXPANSION

CHANGING DEFAULTS FOR DYNAMIC SECONDARY SPACE ADJUSTMENTS

Changing the default values for the secondary space adjustment factor can be easily accomplished by changing the IAM Global Options table with the IAMZAPOPOP program. The keyword relating to the secondary space factor is MAXSECONDARY=(x,y) where x specifies the value used during file loads, and y specifies the value used during file updates. As the product is shipped, the defaults are a factor of ten for file loads, and a factor of five for file updates (10,5). The difference is because generally speaking more data is expected on the file load than during most file updates. However, IAM provides the flexibility for the customer to set the values as necessary to reflect processing needs at their installation. Note that values of zero or one will effectively disable the dynamic space adjustment feature. Below is an example of changing the secondary space adjustment default values to be the same for file loads and file updates:

```
//ZAPOPOP      EXEC  PGM=IAMZAPOPOP
//STEPLIB      DD    DSN=my.i.am.loadlib,DISP=SHR
//SYSLIB       DD    DSN=my.i.am.loadlib,DISP=SHR
//SYSPRINT     DD    SYSOUT=*
//SYSIN        DD    *
                ZAP  MAXSECONDARY=(8,8)
                PRINT
/*
```

Figure 8: Setting Dynamic Space Adjustment Factor Defaults

MULTI- VOLUME CONSID- ERATIONS

With the IAM Version 6.3 Dynamic Secondary Space Adjustment feature, there is an additional option for multivolume files that are defined without guaranteed space. When it appears to IAM that the EOVS request will result in the next volume being allocated, IAM will request the primary space value, rather than the normal secondary quantity that is used for non-VSAM files. This feature can be controlled by either IAM Overrides, or by the IAM Global Options table. This feature is provided to offer an alternative so that IAM space allocations will be similar to VSAM. The new IAM Override keyword is:

1. MULTIVOLUME=PRIMARY or
2. MULTIVOLUME=SECONDARY

These new IAM Override keywords can be used on either the CREATE or ACCESS Override statements. When specified on the CREATE keyword during file definition, the value specified is saved with the file control information and will remain the applicable option, unless overridden by a particular job step.

The rules for IAM files defined on DFSMS managed volumes with the Guaranteed Space attribute are different. This is because DFSMS will automatically allocate the primary space quantity on each volume when the file is defined. The secondary allocation quantity will be modified as described above for single volume files.

Another special multivolume circumstance is when a file is defined with no secondary space specified. For most circumstances, IAM will set the secondary allocation value to the primary value. For data sets on DFSMS managed volumes with Guaranteed Space, the secondary is left as zero, with the primary being allocated on each volume when the file is defined. For systems where DFSMS is active, but the data sets are on non-SMS managed volumes, IAM will leave the secondary as zero, and allocate the primary space on each volume when the file is defined. This is to mimic the DFSMS Guaranteed Space, and will prevent any secondary extents from being taken.

01.30 DYNAMIC FILE EXPANSION

EXAMPLES OF OVERRIDING MULTI- VOLUME ALLOCATIONS

The JCL example below demonstrates how to define an IAM file indicating that IAM is to attempt to obtain the primary space allocation whenever IAM anticipates that MVS will switch to the next volume.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//IAMOVRID   DD    *
              CREATE DD=&ALLDD,MULTIVOLUME=PRIMARY
/*
//SYSIN      DD    *
              DEFINE CLUSTER (.....
                          OWNER($IAM) -
                          VOLUMES(  vol1001, vol1002, vol1003 ) -
                          ..... )
/*
```

Figure 9: Defining a File Indicating to obtain Primary Space on Multiple Volumes

The example below demonstrates the opposite, that is IAM will request the secondary space when IAM anticipates that MVS will switch to the next volume. This causes a more typical non-VSAM type of secondary allocation process, and matches what is done with the Compatible Format files.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//IAMOVRID   DD    *
              CREATE DD=&ALLDD,MULTIVOLUME=SECONDARY
/*
//SYSIN      DD    *
              DEFINE CLUSTER (.....
                          OWNER($IAM) -
                          VOLUMES(  vol1001, vol1002, vol1003 ) -
                          ..... )
/*
```

Figure 10: Using Secondary Allocation for First Extent on Multiple Volume Files

CHANGING IAM DEFAULT FOR MULTI- VOLUME ALLOCATION

IAM is shipped with a default value of MULTIVOLUME=PRIMARY, indicating that IAM will attempt to acquire the primary space quantity as the first extent on each volume. The default value, contained in the IAM Global Options Table, can be easily modified, as demonstrated by the example shown below.

```
//ZAPOPOP    EXEC  PGM=IAMZAPOPOP
//STEPLIB   DD    DSN=my.i am.loadlib,DISP=SHR
//SYSLIB    DD    DSN=my.i am.loadlib,DISP=SHR
//SYSPRINT  DD    SYSOUT=*
//SYSIN     DD    *
              ZAP  MULTIVOLUME=SECONDARY
              PRINT
/*
```

Figure 11: Changing IAM Default for Multiple Volume Allocations

01.40 ONLINE (CICS) ENHANCEMENTS

01.40 ONLINE (CICS) ENHANCEMENTS

HIGHLIGHTS OF ENHANCE- MENTS FOR CICS / ONLINE SYSTEMS

IAM Version 6.3, with the Enhanced File Structure, brings numerous enhancements to improve usability and response time for online systems. The Real Time Tuning routines have been enhanced with the ability to manage more buffers, and also be sensitive to the frequency of I/O requests when adjusting buffers. IAM has improved the handling of asynchronous requests, which will help response times in addition to minimizing the problems encountered with the CICS AMAXTASK parameter. IAM now has the capability to perform concurrent physical I/O, which may help performance on heavily accessed multivolume data sets.

On the usability side, IAM Version 6.3 has reduced the storage requirement for 24-bit addressable storage, by locating more control blocks and code to extended private storage. IAM Enhanced Format files can be left in the CICS LSR Buffer Pools, to ease conversion and testing of IAM. Because of this enhancement, IAM files are now candidates to be used by the CICS Data Table Facility, while IAM still offers it's own version of data in virtual with the IAM Dyncore Feature. The new IAM Dynamic Region Adjustment feature will aid customers with extended private storage management, by not requiring an MVS exit to adjust the extended private storage region.

IAM REAL TIME TUNING ENHANCE- MENTS

With IAM Version 6.0, Innovation introduced to IAM a feature called Real Time Tuning. This feature provided an algorithm for adjusting the number of buffers being used for any particular file based on the I/O demands and buffering benefits. This feature has provided outstanding performance capabilities for many installations. Based on the experience and needs of our customers, Version 6.3 has enhanced the Real Time Tuning concept, to provide an even higher level of performance and responsiveness for online systems. These enhancements include the following:

1. Increased the maximum buffers for any single file from 32 to 255.
2. Added the capability to have multiple concurrent physical I/O active.
3. Consider frequency of I/O requests for buffer adjustments.

Each of these enhancements are further described in the following sections.

Increase Maximum Buffers

While experience has shown that a maximum of 32 buffers is more than sufficient for most files accessed online, there are some exceptions. To improve the responsiveness for these heavily accessed files, the buffer management has been redesigned for the Enhanced File Format to allow for a larger number of buffers. With Version 6.3, the new maximum is set at 255 buffers per file. The value of 255 buffers is an arbitrary limit, and not a design limitation. The current buffer management design will allow for a significantly larger number of buffers without impacting CPU time. With this enhancement, IAM can bring significant performance improvements to even the most highly accessed files. The maximum number of buffers that can be used for a file can be specified either by the IAM MAXBUFNO override, or by the VSAM BUFND parameter. Below is an example of using an IAM Override to increase MAXBUFNO.

```
//CICS      EXEC  PGM=DFHSDIP
//IAMFILE   DD    DSN=my.i am. file ,DISP=OLD
//IAMOVRID  DD    *
              ACCESS DD=IAMFILE ,MAXBUFNO=128
/*
```

Figure 12: JCL Example of Raising MAXBUFNO for an IAM File

01.40 ONLINE (CICS) ENHANCEMENTS

The maximum buffers that IAM can use for a file can also be specified on the JCL DD card for the file, by using the BUFND subparameter of the AMP parameter. The only distinction IAM has between the two, is that if both are specified, the MAXBUFNO override is that value that will be used. As with VSAM, note that the BUFND value specified in JCL will override any BUFND value specified on the VSAM ACB, or for the CICS FCT entry.

```
//UPDATE      EXEC  PGM=mypgm
//IAMFILE     DD    DSN=my.i am. file ,DISP=OLD ,
//            AMP=( ' BUFND=128 ' )
```

Figure 13: JCL Example of Specifying Maximum Buffers with BUFND

Multiple Concurrent Physical I/O

While more buffers will help heavily accessed files substantially, that only represents part of the solution for these files. To further improve responsiveness, many customers have split these heavily accessed files across multiple volumes to achieve concurrent physical I/O to the file. IAM with the Compatible File Format was limited to having only a single physical I/O active at any point in time. In effect, IAM was performing the I/O queuing. With the Enhanced File Format, IAM will now initiate physical I/O requests as they are needed, providing the capability to have concurrent physical I/O's active when a file is distributed across multiple volumes. This capability, combined with the IAM Real Time Tuning algorithm, and the increase in allowable buffers will position IAM as the product to meet the increasing demand for online system responsiveness.

Balancing I/O with Real Storage and Paging Demands

One of the difficulties with Real Time Tuning is how to balance the buffering needs for any particular file, without causing unnecessary real storage and paging demands. In particular, from statistics that customers have sent in, Innovation observed that there were frequently some files within CICS regions that had relatively low I/O activity. However, because CICS was up for many hours, the number of buffers was slowly being increased. This increase in buffers could drive up real storage and paging demands, because by the time buffers are referenced again, they were probably paged out. Other concerns are, of course, being able to efficiently manage a varying number of buffers without eating CPU cycles. Because IAM Real Time Tuning generally runs unauthorized, in problem state, it is limited in its ability to very quickly check on real storage demands and paging. With the Enhanced File Format, when IAM Real Time Tuning is making a decision about changing the number of buffers, the I/O rate and buffer reference frequency are taken into consideration. For example, when physical I/O frequency is very high, the number of buffers will be increased as long as it appears that an additional buffer will provide some benefit. This will aid in IAM being more responsive to sudden high demands against a particular file.

The main reason for looking at frequency is to determine whether the buffering is really providing a benefit, or only causing an impact on real storage. When a decision is being made to adjust buffers, the buffer reference frequency is taken into consideration. For example, if it has been a few minutes since the oldest buffer has been referenced, then in all likelihood it is paged out. While paging I/O is very quick, a typical IAM buffer occupies multiple pages, which may cause multiple page faults. The paging I/O can have more serious negative impact on online system performance, than doing a file I/O. So, in this circumstance, IAM will reduce buffers when buffer reference frequency is very low to attempt to prevent the paging activity. The intent is to free the real storage for those files that have very heavy I/O demand.

Certainly manually adjusting the IAM MINBUFNO and MAXBUFNO values for the low activity files can be done. However the intent of Real Time Tuning is to attempt to minimize the manual effort of adjusting the buffers for overall performance. With all of these above enhancements, IAM's Real Time Tuning is at work, with a goal of improving online system responsiveness with minimal manual effort.

01.40 ONLINE (CICS) ENHANCEMENTS

ENHANCED ASYNCH- RONOUS CAPABILITIES

To further improve online system responsiveness, IAM has enhanced the internal handling of asynchronous I/O requests. Online systems, such as CICS, use an asynchronous method of I/O, so that while the I/O is active, CICS can process other transactions. With the IAM Enhanced File Format, IAM can process multiple concurrent logical I/O requests, in a truly asynchronous manner as seen from the application program. The benefits of the enhanced design are many, but the bottom line is providing faster response to the I/O service requester. For example, I/O requests that are for data residing within the IAM buffer pool, or within the IAM Dyncore area will no longer be delayed by uncompleted prior requests, that are doing physical I/O. This improved asynchronous processing will greatly minimize, and hopefully eliminate, the problems encountered with setting the CICS AMAXTASK parameter to a value lower than the CICS MAXTASK.

Another related new capability is IAM's support of the VSAM UPAD exit. This is a VSAM exit used by CICS to provide lower overhead asynchronous processing when files are in the LSR buffer pool. With the UPAD exit, there is no need to issue the VSAM CHECK macro, eliminating that processing overhead, as well as the overhead of scheduling IRB's to handle the asynchronous requests. Enhanced Format IAM files can now be left in the CICS LSR buffer pools. While IAM will not use any of the buffers in the pool, it does provide enhanced usability and capabilities for CICS users of IAM files. In particular, there is no longer any need to change the FCT entries when converting files from VSAM to IAM, to turn off LSR buffering. Secondly, IAM files can now be used by the CICS Data Table Facility, as an alternative to the IAM Dyncore feature. Thirdly, the CPU overhead is lower using the UPAD exit, as is done by CICS for files in the LSR pool, when compared to the need to issue CHECK macros with asynchronous processing. There may also be some reduction in DASD write I/O for transactions that perform multiple updates. This is because when CICS uses LSR buffering, it defers the write I/O until the updates are to be committed to the file.

If the IAM files are left in the LSR pools under CICS, at some point it may be beneficial to reduce the number of buffers in the affected LSR pools. This is because IAM will not use the buffers, but VSAM will still use them for any files that are left in the pool. This may cause unnecessary demand for real and virtual storage, and have a negative impact on overall CICS performance.

DYNAMIC REGION SIZE ADJUSTMENT

As customers have converted files to IAM, they occasionally hit the MVS default limit of 32 megabytes extended private storage. This has in the past necessitated the modification, or in many circumstances writing, and supplying an IEFUSI exit. Innovation has recently developed, and distributed on request, a sample IEFUSI exit, that frequently can be used with minor modifications. This exit is being distributed in the ICL (Installation Control Library) provided with IAM Version 6.3. The big problem is that installing the new, or modified exit requires an IPL.

Because many customers have unexpectedly hit this limitation, IAM Version 6.3 has the ability to dynamically increase the extended private region limit. When IAM is opening an Enhanced Format File, it estimates the amount of storage that will be required to open the file. Included in that estimate is storage for the prime and overflow indexes, control information, and buffers. IAM then checks to see if that amount of virtual storage is available. If not, IAM will then attempt to increase the extended private storage limit by the quantity required to open the file, rounded to 4 megabytes. By default, IAM will not increase the extended private region to greater than 128 megabytes. This maximum value can be either increased or decreased through the IAM Override facility, or by changing the IAM default in the IAM Global Options Table.

01.40 ONLINE (CICS) ENHANCEMENTS

While processing a file, if a critical storage acquisition fails, that will also drive the dynamic region adjustment, with the same limitations as above. A storage acquisition is considered critical if it is required to successfully complete a request. For example, one critical storage acquisition is to expand the size of the overflow index to complete an insert or update request. If storage can not be obtained, then the insert or update request will fail with a File Full VSAM logical error. With Dynamic Region Adjustment, when a critical getmain fails, IAM performs the region adjustment. It then retries the failing getmain, and if it fails again, then the request is failed. Dynamic Region Adjustment will not be invoked to increase buffers, because processing can continue without failing requests.

Dynamic region adjustment affects only normal file access, it does not function during a file load process. The IAM ACCESS override keyword is MAXREGION, and it specifies the upper limit for the above the line storage, in megabytes. Note that by specifying the keyword DD=&ALLDD, the value will be effective for all Enhanced Format IAM files, unless there is a specific override. In the example below, the IAM files specified by the DDNAMES of IAMFILE1 and IAMFILE4 are handled by the &ALLDD value, so MAXREGION override will be used for those files. The MAXREGION override will also be used for IAMFILE2, because it has MAXREGION explicitly specified. However, the default from the IAM Global Options Table will be in effect for IAMFILE3. This means that once the above the line region exceeds the default value of 128 megabytes, then the region will not be further enlarged for storage for that file. However it will be enlarged for the other files.

```
//CICS      EXEC  PGM=DFHSDIP
//IAMFILE1  DD    DSN=my.i am. file1,DISP=OLD
//IAMFILE2  DD    DSN=my.i am. file2,DISP=OLD
//IAMFILE3  DD    DSN=my.i am. file3,DISP=OLD
//IAMFILE4  DD    DSN=my.i am. file4,DISP=OLD
//IAMOVRID  DD    *
           ACCESS DD=&ALLDD,MAXREGION=256
           ACCESS DD=IAMFILE2,MAXBUFNO=128,MAXREGION=256
           ACCESS DD=IAMFILE3,MAXBUFNO=8
/*
```

Figure 14: Providing MAXREGION Override

The following example demonstrates how to change the IAM default for the MAXREGION specification. In this example, the default is raised to 256 megabytes. Note that by setting the value to 32 megabytes or less will effectively shutdown the IAM Dynamic Region Adjustment feature.

```
//ZAPOPOP   EXEC  PGM=IAMZAPOPOP
//STEPLIB   DD    DSN=my.i am. loadlib,DISP=SHR
//SYSLIB    DD    DSN=my.i am. loadlib,DISP=SHR
//SYSPRINT  DD    SYSOUT=*
//SYSIN     DD    *
           ZAP    MAXREGION=256
           PRINT
/*
```

Figure 15: Changing the IAM Default Value for the Maximum Dynamic Region Adjustment

01.40 ONLINE (CICS) ENHANCEMENTS

IAM VIRTUAL STORAGE CONSTRAINT RELIEF (VSCR FOR OSCORE)

For IAM Version 6.3 with the Enhanced File Format, IAM has moved most of the storage required for control block areas to extended private storage, along with all of the new IAM load modules. Also, a few of the older modules have been moved above the 16 megabyte line as well. Where possible, some of the VIF modules are now loaded above the 16 megabyte line. For an Enhanced Format File, IAM will typically require only 4K of 24-bit addressable storage, also referred to as below the 16-megabyte line storage, under CICS. For CICS customers, this storage area has typically been referred to as OSCORE.

Most of the storage required for open processing is also acquired above the line. No longer is an initial buffer acquired in 24-bit addressable storage. Plus, the table that holds the IAM override values is obtained in extended private. The buffers used for the IAM Override processor are requested in extended private storage, when the customer is running at a DFSMS/MVS level that supports that capability.

THE IAM EXECUTION MONITOR FOR CICS – IAMXMON

Overview

The ability to view current information pertaining to IAM files as they are being used, especially under CICS, has long been desired. IAMXMON is the answer! IAMXMON allows you to view up to the second, detailed information about any open IAM file within a CICS region. You can monitor I/O activity, storage usage, buffer usage, etc.

Usage

After entering the IAMXMON transaction ID (determined by your installation), you will be presented with a list of all open IAM data sets. Each entry in the list contains the IAM files DDNAME, Data set Name, and current storage usage. To view more detailed information about any particular file, select it by placing an 'S' next the desired DDNAME and hitting Enter. Please note that if an IAM file is still under control of the IAM OPEN process, the file will not appear in the current list of open IAM files. Only after an IAM file has completed the OPEN function will it appear on the Data set list map.

The two detail maps give you most of the information found on the IAMINFO reports. The IAMINFO reports are obtained by placing a //IAMINFO DD SYSOUT=* card in any batch step that you require detailed IAM information. The two maps can be reached by using the scroll forward (PF8/20) and scroll backward (PF7/19) keys. To return back to the Data set selection map, use PF3/15. Hitting only Enter on any map will refresh the map with the most current information available.

If, during the viewing of detail information about an IAM file, the file is detected as being closed, all fields will be cleared to blanks, and a message will be displayed at the bottom of the current map stating that the IAM file is no longer valid. Hit any key to view a refreshed Data set list map.

01.40 ONLINE (CICS) ENHANCEMENTS

SAMPLE Summary Screen: DSNLIST SCREENS

The initial screen lists all open IAM data sets within the CICS region. By placing an 'S' next to the IAM DDNAME and pressing ENTER, the first of two detail information screens will be displayed.

Sel	DDNAME	Data set Name	Storage
	KSDS9	IAMV.\$IAM.GAMA.KSDS9	000069632
	KSDS4	IAMV.\$IAM.GAMA.KSDS4	000069632
	KSDS10	IAMV.\$IAM.GAMA.KSDS10	000069632
	ESDS7	IAMV.\$IAM.GAMA.ESDS7	000114688
	ESDS1	IAMV.\$IAM.GAMA.ESDS1	000069632

Figure 16: Display of all Open IAM Files Under CICS

First Detail Screen for a Selected IAM File

The first detail screen consists of execution statistics that can change during the time an IAM file remains open and in use.

IAM Execution Monitor		(c) 1995 Innovation Data Processing	
Command ==>			
DDNAME: KSDS9	DSN: IAMV.\$IAM.GAMA.KSDS9	Page 1 of 2	
IAM Command Execution Summary			
GET Random.....00000000	PUT Update....00000000	ENDREQ.....00000000	
Get Seq.....00000000	PUT Add.....00000000	WRTBFR.....00000000	
Get Prev.....00000000	Point.....00000000	OPEN.....00000001	
GET KGE.....00000000	Point KGE....00000000	CLOSE.....00000000	
GET (skip).....00000000	ERASE.....00000000	CLOSE T.....00000000	
Verify.....00000000	Invalid Req..00000000	Len Change....00000000	
Seq Chain R.....00000000	Seq Chain W...00000000	Flush Bfr.....00000000	
IAM Executions Statistics		Total Records.....0000000900	
Requests Processed.....000000000	Requests Failed.....000000000		
Blocks Read.....000000002	Blocks Written.....000000000		
Dynamic Table Retrievals...000000000	Dynamic Table Records...000000000		
Dynamic Buffer Retrievals...000000000	Minimum Buffers.....0001		
Storage Above 16M.....000065536	Maximum Buffers.....0022		
Storage Below 16M.....000004096	Current Buffers.....0004		
Extended Overflow Information			
Available Overflow Blocks...000000000	Extended PE Blocks Used..000000000		
Extended Data Blocks.....000000000	Extended Data Records...000000000		

Figure 17: First Detail Information Screen for CICS

01.40 ONLINE (CICS) ENHANCEMENTS

Second Detail Screen for a Selected IAM File

The second detail screen consists of static information about the selected IAM file, such as record length, key length, key offset, etc. Active processing options for this IAM file are also displayed.

```
IAM Execution Monitor
DDNAME: KSDS9      DSN: IAMV.$IAM.GAMA.KSDS9
Page 2 of 2

Controlling Program.. DFHSKTSK

Processing Options      Data set Attributes
OPENed for UPDATE..Yes  Data set Type...KSDS   Record Size....00084
Immediate WRITE....No  Shareoptions...2      Block Size....11476
OPTCD=ASY.....No      Extents.....001       Recs/Block....00136
OPTCD=WAITX.....No    Blks/Track....004     Record Fmt....V
DEFER WRITE.....No    Data Comp.....No      Index Comp....No
CI Processing.....No   Key Length....008     RKP.....00004
JRNAD Loaded.....No   Integrated %...000
JRNAD Active.....No
UPAD Loaded.....No
UPAD Active.....No
```

Figure 18: Second Detail Information Screen for CICS

01.50 ADDITIONAL ENHANCEMENTS

01.50 ADDITIONAL ENHANCEMENTS

There are numerous other enhancements with IAM Version 6.3. These improvements will provide performance benefits to various types of batch job processing. Support for the latest versions of MVS and DFSMS, plus enhancements to expand the set of software products that IAM will function with are also provided. IAM provides outstanding performance for batch work, and with these enhancements, certain types of batch processing will now complete even faster. Most of these enhancements are completely automatic, requiring no JCL or programming changes, only defining the IAM file as an Enhanced Format File.

I/O ENHANCEMENTS

There have been several enhancements in I/O and buffer management. As described above in the section on IAM Real Time Tuning Enhancements, on page 13, IAM now has the ability to use even more buffers. For most batch jobs there probably is little need to increase the buffering. The capability is there, however. More importantly for typical sequential batch work is that IAM now has an I/O Read Ahead capability for sequential reads. That capability will provide for the possibility of greater I/O and processor overlap, without going into an asynchronous mode of operation. The most IAM will read ahead is one cylinder. So, except for buffers for overflow blocks, the maximum buffers used by a sequential input program is the number of blocks per cylinder. Typically that value is 60, assuming default of 1/4 track blocking on a 3380 or 3390 device. For mass sequential update applications, IAM now has the ability to write out multiple blocks per I/O. This will help reduce clock time, as well as EXCP's for those types of jobs. The IAMINFO report contains a run time statistic indicating the benefit of this capability, in the field labeled "SEQUENTIAL CHAINED BLOCKS WRITTEN".

The IAM I/O driver will also take advantage of the IBM Record Level Caching, for random I/O's being done on 3990-6 controller, and indicate sequential processing for the 3990-3 and 3990-6 controllers when applicable. The IAM channel programs completely support the ESCON channels, utilizing the ECKD channel commands, for files on control units that support that capability. IAM continues to provide optimal use of the latest hardware in the market, by using blocksizes that fully utilize the device's capability, and by using the latest features and capabilities when performing I/O.

BACKING UP COMPRESSED IAM DATA

IAM Version 6.3, with the Enhanced File Format, now offers the capability to backup and reload compressed data within an IAM file without decompressing or compressing the data. For large files, this is anticipated to allow IAM files to be backed up and reorganized faster than can be done today. Even when the data is compressed at the 3480/3490 control unit level, there is still the overhead of transferring all that data to the controller. With this new feature, both the CPU overhead and that I/O overhead is eliminated. FDRREORG V5.2/50 from Innovation, will automatically use this new IAM feature.

The backup and reload of compressed data is specified for other programs, such as IDCAMS, by the use of the IAM Override facility. The override will have to be specified on both the backup and reload process, because IAM needs to know to not compress the data on the backup side, and that the data is already compressed on the input side. By simply specifying the keyword BACKUP COMPRESSED on the ACCESS and CREATE IAM overrides does the job. IAM adds four bytes to each record when performing this function, so any output file created will have to contain either variable (RECFM=VB) or undefined (RECFM=U) type of file. For variable output files, the record length for the output file (LRECL) will need to be at least 8 bytes more than the defined maximum record length for the file. For example, if the maximum record length for the file is 100, then the output LRECL must be at least 108. For undefined type of records, the maximum LRECL is 104, only 4 bytes more than the file maximum record size. Innovation recommends using RECFM=VB type of output to provide the best output device utilization.

01.50 ADDITIONAL ENHANCEMENTS

Using Backup of Compressed Data with FDRREORG

The FDRREORG utility, available from Innovation, will automatically make use of the new backup compress feature. FDRREORG V5.2/50 is offering another advantage as well, which is the capability to perform Parallel Sequential Reads on multivolume files, which along with the ability to read data without decompressing it will assist in greatly speeding up the file reorganization process for those very large files. FDRREORG offers the capability to only reorganize a file when necessary, based on criteria set by the user. Please contact Innovation for availability dates for these enhancements to FDRREORG. An example of JCL for FDRREORG is provided below.

```
//REORG      EXEC  PGM=FDREORG
//SYSPRINT   DD    SYSOUT=*
//REORGPRT   DD    SYSOUT=*
//REORGRPT   DD    SYSOUT=*
//IAMINFO    DD    SYSOUT=*
//SYSIN      DD    *
              REORG      DSTYPE=IAM
              SELECT     ALLDSN,VOL=(IAM*)
/*
```

Figure 19: JCL to REORG an IAM file with FDRREORG

File Reorganization of Compressed Data with IDCAMS

An example of doing a file reorganization with IDCAMS and the new IAM Backup of Compressed data is shown below. Please take note of the IAM overrides. The buffering specifications for the IAM files will provide for the maximum buffering. The MAXBUFNO=64 on the ACCESS statement will allow one cylinders worth of buffers, presuming that the file has a blocking factor of four (B=4), which is the default for most IAM files. The additional four buffers are to hold any Extended Overflow blocks, without having to steal from the chained sequential buffers. This example is assuming that the file is defined with a maximum record size of 100, so the LRECL for the output file is set to 108.

```
//REORG      EXEC  PGM=IDCAMS
//STEPLIB    DD    DSN=my.i am63.loadlib,DISP=SHR
//SYSPRINT   DD    SYSOUT=*
//IAMINFO    DD    SYSOUT=*
//IAMPRINT   DD    SYSOUT=*
//IAMFILE    DD    DSN=my.i amfile,DISP=OLD
//SEQFILE    DD    DSN=my.seqfile,DISP=(,CATLG),UNIT=TAPE,
//            DCB=(RECFM=VB,LRECL=108,BLKSIZE=32760,BUFNO=30),
//            VOL=SER=SCRATCH
//IAMOVRID   DD    *
              CREATE     DD=IAMFILE,CRBUFOPT=MCYL,BACKUPCOMPRESSED
              ACCESS     DD=IAMFILE,MAXBUFNO=64,BACKUPCOMPRESSED
/*
//SYSIN      DD    *
              LISTCAT ENT(my.i amfile) ALL
              REPRO INFILE(IAMFILE) OUTFILE(SEQFILE)
              REPRO INFILE(SEQFILE) OUTFILE(IAMFILE) REUSE
              LISTCAT ENT(my.i amfile) ALL
/*
```

Figure 20: JCL for Reorganizing an IAM File with BACKUPCOMPRESS

01.50 ADDITIONAL ENHANCEMENTS

Recovery of Data that is Compressed

The IAM file recovery program, IAMRECVR, is enhanced with IAM Version 6.3 to be able to read sequential files created with the compressed IAM data, and put out a sequential file with the data uncompressed. This may be useful for when an application program needs to read the sequential file from FDRREORG or IDCAMS REPRO of the IAM file. This is facilitated with the new command, DECOMPRESS.

DECOMPRESS Command Statement

The DECOMPRESS command allows a compressed backup of an IAM file to be decompressed in the event that IAM is not available. The IAM VSAM Interface (VIF) does not have to be active.

```
DECOMPRESS [ FROMDDNAME=ddname ]
            [ , TODDNAME=ddname ]
            [ , KEYLEN=nnn ]
            [ , RKP=nnnnn ]
            [ , SCAN ]
```

Figure 21: IAMRECVR DECOMPRESS Command Format

01.50.1 OPERAND DESCRIPTIONS

01.50.1 OPERAND DESCRIPTIONS

The table below contains descriptions of the keywords for the DECOMPRESS command of IAMRECVR. The minimal abbreviation for each operand keyword is underlined.

Keyword	Description
<u>FROMDDNAME=</u>	Defines the DDNAME of the compressed IAM backup file. The default is DISKIN.
<u>TODDNAME=</u>	Defines the DDNAME of the output uncompressed flat file. The default is DISKOUT.
<u>KEYLEN=</u>	Specifies the length of the key within the data records in the IAM file. This value can be obtained from run time statistics or a LISTCAT command.
<u>RKP=</u>	Specifies the relative location of the key within the data record in the IAM file. This value can be obtained from run time statistics or a LISTCAT command.
<u>SCAN</u>	Allows a compressed backup file to be read and all records decompressed, but an output file will not be created. SCAN can be used to verify a compressed backup file.

Figure 22: Descriptions of DECOMPRESS Operands

01.50.1 OPERAND DESCRIPTIONS

DECOMPRESS Example

Assume an IAM file with 100 byte records, an 8 byte key length with a relative key position (RKP) of 10 has been backed up without decompression and now must be restored to its uncompressed state to a sequential output file.

```
//DECOMPRS EXEC PGM=IAMRECVR
//SYSPRINT DD SYSOUT=*
//DISKIN DD DSN=my.seqfile,DISP=OLD
//DISKOUT DD DSN=my.uncomp.seqfile,DISP=(,CATLG),
// DCB=(RECFM=VB,LRECL=104,BLKSIZE=32760),UNIT=TAPE
//SYSIN DD *
DECOMPRESS KEYLEN=8,RKP=10
/*
```

Figure 23: Using IAMRECVR to Decompress a Compressed Output File

ENHANCED SINGLE RECORD LOAD PROCESSING

The strategy for handling a single record load is changed with IAM Version 6.3 Enhanced Format files, due to the ability to dynamically increase the file size. With Compatible format files, a single record load is handled by increasing the number of Independent Overflow and PE blocks in the file, to use about 90% of the space allocated to the file during the load. If the single record being loaded has a key of any value other than all high values, x'ff', the record is inserted into an Independent Overflow block, to maximize the use of PE blocks when the subsequent mass insert was done. If the key was high values, then the record was placed in a PE block, because such a record could not exist in an Independent Overflow block. While this strategy generally worked quite well, some difficulties arose due to the need to overallocate the file on the IDCAMS Define. The time to format out the file initially could be quite significant, resulting in longer elapsed times to do the full process of single record load, followed by mass insert.

With the Enhanced Format files and IAM Version 6.3, the strategy of placing the record loaded into an overflow block remains. IAM Enhanced Format files can have a record with a high value key in Extended Overflow, so the special considerations for that key value have been eliminated. Only the few blocks required for control information and the Extended Overflow block will be formatted at file load time. In anticipation of a subsequent mass sequential insert, the allocated DASD space will not be automatically released. Because Extended Overflow and Extended PE blocks are allocated as required, the guesswork of how many blocks each area will need is gone. This also will prevent records from going into overflow unnecessarily. With Compatible Format files, once all the PE blocks were used, the records were then placed in Independent Overflow. For Enhanced Format files, where PE and Overflow blocks are acquired dynamically, this will not occur, saving DASD space. The ability to acquire secondary extents means that the allocations for these single record load files can be set up to handle the typical amount of data, with secondary extents being acquired when necessary to handle increased volume.

It is still always more efficient to do a full file load, as opposed to the single record load and a mass insert. With these enhancements in Version 6.3, the penalty for doing a single record load has been significantly reduced for most circumstances, and the usability has improved with going into extents as needed.

01.50 ADDITIONAL ENHANCEMENTS

ENHANCED SOFTWARE COMPATI- BILITY

IAM Version 6.3 provides support for MVS/ESA 5.2 and DFSMS 1.3. This includes support for UCB's in extended storage (i.e., 31-bit addressable storage), and for the changes in the VSAM control blocks. IAM files are still being processed through a standard DCB as a non-VSAM type of file, for which DFSMS 1.3 requires a captured UCB. IAM system level VSAM Interface, VIF, has been enhanced with a new routine to make sure that the UCB's will be captured when necessary.

It has come to our attention that various VSAM recovery and journaling packages, which can be used equally effectively on IAM files, have incorporated calls to the DFSMS Backup While Open (BWO) function. Prior versions of IAM had not provided any capabilities. With Version 6.3, IAM now has an BWO interface, that is part of the IAM system level VSAM interface. The IAM BWO interface for Version 6.3 will indicate that BWO facilities have not been used on IAM files, and are therefore allowable to go through a recovery process. This capability will eliminate the error messages and failures encountered by these software products for IAM files. Innovation is still evaluating the need for support of BWO with the new Enhanced Format Files, and may provide such support, if needed, in the future.

The IAM support for the VSAM JRNAD Exit has also been enhanced to invoke the exit just prior to an IAM buffer write. This support will enable those software products that are journaling IAM files to insure that all the necessary data is recorded in their journal prior to IAM actually updating the file.

IAM Version 6.3 VSAM Interface (VIF), is enhanced to exist with and where applicable, provide support for the BIM line of products. This is an optional feature, which is enabled by the use of the IAM Global Options change facility, IAMZAPOP. For the support to be in effect, VIF must be started after enabling the option. IAM's VIF can be started either before or after BIM with this feature, however BIM **MUST** be active prior to the first open of an IAM file, otherwise the support is dynamically disabled. The following example demonstrates how to enable the BIM product support.

```
//ZAPOP      EXEC  PGM=IAMZAPOP
//STEPLIB    DD    DSN=my.i am. loadlib,DISP=SHR
//SYSLIB     DD    DSN=my.i am. loadlib,DISP=SHR
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
              ZAP  ENABLE=BIM
              PRINT
/*
```

Figure 24: Enabling IAM Support for BIM

MULTIPLE ACB SUPPORT

For Enhanced Format Files, IAM will now recognize that an ACB has already been opened for a file, within the same address space, and share the index structure and buffers. For applications, or CICS regions, where a file was opened under one DDNAME or ACB for read only, and the other for UPDATE, they occasionally experienced an inability to retrieve an updated data record. This will no longer be the case. This support also reduces the storage requirements for having multiple ACB's opened to the same file, and eliminates the open overhead for this circumstance as well.

IAM will also now perform Dynamic String Acquisition. If an insufficient value for STRNO has been specified, IAM will acquire additional place holders, as needed.

01.50 ADDITIONAL ENHANCEMENTS

MASS SEQUENTIAL DELETES

For Enhanced Format files, IAM has incorporated an ability to temporarily logically delete a record from a data block. Then, when that data block is about to be written out to DASD, the records are physically deleted from the data block. This eliminates the overhead of constantly moving records within a data block, as prior records are deleted, only to end up being deleted as well. In the case where every single record is deleted from a data block, this enhancement eliminates the data movement entirely from that process.

A further difficulty has been rarely encountered after such mass deletes, where certain types of I/O requests result in the empty data blocks being read repeatedly, often unnecessarily. IAM does sequential processing without referring to any index structure, because the internal structure of the file does not require such overhead. IAM has no way to know that a prime data block is empty, without actually reading that data block. With Version 6.3 Enhanced Format files, that are defined with Share Option 1, or that are opened for update with Share Option 2, will now keep track of those empty blocks, and not reread such blocks in sequential modes of processing. This support is limited to particular circumstances, because in other situations there could be another ACB opened for UPDATE against the same file, and therefore there can be no presumption about the contents of a data block.

FILE DATE AND TIME STAMPS

With the Enhanced File Format of Version 6.3, there are three TOD stamps kept in the header of the file. These are when the file is defined, when the file was last loaded, and when the file was last updated. Each of these date and time stamps will be printed on the IAMINFO run time reports, and on the IAMPRINT reports from Listcat. There is also a new IAM Global Option, EURODATE, which will indicate that for Listcat, the date is to be printed as dd/mm/yy instead of mm/dd/yy.

MODEL AND LIKE SUPPORT

IAM Version 6.3 now provides support for IAM files to be the object of MODEL keyword in IDCAMS DEFINE, and the LIKE keyword in JCL and TSO ALLOC command. Both IAM and VSAM files can be defined or allocated using the specified IAM file's attributes. Note that IAM does not keep track of attributes such as IMBED, REPLICATE, SPEED, and others which may result some differences between a file defined with a VSAM model versus an IAM model.

01.50 ADDITIONAL ENHANCEMENTS

PSEUDO LOGICAL RECORD LENGTH SUPPORT

This is a new feature that may help reduce DASD space requirements for some IAM files. For files that have a large theoretical maximum record size, but in actuality never have records that size, the file can now be defined with the smaller record size. With the Pseudo LRECL support, the larger maximum record size is specified on an IAM CREATE override, PSEUDOLRECL, when the file is defined. Whenever the file is opened, the application program will be told that the file can support the larger record size. This allows COBOL programs to access the file, while at the same time reducing the DASD space, particularly for Overflow, required for the file. Also, files that may have been forced into a 32K blocksize, resulting in inefficient use of DASD space, may now be able to use a blocksize that will better utilize the device. Note that any attempts to actually write out a record with a size greater than the defined size will be failed.

For example, a file has a maximum theoretical record size of 32,000 based on the COBOL record layout. However, let's say that the actual largest record in the file is 8,000 bytes. The file can be defined with a maximum record size of 8,000, and a pseudo record size (PSEUDOLRECL) of 32,000. IAM will use a more efficient block size, for example 27,998 on a 3390, which will allow full utilization of each track, instead of a blocksize of 32,760, which will waste 23,236 bytes per track, or 41% of the DASD space. Additionally, only 8,000 bytes will be reserved per record in overflow, instead of the 32,000 bytes, also reducing DASD space requirements. This second type of savings is available to files with smaller record sizes as well. The following example shows how to define this file with the Pseudo LRECL feature:

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//IAMOVRID   DD    *
              CREATE DD=&ALLDD,PSEUDOLRECL=32000
/*
//SYSIN      DD    *
              DEFINE CLUSTER( NAME (my.i.am.cluster) -
                          OWNER($IAM) -
                          RECORDSIZE(2000,8000) -
                          KEYS(12,8) -
                          FREESPACE(10,10) -
                          VOL(myvol) -
                          CYL(10,2) -
                          SHAREOPTIONS(2,3) )
/*
```

Figure 25: Defining an IAM File with PSEUDOLRECL

The Pseudo LRECL feature can be used with either Enhanced Format or Compatible Format files under Version 6.3. This capability was available on prior releases with a custom PTF. For customers that have that PTF, please note that the affected file(s) will have to be redefined with the PSEUDOLRECL option to run under Version 6.3.

01.60 CONVERSION CONSIDERATIONS

01.60 CONVERSION CONSIDERATIONS

NEW IAM GLOBAL OPTIONS

As the new features of IAM Version 6.3 have been presented, examples were given about how to modify the various new default values in the IAM Global Options Table. For complete information on most of the IAM Global Options, please refer to [Section 91](#) of the IAM Version 6.3 Users Guide. That section provides complete information and examples on using IAMZAPOP to change various IAM default options. The table provided below is intended to summarize the new Global Options that are available in Version 6.3, which have been discussed elsewhere in this section.

Keyword	Description
DSORG=	Specifies the default DSORG for IAM files. This option was introduced with IAM Version 6.2. The possible values are DA or PS. The default value with IAM Version 6.3 has been changed to PS.
MAXREGION=	Specifies the default maximum value, in megabytes, that IAM will dynamically adjust the above the line Region value to. To disable feature, set value to 0. Default is 128.
MAXSECONDARY=(x,y)	Specifies the default multiplication factors for IAM Dynamic Secondary Space Adjustment feature, where x is the value for file loads, and y is the value during file updates. To disable this feature, set both values to 0. Default values are (10,5).
MULTIVOLUME=	Specifies either PRIMARY or SECONDARY, indicating which space allocation value IAM will use when it appears that a Data set will take the next extent on the next volume. Default value is PRIMARY.
ENABLE=ENHANCED	Specifies that IAM files will default to the Enhanced File Format when they are defined. Default value is that this option is disabled, which means that IAM files will default to the Compatible Format.
DISABLE=ENHANCED	Specifies that IAM files will default to the Compatible File Format, which is the default value.
ENABLE=EURODATE	Specifies that the TOD stamps from the file will have the date printed in the format of dd/mm/yy in the IAMPRINT Listcat output.
DISABLE=EURODATE	Specifies that the TOD stamps from the IAM file will have the date printed in the format of mm/dd/yy in the IAMPRINT Listcat output.
ENABLE=BIM	Specifies that the IAM VSAM Interface (VIF) will provide support for the BIM product. Default value is that IAM will not provide such support. Requires a restart of VIF to take effect.
DISABLE=BIM	Specifies that the IAM VSAM Interface (VIF) support for the BIM products is not available. Requires a restart of VIF to take effect.

Figure 26: New Global Options for IAM Version 6.3

01.60 CONVERSION CONSIDERATIONS

NEW IAM OVERRIDES

There are a few new IAM Override keywords for Version 6.3, which have been presented as the various new features were discussed. The following table will summarize the new IAM Override keywords. For complete information on the IAM Override processor, and all the previously existing keywords, please refer to [Section 11](#) of the IAM Version 6.3 User's Guide. Certain overrides are not relevant to the Enhanced File Format, and will be ignored if specified. For CREATE, these include OCOREO%, OCOREX%, CRBUFOPT=BSAM, FIXED, and PE. The OVERFLOW (O=) override has a new meaning for the Enhanced Format files, as discussed in the section DASD Space Reserve and Space Release, on page 8. The ACCESS overrides not applicable to Enhanced Format Files include BROWSE, MODE=, OCOREO%, OCOREX%, and OPTCD. The trace processing is different for Enhanced Format Files, as discussed in the section on Enhanced File Structure TRACE Output, on page 43. The following table summarizes the new Override keyword values for Version 6.3. Unless otherwise noted, the keywords are applicable to both CREATE and ACCESS overrides.

Keyword	Description
MAXREGION =	Applies to ACCESS override only. Specifies the maximum value, in megabytes, to which IAM Dynamic Region Adjustment will set the above the line region to. A value of 0 will disable this feature. Default value, from the IAM Global Options Table, is 128.
MAXSECONDARY =	Multiplication factor to be used by IAM Dynamic Region Adjustment feature. Default values are from the IAM Global Options Table. Values from 0 - 10 are permitted. A value of 0 will disable the feature. Default values are 10 for CREATE, and 5 for ACCESS.
MULTIVOLUME =	Indicates whether IAM is to use the PRIMARY or SECONDARY space allocation value, when IAM anticipates that the next extent will be obtained on another DASD volume. Default, from the IAM Global Options Table, is PRIMARY.
PSEUDOLRECL=	A CREATE override, which is applicable when a file is being defined, that specifies the theoretical maximum record size. The actual maximum record size is taken from the DEFINE parameters.
BACKUPCOMPRESSED	For ACCESS, indicates that data records will be left in their compressed format when passed to the calling program. For CREATE, indicates that the application program is passing records already compressed by IAM. Typically only used on IDCAMS REPRO requests, to backup, restore, and reorganize IAM files.

Figure 27: New IAM Overrides for Version 6.3

01.60 CONVERSION CONSIDERATIONS

IAM LISTCAT OUTPUT (IAMPRINT)

There have been a number of revisions to the IAM List Catalog output, as seen on the TSO terminal, or for batch jobs, on the IAMPRINT DD. Additional information that has been saved from the file define is now available, including primary and secondary space allocations, the original control interval size (CI Size), and the original freespace values, both CI and CA. Revised for IAM Version 6.3 Enhanced files is the value printed for Record Size and the value printed for Key Offset. With the Compatible format, if the records were variable in length, or data compressed, IAM would add four to each of those fields for the hidden RDW (Record Descriptor Word) used by IAM. With the Enhanced Format file, the reports print out the actual user supplied values, to eliminate confusion due to the changes IAM had made to those values. There are some new fields on the report due to the Enhanced Format, that present information about the Extended File Areas, including Extended Overflow and Extended PE.

Obtaining Listcat Output

The IAMPRINT Listcat output is easily obtained. A batch job can be run, or a LISTCAT command can be issued from the TSO Command Line. The IAM Listcat processor first checks for a DD name of IAMPRINT. If one is found, the report is written to that DD. Normally, it is a SYSOUT file, however a DASD data set can be used. If an IAMPRINT DD is not found, then under TSO the output is directed to the terminal. For batch, a SYSOUT file is dynamically allocated, as long as there is a SYSPRINT also allocated to SYSOUT. Make sure to specify the ALL operand on the Listcat to obtain the IAM report. Below is an example of a batch job to obtain Listcat output.

```
//LISTCAT      EXEC  PGM=IDCAMS
//SYSPRINT     DD    SYSOUT=*
//IAMPRINT     DD    SYSOUT=*
//SYSIN        DD    *
                LISTCAT  ENTRIES(my.i.am.enhanced.format.file) ALL
/*
```

Figure 28: JCL to Obtain IAM Listcat (IAMPRINT) Output

01.60 CONVERSION CONSIDERATIONS

A Sample of Listcat Output

The following is an example of output from the IAM Listcat processor. Shown is both the SYSPRINT and the IAMPRINT. Note that the SYSPRINT for the IAM file looks like any typical non-VSAM data set, except that there are dots for OWNER. This is actually hexadecimal information stored by IAM, which can be prevented by enabling the NOOWNID IAM Global Option.

```
IDCAMS  SYSTEM SERVICES                                TIME: 16:32:15      01/16/96      PAGE 1
LISTCAT ENTRIES(IAMV.VIT132.CLUSTER) ALL              00680011
NONVSAM ----- IAMV.VIT132.CLUSTER
IN-CAT --- IAMVT.CATALOG
HISTORY
  DATASET-OWNER---.....      CREATION-----1996.016
  RELEASE-----2           EXPIRATION-----0000.000
VOLUMES
  VOLSER-----SCR083      DEVTTYPE-----X'3010200E'  FSEQN-----0
ASSOCIATIONS----- (NULL)
ATTRIBUTES
```

Figure 29: SYSPRINT Output from Listcat

```
IAM400 IAM CATALOG INFORMATION SERVICE -- VER 6.3/02P  DATE 1996.032 TIME 13:58:21
IAM100 IAM FILE ANALYSIS - DSN=IAMV.DCV201.CLUSTER
-----
FILE FORMAT -- = ENHANCED - FILE STATUS ----- = LOADED -
RECORD SIZE -- = 64 - FREESPACE - CI% ----- = 10 -
CI SIZE ----- = 4096 - FREESPACE - CA% ----- = 10 -
BLOCK SIZE --- = 11476 - EXTENDED OVERFLOW ----- = 79 RECS -
BLOCK FACTOR - = 4 - EXTENDED OVERFLOW ----- = 1 BLOCKS -
KEY SIZE ----- = 4 - EXTENDED OVERFLOW USED = 47 % -
KEY OFFSET --- = 12 - EXTENDED PE ----- = 0 BLOCKS -
FILE TYPE ---- = KSDS - EXTENDED ALLOCATED ---- = 36 BLOCKS -
DEVICE TYPE -- = 3380 - EXTENDED AVAILABLE ---- = 33 BLOCKS -
VOLUME COUNT - = 1 - SPACE USED ----- = 7 TRACKS -
VOLSER ----- = SCR083 - SPACE ALLOCATED ----- = 15 TRACKS -
TOTAL EXTENTS = 1 - TOTAL SPACE ALLOCATED - = 15 TRACKS -
PRIMARY SPACE = 6 - SECONDARY SPACE ----- = 2 CYL -
MULTIVOLUME -- = PRIMARY - MAX SECONDARY ----- = 6 CYL -
RELEASE ----- = NO - SHARE OPTIONS ----- = 2 -
DATA COMPRESS = ENABLED - INDEX COMPRESS ----- = NO -
TOTAL RECORDS = 4096 - INSERTS ----- = 0 -
UPDATES ----- = 1024 - DELETES ----- = 0 -
HIGH USED RBA = 309852 - HIGH ALLOCATED RBA ---- = 688560 -
FILE DEFINED - = 1996.032 - 02/01/1996 - 1:57 PM - = 13:57:50 -
FILE LOADED -- = 1996.032 - 02/01/1996 - 1:57 PM - = 13:57:57 -
LAST UPDATED - = 1996.032 - 02/01/1996 - 1:58 PM - = 13:58:09 -
STORAGE REQUIRED FOR KEY STRUCTURE ----- = 84 -
NUMBER OF IAM DATA BLOCKS ----- = 23 -
EXTENDED HIGH ALLOCATED RBN ----- = 59 -
-----
IAM499 IAMLISTC(6.3/02P ) PROCESSING COMPLETED
```

Figure 30: IAMPRINT Listcat Output for the Enhanced File Format

01.60 CONVERSION CONSIDERATIONS

Explanation of Fields on the IAMPRINT Report

IAMPRINT Field	Description of Field Contents:
File Format	Indicates whether the IAM file is an Enhanced format structure, or a Compatible format structure.
Record Size	The maximum amount of data that can be contained in a single record. Value is from the RECORDSIZE parameter on the IDCAMS DEFINE statement.
CI Size	The specified Control Interval (CI) size from the DEFINE. If CI size was not specified on the Define, IAM calculates a valid VSAM CI size based on the maximum record size.
Block Size	The physical block size IAM is using for the file. A block is the amount of data transferred in an I/O operation, and that is stored in a single contiguous stream of data on the actual device
Block Factor	Indicates the number of blocks per track, or the user specified block size from an IAM CREATE Override.
Key Size	For KSDS type of files, indicates the user defined length for the key of each data record.
Key Offset	Specifies the relative position of the key, as an offset from the beginning of the record, where a value of 0 indicates the first byte.
File Type	Indicates if the file is KSDS, keyed sequence data set, or an ESDS, entry sequence data set.
Device Type	Indicates the type of DASD device on which the data set currently resides, for example a 3380 or a 3390.
Volume Count	The number of volumes contained in the catalog entry for this file.
Volser	This indicates the volume(s) to which the data set is cataloged. This line is repeated for each volume that is in the catalog entry for the file.
Total Extents	Indicates the total DASD extents allocated to the file. An extent is a contiguous area of space on the device being used by the data set.
Primary Space	Indicates the requested primary space quantity, as indicated when the file was originally defined.
Multivolume	Has values of PRIMARY or SECONDARY, which indicates from which space parameter the size of the first extent on each DASD volume is allocated.
Release	Indicates if DASD space will be released when the file is loaded again. Generally this is YES before a file has been loaded, and is set to NO after the first load.

01.60 CONVERSION CONSIDERATIONS

IAMPRINT Field	Description of Field Contents:
Data Compress	Indicates whether or not this file can contain IAM data compressed records.
Total Records	The number of user data records in the file, as of the last close.
Updates	The number of user data records that have been updated since the file was last loaded.
High Used RBA	For IAM KSDS files, indicates the amount of space used for the file, in bytes. For IAM ESDS files, indicates the amount of actual user data contained within the file.
File Defined	The date and time when the file was last defined. The date can be in either mm/dd/yy format, or if the EURODATE Global Option is set, will be in the dd/mm/yy format.
File Loaded	The date and time that the last file load was completed. However, if a file load had started, but is not yet complete, this time will be the starting time of the file load.
Last Updated	The date and time of the last close of the file from a program that updated the file.
Storage Required for (Compressed) Key Structure	Indicates the amount of virtual storage that is required to contain the index to the Prime Data area of the file, and whether or not the index has been compressed.
Number of IAM Data Blocks	Indicates the number of blocks in the data set preceding the index area. This value may need to be supplied via the MAXBLKS parameter to IAMRECVR on a RECOVER operation.
Extended High Allocated RBN	Indicates the high allocated block number, to be used when it is necessary to run IAMRECVR to recover a file after it has been damaged. Value is specified via the XTNDEDHARB keyword.
File Status	Indicates if the file is LOADED or UNLOADED. A file will be in the UNLOADED state if it has been defined but has not had a successful load, or if a file load or reorganization has failed or is in progress.
Freespace – CI%	The amount of space to be left free in each block, as a percentage of the block size, as the file is being loaded or extended. Specified on the file definition. For IAM files, this is referred to as the Integrated Overflow Area.
Freespace – CA%	The amount of CA% freespace specified on the file definition. This is used by IAM to calculate an amount of allocated but unused space that is to be retained by the file after the initial load.

01.60 CONVERSION CONSIDERATIONS

IAMPRINT Field	Description of Field Contents:
Extended Overflow – Recs	The number of records in the Extended Overflow area of the file, as of the last time the file was closed.
Extended Overflow – Blocks	The number of blocks assigned to the Extended Overflow area.
Extended Overflow Used	Indicates, as a percentage, the amount of overflow spaced utilized within the currently assigned Extended Overflow blocks.
Extended PE	The number of Extended blocks assigned as Prime Extension (PE) blocks.
Extended Allocated	The maximum number of extended area blocks that can be used within the currently allocated space.
Extended Available	The number of extended area blocks that are available for use, which could be assigned to either Extended Overflow, or Extended PE.
Space Used	The amount of DASD space currently required for the file. (Does not include space required for Extended Available blocks.)
Space Allocated	For each volume, indicates the amount of DASD space allocated, in tracks.
Total Space Allocated	Total DASD space allocated, for all volumes.
Secondary Space	Indicates the amount of space to be requested when a secondary extent is required, as specified on the file define.
Max Secondary	The maximum amount of DASD space IAM will request, when additional DASD space is needed for this data set.
Share Options	Indicates the defined cross region share option.
Index Compress	Indicates whether or not a compressed index structure exists for this file.
Inserts	Indicates the number of records added since the file was last loaded.
Deletes	Indicates the number of records deleted since the file was last loaded.
High Allocated RBA	Total number of bytes of DASD storage allocated to the file.

01.60 CONVERSION CONSIDERATIONS

IAMINFO OUTPUT

There are some revisions to the IAMINFO reports for IAM Version 6.3 Enhanced Format Files. One of the major areas changed has been the IAM Command Execution Summary. The command counters are now based on the VSAM macro request types. It is hoped that this will improve understanding as to what types of requests are being issued as the file is being processed. The Overflow Area Characteristics section, as identified by message IAM364, has been replaced with the IAM Extended Area Characteristics, as identified by message IAM372. Also revised is the IAM361 message. The program name and the type of processing have been revised to better identify the function being performed, and the application program performing the processing.

OBTAINING IAMINFO REPORTS

IAMINFO reports can easily be obtained with the job output for any job step that utilizes IAM files. Simply by adding an IAMINFO DD card to the JCL, IAM will automatically format out a report whenever an IAM file is closed. Normally, this output is assigned to SYSOUT, however it can also be kept in a DASD or tape data set. When having the IAMINFO reports go to a data set, do not code any DCB information, as IAMINFO will supply the appropriate specifications for the report. Below is an example of an IAMINFO DD card, which will have the output go to SYSOUT, to be later printed or viewed.

```
//mypgm      EXEC  PGM=mypgm
//IAMINFO    DD    SYSOUT=*
//iamfile    DD    DSN=my.iamfile,DISP=SHR
```

Figure 31: Obtaining IAMINFO Reports with the IAMINFO DD Card

IAMINFO reports can also be obtained by using the IAMSMF program, supplied with the product. To use this capability, the IAM Global Options must be changed to enable IAM SMF recording, and to specify the SMF record type. Also, make sure that the specified SMF record type are being saved by SMF, both in the active SMF file, and also in being accumulated by the SMF dump program. Below is an example of enabling the IAM SMF recording. IAM will create an SMF record whenever an IAM file is closed. Installations can also produce their own unique reports from the IAM SMF records. A layout of the IAM SMF record is included in the IAM ICL (Installation Control Library).

```
//IAMZAPOP   EXEC  PGM=IAMZAPOP
//STEPLIB    DD    DSN=my.iam63.loadlib,DISP=SHR
//SYSLIB     DD    DSN=my.iam63.loadlib,DISP=SHR
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
              ZAP  SMF=YES,RECTYPE=xxx
              PRINT
/*
```

Figure 32: Enabling the IAM SMF Records

In the example above, you must specify a unique SMF record type, between the values of 128 - 255. Before selecting a record type, make sure that it is not being used by your installation by any other software products, or for your own SMF records. Refer to the appropriate MVS SMF manual for the level of MVS that you are using for complete information on enabling the recording and dumping of the IAM SMF record type that has been selected.

01.60 CONVERSION CONSIDERATIONS

Once the above tasks have been completed, IAMINFO reports can be obtained through the use of the IAMSMF program. The IAMSMF program's IAMINFO command can be used to request IAMINFO reports. This command provides various selection criteria keywords, which can be used if desired to limit the output only to the desired reports. Full information on the IAMSMF IAMINFO command is in [Section 41.10 of the IAM Version 6.3 User Manual](#). The following example below will print out an IAMINFO report for each IAM record contained within the input data set.

```

//IAMINFO      EXEC  PGM=IAMSMF
//SYSPRINT    DD    SYSOUT=*
//SYSMF       DD    DISP=SHR,DSN=my.smf.data
//SYSIN       DD    *
               IAMINFO
/*

```

Figure 33: Obtaining IAMINFO Reports from SMF Data

IAMINFO Report from a file Load

```

IAM400 INNOVATION IAM INFORMATION PRINT ROUTINE -- IAMNINFO VER 6.3/01P -- INNOVATION DATA
PROCESSING DATE - 96.032
IAM360 STEP - DCV201B DDNAME - VSAMCRT1 DATA SET MONITORED - IAMV.DCV201.CLUSTER
IAM361 INFO REQUESTED BY IAMTVSAM PROGRAM PERFORMING FILE CREATION OPENED-96.032.13:57:54
CLOSED-96.032.13:57:57
IAM362 IAM DATA CHARACTERISTICS -
LOGICAL RECORD LENGTH----- 64 - CI SIZE----- 4096
KEY SIZE----- 4 - KEY OFFSET----- 12
AVERAGE RECORD LENGTH----- 46 - LARGEST RECORD PROCESSED-- 64
IAM363 IAM FILE CHARACTERISTICS -
BLOCKING FACTOR----- 4 - BLOCK SIZE----- 11476
TRACKS IN USE----- 6 - DELETE BYTE SUPPORT----- NONE
DATASET TYPE----- KSDS - SHARE OPTIONS----- 2
NUMBER OF IAM DATA BLOCKS-- 23 - HIGH ALLOCATED RBN----- 0
INTEGRATED OVERFLOW (CI%)-- 10 - DASD RESERVE (CA%)----- 10
FILE DEFINED DATE----- 1996.032 - FILE DEFINED TIME----- 13:57:50
FILE LOADED DATE----- 1996.032 - FILE LOADED TIME----- 13:57:57
IAM372 IAM EXTENDED AREA CHARACTERISTICS -
EXT. OVERFLOW RECORDS----- 0 - EXT. OVERFLOW BLOCKS----- 0
EXTENDED BLOCKS ALLOCATED-- 0 - EXTENDED PE BLOCKS----- 0
EXTENDED BLOCKS USED ----- 0 - EXTENDED BLOCKS AVAILABLE-- 0
IAM365 IAM EXECUTION STATISTICS -
TOTAL STORAGE REQUIRED----- 759344 - PRIME INDEX----- 84
STORAGE ABOVE THE LINE---- 739328 - COMPRESSED DATA STRUCTURE-- YES
REQUESTS PROCESSED----- 4098 - REQUESTS FAILED----- 0
DISK BLOCKS READ----- 0 - DISK BLOCKS WRITTEN----- 2
DYNAMIC BUFFER RETRIEVALS-- 0 - MAXIMUM BUFFERS USED----- 60
MINIMUM BUFFERS USED----- 0 - MAXIMUM BUFFERS AVAILABLE-- 0
DATA SPACE USED(M BYTES)-- 1 - DATA SPACE SIZE(M BYTES)-- 128
IAM366 IAM COMMAND EXECUTION SUMMARY -
GET RANDOM----- 0 - PUT UPDATE----- 0
GET SEQUENTIAL----- 0 - PUT ADD----- 4096
GET PREVIOUS----- 0 - POINT (START BROWSE)----- 0
GET KGE/GENERIC----- 0 - POINT KGE/GENERIC----- 0
GET (SKIP SEQUENTIAL)----- 0 - ERASE----- 0
ENDREQ----- 0 - WRTBFR----- 0
IAM STATISTICS----- 0 - IAM FLUSH BUFFER----- 0
CLOSE----- 1 - OPEN----- 1
CLOSE TYPE=T----- 0 - VERIFY----- 0
INVALID REQUESTS----- 0 - RECORD LENGTH CHANGES----- 0
SEQ CHAINED BLOCKS READ---- 0 - SEQ CHAINED BLOCKS WRITTEN= 23

```

Figure 34: IAMINFO Report for a File Load

01.60 CONVERSION CONSIDERATIONS

IAMIFNO Report for File Update

IAM400	INNOVATION IAM INFORMATION PRINT ROUTINE -- IAMNINFO VER 6.3/01P -- INNOVATION DATA		
	PROCESSING DATE - 96.032		
IAM360	STEP - DCV201B	DDNAME - VSAMCRT1	DATA SET MONITORED - IAMV.DCV201.CLUSTER
IAM361	INFO REQUESTED BY IAMTVSAM PROGRAM PERFORMING UPDATE PROCESSING OPENED-96.032.13:58:00		
	CLOSED-96.032.13:58:10		
IAM362	IAM DATA CHARACTERISTICS -		
	LOGICAL RECORD LENGTH-----	64	- CI SIZE-----= 4096
	KEY SIZE-----	4	- KEY OFFSET-----= 12
	TOTAL RECORDS-----	4096	- TOTAL RECORDS DELETED-----= 0
	TOTAL RECORDS UPDATED-----	1024	- TOTAL RECORDS INSERTED-----= 0
IAM363	IAM FILE CHARACTERISTICS -		
	BLOCKING FACTOR-----	4	- BLOCK SIZE-----= 11476
	TRACKS IN USE-----	7	- DELETE BYTE SUPPORT-----= NONE
	DATASET TYPE-----	KSIDS	- SHARE OPTIONS-----= 2
	NUMBER OF IAM DATA BLOCKS--	23	- HIGH ALLOCATED RBN-----= 59
	INTEGRATED OVERFLOW (CI%)--	10	- DASD RESERVE (CA%)-----= 10
	FILE DEFINED DATE-----	1996.032	- FILE DEFINED TIME-----= 13:57:50
	FILE LOADED DATE-----	1996.032	- FILE LOADED TIME-----= 13:57:57
	FILE UPDATE DATE-----	1996.032	- FILE UPDATE TIME-----= 13:58:09
IAM372	IAM EXTENDED AREA CHARACTERISTICS -		
	EXT. OVERFLOW RECORDS-----	79	- EXT. OVERFLOW BLOCKS-----= 1
	EXTENDED BLOCKS ALLOCATED--	36	- EXTENDED PE BLOCKS-----= 0
	EXTENDED BLOCKS USED -----	3	- EXTENDED BLOCKS AVAILABLE--= 33
IAM365	IAM EXECUTION STATISTICS -		
	TOTAL STORAGE REQUIRED-----	167936	- PRIME INDEX-----= 84
	STORAGE ABOVE THE LINE-----	159744	- COMPRESSED DATA STRUCTURE--= YES
	REQUESTS PROCESSED-----	18439	- REQUESTS FAILED-----= 3
	DISK BLOCKS READ-----	92	- DISK BLOCKS WRITTEN-----= 28
	DYNAMIC BUFFER RETRIEVALS--	5650	- MAXIMUM BUFFERS USED-----= 10
	MINIMUM BUFFERS USED-----	1	- MAXIMUM BUFFERS AVAILABLE--= 22
	DYNAMIC TABLE RETRIEVALS--	0	- DYNAMIC TABLE RECORDS-----= 0
IAM366	IAM COMMAND EXECUTION SUMMARY -		
	GET RANDOM-----	5120	- PUT UPDATE-----= 1024
	GET SEQUENTIAL-----	8194	- PUT ADD-----= 0
	GET PREVIOUS-----	4097	- POINT (START BROWSE)-----= 2
	GET KGE/GENERIC-----	0	- POINT KGE/GENERIC-----= 0
	GET (SKIP SEQUENTIAL)-----	0	- ERASE-----= 0
	ENDREQ-----	0	- WRFBFR-----= 0
	IAM STATISTICS-----	0	- IAM FLUSH BUFFER-----= 0
	CLOSE-----	1	- OPEN-----= 1
	CLOSE TYPE=T-----	0	- VERIFY-----= 0
	INVALID REQUESTS-----	0	- RECORD LENGTH CHANGES-----= 835
	SEQ CHAINED BLOCKS READ----	28	- SEQ CHAINED BLOCKS WRITTEN= 0

Figure 35: IAMINFO Report for File Access

01.60 CONVERSION CONSIDERATIONS

IAMINFO Report Field Descriptions

The IAM Run Time Statistics (INFO) Report is produced when the IAM file is closed. There will be one report produced for each OPEN/CLOSE combination a program issues for an IAM file.

IAMINFO Report Heading Lines

IAM400 The IAMINFO Report title containing a module name and version.

IAM360 Provides the name of the job step in execution, the DDNAME and data set name of the IAM file being processed.

IAM361 This identifies the name of the application program that processed the IAM file, the type of processing being performed, and the date and time the IAM file was opened and closed. The type of processing for Enhanced Format files are FILE CREATION, INPUT, or UPDATE processing.

IAMINFO Data Characteristics

IAM362 This heading line always precedes the IAM file Data Characteristics portion of the IAMINFO report. The following fields describing the logical record are provided.

LOGICAL RECORD LENGTH	The maximum user data length of the record.
CI SIZE	The control interval (CI) size specified when file was defined.
KEY SIZE	The length of the key, specified when file was defined.
KEY OFFSET	The relative position of the key within the user data record. (RKP)
AVERAGE RECORD LENGTH	During a load operation IAM calculates an average length for the records that were loaded into the file. For data compressed files, the length is after compression. NOTE: This field only appears when a file is LOADED.
LARGEST RECORD PROCESSED	During a load operation IAM reports the length of the largest record that was written to the file. NOTE: This field only appears when a file is LOADED.
TOTAL RECORDS	The total number of records in the IAM file. NOTE: This field does not appear when the file is initially created.
TOTAL RECORDS DELETED	The total number of records deleted from the IAM file since creation. NOTE: This field does not appear when the file is initially created.
TOTAL RECORDS UPDATED	The total number of records updated in the IAM file since creation. NOTE: This field does not appear when the file is initially created.
TOTAL RECORDS INSERTED	The total number of records inserted into the IAM file since creation. NOTE: This field does not appear when the file is initially created.

Figure 36: IAM Data Characteristics Fields (IAM362)

01.60 CONVERSION CONSIDERATIONS

IAMINFO File Characteristics

IAM363 This heading line always precedes the IAM File Characteristics portion of the IAMINFO report. The following fields describing the physical file are provided.

BLOCKING FACTOR	A value of 1 to 15 is the number of blocks per track (as blocking factor). A larger value indicates the requested block size from the IAM CREATE override.
BLOCKSIZE	Physical blocksize of the IAM file. This is the actual blocksize developed and used by IAM. NOTE: IAM will determine the optimal blocksize for a file based upon the file's record length, the blocking factor requested and the track capacity of the device. The IAM blocksize is transparent to application programs even when the blocksize is altered or the file is relocated to a different device type.
TRACKS IN USE	Number of tracks currently being used by IAM for the file. For Enhanced Format files, which have Extended blocks, this value reflects only the tracks that are occupied by used extended data blocks.
DELETE BYTE SUPPORT	Indicates if ISAM delete byte support is active for the file. This will always be NONE for Enhanced Format files.
DATASET TYPE	Describes the type of data set IAM is simulating. Possible values include KSDS, ESDS, or ESDS/P. The latter is an ESDS type of file with PSEUDORBA specified.
SHARE OPTIONS	Indicates the defined cross region share option for this file.
NUMBER OF IAM DATA BLOCKS	The number of blocks in the file up to where the prime index begins. This field may be needed to run IAMRECVR for recovery if the first block in the file has been damaged.
HIGH ALLOCATED RBN	This the highest allocated block number currently in the IAM file. This field may be needed to run IAMRECVR for recovery if the control information about the file has been damaged.
INTEGRATED OVERFLOW (CI%)	The amount of space, as a percentage, left free in each prime data block as the file is being loaded. The space can subsequently be used for file expansion. This value is specified and is similar in concept to CI Freespace.
DASD RESERVE (CA%)	Indicates the CA Freespace value that was specified when the file was defined. This is used to by IAM to reserve some space for future expansion during automatic space release, which occurs during the first file load.
FILE DEFINED DATE / TIME	The date and time that the file was defined, as taken from the TOD stamp in the file header.
FILE LOADED DATE / TIME	The date and time of the completion of the last file load or reorganization for this file.
FILE UPDATE DATE / TIME	The date and time the file was last closed by a program that updated the file.

Figure 37: IAM File Characteristics Fields (IAM363)

01.60 CONVERSION CONSIDERATIONS

IAMINFO Extended Area Characteristics

IAM372 This heading line always precedes the IAM File Extended area characteristics portion of the IAMINFO report. The following fields describe the extended data and overflow areas that were obtained automatically as needed.

EXT. OVERFLOW RECORDS	Number of records currently contained in extended overflow blocks.
EXT. OVERFLOW BLOCKS	Number of extended overflow blocks.
EXTENDED BLOCKS ALLOCATED	Number of extended overflow blocks that will fit within the current amount of DASD space allocated to the file.
EXTENDED PE BLOCKS	Number of extended blocks assigned as Prime Extension (PE) blocks.
EXTENDED BLOCKS USED	Actual number of extended overflow blocks in use. This includes for Overflow, PE, and Extended Index blocks.
EXTENDED BLOCKS AVAILABLE	The number of extended blocks that are available for use within the currently allocated DASD space.

Figure 38: IAM Extended Area Characteristics Fields (IAM372)

01.60 CONVERSION CONSIDERATIONS

IAMINFO Execution Statistics

IAM365 This heading line always precedes the IAM Execution Statistics portion of the IAM INFO report. The following fields describing system processing and IAM options used during file processing are provided.

TOTAL STORAGE REQUIRED	Total storage acquired by IAM during this execution.
PRIME INDEX [COMPRESSED]	Indicates the amount of virtual storage required for the index to the prime data area of this file. Also, if the index is in a compressed format, the word COMPRESSED appears.
STORAGE ABOVE THE LINE	Total amount of storage acquired above the 16MB line.
COMPRESSED DATA STRUCTURE	Indicates whether data compression was used for this file.
REQUESTS PROCESSED	Number of requests made against the file, since last OPEN, by the application (a breakdown by command type follows in the command execution summary).
REQUESTS FAILED	Number of requests made to IAM which did not complete normally (i.e. EOF on sequential read, no record found on random read)
DISK BLOCKS READ	Number of physical I/O's used to read blocks from the file.
DISK BLOCKS WRITTEN	Number of physical I/O's used to write blocks to the file.
DYNAMIC BUFFER RETRIEVALS	Number of EXCPs saved (not executed). IAM's Real Time Tuning buffer management was able to satisfy this number of I/O requests from buffers that it had maintained in storage, other than its current buffer. These buffer retrievals eliminated the need for I/O to disk.
MAXIMUM BUFFERS USED	The maximum number of data buffers acquired by IAM's Real Time Tuning during this execution. This may be accompanied by an IAM367 or IAM368 informational message indicating that if IAM was allowed to acquire more buffers the number of I/O's required to service the requests against the file could have been reduced. Under MVS/XA and MVS/ESA all buffers are acquired above the 16MB line.
DYNAMIC TABLE RETRIEVALS	Number of random read requests satisfied by IAM's Dynamic Tabling of data records in virtual. Under MVS/XA and MVS/ESA IAM's Dynamic Table (DYNCORE) is maintained in virtual storage above the 16MB line.
DYNAMIC TABLE RECORDS	The total number of data records IAM placed into the Dynamic Table.
DATA SPACE USED(M BYTES)	The amount of area actually used in the data space during file load.
DATA SPACE SIZE(M BYTES)	The size of the data space created during file load.

Figure 39: IAM Execution Statistics Fields (IAM365)

01.60 CONVERSION CONSIDERATIONS

IAMINFO Command Execution Summary

IAM366 indicates the Command Execution Summary. For Enhanced Format files, the section provides the counts of the various types of VSAM macro requests that can be issued for an IAM file.

GET RANDOM	Number of random READ requests with an exact key specified.
PUT UPDATE	Number of UPDATE requests.
GET SEQUENTIAL	Number of sequential GET requests.
PUT ADD	Number of INSERT requests.
GET PREVIOUS	Number of GET PREVIOUS requests.
POINT (START BROWSE)	Number of POINT requests.
GET CI (SEQUENTIAL)	Number of sequential control interval GETS. (ESDS only)
PUT CI (UPDATE)	Number of CONTROL INTERVAL UPDATES (ESDS only)
GET CI (RANDOM)	Number of CONTROL INTERVAL GETS (ESDS only)
PUT CI (ADD)	Number of CONTROL INTERVAL ADDS (ESDS only)
GET KGE/GENERIC	Number of random or skip sequential reads, indicating search for key greater or equal, and /or partial key search. (KSDS only)
POINT KGE/GENERIC	Number of POINT, or START BROWSE, requests indicating a search for key greater or equal, or partial key search. (KSDS only)
GET (SKIP SEQUENTIAL)	Number of GET's issued in Skip Sequential mode, with an exact key specified. (KSDS only)
ERASE	Number of requests to delete records. (KSDS only)
ENDREQ	Number of LSR ENDREQ requests
WRTBFR	Number of LSR WRTBFR requests
IAM STATISTICS	Number of IAM statistic requests
IAM FLUSH BUFFER	Number of IAM flush buffer requests
CLOSE	Number of CLOSE requests
OPEN	Number of OPEN requests
CLOSE TYPE=T	Number of TEMPORARY CLOSE requests
VERIFY	Number of VERIFYs
INVALID REQUESTS	Number of requests with RPL error's
RECORD LENGTH CHANGES	Number of records with length changes
SEQ CHAINED BLOCKS READ	Number of chained buffers during READ
SEQ CHAIN WRITTEN	Number of chained buffers during WRITE

Figure 40: IAM Command Execution Summary Fields (IAM366)

01.60 CONVERSION CONSIDERATIONS

Additional Information Messages

Additional messages may also be present within the report, primarily to indicate if there are some additional tuning considerations, which may improve overall file performance.

IAM367 INSUFFICIENT STORAGE AVAILABLE TO ACQUIRE ADDITIONAL BUFFERS. This is an informational message provided by IAM's Real Time Tuning buffer management. It is provided when an attempt to acquire additional data buffers fails because of insufficient extended private virtual storage. For Enhanced Format files, if MVS returns 24-bit addressable storage for a buffer, IAM will freemain that storage, to prevent virtual storage problems. IAM acquires additional buffers when it had determined that this would improve performance by reducing physical I/O. Increasing the region size will result in fewer disk I/O's to process this file.

IAM368 BUFNO VALUE GREATER THAN nn MAY IMPROVE PERFORMANCE. This informational message is provided whenever IAM determines that a greater number of buffers than provided would result in fewer disk I/O's to process this file. The MAXBUFNO= parameter on the IAM Override statement can be used to increase the number of buffers IAM will use. If the actual number of DISK BLOCKS READ and DISK BLOCKS WRITTEN is small there may be little benefit to be gained.

01.60 CONVERSION CONSIDERATIONS

ENHANCED FILE STRUCTURE TRACE OUTPUT

When processing enhanced IAM KSDS and ESDS files, trace entries can be generated to help with problem determination. This section describes the information generated by the different types of trace requests. The types of activity that can be traced are I/O Start, I/O End, Buffer Manager calls, EXCP requests, and file Extend requests. The I/O Start and I/O End will be of most benefit for tracing requests passed to IAM by application programs, as they indicate the actual VSAM macro requests. The other trace entries may be requested by the IAM Technical Support Department for problem diagnosis

Obtaining Trace Output

IAM Trace Output is obtained through specifying IAM Overrides. On the ACCESS override, indicate the Trace DD name, with the TRACEDD parameter, and the types of entries to trace with the TRACEREQ= keyword. The trace data is normally a SYSOUT file, however it can be placed on DASD or TAPE, and later printed entirely, or selectively with the IAMTSEL, which is documented in [Section 64](#) of the IAM Version 6.3 User Manual. Below is an example of how to specify the overrides to obtain IAM Trace Output:

```
//UPDATE      EXEC  PGM=my p gm
//IAMFILE     DD    DSN=my . iam . file , DISP=OLD
//IAMINFO     DD    SYSOUT=*
//IAMTRPRX    DD    SYSOUT=*
//IAMOVRID    DD    *
              ACCESS DD=IAMFILE , TRACEDD=IAMTRPRX ,
              TRACEREQ=(TRIOS , TRIOE , TRBFR , TREXCP , TRXTND )
/*
```

Figure 41: JCL Example for Obtaining IAM Trace Data

The TRACEREQ keywords, as shown in the above example, indicate the following types of trace records:

TRIOS	Logical I/O Start (VSAM Macro Request)
TRIOE	Logical I/O End
TRBFR	Internal IAM Call for Buffer Management
TREXCP	Physical I/O Request
TRXTND	Internal call for more DASD space

Figure 42: TRACEREQ Keywords

01.60 CONVERSION CONSIDERATIONS

Sample IAM Trace Output

Innovation Access Method (c) Trace VER 6.3/00P									
Date: 96.016 Page: 000001									
DDNAME: VSAMCRT2 DSN: IAMV.VIT132.CLUSTER									
Type	Time	Request	RPL/RBN	PLH	OPTCD	RC/RINFO	RECLEN	Key	
IAM700									
IAM701									
IAM702									
IAM703									
IAM712	BFR 16.32.09.6100	READ	00000001	03B06190	0000				00000000
IAM713	EXCP 16.32.09.6100	READ	00000001	03B06190	0001				
IAM712	BFR 16.32.09.6600	READ	00000074	03B06190	0000				00000000
IAM713	EXCP 16.32.09.6600	READ	00000074	03B06190	0001				
IAM712	BFR 16.32.09.7200	RELEASE	00000074	03B06190	0004				00000000
IAM710	IOS 16.32.09.7300	POINT	0002D508	00032928	2098		0000		00000000
IAM712	BFR 16.32.09.7300	READ	00000073	00032928	0000				00000000
IAM713	EXCP 16.32.09.7300	READ	00000073	00032928	0001				
IAM711	IOE 16.32.09.7600	POINT	0002D508	00032928		00000000	0000		00003FFF
IAM710	IOS 16.32.09.8000	GET	000242E8	00032B4C	4082		0924		0000000F
IAM712	BFR 16.32.09.8000	READ	00000002	00032B4C	0000				0000000F
IAM713	EXCP 16.32.09.8000	READ	00000002	00032B4C	0001				
IAM711	IOE 16.32.09.8200	GET	000242E8	00032B4C		00000000	0924		0000000F
IAM710	IOS 16.32.09.8200	GET	0002C508	00032D70	2080		0000		00000000
IAM712	BFR 16.32.09.8300	READ	00000002	00032D70	1000				00000000
IAM711	IOE 16.32.09.8300	GET	0002C508	00032D70		00000000	0924		0000000F
IAM710	IOS 16.32.09.8300	GET	0002D508	00032928	2098		0000		00000000
IAM711	IOE 16.32.09.8300	GET	0002D508	00032928		00000000	0924		00003FFF
IAM710	IOS 16.32.09.8300	GET	0002E508	00032F94	4081		0000		0000080F
IAM712	BFR 16.32.09.8300	READ	00000010	00032F94	0000				0000080F
IAM713	EXCP 16.32.09.8300	READ	00000010	00032F94	0001				
IAM711	IOE 16.32.09.8500	GET	0002E508	00032F94		00000000	0924		0000080F

VSAM Request I/O Start Trace Entry (TRIOS)

Heading	Format	Description
TYPE	IOS	Identifies the trace type as a logical I/O Start request.
TIME	99.99.99.9999	Time the trace record was generated.
REQUEST	GET	Type of request as found in the RPL. The values can be GET, PUT, CHECK, POINT, ENDREQ, SRCHBFR, MRKBFR or WRTBFR.
RPL/RBN	99999999	Address of RPL used to process the request.
PLH	99999999	Address of Placeholder (PLH) used to process the request.
OPTCD	9999	RPL OPTCD values used during the request.
RC/RINFO	N/A	Not used by IOS.
RECLEN	9999	Length of record being processed.
KEY	AAAAAAA....	First 16 bytes of KEY if KSDS, or the RBA of the ESDS record.

Figure 43: Trace – IOS Entry Field Descriptions

01.60 CONVERSION CONSIDERATIONS

VSAM Request I/O End Trace Entry (TRIOE)

Heading	Format	Description
TYPE	IOE	Identifies the trace type as an I/O End request.
TIME	99.99.99.9999	Time the trace record was generated.
REQUEST	GET	Type of request as found in the RPL. The values can be GET, PUT, CHECK, POINT, ENDREQ, SRCHBFR, MRKBFR or WRTBFR.
RPL/RBN	99999999	Address RPL used to process the request.
PLH	99999999	Address Placeholder (PLH) used to process the request.
OPTCD	9999	RPL OPTCD used during the request.
RC/RINFO	99999999	Return code.
RECLEN	9999	Length of record being processed.
KEY	AAAAAAA....	First 16 bytes of KEY if KSDS, or the RBA of the ESDS record.

Figure 44: Trace – IOE Entry Format (I/O End)

Buffer Manager Call (BFR) Trace Output (TRBFR)

Heading	Format	Description
TYPE	BFR	Identifies the trace type as a Buffer Manager Call request.
TIME	99.99.99.9999	Time the trace record was generated.
REQUEST	READ	Type of Buffer Manager Request. The values can be READ, WRITE, RELEASE or FLUSH.
RPL/RBN	99999999	Relative Block Number (RBN) of the block being processed.
PLH	99999999	Address of Placeholder (PLH) used to process the request.
OPTCD	9999	Status information within the PLH.
RC/RINFO	OV	OV if block being processed is an Overflow block.
RECLEN	N/A	Not used by BFR.
KEY	99999999	Relative Block Number of block within file.

Figure 45: Trace – BFR Entry Format (Buffer Manager Call)

01.60 CONVERSION CONSIDERATIONS

Physical I/O (EXCP) Trace Output (TREXCP)

Heading	Format	Description
TYPE	EXCP	Identifies the trace type as an EXCP request.
TIME	99.99.99.9999	Time the trace record was generated.
REQUEST	READ	Type of EXCP Request. The values can be READ, WRITE, PRFRM, WREOF, WRADD, or UNKN.
RPL/RBN	99999999	Relative Block Number (RBN) of the 1st (or only) block being processed.
PLH	99999999	Address Placeholder (PLH) used to process the request.
OPTCD	9999	Number of blocks being processed by request.
RC/RINFO	N/A	Not used by EXCP.
RECLN	N/A	Not used by EXCP.
KEY	N/A	Not used by EXCP.

Figure 46: Trace EXCP Entry Format

File Extend Request Trace Output (TRXTND)

Heading	Format	Description
TYPE	XTND	Identifies the trace type as an Extend request.
TIME	99.99.99.999	Time the trace record was generated.
REQUEST	EXTEND	Type of Extend processing that was performed. This can be EXTEND, EXPAND, EXTEXP or EXPSRL.
RPL/RBN	99999999	This is the Relative Block Number (RBN) of the first available new block.
PLH	99999999	The number of new index blocks required for the new area.
OPTCD	9999	The first two digits represent the number of VOLUMES that the IAM file occupied BEFORE any new areas or extents were made available. The second two digits represent the number of EXTENTS in the IAM file BEFORE any new areas or extents were made available.
RC/RINFO	99999999	Extend request Return Code.
RECLN	9999	Number of TRACKS obtained in new extent (in HEX).
KEY	99999999	Extend processor flag bytes.

Figure 47: File Extend Trace Entry Format

SECTION 02: IAM TECHNICAL DESCRIPTION

02.00 IAM TECHNICAL DESCRIPTION

02.10 IAM'S VSAM TRANSPARENCY

IAM's system level VSAM Interface (VIF) provides transparency. VIF allows an unaltered application program executing under MVS, MVS/XA and MVS/ESA to access IAM files in place of single index VSAM KSDS or VSAM ESDS files. VIF can be used in conjunction with the common programming languages COBOL, assembler, PL/1, RPG, etc. and any higher level language products which support keyed access to VSAM files.

CICS V/S, COBOL II, and a number of other new program products, restrict key indexed processing to the VSAM ACCESS METHOD. VIF allows these program products to use IAM files in place of VSAM files without modification.

IAM supports programs executing in AMODE(31) and VSAM control blocks (ex: ACB, RPL) residing above the 16MB line.

VIF supports the full range of VSAM file access commands GET, PUT, INSERT, GETPREV, ERASE, POINT, etc. and the file status commands SHOWCB, TESTCB, GENCB, and the VSAM catalog lookup macro, SHOWCAT.

IAM supports the following functions of IDCAMS, as they relate to VSAM KSDS or ESDS file processing: DEFINE, DELETE, LISTCAT, PRINT, REPRO and VERIFY. DELETE, PRINT, REPRO and VERIFY provide the same services for IAM files as they would VSAM clusters.

IDCAMS DEFINE will create an IAM file whenever the OWNER(\$IAM) parameter is specified, '\$IAM' is placed somewhere in the data set name, or \$IAM is part of the Data Class or Storage Class name.

LISTCAT ALL displays IAM files as non-VSAM in its standard SYSPRINT report. LISTCAT ALL also displays the file's IAM characteristics in an IAMPRINT DD report, which will be dynamically allocated if necessary.

[Section 10](#) of this manual documents IAM's support of IDCAMS.

IAM has full support for the SMS environment. This support includes recognizing and using the SMS classes for allocation, honoring file attributes specified in the Data Class, and support for JCL allocation of IAM files; including temporary data set support. Simply place \$IAM in the SMS Data Class or Storage Class name to define a file as IAM. When using the SMS JCL allocation feature specifying \$IAM in the data set name, in the Storage Class name or in the Data Class name on the DD Statement results in an IAM file being allocated.

Activating VIF is a simple procedure. While evaluating IAM, you can activate and deactivate VIF at any time. To start VIF all you need to do is submit the procedure, 'VIFSTART', supplied in the IAM Installation Control Library. Within a few seconds VIF will be active in the system. Once testing has been completed and IAM is in production, you can activate VIF automatically each time the system is IPL'd.

[Section 90.04](#) of this manual documents the VIF Interface.

02.20 IAM FILE ALLOCATION AND SPACE SAVINGS

02.20 IAM FILE ALLOCATION AND SPACE SAVINGS

IAM data files can be allocated on all DASD devices supported by IBM's MVS/XA, and MVS/ESA operating systems. IAM uses a single non-VSAM data set on disk. This single file contains IAM's data records, index and file description.

Non-SMS managed VSAM files must be allocated to specific volumes. IAM files can be allocated using non-specific allocation. (UNIT= see Section 11).

IAM files can be cataloged into any VSAM or ICF catalog. IAM files are cataloged as non-VSAM entries and require a fraction of the space needed to catalog a VSAM cluster. A LISTCAT ALL for an IAM file will simply show it as a non-VSAM data set. However, just by adding an IAMPRINT DD statement to the IDCAMS JCL, LISTCAT displays the characteristics of the IAM file. Unlike VSAM clusters, IAM files are not normally recorded in the VVDS. Under SMS, IAM files are recorded as non-VSAM data sets.

IAM uses an advanced file structure which is far superior to VSAM, and requires 30 to 70% less DASD space than VSAM.

PROBLEM: VSAM's INDEX structure requires a disproportionate amount of the allocated disk space for the following reasons:

1. VSAM's imbedded index occupies one (1) track in every Control Area (CA) of the cluster. For example, an imbedded index for a VSAM cluster, on a 3380, with a 15 track CA takes up 7% of the cluster's allocated space and most of the space on the imbedded index track is wasted.
2. VSAM's index compression technique does not work well for certain types of keys (ex: all numeric) because it only compresses repeating characters from the beginning and end of the key.
3. VSAM's index contains a sizable amount of control information.

SOLUTION: IAM's index requires a fraction of the disk space taken by VSAM for the following reasons:

1. IAM, by loading its entire index into memory eliminates the need for an imbedded index.
2. IAM's index compression technique scans the entire key.
3. IAM's index requires minimal control information to be stored on the disk.

PROBLEM: VSAM's physical blocksize restrictions waste disk space. VSAM's Control Intervals and physical blocks must be a multiple of 512 or 2048 bytes.

SOLUTION: IAM's blocksizes are not restricted. IAM fully utilizes each track. Additionally, IAM automatically calculates the most efficient blocksize for a device.

02.20 IAM FILE ALLOCATION AND SPACE SAVINGS

PROBLEM: VSAM wastes Control Intervals (CI) within every Control Area (CA) in a file, when the key is large or not easily compressed (ex: keys with all numeric data).

An Index Control Interval (CI) must contain all the keys for a Data Control Area. VSAM discontinues loading records into the Data CA when the Index CI is filled. Any remaining CIs in the Data CA are wasted.

SOLUTION: IAM files are not divided into Control Areas. IAM is not affected by large or incompressible keys because the Index is seamless.

PROBLEM: VSAM's FREESPACE concepts waste disk space. VSAM offers two types of FREESPACE: The percentage of each Control Interval (CI%) to be left free and the percentage of the CIs in each CA to be left free (CA%).

VSAM's concept of free CIs within a CA (CA%) can only be used for inserts into that Control Area. Mass inserts to one location will cause the file to split to new CAs. VSAM's CA% usually wastes space as the FREESPACE exists across the entire file and it is very rare that inserts are applied evenly across an entire file.

SOLUTION: IAM's FREESPACE techniques make much more efficient use of disk space. IAM's Integrated Overflow is conceptually the same as CI% free. However, since IAM blocks are usually larger than VSAM, each Integrated Overflow area accepts more records.

In addition to Integrated Overflow, IAM also provides an overflow area that is independent of the individual blocks. An insert to any prime data block which does not contain sufficient free space is placed into this record based overflow area. Since IAM's overflow can accept inserts from anywhere on the file, a much smaller amount of disk space is needed as compared to VSAM's CA%. For Enhanced format files, this area is referred to as Extended Overflow, which is acquired on an as needed basis. For Compatible format files, this area is referred to as Independent Overflow, which is allocated and preformatted when the file is loaded.

This advanced structure is one reason why IAM, even without data compression, takes 20 to 40% less DASD space than VSAM

An estimate of the disk space saving that a conversion of your VSAM clusters to IAM will return is available using the IAM VSAM Space Savings Analysis program IAMSIMVS.

02.30 DATA COMPRESSION

02.30 DATA COMPRESSION

IAM can compress the data in its files and it is completely transparent to the programs that create and use those files. Most data records contain unused fields and repeating sets of characters. IAM compresses all of the data in a record following the key. Enabling IAM's proprietary Data Compression option typically gains an additional 20 to 50% reduction in file size over the 20% to 40% savings that results from simply converting to IAM.

Data compression not only provides disk space saving it also reduces overall processing time. IAM's Data Compression is unique because it reduces the amount of data transferred to and from the disk without materially increasing CPU overhead.

IAM's CPU time is dramatically less than other compression product's. In fact, IAM's CPU time is typically so much less than VSAM's, IAM using data compression in most circumstances will still use less CPU time than normal VSAM processing.

DATA COMPRESSION CONSID- ERATIONS

For ESDS type of files, IAM defaults to always compressing the records. For KSDS type of files, IAM can be enabled to automatically compress a file based on its size. With Auto-Compression, IAM will compress records in any IAM file whose primary space allocation value is equal to or larger than the 'compress size' specified in the installation's IAM Global Option table ([DATACOMPRESS=size see Section 91](#)).

Data Compression can be requested for any individual file when the file is defined or at execution time when the file is loaded using an IAM CREATE Override statement ([DATACOMPRESS=YES see Section 11](#)). Likewise, Auto-Compression can always be manually suspended for an individual file using the IAM Create Override Statement ([DATACOMPRESS=NO see Section 11](#)).

IAM only considers a file for data compression when its maximum record size is at least 10 bytes more than its key length plus the relative key position (RKP).

IAM will only compress individual records when the data following the record key exceeds ten (10) bytes in length.

If compression would make a record larger than the original, IAM leaves the record uncompressed.

Load and sequential processing of compressed files which contain large record sizes (approaching 1K bytes each) may take more CPU time than VSAM.

Data Compression can be used on all your IAM files. If a particular file is found by IAM to be incompressible, there is no penalty in CPU time to process that file after the load. It is as if compression had never been requested for that file. There may also be a few files which just do not show much of a benefit from data compression. For example, SMP/E CSI files have an average record length that is just a bit larger than their key. When there is not much data to work with, there is little Data compression can do to reduce a file's size. In these cases IAM's Data Compression may show little saving, beyond the space reduction that comes with simply converting to IAM. This, however, should really be of little consequence. If a specific file shows only marginal compression there will likewise be only a marginal increase in IAM's CPU time to process that file.

A report on the estimated disk space saving that a conversion of your VSAM clusters to IAM file's with Data Compression is available using the IAM VSAM Space Savings Analysis program IAMSIMVS. ([See Section 42 – IAM Space Savings Analysis Program](#)).

02.40 IAM AUTOMATIC RELEASE

02.40 IAM AUTOMATIC RELEASE

AUTOMATIC RELEASE Allocations are a PROBLEM because in most cases users never know how many records a file will eventually have to contain. IAM's auto-RELEASE eliminates the PROBLEM of user's over allocating VSAM files.

IAM automatically releases unused space. For Enhanced format files, some of the unused space may be retained for future growth, based on CA% Freespace or the Overflow override. For Compatible format files, when a secondary allocation is present, the unused portion of an IAM file's disk allocation is automatically released after the first load, or REPRO into the file. The space required for the IAM internal freespace areas, Independent Overflow, Integrated Overflow and Prime Extension are not released. These areas are preformatted during a file load, and are considered used space by MVS.

RELEASE is an especially important IAM feature. IAM files generally take 30 to 70% less disk space than VSAM. If the original VSAM IDCAMS DEFINE space allocation values are left unchanged without auto- RELEASE a lot of disk space would continue to go to waste.

RELEASE is IAM's default and like most other IAM defaults it can be changed in the IAM Global Options table using the program IAMZAPOP ([RELEASE= see Section 91](#)).

In any installation, there will always be a few files whose owners do not want their space to be auto-RELEASEd. Auto-RELEASE can always be manually suspended for individual files using the IAM CREATE Override Statement ([RELEASE=NO see Section 11](#)). For Enhanced format files, there may be some benefit to not releasing the space, because it can be used by IAM to accommodate file growth. For Compatible format files, the main advantage to specifying RELEASE=NO is to reserve the disk space for future file expansion. IAM will only use this disk space for reloading/reorganizing the IAM file. It will NOT be used to accommodate additional records after the IAM internal freespace areas are filled.

It is recommended that RELEASE, however, should really be left as the default to be used on all your IAM files.

A report on the amount of over allocated space presently held by VSAM clusters as well as an estimate of the potential disk space saving IAM with RELEASE will return to you is available using the IAM VSAM space Savings Analysis program IAMSIMVS.

02.50 IAM'S REAL TIME TUNING

02.50 IAM'S REAL TIME TUNING

REAL TIME TUNING Real Time Tuning buffer management is a concept unique to IAM. One of IAM's most powerful features Real Time Tuning applies the principles of pattern recognition and of learning from experience to the task of buffer management. IAM's Real Time tuning dynamically regulates file processing to match the demands of the job.

Indexed files can be processed in one or more of the following ways:

1. Randomly retrieving records by key throughout the file.
2. Sequentially reading some number of records or the entire file.
3. Updating retrieved records.
4. Inserting and deleting records.

IAM continually monitors logical requests and physical I/O activity during file processing. The combination of request types, buffer usage (active and inactive), average number of records consecutively read and the insert patterns the file is experiencing are just a few of the criteria IAM considers in determining how to apply its Real Time Tuning techniques. Based on this information, IAM will optimize its buffer usage and I/O chaining , making real time adjustments, every 128 I/Os for Random requests and every 32 I/Os for Sequential processing.

RANDOM PROCESSING For random processing, IAM dynamically acquires additional buffers so the most often referenced blocks can be retained in memory. IAM continually tracks the last 32 data blocks retrieved from the file. Monitoring requests in this way IAM can recognize patterns. If a request is for a block already in memory, the record is passed to the user without I/O to the disk. When IAM finds that a sufficient number of requests could have been satisfied without I/O to the disk, if more buffers had been in use, an additional buffer is acquired. This continues until the maximum number of buffers allowed is reached. IAM's buffer management techniques also provide for the release of buffers when it determines they are not providing any benefit.

SEQUENTIAL PROCESSING Buffer management for sequential processing, when consecutive logical records are being retrieved, differs from random in that IAM attempts to determine how many blocks it should be reading ahead in anticipation, to efficiently service the user. The goal is to be able to read as many useful blocks as possible with each real disk I/O. To accomplish this IAM monitors the number of blocks that are read consecutively, increasing or decreasing the number of consecutively chained buffers as appropriate. Both random and sequential operation can occur concurrently against any file. IAM in sequential operations reuses the same set of buffers, in an effort to isolate sequential buffers from those used for random processing.

UPDATE PROCESSING IAM distinguishes the differences between BATCH (SYNCHRONOUS) and ONLINE (ASYNCHRONOUS) processing. To improve batch performance IAM defers the rewrite to disk of an updated buffer until the block residing in the buffer is forced out of the buffer pool. This technique reduces I/O whenever multiple rewrites are requested against data blocks in the buffer pool. Under Online systems (ex: CICS), to insure integrity after a random update, buffers are always immediately rewritten.

CICS SYSTEMS CICS systems can frequently present a dilemma. They typically have to handle large numbers of transactions concurrently using the same file. All making different kinds of requests updates, inserts, some random, some sequential. IAM handles this better than VSAM. IAM's Real Time Tuning balances the buffers used by random and sequential requests, insuring a CICS system the best level of performance for the current activity. As a file experiences an increase in simultaneous requests, IAM's Real Time tuning assigns it more buffers. IAM's buffers are all maintained in extended private storage. When activity starts to let up, IAM gradually releases buffers.

SECTION 02: IAM TECHNICAL DESCRIPTION

02.50 IAM'S REAL TIME TUNING

LSR BUFFER POOLS IAM's Real Time Tuning buffer management is substantially better than CICS's default use of VSAM LSR buffer pools. First, IAM uses considerably less CPU time than VSAM especially when there are large LSR buffer pools to manage. Second, IAM's Real Time tuning dynamically acquires buffers for a file based on its activity and the benefit the buffers provide (how many 'look asides'). With IAM it is almost like having an LSR pool dynamically created for each of your files. Usually only VSAM clusters that need high performance can afford to have their own dedicated pools.

Typically LSR forces clusters to contend for space in fixed size pools that they must share. Only VSAM clusters with the exact same CISIZE can share an LSR pool. IAM's dynamic buffer acquisition frees CICS systems programmers from having to set up and maintain multiple LSR pools for different size CIs, or alternately requiring clusters to be DEFINEd with artificial standard CISIZES so they can share a limited number of pools.

In summary, even though CICS file management defaults to VSAM LSR in attempt to curtail real disk I/O, IAM's Real Time Tuning buffer management will provide a superior level of CICS performance at a fraction of the CPU and manual administrative overhead.

IAM BUFFERS When a file is OPENEd, IAM acquires an initial set of buffers for that file. The number acquired will be equal to the default value specified in the site's IAM Global Option table ([BUFOPNO= see Section 91](#)) or to the number of blocks that reside on one track whichever is greater. Default installation values can be changed in the IAM Global Options table using the program 'IAMZAPOP' ([See Section 91](#)). The initial number of buffers IAM acquires for a specific file can also be overridden at execution time using the IAM Override Statement ([MINBUFNO= see Section 11](#)).

IAM's Real Time Tuning starts with the file's initial buffers and during file processing dynamically adjusts the number of buffers actually used for a file up and down based on demand. More buffers are acquired as IAM determines that having additional buffers can reduce real I/O. The processing is determined by the default value specified in the site's IAM Global Option table ([MAXBUFNO= and BUFSP= see Section 91](#)). These default values can be changed in the IAM Global Options table using program 'IAMZAPOP'.

Normally the default is the number of buffers which will fit into 256K of storage. Whenever a file's ACB specifies a STRNO value larger than IAM's buffer value, STRNO will be used as the maximum number of buffers. Externally a file's maximum number of buffers can be overridden, up to a total of 255 buffers for Enhanced format files, or 32 for Compatible format files. This can be done when the file is DEFINEd or loaded using an IAM CREATE Override Statement and at execution time using the IAM ACCESS Override Statement ([MAXBUFNO= see Section 11](#)).

IAM reduces the number of buffers in use whenever it determines that they are not being referenced and consequently are no longer helping to reduce I/Os. The sample rate for reducing buffers is every 1280 I/Os. As appropriate, IAM continues to reduce the number of buffers it maintains for a file down to a minimum of one (1) or the minimum buffer value specified for the file.

A minimum number of buffers for processing can be established when the file is DEFINEd or loaded using the IAM CREATE Override Statement and at execution time using the IAM ACCESS Override Statement ([MINBUFNO= see Section 11](#)).

The MINBUFNO parameter causes IAM to acquire the specified number of buffers when the file is initially OPENEd and then to retain them as the minimum number of buffers.

NOTE: If the minimum number of buffers for a file exceeds the value specified or calculated to be the maximum value, IAM will set the minimum to be equal to the maximum value in effect.

02.50 IAM'S REAL TIME TUNING

TUNING STATISTICS

IAM's run time INFO report reflects the way an application uses the file. The report includes statistics on requests processed, I/Os to disk, buffer usage and number of I/Os saved by dynamically acquired buffers.

DYNAMIC BUFFER RETRIEVALS – displays the number of I/O requests satisfied from memory other than the current buffer.

MAXIMUM BUFFERS USED – displays the maximum number of buffers used during this OPEN of the file.

An IAM attention message is displayed in the run time INFO report and a flag stored in the IAM SMF record, whenever IAM determines that additional buffers would have been beneficial, but IAM was prevented from acquiring them. IAM will not acquire additional buffers after it reaches a file's maximum number of buffers allowed or whenever there is insufficient storage in the region for additional buffers.

Whenever IAM attention messages are generated applications should be reviewed and either the buffer values or the size of the region should be adjusted.

Applications that consistently show a MAXIMUM BUFFERS USED value that is higher than the file's minimum buffer value will usually benefit by setting the minimum number of buffers equal to this value.

If the number of I/Os IAM uses overall is relatively small, there is little to be gained from a change to IAM buffering values.

FILE LOAD BUFFERING

IAM's load process uses a different buffering concept. File load is exclusively a sequential output process. By default, for files that are allocated on cylinder boundaries, IAM will use enough buffer space to hold an entire cylinder, and write out half a cylinder per physical I/O. The buffer space can be increased, or decreased with the CRBUFOPT override or Global Option. The default for track allocated files is buffer space for two tracks, and writing out one track per I/O. This buffering concept provides for I/O overlap, even for programs using synchronous processing. For MVS/XA and MVS/ESA installations, the buffers are all above the line. The EXCP counts will reflect the actual EXCP (physical I/O) operations done, not the block count as in prior releases.

02.60 IAM'S DYNAMIC TABLING 'DATA RECORDS-IN-VIRTUAL'

02.60 IAM'S DYNAMIC TABLING 'DATA RECORDS-IN-VIRTUAL'

DYNAMIC TABLING IAM optionally tables records retrieved randomly from a file in virtual storage, without any programming changes to existing applications. On random reads, IAM checks to see if the key requested is for a record contained in the virtual table. If the record exists in the table, IAM passes it back to the user, eliminating the I/O to the disk.

NOTE: Random reads which are eligible to use the dynamic tabling feature are identified in the IAMINFO report as R. (READ) commands for Compatible format files, and as GET RANDOM commands for Enhanced format files. Other types of retrievals, such as GET commands and Read-or-GetNext cannot use the dynamic table because the precise key being sought is unknown.

To enable this option, use the IAM Override statement. The keyword on the ACCESS control statement is DYNCORE= (See Section 11). The DYNCORE value is specified in 1024 byte (1K) increments. The following example will reserve 200K of storage for the Dynamic Table.

```
ACCESS DD=iamfile,DYNCORE=200
```

HOW TABLING WORKS IAM's Dynamic Tabling of data records is similar to the Data-In-Virtual concept IBM introduced with Linear Data Sets, without requiring any programming or file format changes. IAM's concept of tabling is more efficient than IBM's because it's oriented to individual records instead of 4K sections. A smaller amount of memory is used requiring fewer real pages to back up the virtual pages.

If the key being requested is not currently in the table, IAM reads the record from the file. If found, the record is passed to the user and tabled for subsequent retrievals. If the record is updated, IAM changes the record in the table and on disk.

If the table fills up, IAM empties a portion of the table insuring the most current retrievals are maintained in the table.

RECOMMENDATIONS Applications which will benefit the most from Dynamic Tabling are those with high file activity where a subset of records in the file are repeatedly being read, with few ever updated. Small files with high random activity and few updates become in core tables without the need for any programming changes.

Under MVS/XA and MVS/ESA, Dynamic Table storage is in extended private.

TABLING STATISTICS IAM's run time INFO report reflects the way an application uses the file. The report includes statistics on requests processed, I/Os to disk, Dynamic Table usage and the number of records retrieved from the Dynamic Table.

DYNAMIC TABLE RETRIEVALS – displays the number of record requests satisfied from IAM's Dynamic Table.

DYNAMIC TABLE RECORDS – displays the number of records in the table when the file is closed. This will also normally be the maximum number of records that were in the table. The exception is when there has been deletes, because when a record is deleted, it is also removed from the table, leaving an available entry.

02.70 SMS SUPPORT IN IAM

02.70 SMS SUPPORT IN IAM

IAM provides support for SMS that is equivalent to the VSAM support, including support for JCL file definition, and temporary data sets. By definition, to be eligible for an SMS managed volume, the file must be assigned a Storage Class. The Storage Class, along with optionally a Data Class and/or Management Class, can be explicitly specified on the DEFINE command, by JCL parameters for JCL defined files, or selected by the ACS routines. IAM files on SMS managed volumes will be cataloged with the class names. As a part of the SMS support, an additional method of triggering an IAM DEFINE is available. Files will be defined as IAM files if the Data Class or Storage Class name contain the \$IAM literal.

ACS ROUTINES

For both IDCAMS DEFINE's and VSAM JCL allocations, the ACS (Automatic Class Selection) routines are called prior to the IAM DEFINE intercept. When IAM intercepts the request, the SMS classes, the SMS Storage Group, and the SMS volumes have already been selected. IAM will then screen the request and determine if the file should use the IAM format. If \$IAM, in the Data Class (DATACLAS) or Storage Class (STORCLAS) name, is being used as the criteria for determining IAM format files, then the class name(s) must have \$IAM in them at this point in the process. They contain either the explicit names from the DEFINE request, or the name(s) selected by the ACS routines. This allows the installation the possibility of controlling IAM files, and IAM usage through the ACS routines. For JCL allocation, these are the classes and volume(s) that will be used.

For IDCAMS DEFINE requests, the ACS routines will be re-entered when IAM issues the dynamic allocation of the file as a non-VSAM data set. The request will specify the SMS classes as received from the intercepted DEFINE request, and the volume(s) that had initially been selected by SMS. While Innovation does not recommend this, the ACS routines can change the SMS classes and the Storage Group, which will change the volume(s) on which the file is placed. The ACS routines can check the &DSORG value, which will be VS (VSAM) on the DEFINE, and PS on the IAM dynamic allocation. At this point, the file must not be switched to a non-SMS volume, because the allocation will fail. However, it can be switched from an unmanaged volume to an SMS managed volume. Changing the DATACLAS at this time will have no effect on the file characteristics, as they were determined by IAM prior to the dynamic allocation. The MGMTCLAS, STORCLAS, and Storage Group can all be effectively changed by the ACS routines on the dynamic allocation request.

IDCAMS DEFINE

IAM provides full support for IDCAMS defines under SMS. The DATACLAS, STORCLAS, and MGMTCLAS can either be explicitly provided on the DEFINE command, or selected by the ACS routines. The Data Class can provide file characteristics for the file being defined, including record length, key length, key offset, share options, free space, and others, eliminating the need to specify those values explicitly on the DEFINE. As per SMS rules, the options in the Data Class will be used, unless explicitly overridden on the DEFINE command. If a Data Class has been specified, IAM will consider the file to contain variable length records. This is because the LRECL in the Data Class is passed as both the average and maximum record length, thus not providing a way to distinguish between fixed and variable length records. If desired, the IAM CREATE override keyword FIX can be used to define a fixed length record file.

If the IAM allocation encounters any errors, the error messages will appear on the JES job log, with the MVS allocation messages (SYSMSGs) and also on the IDCAMS SYSPRINT, if it is available. Due to the manner in which IDCAMS prints messages on SYSPRINT, the error messages from IAM will precede the actual DEFINE command. IDCAMS will also print out additional error messages after the DEFINE, performing an analysis on the return codes set by IAM. Whenever possible, IAM uses the VSAM return codes that most clearly indicate the actual problem, although that is not always possible. Always refer to the IAM and related allocation error messages for the most precise problem determination possible.

02.70 SMS SUPPORT IN IAM

JCL ALLOCATION VSAM files being defined through JCL can also be easily converted to IAM files. This is done by either putting \$IAM in the data set name (DSN), or by using a Storage Class (STORCLAS) or Data Class (DATACLAS) with \$IAM in the name. Both permanent and temporary data sets can be defined, with the restriction that temporary data sets cannot be multi-volume; same as VSAM. The use of a Data Class (DATACLAS) is highly recommended for JCL defined files. By using a Data Class, values for Freespace, CI Size, and Share Options can be specified, which are not available through JCL parameters. While the use of a Data Class is preferable, the IAM override facility can be used to directly provide overflow and block size values. All files defined in JCL will, by default, be capable of handling variable length records, up to the maximum length specified in the DATACLAS or LRECL field.

To allocate IAM files through JCL, IAM must be in the link list. STEPLIB and JOBLIB are ineffective in this case, because it is the initiator that is issuing the allocation, and IAM must have access to various load modules for the define. These modules can also be placed in LPA, contact Innovation for a list of modules that would need to be in LPA for this process to work. Please note that not all of the IAM modules can be in LPA, so care must be taken when placing IAM in LPA.

If the IAM allocation encounters any errors, the error messages will appear on the JES job log and with the MVS allocation messages (SYSMSGs). SMS will also print out additional error messages appearing with the MVS allocation messages, performing an analysis on the return codes set by IAM. Whenever possible, IAM uses the VSAM return codes that most clearly indicate the actual problem, although that is not always possible. Always refer to the IAM and related allocation error messages for the most precise problem determination possible.

Any CREATE overrides for JCL allocated files must be in the job step that loads the file, not necessarily the step allocating the file. The define process does not access the IAMOVRID DD for JCL defines. For example, if the file is allocated in an IEFBR14 step, and then subsequently loaded by an IDCAMS REPRO, the IAM create overrides must be in the IDCAMS REPRO step.

ALLOC COMMAND The TSO ALLOC command, and the MVS DYNALLOC service, have also been enhanced to provide for allocation of new and temporary VSAM files. These requests will also be screened by IAM, and can be converted to an IAM file in the same manner as a JCL allocation can. IAM treats the request just like a JCL request. The new ALLOC keywords are the same as the new JCL keywords, and dynamic allocation has the equivalent text units.

MULTI-VOLUME CONSIDERATIONS IAM files can be spread across multiple SMS managed volumes, both with and without Guaranteed Space. Note that IBM restricts temporary VSAM files to a single volume, this also applies to IAM. When an IAM file is defined with a Storage Class that specifies Guaranteed Space, the primary allocation quantity is allocated on each volume at DEFINE time, as per the SMS non-VSAM rules. When Guaranteed Space is not specified, only the first volume is selected at DEFINE, and the subsequent volumes are cataloged as an '*'. During file load or reorganization, if additional volumes are needed, they will be selected by SMS. SMS has a restriction that within a job step, if a file defined without guaranteed space uses additional volumes, only one DD statement can be used, because any other DD's are not updated to indicate the additional volumes. For IAM users, this is only of concern for job steps loading or reorganizing files and being accessed within the same step.

The following are examples of using the SMS classes for IDCAMS defines, and SMS JCL allocations. Further examples, with full explanations of IDCAMS parameters are in the IDCAMS section of the manual, [Section 10](#).

02.70 SMS SUPPORT IN IAM

SMS ALLOCATION EXAMPLE

In the example below, an IDCAMS DEFINE is done for an IAM file, with all file attributes being supplied from the specified Data Class. The file is allocated as an IAM file because of the \$IAM literal in the Data Class name.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
              DEFINE CLUSTER -
                  (NAME(TSTSMS.SMS202.CLUSTER) -
                  MGMTCLAS(DBSTNDRD) -
                  STORCLAS(BASE) -
                  DATACLAS($IAMTST1))
/*
```

Figure 48: SMS Define of an IAM file

JCL PERMANENT ALLOCATION

Below is an example of a JCL allocation of an IAM file, with explicit specification of a Data Class and a Storage Class. All of the file attributes will be taken from the Data Class, except for SPACE which is explicitly specified. The file is selected to be an IAM file because the literal \$IAM is part of the Data Class name.

```
//ALLOCIAM   EXEC  PGM=anyppgm
//FILE1     DD    DSN=SMS.INDEXED.CLUSTER,DISP=(,CATLG),
//          DATACLAS=DC$IAM01,STORCLAS=BASE,
//          SPACE=(CYL,(6,2))
```

Figure 49: Example of JCL Define of an IAM file

TEMPORARY DATA SET

In the example below, a temporary IAM / VSAM data set is allocated in an IEFBR14 step, followed by an IDCAMS REPRO from a permanent file. The file is defined as an IAM file because of the \$IAM literal in the data set name. In this example, the new SMS JCL parameters are used to define the record attributes for the file.

Note the IAM override in the IDCAMS step to provide for overflow areas in the temporary file. For full information on IAM overrides, refer to [Section 11](#) of the IAM manual.

```
//ALLOCTMP   EXEC  PGM=IEFBR14
//FILE1     DD    DSN=##$IAMFIL,DISP=(,PASS),
//          SPACE=(CYL,(1,1)),
//          RECORGE=KS,LRECL=380,KEYLEN=4,KEYOFF=8,
//          STORCLAS=BASE
//REPRO     EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//TEMPFILE  DD    DSN=##$IAMFIL,DISP=(OLD,PASS)
//IAMOVRID  DD    *
              CREATE DD=TEMPFILE,I=25,O=10000
/*
//SYSIN     DD    *
              REPRO IDS(PROD.INDEXED.FILE) OUTFILE(TEMPFILE)
/*
```

Figure 50: Example of temporary IAM file allocation

02.80 IAM STORAGE REQUIREMENTS

02.80 IAM STORAGE REQUIREMENTS

IAM initially acquires storage for its control blocks, data buffers and index areas when a file is OPENed. Additional storage may be acquired as the file is being processed.

Under MVS/XA and MVS/ESA, IAM requests storage above the 16MB line except for a 4K to 8K block used for its control areas.

IAM's Real Time Tuning acquires additional buffers if performance considerations warrant. Under MVS/XA and MVS/ESA, these buffers are requested above the 16MB line. For Enhanced format files, if a buffer getmain is ever returned in below the line storage, the buffer is immediately free-mained.

IAM's index contains a key for every prime data block (VSAM CI) in the file. IAM however, since it drastically reduces the number of data blocks as compared to VSAM, will have fewer index entries. Additionally, each index entry is much smaller than VSAM because IAM's key is super-compressed. During processing, the entire index resides in storage. Under MVS/XA and MVS/ESA, storage for the index is requested above the 16MB line.

For Compatible format files, IAM requires storage used to maintain pointers for records currently in Independent Overflow. When the file is OPENed for update, sufficient storage is acquired for inserts of up to the OCOREO% value (default 10%) of the total capacity of Overflow. IAM incrementally acquires additional areas equal to the OCOREX% value during processing as needed. This continues until Independent Overflow is exhausted. Under MVS/XA and MVS/ESA, storage for the Independent Overflow table is requested above the 16MB line. To change the overflow percentages used by IAM, specify the IAM Override Statement ([Section 11.01](#)). The keywords are OCOREO%= and OCOREX%=.

For Enhanced format files, the storage for the Overflow index is also acquired above the line, but the amount of storage acquired is more dynamic due to the dynamic expansion capabilities of the Extended Overflow area. With the elimination of the 64,000 blocks in overflow, the size of each index entry for overflow has grown by two bytes, which may cause larger storage requirements for Enhanced format files.

Additional storage may be acquired for special features which do not exist in VSAM, such as Dynamic Tabling. Under MVS/XA and MVS/ESA, this storage is requested above the 16MB line.

02.90 ADDING RECORDS TO AN IAM FILE

02.90 ADDING RECORDS TO AN IAM FILE

VSAM'S INSERT VSAM uses two concepts to manage inserts in a VSAM KSDS file:

CONCEPT

FREESPACE CI% – Specifies the percentage of each control interval that is left empty for future inserts. IAM's Integrated Overflow is basically the same concept.

FREESPACE CA% – Specifies the percentage of CIs within each CA to be left empty. Its purpose is to establish free space areas throughout the file for CIs which are filled. These areas are tied to the file's existing index structure. If an insert does not fit within the CI where it should be placed, VSAM will split the CI, using one of the free CIs within the same CA. If a CA has no remaining free CIs, VSAM splits the CA into an area at the end of the file.

The VSAM technique of splitting record areas has serious drawbacks. If you are adding many records with similar keys (mass inserts), VSAM is forced to split many times. CI and especially CA splits are very time consuming. A file can become unusable if a split does not complete (ex: system crash). If you try to reduce CA splits by increasing CA%, you waste large amounts of disk space as most of the free CIs will never be used.

IAM'S INSERT
CONCEPT

IAM's concept of Overflow is far superior to VSAM's. For Compatible format files, IAM reserves a single general purpose area in the file to be used for inserts which do not fit in the blocks where they should be added. If a block is full, IAM adds the record to Independent Overflow. This concept makes much more efficient use of disk space because it is not tied to the file's existing index structure. For Enhanced format files, this Overflow area is dynamically acquired, on an as needed basis.

IAM never moves data records once they have been placed on the file, unless the record increases in size and no longer fits in the block. In this case, IAM puts the record into Overflow first, before deleting it from the prime area. If an abend occurs, IAM retrieves the correct copy of the record bypassing the duplicate record on subsequent processing.

IAM files do not become unusable because of system crashes or job cancellation. Inserts into IAM files are much faster than VSAM.

It is very easy to tell how much Independent Overflow is used or remaining. The IAM INFO report gives you the exact number of records currently in Independent Overflow and how many records are left for additions. It is next to impossible to get this information from VSAM.

02.100 IAM MISCELLANEOUS FUNCTIONS

02.100 IAM MISCELLANEOUS FUNCTIONS

RUN TIME REPORTS IAM's run time INFO report reflects the way an application uses the file. The report includes statistics on requests processed, I/Os to disk, buffer usage and number of I/O's saved by dynamically acquired buffers. The file's definition and overflow characteristics are also displayed.

If an IAMINFO DD statement is present in the JCL, a run time INFO report is produced each time an IAM file is CLOSEd. IAM can be directed to skip writing INFO reports for one or more files even though the IAMINFO DD statement is present. This may be appropriate when a file is repeatedly OPENed/CLOSEd and unnecessary reports would be generated. To disable the INFO report, use the IAM Override statement ([\(INFO=NO see Section 11\)](#)).

IAM SMF RECORDS IAM can directly record its run time statistics in a system's SMF files. When IAM's optional SMF recording is active an SMF user record is written every time an IAM file is OPENed and CLOSEd.

Member IAMUSMF in the IAM Installation Control Library is a DSECT of the IAM SMF user record's layout. The fields in IAM's SMF user record contain the same information found in IAM's Run Time Statistics (INFO) Report. The IAM (INFO) Run Time Statistics Report is documented in [Section 12.04 of this manual](#).

MESSAGES AND CODES IAM sets return codes compatible to VSAM for error conditions. IAM generates a descriptive error message to the JCL LOG if an error has unique meaning to an IAM file. See Messages and Codes ([Section 80](#)) for details on IAM messages.

IAM uses Route and Descriptor codes on all messages written to the console log. Specific information on descriptor and route codes for WTO console log messages may be found in the IBM manual 'ROUTING AND DESCRIPTOR CODES'. IAM's default is 00 for the descriptor code and 11 for the routing code. The routing and descriptor codes used can be changed in the Global Options Table or by IAM overrides.

GLOBAL OPTIONS IAM obtains many of its defaults when Defining or loading an IAM file from the IAM Global Option table. These installation default values can be changed in the IAM Global Options table using program 'IAMZAPOP' ([Section 91](#)).

FILE RECOVERY IAM provides the ability to recover data from the accessible portions of physically damaged or data checked devices. Program 'IAMRECVR' reads an IAM file, bypassing the unusable portions of the file. A sequential file of the undamaged records is produced. 'IAMRECVR' can even retrieve data from an IAM file with damaged indexes or internal control records.

THIS PAGE WAS INTENTIONALLY LEFT BLANK

SECTION 03: IAM DATA SET STRUCTURE

03.00 IAM DATA SET STRUCTURE

03.10 IAM DATA SET STRUCTURE OVERVIEW

IAM provides improved levels of performance, efficiency and reliability unsurpassed by any other index file processor. To gain these advantages, rather than attempting to manage a VSAM structure more expertly, IAM establishes its own uniquely structured data set.

The organization of an IAM data set is structurally simpler than its VSAM equivalent. An IAM data set is a relative block non-VSAM (DSORG PS) organized and managed by IAM using the EXCP access technique.

Complementing overflow structures within an IAM file can accommodate any type of file growth. A percentage of every block in an IAM file can be left free at load time to accommodate randomly distributed additions. IAM establishes free records within the file to be used when inserts are not randomly distributed but are clustered into groups. For Enhanced format files, free space can be reserved during file load to accommodate file expansion. For Compatible format files, free blocks can be established during file load to accommodate inserts to the end of the file.

Programs which create IAM files are device independent. IAM automatically determines the target device type and, using the file's record length and CISIZE, calculates the best blocksize for that device. Programs which process IAM files are totally insulated from the blocking and structure of an IAM file.

03.20 IAM FILE ALLOCATION

03.20 IAM FILE ALLOCATION

An IAM file is allocated on the disk and cataloged using an IDCAMS DEFINE operation. The IAM file is created when the file is Opened for OUTPUT and loaded with one or more records sequentially by key.

IAM files can be loaded into a new allocation or an existing file allocation with a disposition of OLD or SHR as appropriate. IAM files are always considered to be reusable, unless the IAM Global Option ENABLE=NOREUSE is set, and the file is defined with the NOREUS attribute..

IAM files are simple non-VSAM space allocations managed via the EXCP access technique and can be larger in size than VSAM clusters. IAM datasets can be Defined across multiple volumes. They make full use of secondary space allocation and do not require contiguous (CONTIG) extents. Space for an IAM file can be Defined in records, tracks or cylinders. For Compatible format files, the amount requested should be sufficient to accommodate all of the records to be loaded and the overflow Defined for the file. If a secondary allocation value is Defined, during a load procedure, should the primary allocation be insufficient IAM will acquire additional extents. VSAM's 4.3 Gigabyte (GB) file size limitation does not apply to IAM files.

03.30 IAM COMPATBILE FORMAT FILE AREAS

03.30 IAM COMPATBILE FORMAT FILE AREAS

This section describes the Compatible format IAM file. Information on Enhanced format file structure is contained in [Section 1.2](#), File Structure Overview, on page 3. An IAM file does not exist until the space allocated for the file has been initialized by a load process. Within the space allocated for a file, IAM establishes five distinct areas. These areas are the foundation for management of the file. The following is a description of how these areas are structured.

CONTROL AREA The first area established is IAM's file definition and control area. Control information within an IAM file includes such things as the file characteristics (blocksize, number of blocks in the file, etc.), logical record description (record length, relative key position, etc.), overflow specifications (type of overflow, number of records, etc.), execution time options (core usage, key compression, I/O buffering etc.) and a control area ID. The control area ID serves as an indication that a load of the file has been successfully completed.

INDEPENDENT OVERFLOW AREA The next area established is Independent Overflow. During file load, IAM establishes an area in the file capable of accepting the Independent Overflow records this file was Defined to accommodate. The size of this area is calculated based on the number of Independent Overflow records requested and the block size of the IAM file.

When an IAM file is Defined, the FREESPACE CA% value is used to calculate Independent Overflow. IAM calculates its Independent Overflow value using the CA% and primary allocation values. Since IAM makes more efficient use of overflow, the CA% specified is cut in half for files with more than 10 cylinders.

Based on the size of the file's primary allocation, IAM will not reduce the CA% below the values shown in the following table:

Primary Allocation	Minimum CA%
1 cylinder or less	10%
2 to 3 cylinders	9%
4 to 5 cylinders	8%
6 to 7 cylinders	7%
8 to 9 cylinders	6%
10 cylinders or more	5%

Figure 51: Independent Overflow and CA% Freespace

IAM calculates the approximate number of maximum size records that will fit in the primary allocation and multiplies the result by CA percent. This value will become the number of records in Independent Overflow.

The maximum number of records that IAM will reserve based on CA% freespace is from the IAM Global Options Table MAXOVERFLOW value, which defaults to 50,000. The minimum the lower of 500 records, or the number of records that will fit in half of the primary allocation⁹¹. The number of Overflow records IAM is to reserve in a file can be overridden, when it is Defined and at execution time when it is loaded, using an IAM CREATE Override Statement ([Overflow= see Section 11](#)).

03.30 IAM COMPATIBLE FORMAT FILE AREAS

IAM recognizes SMP/E CSI files by their cluster name (CSI in the last index level) and will use a default of 20% Integrated Overflow and 50,000 Independent Overflow records.

It is very easy to tell how much Independent Overflow is used or remaining in an IAM file. The IAM INFO report or a LISTCAT ALL with IAMPRINT gives you the exact number of records currently in Independent Overflow and how many records are left for additions. It is next to impossible to get this information from VSAM.

If an application loads just one (1) record to a file, IAM formats additional space for Independent Overflow within the primary space allocation. The total combined space for Prime Extension (described below) and Independent overflow is approximately 90% of the primary space requested.

PRIME DATA AREA The next area is called the Prime Data area. When the file is loaded (or created), IAM accepts records to build the file's prime data area. Records passed to IAM load must be in ascending key sequence and are placed into the prime data blocks. During this file load IAM reserves, as free space within every prime data block, the percentage of Integrated Overflow (CI%) specified for the file. IAM continues the prime area load until the user application or IDCAMS REPRO stops passing records and Closes the file.

PRIME EXTENSION AREA The Prime Extension area is established following the prime data blocks. It is used to accept inserts to the file ascendingly higher than the previously existing high key on the file. Prime Extension is specified in blocks. Use the keyword PE= in the IAM Override statement to change the Prime Extension value.

If an application loads just one (1) record to a file, IAM formats additional space for Prime Extension within the primary space allocation. The total combined space for Prime Extension and Independent Overflow will be approximately 90% of the primary space requested.

INDEX AREA The next area established is for the Index blocks. IAM establishes a high key value for each prime data block during file load. IAM uses one key per block for the prime index. These key values are written out to a data space under MVS/ESA, or a temporary work file under MVS/XA, as each prime block is loaded. During program CLOSE processing, these key values are read from the data space and used to create an index for the IAM file. When a data space is not used, IAM dynamically allocates the temporary work file.

IAM compresses the keys in the index using a proprietary compression technique on each key which results in a higher compression ratio than VSAM. If the size of the compressed index is at least 10% smaller than the uncompressed index, the compressed key structure is stored in the file after the uncompressed index.

IAMM

INNOVATION ACCESS METHOD

IDCAMS SUPPORT

THIS PAGE WAS INTENTIONALLY LEFT BLANK

SECTION 10: IAM AND ACCESS METHODS SERVICES

10.01 GETTING STARTED WITH IAM

10.01 GETTING STARTED WITH IAM

OVERVIEW IAM's system level VSAM interface (VIF) allows IDCAMS, CICS, other ONLINE systems, batch programs, and TSO applications to use IAM files in place of single indexed VSAM KSDS files without any change to programs or job step JCL. Both batch and online (synchronous and asynchronous) processing are supported. VIF supports any programming language which provides VSAM single index KSDS key indexed and sequential access or file creation.

The VIF Interface must be active in your system to use IDCAMS to DEFINE or access IAM files ([see Section 90.04](#)).

DEFINING IAM FILES Using IAM files in place of VSAM clusters is simple. For an application currently using VSAM cluster(s), find the IDCAMS steps which DEFINE the VSAM cluster(s). Add the parameter OWNER(\$IAM) to the DEFINE procedure. The next time the cluster is DEFINEd, it will become an IAM file.

A second way to establish an IAM file is to simply include the characters '\$IAM' anywhere in the cluster name. Using '\$IAM' in the cluster name is the recommended way to create new IAM data sets. The advantages of using this method are that IAM files will be easily identifiable from looking at the JCL, on VTOC and DASD management reports, and on SMF reports your installation may use.

SMS users can also establish an IAM file using the DATACLAS and STORCLAS constructs. If '\$IAM' appears in either, the file will be defined as an IAM file. The values for DATACLAS and STORCLAS can be explicitly specified via DEFINE parameters, in the JCL, or by the ACS routines. Selecting a Data Class or Storage Class in the ACS routines can provide a mechanism to automatically convert files to IAM, with minimal manual effort.

EXAMPLE: A cluster name of 'USER.FILE.\$IAM.TEST' will create an IAM file instead of an VSAM cluster.

NOTE: It is a user's option when DEFINEing an IAM file to choose which method they prefer to use when defining an IAM file. Making this choice will be referred to, in this document, as using the '\$IAM parameter'.

CONVERTING VSAM FILES To convert a VSAM cluster to an IAM file with the same name use the following procedure:

1. REPRO the VSAM cluster to a sequential file on tape or disk.
2. DELETE the VSAM cluster.
3. DEFINE the file with the same names and attributes as the VSAM cluster, adding the parameter OWNER(\$IAM) to the owner attributes.
4. REPRO the sequential file into the IAM file.

This is all that is needed to convert a VSAM cluster to an IAM file. [See Section 10.09](#) for an example of this procedure.

IAM files can be loaded by any application program that currently loads VSAM files. For optimal load performance the technique of a single record load subsequently followed by mass inserts should be avoided. COBOL users should note that when a file is opened for loading with ACCESS IS DYNAMIC or RANDOM, COBOL does a single dummy record load, followed by a delete of that record. For optimal COBOL file loading, ACCESS IS SEQUENTIAL should be used.

WARNING: Many applications contain more than one DELETE and re-DEFINE step. You must change all steps which DEFINE the file you are converting. If a DEFINE is executed without the (\$IAM) parameter, the file will revert back to VSAM. If this should occur, jobs will still execute normally, however they will run much slower and all of the IAM benefits will be lost.

10.01 GETTING STARTED WITH IAM

FILE ALLOCATION An IAM file's DASD space allocation can be DEFINEd in records, tracks or cylinders and can extend across multiple volumes. IAM files do not have VSAM's 4.3 gigabyte (GB), file size restriction. The allocation for an IAM file must be sufficient to accommodate all of the records to be loaded with additional allowance for records that will be added before a reorganization.

IAM's advanced file structure and data compression cuts the space requirement of VSAM files by 30 to 70%. DEFINE statement space requirements should be reviewed on all files converted to IAM. It is recommended that a secondary allocation value be specified on all IAM files. This will make the file eligible for IAM's automatic release facility.

IAM reserves the free space areas specified by the DEFINE when the file is loaded, REPROed or reorganized. This free space is distributed throughout the file to accommodate future inserts. If more disk space was requested than is needed, and secondary allocation has been specified, IAM will release the over allocated space.

Careful consideration may need to be given to the space allocation request, particularly when a single record loading technique is used. IAM will only use secondary space on file loads and reorganizations. With single record loads, IAM formats out 90% of the primary space allocation, dividing the space up between the Independent Overflow and Prime Extension areas. ([Refer to Section 3 for a description of the IAM file format.](#)) Secondary space will generally not be used for files that are loaded by a single record, unless explicit overrides are provided for the size of the overflow areas.

ISAM SUPPORT If you wish to convert ISAM files to IAM, use IDCAMS DEFINE to allocate the IAM file. Copy the file to IAM using REPRO. IAM's VSAM Interface (VIF) supports programs coded to use ISAM files. IAM's support for ISAM files is the same as VSAM. Refer to the IBM SRL titled VSAM Administration Guide, for the operating system you are running under, for further explanation. If programs are using functions which are not supported by VSAM's ISAM compatibility, contact INNOVATION.

10.02 ACCESS METHOD SERVICES (IDCAMS) OVERVIEW

10.02 ACCESS METHOD SERVICES (IDCAMS) OVERVIEW

IDCAMS OVERVIEW This section documents the use of IDCAMS with IAM files. Most IDCAMS commands will execute without change. We will indicate for each command if there are special considerations or additional features available when the file being DEFINEd or accessed is an IAM file.

IDCAMS SUPPORT IAM's VSAM Interface (VIF) supports the following IDCAMS commands:

- DEFINE - Use the same parameters you would to define a VSAM KSDS cluster or ESDS file and simply add the parameter '\$IAM'.
- DELETE - Will execute the same for a VSAM cluster or an IAM file
- LISTCAT - LISTCAT output for an IAM file will reflect a non-VSAM format. A LISTCAT ALL will result in IAM providing a full report on the file attributes to the IAMPRINT DD, which if not already allocated will be dynamically allocated to SYSOUT=* in batch, or to the terminal for TSO users. To direct the report elsewhere, code an IAMPRINT DD card in the job's JCL, for example:

```
//IAMPRINT DD SYSOUT=A
```
- PRINT - Will execute the same for a VSAM cluster or IAM file
- REPRO - Will execute the same for a VSAM cluster or IAM file. Use REPRO to copy VSAM clusters to IAM files and IAM files back to VSAM
- VERIFY - IDCAMS will complete normally. However since this command is not applicable to an IAM file it performs no useful function. The file will be Opened and Closed
- ALTER - Is not currently supported by IAM. ALTER can be issued for non-VSAM values, as described in the Access Method Services manual
- IMPORT - Is not currently supported by IAM. See [Section 10.11](#) on how to perform the same function
- EXPORT - Is not currently supported by IAM. See [Section 10.11](#) on how to perform the same function.

10.03 ACCESS METHOD SERVICES (IDCAMS) - JCL REQUIREMENTS

10.03 ACCESS METHOD SERVICES (IDCAMS) - JCL REQUIREMENTS

The VSAM access method services utility program, IDCAMS, executes under IAM's VSAM Interface (VIF) without any JCL modification.

IDCAMS commands to DEFINE, DELETE or access an IAM file follow the same conventions as for a VSAM cluster. There is no change in the format of the command statement. The following documentation will detail each command and any special IAM considerations.

IDCAMS JCL REQUIREMENTS The JCL statements needed when IDCAMS is used to DEFINE, DELETE or access an IAM file are essentially no different from those needed for a VSAM cluster. The following are additional considerations which should be made to take full advantage of IAM's superior capabilities.

EXECUTE STATEMENT Specifies the program IDCAMS and remains unchanged.

STEPLIB/ JOBLIB DD STATEMENT It is recommended that the IAM program load library be included in the system LINKLIST. If it is not, include a STEPLIB/JOBLIB DD statement specifying the library in which the IAM program load modules reside. Because IDCAMS is an authorized program, all libraries in the STEPLIB or JOBLIB concatenation must be APF authorized.

'file' DD STATEMENT This is always an optional DD statement for the DEFINE of an IAM file. It is required by other commands whenever the FILE parameter is specified in a control statement. The 'DDNAME' must match the value coded for FILE.

IAMPRINT DD STATEMENT This is an optional IAM DD statement. A LISTCAT ALL will generate a report detailing the IAM file characteristics; record and key length, blocksize, overflow values, tracks used by the IAM file, etc... to this DD statement. If an IAMPRINT DD is not present, and the SYSPRINT DD is going to SYSOUT, then an IAMPRINT DD is dynamically allocated to SYSOUT=*. For TSO users, the output will be directed to the terminal with TPUT, unless an IAMPRINT DD is present.

EXAMPLE: //IAMPRINT DD SYSOUT=*

IAMINFO DD STATEMENT This is an optional IAM DD statement. If it is present, and the file is IAM, a report is generated each time the IAM file is CLOSEd. This report displays run time statistics (memory used, read and write I/Os), counts of each command issued (GET, PUT, ADD, DELETE, POINT), number of buffers used and additional information on the characteristics of the IAM file. This report is very valuable to you and to INNOVATION in determining what each job is doing.

EXAMPLE: //IAMINFO DD SYSOUT=*

NOTE: It is recommended that this DD statement be included in any job executing against an IAM file. As an alternative, the IAM SMF records can be enabled and the IAMINFO reports can be obtained by running the IAMSMF utility program.

IAMOVRID DD STATEMENT This is an optional IAM DD statement. If it is present, and the file is IAM, IAMOVRID is read for IAM Override control statements. The purpose of this DD statement is to give the user the capability of overriding IAM default values or values specified in the DEFINE CLUSTER command. In addition, IAM gives the user many capabilities that do not exist in VSAM (ex: DYNCORE, RELEASE). To use these options, see IAM Override statements ([Section 11.01](#)).

EXAMPLE: //IAMOVRID DD *
 CREATE DD=i amddnam, OVERFLOW=1000
 CREATE DD=&ALLDD, RELEASE=NO
 ACCESS DD=i amfile, MAXBUFNO=16

10.04 IDCAMS DEFINE CLUSTER COMMAND

10.04 IDCAMS DEFINE CLUSTER COMMAND

DEFINE CLUSTER COMMAND To DEFINE an IAM file instead of an VSAM cluster, use the DEFINE CLUSTER command, add the '\$IAM' parameter, and follow normal VSAM conventions.

When the IAM VSAM Interface (VIF) is active in the system, every DEFINE is analyzed by VIF. If the DEFINE contains the \$IAM parameter, the file is created as an IAM file. If the \$IAM parameter is not coded, a VSAM cluster is created.

IAM allocates a non-VSAM file on disk with a DSORG of DA. This file contains the index, the data records and the file characteristics all incorporated into a single file. The requested overflow areas are established in an IAM file when it is loaded. IAM files must be cataloged in a VSAM or ICF Catalog.

The parameters specified in the DEFINE statement are applied to the IAM file. Because the IAM file and overflow structures are different from VSAM's, some of the parameters specified may be changed or ignored.

When the DEFINE is for an IAM file the following rules will apply:

1. The cluster name will be the name of the IAM file.
2. The data and index component names will be ignored.
3. Index component attributes are not needed and will be ignored.
4. Attribute parameters will be used or ignored as documented in the following paragraphs.

DEFINE PARAMETERS Most VSAM DEFINE parameters are applicable to IAM. The following paragraphs will document the parameters that are necessary for IAM, the ones that do not apply and any differences between IAM and VSAM.

The following parameters must be specified on a DEFINE statement.

CLUSTER - Identifies that a file is to be created. This is a required parameter for an IAM file.

NAME(iamfilename) - This is a required parameter for an IAM file. The entry name specified for the cluster will be the name of the IAM file. The data and index component names are ignored. If '\$IAM' appears anywhere in the name the file will be created as an IAM file.

The following data set name examples will both create an IAM file.

```
DEFINE CLUSTER (NAME(user.file,$IAM,xxxx)
DEFINE CLUSTER (NAME(CICS.TRANSACT.TEST$IAM) -
```

NOTE: For SMS users, the parameters STORCLAS AND DATACLAS can be used to specify \$IAM in lieu of the cluster name or owner parameter.

OWNER(ownerid) - If OWNER(\$IAM) is coded on the CLUSTER statement, the file will be created as an IAM file. When '\$IAM' is not present in the cluster name, and the \$IAM parameter is not otherwise specified, the file will be created as a VSAM cluster.

NOTE: By default, unless changed in the Global Option Table, IAM stores internal information in the OWNER field in the catalog. The information is not printable, and will appear as periods or other characters on a LISTCAT. If your installation uses the OWNER field, please refer to the INSTALLATION section of the manual for information on setting the Global Options Table to preserve the contents of the OWNER field in the catalog.

The following statement will create an IAM file.

```
DEFINE CLUSTER (NAME(user.file) OWNER($IAM) -
```

CYLINDERS(primary secondary) - DEFINES the space in cylinders.

TRACKS (primary secondary) - DEFINES the space in tracks.

RECORDS (primary secondary) - DEFINES the space in records.

Specifies the amount of space to be allocated for the file on the specified volume(s). The first field specifies the primary space allocation value. This space must be available on the first volume. The second field specifies a secondary allocation value. If the primary space is not sufficient to load the data set, IAM will extend the data set by this value.

10.04 IDCAMS DEFINE CLUSTER COMMAND

IAM files require from 30 to 70% less disk space than your existing VSAM clusters. IAM files makes more efficient use of DASD space. Their super compressed index and advanced internal structure usually result in about a 20 to 40% reduction in the amount of disk space compared to a similar VSAM file. IAM's optional Data Compression, in most circumstances, will provide an additional 20 to 50% reduction in disk space. A compressed file requires more CPU time to process than an uncompressed file. IAM's Data Compression techniques however, have proven themselves so efficient, IAM compressed files usually take less system CPU resources to manage than normal VSAM clusters. To request that IAM compress data records, see IAM Override statements ([Section 11.01](#)). The keyword is `DATA COMPRESS=YES`.

In an effort to conserve disk space and prevent over allocation, IAM releases the unused space in the file after it is loaded. This is done automatically when secondary space is specified. If you want to override IAM's default of releasing the over allocated space, see IAM Override statements ([Section 11.01](#)). The keyword is `RELEASE=NO`.

After a file has been loaded (or REPROed), a LISTCAT ALL will show you the exact number of tracks an IAM file is using. See [Section 10.07](#) for a description of the IAMPRINT file attribute report.

IDCAMS allocation rules are followed with the exception that IAM files, because they are non-VSAM data sets, are limited to a maximum of 16 extents per volume. If a DEFINE specifies RECORDS, IAM will allocate sufficient space to contain the specified number of average size records.

As a general rule, when converting VSAM files to IAM, initially retain the original VSAM space allocation values. After observing the IAM space and overflow requirements the space allocation can be adjusted. Changing the primary space quantity also has the effect of altering the size of the Independent Overflow area. For further information on how the size of Independent Overflow is calculated, refer to the FREESPACE parameter.

For files that are initialized with a single record, the primary space quantity must be at least 10% larger than the anticipated space requirement. This is because during a single record load, 90% of the primary allocation is formatted for overflow areas, and no secondary space allocation will be used until the file is reloaded.

VOLUMES(volser volser ...) - Specifies the volume(s) on which the IAM file is to be allocated. The file is initially allocated only on the first volume, with the subsequent volumes cataloged as candidate volumes. These additional volumes may then be used, by the operating system, to satisfy secondary allocation requests for the file. The primary allocation must be available on the first volume.

IAM files can be allocated to non-specific volumes. If you wish to use non-specific allocation, specify `VOLUME(ANYVOL)`. For multi-volume non-specific allocation, specify `VOLUME(ANYVOL ANYVxx ANYVyy ...)`. This will result in IAM issuing a non-specific dynamic allocation request for the IAM file. The first volume will be selected by MVS allocation. Any additional volumes are selected by IAM, which will select volumes from the specified UNIT name that are of the same device type as the first volume selected. IAM builds a list of the eligible volumes, then selects those volumes that have the largest quantity of available contiguous space. All of the volumes must be mounted as STORAGE. Volume selection is done at DEFINE time, and a subsequent LISTCAT will show the volumes selected. The default UNIT name used is SYSDA. To change to a different UNIT, use the IAM CREATE override keyword `UNIT=`. Refer to the IAM Override statements ([Section 11.01](#)) for information on the UNIT override.

When using the IAM non-specific allocation, do NOT specify UNIQUE. If the MODEL parameter is specified, also specify the SUBALLOCATION parameter. This is to prevent IDCAMS from attempting to allocate the volumes before IAM intercepts the request.

The UNORDERED parameter is ignored. IAM always uses the volumes in the order specified.

10.04 IDCAMS DEFINE CLUSTER COMMAND

OPTIONAL DEFINE PARAMETERS

The following parameters are optional on the DEFINE statement.

BUFFERSPACE(bytes) - Specifies the amount of BUFFERSPACE to be used when VSAM processes the file.

IAM ignores the BUFFERSPACE parameter. IAM's Real Time Tuning and buffer management is far superior to VSAM's buffer management. IAM continually adjusts the number of buffers in use as a file is being processed. See IAM Technical Features ([Section 2](#)) for details on IAM's Real Time Tuning.

If you wish to override the IAM default buffer values, see IAM Override statements ([Section 11.01](#)). The keywords are MAXBUFNO=nnn and MINBUFNO=nnn.

CATALOG(catalog name) - Specifies the catalog in which the IAM file is to be cataloged. IAM supports both VSAM and ICF Catalogs. The CATALOG parameter should not be specified for data sets placed on SMS managed volumes.

If CATALOG is not specified, the standard catalog search is used. If STEPCAT/JOBCAT is coded, the IAM file will be cataloged in the STEPCAT/JOBCAT. If STEPCAT/JOBCAT is not present, the file will be cataloged in the Catalog containing the alias for the high level index.

CONTROLINTERVALSIZE(size) - Controls the physical block size of the VSAM file. VSAM restricts CISIZE to multiples of 512 or 2048. This usually results in wasted disk space. IAM takes the CISIZE specified and based on the device type being allocated to and the record size, calculates a blocksize for the file that is a proportion of the track capacity. If the blocksize developed would result in more than 4 blocks on a track, IAM will increase the blocksize until only 4 blocks will fit on a track. The blocksize may be further adjusted if the block cannot contain at least 4 maximum size records.

FOR EXAMPLE ON A 3390:

CISIZE(4096) IAM's BLKSIZE will be 13682 - 4 blocks per track
CISIZE(20480) IAM's BLKSIZE will be 27998 - 2 blocks per track

If the file contains fixed length records, the blocksize is rounded down to be an even multiple of the record length. For variable length records the blocksize is a strict proportion of the track capacity.

The maximum blocksize for an IAM file is 32760. The recommended blocksize is quarter track blocking (4 blocks to a track).

If you wish to override IAM's calculated blocksize, see IAM Override statements ([Section 11.01](#)). The keyword is BLKSIZE=nnnn.

DATACLAS(dataclas) - For SMS installations, this parameter specifies the name of the SMS DATA CLASS construct, which provides the allocation attributes for the new data set. The DATACLAS name must have been previously established by the Storage Administrator. If the DATACLAS name contains the literal \$IAM, the file will be defined as an IAM file. Attributes from the DATACLAS will be used, unless otherwise explicitly specified on the DEFINE statement.

ERASE - Specifies that VSAM is to erase the data set when it is DELETED. IAM ignores this parameter.

EXCEPTIONEXIT - Specifies the name of a load module to get control when an I/O error occurs. IAM ignores this parameter. IAM supports the exception exit list in the ACB.

10.04 IDCAMS DEFINE CLUSTER COMMAND

FILE(ddname) - Specifies a ddname.

FILE can specify the name of a DD statement which identifies the volume DEFINE is to allocate space on. If the DD statement named references a device other than one named in the VOLUMES parameter or if the named statement is not present in the JCL, IDCAMS will allocate to a device based on the VOLUMES parameter.

This parameter is also a way, during DEFINE, to associate a 'ddname' with an IAM file. In this way a DEFINE can also specify additional IAM facilities for a file (i.e. DATACOMPRESS) through an Override statement that references this 'ddname'. For example, by specifying FILE (IAMFILE1), the overrides specified by CREATE DD=IAMFILE1 will be used for this data set. There need not actually be an IAMFILE1 DD statement in the JCL.

NOTE: It is preferable that there not be an IAMFILE1 DD statement in the JCL.

A ddname of '&ALLDD' in the IAM Override statement will apply the override to all of the files DEFINEd in the step, unless the file is explicitly overridden as described above.

EXAMPLE: CREATE DD=&ALLDD , B=6 , OVERFLOW=100000

Default: If the 'FILE' parameter is not specified, IDCAMS will allocate the file on a device listed in the VOLUMES parameter.

FREESPACE(CI% CA%) - Specifies the percentage of free space to be left empty in the VSAM cluster when the file is loaded.

NOTE: FREESPACE is ignored by IDCAMS when defining ESDS files and is unavailable to IAM. There is a new GLOBAL OPTION, ESDSINTEGRATED, that allows you to set a default CI% for ESDS files.

CI% - If the file is IAM, this value is used to calculate an Integrated Overflow percent. IAM's concept of Integrated Overflow is essentially the same as CI%. IAM takes the value coded and leaves this percentage of each block empty, to accommodate adds (inserts) to the block. On fixed length files, IAM will round this percent down to an even multiple of the record size. Unless a file is never updated and never has records added to it, some CI freespace should be specified. This is of particular importance to data compressed files that are updated, because even if the application does not change the length, the stored record may end up being longer after compression.

CA% - If the file is IAM, this value is used to calculate Independent Overflow. IAM reserves an area in the file for adds which do not fit in the blocks where they belong. When a block is full, IAM will then add the record to Independent Overflow. This concept makes much more efficient use of disk space than VSAM's CA%.

IAM makes it very easy to tell how much Independent Overflow is used or remaining. The IAM INFO report or a LISTCAT ALL with an IAMPRINT DD statement gives you the exact number of records currently in Independent Overflow and how many records are left for additions. It is next to impossible to get this information from VSAM.

IAM calculates it's Independent Overflow value using the CA% and primary allocation values. Since IAM makes more efficient use of overflow, the CA% specified is cut in half for files with more than 10 cylinders.

10.04 IDCAMS DEFINE CLUSTER COMMAND

Based on the size of the file's primary allocation, IAM will not reduce the CA% below the values shown in the following table:

Primary Allocation	Minimum CA%
1 cylinder or less	10%
2 to 3 cylinders	9%
4 to 5 cylinders	8%
6 to 7 cylinders	7%
8 to 9 cylinders	6%
10 cylinders or more	5%

IAM calculates the approximate number of maximum size records that will fit in the primary allocation and multiplies the result by CA percent. This value will become the number of records in Independent Overflow.

The maximum number of records that IAM will reserve based on this calculation is 50,000, the minimum is 500 records.

NOTE: If CA% is not specified, or specifies zero (0), IAM will take a default value of the minimum CA% as shown in the above table. If you wish to directly specify the number of records for Independent Overflow, see IAM Override statements (Section 11.01). The keyword is OVERFLOW=nnnnnn.

EXAMPLE: If the file you are creating will never have any records added to it, specify the following IAM Override statement:

```
CREATE DD=&ALLDD , OVERFLOW=0
```

SMP/E DEFAULT: IAM will recognize an SMP/E CSI file and will use defaults of 20% Integrated Overflow and 50,000 Independent Overflow records. If you are doing ZONEMERGES, or other operations which add large numbers of records to the CSI, you may have to increase the Independent Overflow value.

IMBED - Specifies the VSAM index is to be embedded in the data.

NOTE: Not applicable to ESDS files.

REPLICATE - Specifies the index is to be replicated on the track.

NOTE: Not applicable to ESDS files.

IAM files use a Super-Compressed index. This index is always loaded into storage eliminating all I/O to the index. IAM, therefore does not need an embedded index and ignores these parameters.

INDEXED - specifies an indexed KSDS file is to be created. IAM supports single indexed and ESDS files.

NONINDEXED creates a sequential VSAM (ESDS) file. **NUMBERED** creates a relative record VSAM (RRDS) file. **LINEAR** creates sequential files addressed by memory location. RRDS and LINEAR files are not indexed by key and the DEFINE will fail if used with a \$IAM parameter.

DEFAULT: INDEXED

KEYRANGES(low high) - Specifies that VSAM is to load the key values specified into separate components.

IAM files do not have components. IAM will load all the records into one file. IAM will make the secondary allocation value equal to the primary allocation whenever KEYRANGES is coded.

10.04 IDCAMS DEFINE CLUSTER COMMAND

KEYS(length offset) - The first value specifies the length of the file's record key. The second value is the relative key position of that field in the record. This parameter is used by IAM the same as VSAM would.

NOTE: For a variable length format file, IAM reports will show an internal 4 byte key offset added to the relative key position.

Not applicable to ESDS files.

Default: length=64 offset=0

MGMTCLAS(management class name) - For SMS installations, this parameter specifies the name of the SMS Management Class for the new data set.

MODEL(datasetname) - Specifies the attributes of the file being DEFINEd are to be copied from an existing VSAM data set.

NOTE: The model data set can not be an IAM file. Also, when using MODEL with ANYVOL, the SUBALLOCATION parameter must also be specified. Although it is ignored by IAM, it prevents IDCAMS from trying to allocate ANYVOL.

If you are copying an IAM file from an existing VSAM file, this parameter is an easy way to get the attributes of the VSAM cluster. You cannot use this parameter to point to an existing IAM file.

RECATALOG - Specifies that the VSAM file is to be recataloged using information contained in the VVDS.

The IAM file is an OS data set which in non-SMS systems does not exist in the VVDS. If an IAM file becomes uncataloged, but still exists on the disk, this parameter will reestablish the IAM file's VIF control information. A '\$IAM' parameter must also be specified to identify it as an IAM file.

NOTE: Recatalog can be used whenever an IAM file exists on disk but is not cataloged. This parameter supports ICF and VSAM catalogs; it is not applicable to CVOLS.

RECORDSIZE(average maximum) - The first value specifies the average record length in bytes. The second value is the maximum record length. IAM assumes a variable length (VL) record format if the average and the maximum record lengths are different. IAM assumes fixed length (FB) format records when these values are equal. The IAM Global Options Table can be used to override this and indicate that all files are to be eligible for data compression and treated as variable.

WARNING: VSAM considers all files capable of containing variable length records. IAM will not, unless the IAM Global Options have been changed. If a file contains records of varying lengths, ensure RECORDSIZE has an average size value that is less than the maximum size value. This will DEFINE the file to IAM as VARIABLE. When DATA COMPRESS is specified for an IAM file, IAM will internally convert the file to a variable length format. IAM reports will show 4 bytes added to the record length for variable length format files.

NOTE: ESDS files are always defined as variable

Default: average=4089 maximum=4089

REUSE - Specifies the VSAM file can be reloaded.

NOREUSE - Specifies that the VSAM cluster cannot be reloaded. These parameters have no meaning for an IAM DEFINE and are ignored. You never have to DELETE and DEFINE IAM files in order to reload them. IAM files are always re-usable.

NOTE: When IDCAMS REPRO is used to reload an existing IAM file the REUSE parameter must be specified on the REPRO command statement.

NOTE: There is a new GLOBAL OPTION, ENABLE=NOREUSE, which when enabled will cause the NOREUSE option to be honored on an IDCAMS define of an IAM FILE. IAM files defined with NOREUSE must then be deleted and redefined to be reloaded.

10.04 IDCAMS DEFINE CLUSTER COMMAND

SHAREOPTIONS(cross-region cross-system) - Specifies how a file can be shared between different jobs.

The first parameter specifies how a file can be shared in the same system (CPU). The second parameter specifies how a file is shared between systems.

NOTE: IAM supports the cross-region share options the same as VSAM. IAM does not support the cross-system share options. IAM issues enqueues with a major name (QNAME) of IAMENQ and the dataset name as the minor name (RNAME). If you need to enforce enq protection cross-system then you must add the major name of IAMENQ to your MIM or GRS control files or whatever ENQ control product you use.

Share Option

- 1 - Any number of users for read or one user for update.
- 2 - Any number of users for read and one user for update.
- 3 - Any number of users for read or update and users are responsible for integrity.
- 4 - Any number of users for read or update and users are responsible for integrity. Buffers are refreshed for each request.

Because of differences in IAM's file structure Share Options 3 or 4 may not yield the same results, if records are being added to Independent Overflow. Contact INNOVATION if you have any questions.

Default: SHAREOPTIONS(1 3)

SPANNED - Specifies that the VSAM file contains records which are written across multiple CIs.

IAM can have a blocksize up to 32760 bytes. IAM block sizes are not restricted to 512 or 2048 byte multiples. If the CISIZE specified is smaller than the maximum record size, IAM will increase its blocksize to accommodate the record. SPANNED is not needed for IAM files and is ignored.

SPEED - Specifies that VSAM is not to preformat each CA.

RECOVERY - Specifies that VSAM is to preformat each CA on loads.

RECOVERY is a very expensive operation. IAM does not have the concept of Control Areas. IAM ignores this parameter. If an IAM file load is interrupted, there is no need to DELETE the IAM file. The file load may simply be resubmitted.

Default: SPEED, the data set is not preformatted.

STORCLAS(Storage class name) - For SMS installations, this parameter specifies the name of the SMS Storage Class construct. For data sets that are to be placed on SMS volumes, Storage Class must be either implicitly specified by the ACS routines, or explicitly specified on the DEFINE command. If the Storage Class Name contains the literal \$IAM, the file will be defined as an IAM file.

UNIQUE - When the DEFINE is made to a VSAM Catalog (i.e. not ICF), this parameter specifies that VSAM is to establish a space allocations for the exclusive use of the components of this cluster. It is recommended that UNIQUE not be used because IDCAMS will dynamically allocate the volumes indicated by the VOLUME parameter before IAM receives the request.

10.04 IDCAMS DEFINE CLUSTER COMMAND

SUBALLOCATION - When the DEFINE is made to a VSAM Catalog (i.e. not ICF), this parameter specifies that VSAM is to establish an allocation for the components of this cluster within a VSAM data space. VSAM ignores this parameter if the VSAM DEFINE is made to an ICF Catalog as all ICF cataloged clusters are established as unique.

IAM files are always allocated as if they were unique clusters. However it may be necessary to specify this parameter when using the IAM non-specific allocation (ANYVOL) and the MODEL parameter, as SUBALLOCATION does prevent IDCAMS from allocating the volumes indicated in the VOLUME parameter when the MODEL parameter is specified.

NOTE: It is recommended that UNIQUE not be specified in a DEFINE for an IAM file. If UNIQUE is specified, FILE should be used to supply a ddname as described under the FILE parameter.

Default: IAM files are never allocated within a VSAM data space.

TO(date) FOR(days) - Specifies the retention period for the file being DEFINEd. This parameter has the same meaning for an IAM file as a VSAM file.

TO(date) - gives the date in the form YYDDD (two digit year and three digit Julian date), through which the IAM file DEFINEd is to be date protected.

FOR(days) - gives the number of days up to a maximum of 9999, through which the IAM file DEFINEd is to be date protected.

Default: Data set is not date protected.

WRITECHECK - Specifies that VSAM will reread all data written without data transfer. IAM ignores this parameter.

10.05 IDCAMS DEFINE EXAMPLES

10.05 IDCAMS DEFINE EXAMPLES

CONVERTING A DEFINE PROCEDURE 1. The following example shows how an IDCAMS command originally written to DEFINE a VSAM file is changed to create an IAM file.

<u>Original VSAM DEFINE</u>	<u>IAM DEFINE</u>
//S1DEF01 EXEC PGM=IDCAMS	//S1DEF01 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*	//SYSPRINT DD SYSOUT=*
//SYSIN DD *	//SYSIN DD *
DEFINE CLUSTER -	DEFINE CLUSTER -
(NAME(SAMPLE.MASTER) -	(NAME(SAMPLE.MASTER) -
VOLUMES(SAMP01) -	VOLUMES(SAMP01) -
SPEED REUSE) -	SPEED REUSE) -
DATA (RECORDS(5000 500) -	DATA (RECORDS(5000 500) -
FREESPACE(20 10) -	FREESPACE(20 10) -
KEYS(10 0) -	KEYS(10 0) -
RECORDSIZE(256 256)) -	RECORDSIZE(256 256)) -
INDEX (IMBED) -	INDEX (IMBED) -
CATALOG(USER.ICF.CATALOG)	CATALOG(USER.ICF.CATALOG)
	ADD--> OWNER(\$IAM) - <--ADD

10.05 IDCAMS DEFINE EXAMPLES

DEFINE USING \$IAM IN THE NAME 2. The following example shows an IDCAMS DEFINE of an IAM file specifying '\$IAM' in the data set name. The file will be allocated with 100 cylinders on volume 'SAMP01'. FREESPACE is DEFINEd as 10% CI and 5% CA. The file contains variable length records with a maximum record size of 480 bytes and average size of 250.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
DEFINE       CLUSTER -
            (NAME(PAYROLL.$IAM.MASTER) -
            VOLUMES(SAMP01) ) -
DATA -
            (NAME(PAYROLL.MASTER.DATA) -
            CYLINDERS(100 20) -
            FREESPACE(10 5) -
            KEYS(24 8) -
            RECORDSIZE(250 480)) -
INDEX -
            (NAME(PAYROLL.MASTER.INDEX))
```

IDCAMS DEFINE FOR 'PAYROLL.\$IAM.MASTER'			
IDCAMS SYSTEM SERVICES	DATE: 15:05:43	08/29/88	PAGE 1
	DEFINE	CLUSTER -	
		(NAME(PAYROLL.\$IAM.MASTER) -	
		VOLUMES(SAMP01)) -	
	DATA -		
		(NAME(PAYROLL.MASTER.DATA) -	
		CYLINDERS(100 20) -	
		FREESPACE(10 5) -	
		KEYS(24 8) -	
		RECORDSIZE(250 480)) -	
	INDEX -		
		(NAME(PAYROLL.MASTER.INDEX))	
IDC0508I	DATA ALLOCATION STATUS FOR VOLUME SAMP01 IS 0		
IDC0509I	INDEX ALLOCATION STATUS FOR VOLUME SAMP01 IS 0		
IDC0001I	FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0		

NOTE: An IAM file is a single space allocation with a data set name the same as the cluster name given in the DEFINE. The status line for the INDEX allocation in the DEFINE report does not mean that a separate INDEX was allocated. A status of zero (0) means the DEFINE for the IAM file was successful.

10.05 IDCAMS DEFINE EXAMPLES

DEFINE WITH AN OVERRIDE

3. The following example shows an IDCAMS DEFINE of an IAM file using the OWNER parameter. Since this is a read only file, an override statement is used to change the Independent Overflow default for the file and limit it to 100 records.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
  DEFINE      CLUSTER -
              (NAME(MASTER.TABLE.CLUSTER) -
               OWNER($IAM) -
               CYLINDERS(600 100) -
               FREESPACE(0 0) -
               KEYS(20 0) -
               RECORDSIZE(80 200) -
               VOLUMES(SAMP01) )
/*
//IAMOVRID   DD    *
*****      THE FOLLOWING STATEMENT WILL OVERRIDE          *****
*****      THE DEFAULT INDEPENDENT OVERFLOW VALUE        *****
              CREATE DD=&ALLDD,OVERFLOW=100
/*
```

RECATALOG AN IAM FILE WITH DEFINE

4. The following example shows a recatalog of an IAM file using IDCAMS. This procedure may be used in lieu of a DELETE, DEFINE and reload to reestablish an IAM file's VIF control information in the event a file is ever uncataloged. The RECATALOG parameter may be used whenever an IAM file exists on disk but has not been DEFINEd in a catalog for VIF processing; for example existing IAM files established under previous versions of IAM or an IAM file allocated and even cataloged through some means other than an IDCAMS define.

```
//RECAT      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
  DEFINE      CLUSTER -
              (NAME(MASTER.TABLE.CLUSTER) -
               OWNER($IAM) -
               VOLUMES(SAMP01) -
               RECATALOG)
/*
```

NOTE: Recataloging IAM files after they have been moved by utility software such as FDR,(tm) DFDDSS,(tm) or DMS(tm) is no longer necessary. It is still recommended, however, because it will minimize the overhead needed for IAM file identification.

IDCAMS DEFINE WITH SMS ALLOCATIONS

5. In this example, an IDCAMS DEFINE is done for an IAM file, with all file attributes being supplied from the specified Data Class. The file is allocated as an IAM file because of the \$IAM literal in the Data Class name.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
  DEFINE      CLUSTER -
              (NAME(TSTSMS.SMS202.CLUSTER) -
               MGMTCLAS(DBSTNDRD) -
               STORCLAS(BASE) -
               DATACLAS($IAMTST1))
```

10.06 IDCAMS DELETE COMMAND AND EXAMPLES

10.06 IDCAMS DELETE COMMAND AND EXAMPLES

Deleting an IAM file is the same as a VSAM cluster. Use the IDCAMS DELETE or DELETE CLUSTER command. These commands delete and uncatalog the IAM data set. When running under IAM's VSAM Interface (VIF), the DELETE CLUSTER command will delete both VSAM clusters and non-VSAM files.

IAM files can also be DELETED in JCL statements 'DISP=(OLD,DELETE)' and by the program IEHPROGM.

NOTE: If the VIF Interface is not active in your system a DELETE NONVSAM command must be used to DELETE an IAM file.

DELETE EXAMPLE The following example will DELETE two IAM files.

```
//DELETE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
      DELETE  IAM.FILE.ONE CLUSTER
      DELETE  IAM.FILE.TWO
/*
```

IDCAMS DELETE OF 'IAM.FILE.ONE' and 'IAM.FILE.TWO'

```
IDCAMS SYSTEM SERVICES          TIME: 15:05:43          08/29/88          PAGE 1

      DELETE IAM.FILE.ONE CLUSTER
IDC05501 ENTRY (A) IAM.FILE.ONE DELETED
      DELETE IAM.FILE.TWO
IDC05501 ENTRY (A) IAM.FILE.TWO DELETED

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0
```

10.07 IDCAMS LISTCAT COMMAND AND EXAMPLES

10.07 IDCAMS LISTCAT COMMAND AND EXAMPLES

IAM's VSAM Interface (VIF) supports the IDCAMS LISTCAT command. If a LISTCAT is issued against an IAM file, IDCAMS will report the file as non-VSAM. IDCAMS will also report on which volumes the file is cataloged to. The OWNER field is changed by IAM and will be shown as periods, unless IAM usage of the OWNER field is turned off in the IAM Global Options Table.

IAMPRINT specifies an output message data set, usually SYSOUT, and is coded as follows:

```
//IAMPRINT DD SYSOUT=*
```

A report detailing the IAM file characteristics will be generated to the output data set identified by the IAMPRINT DD card. If an IAMPRINT DD card is not currently allocated, IAM will automatically dynamically allocate an IAMPRINT DD to SYSOUT=*, if a SYSPRINT DD is present in the step, and is also going to a SYSOUT data set. Examples of the IAMPRINT report are shown below.

LISTCAT The following example shows a LISTCAT ALL for a set of files. The LISTCAT ALL is used and
LEVEL LEVEL is specified for the data set group USER. IDCAMS will report on all data sets with a high index level of USER. Any IAM file with a high level index of 'USER', will also be included in the IAM file attribute report generated to the optional IAMPRINT DD statement.

```
//LISTCAT EXEC PGM=IDCAMS  
//SYSPRINT DD SYSOUT=*  
//IAMPRINT DD SYSOUT=*  
//SYSIN DD *  
LISTCAT LEVEL(USER) ALL  
/*
```

NOTE: If you are listing out a large number of cataloged datasets, and do not want the IAMPRINT, add the following DD statement:

```
//IAMNOLIC DD DUMMY
```

10.07 IDCAMS LISTCAT COMMAND AND EXAMPLES

LISTCAT ALL The following example shows a LISTCAT ALL for a specific file. The LISTCAT ALL must be used to obtain a report containing the IAM statistics. ENTRIES is specified for the data set 'FILE.\$IAM.CLUSTER'. IDCAMS will report on the data set and IAM will dynamically allocate an IAMPRINT DD to SYSOUT=*, and generate its report.

```
//LISTCAT      EXEC   PGM=IDCAMS
//SYSPRINT    DD     SYSOUT=*
//SYSIN       DD     *
LISTCAT      ENTRIES(FILE.$IAM.CLUSTER) ALL
/*
```

LISTCAT FOR 'FILE.\$IAM.CLUSTER'			
IDCAMS SYSTEM SERVICES	TIME: hh:mm:ss	mm/dd/yy	PAGE 1
LISTCAT ENTRIES(FILE.\$IAM.CLUSTER) ALL			
NONVSAM -----	FILE.\$IAM.CLUSTER		
IN-CAT ---	CATALOG.PICFUSER		
HISTORY			
OWNER-IDENT-----	CREATION-----	yyyy.ddd	
RELEASE-----2	EXPIRATION-----	0000.000	
VOLUMES			
VOLSER-----	SCR083	DEVTYPE-----	X'3010200E' FSEQN-----4
ASSOCIATIONS----- (NULL)			
IDCAMS SYSTEM SERVICES	TIME: 15:05:43	08/29/88	PAGE 2
THE NUMBER OF ENTRIES PROCESSED WAS:			
AIX -----	0		
CLUSTER -----	0		
NONVSAM -----	1		
TOTAL -----	1		
THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0			
IDC00011 FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0			
IDC00021 IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0			

NOTE: An IAM file is always shown as NONVSAM with an OWNERIDENT field that displays as a string of periods and/or characters.

10.07 IDCAMS LISTCAT COMMAND AND EXAMPLES

```

IAMPRINT REPORT FOR 'FILE.$IAM.CLUSTER'
AFTER DEFINE BUT BEFORE LOAD AND USE

IAM400 IAM CATALOG INFORMATION SERVICE -- VERSION 6.x -- DATE yy.ddd    TIME hh:mm:ss
IAM100 IAM FILE ANALYSIS - DSN=FILE.$IAM.CLUSTER

```

RECORD SIZE -- =	670	-	INTEGRATED OVERFLOW --- =	6 PER	-
BLOCK FACTOR - =	4096	-	INDEPENDENT OVFL REQ -- =	612 RECS	-
KEY SIZE ----- =	4	-	INDEPENDENT OVERFLOW -- =	0 BLOCKS	-
REL KEY POS -- =	8	-	INDEPENDENT OPEN ----- =	0 PER	-
RECORD FORMAT. =	FIXED	-	INDEPENDENT FULL ----- =	0 PER	-
FILE TYPE ---- =	KSDS	-	PRIME EXTENSION ----- =	0 BLOCKS	-
DEVICE TYPE -- =	3380	-	PRIME EXTENSION USED -- =	0 BLOCKS	-
VOLUME COUNT - =	1	-	SPACE USED ----- =	N/A TRACKS	-
VOLSER ----- =	SCR083	-	SPACE ALLOCATED ----- =	90 TRACKS	-
TOTAL EXTENTS =	1	-	TOTAL SPACE ALLOCATED - =	90 TRACKS	-
RELEASE ----- =	YES	-	SHARE OPTIONS ----- =	2	-

```

IAM499 IAMLSTC(6.x) PROCESSING COMPLETED

```

NOTE: The IAMPRINT file attributes report will be different for an IAM file that has only been DEFINEd, (that is never been used) when compared to a file that has been loaded with records. In the sample report above is for a file that has not been loaded.

```

IAMPRINT REPORT FOR 'FILE.$IAM.CLUSTER'
AFTER LOAD AND USE

IAM400 IAM CATALOG INFORMATION SERVICE -- VERSION 6.x -- DATE yy.ddd    TIME hh:mm:ss
IAM100 IAM FILE ANALYSIS - DSN=FILE.$IAM.CLUSTER

```

RECORD SIZE -- =	674	-	INTEGRATED OVERFLOW --- =	6 PER	-
BLOCK SIZE --- =	11476	-	ACTUAL INTEGRATED OVFL =	6 PER	-
BLOCK FACTOR - =	4	-	INDEPENDENT OVFL REQ -- =	612 RECS	-
KEY SIZE ----- =	4	-	INDEPENDENT OVERFLOW -- =	36 BLOCKS	-
REL KEY POS -- =	12	-	INDEPENDENT OPEN ----- =	10 PER	-
RECORD FORMAT =	VARIABLE	-	INDEPENDENT FULL ----- =	10 PER	-
FILE TYPE ---- =	KSDS	-	PRIME EXTENSION ----- =	4 BLOCKS	-
DEVICE TYPE -- =	3380	-	PRIME EXTENSION USED -- =	0 BLOCKS	-
VOLUME COUNT - =	1	-	SPACE USED ----- =	53 TRACKS	-
VOLSER ----- =	SCR083	-	SPACE ALLOCATED ----- =	60 TRACKS	-
TOTAL EXTENTS =	1	-	TOTAL SPACE ALLOCATED - =	60 TRACKS	-
RELEASE ----- =	NO	-	SHARE OPTIONS ----- =	2	-
DATA COMPRESS =	ENABLED	-	INDEX COMPRESS ----- =	NO	-
TOTAL RECORDS =	4608	-	INSERTS ----- =	512	-
UPDATES ----- =	64	-	DELETES ----- =	0	-
PERCENT AVAILABLE IN INDEPENDENT OVERFLOW ---- = 58 -					
NUMBER OF RECORDS IN INDEPENDENT OVERFLOW ---- = 256 -					
NO OF UNUSED RECORDS IN INDEPENDENT OVERFLOW - = 356 -					
STORAGE REQUIRED FOR KEY STRUCTURE ----- = 684 BYTES -					

```

IAM499 IAMLSTC(6.x) PROCESSING COMPLETED

```

NOTE: The IAMPRINT report above is for a file that was loaded after it was DEFINEd It contains fields for the number of records in the file as well as a description of how the overflow was used.

10.07 IDCAMS LISTCAT COMMAND AND EXAMPLES

IAMPRINT REPORT FIELDS Fields and their meanings in the IAMPRINT report which accompanies the IDCAMS LISTCAT ALL for an IAM file are as follows:

RECORD SIZE: The logical fixed or maximum variable record length in the IAM file, given in bytes.

NOTE: This value will be 4 higher than the maximum record size specified in the IDCAMS DEFINE command for files with variable length records.

BLOCKSIZE: The actual blocksize of the IAM file, given in bytes.

NOTE: This field is not reported on until the file has been loaded.

BLOCK FACTOR: The blocking factor specified when this file was last DEFINEd, loaded, or reorganized. If the file has been DEFINEd, but not yet loaded, this value will reflect the CISIZE, if any, specified on the DEFINE, which will be converted to a blocking factor when the file is loaded.

KEY SIZE: The length of the record key field, in bytes.

REL KEY POS: The location of the key relative to the start of the record (RKP), in bytes. The first byte in a record is RKP=0. This amount will be 4 higher than the value DEFINEd for files with variable length records, due to IAM's internal RDW.

RECORD FORMAT: The record format used in this file, either FIXED or VARIABLE length records. When a file is DEFINEd, IAM assumes the file to contain fixed length records if the average and maximum record lengths are equal, or variable if average and maximum record lengths are not equal. IAM internally sets DATACOMPRESS files as variable.

FILE TYPE: Indicates the type of file being simulated, a KSDS.

DEVICE TYPE: Indicates the disk device type that the IAM file resides on.

VOLUME COUNT: Indicates the number of volumes that are currently in the catalog for the IAM file.

VOLSER: Indicates the volume on which the IAM file is currently cataloged to. For multi-volume files, additional VOLSERS are printed on subsequent lines. The additional VOLSERS may represent candidate volumes, and may not yet contain any extent of the IAM file.

SPACE ALLOCATED: Indicates the total amount of space, in tracks, which are currently allocated to the IAM file on the volume in the left column.

TOTAL SPACE ALLOCATED: Indicates the cumulative total for all volumes of the space, in tracks, which are currently allocated to the IAM file.

TOTAL EXTENTS: Indicates the total number of extents allocated to the IAM file, across all of the volumes.

RELEASE: Indicates whether or not IAM will automatically release space the next time the file is loaded or reorganized. Once an IAM file is loaded and the unused portion of its allocation is released IAM will automatically change RELEASE=YES to RELEASE=NO, unless overridden by an IAM CREATE override statement, after the first load of the file.

INTEGRATED OVERFLOW: The Integrated Overflow (CI%) percentage specified when the file was last DEFINEd, loaded, or reorganized.

ACTUAL INTEGRATED OVFL: Indicates the percentage used in establishing the Integrated Overflow (CI%) for the IAM file. This figure may be lower than the requested Integrated Overflow due to rounding based on the record length and actual blocksize being used.

NOTE: This is reported as 0 (zero) until the file has been loaded, unless overridden at DEFINE time.

INDEPENDENT OVFL REQ: The number of Independent Overflow records specified when this IAM file was last DEFINEd, loaded, or reorganized.

INDEPENDENT OVERFLOW: The number of physical blocks IAM reserved for Independent Overflow records. This number is developed by IAM during file loading using the number of Independent Overflow records requested, the maximum record size, and the actual blocksize being used.

NOTE: This is reported as 0 (zero) until the file has been loaded, unless overridden at DEFINE time.

10.07 IDCAMS LISTCAT COMMAND AND EXAMPLES

INDEPENDENT OPEN: Specifies the percentage of Independent Overflow keys which will be reserved when the file is OPENed for update operations, as specified by the OCOREO% keyword.

INDEPENDENT FULL: The percentage of Independent Overflow keys for which storage is acquired when the current index tables are full, as specified by the OCOREX% keyword.

NOTE: This field is reported as 0 (zero) until the file has been loaded, unless overridden at DEFINE time.

PRIME EXTENSION: The number of blocks reserved by IAM for resume load type of operations when this file was last DEFINEd, loaded, or reorganized.

NOTE: This field is reported as 0 (zero) until the file has been loaded, unless overridden at DEFINE time.

PRIME EXTENSION USED: This field indicates the maximum number of prime extension blocks that have been used to store data records.

SPACE USED: The actual number of tracks used to accommodate this IAM file in its entirety, including overflow areas.

NOTE: This field is reported as N/A if the file has not yet been loaded, or IAM is unable to determine or calculate the number of tracks used for this file.

SPACE ALLOCATED: Indicates the total amount of space, in tracks, which are currently allocated to the IAM file on the volume in the left column.

TOTAL SPACE ALLOCATED: Indicates the cumulative total for all volumes of the space, in tracks, which are currently allocated to the IAM file.

SHARE OPTIONS: Indicates the VSAM Cross Region share option that was specified when the IAM file was DEFINEd.

DATA COMPRESS: Indicates whether or not the IAM file can contain compressed data. Values are either ENABLED or NO. ENABLED, indicates that the DATA COMPRESS option was either specified as YES or defaulted to YES, based on the global options table, and the IAM file met all other required criteria for data compression. NO, indicates the IAM file does not contain any compressed data records.

NOTE: This field is not reported on until the file has been loaded.

INDEX COMPRESS: Indicates whether or not the IAM file contains a compressed index structure. Values are either YES, indicating a compressed key index structure exists, or NO, indicating that the file does not have a compressed key structure.

NOTE: This field is not reported on until the file has been loaded.

MINBUFNO: Indicates the minimum number of buffers IAM will use when accessing an IAM file.

NOTE: This field is only reported on if a value was specified on a CREATE override statement for MINBUFNO or MAXBUFNO when the file was DEFINEd, loaded, or reorganized. A value of N/A will be printed if no value was specified for MINBUFNO, but a MAXBUFNO value was specified.

MAXBUFNO: Indicates the maximum number of buffers IAM will use when accessing an IAM file.

NOTE: This field is only reported on if a value was specified on a CREATE override statement for MINBUFNO or MAXBUFNO when the file was DEFINEd, loaded, or reorganized. A value of N/A will be printed if no value was specified for MAXBUFNO, but a MINBUFNO value was specified.

10.07 IDCAMS LISTCAT COMMAND AND EXAMPLES

THE FOLLOWING FIELDS WILL NOT BE REPORTED FOR IAM FILES THAT HAVE BEEN DEFINED, BUT NOT YET LOADED, OR HAVE EXPERIENCED A FAILURE DURING FILE LOAD.

TOTAL RECORDS: Indicates the total number of user data records currently residing in the file.

INSERTS: Indicates the total number of records added to the file since it was last loaded or reorganized.

UPDATES: Indicates the total number of records updated in the file since it was last loaded or reorganized.

DELETES: Indicates the total number of records DELETED from the file since it was last loaded or reorganized.

NOTE: THE ABOVE FOUR STATISTICS MAY NOT BE ACCURATE IF AN UPDATE JOB FAILED TO SUCCESSFULLY CLOSE THE IAM FILE. THIS WOULD ONLY OCCUR IF THE ADDRESS SPACE FAILED, AN UPDATE JOB WAS REMOVED FROM THE SYSTEM WITH THE FORCE COMMAND, OR A SYSTEM FAILURE OCCURRED.

PERCENT AVAILABLE IN INDEPENDENT OVERFLOW: Indicates the percentage of the total Independent Overflow records that are currently available. A low value, in the range of 10 to 25, provides an indication that the file needs to be reorganized.

NUMBER OF RECORDS IN INDEPENDENT OVERFLOW: Indicates the number of records currently residing in the Independent Overflow area of the IAM file.

NO OF UNUSED RECORDS IN INDEPENDENT OVERFLOW: The number of additional records the Independent Overflow area can accommodate. This value added to the value in NUMBER OF RECORDS IN INDEPENDENT OVERFLOW may be slightly higher than the number of Independent Overflow Records Requested, because IAM uses a fully blocked overflow area.

STORAGE REQUIRED FOR COMPRESSED KEY STRUCTURE: The amount of virtual storage required for this IAM file's index. This message only appears if the IAM file has a compressed index.

STORAGE REQUIRED FOR KEY STRUCTURE: The amount of virtual storage required for this IAM file's index. This message only appears if the IAM file does not have a compressed index.

10.09 IDCAMS REPRO COMMAND AND EXAMPLES

10.09 IDCAMS REPRO COMMAND AND EXAMPLES

IAM's VSAM Interface (VIF) supports the IDCAMS REPRO command. The REPRO command will function the same way for an IAM file as a VSAM cluster.

Many applications use IDCAMS REPRO to make backup copies of their VSAM files. REPRO will process much faster after the file is converted to IAM format. If you are using some other product to backup your VSAM files, you may want to consider converting these jobs to use IDCAMS REPRO or FDR data set backups.

Use IDCAMS REPRO to copy your existing VSAM files into IAM files. You can also use REPRO to copy IAM files back to VSAM clusters.

COPYING TO A NEW NAME

The following example will copy a VSAM cluster to an IAM file with a new name. First, DEFINE an IAM file, use the existing VSAM file as a model and add the OWNER(\$IAM) parameter. The data set name of the IAM file must be different from that of the VSAM cluster, or it must be cataloged in a separate catalog. Next, REPRO the existing VSAM file to the IAM file.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT    DD    SYSOUT=*
//IAMINFO     DD    SYSOUT=*
//SYSIN       DD    *
  DEFINE      CLUSTER -
    (NAME(new.IAM.file.name) -
    MODEL(existing.VSAM.cluster) -
    OWNER($IAM) -
    VOL(volume) )
  REPRO       INDATASET(existing.VSAM.cluster) -
    OUTDATASET(new.IAM.file.name)
/*
```

10.09 IDCAMS REPRO COMMAND AND EXAMPLES

CONVERTING AN EXISTING CLUSTER If you want to create an IAM file with the same name as an existing VSAM cluster use the following procedure:

1. REPRO the VSAM cluster to a sequential file on tape or disk.
2. DELETE the VSAM cluster.
3. DEFINE the IAM file with the same names and attributes as the VSAM cluster.
4. REPRO the sequential file into the IAM file.

The following sample JCL and IDCAMS control statements will copy a VSAM cluster to an IAM file of the same name. The DD statement 'SEQFILE' specifies a sequential file on disk (or tape) to copy the VSAM file into. It's space allocation must be sufficient to hold the unloaded VSAM file.

```
//DEFINE          EXEC   PGM=IDCAMS
//SYSPRINT       DD     SYSOUT=*
//IAMINFO        DD     SYSOUT=*
//SEQFILE        DD     DSN=existing.VSAM.backup,
//      DISP=(,CATLG),DCB=(RECFM=VB,BLKSIZE=nnnn),
//      UNIT=SYSDA,SPACE=(TRK,(1000,100))
//SYSIN          DD     *
      REPRO          INDATASET(existing.VSAM.name) -
                    OUTFILE(SEQFILE)
      DELETE          existing.VSAM.name
      DEFINE          CLUSTER -
                    (NAME(existing.VSAM.name) -
                    RECORDSIZE(avg max) KEYS(length offset) -
                    IMBED SPEED REUSE -
                    FREESPACE(ci% ca%) -
                    RECORDS(primary secondary) -
                    OWNER($IAM) -
                    VOL(volume) )
      REPRO -
                    INFILE(SEQFILE) -
                    OUTDATASET(existing.VSAM.name)
/*
```

10.10 REORGANIZING AN IAM FILE

10.10 REORGANIZING AN IAM FILE

Files which have records added require periodic reorganization. Reorganization is the process of reloading a file to merge the records added into overflow back into the prime data blocks, freeing up the overflow areas for future inserts.

If a file is not reorganized, it will eventually fill. This condition occurs with both IAM and VSAM files. Within an IAM file, a file full occurs when you attempt to insert a record into a prime data block that has no free space remaining and the Independent Overflow area is exhausted. Inserts to another prime data block may succeed, if there is free space available within that block.

An IAM INFO or LISTCAT ALL (with IAMPRINT) report will show the number of unused Independent Overflow records left in an IAM file.

Use IDCAMS REPRO to reorganize your IAM files. First, REPRO the IAM file to a sequential file. Second, REPRO the sequential file back into the IAM file specifying REUSE. The sequential file created by the REPRO may be retained as a backup.

The following is sample JCL and IDCAMS control statements to reorganize an IAM file. It is recommended that RECFM=VB and a large blocksize (a value equal to or a multiple of the IAM blocksize) be specified for the sequential file created by a REPRO. If there currently is a RECFM and blocksize, they can be used as is. However, specifying a high block size ensures that minimal disk space or tape volumes are used for the sequential copy of the data.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT    DD    SYSOUT=*
//IAMINFO     DD    SYSOUT=*
//SEQFILE     DD    DSN=IAM.file.backup,DISP=(,CATLG),
//      DCB=(RECFM=VB,BLKSIZE=nnnnn),UNIT=TAPE
//SYSIN       DD    *
      REPRO      INDATASET(IAM.file) -
      OUTFILE(SEQFILE)
      REPRO      REUSE INFILE(SEQFILE) -
      OUTDATASET(IAM.file)
/*
```

NOTE: Do not specify a DISP of DELETE on the sequential backup file. If the REPRO fails on the reload (i.e, an abend or cancel) before the IAM file is CLOSEd, the IAM file will be unusable. The IAM file however can simply be reloaded using the sequential backup file.

10.11 IMPORT AND EXPORT OF AN IAM FILE

10.11 IMPORT AND EXPORT OF AN IAM FILE

REPLACEMENT FUNCTIONS As was stated previously in the IDCAMS overview at the beginning of this section, IAM Version 6.2 does not currently support the IDCAMS IMPORT and EXPORT commands. A replacement for these two operations will be needed when ever they are part of an existing system or when a new system is planned that will need to make removable copies of its files. What to use as a replacement will depend on the objective for doing the IMPORT and/or EXPORT.

USE IDCAMS REPRO Frequently, IMPORT and/or EXPORT can be replaced with the IDCAMS REPRO command, coupled with the DELETE and DEFINE commands as necessary. An advantage of replacing IMPORT and EXPORT with REPRO is that IAM, through VIF, will be the data mover, and the file will be reorganized on the import REPRO, just as a VSAM file would be.

A disadvantage of REPRO is that to *'import'* a DELETED IAM file, the IAM file must first be DEFINED before it can be recovered. If this is seen as a difficulty other alternatives can be considered.

USE ONE OF THE FDR FAMILY OF PRODUCTS The best alternative is to use a disk dump/restore utility, such as ABR or DSF from INNOVATION's FDR family of products. An IAM file 'imported' in this manner will not be a reorganized copy of the file. Use of ABR or DSF however, does eliminate the need to retain copies of the IDCAMS parameters originally used to define the file.

FDR, DSF, CPK, and ABR fully support IAM files as simple NONVSAM files. IAM files can be backed up, ARCHIVED, and SUPERSCRATCHED without any special considerations. ABR and FDRDSF can dynamically allocate, restore and catalog IAM files to the same name or a new name. IAM files must be cataloged. CPK can be used to consolidate extents.

UNLIKE DEVICE RESTORE If a utility like FDR/ABR is used to move the file to an unlike device type, the blocksize must not be changed. The file should then be reorganized. This will allow IAM to select a more appropriate blocksize for the device, as well as to recognize that the file is on a different device type. Refer to the preceding [Section 10.10](#) for an example of reorganizing an IAM file.

'export' AN IAM FILE USING FDRDSF The following example will show how to 'export' a data set using INNOVATION's data set utility FDRDSF. FDRDSF will 'export' the file by using a dump to tape. This is as good a removable copy as you can get. The file could then, at the user's option, be DELETED following the DUMP.

```
//EXPORT      EXEC   PGM=FDRDSF
//SYSPRINT    DD     SYSOUT=*
//DISK1       DD     DSN=my.i amfile .cluster ,DISP=OLD
//TAPE1       DD     DSN=my.i amfile .export ,
//           UNIT=3480 ,DISP=( ,CATLG)
//SYSIN       DD     *
              DUMP TYPE=DSF
              S DSN=my.i amfile .cluster
/*
```

10.11 IMPORT AND EXPORT OF AN IAM FILE

'import' This example will 'import' the data set to disk, with a different name and catalog the new name.
TO A Notice there is no requirement for any of the original DEFINE information. The last step then
NEWNAME performs an IDCAMS LISTC ALL and includes an IAMPRINT DD statement to produce an IAM file
WITH FDRDSF attributes report.

```
//IMPORT      EXEC  PGM=FDRDSF
//SYSPRINT   DD    SYSOUT=*
//TAPE1      DD    DSN=my.i amfile . export ,DISP=OLD
//DISK1      DD    UNIT=SYSDA ,DISP=OLD ,VOL=SER=vvvvvv
//SYSIN      DD    *
      RESTORE TYPE=DSF
      S DSN=my.i amfile . cluster ,NEWN=my.i amfile . restored
/*

//LISTC      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//IAMPRINT   DD    SYSOUT=*
//SYSIN      DD    *
      LISTC ENTRIES(my.i amfile . restored) ALL
/*
```

'export' The following example will show how to 'export' a data set using INNOVATION's DASD
AN IAM Management utility FDRABR. ABR will 'export' the file using an ARCHIVE to tape. This will result in
FILE USING a removable copy and automatically DELETE the data set following the backup unless
ABR ARCHIVE SCRATCH=NO is specified.

```
//EXPORT      EXEC  PGM=FDRABR
//SYSPRINT   DD    SYSOUT=*
//SYSPRIN1   DD    SYSOUT=*
//TAPE1      DD    DSN=FDREXP1 ,UNIT=3480 ,DISP=( ,KEEP )
//ARCHIVE    DD    DSN=fdrabr . archive . xxx ,DISP=SHR
//SYSIN      DD    *
      DUMP TYPE=ARC ,BUFNO=MAX ,ONLVOL ,ARCHBACKUP=NO ,SCRATCH
      SELECT CATDSN=my.i amfile . cluster
/*
```

'import' This example will 'import' a data set that has been ARCHIVED. Notice there is no requirement for
FROM any of the original DEFINE information. The last step then performs an IDCAM LISTC ALL and
AN ABR includes an IAMPRINT DD statement to produce an IAM file attributes report.
ARCHIVE

```
//IMPORT      EXEC  PGM=FDRABR
//SYSPRINT   DD    SYSOUT=*
//SYSPRIN1   DD    SYSOUT=*
//ARCHIVE    DD    DSN=fdrabr . archive . xxx ,DISP=SHR
//ABRWORK    DD    UNIT=SYSDA ,SPACE=(CYL ,( 1 , 1 ))
//SYSIN      DD    *
      RESTORE TYPE=ARC ,DYNTAPE ,ONLINE
      SELECT DSN=my.i amfile . cluster
/*

//LISTC      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//IAMPRINT   DD    SYSOUT=*
//SYSIN      DD    *
      LISTC ENTRIES(my.i amfile . restored) ALL
/*
```

IAM

INNOVATION ACCESS METHOD

IAM OVERRIDES

THIS PAGE WAS INTENTIONALLY LEFT BLANK

SECTION 11: IAM OVERRIDE STATEMENTS

11.01 IAM OVERRIDE STATEMENTS OVERVIEW

11.01 IAM OVERRIDE STATEMENTS OVERVIEW

OVERVIEW IAM's file structure is different from VSAM's. Because of this some DEFINE parameters do not have the same meaning to IAM as to VSAM. IAM's Override statements can be used to directly override IAM Global Option Table and IDCAMS DEFINE defaults. Overrides can be applied when a file is DEFINEd at execution time as a file is loaded or is being accessed.

IAM Override statements can also be used to specify special IAM file characteristics and features not available in VSAM.

IAM offers the user two override statements.

The **CREATE** statement can be used when a file is DEFINEd, loaded or reorganized to specify file characteristics and special features.

The **ACCESS** statement can be used when a file is processed to override IAM run time options for a specific job.

CREATE STATEMENT Using the CREATE statement, an IAM user can tailor file processing by changing or overriding file characteristics and options that would otherwise be based on an IDCAMS DEFINE and the defaults in IAM's Global Option table. In some cases IAM's default values may not be sufficient. For example, a very large number of mass inserts may require the specification of additional Independent Overflow. Use a CREATE statement to override the defaults and specify an appropriate Overflow value.

ACCESS STATEMENT Using the ACCESS statement an IAM user, at execution time, can tailor a file's processing. The user can specify options that do not exist in VSAM (ex: Dynamic Tabling of Data Records) or adjust IAM's Real Time Tuning (ex: MAXBUFNO, MINBUFNO). IAM for example dynamically changes the number of buffers acquired for a file as the file is being processed. If you wish to change the limits on the minimum or maximum number of buffers IAM will acquire, use an ACCESS statement.

IAMOVRID DD STATEMENT Specifies the Override control data set. This data set is in card image format. Usually an input stream (i.e. DD *) data set. 'IAMOVRID' may be used in any job that processes IAM files (IDCAMs, CICS, batch applications, etc.). The IAMOVRID file is only read once, when the first IAM file in the step is opened. A table of the overrides is built in storage, and is referenced for other IAM file opens. The table remains in storage until step termination. Up to two hundred override statements can be provided.

OVERRIDE STATEMENT FORMAT IAM Override statements are one or more 80 character records coded in positions 1 through 71. Each control statement starts with a command (ex: CREATE) and is followed by one or more optional keywords. The last keyword in a statement must be followed by one or more blanks. Anything following the blank will be treated as comments.

Comments are not permitted prior to the first keyword in a statement. If * is coded in position 1 of the statement, the entire line will be treated as a comment.

To continue a statement on an additional line the last keyword on the line must be followed by a comma and at least one blank. The next keyword is then begun on the following line. When a keyword requires an optional variable, that keyword and its variable must start and end on the same line.

[See Section 00.03 Control Statement Format](#) for further detail.

11.02 IAM CREATE OVERRIDE STATEMENT - FORMAT

11.02 IAM CREATE OVERRIDE STATEMENT - FORMAT

CREATE COMMAND The CREATE statement can be used to override file characteristics that are stored in a file and processing options that take effect when a file is DEFINEd or loaded.

CREATE

	,MAXBUFNO=nn
DDNAME=ddname	,MINBUFNO=nn
,BLKSIZE=nnnnn	,MINCOMPRESS=nn
,BUFNO=nn	,OCOREO%=nn
,CORELIMIT=nnnnn	,OCOREX%=nn
,CRBUFOPT=xxxx	,OVERFLOW=nnnnnnn
,DATACOMPRESS=YES NO	,PE=nnnnn
,DATASPACE=nnnn	,PSEUDORBA
,DESCRIPTCODE=nn	,RELEASE=YES NO
,FIXED	,RKP=nnnnn
,INDEXCOMPRESS= <u>YES</u> NO	,ROUTECD=nn
,INFO= <u>YES</u> NO	,SMF=YES NO
,INTEGRATED=nn	,TRACEDDNAME=ddname
,KEYLEN=nnn	,UNIT=unitname
,LIMITKEYS=nnn	,VARIABLE
,LOG=YES <u>NO</u>	,VM
,LRECL=nnnnn	,WKDDNAME=ddname

11.02 IAM CREATE OVERRIDE STATEMENT - FORMAT

DDNAME= DD=	<p>Specifies the ddname of the IAM data set that the override is to be applied to. Up to 40 DDNAMEs may be specified. Multiple DDNAMEs must be enclosed in parenthesis.</p> <p>DD=&ALLDD will result in this control statement being applied as a global override to all IAM files being DEFINEd or loaded in this step.</p> <p><u>DDNAME is a required parameter.</u></p>
BLKSIZE= B=	<p>Specifies a blocking factor or block size for the IAM file. For values 1 to 15, IAM will use the file's record length and calculate a blocksize that will best fit that number of blocks on a track for the device in use:</p> <p>1 - full track (limited on 3375, 3380, and 3390 to 1/2 track) 2 - is one-half track blocksize (2 blocks to a track) 3 to 15 - is three to fifteen blocks on a track 300 to track size - is blocksize in bytes.</p> <p>IDCAMS DEFINE default - CISIZE is rounded up to the nearest IAM blocking factor. The minimum being quarter track blocking (4 blocks to a track). If no CISIZE is specified, the default is quarter track blocking or as specified in IAM's Global Option Table. IAM may select a larger blocksize if less than four maximum length records will fit in the default blocksize. The default for other interfaces is 2 (half-track).</p>
BUFNO=	<p>Specifies the number of buffers IAM will use for a file load or REPRO when CRBUFOPT=BSAM is specified. This parameter has no effect on an IDCAMS DEFINE or other file processing.</p> <p>Default is 12 buffers.</p>
CRBUFOPT=	<p>Specifies the buffering option to be used for file loads. For all options except BSAM the buffers are requested above the line on MVS/XA and MVS/ESA systems. Allowable values are:</p> <p>BSAM - BSAM is used, and the number of buffers used is based on BUFNO. CYL - Buffers for one cylinder are getmained, and half a cylinder is written per EXCP. (Effective only for files that are on cylinder boundaries.) MCYL - Buffers for two cylinders are getmained, and one cylinder is written per EXCP. (Effective only for files that are on cylinder boundaries.) MTRK - Buffers for ten tracks are getmained, and five tracks are written per EXCP. (Effective only for files that are on cylinder boundaries.) TRK - Buffers for two tracks are getmained, and one track is written per EXCP.</p> <p>Default is CYL, as indicated in the Global Options Table.</p>
CORELIMIT=	<p>Specifies the size of the smallest index IAM will compress. This value is specified in bytes.</p> <p>Default is an index must be larger than 8000 bytes before IAM will attempt index compression.</p>

11.02 IAM CREATE OVERRIDE STATEMENT - FORMAT

DATACOMPRESS= Specifies if data records are to be compressed. IAM will attempt to compress all data following the record key. When a record does not have 10 bytes of data following its key or if a compressed record would be larger than the original, IAM will leave the record uncompressed. See [Section 02.20](#) IAM File Allocation and Space Savings for further detail on Data Compression.

YES - IAM is to compress the data records in the file.

NO - IAM will not compress the data records.

NOTE: Data compression increases the CPU time needed to load a file over IAM without data compression. There may also be a nominal CPU time increase processing against compressed IAM files. Even with data compression however, IAM will still usually take less CPU time than normal VSAM processing. IAM's data compression option, because of its overall efficiency, should be used on all medium to large files (over 5 CYLS). Some files because of the nature of their data may not benefit as much from compression as others. In those cases, to bypass data compression specify DATACOMPRESS=NO. SMP/E's CSI files are an example of files that usually show only a marginal benefit.

Default is NO unless the file size is equal or greater than the value specified in IAM's Global Option Table.

DATASPACE= Specifies the size, in megabytes, of the Data Space to be used for the temporary storage of the index to the IAM file that is being loaded. This parameter is only valid for MVS systems which support the use of Data Spaces. Valid values are from 0 to 2048. A value of 0 results in the use of a dynamically allocated temporary dataset.

Default is 128 megabytes, or other value as specified in the IAM Global Options Table.

DESCRIPTCODE= Specifies the system 'descriptor code' IAM uses for all IAM messages written to the system/console log.

DESCR=

Default is 0 (X '0000').

FIXED

Specifies that the logical records written to the IAM file are fixed in length. This parameter will also disable data compression for the specified file. The parameter is only honored on the DEFINE or on the first load of the file. It is intended for use when the file actually contains fixed length records, however IAM is defaulting to variable length formats. This can occur when the file is SMS JCL allocated, allocated with a DATA CLAS, or when the Global Options Table indicates the VSAM default is to be variable.

INDEXCOMPRESS= Specifies whether the file's index will be compressed.

INDEXCOMP=

YES - IAM is to compress the index for qualified files.

NO - IAM will not compress the index.

Default is YES.

INFO=

Specifies whether an IAM INFO report is to be generated for this file each time it is CLOSED. The ddname 'IAMINFO' must be present in the step JCL for IAM to produce this report.

YES - a run time INFO report will be produced for the IAM file(s) whenever it is CLOSED.

NO - IAM will not produce a run time INFO report for file(s) named on this control statement. You may want to specify INFO=NO if a file is OPENed and CLOSED many times in one run.

Default is YES.

11.02 IAM CREATE OVERRIDE STATEMENT - FORMAT

INTEGRATED= I=	<p>Specifies the percentage of each prime data block that IAM is to reserve, during the file load, for inserts. A value from 0 to 99 percent may be specified. On fixed length files, IAM will round the Integrated value down to an even multiple of the record length.</p> <p>IDCAMS DEFINE default - the FREESPACE CI% parameter value. The default for other interfaces is 10 percent.</p>
KEYLEN= K=	<p>Specifies the length of the record key field in the output IAM file. This parameter should only be used when the IDCAMS DEFINE is not used.</p> <p>IDCAMS DEFINE default - the first KEYS parameter value.</p>
LIMITKEYS= LIMITK=	<p>Specifies the number of keys that constitute a set during index compression. A minimum of 3 to a maximum of 64 keys in a set may be specified. Changing the number of keys in a set may improve index compression.</p> <p>Default is 32 keys in a set.</p>
LRECL= L=	<p>Specifies the maximum length of the records in the output IAM file. This parameter should only be used when the IDCAMS DEFINE is not used.</p> <p>IDCAMS DEFINE default - the second value of the DEFINE RECORDSIZE parameter.</p>
LOG=	<p>Specifies whether Override statements are to be logged.</p> <p>YES - an image of the Override statement is to be written to the programmer's log.</p> <p>NO - IAM will not log the control statements.</p> <p>Default is NO, unless there is a control statement error.</p>
MAXBUFNO=	<p>Specifies the default maximum number of buffers IAM is permitted to acquire during file processing. IAM's Real Time Tuning will dynamically adjust the number of buffers used for the file as demand warrants up to this maximum.</p> <p>You may specify a value from 1 to 32 buffers. See ACCESS OVERRIDE STATEMENT for details on default buffer values.</p>
MINBUFNO=	<p>Specifies a default minimum number of buffers IAM is to maintain during file processing. IAM's Real Time Tuning dynamically adjusts the number of buffers used for a file as demands change. IAM acquires this specified minimum number of buffers at OPEN and will never reduce the number of buffers maintained below this minimum.</p> <p>You may specify a value from 1 to 32 buffers.</p>
MINCOMPRESS= MINC=	<p>Specifies the minimum reduction of virtual storage IAM will accept to qualify a file for a compressed index. IAM accepts values between 8 and 40 percent.</p> <p>Default is 10 percent.</p>

11.02 IAM CREATE OVERRIDE STATEMENT - FORMAT

- OCOREO%=** Specifies the percentage of the total storage required for Independent Overflow that IAM will obtain when the file is OPENed for update. May specify a value from 0 to 100 percent.
- IAM acquires sufficient storage to contain all keys currently in Independent Overflow plus this percentage of the total capacity of Independent Overflow.
- This keyword is mainly intended for jobs, running on non-MVS/XA and non-MVS/ESA systems, where additional increments of storage may not be available after OPEN. It can be used to reserve the storage needed to manage a large number of inserts, when the file is first OPENed.
- Default is 10 percent.
- OCOREX%=** Specifies the percentage of the total storage, needed for all Independent Overflow keys, that IAM is to obtain when the current tables in storage are exhausted. A value from 1 to 100 percent may be specified.
- Default is 10 percent.
- OVERFLOW=**
O= Specifies the number of records, to be reserved within the IAM file, to accommodate inserts which will not fit into the prime blocks in which they belong. May specify a value from zero (0) to 3,000,000 records. The number of records specified must fit in 64, 000 blocks.
- IDCAMS DEFINE default is as documented in [Section 10.01](#) under IDCAMS FREESPACE CA% specification. The default under any other interface is 500 records.
- NOTE: Under VIF, initially loading an IAM file with only one record is treated as a special circumstance. IAM in this circumstance will attempt to create a file, up to the size of the specified primary allocation, that is capable of growth beyond the specified Overflow amount. The value of the key in the file's initializing record is used by IAM, when it creates the file, to determine the balance that will be established between Overflow and Prime Extension areas. To prevent IAM from preformatting a file in this fashion, specify a Prime Extension value of zero; PE=0. See the note below under PE=.*
- PE=** Specifies the number of blocks IAM is to reserve at the end of a file, as an extension of the prime area, for inserts with keys higher than the existing prime records (OPEN EXTEND). May specify a value from zero (0) to 32,767 blocks. IAM always increments this value by 1.
- NOTE: Under VIF, initially loading an IAM file with only one record is treated as a special circumstance. IAM in this circumstance will attempt to create a file, up to the size of the specified primary allocation, that is capable of growth beyond the specified amount of Prime Extension. The value of the key in the file's initializing record is used by IAM, when it creates the file, to determine the balance that will be established between Overflow and Prime Extension areas. To prevent IAM from preformatting a file in this fashion, specify a Prime Extension value of zero; PE=0.*
- Default is 3, which results in 4 PE blocks.
- PSEUDORBA** Uses non-standard RBAs to allow an IAM ESDS file to be larger than 4.3 Gigabytes.

11.02 IAM CREATE OVERRIDE STATEMENT - FORMAT

RELEASE=	<p>Specifies whether IAM is to release unused space after the file is loaded or REPROed. IAM releases the unused space in a file's allocation after the initial load. The space required for OVERFLOW is not released, as it appears as used space in the VTOC. It will not release space on subsequent reloads or reorganizations of the file, unless RELEASE=YES is specifically coded.</p> <p>YES - IAM will release unused space when a secondary allocation value is specified.</p> <p>NO - IAM will not release unused space in the file. NO should be specified if a file varies greatly in size each time it is loaded. NO should also be specified, if you wish to reserve space on a volume for future growth.</p> <p><i>NOTE: If you delete and reallocate the data set each time the file is reloaded, you should not inhibit IAM's release of the unused space.</i></p> <p>Default is YES, when there is a secondary allocation quantity specified, unless changed to NO in IAM's GLOBAL OPTIONS TABLE.</p>
RKP=	<p>Specifies the relative position of the key in the record. This parameter should only be specified when the IDCAMS DEFINE is not used.</p> <p>IDCAMS DEFINE default - the second value in the KEYS parameter.</p>
ROUTE CODE= ROUTE=	<p>Specifies the system 'route code' used for all IAM messages written to the system/console log.</p> <p>Default is 11 (X '0020').</p>
SMF=	<p>Specifies whether IAM is to write an IAM SMF user record when the file is CLOSEd. The IAM SMF user record will contain all of the same information displayed in the IAM INFO Run Time Statistics Report. Member IAMUSMF in the IAM Installation Control Library is a DSECT of the IAM SMF user record's format.</p> <p>YES - If a RECTYPE value was specified in IAM's GLOBAL OPTION TABLE, a user SMF record will be written when the IAM file is CLOSED. This will override an SMF=NO parameter specified in IAM's GLOBAL OPTION TABLE.</p> <p>NO - IAM will not write a user SMF record. This will override an SMF=YES parameter specified in IAM's GLOBAL OPTION TABLE.</p> <p>Default is as specified in IAM's GLOBAL OPTION TABLE.</p>
TRACEDDNAME= TRACEDD=	<p>Specifies a ddname that a key trace format output file/report will be written to.</p> <p>This keyword activates the IAM Trace Facility. IAM will write a trace record for each request to the IAM file. For further details on the IAM TRACE/DEBUG Facilities refer to the System Analysis Utilities section in this manual (see Section 43)</p> <p>'IAMTRPRC' and 'IAMTRPRx' are special DDNAMEs which specify the trace is to go direct to print.</p> <p>Default is the processing is not traced.</p>

11.02 IAM CREATE OVERRIDE STATEMENT - FORMAT

UNIT=	<p>Specifies a generic unit name for non-specific volume allocation. Disk volume(s) associated with this generic name must be mounted as STORAGE and the file's IDCAMS DEFINE VOLUMES parameter must specify 'ANYVOL' for non-specific allocation to succeed. For more information on IAM non-specific allocations, refer to IDCAMS DEFINE Support, Section 10.04.</p> <p>Default: When ANYVOL is specified in the IDCAMS DEFINE VOLUMES parameter, the default UNIT= will be 'SYSDA'. If ANYVOL is not specified in the DEFINE VOLUMES parameter, IAM will allocate to the volume serial number(s) coded in the DEFINE.</p>
VARIABLE VAR	<p>Specifies that the logical records written to the IAM file are variable in length. This parameter should only be specified when the IDCAMS DEFINE indicates that records are fixed length, but they are actually variable. This parameter is only effective on the first file load or on a DEFINE.</p> <p>IDCAMS DEFINE default - The file will be defined as variable when any of the following conditions are specified:</p> <ul style="list-style-type: none">- File is enabled for Data Compression- Average and Maximum record lengths are different- File is allocated by SMS JCL- File is allocated with an SMS DATACLAS- Global Options Table indicates VSAM default is Variable <p>The default under any other interface is the RECFM coded in the application program.</p>
VM	<p>Specifies that this file is being created under VM/CMS. This is normally not required because IAM determines this dynamically.</p>
WKDDNAME=	<p>Specifies the ddname of a temporary work file required by IAM file load under non-MVS operating systems. Under a non-MVS operating system, i.e. VM/CMS, when multiple IAM files are created simultaneously by a single processor each IAM file must be provided a unique work file ddname.</p> <p>This file is always dynamically allocated by IAM in MVS, MVS/XA or MVS/ESA system.</p> <p>EX: //IAMWKDD DD UNIT=SYSDA,SPACE=(CYL,(1,1))</p> <p>Default name is 'IAMWKDD'.</p>

11.03 IAM ACCESS OVERRIDE STATEMENT - FORMAT

11.03 IAM ACCESS OVERRIDE STATEMENT - FORMAT

ACCESS COMMAND The ACCESS statement is used to override IAM execution time defaults, for a specific step. The ACCESS override statement applies to file access processing only, it does not apply to file DEFINE, load, or create.

ACCESS

DDNAME=ddname	,MODE=x
,BROWSE	,OCOREO%=nn
,DESCRIPTCODE=nn	,OCOREX%=nn
,DYNCORE=nnnnn	,OPTCD=nn
,INDEXCOMPRESS=NO <u>YES</u> NO	,ROUTECD=nn
,INFO=<u>YES</u> NO	,SMF=YES NO
,LOG=<u>YES</u> NO	,TRACEDDNAME=ddname
,LSR=YES <u>NO</u>	,TRACEREQUEST=xxxx
,MAXBUFNO=nn	,UPDATENQ=xxx
,MINBUFNO=nn	

DDNAME= Specifies the ddname of the IAM data set that this control statement will be applied to. A maximum of 40 DDNAMEs may be specified. If multiple DDNAMEs are specified, they must be enclosed in parenthesis.

DD=

A ddname of &ALLDD will result in this control statement being applied as a global override to all IAM files being accessed in this step.

DDNAME is a required parameter.

BROWSE

For compatibility with prior releases and intended only for use with IAM's ISAM Interface (IFI). Specifies that IAM will pass back the next record, if the requested key is not found on the file. A No Record Found condition is still passed to the user.

11.03 IAM ACCESS OVERRIDE STATEMENT - FORMAT

- DESCRIPTCODE=** Specifies the system 'descriptor code' IAM uses for all IAM messages written to the system/console log.
DESCR= Default is 0 (X '0000').
- DYNCORE=** Specifies an amount of memory for IAM's Dynamic Table. The value is specified in 1024 byte (1K) increments. IAM will GETMAIN the value specified, at OPEN for use as an incore table. On all random requests, IAM checks the table for the record. If the record is found in the table, it is returned from memory without I/O. If not currently in the table, the record will be placed in the table after it is retrieved from the file. Random requests eligible for DYNCORE are identified in IAMINFO REPORTS as R.(READ) Commands.
Updates are always made in the table and on the file. Variable length records will be maintained as maximum length entries.
IAM accepts values of from 0 to 16000 (i.e. Dynamic Table up to 16000K bytes, 16 megabytes of storage for tabling). In an MVS/XA or MVS/ESA environment, storage for IAM's Dynamic Table is requested above the 16 megabyte line.
Example: DYNCORE=200 specifies IAM is to acquire 200K of storage for use as a Dynamic Table. [See Section 02.50](#) (Dynamic Table) for details.
Default is zero (0), the Dynamic Table is disabled, no records will be tabled in memory.
- INDEXCOMPRESS=** Specifies whether to use the compressed index.
YES - IAM will use the file's compressed index if one exists.
NO - IAM will not use the compressed index. The uncompressed index will be used.
Default is YES.
- INFO=** Specifies whether an IAM INFO report is to be generated for this file each time it is CLOSEd. The ddname 'IAMINFO' must be present in the step JCL for IAM to produce this report.
YES - a run time INFO report will be produced for the IAM file(s) whenever it is CLOSEd.
NO - IAM will not produce the run time INFO report for the file(s) named on this control statement. You may want to specify INFO=NO if the file is OPENed and CLOSEd many times in one run.
Default is YES.
- LOG=** Specifies whether Override statements are to be logged.
YES - an image of the Override statement is to be written to the programmer's log.
NO - IAM will not log the control statements.
Default is NO, unless there is a control statement error.

11.03 IAM ACCESS OVERRIDE STATEMENT - FORMAT

- LSR=** Specifies whether IAM will let an OPEN that specifies LSR or GSR to complete successfully.
- YES** - an Open that specifies LSR or GSR for an IAM file will be permitted. This will not cause IAM buffers to be put in the LSR pool, nor will IAM support any of the LSR requests. This may be useful however for applications which can not have LSR turned off, but which also do not use any of the LSR services. Any attempt to use an LSR service will result in an SOCx abend. DO NOT SPECIFY YES FOR CICS. LSR must be turned off for IAM files in the FCT.
- NO** - IAM will fail LSR/GSR opens with an IAMW16 error message.
- Default is NO.
- MAXBUFNO=** Specifies the maximum number of buffers IAM is permitted to acquire during file processing. IAM's Real Time Tuning will dynamically adjust the number of buffers used for the file as demand warrants up to this maximum.
- You may specify a value from 1 to 255 buffers.
- MVS/XA and MVS/ESA default: The MAXBUFNO value specified during DEFINE or if not specified the number of buffers which will fit in 256K of storage. For example, for a file with a blocksize of 11,476, IAM will acquire up to a maximum of 22 buffers, if the demand warrants. [See Section 02.50](#) IAM's Real Time Tuning for details.
- NON-MVS/XA and NON-MVS/ESA default: The same as for XA systems with the exception that for online systems (programs using asynchronous processing) the default maximum number of buffers is made equal to the minimum number of buffers (MINBUFNO).
- MINBUFNO=** Specifies a minimum number of buffers IAM is to maintain during file processing. IAM's Real Time Tuning dynamically adjusts the number of buffers used for a file as demands change. IAM acquires this specified minimum number of buffers at OPEN and will never reduce the number of buffers maintained below this minimum.
- You may specify a value from 1 to 32 buffers. If MINBUFNO specifies a value greater than MAXBUFNO, it will be reduced to the MAXBUFNO value.
- MVS/XA and MVS/ESA default: IAM acquires one buffer for each block on a track when the file is OPENed. For example, if the blocksize is 11,000 on a 3380, IAM will get 4 buffers at OPEN time. If IAM determines that buffers are not needed, it will gradually decrease the number of buffers to a minimum of one (1), or the MINBUFNO value specified.
- NON-MVS/XA and NON-MVS/ESA default: IAM acquires one buffer when the file is OPENed.

11.03 IAM ACCESS OVERRIDE STATEMENT - FORMAT

MODE= For compatibility with prior releases, specifies how IAM is to refresh buffers and when updated buffers are rewritten to disk.

For BATCH (NONASYNCHRONOUS) processing the following values apply:

R- specifies read only, no update processing.

U- specifies read and update processing. Blocks are rewritten only when forced out of a buffer by a read for another block.

I- specifies read and update processing. IAM forces the buffer to be rewritten on each update.

J- specifies read and update processing. IAM refreshes the buffer on each read and forces the buffer to be rewritten on each update.

For ONLINE (ASYNCHRONOUS) processing the following values apply:

M- specifies read only, no update processing.

N- specifies read and update processing. Blocks are rewritten only when forced out of a buffer by a read for another block.

O- specifies read and update processing. IAM forces the buffer to be rewritten on each update.

P- specifies read and update processing. IAM refreshes the buffer on each read and forces the buffer to be rewritten on each update.

Defaults are as follows:

Batch(NONASYNCHRONOUS) OPEN for INPUT - MODE=R

Batch(NONASYNCHRONOUS) OPEN for UPDATE - MODE=U

Online(ASYNCHRONOUS) OPEN for INPUT - MODE=M

Online(ASYNCHRONOUS) OPEN for UPDATE - MODE=O

Share Option 4 (NONASYNCHRONOUS) - MODE=J

Share Option 4 (ASYNCHRONOUS) - MODE=P

OCOREO%= Specifies for this execution a percentage of the total storage required for Independent Overflow that IAM will obtain when the file is OPENed for update. May specify a value from 0 to 100 percent.

IAM acquires sufficient storage to contain all keys currently in Independent Overflow plus this percentage of the total capacity of Independent Overflow.

This keyword is mainly intended for jobs, running on non-MVS/XA and non-MVS/ESA systems, where additional increments of storage may not be available after OPEN. It can be used to reserve the storage needed to manage a large number of inserts, when the file is first OPENed.

Default is 10 percent, or the value specified when the file was DEFINEd/loaded.

OCOREX%= Specifies for this execution the percentage, of the total storage needed for all Independent Overflow keys, that IAM is to obtain when the current tables in storage are exhausted. A value from 1 to 100 percent may be specified.

Default is 10 percent, or the value specified when the file was DEFINEd/loaded.

11.03 IAM ACCESS OVERRIDE STATEMENT - FORMAT

OPTCD= For compatibility with prior releases, specifies ISAM file compatibility and optional special processing codes. These codes direct how IAM will handle ISAM compatibility for non/ISAM interface processors including VIF and 'programmed real delete' requests to the IAM ISAM File Interface (IFI).

NOTE: IFI requires the ISAM OPTCD processing options be specified in the program DCB, as a JCL DCB parameter or in the file's model DSCB.

L - Specifies 'logical ISAM delete' support. Non/IFI access will process as an ISAM file marked with 'OPTCD=L'. In an IAM file so marked, an X 'FF' in the first byte of a record will be treated as a flag requesting the following special processing:

Sequential Get - IAM bypasses the record.

Random reads and positioning commands the record will be processed normally.

Add - the record is replaced instead of the add failing as a duplicate.

Update - the record is rewritten to the file.

LI or IL - Specifies that an update, when the first byte of a record contains the value X 'FF', will delete the record from the file.

M - Specifies 'mixed logical and real delete' support. This override will result in any access to the file for a record with a X 'FF' in the first byte being treated as if the file were marked with both an 'L' and an 'I'. If a record is processed with X 'FF' in the first byte, the following will happen:

Sequential Get - IAM bypasses the record.

Random reads and positioning commands the record will be processed normally.

Add - the record is replaced instead of the add failing as a duplicate.

Update - the record is deleted from the file.

ND - Suppress delete byte support, IAM will ignore a delete byte (X 'FF' in the first byte).

Default: for VIF processing is no ISAM compatibility support and for IFI processing is as specified in the program DCB or JCL.

ROUTE CODE= Specifies the system 'route code' IAM will use for all IAM messages written to the system/console log.

ROUTE= Default is 11.

SMF= Specifies whether IAM is to write user IAM SMF records when the file is CLOSED. The IAM SMF user record will contain all of the same information displayed in the IAM INFO Run Time Statistics Report. Member IAMUSMF in the IAM Installation Control Library is a DSECT of the IAM SMF user record format.

YES - If a RECTYPE value was specified in IAM's GLOBAL OPTION TABLE a user SMF record will be written when the IAM file is CLOSED. This will override an SMF=NO parameter specified in IAM's GLOBAL OPTION TABLE.

NO - IAM will not write a user SMF record. This will override an SMF=YES parameter specified in IAM's GLOBAL OPTION TABLE.

Default is as specified in IAM's GLOBAL OPTION TABLE.

11.03 IAM ACCESS OVERRIDE STATEMENT - FORMAT

TRACEDDNAME= Specifies the ddname that a key trace format output file/report will be written to.
TRACEDD= TRACEDD activates the trace facility. IAM writes a trace record for each request to the IAM file.

For further details on IAM TRACE/DEBUG Facilities refer to the System Analysis Utilities portion of this manual ([see Section 43](#)).

'IAMTRPRC' and 'IAMTRPRx' are special DDNAMEs which specify the trace is direct to print.

ESDS only - The IAMTRPRC DDNAME must be specified for tracing to occur.
Default is the processing is not traced.

TRACEREQUEST= Specifies the output format for a trace report.

There are several report formats available. For further details on IAM TRACE/DEBUG Facilities refer to System Analysis Utilities ([Section 43](#)).

TRACE - IAM will write a TRACE record for each request to the IAM file. The key, command issued and completion code are recorded.

MONITOR - IAM will write a MONITOR format record for each request to the IAM file.

Default is TRACE.

ESDS only - The following trace types can be requested when processing IAM ESDS files:

TRIOS - Trace I/O Start requests

TRIOE - Trace I/O End requests

TRBFR - Trace IAM Buffer Manager requests

TREXCP - Trace EXCP requests

TRXTND - Trace EXTEND dataset requests

UPDATENQ= Overrides the enqueue protection specified for the file by the IDCAMS DEFINE SHAREOPTIONS (or in the IAM defaults) during this execution.

The UPDATENO, override can be used in lieu of using an IDCAMS ALTER to temporarily change the share options. As a rule, to read or update a file that is currently open to another job with enqueue protection, you will need to specify UPDATENQ=NONE. Using the override will only effectively reduce enqueue protection for a file, it will not increase protection.

EXCL - Specifies that IAM will only allow one user to have this file OPENed for update but any number will be able to concurrently read the file. This option will NOT cause a Share Option 1 file to have a Share Option 2.

SHR - Specifies that IAM will allow other users to have this file OPENed, for read or update. All concurrent users must specify SHR for this option to be effective.

NONE - Specifies that during this execution IAM will ignore enqueue protection on this file. Other jobs will be able to access and update the file at the same time. Same as VSAM's Share Option 3.

IDCAMS DEFINE default - the Share Options specified in the DEFINE.

Other IAM Interfaces default to EXCL.

11.04 IAM OVERRIDE STATEMENT EXAMPLES

11.04 IAM OVERRIDE STATEMENT EXAMPLES

DEFINE USING OVERRIDES 1. The following example shows how an IAM Override statement can be used during an IDCAMS DEFINE of an IAM file. Many records are inserted into this file, therefore Independent Overflow is being increased to 100,000 records. The file changes in size daily, so RELEASE=NO is coded to prevent the release of unused space.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//IAMOVRID   DD    * OVERRIDE STATEMENT FOLLOWS *
              CREATE DD=&ALLDD,OVERFLOW=100000,RELEASE=NO
/*
//SYSIN      DD    *
              DEFINE CLUSTER -
                  (NAME(MASTER.INVOICE.$IAM) -
                  CYLINDERS(100 15) -
                  FREESPACE(10 5) -
                  KEYS(9 0) -
                  RECORDSIZE(500 500) -
                  VOLUMES(MVS001) )
/*
```

DEFINE WITH DATA COMPRESS 2. The following example shows how an IAM Override statement is used during an IDCAMS DEFINE of an IAM file, to request the file's data records be compressed. The IAM file's DEFINE will use the attributes of an existing VSAM cluster.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
              DEFINE CLUSTER -
                  (NAME(CICS.MASTER.$IAM) -
                  MODEL(CICS.MASTER.CLUSTER) -
                  VOLUMES(MVS001) )
/*
//IAMOVRID   DD    * OVERRIDE STATEMENT FOLLOWS *
              CREATE DD=&ALLDD,DATACOMPRESS=YES
/*
```

11.04 IAM OVERRIDE STATEMENT EXAMPLES

DEFINE USING NON-SPECIFIC ALLOCATION

3. The following example shows an IDCAMS DEFINE of an IAM file using non-specific allocation. The file will be allocated to any volume with the generic unit of 'TSO'. The VOLUMES parameter in the IDCAMS DEFINE must specify 'ANYVOL'.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
              DEFINE CLUSTER -
                (NAME(TSOUSER.TEST.FILE) -
                OWNER($IAM) ) -
                DATA (CYLINDERS(100 20) -
                FREESPACE(10 5) -
                KEYS(24 8) -
                RECORDSIZE(250 480) -
                VOLUMES(ANYVOL) )
                INDEX (TRACKS(1 1) -
                VOLUMES(ANYVOL) )
/*
//IAMOVRID   DD    * OVERRIDE STATEMENT FOLLOWS *
              CREATE DD=&ALLDD,UNIT=TSO
/*
```

OVERRIDE PROCESSING AT EXECUTION TIME

4. The following example shows an IAM Override for files being processed under CICS. Two of the files will have their maximum buffers increased from the default to 16 buffers. Another file will use the Dynamic Table (DYNCORE) option to table records in memory. A size of 100K of memory is specified for the table.

```
//CICS       EXEC  PGM=DFHSP
//.....
//FILE1     DD    DSN=IAM.FILE1,DISP=SHR
//FILE2     DD    DSN=IAM.FILE2,DISP=SHR
//FILE3     DD    DSN=IAM.FILE3,DISP=SHR
//.....
//IAMINFO   DD    SYSOUT=*      REPORTS RUN TIME STATS
//IAMOVRID   DD    * OVERRIDE STATEMENT FOLLOWS *
              ACCESS DD=(FILE1,FILE2),MAXBUFNO=16
              ACCESS DD=FILE3,DYNCORE=100
/*
```

TRACING REQUESTS MADE AGAINST AN IAM FILE

5. The following example shows how to trace the requests to an IAM file. IAM will report for each request to the file the key used, the command being requested and the completion code. The DDNAME 'IAMTRPRX' will cause the trace report to go direct to print. The key will be displayed in hexadecimal. To have the key printed in character format, change the DDNAME 'IAMTRPRX' in the JCL to 'IAMTRPC'.

```
//COBOL     EXEC  PGM=PAYROLL
//.....
//MASTER    DD    DSN=MASTER.$IAM.FILE,DISP=OLD
//IAMINFO   DD    SYSOUT=*
//IAMTRPRX  DD    SYSOUT=*
//IAMOVRID   DD    * OVERRIDE STATEMENT FOLLOWS *
              ACCESS DDNAME=MASTER,TRACEDD=IAMTRPRX
/*
```

SECTION 12: IAM RUN TIME STATISTICS REPORT AND SMF RECORDING

12.01 IAM RUN TIME STATISTICS (INFO) OVERVIEW

12.01 IAM RUN TIME STATISTICS (INFO) OVERVIEW

OVERVIEW IAM's Run Time Statistics (INFO) Report provides a means of accurately tracking a program's file processing workload and the overflow status of the IAM files it uses.

TECHNICAL SUMMARY IAM maintains run time statistics during execution for all IAM files in use. These statistics will be put to a SYSOUT data set in report form whenever an IAMINFO DD statement is present in the JCL and/or independently they can be written out to SMF in an IAM user record.

FUNCTIONAL SUMMARY IAM can report on the processing performed against IAM files each time they are OPENed and CLOSEd. Available for both the batch and CICS user, IAM's Run Time Statistics (INFO) Report provides a complete summary of an applications IAM processing. IAM's run time statistics includes counts of requests by command type, number of disk reads and writes, overflow records used, memory required, tracks used, I/O buffer usage, buffering recommendations, and other file characteristics.

RECORDING IAM STATS IN SMF IAM can directly record its run time statistics in a system's SMF files each time an IAM file is OPENed and CLOSEd. When IAM's optional SMF recording is active an SMF user record will be written containing the same information found in IAM's Run Time Statistics (INFO) Report.

IAM SMF USER RECORD Member IAMUSMF in the IAM Installation Control Library is a DSECT of the IAM SMF user record's layout, it is approximately 500 bytes in length. This record can be selected from an SMF file and displayed using the IAM SMF utility IAMSMF PRINT command ([See Section 41.07](#)). The information in the IAM SMF user record can be formatted and printed by any utility that can read SMF records. The fields in IAM's SMF user record contain the same values as displayed in the IAM (INFO) Run Time Statistics Report which is documented in [Section 12.04](#) of this manual.

IAM INFO REPORTS FROM SMF DATA IAM Run Time Statistics (INFO) Reports can be produced from information recorded in IAM's optional SMF User Records. IAM also provides a number of other reports on the use of IAM files that can be produced from data recorded in a system's SMF data sets.

The IAM SMF Job Resource Analysis program (IAMSMF) has a specific command that will generate IAM INFO reports from IAM SMF User Records. [See Section 41.10](#) for information on the IAMSMF analysis program's IAMINFO command.

The IAM VSAM Usage Analysis program (IAMSMFVS) provides Space, Activity and Summary reports on IAM file usage, again based on IAM SMF User Records. [See Section 40](#) for detail information on the summary and detail reports that are available.

12.02 REQUIREMENTS TO ACTIVATE IAM'S RUN TIME STATISTICS

12.02 REQUIREMENTS TO ACTIVATE IAM'S RUN TIME STATISTICS

ACTIVATING IAM INFO THE REPORT To activate the INFO report, include an IAMINFO DD statement in any run which processes IAM files.

THE REPORT // IAMINFO DD SYSOUT=*

There are no other changes needed.

To exclude one or more files in a Jobstep from IAM Run Time INFO Report processing, use the IAM Override Statement keyword INFO=NO. See [Section 11](#), IAM Override Statements for further detail.

IAMINFO DD STATEMENT Specifies the data set to be used for the IAM run time statistics report. Usually a SYSOUT data set. If blocking is required a BLKSIZE which is a multiple of 121 must be specified.

IAM SMF STATS The IAM SMF recording facility must be enabled in a system before user SMF records can be written. [Section 91.03](#), Global Option Change Facility explains how to enable SMF recording using the ZAP command's parameters RECTYP= and SMF=. When IAM SMF recording is enabled in a system it can be activated or deactivated at execution time using the SMF= parameter on an IAM Override Statement. See [Section 11](#), IAM Override Statements for details on the SMF= keyword.

12.03 EXAMPLE OF RUN TIME STATISTICS REPORTING

12.03 EXAMPLE OF RUN TIME STATISTICS REPORTING

RUN TIME STATISTICS REPORT Activate run time statistics reporting on the IAM data set 'IAMV.DCM201.CLUSTER' when the program 'USERPGM' is executing. Print the INFO report on SYSOUT=A.

```
//DCM201B EXEC PGM=USERPGM
//VSAMCRT1 DD DSN=IAMV.DCM201.CLUSTER,
DISP=OLD...
//IAMINFO DD SYSOUT=A
```

SAMPLE CREATE STATISTICS 1. IAM RUN TIME STATISTICS REPORT FROM FILE LOAD

```
IAM400 INNOVATION IAM INFORMATION PRINT ROUTINE -- IAMNINFO VER 6.X -- INNOVATION DATA PROCESSING
DATE - 90.208
IAM360 STEP - IAMVST1 DDNAME - SYS00003 DATA SET MONITORED - IAMV.MARK$IAM.IAMFILE1
IAM361 INFO REQUESTED BY IAMCRT PROGRAM PERFORMING FILE CREATION OPENED-90.208.09:07:42
CLOSED-90.208.09:08:12
IAM362 IAM DATA CHARACTERISTICS -
LOGICAL RECORD LENGTH-----= 1020 - RECORD FORMAT-----= VARIABLE
KEY LENGTH-----= 4 - RELATIVE KEY POSITION-----= 8
AVERAGE RECORD LENGTH-----= 690 - LARGEST RECORD PROCESSED--= 1016
IAM363 IAM FILE CHARACTERISTICS -
BLOCKING FACTOR-----= 2 - BLOCK SIZE-----= 23476
TRACKS IN USE-----= 325 - DELETE BYTE SUPPORT-----= NONE
DATASET TYPE-----= KSDS - SHARE OPTIONS-----= 2
NUMBER OF IAM DATA BLOCKS--= 648
IAM364 IAM OVERFLOW CHARACTERISTICS -
INDEPENDENT RECORDS REQ---= 1320 - REQUESTED INTEGRATED PER--= 0
INDEPENDENT RECORDS ALLOC--= 1334 - ACTUAL INTEGRATED PER-----= 0
INDEPENDENT RECORDS USED--= 0 - PE BLOCKS REQUESTED-----= 3
INDEPENDENT RECORDS EMPTY--= 1334 - PE BLOCKS - ALLOCATED,USED= 4,0
INDEPENDENT BLOCKS ALLOC--= 58 - INDEPENDENT PCT-OPEN,FULL--= 10,10
IAM365 IAM EXECUTION STATISTICS -
TOTAL STORAGE REQUIRED-----= 769216 - COMPRESSED KEY STRUCTURE--= NO
STORAGE ABOVE THE LINE-----= 737280 - COMPRESSED DATA STRUCTURE--= YES
REQUESTS PROCESSED-----= 24641 - REQUESTS FAILED-----= 0
DISK BLOCKS READ-----= 2 - DISK BLOCKS WRITTEN-----= 61
DYNAMIC BUFFER RETRIEVALS--= 0 - MAXIMUMBUFFERS USED-----= 30
DATA SPACE USED(M BYTES)--= 1 - DATA SPACE SIZE(M BYTES)--= 128
IAM366 IAM COMMAND EXECUTION SUMMARY -
R.(READ)-----= 0 - W.(WRITE)-----= 24640
F.(GETNEXT)-----= 0 - A.(ADD)-----= 0
D.(DELETE)-----= 0 - N.(READOR GETNEXT)-----= 0
G.(GET)-----= 0 - S.(START KEY EQUAL)-----= 0
K.(START KC OR POINT)-----= 0 - P.(GET PREVIOUS)-----= 0
U.(READ FOR UPDATE)-----= 0 - RESERVED-----= 0
I.(INFORMATION)-----= 0 - E.(STATISTICS)-----= 0
C.(CLOSE)-----= 1 - O.(OPEN)-----= 0
X.(FLUSH)-----= 0 - VARIABLE LENGTH CHANGES---= 0
DELETED RECORDS BYPASSED--= 0 -
```

12.03 EXAMPLE OF RUN TIME STATISTICS REPORTING

SAMPLE 2. IAM RUN TIME STATISTICS REPORT FROM FILE ACCESS ACCESS STATISTICS

```

IAM400 INNOVATION IAM INFORMATION PRINT ROUTINE -- IAMNINFO VER 6.X -- INNOVATION DATA PROCESSING
      DATE - 90.208
IAM360 STEP - IAMVTST1 DDNAME - SYS00003 DATA SET MONITORED - IAMV.MARK$IAM.IAMFILE1
IAM361 INFO REQUESTED BY IAMACC PROGRAM PERFORMING RANDOM/GET-NEXT OPENED-90.208.09:07:11
      CLOSED-90.208.09:07:42
IAM362 IAM DATA CHARACTERISTICS -
      LOGICAL RECORD LENGTH-----= 1020 - RECORD FORMAT-----= FIXED
      KEY LENGTH-----= 4 - RELATIVE KEY POSITION-----= 8
      TOTAL RECORDS-----= 24640 - TOTAL RECORDS DELETED-----= 0
      TOTAL RECORDS UPDATED-----= 0 - TOTAL RECORDS INSERTED-----= 0
IAM363 IAM FILE CHARACTERISTICS -
      BLOCKING FACTOR-----= 2 - BLOCK SIZE-----= 23460
      TRACKS IN USE-----= 569 - DELETE BYTE SUPPORT-----= NONE
      DATASET TYPE-----= KSDS - SHARE OPTIONS-----= 2
      NUMBER OF IAM DATA BLOCKS--= 648
IAM364 IAM OVERFLOW CHARACTERISTICS -
      INDEPENDENT RECORDS REQ--= 1320 - REQUESTED INTEGRATED PER--= 0
      INDEPENDENT RECORDS ALLOC--= 1334 - ACTUAL INTEGRATED PER-----= 0
      INDEPENDENT RECORDS USED--= 0 - PE BLOCKS REQUESTED-----= 3
      INDEPENDENT RECORDS EMPTY--= 1334 - PE BLOCKS - ALLOCATED,USED= 4,0
      INDEPENDENT BLOCKS ALLOC--= 58 - INDEPENDENT PCT-OPEN,FULL--= 10,10
IAM365 IAM EXECUTION STATISTICS -
      TOTAL STORAGE REQUIRED-----= 799120 - COMPRESSED KEY STRUCTURE--= NO
      STORAGE ABOVE THE LINE-----= 770048 - COMPRESSED DATA STRUCTURE--= NO
      ACCESS PROCESSING MODE-----= R - KEY STRUCTURE STORAGE REQ--= 4304
      REQUESTS PROCESSED-----= 24642 - REQUESTS FAILED-----= 1
      DISK BLOCKS READ-----= 55 - DISK BLOCKS WRITTEN-----= 0
      DYNAMIC BUFFER RETRIEVALS--= 1023 - MAXIMUMBUFFERS USED-----= 32
      DYNAMIC TABLE RETRIEVALS--= 0 - DYNAMICTABLE RECORDS-----= 0
IAM366 IAM COMMAND EXECUTION SUMMARY -
      R.(READ)-----= 0 - W.(WRITE)-----= 0
      F.(GETNEXT)-----= 0 - A.(ADD)-----= 0
      D.(DELETE)-----= 0 - N.(READOR GETNEXT)-----= 0
      G.(GET)-----= 24641 - S.(START KEY EQUAL)-----= 0
      K.(START KC OR POINT)-----= 0 - P.(GET PREVIOUS)-----= 0
      U.(READ FOR UPDATE)-----= 0 - RESERVED-----= 0
      I.(INFORMATION)-----= 0 - E.(STATISTICS)-----= 0
      C.(CLOSE)-----= 1 - O.(OPEN)-----= 0
      X.(FLUSH)-----= 0 - VARIABLE LENGTH CHANGES--= 0
      DELETED RECORDS BYPASSED--= 0 -

```

12.04 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

12.04 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

The IAM Run Time Statistics (INFO) Report is produced when the IAM file is CLOSED. There will be one report produced for each OPEN/CLOSE combination a program issues for an IAM file.

A description of the report and its major fields follows.

IAM400 The IAMINFO Report title containing a module name and version.

IAM360 Provides the name of the job step in execution, the DD statement, DDNAME and data set name of the IAM file being processed.

IAM361 Names the IAM processing module that requested the INFO Report, the type of processing being performed, and the date and time the IAM file was opened and closed. The type of processing maybe CREATE, SEQUENTIAL, RANDOM or RANDOM/GETNEXT. CREATE is shown when a file is loaded. RANDOM/GETNEXT will be shown whenever a file is processed with random and/or sequential commands. Sequential access is reported only when after open the PGM starts with the first record then proceeds record by record through the file (i.e., on start/point command specified).

IAM362 This heading line always precedes the IAM file Data Characteristics portion of the IAM INFO report. The following fields describing the logical record are provided.

LOGICAL RECORD LENGTH - The maximum length of the record. For variable length records IAM adds an internal 4 bytes (see Note: below).

RECORD FORMAT - Fixed or Variable format indication.

KEY LENGTH - The length of the key.

RELATIVE KEY POSITION - The position of the key in the record relative to zero (0). For variable length records, IAM internally adds 4 bytes (see Note: below).

NOTE: Internally IAM maintains variable length records as true OS format variable records with a record descriptor word (RDW) in the first four bytes of the record. Consequentially, the record length shown in the report will be four bytes longer than an equivalent VSAM variable length record. The RKP, in this case, will also be offset by four bytes.

AVERAGE RECORD LENGTH - During a load operation IAM calculates an average length for the records that were loaded into the file.

NOTE: This field only appears when a file is LOADED.

LARGEST RECORD PROCESSED - During a load operation IAM reports the length of the largest record that was written to the file.

NOTE: This field only appears when a file is LOADED.

TOTAL RECORDS - The total number of records in the IAM file.

NOTE: This field does not appear when the file is initially created.

TOTAL RECORDS DELETED - The total number of records deleted from the IAM file since creation.

NOTE: This field does not appear when the file is initially created.

TOTAL RECORDS UPDATED - The total number of records updated in the IAM file since creation.

NOTE: This field does not appear when the file is initially created.

TOTAL RECORDS INSERTED - The total number of records inserted into the IAM file since creation.

NOTE: This field does not appear when the file is initially created.

12.04 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

IAM363 This heading line always precedes the IAM File Characteristics portion of the IAM INFO report. The following fields describing the physical file are provided.

BLOCKING FACTOR - A value of 1-9 is the number of blocks per track (as blocking factor). A larger value is the blocksize if it was specified in bytes.

BLOCKSIZE - Physical blocksize of the IAM file. This is the actual blocksize developed and used by IAM.

NOTE: IAM will determine the optimal blocksize for a file based upon the file's record length, the blocking factor requested and the track capacity of the device. The IAM blocksize is transparent to application programs even when the blocksize is altered or the file is relocated to a different device type.

TRACKS IN USE - Number of tracks actually occupied by the file. Not the total allocation.

NOTE: Any disk space allocated beyond the value identified by the 'tracks used' field are an over allocation. They will not be used by IAM for any file processing. Over allocated space can be released automatically when an IAM file is loaded. See Section 11, the RELEASE keyword on the CREATE statement.

DELETE BYTE SUPPORT - Indicates if ISAM delete byte support is active for the file and the option in effect, see Section 11 OPTCD keyword on the ACCESS statement.

DATASET TYPE - Describes the type of data set IAM is simulating. When accessing IAM files through VIF, this field will contain KSDS. When accessing files through the ISAM interface, this field will contain ISAM. When accessing through native IAM calls, this field will contain NATIVE.

SHARE OPTIONS - Describes the defined IAM share options for this data set.

NUMBER OF IAM DATA BLOCKS - The number of blocks in the file, except for the blocks for the index. This figure may be necessary if a file recovery is required and the first block or track of the file is bad. This is the value to plug into the IAMRECVR MAXBLKS operand.

12.04 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

IAM364 This heading line always precedes the IAM File Overflow Status portion of the IAM INFO report. The following fields describing the overflow profile requested for the file when it was created as well as the status of the file's independent overflow (the overflow of last resort) at file close time are provided.

INDEPENDENT RECORDS REQ	Number of Independent Overflow records requested or calculated based on the IDCAMS DEFINE for the file. See Section 10.04 under IDCAMS FREESPACE CA%.
REQUESTED INTEGRATED PER	Percentage of Integrated Overflow (CI%) requested.
INDEPENDENT RECORDS ALLOC	Actual number of Independent Overflow records allocated by IAM.
ACTUAL INTEGRATED PER	Actual Percentage of Integrated Overflow (CI%) allocated by IAM. On fixed length files Integrated Overflow is rounded down to an even multiple of the record length.
INDEPENDENT RECORDS USED	Number of records currently in use in Independent Overflow.
PE BLOCKS REQUESTED	The number of Prime Extension blocks requested. If a specific amount of Prime Extension was not requested, this field will contain the default value.
INDEPENDENT RECORDS EMPTY	Number of records currently available in Independent Overflow.
<i>NOTE: As the 'INDEPENDENT RECORDS EMPTY' field approaches zero, the file is nearing a condition when reorganization is necessary. This should be done to prevent a file full condition for INSERT requests. A relatively high percentage of inserts going into independent overflow would indicate a need to evaluate the file's integrated overflow characteristics. Integrated overflow may have been set to low to allow inserts into the file where they logically belong.</i>	
PE BLOCKS – ALLOCATED,USED	The actual number of Prime Extension blocks allocated followed by the number of Prime Extension blocks in use.
INDEPENDENT BLOCKS ALLOC	The number of blocks used for Independent Overflow.
INDEPENDENT PCT – OPEN,FULL	The percentage of the Overflow keys to be reserved at OPEN for inserts followed by the percentage of the Overflow keys to be acquired when the incore tables are filled up.

12.04 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

IAM365 This heading line always precedes the IAM Execution Statistics portion of the IAM INFO report. The following fields describing system processing and IAM options used during file processing are provided.

- TOTAL STORAGE REQUIRED** - Total storage acquired by IAM during this execution.
- COMPRESSED KEY STRUCTURE** - Indicates whether a compressed index was used for this file. Under MVS/XA and MVS/ESA systems IAM's compressed index is maintained above the 16MB line.
- STORAGE ABOVE THE LINE** - Total amount of storage acquired above the 16MB line.
- COMPRESSED DATA STRUCTURE** - Indicates whether data compression was used for this file.
- ACCESS PROCESSING MODE** - Type of processing for this OPEN.
 - R - READ only
 - U - UPDATE
 - O - ONLINE
- KEY STRUCTURE STORAGE REQ.** - Total amount of storage required for the IAM Index Structure.
- REQUESTS PROCESSED** - Number of requests made against the file, since last OPEN, by the application (a breakdown by command type follows in the command execution summary).
- REQUESTS FAILED** - Number of requests made to IAM which did not complete normally (i.e. EOF on sequential read, no record found on random read)
- DISK BLOCKS READ** - Number of physical I/O's used to read blocks from the file.
- DISK BLOCKS WRITTEN** - Number of physical I/O's used to write blocks to the file.
- DYNAMIC BUFFER RETRIEVALS** - Number of EXCPs saved (not executed). IAM's Real Time Tuning buffer management was able to satisfy this number of I/O requests from buffers that it had maintained in storage, other than its current buffer. These buffer retrievals eliminated the need for I/O to disk.
- MAXIMUM BUFFERS USED** - The maximum number of data buffers acquired by IAM's Real Time Tuning during this execution. This may be accompanied by an IAM367 or IAM368 informational message indicating that if IAM was allowed to acquire more buffers the number of I/O's required to service the requests against the file could have been reduced. Under MVS/XA and MVS/ESA all buffers are acquired above the 16MB line.
- DYNAMIC TABLE RETRIEVALS** - Number of random read requests satisfied by IAM's Dynamic Tabling of data records in virtual. Under MVS/XA and MVS/ESA IAM's Dynamic Table (DYNCORE) is maintained in virtual storage above the 16MB line.
- DYNAMIC TABLE RECORDS** - The total number of data records IAM placed into the Dynamic Table.
- DATA SPACE USED(M BYTES)** - Indicates the amount of storage in the Data Space that was used, in megabytes, for the temporary storage of the index structure.
- DATA SPACE SIZE(M BYTES)** - Indicates the amount of storage requested for a Data Space, in megabytes, for the temporary storage of the index structure.

12.04 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

IAM366 This heading line always precedes the IAM Command Execution Summary portion of the IAM INFO report. The following fields provide a count of the logical application and internal IAM processing commands issued against the file since it was last OPENed.

R.(READ)	- Number of random Read requests processed
W.(WRITE)	- Number of Update requests processed
F.(GETNEXT)	- Number of GETNEXT commands processed
A.(ADD)	- Number of Insert requests processed
D.(DELETE)	- Number of Delete requests processed
N.(READ OR GETNEXT)	- Number of READ-or-GETNEXT commands
G.(GET)	- Number of sequential Get requests
S.(START KEY EQUAL)	- Number of IAM Start Key Equal commands
K.(START KEY CLASS)	- Number of Point requests processed
P.(GET PREVIOUS)	- Number of Get Previous requests processed
U.(READ FOR UPDATE)	- Number of Read for Update requests
RESERVED	- Reserved
I.(INFORMATION)	- Number of IAM Information commands
E.(STATISTICS)	- Number of IAM statistic commands issued
C.(CLOSE)	- Number of CLOSE requests processed
O.(OPEN)	- Number of IAM OPEN commands issued
X.(FLUSH)	- Number of requests with RPL errors
VARIABLE LENGTH CHANGES	- Number of records that had their record length changed.
DELETED RECORDS BYPASSED	- Count of records bypassed because the first byte contained X'FF'.

IAM367 . INSUFFICIENT STORAGE AVAILABLE TO ACQUIRE ADDITIONAL BUFFERS

This is an informational message provided by IAM's Real Time Tuning buffer management. It is provided when an attempt to acquire additional data buffers fails because of insufficient virtual storage. IAM acquires additional buffers when it had determined that this would improve performance. Increasing the region size will result in fewer disk I/O's to process this file.

IAM368 ... BUFNO VALUE GREATER THAN nn MAY IMPROVE PERFORMANCE

This informational message is provided whenever IAM determines that a greater number of buffers than provided would result in fewer disk I/O's to process this file. The MAXBUFNO= parameter on the IAM Override statement can be used to increase the number of buffers IAM will use. If the actual number of DISK BLOCKS READ and DISK BLOCKS WRITTEN is small there may be little benefit to be gained.

IAM371 INCREASING THE NUMBER OF PRIME EXTENSION BLOCKS MAY IMPROVE PERFORMANCE.

This informational message is provided whenever IAM determines that a record has been added to the INDEPENDENT OVERFLOW instead of being added to the PRIME EXTENSION because the PRIME EXTENSION is full.

12.04 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

IAMM

INNOVATION ACCESS METHOD

SYSTEMS ANALYSIS UTILITIES



THIS PAGE WAS INTENTIONALLY LEFT BLANK

SECTION 40: IAMSMFVS - VSAM ANALYSIS PROGRAM

40.01 IAMSMFVS - IAM'S VSAM USAGE SMF ANALYSIS PROGRAM

40.01 IAMSMFVS - IAM'S VSAM USAGE SMF ANALYSIS PROGRAM

OVERVIEW IAMSMFVS is a special purpose SMF data analysis program, intended to provide useful information in a concise format about indexed data set activity. The reports can be used to identify VSAM KSDS clusters that are candidates for conversion to IAM, to compare the results once the files are converted to IAM, and as a source for tracking IAM data set activity to monitor IAM data set usage, and determine when IAM files need to be reorganized.

With just a glance at IAMSMFVS's single page Summary report you will be able to see the extent of the burden imposed by VSAM KSDS clusters on your system's available DASD space and I/O activity. At a more detailed level IAMSMFVS's Activity Report points out your most active KSDS clusters and provides information on the jobs that uses them. The Cluster Size Report which ranks your KSDS clusters by size also provides information on file attributes and usage statistics. Optionally you can request detailed information on the names of jobs and programs accessing your VSAM clusters, with details on their specific file activity.

IAMSMFVS produces similar reports for IAM data sets, which are accessed through IAM's VIF interface. There is a separate Activity Report for IAM files, which includes information on the type and quantity of activity against the file, and an IAM Size Report, which includes information on how much space is required, the number of records in the file, and information on the usage of the IAM overflow areas. The IAM reports can be used to compare results after files have been converted to IAM, and can be used to monitor IAM data set activity.

TECHNICAL DESCRIPTION The IAMSMFVS program analyzes data set usage statistics gathered from two sources the Systems Management Facility (SMF records) and Catalog Management. SMF data may be taken from either the active SMF data sets or sequential SMF history data sets.

IAMSMFVS has a series of commands which control the available functions. Help is available for all commands. Help is also available for JCL or CLIST showing the required allocations for background or foreground processing. Use 'HELP HELP' to obtain information on the IAMSMFVS 'HELP' syntax.

Most commands and operands have abbreviations. The help information shows the shortest abbreviation allowed for each command or operand. It is always okay to include more characters.

REPORT TYPES IAMSMFVS summary and special purpose detail reports.

ACTIVITY REPORT The first report is the ACTIVITY report. It identifies the MDBO MDNM data sets having the highest EXCP counts. A separate Activity report is produced for each type of data set, that is VSAM clusters are in one report, and IAM files are grouped separately in their own Activity report. Each report sorts the clusters or data sets based upon the total number of EXCPs issued against the file (DATA and INDEX components) listing those clusters with the highest activity first. The report includes logical file activity statistics, such as records read, inserted, updated, and deleted. The name of each cluster is reported followed by its components listed in EXCP activity order. The report contains the space allocated for each component reported on along with the number of CI/CA splits the component may have gone through. A use count is included in the report showing the number of times the cluster was OPENed and CLOSEd during the report period.

High activity clusters are attractive candidates for conversion to IAM because compared to VSAM IAM reduces disk I/O 50% to 80%. By default, IAMSMFVS selects the 100 most active clusters but actually the number of VSAM clusters reported upon is a user option.

40.01 IAMSMFVS - IAM'S VSAM USAGE SMF ANALYSIS PROGRAM

**CLUSTER/
DATA SET
NAME
REPORT** The second report, is presented in data set name sequence. It lists file attributes such as key length, relative key position (RKP), and control interval size. A cluster name is always followed by its DATA component which is next followed by the INDEX.

If more than two components are reported for a cluster, the cluster contains alternate indexes. **If no cluster name is shown the components belong to a cluster that's no longer cataloged in any available catalog.**

NOTE: The RECFM and RKP will only be filled in if the cluster is currently cataloged.

**CLUSTER
SIZE
REPORT** The third report lists VSAM clusters or data sets in size order, starting with the largest. Like the Activity Report, separate Size Reports are provided for VSAM and IAM data sets. The data set name is reported on the first line with the total allocation for the cluster, or the amount of space used for IAM files, given in tracks. For VSAM clusters, the DATA and INDEX components follow, listed in size order by their individual allocations. This report also gives the number of extents the components occupy, their CISIZE and file attributes.

The cluster size report identifies the largest VSAM clusters in your installation. IAM's internal file structure saves 20% to 40% in disk space versus VSAM. With data compression, IAM can reduce a file's disk space an additional 20% to 50%.

**IAMSMFVS
SUMMARY
REPORT** The summary identifies the period of time the reports cover and summarizes the impact VSAM had on the system during that period in terms of I/O activity and disk space utilization. This is done by comparing the number of EXCPs issued against VSAM KSDS clusters against the total number of disk EXCPs executed during the report period. A summary of the total disk space allocated to VSAM KSDS clusters is also produced.

**IAMSMFVS
INPUT** IAMSMFVS takes as input SMF data records from either accumulated MDBO MDNM sequential SMF history files or from the live system SMF data sets. SMF record type 30 subtype 4 (or SMF record type 4) and SMF record type 64 must be collected and accumulated for IAMSMFVS to produce reports on VSAM clusters. For the IAM data set reports, the IAM Global Options must be set to enable recording of the IAM SMF records, and the SMF parameters and dump procedures must be set up to collect the IAM SMF records. To obtain a meaningful analysis, IAMSMFVS should be run against at least a few weeks worth of SMF data.

**COMMAND
SUMMARY** IAMSMFVS has the following commands:

DEFAULT - Set global execution defaults.
END - Terminate processing.
HELP - Provide help for documented services.
EXTRACT - Extract cluster/data set utilization records from SMF data for later REPORT processing.
REPORT - Report cluster/data utilization.

40.02 IAMSMFVS JCL REQUIREMENTS

40.02 IAMSMFVS JCL REQUIREMENTS

JCL REQUIREMENTS The JCL statements required to execute IAMSMFVS are as follows:

EXECUTE STATEMENT Specifies the IAM VSAM SMF ANALYSIS program name – IAMSMFVS.

STEPLIB JOBLIB DD STATEMENT This statement may be omitted if the IAM modules are in a system connected library (e.g.: SYS1.LINKLIB). It is recommended that the IAM program load library be included in the system LINKLIST.

SYSUDUMP DD STATEMENT Specifies the ABEND data set used if major errors are detected. Usually a SYSOUT dataset.

SYSPRINT DD STATEMENT Specifies where the IAMSMFVS reports are to be printed. Usually a SYSOUT data set.

SYSMF DD STATEMENT Identifies the file containing the input SMF data. This may be an active SMF dataset or any file containing off loaded SMF data.

SYSEXTIN DD STATEMENT Identifies the input for an IAMSMFVS 'REPORT' to be run from a file containing records selected by the IAMSMFVS 'EXTRACT' command.

SYSIN DD STATEMENT Specifies the input control statement data set. Usually a DD * data set.

SORTIN DD STATEMENT Specifies a work file, created by IAMSMFVS, which will be passed to an external sort routine.

SORTOUT DD STATEMENT Specifies a work file, which will be returned from the external sort routine.

OPTIONAL SORT DD STATEMENTS IAMSMFVS always invokes an external sort. If there are any DD statements needed by your sort program (i.e. SORTLIB, SORTWKnn, etc.), they must be included. Refer to documentation for your sort.

OPTIONAL SYSPUNCH DD IAMSMFVS can optionally generate IAMSIMVS SELECT statements. Whenever IAMSMFVS finds a SYSPUNCH DD statement is present in the JCL it will use the cluster names provided in the Cluster Size Report and create IAMSIMVS SELECT command statements. This DD statement should specify a sequential data set or a member in a PDS capable of accepting 80 character records.

```
ex. //SYSPUNCH DD DSN=SIMVS.INPUT,DISP=(,CATLG),
// UNIT=SYSDA,SPACE=(TRK,(15,15))
```

TSO ALLOCATION The allocations required for the foreground execution of IAMSMF are as follows:

```
ALLOC          F(SYSMF) DA(...smf history file...) SHR
ALLOC          F(SYSPRINT) DA(*)
ALLOC          F(SYSIN) DA(*)
IAMSMFVS
```

NOTE: If the library containing the program IAMSMFVS is not in the LINKLIST, the 'CALL' command or its equivalent must be used to invoke the program. An example of the 'CALL' command is: CALL 'user.load.library(IAMSMFVS)'

SECTION 40: VSAM USAGE ANALYSIS PROGRAM

40.03 IAMSMFVS DEFAULT COMMAND

40.03 IAMSMFVS - DEFAULT COMMAND

The DEFAULT command establishes default values for the IAMSMFVS processing options.

DEFAULT

ABNDFE	,LINECNT=nnnn
,NOABNDFE	,MAXRECLENGTH=nnnn
,DFEFERRPRT=cccccc	,PRTLENGTH=nnnn
,DISABLE=cccccc	,RECSIZE(rrr)=nnnn
,ENABLE=cccccc	,SORTCORE=nnnnnnn

OPERANDS The following operands may be specified on the DEFAULT command.

ABNDFE Specifies abnormal termination if any disk error is encountered reading the SMF data set.

ABN
The default is NOABNDFE.

NOABNDFE Specifies that disk errors are to be documented and bypassed.

NOABN
The default is NOABNDFE.
Disk errors will be documented and bypassed.

DFEFERRPRT= Indicates whether or not error messages from the VSAM ICF identification processor should be printed.

DFEFERR=
AC - All messages to the console
NO - Error messages are not to be printed.
YES - Error messages are to be printed.

The default is NO.

DISABLE= Specifies the options which are to be disabled. The available options are documented under the 'ENABLE' operand. Multiple options may be DISABLED with a single command if entered as follows:

DIS=
DISABLE=(c...c,....,c...c)

ENABLE= Specifies the options that are to be enabled. Available options include:

ENA=

CHECKLENGTH Check the length of each SMF record read. Specifies that SMF data records are to be validated against a table of minimum record lengths. If the user has modified the minimum record length or is executing in a non-compatible system, the correct record lengths may be specified by the 'RECSIZE' operand.

CHE

The CHECKLENGTH default is SMF record lengths are not validated.

40.03 IAMSMFVS DEFAULT COMMAND

ERRORPRINT ERR	<p>Print any SMF record which fails the 'CHECKLENGTH' function. Specifies that any SMF data record which causes an error during processing or fails length verification is to be printed.</p> <p>The ERRORPRINT default is records in error are not printed.</p>
TEMPDSNAMES TEM	<p>The 'EXTRACT' and 'REPORT' commands are to process temporary as well as permanent datasets and VSAM clusters. Specifies that reports produced by the 'REPORT' command will contain information on temporary as well as permanent data sets.</p> <p>The TEMPDSNAMES default is temporary data sets are ignored.</p>
	<p>Multiple options may be ENABLED with a single command if entered as follows: ENABLE=(c...c,...,c...c)</p>
FROMDDNAME= FROMDD=	<p>Specifies the DDNAME of the SMF file to be used as input to IAMSMFVS.</p> <p>The default input DDNAME is SYSMF.</p>
LINECNT= LINE=	<p>Specifies the maximum number of lines to be printed on any report page, may be from 28 to 99 inclusive.</p> <p>The default is 58.</p>
MAXRECLENGTH= MAXR=	<p>Specifies the largest SMF record that the program will process. The number may be any value from 16384 to 65536, inclusive.</p> <p>The default is 16384 bytes.</p>
PRTLENGTH= PRTL=	<p>Limits to 'nnnn' bytes the amount of data to be printed if 'ERRORPRINT' is indicated. The number may be any value from 32 to 65536, inclusive.</p> <p>The default is 32768 bytes.</p>
RECSIZE(rrr)= RECS(rrr)=	<p>Establishes the minimum length of the SMF record type 'rrr' as the value 'nnnn'.</p> <p>The default minimum record lengths for system generated SMF records are documented in IBM SPL: Systems Management Facilities (GC28-0706).</p>
SORTCORE=	<p>Specifies the amount of storage the program 'SORT' is to use if external sorting is required. The number may be any value from 10000 to 8000000 inclusive.</p> <p>The default is 100000.</p>
SORTMSG=	<p>Specifies the message option to be used by the program 'SORT' if external sorting is required.</p> <p>AC - All messages to the console AP - All messages to the printer (SYSOUT) CC - Critical messages to the console CP - Critical messages to the printer NO - No messages to be produced PC - Critical messages to both console and printer</p> <p>The default is CC.</p>
SORTPFX=	<p>Specifies the DDNAME prefix to be used by the program 'SORT' if external sorting is required. If the string specified is less than 4 characters, a dollar sign(\$) fill character will be used.</p> <p>The default is SORT.</p>

40.04 IAMSMFVS - END COMMAND

40.04 IAMSMFVS - END COMMAND

END terminates IAMSMFVS processing it is intended for TSO users.

END

OPERANDS There are no operands for the END command.

40.05 IASMFVS - EXTRACT COMMAND

40.05 IASMFVS - EXTRACT COMMAND

The EXTRACT command is used to extract data set utilization statistics from an SMF history file (RECFM=VBS) or an active SMF data recording file in either SAM or VSAM format. The file created can then be used as input to the 'REPORT' command.

EXTRACT

CHECKLENGTH	,MAXJOBS=nnnn
,DFEFERRPRT=ccc	,MAXRECLENGTH=nnnnnn
,DSGROUPS=ccccc	,MAXREPORTS=nnnn
,DSNAMES=ccccc	,PRTLENGTH=nnnn
,DSORG=ccc	,RECSIZE(rrr)=nnnn
,ERRORPRINT	,SORTCORE=nnnnnn
,FROMDATE=yyddd	,SORTMSG=cc
,FROMDDNAME=ddname	,SORTPFX=ccccc
,FROMTIME=hhmmss	,TEMPDSNAMES
,GROUPNAMES=ccccc	,TODATE=yyddd
,JOBNAMES=ccccc	,TODDNAME=ddname
,MAXDSNS=nnnn	,TOTIME=hhmmss

OPERANDS The following operands may be used with the EXTRACT command.

CHECKLENGTH Specifies that SMF data records are to be validated against a table of minimum record lengths. If the user has modified the minimum record length or is executing in a non-compatible system, the correct record lengths may be specified by the 'RECSIZE' operand.

The default is SMF record length is not validated.

DFEFERRPRT= Indicates whether or not error messages from the VSAM ICF identification processor should be printed.

DFEFERR=

AC - All messages to the console

NO - Error messages are not to be printed.

YES - Error messages are to be printed.

The default is NO.

40.05 IAMSMFVS - EXTRACT COMMAND

- DSGROUPS=** Specifies that only records having a data set name which begin with the given character string(s) will be processed. This operand specifies a partial dataset name from 1 to 44 characters in length. Up to 50 data set groups may be specified for a single command if entered as follows:
DSG=
DSGROUPS=(c...c,...,c...c)
If neither the 'DSGROUPS' nor 'DSNAMES' operand is specified, data set selection will be based upon the value specified for 'DSORG' (or its default).
- DSNAMES=** Specifies that only records having a data set name which match the dataset name(s) specified will be processed. This operand specifies a complete data set name from 1 to 44 characters in length. Up to 50 data set names may be specified for a single command if entered as follows:
DSN=
DSNAMES=(c...c,...,c...c)
If neither the 'DSGROUPS' nor 'DSNAMES' operand is specified, data set selection will be based upon the value specified for 'DSORG' (or its default).
- DSORG=** Identifies the data set organization(s) that is(are) to be processed.
DSO=
ALL - All applicable data set organizations (IAM, BDAM, ISAM, and VSAM).
AM - VSAM clusters, both ICF/VSAM and VSAM.
DA - BDAM data sets.
IAM - IAM data sets.
IN - Indexed data sets, both ISAM and VSAM.
IS - ISAM data sets.
NOTE: To extract data for IAM data sets, SMF recording of IAM SMF records must be enabled, as described in the IAM Users Manual in [section 91](#). To extract data for NONVSAM data sets, SMF record types 14 and 15 must be present.
The default is 'AM' (VSAM) and IAM.
- ERRORPRINT** Specifies that any SMF data record which causes an error during processing or fails length verification is to be printed.
ERR
The default is records in error are not printed.
- FROMDATE=** Specifies the lower date limit of the SMF records which are to be copied.
FROMDA=
The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.
- FROMDDNAME=** Specifies the DDNAME of the SMF file to be used as input to IAMSMF.
FROMDD=
The default input DDNAME is SYSMF.
- FROMTIME=** Specifies the lower time limit of the SMF records which are to be copied.
FROMTI=
]The default is the time of the records does not participate in the selection criteria, unless the operands 'FROMTIME' and/or 'TOTIME' were specified. When used with FROMDATE, forms a combined starting point of date and time. When used without FROMDATE, the FROMTIME applies to all days for which SMF records are being processed.
- GROUPNAMES=** Specifies that only records having a job name which begin with the specified character string(s) will be copied. This operand specifies a partial job name from 1 to 8 characters in length. Up to 50 job groups and/or names may be specified for a single command if entered as follows:
G=
GROUPNAMES=(c...c,...,c...c)
The default, if neither the 'GROUPNAMES' nor 'JOB NAMES' operand is specified, is that the job name will not participate in SMF record selection.

40.05 IAMSMFVS - EXTRACT COMMAND

JOBNAMES= J=	<p>Specifies that only records having a job name which match the jobname(s) specified will be copied. This operand specifies a complete job name from 1 to 8 characters in length. Up to 50 job names may be specified in a single command if entered as follows:</p> <p>JOBNAMES=(c...c,...,c...c)</p> <p>The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the job name will not participate in SMF record selection.</p>
MAXDSNS= MAXD=	<p>Specifies the maximum number of unique data sets and/or clusters which will be tabled during this execution. The number can be any value from 20 to 32000, inclusive.</p> <p><i>NOTE: Each non-VSAM data set takes one (1) entry while each VSAM cluster takes three (3) or more entries, depending upon the number of components.</i></p> <p>The default is 1500 data set table entries.</p>
MAXJOBS= MAXJ=	<p>Specifies the maximum number of unique job names which will be tabled during this execution. The number can be any value from 20 to 32000, inclusive.</p> <p>The default is 5000 job name table entries.</p>
MAXRECLENGTH= MAXR=	<p>Specifies the largest SMF record that the program will process. The number may be any value from 16384 to 65536, inclusive.</p> <p>The default is 16384 bytes.</p>
MAXREPORTS= MAXREP=	<p>Specifies the maximum number of tabled data set and/or cluster names which will be extracted during this execution. The number can be any value from 20 to 32000, inclusive.</p> <p>The default is the top 100 data set and/or cluster names, ranked by EXCP count.</p>
PRTLENGTH= PRTL=	<p>Limits to 'nnnn' bytes the amount of data to be printed if 'ERRORPRINT' is indicated. The number may be any value from 32 to 65536, inclusive.</p> <p>The default is 32768 bytes.</p>
RECSIZE(rrr)= RECS(rrr)=	<p>Establishes the minimum length of the SMF record type 'rrr' as the value 'nnnn'.</p> <p>The default minimum record lengths for system generated SMF records are documented in IBM SPL: Systems Management Facilities (GC28-0706).</p>
SORTCORE=	<p>Specifies the amount of storage the program 'SORT' is to use if external sorting is required. The number may be any value from 10000 to 8000000 inclusive.</p> <p>The default is 100000.</p>
SORTMSG=	<p>Specifies the message option to be used by the program 'SORT' if external sorting is required.</p> <p>AC - All messages to the console AP - All messages to the printer (SYSOUT) CC - Critical messages to the console CP - Critical messages to the printer NO - No messages to be produced PC - Critical messages to both console and printer</p> <p>The default is CC.</p>

40.05 IAMSMFVS - EXTRACT COMMAND

SORTPFX=	Specifies the DDNAME prefix to be used by the program 'SORT' if external sorting is required. If the string specified is less than 4 characters, a dollar sign(\$) fill character will be used. The default is SORT.
TEMPDSNAMES TEM	Information about temporary as well as permanent data sets is to be EXTRACTed. By default, temporary data sets are ignored.
TODATE= TODA=	Specifies the upper date limit of the SMF records which are to be copied. The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.
TODDNAME= TODD=	Specifies the DDNAME of the output file for the copied SMF records. Normally, copies of the EXTRACTed data set utilization records are not made. The records are instead contained in the data set identified by the SORTOUT DD.
TOTIME= TOTI=	Specifies the upper time limit of the SMF records which are to be copied. The default is that the time of the records does not participate in the selections criteria, unless the operands 'TOTIME' and/or 'FROMTIME' are specified. When used with the 'TODATE' keyword, this forms a combined ending point of the date and time specified. When used without the 'TODATE' keyword, the 'TOTIME' is applied to each day for which SMF records are being processed.

40.06 IAMSMFVS - HELP COMMAND

40.06 IAMSMFVS - HELP COMMAND

HELP COMMAND The HELP command prints help by function or by command.

HELP

ALL

,SUBCOMMAND(ccccc)

,SYNTAX

OPERANDS The following operands may be specified on the HELP command.

ALL Lists the entire contents of the help member, for all commands and field name descriptions.

SUBCOMMAND Requests full help for the command 'cccc'.

SU *NOTE: 'JCL' and 'CLIST' are also valid operands for the SUBCOMMAND parameter.*

SYNTAX When specified with the operand SUBCOMMAND(ccccc), requests help for only the syntax of 'cccc'. When specified without the operand SUBCOMMAND(ccccc), requests help for the syntax of the command or program.

40.07 IAMSMFVS - REPORT COMMAND

40.07 IAMSMFVS - REPORT COMMAND

REPORT COMMAND The REPORT command is used to request a data set utilization report. The input may be a SMF history file (RECFM=VBS), a system SMF data set recording in either SAM or VSAM format, or a data set created by the 'EXTRACT' command of IAMSMFVS. All operands are optional.

It is recommended that initially no operands be specified. The default reports would then list the 100 most active and the 100 largest VSAM clusters.

REPORT

CHECKLENGTH	,JOBNAMES=cccc
,CURRENT	,MAXDSNS=nnnn
,DATATYPE=ccc	,MAXJOBS=nnnn
,DETAIL	,MAXRECLENGTH=nnnnn
,DFEFERRPRT=ccc	,MAXREPORTS=nnnn
,DSGROUPS=cccc	,PRTLENGTH=nnnnn
,DSNAMES=cccc	,RECSIZE(rrr)=nnnnn
,DSORG=ccc	,SORTCORE=nnnnn
,ERRORPRINT	,SORTMSG=cc
,FROMDATE=yyddd	,SORTPFX=cccc
,FROMDDNAME=ddname	,TEMPDSNAMES
,FROMTIME=hhmmss	,TODATE=yyddd
,GROUPNAMES=cccc	,TOTIME=hhmmss

OPERANDS The following operands may be specified on the REPORT command.

CHECKLENGTH Specifies that SMF data records are to be validated against a table of minimum record lengths. If the user has modified the minimum record length or is executing in a non-compatible system, the correct record lengths may be specified by the 'RECSIZE' operand.

By default, the length of an SMF record is not validated.

CURRENT Specifies that only the most recent values for overflow use in IAM files will be reported on, rather than the maximum amount used. This may be useful for determining when an IAM file needs to be reorganized.

By default, the maximum amount of overflow use is reported on.

40.07 IAMSMFVS - REPORT COMMAND

DATATYPE DATA	Identifies the type of input data to be processed. EXTRACT - Output from the 'EXTRACT' command. SMF - SMF data from either an active SMF data set or a sequential history file. The default is SMF.
DETAIL DET	Data set utilization reports are to include jobname and program statistics. By default, only data usage is reported.
DFEFERRPRT= DFEFERR=	Indicates whether or not error messages from the VSAM ICF identification processor should be printed. AC - All messages to the console NO - Error messages are not to be printed. YES - Error messages are to be printed. The default is NO.
DSGROUPS= DSG=	Specifies that only records having a data set name which begin with the given character string(s) will be processed. This operand specifies a partial dataset name from 1 to 44 characters in length. Up to 100 data set groups may be specified for a single command if entered as follows: DSGROUPS=(c...c,...,c...c) If neither the 'DSGROUPS' nor 'DSNAMES' operand is specified, data set selection will be based upon the value specified for 'DSORG' (or its default).
DSNAMES= DSN=	Specifies that only records having a data set name which match the dataset name(s) specified will be processed. This operand specifies a complete data set name from 1 to 44 characters in length. Up to 100 dataset names may be specified for a single command if entered as follows: DSNAMES=(c...c,...,c...c) If neither the 'DSGROUPS' nor 'DSNAMES' operand is specified, data set selection will be based upon the value specified for 'DSORG' (or its default).
DSORG= DSO=	Identifies the data set organization(s) that is(are) to be processed. ALL - All applicable data set organizations (IAM, BDAM, ISAM, and VSAM). AM - VSAM clusters, both ICF/VSAM and VSAM. DA - BDAM data sets. IAM - IAM data sets. IN - Indexed data sets, both ISAM and VSAM. IS - ISAM data sets. The default is 'AM' (VSAM) and IAM. <i>NOTE: To extract data for IAM data sets, SMF recording of IAM SMF records must be enabled, as described in the IAM Users Manual in section 91. To extract data for NONVSAM data sets, SMF record types 14 and 15 must be present.</i>
ERRORPRINT ERR	Specifies that any SMF data record which causes an error during processing or fails length verification is to be printed. The default is records in error are not printed.

40.07 IAMSMFVS - REPORT COMMAND

FROMDATE= FROMDA=	Specifies the lower date limit of the SMF records which are to be copied. The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.
FROMDDNAME= FROMDD=	Specifies the DDNAME of the SMF file to be used as input to IAMSMF. The default input DDNAME is SYSMF.
FROMTIME= FROMTI=	Specifies the lower time limit of the SMF records which are to be analyzed. The default is the time of the records does not participate in the selection criteria, unless the operands 'FROMTIME' and/or 'TOTIME' were specified. When used with FROMDATE, forms a combined starting point of date and time. When used without FROMDATE, the FROMTIME applies to all days for which SMF records are being processed.
GROUPNAMES= G=	Specifies that only records having a job name which begin with the specified character string(s) will be copied. This operand specifies a partial job name from 1 to 8 characters in length. Up to 50 job groups and/or names may be specified for a single command if entered as follows: GROUPNAMES=(c...c,....c...c) The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the job name will not participate in SMF record selection.
JOBNAMES= J=	Specifies that only records having a jobname which match the jobname(s) specified will be copied. This operand specifies a complete jobname from 1 to 8 characters in length. Up to 50 jobnames may be specified in a single command if entered as follows: JOBNAMES=(c...c,....c...c) The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the job name will not participate in SMF record selection.
MAXDSNS= MAXD=	Specifies the maximum number of unique data sets and/or clusters which will be tabled during this execution. The number can be any value from 20 to 32000, inclusive. <i>NOTE: Each non-VSAM data set takes one (1) entry while each VSAM cluster takes three (3) or more entries, depending upon the number of components.</i> The default is 1500 data set table entries.
MAXJOBS= MAXJ=	Specifies the maximum number of unique job names which will be tabled during this execution. The number can be any value from 20 to 32000, inclusive. The default is 5000 job name table entries.
MAXRECLENGTH= MAXR=	Specifies the largest SMF record that the program will process. The number may be any value from 16384 to 65536, inclusive. The default is 16384 bytes.
MAXREPORTS= MAXREP=	Specifies the maximum number of tabled data set and/or cluster names to be reported upon during this execution. The number can be any value from 20 to 32000, inclusive. The default is: first an EXCP Activity Report will list the 100 most active data sets and/or clusters ranked according to their total EXCP counts. Next a Data Set Summary Report will be provided which will list alphabetically the names of all tabled data sets and clusters. Finally a Size Report will list the 100 largest data sets and/or clusters ranked according to the size of their total allocation.

40.07 IAMSMFVS - REPORT COMMAND

- PRTLENGTH=** Limits to 'nnnn' bytes the amount of data to be printed if 'ERRORPRINT' is indicated. The number may be any value from 32 to 65536, inclusive.
- PRTL=** The default is 32768 bytes.
- RECSIZE(rrr)=** Establishes the minimum length of the SMF record type 'rrr' as the value 'nnnn'.
- RECS(rrr)=** The default minimum record lengths for system generated SMF records are documented in IBM SPL: Systems Management Facilities (GC28-0706).
- SORTCORE=** Specifies the amount of storage the program 'SORT' is to use if external sorting is required. The number may be any value from 10000 to 8000000 inclusive.
- The default is 100000.
- SORTMSG=** Specifies the message option to be used by the program 'SORT' if external sorting is required.
- AC - All messages to the console**
AP - All messages to the printer (SYSOUT)
CC - Critical messages to the console
CP - Critical messages to the printer
NO - No messages to be produced
PC - Critical messages to both console and printer
- The default is CC.
- SORTPFX=** Specifies the DDNAME prefix to be used by the program 'SORT' if external sorting is required. If the string specified is less than 4 characters, a dollar sign(\$) fill character will be used.
- The default is SORT.
- TEMPDSNAMES** Specifies that reports produced by the 'REPORT' command will contain information on temporary as well as permanent data sets.
- TEM** By default, temporary data sets are ignored.
- TODATE=** Specifies the upper date limit of the SMF records which are to be copied.
- TODA=** The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.
- TOTIME=** Specifies the upper time limit of the SMF records which are to be analyzed.
- TOTI=** The default is that the time of the records does not participate in the selections criteria, unless the operands 'TOTIME' and/or 'FROMTIME' are specified. When used with the 'TODATE' keyword, this forms a combined ending point of the date and time specified. When used without the 'TODATE' keyword, the 'TOTIME' is applied to each day for which SMF records are being processed.

40.08 IAMSMFVS SAMPLE REPORTS

40.08 IAMSMFVS SAMPLE REPORTS

IAMSMFVS supplies the user with a SUMMARY and three reports. Those reports are the EXCP ACTIVITY REPORT, the CLUSTER NAME REPORT, and the CLUSTER SIZE REPORT.

IAMSMFVS SUMMARY

The SUMMARY REPORT shows the period of time the reports cover (JOBS – DATED yy.ddd THRU yy.ddd) and it summarizes the impact that VSAM and IAM had on the system during that period in terms of I/O activity and disk space utilization. It does this for I/O activity by comparing the number of EXCP's issued against VSAM KSDS clusters, to the total number of disk EXCP's executed during the report period. A summary of the total disk space allocated to VSAM KSDS clusters and IAM files is also included.

IAM400 SMF REPORT/DATA EXTRACT PROGRAM--IAMSMFVS VER 6.X-INNOVATION DATA PROCESSING	
IAMSMFVS SUMMARY REPORT	
IAM601 SMF RECORDS -- READ...151867 USED...54582 DROPPED.....0	
546 DATA SETS TABLED REPRESENTING 1051 JOBS -- DATED yy.ddd THRU yy.ddd	
SPACE UTILIZATION SUMMARY --	
DEVICE TYPE.....3380	TOTAL VSAM CYL/TRK.....45072/00
TOTAL DISK EXCPS.....593097575	TOTAL VSAM EXCPS.....193366518

40.08 IAMSMFVS SAMPLE REPORTS

EXCP ACTIVITY REPORT

The ACTIVITY report is the first report type produced, with separate reports for VSAM and IAM files. These reports identifies those VSAM clusters or IAM files that have the highest EXCP counts. This report lists clusters (DATA COMPONENT \$PL INDEX COMPONENT) with the highest activity (total EXCPs) first. The cluster name is reported on the first line. The components are reported next in the order of the component which has the most EXCPs (DATA or INDEX). A USE COUNT is included in the report which shows the number of times the cluster was OPENed and CLOSEd during the report period.

The default is to report on the 100 most active data sets or for each data set organization.

IAM400 SMF REPORT/DATA EXTRACT PROGRAM-- IAMSMFVS VER 6.x-- INNOVATION DATA PROCESSING DATE-yy.ddd										
VSAM EXCP REPORT										
DATA SET NAME	USE COUNT	TOTAL EXCPS	RECORDS	READS	INSERTS	UPDATES	DELETES	SPLITS		ALLOC
								CI	CA	TRKS
COMPANY.MASTER	437	23396841								
COMPANY.MASTER.INDEX	437	11736876	1310	0	0	0	0	39	6	15
COMPANY.MASTER.DATA	437	11659965	513571	42389011	61571	72990	20074	2143	24	8151
ACCOUNT.MASTER	101	9468410								
ACCOUNT.MASTER.DATA	101	7022603	373571	9799323	30571	46580	19446	1743	17	6150
ACCOUNT.MASTER.INDEX	101	2445807	912	0	0	0	0	49	4	26
SMPE.TDFP223.CSI	345	2150620								
SMPE.TDFP223.DATA	345	1859520	537085	13199605	21089	5460	5534	128	14	899
SMPE.TDFP223.INDEX	345	291100	40	0	0	0	0	0	0	16

The EXCP report identifies your high impact VSAM clusters. High activity clusters are attractive candidates for conversion to IAM because, **compared to VSAM, IAM reduces disk I/O 50% to 80%.**

NOTE: If components are listed without an associated cluster name IAM could not find the named components in any available catalog.

40.08 IAMSMFVS SAMPLE REPORTS

CLUSTER/ DATA SET NAME REPORT The second report lists VSAM clusters and IAM files in name order. The data set or cluster name is reported on the first line. The DATA component is reported next followed by the INDEX. This report details the component's file and record characteristics. If a cluster is not currently cataloged, some of the file and record characteristic information may not appear in the report.

If more than two components are reported for a cluster, the cluster contains alternate indexes.

The default is to report on the first 1500 VSAM clusters reported by SMF.

IAM400 SMF REPORT/DATA EXTRACT PROGRAM--IAMSMFVS VER 6.X--INNOVATION DATA PROCESSING											
DATA SET SUMMARY REPORT											
DATA SET NAME	USE COUNT	TOTAL EXCPS	DSORG	RECFM	AVG LRECL	MAX LRECL	KEY LEN	RKP	BLK OR C I SIZE	FRSPC CI%	CA%
COMPANY.MASTER	437	23396841	VSAM								
COMPANY.MASTER.DATA	437	11659965	VSAM	VB	383	580	23	0	5800		
COMPANY.MASTER.INDEX	437	11736876	VSAM	IMBED		1529	23	0	1536	10	5
ACCOUNT.MASTER	101	9468410	VSAM								
ACCOUNT.MASTER.DATA	101	7022603	VSAM	FB	256	256	18	0	4096	10	5
ACCOUNT.MASTER.INDEX	101	2445807	VSAM	NOIMB		1124	18	0	1131		
SMPE.TDFP223.CSI	345	2150620	VSAM								
SMPE.TDFP223.DATA	345	1859520	VSAM	VB	23	142	23	0	4096	10	5
SMPE.TDFP223.INDEX	345	291100	VSAM	IMBED		1529	23	0	1536		

NOTE: If components are listed without an associated cluster name IAM could not find the named components in any available catalog.

40.08 IAMSMFVS SAMPLE REPORTS

CLUSTER SIZE REPORT The third report is a size report, with separate reports for VSAM and IAM files, each in size order with the largest first.

In the VSAM Size Report, cluster name is reported on the first line with the total allocation for the cluster given in tracks. The DATA component and INDEX follow, listed based on the size of their individual allocations. This report also gives the number of extents the components occupy, their file and record characteristics. If a cluster is not currently cataloged some of the file and record characteristic information may not appear in the report.

The IAM size report is similar to the VSAM Size Report, except that the size is the amount of tracks actually used by the IAM files. The report also provides information on the use of the independent overflow area, and the size of the integrated overflow and prime extension areas.

The default is to report on the 100 largest VSAM clusters.

IAM400 SMF REPORT/DATA EXTRACT PROGRAM--IAMSMFVS VER 6.x--INNOVATION DATA PROCESSING DATE-yy.ddd									
CLUSTER SIZE REPORT									
DATA SET NAME	ALLOC TRKS	TOTAL EXCPS	USE COUNT	EXTENTS	AVG LRECL	MAX LRECL	KEY LEN	RKP	CISIZE
BIG.CLUSTER	37155	2507803	24						
BIG.CLUSTER.DATA	37100	2105001	24	5	223	580	28	0	4096
BIG.CLUSTER.INDEX	55	402802	24	1		1529	28	0	1536
A.FILE.SMALLER	16540	679216	159						
A.FILE.SMALLER.DATA	16500	270501	159	2	208	208	9	0	8192
A.FILE.SMALLER.INDEX	40	408715	159	1		4059	9	0	4096
SMPE.TDFP230.CSI	12315	3980211	15						
SMPE.TDFP230.CSI.DATA	12300	3075021	15	1	23	142	23	0	4096
SMPE.TDFP230.CSI.INDEX	15	805190	15	1		1529	23	0	4096
IAM499 IAMSMFVS(6.x) PROCESSING COMPLETED									

Use the cluster size report to identify the largest VSAM clusters in your installation. **IAM internal file structure saves 20% to 40% in disk space versus VSAM. With data compression, IAM can reduce a file's disk space an additional 20% to 50%.**

NOTE: If components are listed without an associated cluster name IAM could not find the named components in any available catalog. Additionally in this case the report will show an exaggerated number of tracks allocated for index components with the IMBED attribute. The allocation shown will include the extents of the data component containing the imbedded index.

40.09 IAMSMFVS EXAMPLES

40.09 IAMSMFVS EXAMPLES

1. List the 100 most active and the 100 largest VSAM clusters

DEFAULT REPORT This is an example of how to initially run IAMSMFVS. In this example, all the defaults are taken. The default reports list the top 100 most active clusters (by EXCP count), followed by a report of all the tabled names in cluster component name order and finally a list of the 100 largest clusters in size order.

This sample JCL is in the IAM Installation Control Library member named 'IAMSMFVS'.

```
//IAMSMFVS EXEC PGM=IAMSMFVS,REGION=1024K
//STEPLIB DD DISP=SHR,DSN=IDP.MODIAM60<=POINTS TO IAMLIBRARY
//SYSMF DD DISP=SHR,DSN=..... <=POINTS TO SMF DATA
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SORTIN DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTOUT DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTLIB DD DISP=SHR,DSN=SYS1.SORTLIB
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SYSOUT DD SYSOUT=*
//SYSIN DD *
REPORT
/*
```

2. Produce a selective DETAIL Report

DETAIL REPORT In this example IAMSMFVS is to produce a detail report, of all clusters that have names beginning with CICS, as set by the DSG= keyword. DETAIL is used to request a report with job/step/program information. This would normally be the second step in selecting candidate VSAM clusters for conversion to IAM. Only a selected data set group is requested because a detail report can be quite long.

```
//IAMSMFVS EXEC PGM=IAMSMFVS,REGION=1024K
//STEPLIB DD DISP=SHR,DSN=IDP.MODIAM60<=POINTS TO IAMLIBRARY
//SYSMF DD DISP=SHR,DSN=..... <=POINTS TO SMF DATA
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SORTIN DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTOUT DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTLIB DD DISP=SHR,DSN=SYS1.SORTLIB
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(10,5))
//SYSOUT DD SYSOUT=*
//SYSIN DD *
REPORT DSG=CICS,DETAIL
/*
```

40.09 IAMSMFVS EXAMPLES

3. Generate IAMSIMVS SELECT command statements for the largest VSAM clusters, Change IAMSMFVS capacity defaults.

INCREASE CLUSTER TABLE CAPACITY If you have a very large number of VSAM clusters, and/or large SMF records, it may be necessary to change some of the default parameters in IAMSMFVS. There are a number of parameters that directly control the number of clusters IAMSMFVS can report on and the total number of SMF records it can process.

If a non-zero value for RECORDS DROPPED appears in message IAM601, it is possible some SMF data may have been bypassed. The buffer area for the SMF records needs to be increased. This is specified by the MAXRECLENGTH keyword. In this example, the value is increased to 32768 bytes (32K).

If either message IAM621 and/or IAM622 appear, IAMSMFVS's data table is not large enough to hold all of the VSAM clusters that meet the selection criteria. VSAM clusters typically need three entries in the EXTRACTed data table, one for each component plus one for the cluster level. The number of table entries is changed by the MAXDSNS parameter. In this example, up to 2000 VSAM clusters are desired, so MAXDSNS is set to three times that number, which is 6000.

Changing these values will result in larger region requirements, as well as additional DASD work space. In this example, the REGION size has been increased to 2048K, and the space allocations for the work files, SORTWKxx, SORTIN, and SORTOUT have been increased. These sizes are estimates only, and may need to be revised dependent on the type of work DASD and the SORT product used.

CREATE IAMSIMVS SELECT STATEMENTS Additionally a SYSPUNCH DD statement is added. This will cause IAMSMFVS to create SELECT command statements that can be used as input by IAMSIMVS. These statements will contain the same cluster names as are shown on the Cluster Size Report. This output dataset may be either a sequential dataset or a member in a PDS as long as it is capable of accepting 80 character records.

```
// IAMSMFVS EXEC PGM=IAMSMFVS,REGION=2048K
// STEPLIB DD DISP=SHR,DSN=i.am.library <=USER CHANGE
// SYSMF DD DISP=SHR,DSN=... <=POINTS TO SMF DATA
// SYSPRINT DD SYSOUT=*
// SORTIN DD UNIT=SYSDA,SPACE=(CYL,(25,5))
// SORTOUT DD UNIT=SYSDA,SPACE=(CYL,(25,5))
// SORTLIB DD DISP=SHR,DSN=SYS1.SORTLIB
// SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(25,5))
// SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(25,5))
// SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(25,5))
// SYSOUT DD SYSOUT=*
// SYSPUNCH DD DSN=VSAM.CLUSTER.NAMES
// SYSIN DD *
REPORT MAXRECLENGTH=32768,MAXDSNS=6000
/*
```

40.09 IAMSMFVS EXAMPLES

4. Create an EXTRACT file of the VSAM records in an SMF history file

EXTRACT RECORDS In this example IAMSMFVS is to first create an EXTRACT of the VSAM related data from an SMF history file, then use that EXTRACT to produce a report listing the most active and largest clusters. This EXTRACT file can then be used later to produce DETAIL Reports on the selected VSAM clusters. This technique eliminates having to pass a multi-volume SMF history file subsequent times.

```
//EXTRACT      EXEC   PGM=IAMSMFVS,REGION=1024K
//STEPLIB      DD     DISP=SHR,DSN=IDP.MODIAM60<=POINTS TO IAMLIBRARY
//SYSMF        DD     DISP=SHR,DSN=.....          <=POINTS TO SMFDATA
//SYSPRINT     DD     SYSOUT=*
//SYSUDUMP     DD     SYSOUT=*
//SORTIN       DD     UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTOUT      DD     UNIT=SYSDA,SPACE=(CYL,(10,5))
//SYSEXTIN     DD     DSN=EXTRACT.DATA,          <=POINTS TO EXTRACT FILE
//              DD     UNIT=SYSDA,SPACE=(CYL,(10,5)),
//              DD     DISP=(,CATLG,DELETE)
//SORTLIB      DD     DISP=SHR,DSN=SYS1.SORTLIB
//SORTWK01     DD     UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTWK02     DD     UNIT=SYSDA,SPACE=(CYL,(10,5))
//SORTWK03     DD     UNIT=SYSDA,SPACE=(CYL,(10,5))
//SYSOUT       DD     SYSOUT=*
//SYSIN        DD     *
                EXTRACT   TODD=SYSEXTIN
                REPORT    DATATYPE=EXTRACT
/*
```

NOTE: The JOBS processed and TOTAL EXCP statistics in the Heading Summary of an IAMSMFVS report are based on the input data. When an EXTRACT file is used as input these statistics will reflect the fact that the input was a subset of the original SMF history file. They will be proportionally lower than the Summary Statistics for the original history file.

SECTION 41: SMF JOB RESOURCE USAGE ANALYSIS PROGRAM

41.01 SMF JOB RESOURCE USAGE ANALYSIS OVERVIEW

41.01 SMF JOB RESOURCE USAGE ANALYSIS OVERVIEW

OVERVIEW A breakout by job of resource usage is vital if the overall impact of an application is to be evaluated. System utilization reports based on SMF data are produced by the utility IAMSMF. IAMSMF reports job oriented SMF information based upon a jobname or job group. IAMSMF can also be used to query as to which jobs are referencing specific data sets and to identify which job/step(s) DEFINE and/or DELETE VSAM clusters.

If IAM's SMF recording has been enabled (SMF=YES in the IAM Global Option Table), IAMSMF can be used to generate IAMINFO reports from those SMF records.

TECHNICAL DESCRIPTION IAMSMF is a batch program, which based upon simple control statements, will extract resource utilization information from SMF job/step and data set records. Information can be selected directly from the active system SMF datasets, from SMF format archival files, or from extract files created by the COPY function of the program itself.

Help is available for the IAMSMF commands. In addition, the user can request help for JCL or CLIST which will show the required allocations for background or foreground processing. Use 'HELP HELP' to obtain information on the IAMSMF 'HELP' syntax.

Under normal circumstances, IAMSMF requires approximately 512K to execute.

FUNCTIONAL SUMMARY The IAMSMF program has been designed to give users the ability to extract data and statistics from Systems Management Facility (SMF) records. The functions include data set query, job query, record copy, and record print.

IAMSMF reports: Jobname - Step Name - Executing Program Name
Step Elapsed Time - Step CPU (and MVS SRB) Time
Memory Utilization - Paging Activity
Summarized EXCPs for Disk and Tape

Optional data set utilization information can show in detail:

DATASET NAME:
DDNAME - DSORG - EXCP Count
ACTIVITY/USE: DEFINE / OUTPUT / INPUT / DELETE

Additionally users can display various MVS/SRM statistics:

Task Active Time - Task Inactive Time
Task Residency Time - Service Units
Swap Count - Swap-In - Swap-Out

COMMAND SUMMARY IAMSMF has the following commands:

COPY - Copy selected SMF records to a sequential file creating a subset of the data for later use in detailed analysis.

DEFAULT - Set global execution defaults.

END - Terminate processing under TSO.

HELP - Provide help for documented services.

IAMINFO - Produce IAMINFO reports from IAM generated SMF records.

PRINT - Print selected SMF records (Hex print format).

QUERY - Report on which job/step(s) use specific data sets and how those data sets are used.
Can be requested for a specific Data Set Name(s) or by Data Set Group Name(s).

REPORT - Report on job/step/TSO session system resource usage in chronological order from the available SMF data. Can be requested for a specific Jobname(s) or by Job Group Name(s).

41.02 JCL REQUIREMENTS FOR IAMSMF

41.02 JCL REQUIREMENTS FOR IAMSMF

JCL REQUIREMENTS The JCL statements required to execute IAMSMF are as follows:

EXECUTE STATEMENT Specifies the name of the IAM SMF ANALYSIS program - IAMSMF. The region size used by IAMSMF is at least 512K.

STEPLIB or JOBLIB DD STATEMENT It is recommended that the IAM program load library be included in the system LINKLIST. If it is not you must include a STEPLIB/JOBLIB DD statement specifying the library in which the IAMSMF program load modules reside.

May be omitted if these modules are in a system connected library (e.g.: SYS1.LINKLIB).

SYSUDUMP DD STATEMENT Specifies the ABEND data set used if major errors are detected. Usually a SYSOUT data set.

SYSPRINT DD STATEMENT Specifies where the IAMSMF control statements, messages and reports are to be printed. Usually a SYSOUT data set.

SYSMF DD STATEMENT Specifies the SMF source data set. This DD statement may point to one or more of the active SMF data sets on disk (ex: SYS1.MANx), an SMF format archival file, or a sequential file produced by the COPY function of this program. An alternate DD name may be specified via a control statement.

SYSUT2 DD STATEMENT Specifies the output data set for a COPY operation. Usually a sequential file on tape or disk. An alternate DD name may be specified via a control statement.

SYSIN DD STATEMENT Specifies the input control statement data set. Usually a DD * data set.

TSO ALLOCATION The allocations required for the foreground execution of IAMSMF are as follows:

```
ALLOC      F(SYSMF) DA( smf file name ) SHR
ALLOC      F(SYSPRINT) DA(*)
ALLOC      F(SYSIN)  DA(*)
IAMSMF
```

NOTE: If the library containing the program IAMSMF is not in the LINKLIST, the 'CALL' command or its equivalent must be used to invoke the program. An example of the 'CALL' command is: CALL 'user.load.library(IAMSMF)'

41.03 IAMSMF - COPY COMMAND

41.03 IAMSMF - COPY COMMAND

COPY COMMAND The COPY command is used to copy selected SMF records to a sequential file from either a history file (RECFM=VBS) or an active SMF data recording file.

COPY

ALLRECORDS	,JOBNAMES=cccc
,CHECKLENGTH	,MAXRECLENGTH=nnnn
,ERRORPRINT	,PRTLENGTH=nnn
,FROMDATE=yyddd	,RECSIZE(rrr)=nnnn
,FROMDDNAME=ddname	,RECTYPE=nnn
,GROUPNAMES=cccc	,TODATE=yyddd
,IAMRECORDS	,TODDNAME=ddname

OPERANDS The following operands may be specified on the COPY command.

ALLRECORDS Specifies that the SMF record type does not participate in record selection.

ALLR The default is deferred to the operand 'RECTYPE'.

NOTE: This operand conflicts with the operand 'RECTYPE'.

CHECKLENGTH Specifies that SMF data records are to be validated against a table of minimum record lengths. If the user has modified the minimum record length or is executing in a non-compatible system, the correct record lengths may be specified by the 'RECSIZE' operand.

CHE

The default is SMF record length is not validated.

ERRORPRINT Specifies that any SMF data record which causes an error during processing or fails length verification is to be printed.

ERR

The default is records in error are not printed.

FROMDATE= Specifies the lower date limit of the SMF records which are to be copied.

FROMDA=

The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.

FROMDDNAME= Specifies the DDNAME of the SMF file to be used as input to IAMSMF.

FROMDD=

The default input DDNAME is SYSMF.

41.03 IAMSMF - COPY COMMAND

GROUPNAMES= G=	<p>Specifies that only records having a jobname which begin with the specified character string(s) will be copied. This operand specifies a partial jobname from 1 to 8 characters in length. Up to 50 job groups and/or names may be specified for a single command if entered as follows:</p> <p>GROUPNAMES=(c...c,...,c...c)</p> <p>The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the jobname will not participate in SMF record selection.</p>
IAMRECORDS IAMR	<p>Specifies that SMF records generated by IAM are to be selected.</p> <p>The default is deferred to the operand 'RECTYPE'.</p> <p><i>NOTE: This operand conflicts with the operands 'RECTYPE' and 'ALLRECORDS'</i></p>
JOBNAMES= J=	<p>Specifies that only records having a jobname which match the jobname(s) specified will be copied. This operand specifies a complete jobname from 1 to 8 characters in length. Up to 50 jobnames may be specified in a single command if entered as follows:</p> <p>JOBNAMES=(c...c,...,c...c)</p> <p>The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the jobname will not participate in SMF record selection.</p>
MAXRECLENGTH= MAXR=	<p>Specifies the largest SMF record that the program will process. The number may be any value from 16384 to 65536, inclusive.</p> <p>The default is 16384 bytes.</p>
PRTLLENGTH= PRTL=	<p>Limits to 'nnnn' bytes the amount of data to be printed if 'ERRORPRINT' is indicated. The number may be any value from 32 to 65536, inclusive.</p> <p>The default is 32768 bytes.</p>
RECSIZE(rrr)= RECS(rrr)=	<p>Establishes the minimum length of the SMF record type 'rrr' as the value 'nnnn'.</p> <p>The default minimum record lengths for system generated SMF records are documented in IBM SPL: Systems Management Facilities (GC28-0706).</p>
RECTYPE= RECT=	<p>Identifies the specific record type(s) to be copied. Up to 50 record types may be specified for a single command if entered as follows:</p> <p>RECTYPE=(rrr,...,rrr)</p> <p>The SMF record types which will be copied by default (if 'ALLRECORDS' is not specified) are as follows:</p> <ul style="list-style-type: none">4 - Step termination5 - Job termination14 - NON-VSAM Dataset CLOSEd (input)15 - NON-VSAM Dataset CLOSEd (output/update)20 - Job initiation30 - Common Address Space Work Record34 - TSO session termination64 - VSAM Dataset CLOSEd <p><i>NOTE: This operand conflicts with the operand 'ALLRECORDS'.</i></p>
TODATE= TODA=	<p>Specifies the upper date limit of the SMF records which are to be copied.</p> <p>The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.</p>
TODDNAME= TODD=	<p>Specifies the DDNAME of the output file for the copied SMF records.</p> <p>The default output DDNAME is SYSUT2.</p>

41.04 IAMSMF - DEFAULT COMMAND

41.04 IAMSMF - DEFAULT COMMAND

DEFAULT COMMAND The DEFAULT command is used to establish default values for IAMSMF processing options. If permanent defaults are to be established, a SYSLIB DD must be present pointing to the load module library (PDS) containing the member 'IAMSMFOP'. The user can override the DDNAME of SYSLIB by specifying the 'SYSLIB' operand with a DDNAME of their own choosing.

DEFAULT

ABNDFE	,MAXRECLENGTH=nnnn
,NOABNDFE	,MAXSTACK=nnnn
,CHANGE=cccc	,PRTLENGTH=nnnn
,DISABLE= cccc	,RECSIZE(rrr)=nnnn
,ENABLE=cccc	,SELTERR=nnnn
,FORMAT=cccc	,SORTCORE=nnnn
,FROMDDNAME=ddname	,SORTMSG=cc
,LINECNT=nn	,SORTPFX=ccc
,MAXCORE=nnnn	,SYSLIB=ddname
,MAXPRINT=nnnn	,TODDNAME=ddname

OPERANDS The following operands may be specified on the DEFAULT command.

ABNDFE	Specifies that IAMSMF is to abnormally terminate if any error is encountered reading the SMF data set.
ABN	
NOABNDFE	Specifies that disk errors are to be documented and bypassed.
NOABN	The default is NOABNDFE, report all errors and continue processing.

41.04 IAMSMF - DEFAULT COMMAND

- CHANGE=** Specifies the duration of the defaults.
- NOTE: If you wish to permanently change the options in the IAM library you must have sufficient authority to do so.*
- PERMANENT** - Requests that the changes be written back to disk, utilizing the DDNAME specified by the 'SYSLIB' operand.
- RESET** - Requests that the defaults be reset to the level provided on the distribution tape and written back to disk, utilizing the DDNAME specified by the 'SYSLIB' operand.
- TEMPORARY** - Indicates that the defaults to be altered only for the duration of this job. No changes will be written back to disk.
- The default is 'TEMPORARY'.
- DISABLE=** Specifies the options that are to be disabled. The available options are documented under the 'ENABLE' operand.
- Multiple options may be disabled with a single command if entered as follows:
DISABLE=(c...c,....c...c)
- ENABLE=** Specifies the options that are to be enabled. Available options include:
- ALLDSNAMES** Specifies that the IAMSMF report will show all the occurrences of any referenced dataset name. The EXCP count displayed will reflect each OPEN/CLOSE.
ALLD As a default, data set usage counts (EXCPs) are accumulated for each permanent data set, and the data set is only shown once. Temporary data sets are ignored.
- NOTE: This operand conflicts with the operand 'NODSNAMES'.*
- ALLRECORDS** Record type does not participate in record selection with the
ALLR 'COPY' and 'PRINT' commands.
- The SMF record types which will be copied by default (if 'ALLRECORDS' is not specified) are as follows:
- 4 - Step termination**
 - 5 - Job termination**
 - 14 - NON-VSAM Dataset CLOSEd (input)**
 - 15 - NON-VSAM Dataset CLOSEd (output/update)**
 - 20 - Job initiation**
 - 30 - Common Address Space Work Record**
 - 34 - TSO session termination**
 - 64 - VSAM Dataset CLOSEd**
- CHECKLENGTH** Specifies that SMF data records are to be validated against
CHE a table of minimum record lengths. If the site has modified the minimum record length or is executing in a noncompatible system, the correct record lengths may be specified by the 'RECSIZE' operand.
- By default, the length of an SMF record is not validated.
- DISPLAY** Display the current SMF related defaults after each
DISP DEFAULT command issued.

41.04 IAMSMF - DEFAULT COMMAND

DISRECSIZE DISR	Display the current SMF minimum record size table after each DEFAULT command issued.
ERRORPRINT ERR	Specifies that any SMF data record which causes an error during processing or fails length verification is to be printed. By default, records in error are not printed.
NODSNAMES NOD	Specifies that the report produced will not show dataset oriented information. By default, data set usage information is displayed only once for each permanent data set with accumulated statistics for multiple OPEN/CLOSE. Temporary data sets are ignored. <i>NOTE: This operand conflicts with 'ALLDSNAMES' and 'TEMPDSNAMES'.</i>
TEMPDSNAMES TEM	Specifies that reports produced by the 'REPORT' command will contain information on temporary as well as permanent datasets. By default, temporary data sets are ignored. <i>NOTE: This operand conflicts with 'NODSNAMES'.</i>

Multiple options may be enabled with a single command if entered as follows:

ENABLE=(c...c,...,c...c)

FORMAT=	Specifies the report is to be prepared using other than the default selected by the program.
CRT	Produces a report without heading lines using a line length of 78, characters for use on a terminal.
DEVICE	Produces a report having the format of either CRT or PRINT, based on the assigned output device. Device determination is based upon the bit setting of the terminal indicator flag in the TIOT.
PRINT	Produces a report with multiple column heading and a line length of 120, for use on hard-copy printers.
TSO	Produces a report with multiple column heading lines using a line length of 78, for use on terminals. The default for format is DEVICE.
FROMDDNAME= FROMDD=	Specifies the DDNAME of the SMF file to be used as input to IAMSMF. The default input DDNAME is SYSMF.
LINECNT= LINE=	Specifies the maximum number of lines to be printed on any report page. The number may be from 28 to 99 inclusive. The default is 58.
MAXCORE= MAXC=	Specifies the maximum number of bytes of working storage available for various commands. The number maybe any value from 1000 to 120000, inclusive. The default is 2400 bytes.

41.04 IAMSMF - DEFAULT COMMAND

MAXPRINT= MAXP=	Specifies the maximum number of records the program will print. The number may be any value from 1 to 65536, inclusive. The default is 20 records.
MAXRECLNGTH= MAXR=	Specifies the largest SMF record that the program will process. The number may be any value from 16384 to 65536, inclusive. The default is 16384 bytes.
MAXSTACK= MAXS=	Specifies the maximum number of compressed SMF records which can be retained in storage. The number can be any value from 100 to 50000, inclusive. The default is 2500 records.
PRTLNGTH= PRTL=	Limits to 'nnnn' bytes the amount of data to be printed if 'ERRORPRINT' is indicated. The number may be any value from 32 to 65536, inclusive. The default is 32768 bytes.
RECSIZE(rrr)= RECS(rrr)=	Establishes the minimum length of the SMF record type 'rrr' as the value 'nnnn'. The default minimum record lengths for system generated SMF records are documented in IBM SPL: Systems Management Facilities (GC28-0706).
SELTERR= SELT=	Specifies that unsatisfied selection criteria will cause the condition code to be set to a non-zero value. YES - Produce a non-zero condition code when selection criteria is not met. NO - Ignore selection errors. The default is YES.
SORTCORE=	Specifies the amount of storage the program 'SORT' is to use if external sorting is required. The number may be any value from 10000 to 8000000 inclusive. The default is 100000.
SORTMSG=	Specifies the message option to be used by the program 'SORT' if external sorting is required. AC - All messages to the console AP - All messages to the printer (SYSOUT) CC - Critical messages to the console CP - Critical messages to the printer NO - No messages to be produced PC - Critical messages to both console and printer The default is CC.
SORTPFX=	Specifies the DDNAME prefix to be used by the program 'SORT' if external sorting is required. If the string specified is less than 4 characters, a dollar sign(\$) fill character will be used. The default is SORT.
SYSLIB=	Specifies the DDNAME to be used to update the SMF option module, IAMSMFOP. The default update DDNAME is SYSLIB.
TODDNAME= TODD=	Specifies the DDNAME of the output file for the SMF records created by the COPY function. The default output DDNAME is SYSUT2.

41.05 IAMSMF - END COMMAND

41.05 IAMSMF - END COMMAND

END The END command terminates IAMSMF processing. This command is intended for TSO users.
COMMAND

END

OPERANDS There are no operands for the END command.

41.06 IAMSMF - HELP COMMAND

41.06 IAMSMF - HELP COMMAND

HELP COMMAND The HELP command prints help by function or by command.

HELP ALL
,SUBCOMMAND(ccccc)

OPERANDS The following operands may be specified on the HELP command.

ALL Lists the entire contents of the help member, for all commands and field name descriptions.

SUBCOMMMAND Requests full help for the command 'cccc'.

SU *NOTE: 'JCL' and 'CLIST' can both be used as operands for the SUBCOMMAND parameter.*

41.07 IAMSMF - PRINT COMMAND

41.07 IAMSMF - PRINT COMMAND

PRINT COMMAND The PRINT command is used to print selected SMF records from either a history file (RECFM=VBS) or an active SMF data recording file.

PRINT

ALLRECORDS ,**MAXPRINT**=nnnn
,CHECKLENGTH ,**MAXRECLENGTH**=nnnn
,ERRORPRINT ,**PRTLENGTH**=nnn
,FROMDATE=yyddd ,**RECSIZE**(rrr)=nnnnn
,FROMDDNAME=cccc ,**RECTYPE**=nnn
,GROUPNAMES=cccc ,**TODATE**=yyddd
,JOBNAMES=cccc

OPERANDS The following operands may be specified on the PRINT command.

ALLRECORDS Specifies that the SMF record type does not participate in record selection.

ALL

NOTE: This operand conflicts with the operand 'RECTYPE'.

The default is deferred to the operand 'RECTYPE'.

CHECKLENGTH Specifies that SMF data records are to be validated against a table of minimum record lengths. If the user has modified the minimum record length or is executing in a non-compatible system, the correct record lengths may be specified by the 'RECSIZE' operand.

CHE

By default, the length of an SMF record is not validated.

ERRORPRINT Specifies that any SMF data record which causes an error during processing or fails length verification is to be printed.

ERR

The default is records in error are not printed.

FROMDATE= Specifies the lower date limit of the SMF records which are to be printed.

FROMDA=

The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.

FROMDDNAME= Specifies the DDNAME of the SMF file to be used as input to IAMSMF.

FROMDD=

The default input DDNAME is SYSMF.

41.07 IAMSMF - PRINT COMMAND

GROUPNAMES= G=	<p>Specifies that only records having a jobname which begin with the specified character string(s) will be printed. This operand specifies a partial jobname from 1 to 8 characters in length. Up to 50 job groups and/or names may be specified for a single command if entered as follows:</p> <p>GROUPNAMES=(c...c,....,c...c)</p> <p>The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the jobname will not participate in SMF record selection.</p>
JOBNAMES= J=	<p>Specifies that only records having a jobname which match the jobname(s) specified will be printed. This operand specifies a complete jobname from 1 to 8 characters in length. Up to 50 jobnames may be specified in a single command if entered as follows:</p> <p>JOBNAMES=(c...c,....,c...c)</p> <p>The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the jobname will not participate in SMF record selection.</p>
MAXPRINT= MAXP=	<p>Specifies the maximum number of records the program will print. The number may be any value from 1 to 65536, inclusive.</p> <p>The default is 20 records.</p>
MAXRECLENGTH= MAXR=	<p>Specifies the largest SMF record that the program will process. The number may be any value from 16384 to 65536, inclusive.</p> <p>The default is 16384 bytes.</p>
PRTLLENGTH= PRTL=	<p>Limits to 'nnnn' bytes the amount of data to be printed if 'ERRORPRINT' is indicated. The number may be any value from 32 to 65536, inclusive.</p> <p>The default is 32768 bytes.</p>
RECSIZE(rrr)= RECS(rrr)=	<p>Establishes the minimum length of the SMF record type 'rrr' as the value 'nnnn'.</p> <p>The default minimum record lengths for system generated SMF records are documented in IBM SPL: Systems Management Facilities (GC28-0706).</p>
RECTYPE= RECT=	<p>Identifies the specific record type(s) to be copied. Up to 50 record types may be specified for a single command if entered as follows:</p> <p>RECTYPE=(rrr,....,rrr)</p> <p>The SMF record types which will be copied by default (if 'ALLRECORDS' is not specified) are as follows:</p> <ul style="list-style-type: none">4 - Step termination5 - Job termination14 - NON-VSAM Dataset CLOSEd (input)15 - NON-VSAM Dataset CLOSEd (output/update)20 - Job initiation30 - Common Address Space Work Record34 - TSO session termination64 - VSAM Dataset CLOSEd <p><i>NOTE: This operand conflicts with the operand 'ALLRECORDS'.</i></p>
TODATE= TODA=	<p>Specifies the upper date limit of the SMF records which are to be printed.</p> <p>The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.</p>

41.08 IASMF - QUERY COMMAND

41.08 IASMF - QUERY COMMAND

QUERY COMMAND The QUERY command is used to report on data set utilization as requested by data set name or data set group name from a history file (RECFM=VBS), an active SMF data recording file, or from a sequential file produced by the COPY command of the program itself.

QUERY

CHECKLENGTH	,JOBNAMES=cccc
,DSGROUPS=cccc	,MAXCORE=nnnn
,DSNAMES=cccc	,MAXRECLENGTH=nnnn
,ERRORPRINT	,MAXSTACK=nnnn
,FROMDATE=yyddd	,PRTLENGTH=nnn
,FROMDDNAME=ddname	,RECSIZE(rrr)=nnnnn
,GROUPNAMES=cccc	,TODATE=yyddd

OPERANDS The following operands may be specified on the QUERY command.

CHECKLENGTH Specifies that SMF data records are to be validated against a table of minimum record lengths. If the user has modified the minimum record length or is executing in a non-compatible system, the correct record lengths may be specified by the 'RECSIZE' operand.

By default, the length of an SMF record is not validated.

DSGROUPS= Specifies that only records having a data set name which begin with the given character string(s) will be processed. This operand specifies a partial dataset name from 1 to 44 characters in length. Up to 50 data set groups may be specified for a single command if entered as follows:

DSG=
DSGROUPS=(c...c,...,c...c)

NOTE: If neither the 'DSGROUPS' nor 'DSNAMES' operand is specified, the 'QUERY' command will be marked in error.

DSNAMES= Specifies that only records having a data set name which match the dataset name(s) specified will be processed. This operand specifies a complete data set name from 1 to 44 characters in length. Up to 50 data set names may be specified for a single command if entered as follows:

DSN=
DSNAMES=(c...c,...,c...c)

NOTE: If neither the 'DSGROUPS' nor 'DSNAMES' operand is specified, the 'QUERY' command will be marked in error.

41.08 IAMSMF - QUERY COMMAND

ERRORPRINT ERR	<p>Specifies that any SMF data record which causes an error during processing or fails length verification is to be printed.</p> <p>By default, records in error are not printed.</p>
FROMDATE= FROMDA=	<p>Specifies the lower date limit of the SMF records which are to be selected.</p> <p>The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.</p>
FROMDDNAME= FROMDD=	<p>Specifies the DDNAME of the SMF file to be used as input to IAMSMF.</p> <p>The default input DDNAME is SYSMF.</p>
GROUPNAMES= G=	<p>Specifies that only records having a jobname which begin with the given character string(s) will be processed. This operand specifies a partial jobname from 1 to 8 characters in length. Up to 50 job groups and/or names may be specified for a single command if entered as follows:</p> <p>GROUPNAMES=(c...c,....c...c)</p> <p>The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the jobname will not participate in SMF record selection.</p>
JOBNAMES= J=	<p>Specifies that only records having a jobname which match the jobname(s) specified will be processed. This operand specifies a complete jobname from 1 to 8 characters in length. Up to 50 jobnames may be specified for a single command if entered as follows:</p> <p>JOBNAMES=(c...c,....c...c)</p> <p>The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the jobname will not participate in SMF record selection.</p>
MAXCORE= MAXC=	<p>Specifies the maximum number of bytes of working storage available for various commands. The number may be any value from 1000 to 120000, inclusive.</p> <p>The default is 2400 bytes.</p>
MAXRECLNGTH= MAXR=	<p>Specifies the largest SMF record that the program will process. The number may be any value from 16384 to 65536, inclusive.</p> <p>The default is 16384 bytes.</p>
MAXSTACK= MAXS=	<p>Specifies the maximum number of compressed SMF records which can be retained in storage. The number can be any value from 100 to 50000, inclusive.</p> <p>The default is 2500 records.</p>
PRTLNGTH= PRTL=	<p>Limits to 'nnnn' bytes the amount of data to be printed if 'ERRORPRINT' is indicated. The number may be any value from 32 to 65536, inclusive.</p> <p>The default is 32768 bytes.</p>
RECSIZE(rrr)= RECS(rrr)=	<p>Specifies the minimum length of the SMF records type 'rrr' is to be set to the value 'nnnn'.</p> <p>The default minimum record lengths for system generated SMF records are documented in IBM SPL:Systems Management Facilities (GC28-0706).</p>
TODATE= TODA=	<p>Specifies the upper date limit of the SMF records which are to be selected.</p> <p>The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.</p>

41.09 IAMSMF - REPORT COMMAND

41.09 IAMSMF - REPORT COMMAND

REPORT COMMAND The REPORT command is used to print job and/or TSO session related information as requested by job or job group name, with optional data set usage information, from a history file (RECFM=VBS), an active SMF recording file, or from a sequential file produced by the COPY command of the program itself.

REPORT

ALLDSNAMES	,MAXRECLENGTH=nnnn
,CHECKLENGTH	,MAXSTACK=nnnn
,ERRORPRINT	,NODSNAMES
,FROMDATE=yyddd	,PRTLENGTH=nnnn
,FROMDDNAME=ddname	,RECSIZE(rrr)=nnnn
,GROUPNAME=cccc	,TEMPDSNAMES
,JOBNAMES=cccc	,TODATE=yyddd

OPERANDS The following operands may be specified on the REPORT command.

ALLDSNAMES ALLD Specifies that the IAMSMF report will show all the occurrences of any referenced data set name. The EXCP count displayed will reflect each OPEN/CLOSE.

As a default, data set usage counts (EXCPs) are accumulated for each permanent data set, and the data set is only shown once. Temporary data sets are ignored.

NOTE: This operand conflicts with the operand 'NODSNAMES'.

CHECKLENGTH CHE Specifies that SMF data records are to be validated against a table of minimum record lengths. If the user has modified the minimum record length or is executing in a non-compatible system, the correct record lengths may be specified by the 'RECSIZE' operand.

The default is the length of an SMF record is not validated.

ERRORPRINT ERR Specifies that any SMF data record which causes an error during processing or fails length verification is to be printed.

The default is records in error are not printed.

FROMDATE= FROMDA= Specifies the lower date limit of the SMF records which are to be selected.

The default if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.

41.09 IAMSMF - REPORT COMMAND

FROMDDNAME= FROMDD=	Specifies the DDNAME of the SMF file to be used as input to IAMSMF. The default input DDNAME is SYSMF.
GROUPNAMES= G=	Specifies that only records having a jobname which begin with the specified character string(s) will be processed. This operand specifies a partial jobname from 1 to 8 characters in length. Up to 50 job groups and/or names may be specified for a single command if entered as follows: GROUPNAMES=(c...c,...,c...c) The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the jobname will not participate in SMF record selection.
JOBNAMES= J=	Specifies that only records having a jobname which match the jobname(s) specified will be processed. This operand specifies a complete jobname from 1 to 8 characters in length. Up to 50 jobnames may be specified for a single command if entered as follows: JOBNAMES=(c...c,...,c...c) The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the jobname will not participate in SMF record selection.
MAXSTACK= MAXS=	Specifies the maximum number of compressed SMF records which can be retained in storage. The number can be any value from 100 to 50000, inclusive. The default is 2500 records.
MAXRECLENGTH= MAXREC=	Specifies the largest SMF record that the program will process. The number may be any value from 16384 to 65536, inclusive. The default is 16384 bytes.
NODSNAMES NOD	Specifies that the report produced will not show data set oriented information. By default, data set usage information is displayed only once for each permanent data set with accumulated statistics for multiple OPEN/CLOSE. Temporary data sets are ignored. <i>NOTE: This operand conflicts with the operands 'ALLDSNAMES' and 'TEMPDSNAMES'.</i>
PRTLENGTH= PRTL=	Limits to 'nnnn' bytes the amount of data to be printed if 'ERRORPRINT' is indicated. The number may be any value from 32 to 65536, inclusive. The default is 32768 bytes.
RECSIZE(rrr)= RECS(rrr)=	Specifies the minimum length of the SMF record type 'rrr' is to be set to the value 'nnnn'. The default minimum record lengths for system generated SMF records are documented in IBM SPL:Systems Management Facilities (GC28-0706).
TEMPDSNAMES TEM	Specifies that reports produced will contain information on temporary as well as permanent datasets. By default, temporary data sets are ignored. <i>NOTE: This operand conflicts with the operand 'NODSNAMES'.</i>
TODATE= TODA=	Specifies the upper date limit of the SMF records which are to be selected. The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.

41.10 IAMSMF - IAMINFO COMMAND

41.10 IAMSMF - IAMINFO COMMAND

IAMINFO COMMAND The IAMINFO command is used to produce IAMINFO reports from IAM generated SMF records.

IAMINFO

ATTRIBUTE=cccccc	,MAXSTACK=nnnnn
,BLOCKS=nnnnnn	,MERGE=cccccc
,CYLS=nnnnn	,NUMCMDS=nnnnn
,DSGROUPS=cccccc	,NUMRECS=nnnnn
,DSNAMES=cccccc	,OFULL=nn
,DYNRET=nnnnn	,ORECS=nnnnn
,EXCP=nnnnn	,PEBLKS=nnnnn
,FROMDATE=yyddd	,PEFULL=nn
,FROMDDNAME=ddname	,RECTYPE=nnn
,GROUPNAMES=cccccc	,TODATE=yyddd
,JOBNAMES=cccc	,TRACKS=nnnnn
,LRECL=nnnnn	

41.10 IAMSMF - IAMINFO COMMAND

OPERANDS The following operands may be specified on the IAMINFO command.

ATTRIBUTE= ATTR=	Identifies which attributes an IAM file must have to participate in record selection.
DATACOMP	- The IAM file contains compressed data.
KEYCOMP	- The IAM file contains compressed keys.
INCPEBLK	- IAM found that a record which could have gone into PRIME EXTENSION had to be placed into INDEPENDENT OVERFLOW because the PRIME EXTENSION was full. Redefining this file with a larger PRIME EXTENSION may improve performance. An IAM371 message was issued if an IAMINFO DD was present during the job's execution.
MOREBUFFER	- IAM found that additional buffers could have been used to improve performance, however the MAXBUF or BUFNO option prevented IAM from acquiring more buffers. An IAM368 message was issued if an IAMINFO DD was present during the job's execution.
NOCORE	- There was not enough storage available in the region for IAM to acquire additional buffers. An IAM367 message was issued if an IAMINFO DD was present during the job's execution. This operand supports the following logical operators: =, ≠ .

By default, the data set attributes do not participate in the selection criteria.

BLOCKS= BL=	Specifies the size, in blocks, of the files to be selected. By default, the size of the IAM data set does not participate in the selection criteria. This operand supports the following logical operators: =, ≠, >, >=, <, <= .
------------------------	---

CYLS=	Specifies the size, in cylinders, of the files to be selected. By default, the size of the IAM data set does not participate in the selection criteria. This operand supports the following logical operators: =, ≠, >, >=, <, <= .
--------------	---

DSGROUPS= DSG=	Specifies that only records having a data set name which begin with the given character string(s) will be processed. This operand specifies a partial dataset name from 1 to 44 characters in length. Up to 50 data set groups may be specified for a single command if entered as follows:
---------------------------	---

DSGROUPS=(c...c,...,c...c)

By default, the name of the IAM data set does not participate in the selection criteria.

DSNAMES= DSN=	Specifies that only records having a data set name which match the dataset name(s) specified will be processed. This operand specifies a complete data set name from 1 to 44 characters in length. Up to 50 data set names may be specified for a single command if entered as follows:
--------------------------	---

DSNAMES=(c...c,...,c...c)

By default, the name of the IAM data set does not participate in the selection criteria.

41.10 IAMSMF - IAMINFO COMMAND

DYNRET=	Establishes the limit of record retrievals from IAM's Dynamic Table.
DYN=	<p>By default, the number of dynamic retrievals does not participate in the selection criteria.</p> <p>This operand supports the following logical operators: =, !=, >, >=, <, <= .</p>
EXCP=	<p>Establishes the limit of EXCPs (physical read and writes) to the IAM file.</p> <p>By default, the number of EXCPs does not participate in the selection criteria.</p> <p>This operand supports the following logical operators: =, !=, >, >=, <, <= .</p>
FROMDATE=	Specifies the lower date limit of the SMF records which are to be copied.
FD=	The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.
FROMDDNAME=	Specifies the DDNAME of the SMF file to be used as input to IAMSMF.
FROMDD=	The default input DDNAME is SYSMF.
GROUPNAMES=	Specifies that only records having a job name which begin with the specified character string(s) will be selected. This operand specifies a partial job name from 1 to 8 characters in length. Up to 50 job groups and/or names may be specified for a single command if entered as follows:
G=	<p>GROUPNAMES=(c...c,...,c...c)</p> <p>The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the job name will not participate in SMF record selection.</p>
JOBNAMES=	Specifies that only records having a job name which match the jobname(s) specified will be selected. This operand specifies a complete job name from 1 to 8 characters in length. Up to 50 job names may be specified in a single command if entered as follows:
J=	<p>JOBNAMES=(c...c,...,c...c)</p> <p>The default, if neither the 'GROUPNAMES' nor 'JOBNAMES' operand is specified, is that the job name will not participate in SMF record selection.</p>
LRECL=	<p>Establishes the limit for the length of records (maximum length for variable) within an IAM file.</p> <p>By default, the record length does not participate in the selection criteria.</p> <p>This operand supports the following logical operators: =, !=, >, >=, <, <= .</p>
MAXSTACK=	Specifies the maximum number of compressed SMF records which can be retained in storage. The number can be any value from 100 to 50000, inclusive.
MAXS=	The default is 2500 records.

41.10 IAMSMF - IAMINFO COMMAND

- MERGE=** Identifies how IAMSMF will process multiple SMF records created for the same dataset by the same job.
- JOB** - All SMF records referencing the same IAM data set within a job will be merged into a single report.
- NO** - No SMF records are to be merged. An IAMINFO report will be generated for each SMF record found.
- STEP** - All SMF records referencing the same IAM data set within a job step (i.e. multiple OPEN/CLOSE) will be merged into a single report.
- The default is STEP.
- NUMCMDS=**
NUMC= Establishes the limit of the total number of commands issued against the IAM file. When used in conjunction with the TYPECMD operand, the selection will be limited to those files having a command count of that type.
- By default, the number of commands issued against the IAM file does not participate in the selection criteria.
- This operand supports the following logical operators:
=, \neq , >, >=, <, <= .
- OFULL=nn** Establishes the percent of overflow used limit within an IAM file.
- By default, the percent of overflow used does not participate in the selection criteria.
- This operand supports the following logical operators:
=, \neq , >, >=, <, <= .
- ORECS=nnnnn** Establishes the number of overflow records limit within an IAM file.
- By default, the number of overflow records does not participate in the selection criteria.
- This operand supports the following logical operators:
=, \neq , >, >=, <, <= .
- PEBLKS=nnnnn** Establishes the number of prime extension blocks limit within an IAM file.
- By default, the number of prime extension blocks does not participate in the selection criteria.
- This operand supports the following logical operators:
=, \neq , >, >=, <, <= .
- PEFULL=nn** Establishes the percent of prime extension blocks used limit within an IAM file.
- By default, the percent of prime extension blocks used does not participate in the selection criteria.
- This operand supports the following logical operators:
=, \neq , >, >=, <, <= .
- RECTYPE=**
RECT= Identifies the record type of the IAM generated SMF record.
- The default record type is that which is defined in the IAM option table. Refer to [section 91](#) for further documentation on the IAM option table.
- TODATE=**
TODA= Specifies the upper date limit of the SMF records which are to be copied.
- The default, if the 'FROMDATE' and/or 'TODATE' operands are not specified, is that the date of the SMF record will not participate in the selection criteria.

41.10 IAMSMF - IAMINFO COMMAND

TRACKS= Specifies the size, in tracks, of the files to be selected.
TR= By default, the size of the IAM data set does not participate in the selection criteria.

This operand supports the following logical operators:
=, !=, >, >=, <, <= .

TYPECMD= Identifies the type of command(s) which must have been issued against the IAM file. If NUMCMD is also specified, the SMF record will be selected if that number of commands of that type was issued. If NUMCMD is not specified, the SMF record will be selected if ANY command of that type was issued. Up to 16 command types may be specified if entered as follows:
TYPEC=

TYPECMD=(c,c,...,c)

The following command types (which correspond to those generated by IAMINFO) are supported:

R - Read	W - Write
F - Getnext	A - Add
D - Delete	N - Read or Getnext
G - Get	S - Start Key Equal
K - Start KC or Point	P - Get previous
U - Read for Update	E - Stastics
C - Close	O - Open
I - Information	X - Flush

By default, the type of commands issued against the IAM file does not participate in the selection criteria.

41.11 IAMSMF - EXAMPLES

41.11 IAMSMF - EXAMPLES

REPORT EXAMPLE 1. In this example, IAMSMF will scan SMF history data for all occurrences of jobs which start with the name 'R1D3W3'. A detail report will be generated for each DDNAME used in the jobs.

```
//IAMSMF      EXEC  PGM=IAMSMF
//STEPLIB    DD   DISP=SHR,DSN=iam.library  <==USER CHANGE
//SYSPRINT   DD   SYSOUT=*
//SYSUDUMP   DD   SYSOUT=*
//SYSMF      DD   DISP=SHR,DSN=SMF.HISTORY.MONTH
//SYSIN      DD   *
              REPORT GROUPNAMES=R1D3W3,ALLDSNAMES
/*
```

IAM400 SMF REPORT/DATA EXTRACT PROGRAM -- IAMSMF VER 6.x -- INNOVATION DATA PROCESSING					
IAM303 CARD IMAGE - * REPORT GROUPNAME=R1D3W3					*
IAM491 SMF REPORT FUNCTION STARTED - 14.54.54					
JOB NAME	R1D3W31	DATE	88.218		
STEP NAME	DEFINE	PGM	IDCAMS	REGION REQ	.8192
START	19.19.47	STOP	19.19.50	WALL	00.00.03.37
				TAT	00.00.03.09
				CPU	00.00.00.85
				TNA	00.00.00.28
				TRT	00.00.03.09
				COMP CODE	C-0000
SUMMARY OF I/O ACTIVITY.....					
		PAGE IN	0	PAGE OUT	0
		SWAP IN	0	SWAP OUT	0
STEP NAME	EXTRACT	PGM	PR13W12	REGION REQ	.8192
START	19.18.50	STOP	03.23.54	WALL	08.05.03.26
				TAT	07.56.35.97
				CPU	00.21.49.56
				TNA	00.08.27.29
				TRT	06.23.36.53
DD	INFILE	DSNAME	R1D3.MASTER.DATA	AM USE	UPDATE
	INFILE		R1D3.MASTER.INDEX	AM	UPDATE
	OUTFILE		R1D3.EX415.DATA	AM	UPDATE
	OUTFILE		R1D3.EX415.INDEX	AM	UPDATE
	STEPLIB		R1D3.PROD.LOAD	PO	INPUT
				DISK	29841
				DISK	17654
				DISK	31112
				DISK	20210
				DISK	7
				TRT	06.23.36.53
SUMMARY OF I/O ACTIVITY.....					
		PAGE IN	590	PAGE OUT	500
		SWAP IN	2273	SWAP OUT	2273
				SWAP COUNT	1944
				SRVU	4718642
JOB NAME	R1D3W32	DATE	88.218		
STEP NAME	REPORT	PGM	PR13W22	REGION REQ	.8192
START	03.23.54	STOP	03.53.55	WALL	00.30.04.96
				TAT	00.27.04.18
				CPU	00.01.32.81
				TNA	00.03.00.78
				TRT	00.25.04.18
DD	R22FILE	DSNAME	R1D3.EX415.DATA	AM USE	UPDATE
	R22FILE		R1D3.EX415.INDEX	AM	UPDATE
	STEPLIB		R1D3.PROD.LOAD	PO	INPUT
				DISK	21935
				DISK	11829
				DISK	12
SUMMARY OF I/O ACTIVITY.....					
		PAGE IN	93	PAGE OUT	64
		SWAP IN	170	SWAP OUT	170
				SWAP COUNT	134
				SRVU	51919

41.11 IAMSMF - EXAMPLES

QUERY EXAMPLE 2. In this example, IAMSMF will scan the active SMF data set for all jobs which reference the VSAM cluster 'ACCOUNT.FILE'.

```
//IAMSMF      EXEC  PGM=IAMSMF
//STEPLIB    DD   DISP=SHR,DSN=i am.libr ary   <==USER CHANGE
//SYSPRINT   DD   SYSOUT=*
//SYSUDUMP   DD   SYSOUT=*
//SYSMF      DD   DISP=SHR,DSN=SYS1.MANx
//SYSIN      DD   *
              QUERY  DSNNAME=ACCOUNT.FILE
/*
```

IAM303 CARD IMAGE - *		QUERY DSNNAME=ACCOUNT.FILE	
IAM491 DATA SET QUERY		FUNCTION STARTED - 11.03.05	
DSN---ACCOUNT.FILE		DATE-----88.216	TIME-----06.44.42
JOB-----R1D3W3A	DDNAME----ACCNT1	DSORG-----AM	USE-----OUTPUT
EXCP COUNTS----	DISK-----132525		
DSN---ACCOUNT.FILE		DATE-----88.216	TIME-----22.34.37
JOB-----R1DA31	DDNAME----FINAF1	DSORG-----AM	USE-----INPUT
EXCP COUNTS----	DISK-----41321		
DSN---ACCOUNT.FILE		DATE-----88.216	TIME-----23.51.41
JOB-----R1DA45	DDNAME----ACTFILE	DSORG-----AM	USE-----INPUT
EXCP COUNTS----	DISK-----23442		
DSN---ACCOUNT.FILE		DATE-----88.217	TIME-----02.11.59
JOB-----R1DC01	DDNAME----APPLY05	DSORG-----AM	USE-----OUTPUT
EXCP COUNTS----	DISK-----11315		
DSN---ACCOUNT.FILE		DATE-----88.217	TIME-----06.51.03
JOB-----MAP	DDNAME----ACCNT1	DSORG-----AM	USE-----OUTPUT
EXCP COUNTS----	DISK-----140015		

41.11 IAMSMF - EXAMPLES

COPY EXAMPLE 3. IAMSMF can be used to determine all of the jobs which DEFINEd selected VSAM cluster(s). The input SMF data can be from an active system SMF file, from an SMF history or from a sequential file produced by the COPY command of the program itself.

NOTE: Record types 61 for ICF catalog users, and 63 for VSAM catalog users must be collected. These records are written by catalog management after a VSAM cluster or component has been DEFINEd.

In this example, IAMSMF will be used to scan SMF history data on tape. All jobs DEFINEing VSAM clusters whose name starts with the high level index of 'R1D3W3' will be reported. The first command (COPY) extracts the type 61 and 63 records from the SMF file. The QUERY command will then report on the jobs which did the DEFINE.

```
//IAMSMF      EXEC  PGM=IAMSMF
//STEPLIB    DD    DISP=SHR,DSN=iam.library  <==USER CHANGE
//SYSPRINT   DD    SYSOUT=*
//SYSUDUMP   DD    SYSOUT=*
//SYSMF     DD    DISP=SHR,DSN=SMF.HISTORY.FILE
//SYSUT2     DD    UNIT=SYSDA,SPACE=(CYL,(5,1))
//SYSIN      DD    *
              COPY  RECTYPE=(61,63)
              QUERY FROMDD=SYSUT2,DSG=R1D3W3
/*
```

```
IAM303 CARD IMAGE - *          COPY  RECTYPE=(61,63)
IAM491 RECORD COPY           FUNCTION STARTED - 11.02.32

IAM601 SMF RECORDS -- READ.....31355  USED.....1240 DROPPED.....0
IAM492 RECORD COPY           FUNCTION ENDED   - 11.02.59 - CONDITION CODE 000

IAM303 CARD IMAGE - *          QUERY FROMDD=SYSUT2,DSG=R1D3W3
IAM491 DATA SET QUERY       FUNCTION STARTED - 11.02.59

DSN---R1D3W3.TRANS.LOG1.D88217  DATE-----88.217  TIME-----06.32.39
JOB-----R1D3W3A  DDNAME--LOGFILE  DSORG-----AM  USE-----DEFINE

DSN---R1D3W3.TRANS.LOG2.D88217  DATE-----88.217  TIME-----06.35.02
JOB-----R1D3W3A  DDNAME--LOGFILE  DSORG-----AM  USE-----DEFINE

DSN---R1D3W3.EAST.MASTER  DATE-----88.219          TIME-----01.59.02
JOB-----R3EDFWK  DDNAME--EMASTER  DSORG-----AM  USE-----DEFINE

DSN---R1D3W3.REPORT.W8832  DATE-----88.218          TIME-----22.28.52
JOB-----R3RPWK  DDNAME--**ABSENT  DSORG-----AM  USE-----DEFINE

DSN---R1D3W3.WEST.MASTER  DATE-----88.219          TIME-----01.11.02
JOB-----R3WDFWK  DDNAME--WMASTER  DSORG-----AM  USE-----DEFINE
```

41.11 IAMSMF - EXAMPLES

EXAMPLE 4 In this example, IAMSMF will be used to generate IAMINFO reports for IAM files which are beginning to run out of space in Independent Overflow. Only those files which have more than 1000 total records in independent overflow of which 90% are used will be selected.

```
// IAMSMF      EXEC   PGM=IAMSMF
// STEPLIB    DD     DISP=SHR,DSN=i am . l i b r a r y           <==USER CHANGE
// SYSPRINT   DD     SYSOUT=*
// SYSUDUMP   DD     SYSOUT=*
// SYSMF      DD     DISP=SHR,DSN=SMF.HISTORY.FILE           <==USER CHANGE
// SYSIN      DD     *
                IAMINFO  ORECS<1000,OFULL<=90
/*
```

EXAMPLE 5 In this example, IAMSMF will be used to generate IAMINFO reports for IAM files which either didn't have enough space in the region to acquire more buffers for processing or was stopped from getting more buffers by the MAXBUF/BUFNO operands.

```
// IAMSMF      EXEC   PGM=IAMSMF
// STEPLIB    DD     DISP=SHR,DSN=i am . l i b r a r y           <==USER CHANGE
// SYSPRINT   DD     SYSOUT=*
// SYSUDUMP   DD     SYSOUT=*
// SYSMF      DD     DISP=SHR,DSN=SMF.HISTORY.FILE           <==USER CHANGE
// SYSIN      DD     *
                IAMINFO  ATTR=MOREBUFFER,NOCORE
/*
```

EXAMPLE 6 In this example, IAMSMF will first be used to copy all IAM generated SMF records to a temporary data set. This file will then be used to find those IAM files which have had any commands flushed (for RPL errors etc.) when accessed. Also, IAMSMF will report on those files which had more than 1000 random reads issued against them and are not taking advantage of IAM's Dynamic Tabling (i.e. the number of dynamic retrievals was zero) .

```
// IAMSMF      EXEC   PGM=IAMSMF
// STEPLIB    DD     DISP=SHR,DSN=i am . l i b r a r y           <==USER CHANGE
// SYSPRINT   DD     SYSOUT=*
// SYSUDUMP   DD     SYSOUT=*
// SYSMF      DD     DISP=SHR,DSN=SMF.HISTORY.FILE           <==USER CHANGE
// SYSUT2     DD     UNIT=SYSDA,SPACE=(CYL,(5,1))
// SYSIN      DD     *
                COPY      IAMRECORDS
                IAMINFO  TYPECMD=X,FROMDD=SYSUT2
                IAMINFO  TYPECMD=R,NUMCMD<1000,DYNRET=0,
                        FROMDD=SYSUT2
/*
```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

SECTION 42: IAM SPACE SAVINGS ANALYSIS

42.01 IAMSIMVS - IAM'S SPACE SAVINGS ANALYSIS PROGRAM

42.01 IAMSIMVS - IAM'S SPACE SAVINGS ANALYSIS PROGRAM

OVERVIEW VSAM data sets usually account for a large portion of the DASD space allocated within any shop. IAM's advanced file structure takes 20% to 40% less disk space than an equivalent VSAM data set. Using IAM's proprietary compression techniques an additional 20% to 50% savings can be realized. IAMSIMVS projects the DASD space savings and resultant cost reductions a VSAM to IAM conversion can provide for your installation. IAMSIMVS's reports give estimated savings for individual data sets and then summarize these savings by device type for the installation as a whole. IAMSIMVS can also be used to help plan a conversion to IAM by identifying data sets whose conversion would result in the greatest savings.

FUNCTIONAL SUMMARY IAMSIMVS selectively reads the records contained in a VSAM cluster and simulates the conversion of that cluster to an IAM data set. The simulation results in two sets of figures. One showing the results if the file were loaded in IAM's standard (uncompressed) format and the other the results if the file were loaded in IAM's compressed file format. The simulated IAM file is assumed to have been DEFINEd using the same IDCAMS parameters (SPACE, RECORDSIZE, FREESPACE, etc.) that would have established the VSAM cluster being analyzed.

DATA COMPRESSION IAM's Data Compression techniques result in a 20% to 50% reduction in the size of the individual records contained in a file. Compression however, is not accomplished at the expense of increasing the CPU time it would have taken to process the file with VSAM normally.

File processing with IAM's Data Compression takes considerably less CPU time than other vendor's VSAM compression packages. In fact, IAM's CPU time is so much less than VSAM's, IAM with data compression typically takes less CPU time than normal VSAM processing without data compression. Compare this to the fact that other data compression packages must add their CPU time to compress and decompress on top of VSAM's processing time.

IAM can automatically compress data records based on a qualifying default file size, or individually on request. IAM attempts to compress the data following the key in each record. The record must contain at least 10 bytes of data following the key to qualify for compression. If a compressed record is larger than the original, IAM leaves the record uncompressed.

DATA SAMPLING The most accurate way to simulate a conversion is of course to read and process every record in a file. Reading every record in a large file however can be relatively time consuming. IAMSIMVS attempts to achieve a balance between exacting accuracy and reasonable run time by using only a sample of the records in a VSAM cluster. IAMSIMVS by default, limits its selection of the control areas it will read records from to a percentage of the control areas in the cluster. This technique limits the records actually read to a relatively small percentage of the total records the file actually contains. IAMSIMVS in this way is able to obtain both a good distribution in its sample of records and maintain an optimum level of performance. The sample rate IAMSIMVS will use is based upon the allocated size of the VSAM cluster whose conversion is to be simulated and may be anywhere from 10% to 100% of the records in the VSAM cluster, as shown in the following table.

<u>Cylinders</u>	<u>Sampling Rate</u>
1 - 9	100%
10 - 49	50%
50 - 99	20%
100 - nnnn	10%

In some unique circumstances, sampling may not be appropriate. For example, if a file was created with 100,000 records and all but the last 5,000 have since been deleted, sampling 10% of the file's **Control Areas** may not give the most accurate picture of the data that this file usually contains (i.e. most control areas are presently empty).

42.01 IAMSIMVS - IAM'S SPACE SAVINGS ANALYSIS PROGRAM

AUTOMATIC RELEASE IAMSIMVS will show you exactly how much of their allocated space your VSAM clusters are using for the data they contain. The remaining space, usually the result of over allocation, is wasted and VSAM has no facility to release it. IAM releases unused space when a file is initially loaded, automatically when the file is DEFINEd with secondary allocation.

SIZE ESTIMATES Care has been taken to give you the most accurate estimate possible. However, IAMSIMVS estimates on the size of a converted IAM file are still only estimates and may even vary by as much as + or – five percent from a true IAM file for the following reasons:

- The VSAM catalog record contains the results of a DEFINE not the exact parameters used for the DEFINE.
- The VSAM cluster may contain records which vary greatly in size.

IAMSIMVS size estimates could be off by more than five percent from a true IAM file for the following reasons:

The VSAM file varies greatly in size from day to day. When the file is converted to IAM it may contain a different number of records.

The VSAM file contains a large number of deleted records when compared to the number of records initially loaded or inserted. Since these records are no longer on the file, IAM cannot determine their average record sizes. In addition, the maximum number of records the file may have contained cannot be determined.

The VSAM file uses a small number of tracks as compared to the allocated space. For example, the VSAM files was allocated with 1,000 tracks but is using only 20 tracks. IAM does not know if the over-allocation was accidental or in anticipation of file growth.

REPORT FORMATS IAMSIMVS's Data Set Report displays for each cluster reported on:

- The number of tracks allocated for the VSAM cluster.
- The number of tracks actually used by the VSAM cluster.
- An estimated number of tracks an equivalent IAM file will occupy.
- An estimated number of tracks a compressed IAM file will occupy.
- An estimated percent of the IAM savings over the VSAM used space. (or optionally the allocated space).
- The average and largest record size encountered sampling the file.
- Various VSAM file attributes (total records, key length, etc.)

IAMSIMVS also produces a Summary Report by DASD device type on:

- The total number of datasets processed.
- The total number of VSAM tracks allocated.
- The total number of VSAM tracks in use.
- VSAM's percent of tracks used as compared to allocated.
- The total savings, in tracks, a conversion to IAM can return.
- The total savings, in tracks, IAM's Data Compression can return.
- The dollar value, of the savings which can be realized by converting to IAM (with and without Data Compression).

ADDITIONAL CONSIDERATIONS IAMSIMVS must open and read records from the specified VSAM clusters. Under a security package (RACF, Top Secret, ACF2, etc.) you must have proper access authorization to select a cluster for conversion simulation. If the IAMSIMVS job does not have authorization to read the specified VSAM clusters it will fail.

IAMSIMVS does not need or use IAM's VSAM Interface (VIF).

COMMAND SUMMARY IAMSIMVS has the following commands:

REPORT - Set processing defaults

SELECT - Identify those data sets which are to be used for simulation.

SECTION 42: IAM SPACE SAVINGS ANALYSIS

42.02 JCL REQUIREMENTS FOR IAMSIMVS

42.02 JCL REQUIREMENTS FOR IAMSIMVS

JCL REQUIREMENTS The JCL statements required to execute IAMSIMVS are as follows:

EXECUTE STATEMENT Specifies the name of the IAM simulation program - IAMSIMVS. For optimum VSAM performance, the region size for IAMSIMVS should be at least 8192K.

STEPLIB or JOBLIB DD It is recommended that the IAM program load library be in the system LINKLIST. If it is not, you must include a STEPLIB/JOBLIB DD statement specifying the library in which IAMSIMVS program load modules reside.

May be omitted if these modules are in a system connected library (i.e. SYS1.LINKLIB).

SYSUDUMP DD STATEMENT Specifies the ABEND data set used if major errors are detected. Usually a SYSOUT dataset.

SYSPRINT DD STATEMENT Specifies where the IAMSIMVS control statements, messages, and reports are to be printed. Usually a SYSOUT data set.

SYSIN DD STATEMENT Specifies the input control statement data set. Usually a DD * data set.

NOTE: *The IAMSMFVS VSAM Usage Analysis program can optionally generate IAMSIMVS SELECT command statements corresponding to the VSAM cluster's named in its Cluster Size Report. (See Section 40.02, JCL Requirements - optional SYSPUNCH DD Statement). To have a Space Savings Report produced by IAMSIMVS using these generated SELECT statements simply assign the IAMSIMVS SYSIN DD statement to the data set that was written to by the IAMSMFVS SYSPUNCH DD Statement.*

42.03 IAMSIMVS - REPORT COMMAND

42.03 IAMSIMVS - REPORT COMMAND

REPORT COMMAND The IAMSIMVS REPORT command is used to override processing defaults.

```
REPORT          COSTPERMB=nn          ,MAXDSN=nnnn
                ,COST3380=nn          ,TYPE=cccccccc
                ,COST3390=nn          ,$SIGN=xxxx
                ,ESTIMATE=cccccc
```

OPERANDS The following operands may be specified on the REPORT command.

COSTPERMB= Specifies the cost per megabyte for a 3380. In the Summary Report, IAMSIMVS will generate an estimate of the cost savings which could result if the selected VSAM clusters were converted to IAM.

The default is \$17 per megabyte. This value was derived from the average cost for all 3380 models available from both IBM and the third party market, weighted to reflect controller and operating costs.

COST3390= Specifies the cost per megabyte for a 3390. In the Summary Report, IAMSIMVS will generate an estimate of the cost savings which could result if the selected VSAM clusters were converted to IAM.

The default is \$15 per megabyte. This value was derived from the average cost for all 3390 models available from both IBM and the third party market, weighted to reflect controller and operating costs.

ESTIMATE= Specifies the base from which IAM is to determine its space savings estimates.

EST=

ALLOC - Savings estimates are to be based upon the number of tracks allocated for the VSAM cluster. IAM automatically releases any unused space within the allocation.

USED - Savings estimates are to be based only upon the number of tracks actually used by the data in the VSAM cluster.

The default is USED.

MAXDSN=

MAXD=

Specifies the maximum number of data sets IAMSIMVS is prepared to simulate conversions of in one report.

NOTE: If the MAXDSN value is exceeded, IAM will split the datasets among multiple reports.

The default is 1000 datasets.

TYPE=

Identifies the format of IAM file the simulated conversion is to be for. The following options are supported:

IAM - Simulate conversion to standard format

COMPRESS - Simulate conversion compressed format.

BOTH - Simulate conversion to both IAM standard format and IAM compressed formats.

The default is BOTH.

\$SIGN=

Specifies the unit of currency to be used when displaying the value of the disk space saved in IAMSIMVS's summary report. Specify 1 to 4 characters. If you change the currency sign, you should also adjust the cost per MB value to reflect the currency change.

The default is '\$'.

42.04 IAMSIMVS - SELECT COMMAND

42.04 IAMSIMVS - SELECT COMMAND

SELECT COMMAND The IAMSIMVS SELECT command is used to identify those files for which a conversion simulation is to be performed. Additionally the SELECT command can be used to specify IAMSIMVS's sample rate and IAM file definition options. One or more SELECT statements must be specified.

```
SELECT      DSNAMES=dsname                ,OVERFLOW=nnnnnn  
           ,DSTYPE=KSDS]ESDS]ALL        ,PE=nnnn  
           ,BLOCKSIZE=nnnnn            ,SAMPLE=nnn  
           ,BUFND=nnn                  ,VARIABLE  
           ,INTEGRATED=nn
```

OPERANDS The following operands may be specified on the SELECT command.

DSNAMES= Identifies those data sets for which a conversion to IAM is to be simulated. Up to 50 data set names may be specified on each SELECT if entered as follows:

DSN=

DSN=(dsn,dsn,....,dsn)

There is no default value for DSNAMES and it must be specified.

NOTE: Although only 50 dataset names may be identified on each SELECT, all entries are tabled and will appear on the same report (up to the MAXDSN=value).

DSTYPE= Specifies the type of VSAM files to be processed.

KSDS - Process VSAM KSDS files only.

ESDS - Process VSAM ESDS files only.

ALL - Process KSDS and ESDS files.

The default is ALL.

BUFND= Sets the number of VSAM data buffers to be used when accessing the specified data set.

To obtain optimum performance, a default value will be set to a number large enough for VSAM to read an entire control area at a time.

BLOCKSIZE= Establishes the blocksize or blocking factor to be used when simulating conversion to IAM.

B=

By default, IAMSIMVS will base its calculations on the VSAM cluster Control Interval (CI) size. IAM rounds the CI size value up to develop a blocksize that will fit most efficiently on the track a multiple number of times. A minimum of quarter track blocking (i.e. four blocks on a track) is used. See Section 10 of the IAM manual for further details.

42.04 IAMSIMVS - SELECT COMMAND

INTEGRATED= Establishes the amount of free space to be simulated in each block to accommodate record inserts. This keyword is comparable to the IDCAMS FREESPACE(CI% ...) option.

I=

The default is to use the VSAM Cluster's Control Interval Percentage Free value.

OVERFLOW= Specifies the simulated number of blocks to be set aside for overflow records.

O=

The default overflow value is established by taking the VSAM Cluster's Control Area Percentage Free value (FREESPACE (... CA%)) and multiplying it by the estimated number of records in the cluster's primary allocation. See Section 10 of the IAM manual for further details.

PE= Establishes the simulated number of blocks IAM is to set aside for adds to the end of the file. The default is 3 blocks.

SAMPLE= Specifies the percentage of records to be read from the VSAM cluster when simulating. Any value from 10 to 100 may be used.

The default sampling rates are based upon the size of the VSAM cluster as follows:

Cylinders	Sampling rate
1 - 9	100%
10 - 49	50%
50 - 99	20%
100 - nnnn	10%

VARIABLE Identifies the file as containing variable length records.

VAR

Under normal processing, a file is assumed to have variable length records if the average and maximum record lengths (RECORDSIZE) specified for the VSAM file are not equal.

If message IAM318 is presented during processing the file was assumed to contain fixed length records, but the record lengths were not all equal. Should this message appear, rerun the simulation specifying the VARIABLE keyword. The values established by IAMSIMVS (overflow, blocking, etc.) will be different for a file containing fixed length records and one containing variable length records.

NOTE: If you decide to convert this file to an IAM file, the average record length in the IDCAMS DEFINE RECORDSIZE parameter must be changed to a value which is one less than maximum record length. If this is not done the load of the IAM file will fail.

42.05 IAMSIMVS - SAMPLE REPORTS

42.05 IAMSIMVS - SAMPLE REPORTS

IAMSIMVS supplies the user with two basic reports;
 A. The Conversion Simulation Cluster/Data Set Report
 B. The Savings Summary Reports.

A. The Conversion Simulation Cluster/Data Set Report

IAM400 VSAM ALLOCATION ANALYSIS - IAMSIMVS VER 6.0 - INNOVATION DATA PROCESSING DATE-yy.ddd PAGE nn															
IAM VS. VSAM ALLOCATION SIMULATION															
DATA SET NAME	VSAM ALLOC	VSAM TRKS USED	IAM TRKS STD	% COMP	SAVINGS STD	% COMP	TOTAL RECORDS	AVERAGE LARGEST	MAX LRECL	KEYLN RKP	CISIZ BLKSZ	FLAGS	C1% CA%	CISPL CASPL	
SMALLER.VSAM.DATASET	451	346	270	90	22	74	40496	242	250	20	4096	2	10	0	
							SAMPLE=50%	250			0	4096	IMBED	10	0
MID.SIZED.VSAM.KSDS.DATASET	1126	1111	780	510	30	54	46118	502	679	15	4096		10	0	
							SAMPLE=20%	670			1	4096		10	0
LARGER.VSAM.DATASET	2964	2949	2220	915	25	69	157405	600	600	20	4096		0	0	
							SAMPLE=10%	600			1	4096	IMBED	0	0
VSAM.KSDS.DATASET.WITH.SPANNED.RECS	466	176	90	60	49	66	1001	1090	2048	8	1024	4,5	20	0	
								2048			4	1024	IMBED	15	0

FLAGS: 1) UNABLE TO ESTIMATE STANDARD IAM FILE ALLOCATION.
 2) FILE WAS ASSUMED TO BE FIXED BUT WAS FOUND TO CONTAIN VARIABLE LENGTH RECORDS.
 3) FILE INELIGIBLE FOR COMPRESSION BECAUSE THE LENGTH OF THE DATA PAST THE KEY WAS NOT MORE THAN 10 BYTES.
 4) FILE HAS A HIGH DELETION RATE WHICH MAY RESULT IN THE IAM ESTIMATES BEING UNDERSTATED BY MORE THAN 5%.
 5) FILE CONTAINS SPANNED RECORDS, IF SAMPLING WAS DONE, RECORDS WERE READ FROM THE BEGINNING OF THE FILE ONLY.

B. The Savings Summary Reports

IAM VS. VSAM ALLOCATION SUMMARY												
DEVICE TYPE	TOTAL DSNS	VSAM TRACKS			IAM STANDARD TRACKS \$				IAM COMPRESSED TRACKS \$			
		ALLOCATED	USED	%USED	ALLOCATED	SAVED	%SAVED	SAVED	ALLOCATED	SAVED	%SAVED	SAVED
3380	4	5,007	4,582	92	3,360	1647	36	\$1,120	1,575	3,432	75	\$2,334

NOTE: THE ESTIMATED SAVINGS ARE BASED UPON THE USED NUMBER OF VSAM TRACKS AT A RATE OF \$17 PER MB.
 THE ACTUAL TRACKS USED BY AN IAM FILE MAY VARY BY + OR - FIVE PERCENT FROM THE ESTIMATES SHOWN.

42.06 IAMSIMVS - EXAMPLES

42.06 IAMSIMVS - EXAMPLES

EXAMPLE 1 Simulate the conversion of a number of VSAM files to IAM. Use the default values for the simulation.

```
//SIMULATE EXEC PGM=IAMSIMVS
//STEPLIB DD DISP=SHR,DSN=iam.library <==USER CHANGE
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
SELECT DSN=(VSAM.KSDS.FILE1,VSAM.KSDS.FILE2,
VSAM.KSDS.FILE3,VSAM.KSDS.FILEN)
SELECT DSN=(PAYROLL.VSAM.FILE,CICS.VSAM.FILE)
...
SELECT DSN=(MASTER.VSAM.FILE,TABLE.VSAM.FILE)
/*
```

EXAMPLE 2 Simulated the conversion of four large VSAM clusters. The default values will be used for the first two clusters. The third cluster if converted to IAM, will need a larger amount of overflow set aside for inserts. Use a sample rate of 100% instead of the default of 10% for the fourth file, which is known to have variable length records which vary greatly is size.

```
//SIMULATE EXEC PGM=IAMSIMVS
//STEPLIB DD DISP=SHR,DSN=iam.library <==USER CHANGE
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
SELECT DSN=(LARGE.VSAM.KSDS.FILE1,LARGE.VSAM.KSDS.FILE2)
SELECT DSN=FILE.WITH.MANY.ADDITION,OVERFLOW=100000
SELECT DSN=LARGE.FILE.WITH.STRANGE.DATA,SAMPLE=100
/*
```

EXAMPLE 3 Simulate the conversion of two large KSDS clusters. Since IAM normally releases the unused portion of its allocation, have IAMSIMVS estimate its savings based upon the number of the allocated VSAM tracks instead of the IAMSIMVS default value of tracks used. This will result in a maximum savings figure for the amount of disk space IAM can return. Using this savings figure allows a direct comparison of IAM's disk space savings to that provided by any other technique. In addition the cost of the DASD per million bytes (MB) is changed from \$17 per MB to \$25 per MB. Use the 10% default sampling rate and defaults for all other values.

```
//SIMULATE EXEC PGM=IAMSIMVS
//STEPLIB DD DISP=SHR,DSN=iam.library <==USER CHANGE
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
REPORT ESTIMATE=ALLOC,COST=25
SELECT DSN=(CICS.MASTER.FILE1,CICS.MASTER.FILE2)
/*
```

42.06 IAMSIMVS - EXAMPLES

EXAMPLE 4 Use IAMSMFVS to produce a Cluster Size Report ranking all (up to 500) of the KSDS Clusters identified in the 'currently available' SMF history file. Then use IAMSIMVS to produce a detail Data Set Report showing the amount of savings a conversion to IAM would produce for each of these clusters and a Summary Saving Report - showing the overall saving the site would receive if all of these clusters were converted to IAM. Since many of the clusters are over allocated and this space can be recovered by IAM, estimate IAM's savings based on tracks allocated instead of tracks used. IAM will RELEASE the unused portion of an allocation, VSAM can not. RELEASE can save hundreds of cylinders, or more, depending on the number VSAM clusters a site has permanently allocated.

```
CREATE //IAMSMFVS EXEC PGM=IAMSMFVS,REGION=2048K
IAMSIMVS //STEPLIB DD DISP=SHR,DSN=iam.library <==USER CHANGE
SELECT //SYSMF DD DISP=SHR,DSN=... <==POINTS TO SMF DATA
STATEMENTS //SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SORTIN DD UNIT=SYSDA,SPACE=(CYL,(25,5))
//SORTOUT DD UNIT=SYSDA,SPACE=(CYL,(25,5))
//SORTLIB DD DISP=SHR,DSN=SYS1.SORTLIB
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(25,5))
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(25,5))
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(25,5))
//SYSOUT DD SYSOUT=*
//SYSPUNCH DD DSN=VSAM.CLUSTER.NAMES,DISP=(,CATLG),
// UNIT=SYSDA,SPACE=(CYL,(5,5))
//SYSIN DD *
REPORT MAXDSNS=6000,MAXREPORT=500
/*

REPORT //SAVINGS EXEC PGM=IAMSIMVS
THE //STEPLIB DD DISP=SHR,DSN=iam.library <==USER CHANGE
SAVINGS //SYSPRINT DD SYSOUT=*
AN IAM //SYSUDUMP DD SYSOUT=*
CONVERSION //SYSIN DD *
WILL PROVIDE REPORT ESTIMATE=ALLOC
// DD DSN=VSAM.CLUSTER.NAMES,DISP=SHR <==SEE NOTE 1
//
```

NOTE1: The first step executes IAMSMFVS to create a data set containing formatted IAMSIMVS Select Command statements for the 500 largest VSAM clusters. The second step then reads this data set to determine the names of the VSAM clusters to simulate a conversion for this data set is concatenated after the instream REPORT Command Statement, which is used to establish default processing parameters.

THIS PAGE WAS INTENTIONALLY LEFT BLANK

SECTION 43: TRACE/DEBUG FACILITY

43.00 IAM V6 TRACE/DEBUG

43.01 IAM TRACE/DEBUG OVERVIEW

OVERVIEW IAM's Trace/Debug allows an IAM user to trace all of an application's IAM I/O requests. This facility reports the users requested IAM command, IAM completion status, key value, and time of the request.

FUNCTIONAL SUMMARY The Trace/Debug facility of IAM can be activated for any batch or CICS (asynchronous) task accessing one or more IAM files. Trace information is recorded in a trace data set for each IAM file during the execution of the application program. This data set can then be interrogated. The IAM trace report generator will selectively produce reports based on user supplied reporting criteria. The select options include - command(s) used, completion status, key value, request time, sense information or any combination thereof. Optionally IAM can at execution time directly produce a Run Time Trace Report of all requests in chronological sequence.

TECHNICAL DESCRIPTION TRACE/DEBUG is invoked through the execution time override facility using IAM override control statement keyword TRACEDD=. The Trace/Debug recording data set is opened and closed in conjunction with the opening and closing of the IAM file. DISP=MOD is not supported for the trace data set.

ACTIVATING TRACE/DEBUG The IAM TRACE facility is activated using IAM override control statements. Use of the TRACE facility requires the addition of a trace recording data set DD statement to the JCL. The override control statement(s) must identify the function being performed (i.e. CREATE or ACCESS) and contain a keyword TRACEDD= identifying the DD name of a trace recording data set. Other override keywords may still be used. Refer to [Section 11](#) - 'IAM OVERRIDE STATEMENTS' for a full description of overrides and their usage.

COMPATIBILITY For those who have used the trace facility in prior releases, the technique of supplying a pseudo-member name on the IAM file to activate IAM trace processing is still supported. However, due to its increased flexibility and consistency, the use of overrides is recommended.

43.02 JCL REQUIREMENTS FOR ACTIVATING IAM TRACE/DEBUG

43.02 JCL REQUIREMENTS FOR ACTIVATING IAM TRACE/DEBUG

IAMTRPRx STATEMENT/ RUN TIME TRACE REPORT

A Run Time Trace Report of IAM processing can be produced directly at program execution time. IAM's Run Time Trace Report is a chronological report with one line for each request to IAM. A substantial report will result from any processing of more than a few hundred I/O requests. To activate Run Time Trace, use an override TRACEDD= value beginning with the seven character prefix 'IAMTRPR'. The last (eighth) character may be user defined. IAMTRPRx is usually a SYSOUT dataset. This data set cannot be used as input to the report processor program IAMONRPT.

NOTE: Using a 'C' as the eighth character of the ddname (i.e. IAMTRPRC) indicates that the key of the records reported should be printed in character format instead of the default hexadecimal.

For example: To print the trace in Hex

```
//USERJOB      EXEC PGM=userpgm
.
.
//IAMFILE      DD DSN=...
//IAMTRPRX     DD SYSOUT=A
//IAMOVRID     DD *
                ACCESS DDNAME=IAMFILE,TRACEDD=IAMTRPRX
/*
```

TRACE RECORDING DD

The ddname used for the IAM TRACE recording file must correspond to the ddname specified in the IAM override statement. If the recording DD data set is on disk, sufficient space should be allocated to record the desired amount of data. Secondary allocation is ignored. When the primary space allocation is exhausted, Trace/Debug recording will be discontinued.

For example:

```
//USERJOB      EXEC PGM=userpgm
.
.
//USERIAM      DD DSN=name.of.IAM.file,DISP=...
//TRACE        DD DSN=name.of.trace.file,
//              UNIT=SYSDA,SPACE=(TRK,nnn),DISP=(,CATLG)
//IAMOVRID     DD *
                ACCESS DDNAME=USERIAM,TRACEDDNAME=TRACE
/*
```

NOTE: Trace/Debug data sets are opened and closed in conjunction with the IAM file. Once the data set is reopened, the previously recorded information will be overlaid. 'DISP=MOD' is not supported.

43.03 JCL REQUIREMENTS FOR THE IAM TRACE REPORT GENERATOR-IAMONRPT

43.03 JCL REQUIREMENTS FOR THE IAM TRACE REPORT GENERATOR-IAMONRPT

JCL REQUIREMENTS After a TRACE data set has been created detailed reports can be generated from the recorded data using the IAM trace report program IAMONRPT.

The JCL statements required to execute IAMONRPT are as follows:

EXECUTE STATEMENT Specifies the name of the IAM trace report program - IAMONRPT.

STEPLIB or JOBLIB DD STATEMENT Specifies the library containing the IAM program load library. It is recommended that the IAM program load library be included in the system LINKLIST. If it is not, include a STEPLIB/JOBLIB DD statement specifying the library in which the IAM program load modules reside. May be omitted if these modules are in a system connected library (e.g.: SYS1.LINKLIB).

SYSUDUMP DD STATEMENT Specifies the ABEND data set used if major errors are detected. Usually a SYSOUT dataset.

SYSPRINT DD STATEMENT Specifies where the IAMONRPT reports are to be printed. Usually a SYSOUT data set. If this file is to be written to a dataset, the blocksize must be a multiple of 121.

SYSIN DD STATEMENT Specifies the input control statement data set. Usually a DD * data set.

IAMNxxxx DD STATEMENT Identifies the IAM trace recording data set(s) from which the report data is to be taken. The four (4) characters in the ddname IAMNxxxx represented by xxxx may be any combination of from 1 to 4 alphanumeric characters. IAMONRPT will serially process any number of IAMNxxxx data sets.

43.04 IAMONRPT TRACE REPORT CONTROL STATEMENT FORMAT

43.04 IAMONRPT TRACE REPORT CONTROL STATEMENT FORMAT

TRACE COMMAND The TRACE statement directs the IAM Report Generator to prepare a report on the IAM I/O processing events that were recorded in the Event recording file. The content and format of the report will be based on the default options in effect or as specified on the TRACE control statement itself.

TRACE	CMD=(c,...,c)	,KEY=xxxxxx
	,COMP=xx	,RTYPE=cccccc
	,FROMHOUR=hh	,SENSEBYTE=cccc
	,FROMMINUTE=mm	,TOHOUR=hh
	,FROMREQUEST=nnnn	,TOMINUTE=mm
	,HEX	,TOREQUEST=nnnnnn

OPERANDS **CMD=** Will limit the report to only those transactions with commands matching those specified. From 1 to 8 file access commands may be specified as a character string.

A - Add (random)
K - Start class (random)
C - Close
N - Read or getnext (random)
D - Delete
R - Read
F - Getnext (random)
S - Start key (random)
G - Get (sequential)
W - Rewrite

By default, all commands are selected.

COMP= Will limit the report to transactions to which IAM returned a completion code of 'xx'. 'xx' is specified as a two (2) digit hexadecimal value.

Some of the more common completion codes are:

X'00' - Function successfully completed.
X'FF' - Function unsuccessful or end of file.
X'C6' - File area full (C'F').

By default, all completion codes are selected.

43.04 IAMONRPT TRACE REPORT CONTROL STATEMENT FORMAT

- FROMHOUR=** The starting hour to be used by IAMONRPT in selection of transactions. When an hour is specified, events recorded during or after that hour are selected.
FH=
- The default starting hour is the time the task was activated.
- FROMMINUTE=** The starting minute value to be used by IAMONRPT in selection of transactions.
FM= This operand is normally used in conjunction with 'FROMHOUR' option to further define the time selection criteria. When used in conjunction with the 'TOMINUTE' keyword, IAMONRPT will select only the transactions which occurred between those specified minutes.
- The default is 00.
- FROMRE-
QUEST=** A starting transaction number, from 1 to 8 digits, to be used by IAMONRPT in selection of transactions. Each transaction within the trace data set is numbered beginning with 1. The combination of the 'FROMREQUEST' and 'TOREQUEST' operands can be used to produce reports based upon this sequence number.
FR=
- The default is 0.
- HEX** Directs IAMONRPT to print the key value in hexadecimal. A maximum of thirty (30) bytes of the key can be printed when 'HEX' is specified.
- By default, the key value print is printed in character.
- KEY=** Directs IAMONRPT to process only those transactions which have 'xxxxxx' in the left most portion of the key. The key is specified in hexadecimal, with a maximum length of 50 hexadecimal digits (25 data bytes). Up to 200 keys may be specified during any given execution. Multiple key prefixes may be entered in a single command as follows:
- KEY=(xxxxxx,xxxxxx,.....,xxxxxx)
- By default, all keys are selected.
- RTYPE=** Identifies the type of report to be generated.
- SUMMARY** - Print only the summary statistics.
USER - Provide a detailed list of commands, completion codes and other information along with summary statistics.
- The default RTYPE is USER.
- SENSEBYTE=** Only those transactions having an IAM sense byte SNS=(expanded) matching the value 'cccc' are to be selected. Acceptable values for 'cccc' are:
- IND** - Select records read from or written to independent overflow.
LIND - Select records read from or written to the last block of independent overflow.
DYN - Select records read from or written to IAM Dynamic Table.
- By default, the value in the sense byte of a record is ignored for selection purposes.

43.04 IAMONRPT TRACE REPORT CONTROL STATEMENT FORMAT

TOHOUR= TH=	Identifies the ending hour to be used in selection of transactions. When an ending hour is specified, events recorded before or during the hour specified are selected. The default is EOF on the trace data set.
TOMINUTE= TM=	The ending minute value to be used by IAMONRPT in selection of transactions. When used in conjunction with the 'TH' option it further defines the ending time for transaction selection. When used in conjunction with the 'FM' option, IAMONRPT will select transactions occurring during those minutes of each hour. The default is 59.
TOREQUEST= TR=	Identifies the ending transaction number to be selected. This value may be from 1 to 8 digits. Each transaction within the trace data set is numbered beginning with 1. The 'TR' option used in conjunction with the 'FR' option can be used to produce reports based upon the sequence number. IAMONRPT will normally go to EOF on the trace data set.

43.05 TRACE/DEBUG EXAMPLES

43.05 TRACE/DEBUG EXAMPLES

TRACE TO A DISK 1. Activate Trace/Debug for an IAM data set when the program 'userpgm' is executed. Record the trace information to a disk data set for later interrogation.

```
TRACE //STEP1          EXEC  PGM=userpgm
DATA SET .
          .
          .
          //USERDD          DD  DSN=any.i am. file, DISP=OLD
          //TRACEDD         DD  DSN=trace. file, DISP=(,CATLG),
          //                  UNIT=SYSDA, SPACE=(CYL,5)
          //IAMOVRID        DD  *
          ACCESS          DDNAME=USERDD, TRACEDDNAME=TRACEDD
/*
```

PRINT A TRACE REPORT FROM A DISK 2. Print a trace report of all ADD requests made to the file from the example above.

```
TRACE //STEP2          EXEC  PGM=IAMONRPT
DATA SET //SYSPRINT        DD  SYSOUT=*
          //IAMN01         DD  DSN=trace. file, DISP=SHR
          //SYSIN          DD  *
          TRACE          CMD=A
/*
```

43.05 TRACE/DEBUG EXAMPLES

IAMINFO 3. Activate run time Trace/Debug for an IAM data set accessed by the program 'userpgm'. The
THE information contained in the key of the records should be printed in hexadecimal format.
RUN TIME
REPORT

```
//STEP1      EXEC   PGM=userpgm
.
.
//USERDD1    DD     DSN=CICS.MASTER.IAMFILE,DISP=OLD
//IAMTRPR1   DD     SYSOUT=A
//IAMOVRID   DD     *
              ACCESS DDNAME=USERDD1,TRACEDDNAME=IAMTRPR1
/*
```

NOTE: To produce the trace key in character format use IAMTRPRC

```
//IAMTRPRC   DD     SYSOUT=A
              .
              .
              ACCESS DDNAME=USERDD1,TRACEDDNAME=IAMTRPRC
              .
              .
```

SAMPLE TRACE OUTPUT						
IAM400 IAM TRACE FACILITY--IAMTRACE VER 6.0 INNOVATION DATA PROCESSING						
IAM332 DATA SET TRACED: CICS.MASTER.IAMFILE TIME: 14.00.21.003951						
COMMAND	COMP CODE	REQUEST	CMPLTED	SNS	NUMBER	IAM REQUEST KEY INUSE...HEXADECIMAL
-----	----	-----	-----	---	-----	1...5...10...15...20...25...30.
ADD	00	14.00.21.128800		IND	1	0000003A
READ	FF	14.00.21.129475			2	00000030
READ	FF	14.00.21.130210			3	00000000
ADD	00	14.00.21.130474		IND	4	0000003C
SETL KC	00	14.00.21.134994			5	00000020
GET	00	14.00.21.136305			6	00000020
DELETE	00	14.00.21.137994			7	00000020
GET	00	14.00.21.139014		IND	8	00000021
DELETE	00	14.00.21.140615		IND	9	00000021
READ	FF	14.00.21.193803			10	00000022
READ	FF	14.00.21.194839			11	00000000
READ	00	14.00.21.196047			12	00000024
READ	FF	14.00.21.198554			13	00000020
READ	FF	14.00.21.208278			14	00000000
READ	00	14.00.21.211718			15	00000032
REWRITE	00	14.00.21.218336			16	00000032
READ	FF	14.00.21.218475			17	00000030
READ	00	14.00.21.218224		IND	19	00000036
REWRITE	00	14.00.21.219336		IND	20	00000036

SECTION 44: IAM EVENT RECORDING MONITOR

44.01 IAM EVENT RECORDING MONITOR OVERVIEW

44.01 IAM EVENT RECORDING MONITOR OVERVIEW

OVERVIEW IAM's Event Recording Monitor facility provides both batch and CICS (asynchronous) users with a means of reviewing detailed information on the functions performed by IAM to satisfy a user's I/O requests. This IAM facility allows the user to measure the performance of an individual application's I/O requests. When the monitor is active, the following data is recorded; IAM command, completion code, sense information, time of the request, time of completion, number of outstanding concurrent requests and the event control block in use. The data can then be post-processed for presentation in general or specific purpose summary, graphic and detail reports.

FUNCTIONAL SUMMARY IAM's Event Recording Monitor can be activated for any batch or CICS (asynchronous) task accessing IAM files. Monitor information is recorded in a separate monitor data set for each IAM file during execution of an application program. The monitor data set may then be selectively interrogated by the IAM monitor report program - IAMONRPT. Reports can be generated based upon command, time of request, duration of request, duration of I/O, or any combination thereof. IAM's Event Recording Monitor can be especially valuable in providing statistics which assist in the tuning of a CICS system's I/O processing.

TECHNICAL SUMMARY IAM's Event Recording Monitor is invoked through the execution time override facility using the two IAM override control statement keywords TRACEDD= and TRACEREQUEST=. Monitor is only available for ACCESS processing. The data set created by monitor processing is OPENed and CLOSEd in conjunction with the opening and closing of the IAM FILE. DISP=MOD is not supported for this data set.

44.02 EXECUTION TIME JCL REQUIREMENTS FOR ACTIVATING IAM EVENT RECORDING

44.02 EXECUTION TIME JCL REQUIREMENTS FOR ACTIVATING IAM EVENT RECORDING

ACTIVATING MONITOR The IAM event recording monitor is activated using IAM override control statements. Use of the MONITOR facility requires the addition of a monitor recording data set DD statement to the JCL. The override control statement(s) in addition to identifying the function to be traced as ACCESS must also contain the keywords TRACEDD= identifying the ddname of the monitor recording data set and TRACEREQUEST=MONITOR to request Monitor format output. Other override keywords may still be used. Refer to [Section 11](#) - 'IAM OVERRIDE STATEMENTS' for a full description of overrides and their usage.

NOTE: Monitor is not available for CREATE.

MONITOR RECORDING The ddname used for the IAM Event recording data set must correspond to the name specified by the TRACEDD keyword in the override control statement. If the recording data set is on disk, sufficient space should be allocated to record the desired amount of data. Secondary allocation is ignored. When the primary space allocation is exhausted, monitor recording will be discontinued.

For example:

```
//STEP1      EXEC   PGM=USERPGM
.
.
.
//FILEDD     DD     DSN=name.of.IAM.file,DISP=...
//TRACE      DD     DSN=name.of.monitor.file,
//           UNIT=SYSDA,SPACE=(TRK,nnn),DISP=(,CATLG)
//IAMOVRID   DD     *
           ACCESS DDNAME=FILEDD,TRACEDDNAME=TRACE,
           TRACE  REQUEST=MONITOR
/*
```

NOTE: Monitor data sets are OPENed and CLOSEd in conjunction with the IAM file. Once the data set is reopened, the previously recorded information is no longer available. 'DISP=MOD' is not supported.

COMPATIBILITY For those who have used the monitor facility in prior releases, the technique of supplying a pseudo-member name on the IAM file to activate the IAM monitor is still supported. However, due to its increased flexibility and consistency, the use of overrides is recommended.

44.03 JCL REQUIREMENTS FOR THE MONITOR ANALYSIS PROGRAM - IAMONRPT

44.03 JCL REQUIREMENTS FOR THE MONITOR ANALYSIS PROGRAM - IAMONRPT

After a MONITOR data set has been created detailed reports can be generated from the recorded data.

JCL REQUIREMENTS The JCL statements required to execute IAMONRPT are as follows:

EXECUTE STATEMENT Identifies the IAM monitor analysis program - IAMONRPT.

STEPLIB or JOBLIB DD STATEMENT Specifies the library containing the IAM program load library. It is recommended that the IAM program load library be included in the system LINKLIST. If it is not, include a STEPLIB/JOBLIB DD statement specifying the library in which the IAM program load modules reside.

May be omitted if these modules are in a system connected library (e.g.: SYS1.LINKLIB).

SYSUDUMP DD STATEMENT Specifies the ABEND data set used if major errors are detected. Usually a SYSOUT dataset.

SYSPRINT DD STATEMENT Specifies where the IAMONRPT reports are to be printed. Usually a SYSOUT data set. If this file is to be written to a dataset, the blocksize must be a multiple of 121.

SYSIN DD STATEMENT Specifies the input control statement data set. Usually a DD * data set.

IAMNxxxx DD STATEMENT Identifies the IAM trace recording data set(s) from which the recorded data is to be taken. The four (4) characters in the ddname IAMNxxxx represented by xxxx may be any combination of from 1 to 4 alphanumeric characters. IAMONRPT will serially process any number of IAMNxxxx data sets.

44.04 IAMONRPT MONITOR REPORT COMMAND STATEMENT FORMAT

44.04 MONITOR REPORT COMMAND STATEMENT FORMAT

MONITOR COMMAND The MONITOR statement directs the IAM Report Generator to prepare a report on the IAM I/O processing events that were recorded in the Monitor recording data set. The content and format of the report will be based on the default options in effect or as specified on the MONITOR control statement itself.

MONITOR	CMD=(C,...,C)	,RTYPE=CCCCC
	,COMP=xx	,SENSEBYTE=CCCC
	,FROMHOUR=hh	,TOHOUR=hh
	,FROMMINUTE=mm	,TOMINUTE=mm
	,FROMREQUEST=nnnn	,TOREQUEST=nnnnn
	,IOTIME=nnnnnn	,TRTIME=nnnnnn
	,MAXCONCUR=nn	,XAXIS=nnnnn
	,NOGRAPHS	,YAXIS=nnn
	,QUTIME=nnnnnn	

OPERANDS **CMD=** Will limit the report to only those transactions with commands matching those specified. From 1 to 8 file access commands may be specified as a character string.

- A - Add (random)**
- K - Start class (random)**
- C - Close**
- N - Read or get next (random)**
- D - Delete**
- R - Read**
- F - Get next (random)**
- S - Start key (random)**
- G - Get (sequential)**
- W - Rewrite**

By default, all commands are selected.

44.04 IAMONRPT MONITOR REPORT COMMAND STATEMENT FORMAT

COMP=	<p>Will limit the report to transactions to which IAM returned a completion code of 'xx'. 'xx' is specified as a two (2) digit hexadecimal value.</p> <p>Some of the more common completion codes are:</p> <p>X'00' - Function successfully completed. X'C6' - File area full (C'F'). X'FF' - Function unsuccessful or end of file.</p> <p>By default, all completion codes are selected.</p>
FROMHOUR= FH=	<p>The starting hour to be used by IAMONRPT in selection of transactions. When an hour is specified, events recorded during or after that hour are selected.</p> <p>The default starting hour is the time the task was activated.</p>
FROMMINUTE= FM=	<p>The starting minute value to be used by IAMONRPT in selection of transactions. This operand is normally used in conjunction with 'FROMHOUR' option to further define the time selection criteria. When used in conjunction with the 'TOMINUTE' keyword, IAMONRPT will select only the transactions which occurred between those specified minutes.</p> <p>The default is 00.</p>
FROMREQUEST= FR=	<p>A starting transaction number from 1 to 8 digits to be used by IAMONRPT in selection of transactions. Each transaction within the trace data set is numbered beginning with 1. The combination of 'FROMREQUEST' and 'TOREQUEST' can be used to produce reports based on this sequence number.</p> <p>The default is 0.</p>
IOTIME= IOT=	<p>Specifies a limiting I/O time from one (1) to six (6) digits expressed in milliseconds. Only those transactions produced by the IAM event recording monitor having an I/O time greater than or equal to the value specified will be selected.</p> <p>By default, I/O time does not participate in transaction selection.</p>
MAXCONCUR= MAXC=	<p>Specifies a limiting concurrent user value from one (1) to two (2) digits. Only those transactions produced by the IAM event recording monitor having a maximum concurrent used count greater than or equal to the value specified will be selected.</p> <p>Maximum concurrent used count does not normally participate in transaction selection.</p>
NOGRAPHS NOG	<p>Reports produced from the IAM event recording monitor data are not to be accompanied by graphic displays.</p> <p>By default, graphic displays are produced.</p>
QUTIME= QUT=	<p>Specifies a limiting queue time from one (1) to six (6) digits expressed in milliseconds. Only those transactions recorded by the IAM monitor having a queue time greater than or equal to the value specified will be selected.</p> <p>Queue time does not normally participate in transaction selection.</p>

44.04 IAMONRPT MONITOR REPORT COMMAND STATEMENT FORMAT

RTYPE= RT=	Identifies the type of report to be generated. DETAIL - Print a complete analysis of IAM event recording monitor data. SUMMARY - Print only the summary statistics. USER - Print the summary statistics and a detailed report for the data being analyzed. The default RTYPE is USER.
SENSEBYTE= SNS=	Only those transactions having an IAM sense byte (expanded) matching the value 'cccc' are to be selected. Acceptable values for 'cccc' are: IND - Select records read from or written to independent overflow. LIND - Select records read from or written to the last block of independent overflow. DYN - Select records read from or written to an IAM Dynamic Table. By default, the value of the sense byte of a record is ignored.
TOHOUR= TH=	Identifies the ending hour to be used in selection of transactions. When an ending hour is specified, events recorded before or during the hour specified are selected. The default is EOF on the trace data set.
TOMINUTE= TM=	The ending minute value to be used by IAMONRPT in selection of transactions. When used in conjunction with the 'TOHOUR' option it further defines the ending time for transaction selection. When used in conjunction with the 'FROMMINUTE' option, IAMONRPT will select transactions occurring during those minutes of each hour. The default is 59.
TOREQUEST= TR=	Identifies the ending transaction number to be selected. This value may be from 1 to 8 digits. Each transaction within the trace data set is numbered beginning with 1. The 'TOREQUEST' option used in conjunction with the 'FROMREQUEST' option can be used to produce reports based upon the sequence number. IAMONRPT will normally go to EOF on the trace dataset.
TRTIME= TRT=	Specifies a limiting transaction (REQUEST) time from one (1) to six (6) digits expressed in milliseconds. Only those transactions produced by the IAM event recording monitor having a transaction time greater than or equal to the value specified will be selected. By default, transaction (REQUEST) time does not participate in record selection.
XAXIS= XA=	Specifies a scaler value from one (1) to six (6) digits expressing the completion of an I/O request in milliseconds. Used as a maximum limit on the horizontal axis for all graphs depicting IAM I/O request time and physical I/O time. The value may be any number from 10 to 10,000, inclusive. The default is 200 milliseconds.
YAXIS= YA=	Specifies a scaler value from two (2) to three (3) digits expressing the maximum percentage of I/O requests completing. Used as an upper limit on the vertical axis of all graphs depicting IAM I/O request time and physical I/O time. The value may be any number from 10 and 100, inclusive. The default is 50 percent.

44.05 IAMONRPT EXAMPLES OF MONITOR PROCESSING

44.05 IAMONRPT EXAMPLES OF MONITOR PROCESSING

MONITOR TO A DISK RECORDING DATA SET 1. Activate a monitor on an IAM data set when the program CICS is executed. Record the monitor information to a disk data set.

```
//STEP1      EXEC  PGM=CICS
.
.
//USERDD     DD    DSN=any.iAM.file,DISP=OLD
//MNTRFILE   DD    DSN=IAM.monitor.file,DISP=(,
                CATLG),
//          UNIT=SYSDA,SPACE=(CYL,5)
//IAMOVRID   DD    *
                ACCESS DD=USERDD,TRACEDDNAME=MNTRFILE,
                TRACE  REQUEST=MONITOR
/*
```

PRINT A MONITOR REPORT FROM A DISK RECORDING DATA SET 2. Print a summary report and I/O graphs using the transaction data recorded from the example above.

```
//STEP2      EXEC  PGM=IAMONRPT
//SYSPRINT   DD    SYSOUT=A
//IAMN01     DD    DSN=IAM.monitor.file,DISP=SHR
//SYSIN      DD    *
                MONITOR
/*
```

44.06 SAMPLE MONITOR REPORTS

44.06 SAMPLE MONITOR REPORTS

SUMMARY REPORT

The MONITOR SUMMARY is a concise report on the IAM I/O events that took place during the time period the report is requested for. It gives totals on the number of transactions (requests) processed, the total I/O time, total inactive time (no requests to process) and total time transactions were held on a queue pending service. There are averages given based on the totals and the number of transactions processed. Averages are given for TRANSACTION TIME (the time it took to satisfy a user request from the time IAM received the request until IAM returned control of the record to the user), I/O TIME (the time from IAM's I/O request to the operating system until the I/O is posted), INACTIVE TIME (the average time IAM was waiting for a user request with no work outstanding) and QUEUE TIME (the average time a users' request was held by IAM after it was accepted before IAM was able to issue an I/O for the users' record or obtain it out of a buffer).

```

IAM400 IAM MONITOR REPORT CONTROLLER -- IAMONRPT VER 6.x -- INNOVATION DATA PROCESSING

IAM312 DATASET QUALIFIED BY DDNAME IAMNTR1 IS A MONITOR DATA SET
IAM303 CARD IMAGE - * MONITOR RTYPE=SUMMARY*
IAM400 IAM MONITOR GENERATED REPORTS -- IAMONRPT VER 6.x -- INNOVATION DATA PROCESSING

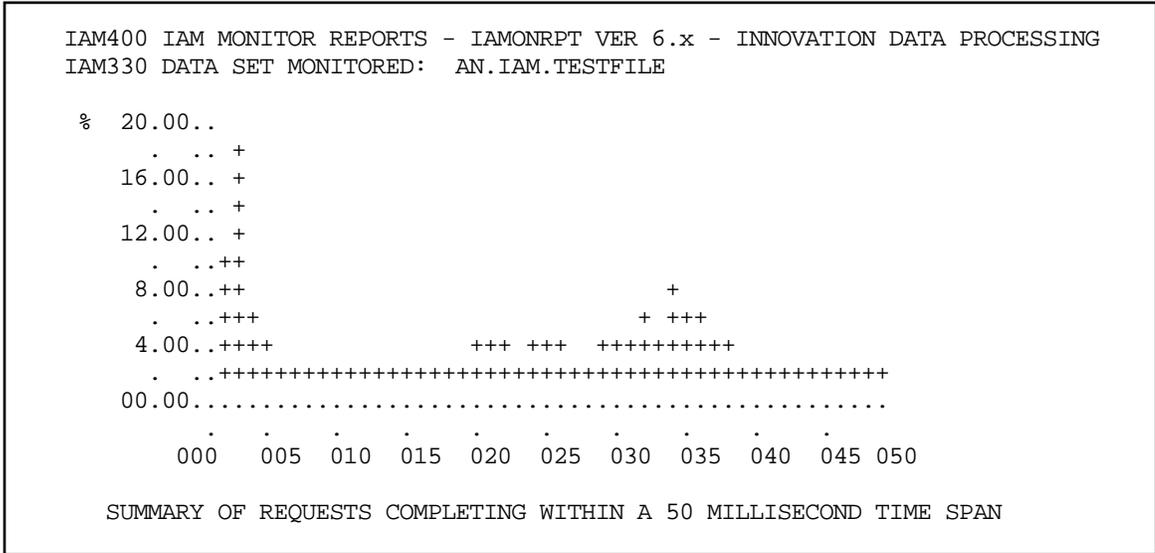
IAM339 SUMMARY REPORT FOR 01 MONITORED DATA SETS

IAM307      DSN ID...01      DATA SET NAME...ANY.IAM.FILE                        IAM MODE...R IAM DPC...U

                IAM ACTIVATED.....10.22.58.553480  IAM DEACTIVATED...10.24.06.582713
REQUESTS PROCESSED...30711  TOTAL TRANS TIME...00.02.21.412881  AVG TRANS TIME...00.00.00.004604
I/O OPERATIONS.....270    TOTAL I/O TIME.....00.00.07.664716  AVG I/O TIME.....00.00.00.028387
INACTIVE PERIODS.....3022  TOTAL INACTIVE TIME.00.00.05.050207  AVG INACTIVE TIME.00.00.00.001671
QUEUE OPERATIONS....27189  TOTAL TIME ON QUEUE.00.02.03.162429  AVG QUEUE TIME...00.00.00.004529
MONITOR TRANSACTION TOTALS --30210      270      270      30710  27189
                                27189  29940      0      0      0
                                0      0      0      0      0
                                0      0      0      0      0
    
```

REQUESTS COMPLETING PERCENTAGE REPORT

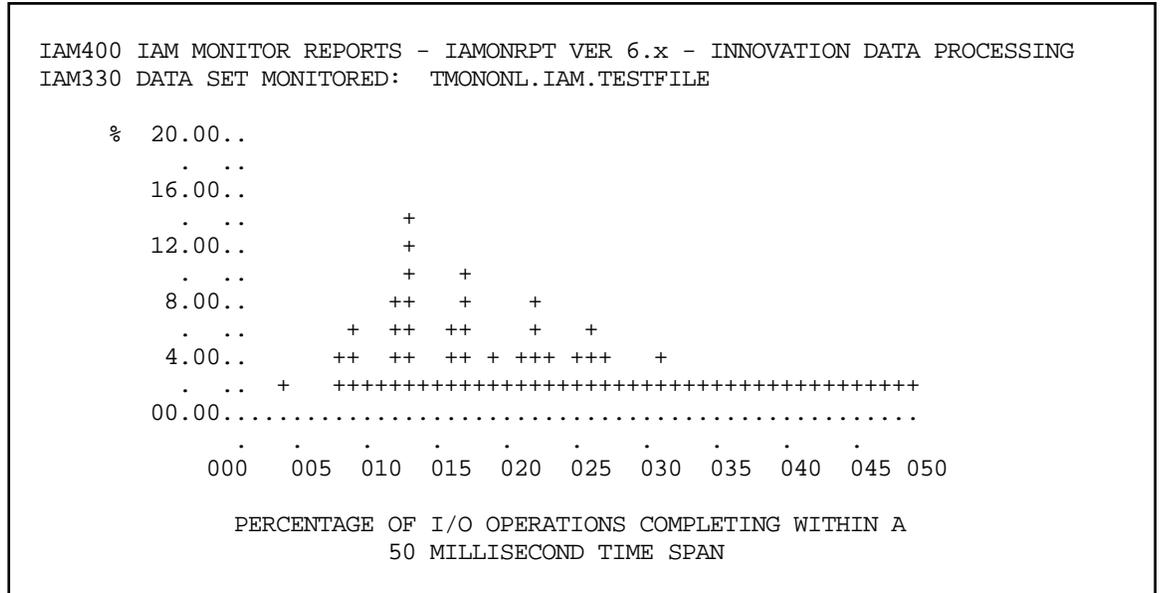
The REQUESTS COMPLETING PERCENTAGE REPORT is a graphic depiction of the individual transaction's request time expressed as to what percentage of requests were able to be satisfied in a particular time span. The time it took to satisfy a user request from the time IAM received the request until IAM returned control of the record to the user.



44.06 SAMPLE MONITOR REPORTS

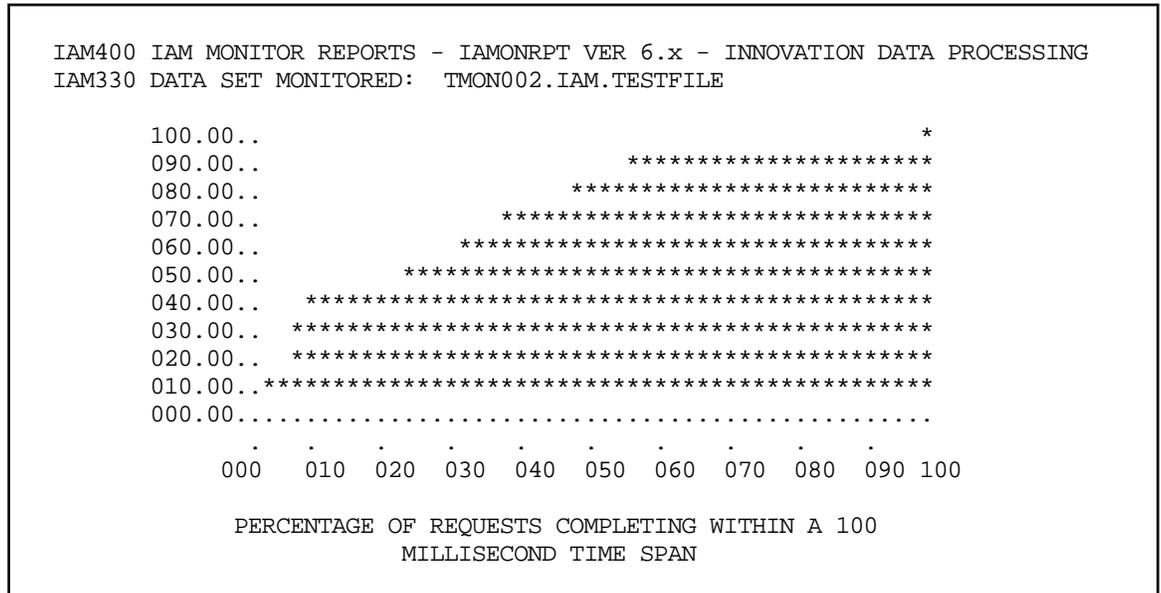
I/O OPERATIONS COMPLETING PERCENTAGE REPORT

The I/O OPERATIONS COMPLETING PERCENTAGE REPORT is a graphic depiction of the individual transaction's I/O time expressed as what percentage of I/O operations that IAM started were completed in a particular time span. The time from IAM's I/O request to the operating system until the I/O was posted as complete.



REQUESTS COMPLETING SUMMARY REPORT

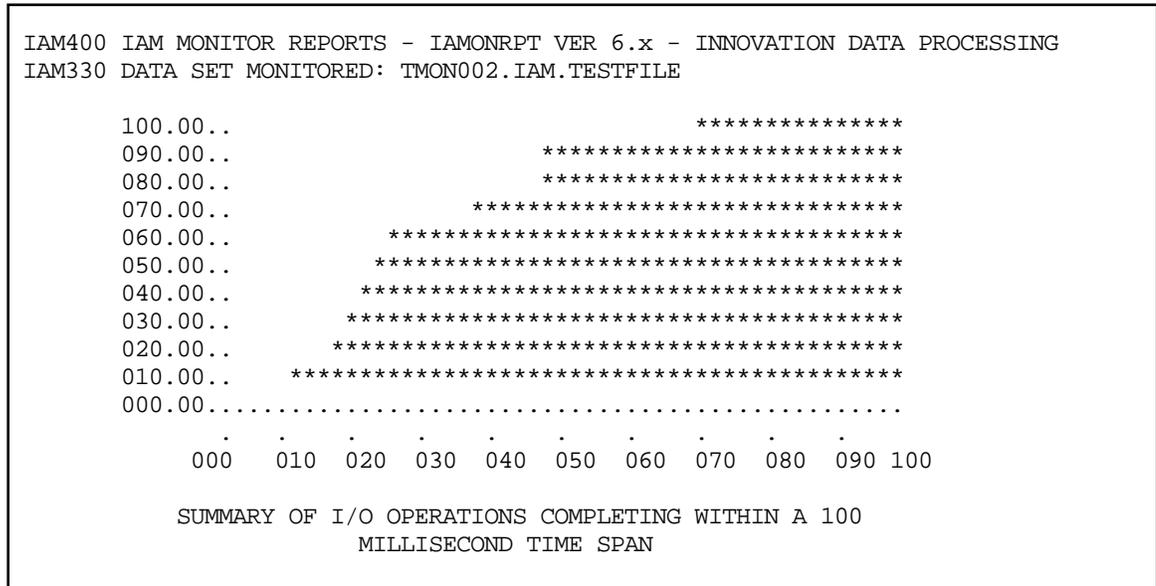
The REQUESTS COMPLETING SUMMARY REPORT is a graphic depiction of the individual transaction's request time expressed as to what cumulative percentage of requests were able to be satisfied in a particular time span. The time it took to satisfy a user request from the time IAM received the request until IAM returned control of the record to the user.



44.06 SAMPLE MONITOR REPORTS

I/O OPERATIONS COMPLETING SUMMARY REPORT

The I/O OPERATIONS COMPLETING SUMMARY REPORT is a graphic depiction of the individual transaction's I/O time expressed as what cumulative percentage of I/O operations that IAM started were completed in a particular time span. The time from IAM's I/O request to the operating system until the I/O was posted as complete.



DETAIL REPORT

The DETAIL REPORT is an exact report on all of the IAM I/O events that took place during the time period the report is requested for. It gives for each transaction (request) processed, the I/O command requested by the user, the time the request was accepted by IAM and the time it was completed. Additionally for each request the report shows the following the REQUEST TRANSACTION TIME (the time it took to satisfy a user request for the time IAM received the request until IAM returned control of the record to the user), QUEUE TIME if any (the time a users' request was held by IAM after it was accepted before IAM was able to issue an I/O for the user's record or obtain it out of a buffer), I/O TIME (the time from IAM's I/O request to the operating system until the I/O was posted....there may be no I/O time if IAM was able to satisfy a users request from a buffer maintained in storage or from a DYNAMIC TABLE), IAM SENSE (an indication if the record went to or came from overflow or a Dynamic Table, a sequential request number maintained by the report processor and finally the number of requests active at the time the transaction was completed (an indication as to the number of concurrent requests queued by IAM).

```

IAM400 IAM MONITOR GENERATED REPORTS -- IAMONRPT VER 6.x -- INNOVATION DATA PROCESSING DATE - yy.ddd
IAM330 DATA SET MONITORED: TMON002.IAM.TESTFILE

```

COMMAND	COMP CODE	REQUEST ISSUED	REQUEST CMLPTED	REQUEST TIME	TIME ON QUEUE	I/O TIME USED	IAM SNS	REQ NUM	REQ ACT
GET	00	16.21.29.562837	16.21.29.580849	00.00.018012		00.00.017584		28801	5
READ	00	.565057	.597181	.032124	00.00.015825	.015976		28802	5
READ	00	.566003	.638803	.072800	.031213	.041246		28803	4
READ	00	.566963	.679847	.112884	.071879	.040670		28804	4
READ	00	.567934	.771111	.203177	.111953	.089596		28805	3
ADD	00	.583641	.771229	.187588	.187512			28806	2
ADD	00	.640958	.773456	.132498	.132419			28807	3
ADD	00	.772096	.774209	.002113	.002042			28808	2
READ	00	.772916	.775165	.002249	.001962			28809	1
READ	00	.776002	.776338	.000336	.000258			28810	1
GET	00	.777776	.800371	.022595		00.00.022166		28811	5
READ	00	.780004	.812131	.032127	00.00.020402	.011404		28812	5
READ	FF	.780957	.913161	.132204	.031208	.100669		28813	4
READ	00	.781880	.927898	.146018	.131358	.014322		28814	4
READ	FF	.782878	.952971	.170093	.145054	.024710		28815	3

SECTION 45: IAM RECOVERY PROGRAM (IAMRECVR)

45.01 RECOVERY PROGRAM OVERVIEW

45.01 RECOVERY PROGRAM OVERVIEW

OVERVIEW The IAMRECVR program has been designed to permit users to recover the contents of IAM files which have become unusable due to data checks on disk, having been partially over written or some physical damage. Both standard and compressed format files can be processed by IAMRECVR.

Data checks on disk and other causes of physical damage to disk data sets are infrequent, but they do occur. Files unfortunately can be OPENed accidentally and partially written over by unauthorized programs, or simply OPENed for output and CLOSEd. If this occurs, files become unusable to the applications depending on them. To assist in recovery of valuable data, IAM includes a special function program, IAMRECVR, which can often create a backup copy of the data in an IAM file which is unreadable by normal means.

IAMRECVR reads the damaged IAM file using a high performance EXCP technique when it is run on MVS systems, the standard BSAM access method when run on VS1 or VM/CMS systems. BSAM can be forced for MVS systems by specifying the BSAM keyword on the DIAGNOSE, PRINT, and RECOVER commands. Only blocks determined to match the IAM file specifications are processed. Any block determined to be unreadable due to physical damage to the device (e.g.: data check) will be skipped. The logical records accepted are written to a sequential backup file in an uncompressed format. These records are sequence checked and sequence errors reported. IAMRECVR can also be used to validate the condition of any IAM file.

IAMRECVR determines an IAM file's specifications directly from its IAM control record. Recovery processing of a file with an unreadable control record is also supported. The required file specifications can be supplied using the operands on the 'RECOVER' command. File specifications can be obtained from any IAM Run Time Statistics (INFO) report, IDCAMS LISTCAT ALL which had an IAMPRINT DD statement included, or an IAMRECVR 'LIST', produced before the file's control record was destroyed.

45.02 JCL REQUIREMENTS FOR IAM FILE RECOVERY

45.02 JCL REQUIREMENTS FOR IAM FILE RECOVERY

JCL REQUIREMENTS	The JCL statements required to execute IAMRECVR are as follows:
EXECUTE STATEMENT	Specifies the name of the IAM recovery program -- IAMRECVR.
STEPLIB or JOBLIB DD STATEMENT	Specifies the library containing the IAM program load library. It is recommended that the IAM program load library be included in the system LINKLIST. If it is not, include a STEPLIB/JOBLIB DD statement specifying the library in which the IAM program load modules reside. May be omitted if these modules are in a system connected library (e.g.: SYS1.LINKLIB).
SYSUDUMP DD STATEMENT	Specifies the ABEND data set used if major errors are detected. Usually a SYSOUT dataset.
SYSPRINT DD STATEMENT	Specifies where the IAMRECVR messages are to be printed. Usually a SYSOUT data set.
DISKIN DD STATEMENT	Specifies the IAM data set to be recovered, listed or printed.
DISKOUT DD STATEMENT	Specifies the new IAM file to be created from the damaged IAM file. The new IAM file will have the same characteristics as the IAM file being recovered.
TAPEOUT DD STATEMENT	Specifies the sequential recovery file to be created by IAMRECVR on either a tape or direct access device. If a direct access device is used, enough space must be allocated to contain a sequential copy of the IAM file. The use of secondary allocation values is permitted and encouraged.
LOG DD STATEMENT	For the RECOVER option, specifies a sequential file where the duplicates, if any, are to be stored. For the APPLY option, specifies the sequential file to be used as input.
VSAMOUT DD STATEMENT	Specifies an IAM or VSAM file to be loaded from the damaged IAM file. If a new IAM or VSAM file is being created, it must be DEFINED using IDCAMS prior to the recover.
SYSIN DD STATEMENT	Specifies the input control statement data set. Usually a DD * data set.

45.03 APPLY COMMAND STATEMENT - FORMAT

45.03 APPLY COMMAND STATEMENT - FORMAT

APPLY COMMAND STATEMENT The APPLY command reads the log data set created by the RECOVER command and either adds records to or replaces records in an IAM file. Every record in the log data set will be processed.

```
APPLY          AUDIT=cccccc  
  
              ,IAMDDNAME=ddname  
  
              ,LOGDDNAME=ddname  
  
              ,OUTPUTFILE=cccc  
  
              ,PRTLENGTH=nnnnn  
  
              ,VSAMDDNAME=ddname
```

OPERANDS The following operands may be specified with the APPLY command.

AUDIT= Defines the audit trail requirements for records processed by the APPLY command.

KEY - Print the key from the data record.

NONE - No audit trail is produced.

RECORD- Print the entire data record.

The default is KEY.

IAMDDNAME= Defines the DDNAME of the IAM file that is to be updated.

I= The default is DISKOUT.

LOGDDNAME= Defines the DDNAME of the log file to be used.

The default DDNAME is LOG.

OUTPUTFILE= Specifies which access method interface to use to apply updates.

OU=

IAM - Use the native IAM interface.

VSAM - Use standard VSAM GET/PUT macro's. This option is required if the file to be updated is a real VSAM file. If the file to be updated is an IAM file, IAM VSAM INTERFACE (VIF) must be active.

PRTLENGTH= Limit the amount of information printed in the AUDIT trail to this value or the length of the key or data, whichever is smaller.

PR=

The default is 32768.

VSAMDDNAME= Defines the ddname of the IAM or VSAM file that is to be updated when **OUTPUTFILE=VSAM** is specified.

V=

The default is VSAMOUT.

45.04 DEFAULT COMMAND STATEMENT - FORMAT

45.04 DEFAULT COMMAND STATEMENT - FORMAT

DEFAULT COMMAND STATEMENT The DEFAULT command is used to assign defaults to recovery options.

DEFAULT	FROMDDNAME=ddname	,SORTMSG=cc
	,IAMDDNAME=ddname	,SORTPFX=cccc
	,LOGDDNAME=ddname	,TODDNAME=ddname
	,PRTLENGTH=nnnnn	,VSAMDDNAME=ddname
	,SORTCORE=nnnnn	,WORKDDNAME=ddname

OPERAND The following operands may be specified with the DEFAULT command.

FROMDDNAME= Defines the DDNAME of the IAM file to be recovered.

FR= The default DDNAME is DISKIN.

IAMDDNAME= Defines the DDNAME of the IAM file to be updated.

I= The default DDNAME is DISKOUT.

LOGDDNAME= Defines the DDNAME of the LOG file to be used.

The default DDNAME is LOG.

PRTLENGTH= A limit for the amount of information printed for each record to this value.

PR= The default is 32768.

SORTCORE= Specifies the amount of storage 'SORT' is to use. The number maybe from 10000 to 8000000 inclusive.

The default is 100000.

SORTMSG= Specifies the message option to be used by 'SORT'

AC - All messages to the console

AP - All messages to the printer (SYSOUT)

CC - Critical messages to the console

CP - Critical messages to the printer

NO - No messages to be produced

PC - Critical messages to both console and printer

The default is CC.

SORTPFX= Specifies the DDNAME prefix to be used by 'SORT'. If the string specified is less than 4 characters, a dollar sign(\$) fill character will be used.

The default is SORT.

TODDNAME= Defines the DDNAME of the sequential output data set created during recovery.

The default is TAPEOUT.

VSAMDDNAME= Defines the DDNAME of the IAM or VSAM file that is to be updated.

V=

WORKDDNAME= Defines the DDNAME of the work file used during a IAM load function.

The default DDNAME is WORK.

45.06 HELP COMMAND STATEMENT - FORMAT

45.06 HELP COMMAND STATEMENT - FORMAT

**HELP
COMMAND
STATEMENT** The HELP command prints help by function or command for IAMRECVR processing. Note that JCL and CLIST are also considered SUBCOMMANDS by the HELP command.

HELP **ALL**

 ,OPERANDS

 ,SUBCOMMAND(command)

OPERANDS The following operands may be specified with the HELP command.

ALL Lists the entire contents of the HELP member, including all commands, field names and descriptions.

OPERANDS When specified with the operand SUBCOMMAND, requests HELP for only the operands of that command. When specified without the operand SUBCOMMAND, requests HELP for the operands of the command or program.

SUBCOMMAND Requests HELP for a specific COMMAND.

SU EXAMPLE: HELP SU(RECOVER) will explain the operands available for the RECOVER command.

45.07 LIST COMMAND STATEMENT - FORMAT

45.07 LIST COMMAND STATEMENT - FORMAT

LIST The LIST command is used to list the characteristics of an IAM file.
COMMAND STATEMENT **LIST** **FROMDDNAME=ddname**

OPERANDS The following operand may be specified with the LIST command.
FROMDDNAME= Defines the DDNAME of the IAM file that is to be listed.
FR= The default is DISKIN.

45.08 PRINT COMMAND STATEMENT - FORMAT

45.08 PRINT COMMAND STATEMENT - FORMAT

PRINT COMMAND STATEMENT The PRINT command is used to print various areas or blocks from an IAM file. A list of the file characteristics is produced upon completion.

PRINT	ALLBLKS	,KEYS
	,BSAM	,MAXBLKS=nnnnn
	,DATA	,OFLOW
	,FBLK=nnnnnn	,PE
	,FROMDDNAME=ddname	,PRTLENGTH=nnnnn
	,IDPINQ	,TBLK=nnnnn

OPERANDS The following operands may be specified with the PRINT command.

ALLBLKS	Specifies that all (or portions of all) blocks in the IAM file are to be printed. <i>NOTE: Use of this operand conflicts with 'DATA', 'FBLK', 'KEYS', 'OFLOW', 'PE', and 'TBLK'.</i>
BSAM	Forces the use of the BSAM access method on MVS systems.
DATA	Specifies that all (or portions of all) prime data blocks in the IAM file are to be printed. <i>NOTE: Use of this operand conflicts with 'ALLBLKS', 'FBLK', and 'TBLK'.</i>
FBLK=	Specifies the block number, relative to 1, from which printing is to begin. <i>NOTE: Use of this operand conflicts with 'ALLBLKS', 'DATA', 'KEYS', 'OFLOW', and 'PE'.</i>
FROMDDNAME=	Defines the DDNAME of the IAM file that is to be listed.
FR=	The default is DISKIN.
IDPINQ	Specifies the IAM control block is to be printed.
KEYS	Specifies that all (or portions of all) key blocks in the IAM file are to be printed. <i>NOTE: Use of this operand conflicts with 'ALLBLKS', 'FBLK', and 'TBLK'.</i>
MAXBLKS=	Specifies the maximum number of blocks to be printed from each area selected in the IAM file. <i>NOTE: Use of this operand conflicts with 'TBLK'.</i> The default is the number of blocks in the area selected or the entire file if 'ALLBLKS' is specified, excluding the IAM control block.

45.08 PRINT COMMAND STATEMENT - FORMAT

OFLOW	Specifies that all (or portions of all) Independent Overflow blocks in the IAM file are to be printed. <i>NOTE: Use of this operand conflicts with 'ALLBLKS', 'FBLK', and 'TBLK'.</i>
PE	Specifies that all (or portions of all) Prime Extension blocks in the IAM file are to be printed. <i>NOTE: Use of this operand conflicts with 'ALLBLKS', 'FBLK', and 'TBLK'.</i>
PRTLENGTH= PR=	Limit the amount of data printed for each block to this value or the length of block, whichever is smaller. The default is 32768.
TBLK=	Specifies the block number, relative to 1, at which printing is to end. <i>NOTE: Use of this operand conflicts with 'ALLBLKS', 'DATA', 'KEYS', 'MAXBLKS', 'OFLOW', and 'PE'.</i>

45.09 RECOVER COMMAND STATEMENT - FORMAT

45.09 RECOVER COMMAND STATEMENT - FORMAT

RECOVER COMMAND STATEMENT The RECOVER command is used to produce either an IAM file or a sequential copy (or both) from a damaged IAM file.

RECOVER	AUDIT=cccccc	,OVERFLOW=nnnnn
	,BLKSIZE=nnnnn	,PRTLENGTH=nnnnn
	,BSAM	,RKP=nnnn
	,COMPRESSED	,SORT=cccccc
	,DUPLICATES=cccccc	,SORTCORE=n...n
	,FROMDDNAME=ddname	,SORTMSG=cc
	,IAMDDNAME=ddname	,SORTPFX=cccc
	,KEYLEN=nnn	,TODDNAME=ddname
	,LOGDDNAME=ddname	,VARIABLE
	,LRECL=nnnnn	,VSAMDDNAME=ddname
	,MAXBLKS=n...n	,WORKDDNAME=ddname
	,OUTPUTFILES=ccc	

OPERANDS The following operands may be specified with the RECOVER subcommand.

AUDIT= Defines the audit trail requirements for duplicate records processed by the RECOVER command if 'SORT=IFREQ' or 'SORT=YES' and 'DUPLICATE=APPLY' or if 'DUPLICATE=PRINT' or 'DUPLICATE=LOG' are specified.

KEY - Print the key from the data record.

NONE - No audit trail is produced.

RECORD - Print the entire data record.

The default is KEY.

BLKSIZE=
B= Specifies the actual blocksize of the IAM file. This value can be obtained from the run time statistics or a 'LIST' command.

NOTE: This value is ignored unless the IAM control record is destroyed, at which time it is required.

BSAM Forces the use of the BSAM access method on MVS systems.

45.09 RECOVER COMMAND STATEMENT - FORMAT

- COMPRESSED** Identifies the file as containing compressed records.
- NOTE: This value is ignored unless the IAM control record is destroyed. This value is optional and is only used to request that the new IAM file is to have a compressed data structure. IAMRECVR can always detect a compressed record and decompress it.*
- DUPLICATES=** Defines the processing requirements for any duplicate records processed by the RECOVER command if 'SORT=IFREQ' or 'SORT=YES' is specified.
- D=**
- APPLY** - Update the IAM file being created with the duplicate records. Ignored unless 'OUTPUTFILES=IAM' or 'OUTPUTFILES=BOTH' is specified.
 - IGNORE** - Ignore duplicate records.
 - LOG** - Create a log data set of any duplicate records. This data set may later be used as input the 'APPLY' command.
 - PRINT** - Print any duplicate records.
- The default is PRINT.
- FROMDDNAME=** Defines the DDNAME of the IAM file that is to be recovered.
- FR=** The default is DISKIN.
- IAMDDNAME=** Defines the DDNAME of the IAM file that is to be created when OUTPUTFILES=IAM or BOTH is specified.
- IAMNAME=** The default is DISKOUT.
- KEYLEN=** Specifies the length of the key within the data records in the IAM file. This value can be obtained from the run time statistics or a 'LIST' command.
- K=**
- NOTE: This value is ignored unless the IAM control record is destroyed, at which time it is required.*
- LOGDDNAME=** Defines the DDNAME of the log file to be created.
- The default is LOG.
- LRECL=** Specifies the logical record length of the data records in the IAM file. This value can be obtained from the run time statistics or a 'LIST' command.
- L=**
- NOTE: This value is ignored unless the IAM control record is destroyed, at which time it is required.*
- MAXBLKS=** Specifies the number of data blocks in the IAM file. This value can be obtained from the run time statistics or a 'LIST' command.
- NOTE: This value is ignored unless the IAM control record is destroyed, at which time it is required.*
- OUTPUTFILES=** Defines the output requirements for the RECOVER subcommand.
- BOTH** - Create both an IAM file and a sequential copy.
 - BOTHV** - Create both an IAM file (using the VSAM interface) or a VSAM file, and a sequential copy.
 - IAM** - Create only an IAM file.
 - SEQ** - Create only a sequential copy of the recoverable data remaining in the IAM file.
 - VSAM** - Create only an IAM file (using the VSAM interface) or a VSAM file.
- The default is SEQ.

45.09 RECOVER COMMAND STATEMENT - FORMAT

- OVERFLOW=** Specifies the number of Independent Overflow blocks in the IAM file. This value can be obtained from the run time statistics or a 'LIST' command.
- O=**
- NOTE: This value is ignored unless the IAM control record is destroyed, at which time it is required.*
- PRTLENGTH=** Limit the amount of data printed for each block to this value or the length of block, which ever is smaller. The default is 32768.
- PR=**
- RKP=** Specifies the relative location of the key within a data record in the IAM file. This value can be obtained from the run time statistics or a 'LIST' command.
- NOTE: This value is ignored unless the IAM control record is destroyed, at which time it is required.*
- SORT=** Defines the output sorting requirements for the RECOVER command.
- IFREQ** - Sort the records only if sequence checks are encountered in the file.
NO - Do not sort the records.
YES - Sort the records.
- The default is NO.
- SORTCORE=** Specifies the amount of storage the program 'SORT' is to use if external sorting is required. The number may be from 10000 to 8000000 inclusive.
- The default is 100000.
- SORTMSG=** Specifies the message option to be used by the program 'SORT' if external sorting is required.
- AC** - All messages to the console
AP - All messages to the printer (SYSOUT)
CC - Critical messages to the console
CP - Critical messages to the printer
NO - No messages to be produced
PC - Critical messages to both console and printer
- The default is CC.
- SORTPFX=** Specifies the DDNAME prefix to be used by the program 'SORT' if external sorting is required. If the string specified is less than 4 characters, a dollar sign(\$) fill character will be used.
- The default is SORT.
- TODDNAME=** Defines the DDNAME of the sequential output data set created during recovery.
- The default is TAPEOUT.
- VARIABLE** Identifies the IAM file as having variable length records.
- NOTE: This value is ignored unless the IAM control record is destroyed, at which time it is required.*
- VSAMDDNAME=** Defines the DDNAME of the IAM or VSAM file to be created when OUTPUTFILES=VSAM or BOTHV is specified.
- The default is VSAMOUT.
- WORKDDNAME=** Defines the DDNAME of the work file used during an IAM load function.
- The default is WORK. Under MVS systems this DD will be dynamically allocated for the user.

45.10 IAMRECVR JCL EXAMPLES

45.10 IAMRECVR JCL EXAMPLES

LIST FILE ATTRIBUTES

1. The user wishes to display the current attributes of an IAM file.

```
//LIST          EXEC  PGM=IAMRECVR
//SYSPRINT     DD    SYSOUT=A
//DISKIN       DD    DSN=i am . f i l e , D I S P = S H R
//SYSIN        DD    *
LIST
/*
```

DIAGNOSE A FILE FOR DAMAGE

2. The user is getting data checks accessing an IAM file and wishes to determine the extent of the damage. From the report, the user will be able to identify the types of blocks in the file in error (i.e. Prime Data, Index, Overflow), how many blocks are recoverable and other important information.

```
//DIAGNOSE     EXEC  PGM=IAMRECVR
//SYSPRINT     DD    SYSOUT=A
//DISKIN       DD    DSN=data . c h e c k e d . i a m . f i l e , D I S P = S H R
//SYSIN        DD    *
DIAGNOSE
/*
```

RECOVER A DAMAGED FILE's DATA

3. Any recoverable data in a damaged IAM file is to be extracted by IAMRECVR and placed in a new file.

```
//DEFNEW       EXEC  PGM=IDCAMS
//SYSPRINT     DD    SYSOUT=A
//SYSIN        DD    *
DEFINE        CLUSTER(NAME(new.i am . f i l e) -
              .....
              (attributes taken from the existing file)
              .....
/*
//RECOVER      EXEC  PGM=IAMRECVR
//SYSPRINT     DD    SYSOUT=A
//DISKIN       DD    DSN=old . i a m . f i l e , D I S P = S H R
//DISKOUT      DD    DSN=new . i a m . f i l e , D I S P = O L D
//SYSIN        DD    *
RECOVER      OUTPUTFILES=IAM
/*
```

45.10 IAMRECVR JCL EXAMPLES

REBUILD A FILE WITH RECOVERED DATA

4. The user has been getting data checks against an IAM file. IAMRECVR will be run to extract as many valid data records from the file as possible and write this file to tape. The damaged IAM file will be deleted by IDCAMS and a new file DEFINED on another volume. The records obtained from the first step will then be REPROed to this newly defined IAM file.

```
//UNLOAD      EXEC  PGM=IAMRECVR
//SYSPRINT   DD    SYSOUT=A
//DISKIN     DD    DSN=some.i am. file,DISP=SHR
//TAPEOUT    DD    DSN=seqntl.copy.of.i am. file,
//           UNIT=TAPE,DISP=NEW
//SYSIN      DD    *
  RECOVER
/*
//RECREATE   EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=A
//SEQIAM     DD    DSN=seqntl.copy.of.i am. file,
                DISP=(OLD,KEEP)
//SYSIN      DD    *
  DELETE some.i am. file
  DEFINE CLUSTER(NAME(some.i am. file) -
                .....
                (attributes taken from the existing file)
                .....
  REPRO INFILE(SEQIAM)
        OUTDATASET(some.i am. file)
/*
```

INSERT RECOVERED RECORDS INTO A BACKUP FILE

5. The user has been getting data checks against an IAM file. A backup copy of the file exists and will be restored. In an attempt to create as current a file as possible, all valid data records will be extracted from the damaged file and placed into a sequential dataset. These valid records will then be applied to the restored backup file. As a final step, the data entry area will validate the contents of the updated file to insure all transactions issued since the backup was created are in place (record changes, deletions, etc.)

```
//RECOVER     EXEC  PGM=IAMRECVR
//SYSPRINT   DD    SYSOUT=A
//BADFILE    DD    DSN=fried.i am. file,DISP=SHR
//SEQFILE    DD    DSN=&&tempfile,DISP=(,PASS),
//           UNIT=SYSDA,SPACE=(TRK,(nnn,nnn))
//GOODFILE   DD    DSN=backup.i am. file,DISP=OLD
//SORTWK01   DD    SPACE=(CYL,(1)),UNIT=SYSDA
//SORTWK02   DD    SPACE=(CYL,(1)),UNIT=SYSDA
//SORTWK03   DD    SPACE=(CYL,(1)),UNIT=SYSDA
//SYSOUT     DD    SYSOUT=A
//SYSIN      DD    *
  RECOVER OUTPUTFILES=SEQ, SORT=IFREQ,
          FROMDDNAME=BADFILE, TODDNAME=SEQFILE
  APPLY   LOGDDNAME=SEQFILE, IAMDDNAME=GOODFILE
/*
```

IAMISPF — IAM'S ISPF INTERFACE

50.01 IAMISPF — IAM'S ISPF INTERFACE

50.01 IAMISPF — IAM'S ISPF INTERFACE

IAM now includes ISPF panels that can be used to perform many utility functions against IAM and most VSAM datasets. The panels also include a complete ISPF tutorial that describes how to use the panels, as well as a complete description of the key IAM OVERFLOW statistics and OVERRIDE parameters.

FEATURES The key features included in the IAM ISPF panels are:

- Allocation of new IAM datasets with full Override support.
- Allocation of new VSAM clusters.
- Allocation with an IAM or VSAM model dataset.
- Multi-Volume dataset support.
- SMS support.
- Deletion of Datasets, Clusters, Paths, and Alternate Indexes.
- Renaming of Datasets, Clusters, Paths, and Alternate Indexes.
- Copy/Move support of IAM datasets, VSAM ESDS's, and VSAM KSDS's.
- Full IAM dataset information.
- VSAM cluster information.
- Interactive execution of selected IAM utility functions.

50.01 IAMISPF — IAM'S ISPF INTERFACE

Figure 1 shows how the ISPF/PDF Primary Option Menu can be customized to include a selection for the IAM ISPF Dialog. By selecting option I, the user will be presented with the IAM Primary Option Menu.

FIGURE 1
ISPF/PDF
PRIMARY
OPTION MENU

```
----- ISPF/PDF PRIMARY OPTION MENU   ISPF 3.2 -----
OPTION  ==> _

      0  ISPF PARMs  - Specify terminal and user parameters      USERID  - MARK
      1  BROWSE     - Display source data or output listings    TIME     - 07:14
      2  EDIT       - Create or change source data             TERMINAL - 3278
      3  UTILITIES  - Perform utility functions                PF KEYS  - 24
      4  FOREGROUND - Invoke language processors                DATE     - 91/10/15
      5  BATCH      - Submit job for language processing        JULIAN   - 91.288
      6  COMMAND    - Enter TSO command or CLIST               APPLID   - ISR
      7  DIALOG TEST - Perform dialog testing
      8  LM UTILITIES- Library management utility functions
      9  IBM PRODUCTS- Additional IBM program development products
      I  IAM        - IAM utility functions
      C  CHANGES   - Display summary of changes for this release
      T  TUTORIAL   - Display information about ISPF/PDF
      X  EXIT       - Terminate ISPF using log and list defaults

Enter END command to terminate ISPF.
```

Figure 2 shows the IAM primary option menu. For most functions, the user can provide a dataset name and select the desired function.

FIGURE 2
IAM PRIMARY
OPTION MENU

```
----- IAM PRIMARY OPTION MENU -----
OPTION  ==> _

      I  - Allocate (DEFINE) a new IAM Dataset
      V  - Allocate (DEFINE) a new VSAM Cluster
      D  - Delete a Dataset, Cluster, Path, or Alternate Index
      C  - Copy an IAM Dataset, VSAM KSDS, or VSAM ESDS
      M  - Move an IAM Dataset, VSAM KSDS, or VSAM ESDS
      R  - Rename a Dataset, Cluster, Path, or Alternate Index
      U  - Invoke an IAM Utility program
      blank - Display IAM Dataset or VSAM Cluster Information

Enter dataset name (required for all options except U)
Dataset Name      ==>

Enter Model or New dsname (optional for options I and V, required for option R)
Model|Newname     ==>

Delete Confirmation ==> YES  Yes|No
```

50.01 IAMISPF — IAM'S ISPF INTERFACE

Figure 3 shows the IAM dataset define panel. This panel supports both SMS and nonSMS allocations. Additionally, most IAM overrides can be specified by providing the appropriate values in the fields on the right hand side of the panel. By specifying YES in the Multi-Volume Allocation field, an additional panel will be displayed where additional volumes can be entered.

FIGURE 3
DEFINE AN
IAM FILE

```

----- DEFINE AN IAM FILE -----
COMMAND ==> _
Data Set Name: IAMV.TEST.$IAMFILE

ALLOCATION                                     Multi-Volume Allocation ==> NO
Volume                                     ==>
SMS Storage Class ==>
SMS Data Class ==>
SMS Mgmt Class ==>
Cyls|Recs|Trks ==>
Primary Space ==>
Secondary Space ==>

ATTRIBUTES
KSDS|ESDS ==>
Max Recordsize ==>
Avg Recordsize ==>
Key Length ==>
Key Offset ==>
CI Size ==>
CI/CA Free % ==> /
Shareoption ==>

IAM OVERRIDES
ANYVOL Unit ==>
Blocking Factor ==> 1-15,>300
Overflow Records ==> 0-2000000
Prime Extension ==> 0-32767
Space Release ==> Yes|No
Data Compress ==> Yes|No
Record Format ==> F|V
Minbufno ==> 1-32
Maxbufno ==> 1-32
Ocoreo % ==> 0-100
Ocorex % ==> 10-100
PSEUDORBA (ESDS) ==> Yes|No
RETENTION
DAYS ==> 0-9999
EXPIRATION DATE ==> YYYY.DDD
    
```

Figure 4 shows the IAM file characteristics panel. Full file information is provided in addition to a full Overflow analysis.

For multi-volume datasets, a secondary panel will be displayed which contains all the volumes the dataset has been allocated on.

FIGURE 4
IAM FILE
CHARAC-
TERISTICS

```

----- IAM FILE CHARACTERISTICS -----
COMMAND ==> _
Dataset Name: IAMV.MARK$IAM.IAMFILE1

Record Length:      1100   Volume:      IDPLB5   Creation:      1991.050
Record Format:  VARIABLE   Device Type:    3380   Expiration:    0000.000
Key Length:         8     Tracks in use:  447   Last Reference: 1991.287
Key Offset:         4     Block Size:    11476  Records:      40378
Dataset Type:      KSDS   Blocking Factor:  4   Deletes:      0
Share Option:      2     Alloc Type:    CYLINDERS  Inserts:      0
Release:          NO     Primary Alloc:  30   Updates:      0
Storage Class:    Secondary Alloc:  5   Minimum Buffers:
Data Class:       Compressed Keys:  YES  Maximum Buffers:
Mgmt Class:       Compressed Data:  NO   Key Storage:   10692

----- IAM OVERFLOW CHARACTERISTICS -----

Records Req:      900   PE Blocks Req:  3   Overflow Open %:  10
Records Alloc:   900   PE Blocks Alloc: 4   Overflow Full %:  10
Blocks Alloc:    90    PE Blocks Used:  0   CI% Requested:   0
Records Used:    0     PE Blocks Empty: 4   CI% Actual:      0
Records Empty:   900   % PE Available: 100
% Available:     100
    
```

50.01 IAMISPF — IAM'S ISPF INTERFACE

Figure 5 shows the VSAM cluster define panel. This panel can be used to define a VSAM ESDS, KSDS, LDS, or RRDS. To define a multi-volume cluster, specify YES in the Multi-volume Allocation field. An additional panel will be displayed where the additional volumes can be entered.

FIGURE 5

DEFINE A VSAM CLUSTER

```

----- DEFINE A VSAM CLUSTER -----
COMMAND ===> _

Cluster Name : IAMV.TEST.VSAMFILE
Data Name ===> 'IAMV.TEST.VSAMFILE.DATA'

ALLOCATION                                     Multi-Volume Allocation ===> NO
Volume                                     ===>
SMS Storage Class                         ===>
SMS Data Class                           ===>
SMS Mgmt Class                           ===>
Cyls|Recs|Trks                           ===>
Primary Space                             ===>
Secondary Space                           ===>
ATTRIBUTES
ESDS|KSDS|RRDS|LDS                       ===>
Max Recordsize                           ===>
Avg Recordsize                            ===>
Key Length                                ===>
Key Offset                                ===>
CI Size                                   ===>
CI Freespace %                            ===>
CA Freespace %                            ===>

MISC. ATTRIBUTES
Owner                                     ===>
X Region Share                           ===>          1-4
X System Share                            ===>          3-4
Bufferspace                              ===>
Reuse                                     ===>          Yes|No
Speed                                     ===>          Yes|No
Spanned                                   ===>          Yes|No
Erase                                     ===>          Yes|No

RETENTION
DAYS                                     ===>          0-9999
EXPIRATION DATE                          ===>          YYYY.DDD
    
```

Figure 6 shows the additional panel that will be displayed if the VSAM cluster being defined is a KSDS. This panel can be used to provide index specific parameters.

FIGURE 6

INDEX ALLOCATION FOR A KSDS

```

----- INDEX ALLOCATION FOR A VSAM KSDS -----
COMMAND ===> _

Cluster Name : IAMV.TEST.VSAMFILE
Data Name   : IAMV.TEST.VSAMFILE.DATA
Index Name  ===> 'IAMV.TEST.VSAMFILE.INDEX'

ALLOCATION
Volume(s)   ===> _____
Cyls|Recs|Trks ===> _____
Primary Space ===> _____
Secondary Space ===> _____

ATTRIBUTES
CI Size     ===>
Imbed       ===>          Yes|No
Replicate   ===>          Yes|No
    
```

50.01 IAMISPF — IAM'S ISPF INTERFACE

Figure 7 shows the VSAM data component information panel.

This is a condensed display of the most relevant fields from the IDCAMS LISTCAT command. For multi-volume data components, an additional panel will be displayed which contains all the volumes the data component has been allocated on.

FIGURE 7
VSAM DATA COMPONENT INFORMATION

```

----- VSAM DATA COMPONENT INFORMATION -----
COMMAND ==> _

Cluster Name: IAMV.MARKVSAM.CLUSTER1
Data Name:   IAMV.MARKVSAM.CLUSTER1.DATA

ATTRIBUTES                ALLOCATION                STATISTICS
Dataset Type:           KSDS      Volume:           SCR083      Records:           616
CI Size:                22528    Device Type:           3380      Deletes:           0
CI's per CA:            28      Alloc Type:           CYLINDERS    Inserts:           0
Avg Recordsize:        1020    Primary Alloc:         2      Updates:           0
Max Recordsize:        1020    Secondary Alloc:       1      Retrievals:        3696
Key Length:             4      Tracks Alloc:          30     Excps:             15
Key Offset:             8      Tracks Used:           15     CI Splits:         0
Cross Region Shropt:    2      Extents:               1      CA Splits:         0
Cross System Shropt:    3      Storage Class
Bufferspace:            47104    Data Class:
Reuse:                  YES     Mgmt Class:
Speed:                  YES
Spanned:                NO     CLUSTER HISTORY
Erase:                  NO     Owner:
                               Creation:           1991.274
                               Expiration:         0000.000
FREE SPACE
CI Freespace %:         0
CA Freespace %:         0
Bytes:                  630784
Tracks:                 15

```

Figure 8 shows the VSAM index component information panel. This panel will only be displayed for VSAM KSDS's. This is condensed display of the most relevant fields from the IDCAMS LISTCAT command. For multi-volume data components, an additional panel will be displayed which contains all the volumes the index component has been allocated on.

FIGURE 8
VSAM INDEX COMPONENT INFORMATION

```

----- VSAM INDEX COMPONENT INFORMATION -----
COMMAND ==> _

Cluster Name: IAMV.MARKVSAM.CLUSTER1
Index Name:   IAMV.MARKVSAM.CLUSTER1.INDEX

ATTRIBUTES                ALLOCATION                STATISTICS
CI Size:                2048    Volume:           SCR083      Records:           1
CI's per CA:            18     Device Type:           3380      Deletes:           0
Max Recordsize:        2041    Alloc Type:           TRACKS     Inserts:           0
Key Length:             4      Primary Alloc:         2      Updates:           0
Key Offset:             8      Secondary Alloc:       1      Retrievals:        0
Imbed:                  YES     Tracks Alloc:         2      Excps:             11
Replicate:              NO     Tracks Used:           0      CI Splits:         0
                               Seq-Set Tracks:       1      CA Splits:         0
CLUSTER HISTORY          Total Used:           1      Entries/Section:   5
Owner:                  Extents:              2      Index Levels:       1
Creation:               1991.274  Storage Class:
Expiration:              0000.000  Data Class:
                               Mgmt Class:
                               Seq-Set RBA:         73728
                               Hi-Level RBA:        73728

```

50.01 IAMISPF — IAM'S ISPF INTERFACE

Figure 9 shows the copy panel. If the target dataset does not exist, you may specify YES in the define target dataset field and specify the target dataset type. The define panel will then be displayed with all fields filled in using information from the source dataset as a model.

FIGURE 9
COPY AN IAM DATASET, VSAM ESDS OR KSDS

```
----- COPY AN IAM DATASET, VSAM ESDS, OR VSAM KSDS -----  
COMMAND ===> _  
  
Source IAM Dataset Name:   IAMV.MARK$IAM.IAMFILE1  
  
Enter target dataset or cluster name  
Target Dataset Name ===>  
  
Define Target Dataset ===> NO   (Yes|No   - If target does not exist)  
Target Dataset Type   ===>     (IAM|VSAM - If target is to be defined)
```

Figure 10 shows the move panel. The move panel is basically the same as the copy panel with one additional option. After the move is complete, you have the option of renaming the target dataset using the source dataset name as the newname.

FIGURE 10
MOVE AN IAM DATASET, VSAM ESDS or KSDS

```
----- MOVE AN IAM DATASET, VSAM ESDS, OR VSAM KSDS -----  
COMMAND ===> _  
  
Source IAM Dataset Name:   IAMV.MARK$IAM.IAMFILE1  
  
Enter target dataset or cluster name  
Target Dataset Name ===>  
Rename after Move   ===> NO   (Yes|No   - To rename target to source dsn)  
  
Define Target Dataset ===> NO   (Yes|No   - If target does not exist)  
Target Dataset Type   ===>     (IAM|VSAM - If target is to be defined)  
  
Warning - Source dataset will be deleted after move !!!
```

Figure 11 shows the IAM utility selection panel. Since all the IAM utilities are really batch utilities, the additional panels simply build and pass an in storage command stream to the requested utility. The utility output can be directed to either a temporary print dataset or a new print dataset. ISPF BROWSE is called to view the utility print output after the utility executes.

FIGURE 11
IAM UTILITY PROGRAM SELECTION MENU

```
----- IAM UTILITY PROGRAM SELECTION MENU -----  
OPTION ===> _  
  
S Specify utility print dataset allocation parameters  
  
R IAMRECVR - IAM file diagnostic and recovery utility  
V IAMSTATS - IAM VSAM Interface (VIF) module information  
Z IAMZAPOP - IAM options table utility
```

50.01 IAMISPF — IAM'S ISPF INTERFACE

Figure 12 shows the utility print dataset allocation parms panel. The utility print datasets will be allocated using the values specified in this panel. For most utilities, the default parameters should be sufficient. However, if the IAMRECVR PRINT function is used, you may wish to use this panel to provide additional space since this function can generate a significant amount of output depending on the options selected.

FIGURE 12
IAM UTILITY
PRINT
DATASET
ALLOCATION
PANEL

```
----- IAM UTILITY PRINT DATASET ALLOCATION PANEL -----  
COMMAND ===> _  
  
Blks|Cyls|Trks   ===> CYLINDERS  
Primary Space   ===> 1  
Secondary Space ===> 1  
  
Note - If blocks is specified, an average blksize of 6171 will be used
```

Figure 13 shows the IAMRECVR function selection panel. For execution under ISPF, the DIAGNOSE, LIST, and PRINT functions are supported. The other IAMRECVR functions are only supported in batch.

FIGURE 13
IAMRECVR
FUNCTION
SELECTION
MENU

```
----- IAMRECVR FUNCTION SELECTION MENU -----  
OPTION ===> _  
  
D  DIAGNOSE - Scan an IAM file for errors  
L  LIST     - List file characteristics  
P  PRINT    - Dump selected areas or blocks from an IAM file
```

Figure 14 shows the IAMRECVR DIAGNOSE options menu. DIAGNOSE can be used to validate the structural integrity of an IAM file.

As with all of the utility option panels, if you leave the print DSN field blank, the print output will be directed to a temporary print dataset. If you provide a print dataset name, the name must be new name. The print dataset will be allocated for you.

FIGURE 14
IAMRECVR
DIAGNOSE
OPTIONS
MENU

```
----- IAMRECVR - DIAGNOSE OPTIONS MENU -----  
COMMAND ===> _  
  
Dataset Name ===> 'IAMV.MARK$IAM.IAMFILE1'  
Print DSN     ===>                                     (optional)  
  
BSAM          ===> NO          (Yes|No - Use BSAM instead of EXCP)
```

Figure 15 shows the IAMRECVR LIST OPTIONS menu. The LIST function produces output similar to the IAMPRINT report provided by the IDCAMS LISTCAT command ([see section 10.07 for a sample IAMPRINT report](#)).

FIGURE 15
IAMRECVR
LIST OPTIONS
MENU

```
----- IAMRECVR - LIST OPTIONS MENU -----  
COMMAND ===> _  
  
Dataset Name ===> 'IAMV.MARK$IAM.IAMFILE1'  
Print DSN     ===>                                     (optional)
```

50.01 IAMISPF — IAM'S ISPF INTERFACE

Figure 16 shows the IAMRECVR PRINT OPTIONS menu. The PRINT function produces a formatted dump of areas you specify within an IAM file.

FIGURE 16

IAMRECVR
PRINT
OPTIONS
MENU

```

----- IAMRECVR - PRINT OPTIONS MENU -----
COMMAND ==> _

Dataset Name ==> 'IAMV.MARK$IAM.IAMFILE1'
Print DSN      ==>                                     (optional)
BSAM          ==> NO          (Yes|No - Use BSAM instead of EXCP)
IDPINQ       ==> NO          (Yes|No - Print IAM control block)
PRINT LENGTH ==>             (32-32768 - Maximum data to print per block)
MAX BLOCKS   ==>             (Maximum blocks to print from each area)

ALL BLOCKS   ==>             (Yes|No - Print blocks from all areas)

OR

DATA         ==>             (Yes|No - Print data blocks)
KEYS         ==>             (Yes|No - Print key blocks)
OVERFLOW     ==>             (Yes|No - Print overflow blocks)
PE           ==>             (Yes|No - Print prime extension blocks)

OR

FROM BLOCK   ==>             (First block number to print)
TO BLOCK     ==>             (Last block number to print)
    
```

Figure 17 shows the IAMZAPOP function selection menu. The ZAP function is only supported in batch and not under ISPF.

The LIST function produces a report of the options from the IAM Global Options Table in the IAM library specified.

The RESET function is used to reset the IAM Global Options Table in the library specified back to its distributed default values.

The COPY function will copy the IAM Global Options Table settings from the IAM Library Name specified to the IAM Global Options Table in the COPY to DSNAME specified. Only the changed options will be copied to the new table.

The AUDIT function will list all of the options in the IAM Global Options Table that have been changed from their distributed default values.

You may provide the name of an optional new print dataset. If this field is left blank, a temporary dataset will be used.

FIGURE 17

IAMZAPOP
FUNCTION
SELECTION
MENU

```

----- IAMZAPOP FUNCTION SELECTION MENU -----
OPTION ==> _

L LIST      - List IAM options
R RESET     - Reset IAM options
C COPY      - COPY IAM options
A AUDIT     - AUDIT IAM options

IAM Library Name ==>
Print DSNAME     ==>                                     (optional)
COPY to DSNAME   ==>                                     (for copy)
    
```

50.01 IAMISPF — IAM'S ISPF INTERFACE

Figure 18 shows the results of the IAMSTATS request.

The IAM version shows the version of the IAM options table that was found in core. The VIF key is an internal indicator used to allow multiple versions of IAM to run concurrently and not interfere with each other. There is a different VIF key for each IAM version, and sometimes a different key for different releases within the same version.

The IAM Vector Table status can be INSTALLED or NOT INSTALLED. If the status is NOT INSTALLED, the job used at your site to start IAM has NOT been run. The status of IAM itself can be ACTIVE or NOT ACTIVE. If the vector table is NOT INSTALLED, the status will always be NOT ACTIVE. The status may also be NOT ACTIVE if IAM has been started, then stopped. Stopping IAM does not remove the IAM vector table from memory.

The rest of the display lists all the IAM modules that are resident, where they are currently loaded, their length and their version/level information if available.

If there are multiple versions and/or levels of IAM active, hitting the ENTER key will display information about each IAM vector table found.

FIGURE 18
RESULTS OF
IAMSTATS
REQUEST

```
----- IAM VSAM Interface (VIF) Statistics -----  
OPTION ===> _  
  
IAM Version: 6.2  
VIFKEY: FFFE0020  
  
IAM Vector Table is INSTALLED  
and IAM is ACTIVE  
  
Module Address Len Version  
IAMVECTB 00CCC0E8 000628  
IAM0192A 00CC5B10 0014F0 6.2/00T  
IAM0200T 00CC5098 000A78 6.2/05T  
IAM0231T 00CC8508 000AF8 6.2/00T  
IGG0CLA0 00CC8000 000508  
IAM0CLA0 80CC48F8 000708 6.2/05T  
IAM026DU 80CC18B8 002748 6.2/00T  
IAMCSMF 00CC4318 0005E0 6.2/05T  
  
----- Figure 18 -----
```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

IAMM

INNOVATION ACCESS METHOD

ESDS SUPPORT

THIS PAGE WAS INTENTIONALLY LEFT BLANK

SECTION 60: IAM ESDS SUPPORT

60.00 IAM ESDS SUPPORT

60.10 IAM ESDS SUPPORT OVERVIEW

This section will explain how to use IAM files in place of VSAM ESDS files, including how to utilize the special features that IAM brings to ESDS files, along with any special considerations when converting ESDS files to the IAM format. As a general guideline, all that is normally REQUIRED to be changed to use IAM is some indication on the IDCAMS DEFINE statement that the file should be an IAM file. For example, this might be done by either changing the datasetname to have the characters \$IAM in it, or adding an OWNER(\$IAM) keyword to the DEFINE. There might be some circumstances when there will be a need to change other DEFINE parameters, or possibly JCL, even though the IAM product has gone to great lengths to minimize those circumstances. Additionally, there might be some JCL changes in order to utilize some of the special features that are unique to IAM.

IAM Version 6.2 brings the IAM advantages of reduced DASD space utilization, reduced I/O, lower elapsed time for batch jobs, and improved response time for online systems to applications that utilize VSAM ESDS type of files. These advantages are the result of IAM's automatic data compression capability, IAM's file structure, which allows for more effective utilization of DASD space, and the IAM Real Time Tuning buffer management capabilities, which have been enhanced for ESDS files with IAM Version 6.2. Additionally, applications which do not have a dependency on true VSAM RBA (Relative Byte Addressess) for the data records can use an IAM option to have IAM ESDS files exceed the VSAM limitation of 4 Gigabytes. The theoretical maximum size of an IAM file is approximately 860 gigabytes of stored data, the actual user data provided may be larger due to the benefits of IAM data compression. (This estimate of maximum file size is based on using a 3390. The maximum size of an IAM file is 16,711,425 tracks contained within 255 extents, each being 65,535 tracks.)

IAM ESDS files are indexed files, with a key added to each record that is invisible to any application program. This "key" is a value determined by IAM as what would have been the VSAM RBA (Relative Byte Address) of the record if it had been in a real VSAM ESDS. This 4 byte RBA value that VSA uses to identify records in an ESDS is the cause of the 4 gigabyte file size restriction that VSAM has. IAM can support larger size ESDS files, providing that the application software using the file does not have a dependency on the actual VSAM RBA. This capability can be used by specifying the PSEUDORBA keyword on an IAM CREATE override when the file is defined or loaded. When this option is utilized, the IAM "key" for the file is in fact a relative record number, which will be passed back to the application as the record's RBA.

A number of enhancements were made to IAM for ESDS files to make the conversion process even easier for ESDS type of files. IAM ESDS files require no preformatted overflow areas, as are required for IAM KSDS files. The size of IAM ESDS files can grow, and acquire SECONDARY EXTENTS as needed, to handle the addition of records, as well as to handle situations where overflow is needed due to record size changing because of different compression after the update. Please note that due to IAM files being stored as a non-VSAM file type on DASD, that there is an MVS limitation of only 16 extents per dasd volume, and a maximum of 65,535 tracks per volume. These MVS limitations may require some adjustment in the space allocation quantities for ESDS files when they are converted to IAM, particularly for ESDS files that take more than 16 extents per volume as VSAM ESDS files. IAM ESDS files can be left in LSR buffer pools. While IAM will not utilize the LSR buffers, it does provide support for the LSR macros, and the UPAD exit used for LSR requestors. IAM ESDS files default to the IAM variable length data compressed format, eliminating problems where the IDCAMS define indicates that the files contain fixed length records, and automatically providing the benefits of reduced DASD space utilization and reduced I/O that data compression provides.

60.20 IAM ESDS FEATURES

60.20 IAM ESDS FEATURES

IAM has many features and capabilities that rise above what is normally expected of an access method. Many of these features are automatically provided, such as the IAM data compression capability and the automatic DASD space release of unused DASD space after the file is initially loaded. IAM has various other features, which can be controlled on an installation wide basis through the IAM Global Options facility, as described in [section 91](#) of the IAM Users Guide. Control of these features is also offered on a file by file basis through the use of the IAM Override facility.

REAL TIME TUNING

IAM is continually monitoring on a file by file basis, the amount and type of I/O activity that is taking place for the file. Based on this data, IAM will dynamically adjust the number of buffers in use for the file along with the buffering technique. The IAM Real Time Tuning considers several factors in controlling the I/O and buffering being used, including frequency of I/O requests, type of requests, such as random, sequential, input, output, batch or online (asynchronous) requests, success or lack of success in buffering decisions made previously. The intent of the IAM Real Time Tuning is to reduce the amount of physical I/O required to process the file, based on the type of I/O requests that are being processed.

IAM Real Time Tuning works within a specifyable number of buffers, by default using 1 as the minimum, referred to as MINBUFNO, and calculating the number of buffers that will fit in 256K as the maximum, referred to as MAXBUFNO. When an IAM file is opened, it generally starts out with 4 buffers, and will then be adjusted from that point based on the I/O activity. These values can be overridden globally in the IAM Global Options Table, or on a file basis by the use of CREATE overrides when the file is defined, or an ACCESS override when the file is accessed. MAXBUFNO can also be increased through the use of the VSAM BUFND parameter, either in JCL, or in the ACB.

The IAM Real Time Tuning and buffer management for ESDS files has the following capabilities:

1. Manage up to 255 buffers per file.
2. Sequential read ahead (up to 1 cyl) with I/O and processing overlap.
3. On sequential updates and inserts, accumulate updated blocks, and write multiple blocks per I/O, again with processing overlap.
4. Supports multiple concurrent active I/O requests to improve responsiveness on multivolume files with heavy I/O activity.
5. Supports ESCON channels, and the ECKD channel command architecture.
6. Supports the sequential access caching on 3990-3 and 3990-6 when processing sequential requests, and record level caching on the 3990-6 for random I/O requests.

IAM does not lock buffers to requests except when there is an actual I/O in progress. This design eliminates the problems with Exclusive Control on control intervals with VSAM files that can be troublesome especially when using LSR buffering. IAM features record level lockout under a single ACB, unlike the Control Interval locking that VSAM provides. These capabilities improve response time to online transactions due to elimination of delays in processing that can occur with VSAM. Additionally, the IAM support for ESDS includes a significant improvement in handling concurrent I/O requests, which will improve response time for online transaction processing as well.

60.20 IAM ESDS FEATURES

SECONDARY EXTENTS

The IAM ESDS support now includes the capability for IAM ESDS files to take additional extents as the file grows. The preformatted Independent Overflow area and Prime Extension Area that are an integral part of IAM KSDS files have been eliminated from IAM ESDS files. File out of space conditions will now only occur when IAM is not able to obtain additional DASD space. For the ESDS files, it is anticipated that the largest need for additional space will be to accommodate records being added to the file. However, in a data compressed file, there may also be a need for additional space when a record is updated and is not as compressible as it had been, requiring more space. If the record can not fit within the block, it will then be moved to an overflow block in the extended area.

Because IAM files reside on DASD as NONVSAM, they are restricted by MVS to 16 extents per volume, and 65,535 tracks per volume. Because of this limitation, the space values on the DEFINES of IAM ESDS files may have to be increased, because IAM files are not able to take as many extents per volume as a VSAM ESDS. The necessity for this can be determined from either a LISTCAT output, or from the IAMSMFVS SIZE report, which has space allocated and the number of extents the dataset is in.

EXCEEDING 4 GIGABYTE SIZE FILES

Because of the VSAM architecture of using a four byte Relative Byte Address (RBA) for identifying records within a VSAM dataset, the VSAM file size is limited to 4 gigabytes of data. For complete compatibility with VSAM ESDS files, which is necessary for some applications, IAM has inherited this limitation. However, for application programs that do not have a dependency on the RBA architecture, IAM can exceed the 4 gigabyte file size limitation. This is done by using a relative record number in place of an relative byte address for the data records. This will allow an IAM ESDS file to handle a maximum of four billion records. This optional capability is activated on a file by file basis through the use of the PSEUDORBA keyword on the IAM CREATE Override statement, which can be specified when the file is defined or loaded.

DYNAMIC TABLING OF RECORDS IN STORAGE

IAM has a Dynamic Tabling facility, which will store records in a table in storage, for use by random reads. This capability requires no application programming changes to use, and is intended for use primarily by online transaction processing systems that have a heavy random read activity, where in effect buffering records in storage may be more effective than buffering blocks of data. This facility is activated on a file by file basis using the ACCESS DYNCORE override, where the amount of storage to be used for the table is specified in K bytes. As records are randomly read, IAM will first search the in storage table to see if the record is there, and if it is, the record will be returned without I/O. However, if the record is not in the table, it will be read into a buffer, if the block containing the record is not already in a buffer, and the requested record will be placed into the table, possibly replacing another record if the table has reached it's capacity. The IAMINFO report contains statistics on how many records were contained in the dynamic table, and how many times records were retrieved from the table, so the effectiveness can be determined. Selecting the best size for the table must be done by trying various sizes, to see where the most benefit is obtained. Using this capability, small files could be kept entirely in storage.

60.20 IAM ESDS FEATURES

IAMINFO RUN TIME REPORTING AND SMF RECORDS

IAM will optionally produce a report every time an IAM file is closed, containing various statistics about what type of requests were processed, statistics on buffer use, DYNCORE use, storage use, EXCP's, and the file attributes. This feature is activated by adding an IAMINFO DD card to any job step using IAM files. While normally these reports are written out to SYSOUT, the data can also be written out to a dataset.

The IAMINFO reports are very useful to answer questions about the file and processing that was done, so the use of this facility is highly recommended. However, if there is too much effort required to add IAMINFO to all the jobs using IAM, an alternative is to collect the optional IAM SMF records. Then the IAMINFO reports can be produced using the IAMSMF program's IAMINFO command. There is also summary type of reports from the IAMSMFVS program, plus an installation can write their own program(s) to report on IAM file usage in their own unique way. Refer to the IAM GLOBAL OPTIONS topic in the INSTALLATION section of the manual for instructions on enabling this feature.

IAMOVRID — MODIFYING IAM PARAMETERS

The use of various IAM features is activated or modified through the use of the IAM Override facility. Basically, this facility provides a way to specify parameters for IAM to use for either selected and/or all of the IAM files used within a job step. IAM Overrides are specified on control card input passed either as instream card images via DD *, or as an input card image dataset. There are CREATE overrides, which are relevant to defining and loading IAM files, and the ACCESS overrides, which are used for general file access type of programs.

For example, the CREATE override can be used to specify that some freespace is to be left in each block, which is useful for files that are frequently updated, to revise the file load buffering, to alter the file's blocksize, or to disable the automatic space release feature. The ACCESS override can be used to alter the number of buffers used, to use the IAM data record tabling feature (DYNCORE), or to specify IAM request tracing. For complete information on using the IAM Override capability, please refer to the IAM Overrides section of the IAM manual.

SECTION 61: USING IAM ESDS FILES

61.00 USING IAM ESDS FILES

61.10 DEFINING IAM ESDS FILES WITH IDCAMS

DEFINING IAM ESDS FILES Using IAM ESDS files in place of VSAM ESDS is the same as any IAM file definition. Just add the \$IAM parameter to either the OWNER, file name, or the DATACLAS or STORCLAS SMS fields in the DEFINE procedure. The next time the file is DEFINED, it will become an IAM file.

CONVERTING A DEFINE PROCEDURE 1. The following example shows how an IDCAMS command originally written to DEFINE a VSAM ESDS is changed to create an IAM ESDS.

Original VSAM DEFINE	IAM DEFINE
//S1DEF01 EXEC PGM=IDCAMS	//S1DEF01 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*	//SYSPRINT DD SYSOUT=*
//SYSIN DD *	//SYSIN DD *
DEFINE CLUSTER -	DEFINE CLUSTER -
(NAME(PAYROLL.MASTER) -	(NAME(PAYROLL.MASTER) -
VOL(SAMP01) -	VOL(SAMP01) -
NONINDEXED -	NONINDEXED -
TRK(45 10) -	TRK(45 10) -
RECORDSIZE(256 256) -	ADD -> OWNER(\$IAM) - <- -ADD
SHAREOPTIONS(2,3)	RECORDSIZE(256 256) -
	SHAREOPTIONS(2,3)

DEFINE USING \$IAM IN THE NAME 2. The following example shows an IDCAMS DEFINE of an IAM file specifying '\$IAM' in the data set name. The file will be allocated with 100 cylinders on volume 'SAMP01'. The file contains variable length records with a maximum record size of 480 bytes and average size of 250.

```
//DEFINE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
DEFINE CLUSTER -
    (NAME(PAYROLL.$IAM.MASTER) -
    VOL(SAMP01) -
    NONINDEXED -
    CYL(100 20) -
    RECORDSIZE(250 480) -
    SHAREOPTIONS(2,3))
```

NOTE: With ESDS files, VSAM ignores any freespace specifications on the IDCAMS DEFINE, so they are not available to IAM to consider even if they are coded on the IDCAMS DEFINE statement. If the file is heavily updated, then specification of some Integrated overflow, which is very similar to CI% Freespace on KSDS files, is recommended.

DEFINE WITH AN OVERRIDE 3. The following example shows an IDCAMS DEFINE of an IAM file using the OWNER parameter. Since this file is updated, an OVERRIDE statement is used to change the INTEGRATED OVERFLOW value to 5%, which will allow some room for records which do not compress as well after being updated.

```
//DEFINE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//IAMOVRID DD *
***** THE FOLLOWING STATEMENT WILL OVERRIDE *****
***** THE DEFAULT INTEGRATED OVERFLOW VALUE *****
CREATE DD=&ALLDD,I=5
/*
//SYSIN DD *
DEFINE CLUSTER -
    (NAME(MASTER.TABLE.CLUSTER) -
    OWNER($IAM) -
    CYLINDERS(600 100) -
    NONINDEXED -
    RECORDSIZE(80 200) -
    VOLUMES(SAMP01) )
/*
```

61.00 USING IAM ESDS FILES

DEFINE FILE FOR OVER 4 GIG 4. The following example shows an IDCAMS DEFINE of an IAM file with an override of PSEUDORBA so that the amount of data can exceed 4 gigabytes.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//IAMOVRID   DD    *
      CREATE   DD=&ALLDD,PSEUDORBA
/*
//SYSIN      DD    *
      DEFINE   CLUSTER -
              (NAME(MASTER.TABLE.CLUSTER) -
               OWNER($IAM) -
               CYLINDERS(1100 1100) -
               NONINDEXED -
               RECORDSIZE(200 4080) -
               VOLUMES(V33901 V33902 V33903 V33904) )
/*
```

DEFINE WITH MULTIPLE OVERRIDE 5. The following example shows an IDCAMS DEFINE of multiple IAM files, each with a different override. This is accomplished through the use of the FILE parameter on the DEFINE, which matches the DDNAME parameter on the CREATE override card. Note that as long as the file is an IAM file, there is no need for an actual DD card of the specified name to be in the JCL, although if the file is switched to VSAM, then it is required.

```
//DEFINE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//IAMOVRID   DD    *
***** THE FOLLOWING STATEMENT WILL OVERRIDE *****
***** THE DEFAULT INTEGRATED OVERFLOW VALUE *****
      CREATE   DD=MASTER,I=5
      CREATE   DD=BIGGER,I=10
/*
//SYSIN      DD    *
      DEFINE   CLUSTER -
              (NAME(MASTER.TABLE.CLUSTER) -
               OWNER($IAM) -
               FILE(MASTER) -           <--- INDICATES OVERRIDE DD
               CYLINDERS(100 10) -
               NONINDEXED -
               RECORDSIZE(60 100) -
               VOLUMES(SAMP01) )
      DEFINE   CLUSTER -
              (NAME(BIGGER.TABLE.CLUSTER) -
               OWNER($IAM) -
               FILE(BIGGER) -           <--- INDICATES OVERRIDE DD
               CYLINDERS(600 100) -
               NONINDEXED -
               RECORDSIZE(80 200) -
               VOLUMES(SAMP01) )
/*
```

61.20 DELETING IAM ESDS FILES WITH IDCAMS

61.20 DELETING IAM ESDS FILES WITH IDCAMS

DELETING IAM ESDS FILES IAM ESDS files can be deleted just as the VSAM ESDS files. No changes are required to the IDCAMS DELETE control statements, as long as the IAM VIF interface is active. The IAM VIF interface will intercept the DELETE CLUSTER request, and if the file is IAM, convert it to a DELETE NONVSAM request. The PURGE and ERASE keywords of the DELETE command are fully supported for IAM files.

One difference in the DELETE processing for IAM files is that due to MVS dataset integrity control, IAM files, like any other MVS non-VSAM dataset, can not be deleted if it is allocated to any other job or user. This is a difference between IAM and VSAM, as VSAM datasets can be deleted even if they are allocated to another job or user. Unless the space allocation for an IAM ESDS file needs to be changed, there should not generally be a need to delete IAM files to reload them, as by default IAM files are always reusable.

DELETE AN IAM FILE 1. The following example shows an IDCAMS DELETE if an IAM file using the CLUSTER attribute.

```
//DELETE      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=*
//SYSIN      DD    *
DELETE      MASTER.TABLE.CLUSTER -
           CLUSTER
/*
```

61.30 LISTCAT OF IAM ESDS FILES WITH IDCAMS

61.30 LISTCAT OF IAM ESDS FILES WITH IDCAMS

LISTCAT IAM FILES A LISTCAT of an IAM file will result in IDCAMS producing output for a nonVSAM entry type. When the IAM VIF interface is active, IAM will produce a separate report of the file attributes. For batch jobs, when SYSPRINT is allocated to SYSOUT, IAM will dynamically allocate a file, IAMPRINT, to SYSOUT and produce a report for each IAM file that the LISTCAT processes, providing that the ALL keyword is specified. If the LISTCAT is issued by a TSO user, again with ALL specified, the report is routed directly to the user's terminal. If the user desires the IAMPRINT output to go somewhere else, for example to a specific SYSOUT class, or to a dataset, this can be done by the user providing an IAMPRINT DD card, indicating where the report should go.

LISTCAT AN IAM FILE 1. The following example shows an IDCAMS LISTCAT of an IAM file using the ALL keyword, which will result in IAM allocating a SYSOUT file to a DD name of IAMPRINT, and producing a report to that SYSOUT dataset.

```
//DELETE      EXEC  PGM=IDCAM5
//SYSPRINT    DD    SYSOUT=*
//SYSIN       DD    *
              LISTCAT ENT(IAMV.ESD62001.CLUSTER) ALL
/*
```

```
IAM CATALOG INFORMATION SERVICE -- VER v.m/llp      DATE yy.ddd TIME hh:mm:ss
IAM FILE ANALYSIS - DSN=IAMV.ESD62001.CLUSTER
```

RECORD SIZE -- =	510	-	INTEGRATED OVERFLOW --- =	20	PER	-
BLOCK SIZE --- =	11476	-	EXTENDED OVERFLOW ----- =	726	RECS	-
BLOCK FACTOR - =	4	-	EXTENDED OVERFLOW ----- =	33	BLOCKS	-
FILE TYPE ---- =	ESDS	-	EXTENDED ALLOCATED ---- =	36	BLOCKS	-
DEVICE TYPE -- =	3380	-	EXTENDED USED ----- =	36	BLOCKS	-
VOLUME COUNT - =	1	-	SPACE USED ----- =	1875	TRACKS	-
VOLSER ----- =	SCR083	-	SPACE ALLOCATED ----- =	1875	TRACKS	-
TOTAL EXTENTS =	6	-	TOTAL SPACE ALLOCATED - =	1875	TRACKS	-
RELEASE ----- =	NO	-	SHARE OPTIONS ----- =	2		-
DATA COMPRESS =	ENABLED	-	INDEX COMPRESS ----- =	YES		-
TOTAL RECORDS =	200000	-	INSERTS ----- =	0		-
UPDATES ----- =	6500	-	DELETES ----- =	0		-
CI SIZE ----- =	4096	-	HIGH USED RBA ----- =	102400000		-
STORAGE REQUIRED FOR COMPRESSED KEY STRUCTURE =	16722					-
NUMBER OF IAM DATA BLOCKS ----- =	7459					-
EXTENDED HIGH ALLOCATED RBN ----- =	7499					-

```
IAMLISTC(v.m/llp ) PROCESSING COMPLETED
```

61.35 EXPLANATION OF LISTCAT OUTPUT

61.35 EXPLANATION OF LISTCAT OUTPUT

RECORD SIZE – This is the maximum record size for the file, as specified on the IDCAMS DEFINE for the dataset.

INTEGRATED OVERFLOW – This is the amount of free space, as a percentage, that is left free in each data block as the file is being loaded and extended, and is similar to VSAM CI% Freespace in a KSDS type of file. For IAM ESDS files, this free space is very useful when files are updated, because frequently updated records do not data compress as much as before the update, and therefore require more space.

BLOCK SIZE – The physical block size IAM is using for the file. The size selected is based on maximizing the data utilization of the device.

EXTENDED OVERFLOW – The number of records that have been moved to extended overflow blocks due to an increase in size after being updated. A large number of records in Extended Overflow indicate that the file should be reorganized, as overflow records cause increased virtual storage usage due to a record based index, and may cause performance degradation on sequential I/O operations. If there is frequently a large number of Extended Overflow records, then increasing Integrated Overflow should reduce the amount of Extended Overflow records.

BLOCK FACTOR – Indicates the number of blocks per track or the user requested block size from an IAMOVRID CREATE override.

EXTENDED OVERFLOW – The number of blocks being used to hold the the Extended Overflow records.

FILE TYPE – Indicates the type of VSAM file the IAM file is being used for. Values are ESDS, ESDS/P (when PSEUDORBA is being used), or KSDS.

EXTENDED ALLOCATED – The number of blocks that have been added to the end of the file since it was last loaded. These will include data blocks, overflow blocks, index blocks for the extended area, and available blocks to the end of the file.

DEVICE TYPE – Indicates the type of device that the dataset currently resides on.

EXTENDED USED – The number of blocks used beyond the original size of the file. Generally these blocks will reside in secondary extents that are obtained as needed.

VOLUME COUNT – The number of volumes contained in the catalog entry for the file.

SPACE USED – The amount of DASD space required for the file, which includes the Extended Allocated blocks.

VOLSER – Identifies the volume(s) to which the file is cataloged.

SPACE ALLOCATED – Indicates the amount of DASD space, in tracks, allocated on the volume identified to the left of this field.

TOTAL EXTENTS – The total number of extents allocated to the IAM file across all volumes the dataset is cataloged to.

TOTAL SPACE ALLOCATED – The sum of the space allocated for all volumes on which the file is cataloged.

RELEASE – Indicates whether or not the automatic space release feature will be used the next time the file is loaded or reorganized without being deleted and redefined. By default, files have excess DASD space released on the first load after a define.

SHARE OPTIONS – Corresponds to the VSAM cross region share option, as specified on the IDCAMS DEFINE.

DATA COMPRESS – Indicates whether or not the data records in the file are eligible for the IAM data compression.

INDEX COMPRESS – Indicates whether the file has a compressed index, which is used to reduce the virtual storage requirements. The index to files is automatically compressed when the index size exceeds 8,000 bytes, and is used only when there is a savings in storage from the noncompressed index format.

61.35 EXPLANATION OF LISTCAT OUTPUT

The following statistics depend on the file always being properly closed after being updated. If the dataset has not been properly closed, due to system failure, then these statistics may be incorrect.

TOTAL RECORDS – The total number of records contained in the dataset.

INSERTS – The number of records added to the dataset since it had last been loaded or reorganized.

UPDATES – The number of updates that have been performed on existing records.

DELETES – The number of records deleted. For ESDS files, this should always be 0, because VSAM does not support deleting records in an ESDS type of file.

CI SIZE – The simulated CI Size used for this dataset. Normally, this is specified on the IDCAMS DEFINE. However, if no size was specified, IAM determines one to use.

HIGH USED RBA – The RBA of the first empty CA in the dataset. (CA and CI are only logical entities, they do not actually exist for IAM files.

STORAGE REQUIRED FOR (COMPRESSED) KEY STRUCTURE – indicates the amount of virtual storage required for the index to the non-extended area of the file, which represents the data that was initially loaded into the dataset. COMPRESSED appears in the text of the message when a compressed index has been built to reduce the virtual storage requirements of the index. Unlike VSAM ESDS files, IAM ESDS files have an internal index, consisting of the assigned RBA (Relative Byte Address) of the last record in each data block.

NUMBER OF IAM DATA BLOCKS – Indicates the number of blocks in the dataset, preceding the index. This value is provided to be used when IAMRECVR has to be used to recover the dataset, and the self-describing record in the front of the file has been damaged. Specify this value for the MAXBLKS keyword of the RECOVER command only when necessary for the RECOVER process to work.

EXTENDED HIGH ALLOCATED RBN – Indicates the high allocated block number, to be used when it is necessary to run IAMRECVR to recover the file after it has been damaged. This value will be used for the XTENDEDHARBAN keyword of the recover command.

61.40 COPYING AND REORGANIZING IAM ESDS FILES WITH IDCAMS

61.40 COPYING AND REORGANIZING IAM ESDS FILES WITH IDCAMS

IAM files can be easily copied and reorganized using Innovation's FDRREORG utility, or the IDCAMS utility. Because IAM ESDS files are indexed files, they may require occasional reorganization if there have been substantial updates that cause the use of extended overflow blocks, or very significant quantity of records added to the file. FDRREORG or IDCAMS can be used to make a sequential copy of the ESDS on tape or dasd for backup, reorg, or other processing.

COPYING IAM ESDS FILES

IAM ESDS files can be copied into sequential datasets, or into other IAM or VSAM datasets, just as can be done with regular VSAM ESDS files. Existing jobs to copy VSAM ESDS files can be run generally without any changes with IAM ESDS files.

The following examples are for illustrative purposes only. Please refer to the appropriate product's documentation for complete information on using the product.

The following is an example of copying an IAM file to a sequential file on tape. The file on tape can be used for backup purposes, for transferring data to another site or, for other sequential processing. In this example, FDRREORG is used to make the sequential output file. FDRREORG will dynamically allocate the necessary files, based on the installation's global options and the control statement input.

```
//COPYFILE EXEC PGM=FDRREORG,REGION=6M
//SYSPRINT DD SYSOUT=*
//REORGPRPT DD SYSOUT=*
//REORGRPT DD SYSOUT=*
//IAMINFO DD SYSOUT=*
//SYSIN DD *
    REORG ALWAYSBACKUP, BACKUP=PERM, BACKUPUNIT=3480
    SELECT PCATDSN=master.$iam.cluster, NOREORG
/*
```

The following is an example of copying an IAM file to a sequential file on tape using IDCAMS. The file on tape can be used for backup purposes, for transferring data to another site or, for other sequential processing. A RECFM and BLKSIZE is specified for the sequential output file because the IDCAMS default is RECFM=U, and results in a single record per block, causing wasted media space.

```
//COPYFILE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//IAMINFO DD SYSOUT=*
//TAPEFILE DD DSN=SEQ.COPY.DATA,UNIT=3480,DISP=(,CATLG),
//          DCB=(RECFM=VB,BLKSIZE=32760)
//SYSIN DD *
    REPRO IDS(MASTER.$IAM.CLUSTER) OUTFILE(TAPEFILE)
/*
```

61.40 COPYING AND REORGANIZING IAM ESDS FILES WITH IDCAMS

REORGING IAM ESDS FILES

IAM ESDS files can be reorganized using FDRREORG or IDCAMS REPRO.

In this example, the IAM ESDS file is reorganized using FDRREORG, with the installation's default options. The reorg is done only if necessary, using the installation's default options. The sequential output copy of the file will be automatically kept if the reorg fails for any reason, and can be used to recover the file with the FDRREORG RECOVER command. If the REORG is successful, the default is that the sequential output copy is deleted.

```
//REORGFIL EXEC PGM=FDRREORG,REGION=6M
//SYSPRINT DD SYSOUT=*
//REORCPRT DD SYSOUT=*
//REORCRPT DD SYSOUT=*
//IAMINFO DD SYSOUT=*
//SYSIN DD *
REORG
SELECT PCATDSN=master.$iam.cluster
/*
```

In this example, the IAM ESDS file is copied to tape, then the file is reloaded from the tape. IAM overrides are used to minimize physical I/O. When doing a file reorganization, NEVER use a temporary data set for the sequential copy of the data, because if the reload fails for any reason, the data will be lost.

```
//REORG EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//IAMINFO DD SYSOUT=*
//IAMOVRID DD *
ACCESS DD=&ALLDD,MINBUFNO=16,MAXBUFNO=62
CREATE DD=&ALLDD,CRBUFOPT=MCYL
/*
//TAPEFILE DD DSN=SEQ.COPY.DATA,UNIT=3480,DISP=(,CATLG),
// DCB=(RECFM=VB,BLKSIZE=32760)
//SYSIN DD *
REPRO IDS(MASTER.$IAM.CLUSTER) OUTFILE(TAPEFILE)
REPRO ODS(MASTER.$IAM.CLUSTER) INFILE(TAPEFILE) REUSE
/*
```

An alternative to using a sequential file for intermediate placement of the data is to REPRO the file into a another IAM file, providing that there is sufficient DASD space available to do so. Then, if desired RENAMES can be done. This type of REORG will be a time saver, particularly if the file is large, and there is no need to have a sequential backup copy of the data.

61.50 JCL CONSIDERATIONS WITH IAM

61.50 JCL CONSIDERATIONS WITH IAM

In most circumstances, there are no JCL changes required to use IAM files in place of VSAM files. IAM does have certain capabilities that do require JCL changes to utilize. The IAM unique JCL which can be used is identified below, by DDNAME and the purpose or significance.

DDNAME	PURPOSE
IAMINFO	Identifies a sequential output dataset, which is normally SYSOUT, into which IAM produces a report every time an IAM file is closed in the job step which has this DD card. Use of this capability is highly recommended, and may be requested by the IAM technical support staff when trying to resolve various problems. An alternative to changing JCL is to collect the IAM SMF records. Then the IAMINFO reports can be produced from the SMF data using the IAMSMF program's IAMINFO command.
IAMPRINT	Identifies a sequential output dataset, which is normally SYSOUT, into which IAM produces a report whenever an IDCAMS LISTCAT ALL is done for an IAM file. IAM will normally dynamically allocate this file to the default SYSOUT class. The DD card can optionally be provided by the user to route the report to a dataset, or to a different SYSOUT class.
IAMDEBUG	When specified as a DD DUMMY, will result in IAM issuing a U0184 ABEND on various error conditions, so that a dump can be obtained at the point in time various error conditions are actually encountered. This DD card should only be provided when the IAM technical support staff so requests, otherwise a normally running job may abend unexpectedly.
IAMNOLIC	When specified as a DD DUMMY on IDCAMS LISTCAT requests, will prevent the production of an IAMPRINT report for IAM files. This is primarily intended for use when a job is doing a LISTCAT of an entire catalog, and it is desirable to eliminate the overhead of IAM processing, which can be significant when thousands of datasets are being processed.
IAMOVRID	Specifies a card image input dataset, which contains the control cards and keywords to use or modify the use of various IAM features. This is normally a DD *, however it can point to a sequential dataset, or to a member of a card image PDS. Refer to the IAM Overrides section 11 of the IAM manual for complete information on using IAM overrides.

61.60 IAM OVERRIDES FOR ESDS FILES

61.60 IAM OVERRIDES FOR ESDS FILES

This section is intended to briefly introduce the IAM Override capability, particularly as it relates to those overrides that will be of most use for IAM ESDS files. For complete information on using the IAM Overrides, please refer to [section 11](#) of the IAM manual.

There are two types of override statements, the CREATE override and the ACCESS override. The CREATE override is used to specify file attributes and various processing options on file DEFINES and file loads. The ACCESS override is used to specify various file processing options during non-load file access operations.

The following subset of the CREATE Override keywords are those which are of most use for IAM ESDS file users.

KEYWORD	PURPOSE
DDNAME=	Specifies the DDNAME of the IAM file that is being created or defined. A value of &ALLDD relates to all IAM files that go through a DEFINE or CREATE in the jobstep with the override. To indicate a specific file on the DEFINE, use a DDNAME value that corresponds to the FILE(ddname) value on the IDCAMS DEFINE statement for the file. Please note as long as the file being DEFINED is an IAM file, that an actual DD card by that name need not be in the JCL for the jobstep. For a file load done with IDCAMS, the DDNAME if other than &ALLDD must match the OUTFILE(ddname) specification to get a match.
CRBUFOPT=	Indicates the buffering option to be used during the file load. This override is only applicable for a file load, during file definition it is ignored. Valid values include: CYL: Get enough buffers to hold 1 cylinders worth of data blocks, and write out 1/2 a cylinder per I/O. This is the general default, unless specified to use a different value in the Global Options Table. MCYL: Get enough buffers to hold 2 cylinders worth of data blocks, and write out 1 cylinder per I/O. This option may provide an elapsed time savings during a file load. TRK: Get enough buffers to hold 2 tracks worth of data blocks, and write out 1 track per I/O. MTRL: Get enough buffers to hold 10 tracks worth of data blocks, and write out 5 tracks per I/O.
DATACOMPRESS=	Can be used to turn off datacompression for a file. For ESDS files, the default is to always use the IAM Data Compression feature, unless otherwise specified. Valid values are: YES: Enable IAM data compression for the specified file(s). NO: Do not data compress any of the records. This may be of value when the critical performance need is to reduce CPU time. With this option off, many of the benefits of IAM data compressed files, such as reduction in file size and reduction in file I/O activity will be less, however there should be some savings in the CPU time due to the elimination of the compression and decompression overhead.

61.60 IAM OVERRIDES FOR ESDS FILES

KEYWORD	PURPOSE
INTEGRATED=	Specifies, as a percentage, the amount of free space to be left in each block as IAM ESDS files are loaded and/or extended. For data compressed files which are frequently updated, specifying a non-zero value here will likely provide a benefit of keeping data out of Extended Overflow. For ESDS files that are not data compressed, or that are seldom updated, then this value is best left at 0 to conserve DASD space.
MAXBUFNO=	Specifies a default for the indicated file for the maximum number of buffers that IAM can use when processing this file for other than load operations. Valid values are 1 to 255 for the IAM ESDS files.
MINBUFNO=	Specifies a default for the indicated file for the minimum number of buffers that IAM can use when processing this file for other than load operations. Valid values are 1 to 255 for the IAM ESDS files. When specified, will also be used as the initial number of buffers. Note that by specifying the same value for MINBUFNO and MAXBUFNO, IAM will not do any dynamic changes in the number of buffers, however the other buffering related capabilities provided by IAM will still be effective.
PSEUDORBA	This is intended for use when the file size is expected to exceed 4.3 gigabytes on VSAM ESDS type of files. By indicating this keyword, IAM still uses an index to locate records, however rather than using a vaule that would be identical to VSAM for identifying records, IAM will pass back a unique number that is not compatible with the VSAM RBA. Applications that desire to use this capability are encouraged to do a lot of testing with the subject application and files to ensure that the application code can use this capability, and will process properly.
RELEASE	This value can be used to change the default of automatic space release on first load for files with secondary space allocations. Turning off this function may be useful for circumstances where there is expected to be a significant file growth, and you want to ensure that the DASD space remains available for use by the dataset. YES: Will result in unused DASD space always being released, not only on the first load but on any subsequent loads that are done without redefining the file. NO: Do not automatically release unused space at the end of the file load. This is useful for reserving DASD space for file growth of the particular file.

61.60 IAM OVERRIDES FOR ESDS FILES

The following subset of the ACCESS Override keywords are those which are of most use for IAM ESDS file users.

KEYWORD	PURPOSE
DDNAME=	Specifies the DDNAME of the IAM file that is being accessed. A value of &ALLDD relates to all IAM files that are accessed in the job step, unless an additional override with an explicit DDNAME is provided for that file.
DYNCORE=	Specifies the amount of storage, in 1K (1024) increments, to be used as an area to retain randomly read records in storage for future reference. This serves as a mechanism to buffer records in storage, as opposed to blocks. Additional information on IAM's Dynamic tabling is available in section 02.60 of the manual.
MAXBUFNO=	Specifies the maximum number of buffers that IAM can use for the specified file. Acceptable values are from 1 to 255. The default, unless changed in the IAM Global Option Tables, is the maximum is set to the number of buffers that will fit in 256K. A message will appear on the IAMINFO reports if raising MAXBUFNO will help reduce I/O.
MINBUFNO=	Specifies the minimum number of buffers that IAM will use for the specified file(s). When specified, this is also the value used for the number of buffers that IAM starts out with for the file at OPEN time. By setting MINBUFNO to the same value as MAXBUFNO will cause IAM to use a fixed number of buffers.

SECTION 62: SPECIAL CONSIDERATIONS

62.00 SPECIAL CONSIDERATIONS WHEN USING IAM FOR ESDS FILES

62.10 FREESPACE AND IAM ESDS FILE REORGANIZATIONS

With ESDS files, VSAM ignores any freespace specifications on the IDCAMS DEFINE, so they are not available for IAM to consider even if they are coded on the IDCAMS DEFINE statement. Due to the nature of data compression, when records are updated in an IAM ESDS file, the compressed size of the updated record may be smaller or larger than it was previously. If the record becomes larger, due to changes in the data not offering quite as good compression, then the record may have to be moved to a dynamically acquired overflow block, if there is insufficient room left in the original block to hold the enlarged record. As more and more records may end up using overflow blocks, sequential I/O performance will slowly deteriorate, and more virtual storage will be required for the record based index into the overflow area. Eventually, the file may need to be reorganized to correct these performance issues. One way to delay the need for reorganizations is to specify an amount of freespace to be left in each block. This is similar to the VSAM CI% Freespace concept for KSDS files. IAM terminology refers to this as Integrated Overflow.

IAM Integrated freespace can be specified through the IAM Override facility, using the I= keyword on the CREATE override card, or an installation default value can be set with the IAM ESDSINTEGRATED= global option. If most of the ESDS files are frequently updated, then the use of the global option is recommended, otherwise the use of the IAM Override (I=) would be preferable. Selecting a value for this freespace may be arbitrary to begin with, but as experience is gained with particular files, more reasonable values can be used. Suggested starting values for Integrated Freespace are between 5% to 10%. The IAMSMFVS reports can be used to see the frequency of updates against existing VSAM ESDS files, so that a judgment can be made to see how IAM will best fit into each individual site.

62.20 STORAGE USE

62.20 STORAGE USE

Storage requirements for IAM ESDS files will vary depending on a number of factors. The requirement for below the line storage will usually be 8K, when the VSAM control blocks are to reside in 24-bit addressable storage, or 4K when VSAM control blocks can reside in 31-bit addressable storage, such as occurs with CICS. These storage requirements may be larger when a larger number of strings are specified, through the STRNO parameter, or when additional strings are dynamically acquired because the STRNO value was too small. Requirements for above the line 31-bit addressable storage is a minimum of 8K plus storage for buffers and storage for the index. IAM will always keep the index for each open file in virtual storage. As records are added to the file, the size of the index will necessarily grow. The index to the extended data blocks is not presently compressed, so if a significant number of extended data blocks are in use, the storage requirements can be reduced by reorganizing the file. The index to the extended overflow blocks is record oriented, therefore it can end up requiring a lot of virtual storage as well. Here again, a file reorganization will reduce the storage requirements, although there is not anticipated to be a significant need for extended overflow blocks, particularly if some Integrated freespace is specified for the file.

62.30 VSAM COMPATIBILITY

62.30 VSAM COMPATIBILITY

IAM provides a VSAM compatible application program interface to the unique IAM file structure. Except where otherwise identified, IAM provides support for the documented VSAM macro calls, functions, and exits. IAM provides support for the VSAM control blocks that are commonly referenced by application programs through the documented TESTCB, SHOWCB, and MODCB macros. Programs and products which use undocumented capabilities or which utilize internal VSAM control blocks may encounter problems when attempting to utilize IAM files. When such problems are encountered IAM technical support is available to identify, and whenever possible, correct these problems.

In general, IAM files can be used without changes to JCL or application programs, other than to have the file DEFINE modified to indicate that the file is to be an IAM file. In some circumstances, there may be some JCL changes recommended to better utilize IAM. There may arise some circumstances where JCL changes may be required. IAM technical support is available to help identify these circumstances, and rectify any problems encountered.

IAM supports ESDS Control Interval processing for loading files, and for reading files. Control Interval updates are not presently supported.

IAM supports the VSAM LERAD, SYNAD, EODAD and UPAD exits. IAM provides support for the non-LSR entry calls to the JRNAD exit. IAM does not presently support the use of the EXCEPTIONEXIT, which is specified on the file IDCAMS DEFINE of datasets, or the USVR (User Security Verification Routine).

Application programs can issue the VSAM LSR macros, MRKBFR, SCHBFR, and WRTBFR. The WRTBFR will effectively serve the same purpose as it does with VSAM, which is to force the specified group of updated buffers out to DASD. The MRKBFR has no equivalent function for IAM, because IAM does not lock buffers to RPL ownership, except when there is actual I/O in progress. The SCHBFR request will always fail with a VSAM logical error, as IAM has no way to handle this type of request, because IAM is not using the VSAM LSR buffer pool, and IAM is not buffering what the application would see as being a CI buffer.

JCL which has been setup to utilize the IBM Batch LSR feature can be used with IAM ESDS files without change. IAM will not use any of the LSR buffers, so while it is recommended that the JCL be changed to eliminate the usage of Batch LSR, it is not a requirement to do so.

IAM files can not use the IBM HIPERBATCH capability that is available for VSAM files.

IAM does not support SPANNED records, either in the simulated CI's or within the IAM data blocks. IAM will force a blocksize that is large enough to hold the maximum record size specified by the DEFINE. The maximum size user data record supported by IAM for ESDS files is 32,752.

SECTION 63: IAM ESDS FILE RUN TIME STATISTICS (INFO) REPORT

63.00 IAM ESDS FILE RUN TIME STATISTICS (INFO) REPORT

63.10 EXAMPLES OF RUN TIME STATISTICS REPORTING

RUN TIME STATISTICS REPORT Activate run time statistics reporting on the IAM data set 'IAMV.DCM201.CLUSTER' when the program 'USERPGM' is executing. Print the INFO report on SYSOUT=A.

```
//DCM201B EXEC PGM=USERPGM
//VSAMCRT1 DD DSN=IAMV.DCM201.CLUSTER,
DISP=OLD...
//IAMINFO DD SYSOUT=A
```

SAMPLE CREATE STATISTICS 1. IAM RUN TIME STATISTICS REPORT FROM FILE LOAD

```
IAM400 INNOVATION IAM INFORMATION PRINT ROUTINE -- IAMNINFO VER 6.X --
INNOVATION DATA PROCESSING DATE - 94.074
IAM360 STEP - IAMVTST1 DDNAME - SYS00003 DATA SET MONITORED -
IAMV.GFM$IAM.IAMFILE1
IAM361 INFO REQUESTED BY IAMCRT PROGRAM PERFORMING FILE CREATION
OPENED-94.074.09:07:42 CLOSED-94.074.09:08:12
IAM362 IAM DATA CHARACTERISTICS -
LOGICAL RECORD LENGTH-----= 500 - RECORD FORMAT-----= VARIABLE
SIMULATED CI SIZE -----= 4096 - SIMULATED HIGH USED RBA---= 397312
AVERAGE RECORD LENGTH-----= 500 - LARGEST RECORD PROCESSED--= 500
IAM363 IAM FILE CHARACTERISTICS -
BLOCKING FACTOR-----= 4 - BLOCK SIZE-----= 11476
TRACKS IN USE-----= 10 - DELETE BYTE SUPPORT-----= NONE
DATASET TYPE-----= ESDS - SHARE OPTIONS-----= 2
NUMBER OF IAM DATA BLOCKS--= 37 - HIGH ALLOCATED RBN -----= 0
IAM372 IAM EXTENDED AREA CHARACTERISTICS -
EXTENDED OVERFLOW RECORDS--= 0 - EXTENDED OVERFLOW BLOCKS--= 0
EXTENDED BLOCKS ALLOCATED--= 0 - EXTENDED BLOCKS USED-----= 0
IAM365 IAM EXECUTION STATISTICS -
TOTAL STORAGE REQUIRED-----= 757032 - COMPRESSED KEY STRUCTURE--= NO
STORAGE ABOVE THE LINE-----= 737280 - COMPRESSED DATA STRUCTURE--= NO
REQUESTS PROCESSED-----= 771 - REQUESTS FAILED-----= 0
DISK BLOCKS READ-----= 0 - DISK BLOCKS WRITTEN-----= 3
DYNAMIC BUFFER RETRIEVALS--= 0 - MAXIMUM BUFFERS USED-----= 60
DATA SPACE USED (M BYTES)--= 1 - DATA SPACE SIZE(M BYTES)--= 128
IAM366 IAM COMMAND EXECUTION SUMMARY -
R.(READ)-----= 0 - W.(WRITE)-----= 770
F.(GETNEXT)-----= 0 - A.(ADD)-----= 0
D.(DELETE)-----= 0 - N.(READOR GETNEXT)-----= 0
G.(GET)-----= 0 - S.(START KEY EQUAL)-----= 0
K.(START KC OR POINT)-----= 0 - P.(GET PREVIOUS)-----= 0
U.(READ FOR UPDATE)-----= 0 - RESERVED-----= 0
I.(INFORMATION)-----= 0 - E.(STATISTICS)-----= 0
C.(CLOSE)-----= 1 - O.(OPEN)-----= 0
X.(FLUSH)-----= 0 - VARIABLE LENGTH CHANGES---= 0
SEQ CHAIN BUFF/DEL REC BYP= 0 -
```

63.00 IAM ESDS FILE RUN TIME STATISTICS (INFO) REPORT

SAMPLE 2. IAM RUN TIME STATISTICS REPORT FROM FILE ACCESS
ACCESS
STATISTICS

IAM400	INNOVATION IAM INFORMATION PRINT ROUTINE -- IAMNINFO VER 6.X		
	-- INNOVATION DATA PROCESSING DATE - 94.074		
IAM360	STEP - IAMVTST1 DDNAME - SYS00003 DATA SET MONITORED -		
	IAMV.GFM\$IAM.IAMFILE1		
IAM361	INFO REQUESTED BY IAMAVSOC PROGRAM PERFORMING RANDOM/GETNEXT		
	OPENED-94.074.09:07:11 CLOSED-94.074.09:07:42		
IAM362	IAM DATA CHARACTERISTICS -		
	LOGICAL RECORD LENGTH-----= 500 - RECORD FORMAT-----=	VARIABLE	
	SIMULATED CI SIZE -----= 4096 - SIMULATED HIGH USED RBA --=	2457600	
	TOTAL RECORDS-----= 1270 - TOTAL RECORDS DELETED-----=	0	
	TOTAL RECORDS UPDATED-----= 0 - TOTAL RECORDS INSERTED-----=	50	
IAM363	IAM FILE CHARACTERISTICS -		
	BLOCKING FACTOR-----= 4 - BLOCK SIZE-----=	11476	
	TRACKS IN USE-----= 16 - DELETE BYTE SUPPORT-----=	NONE	
	DATASET TYPE-----= ESDS - SHARE OPTIONS-----=	2	
	NUMBER OF IAM DATA BLOCKS--= 37 - HIGH ALLOCATED RBN -----=	xxxx	
IAM372	IAM EXTENDED AREA CHARACTERISTICS -		
	EXTENDED OVERFLOW RECORDS--= 0 - EXTENDED OVERFLOW BLOCKS--=	0	
	EXTENDED BLOCKS ALLOCATED--= 0 - EXTENDED BLOCKS USED-----=	0	
IAM365	IAM EXECUTION STATISTICS -		
	TOTAL STORAGE REQUIRED-----= 315392 - COMPRESSED KEY STRUCTURE--=	NO	
	STORAGE ABOVE THE LINE-----= 307200 - COMPRESSED DATA STRUCTURE--=	NO	
	REQUESTS PROCESSED-----= 502 - REQUESTS FAILED-----=	0	
	DISK BLOCKS READ-----= 4 - DISK BLOCKS WRITTEN-----=	0	
	DYNAMIC BUFFER RETRIEVALS--= 25 - MAXIMUMBUFFERS USED-----=	32	
	DYNAMIC TABLE RETRIEVALS--= 0 - DYNAMICTABLE RECORDS-----=	0	
IAM366	IAM COMMAND EXECUTION SUMMARY -		
	GET RANDOM-----= 0 - PUT UPDATE-----=	0	
	GET SEQUENTIAL-----= 0 - PUT ADD-----=	500	
	GET PREVIOUS-----= 0 - POINT (START BROWSE)-----=	0	
	GET CI (SEQUENTIAL)-----= 0 - PUT CI (UPDATE)-----=	0	
	GET CI (RANDOM)-----= 0 - PUT CI (ADD)-----=	0	
	ENDREQ-----= 0 - WRTBFR-----=	0	
	IAM STATISTICS-----= 0 - IAM FLUSH BUFFER-----=	0	
	CLOSE-----= 1 - OPEN-----=	0	
	CLOSE TYPE=T-----= 0 - VERIFY-----=	0	
	INVALID REQUESTS-----= 0 - RECORD LENGTH CHANGED-----=	0	
	SEQ CHAINED BLOCKS READ---= 0 - SEQ CHAINED BLOCKS WRITTEN=	18	

63.20 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

63.20 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

The IAM Run Time Statistics (INFO) Report is produced when the IAM file is CLOSED. There will be one report produced for each OPEN/CLOSE combination a program issues for an IAM file.

A description of the report and its major fields follows.

- IAM400** The IAMINFO Report title containing a module name and version.
- IAM360** Provides the name of the job step in execution, the DD statement, DDNAME and data set name of the IAM file being processed.
- IAM361** Names the IAM processing module that requested the INFO Report, the type of processing being performed, and the date and time the IAM file was opened and closed. The type of processing maybe CREATE, SEQUENTIAL, RANDOM or RANDOM/GETNEXT. CREATE is shown when a file is loaded. RANDOM/GETNEXT will be shown whenever a file is processed with random and/or sequential commands. Sequential access is reported only when after open the PGM starts with the first record then proceeds record by record through the file (i.e., on start/point command specified).
- IAM362** This heading line always precedes the IAM file Data Characteristics portion of the IAM INFO report. The following fields describing the logical record are provided.

- LOGICAL RECORD LENGTH** – The maximum length of the record.
- RECORD FORMAT** – Fixed or Variable format indication.
- AVERAGE RECORD LENGTH** – During a load operation IAM calculates an average length for the records that were loaded into the file.
NOTE: This field only appears when a file is LOADED.
- LARGEST RECORD PROCESSED** – During a load operation IAM reports the length of the largest record that was written to the file.
NOTE: This field only appears when a file is LOADED.
- TOTAL RECORDS** – The total number of records in the IAM file.
NOTE: This field does not appear when the file is initially created.
- TOTAL RECORDS DELETED** – The total number of records deleted from the IAM file since creation.
NOTE: This field does not appear when the file is initially created.
- TOTAL RECORDS UPDATED** – The total number of records updated in the IAM file since creation.
NOTE: This field does not appear when the file is initially created.
- TOTAL RECORDS INSERTED** – The total number of records inserted into the IAM file since creation.
NOTE: This field does not appear when the file is initially created.
- SIMULATED HIGH USED RBA** – This is the value IAM will return if a request is made for the high used RBA of an IAM ESDS file.
NOTE: This field will only appear after a file is loaded.

63.20 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

IAM363 This heading line always precedes the IAM File Characteristics portion of the IAM INFO report. The following fields describing the physical file are provided.

BLOCKING FACTOR – A value of 1-9 is the number of blocks per track (as blocking factor). A larger value is the blocksize if it was specified in bytes.

BLOCKSIZE – Physical blocksize of the IAM file. This is the actual blocksize developed and used by IAM.

NOTE: IAM will determine the optimal blocksize for a file based upon the file's record length, the blocking factor requested and the track capacity of the device. The IAM blocksize is transparent to application programs even when the blocksize is altered or the file is relocated to a different device type.

TRACKS IN USE – Number of tracks actually occupied by the file. Not the total allocation.

NOTE: Any disk space allocated beyond the value identified by the 'tracks used' field are an over allocation. They will not be used by IAM for any file processing. Over allocated space can be released automatically when an IAM file is loaded. See Section 11, the RELEASE keyword on the CREATE statement.

DELETE BYTE SUPPORT – Indicates if ISAM delete byte support is active for the file and the option in effect, see Section 11 OPTCD keyword on the ACCESS statement.

DATASET TYPE – Describes the type of data set IAM is simulating. The default type for ESDS files is ESDS. If PSEUDORBA was specified during the definition of the file, this field will be ESDS/P. PSEUDORBA will allow an IAM ESDS file to be larger than the 4.3G. This is done by using nonstandard RBAs during processing.

SHARE OPTIONS – Describes the defined IAM share options for this data set.

HIGH ALLOCATED RBN – This the highest allocated block number currently in the IAM file.

IAM372 This heading line always precedes the IAM File Extended Area characteristics portion of the IAM INFO report. The following fields describe the extended data and overflow areas that were obtained automatically as needed.

EXT. OVERFLOW RECORDS Number of records currently contained in extended overflow.

EXT. OVERFLOW BLOCKS Number of extended overflow blocks available.

EXTENDED BLOCKS ALLOCATED Actual number of data/extended overflow blocks allocated by IAM.

EXTENDED BLOCKS USED Actual number of data/extended overflow blocks in use.

63.20 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

IAM365 This heading line always precedes the IAM Execution Statistics portion of the IAM INFO report. The following fields describing system processing and IAM options used during file processing are provided.

- TOTAL STORAGE REQUIRED** – Total storage acquired by IAM during this execution.
- COMPRESSED KEY STRUCTURE** – Indicates whether a compressed index was used for this file. Under MVS/XA and MVS/ESA systems IAM's compressed index is maintained above the 16MB line.
- STORAGE ABOVE THE LINE** – Total amount of storage acquired above the 16MB line.
- COMPRESSED DATA STRUCTURE** – Indicates whether data compression was used for this file.
- REQUESTS PROCESSED** – Number of requests made against the file, since last OPEN, by the application (a breakdown by command type follows in the command execution summary).
- REQUESTS FAILED** – Number of requests made to IAM which did not complete normally (i.e. EOF on sequential read, no record found on random read)
- DISK BLOCKS READ** – Number of physical I/O's used to read blocks from the file.
- DISK BLOCKS WRITTEN** – Number of physical I/O's used to write blocks to the file.
- DYNAMIC BUFFER RETRIEVALS** – Number of EXCPs saved (not executed). IAM's Real Time Tuning buffer management was able to satisfy this number of I/O requests from buffers that it had maintained in storage, other than its current buffer. These buffer retrievals eliminated the need for I/O to disk.
- MAXIMUM BUFFERS USED** – The maximum number of data buffers acquired by IAM's Real Time Tuning during this execution. This may be accompanied by an IAM367 or IAM368 informational message indicating that if IAM was allowed to acquire more buffers the number of I/O's required to service the requests against the file could have been reduced. Under MVS/XA and MVS/ESA all buffers are acquired above the 16MB line.
- DYNAMIC TABLE RETRIEVALS** – Number of random read requests satisfied by IAM's Dynamic Tabling of data records in virtual. Under MVS/XA and MVS/ESA IAM's Dynamic Table (DYNCORE) is maintained in virtual storage above the 16MB line.
- DYNAMIC TABLE RECORDS** – The total number of data records IAM placed into the Dynamic Table.
- DATA SPACE USED(M BYTES)** – The amount of area actually used in the data space during file load.
- DATA SPACE SIZE(M BYTES)** – The size of the data space created during file load.

63.20 KEY INFORMATION IN THE RUN TIME STATISTICS (INFO) REPORT

IAM366 IAM COMMAND EXECUTION SUMMARY

GET RANDOM	– Number of random READ requests
PUT UPDATE	– Number of UPDATE requests
GET SEQUENTIAL	– Number of sequential GET requests
PUT ADD	– Number of INSERT requests
GET PREVIOUS	– Number of GET PREVIOUS requests
POINT (START BROWSE)	– Number of POINT requests
GET CI (SEQUENTIAL)	– Number of sequential control interval GETS
PUT CI (UPDATE)	– Number of CONTROL INTERVAL UPDATES
GET CI (RANDOM)	– Number of CONTROL INTERVAL GETS
PUT CI (ADD)	– Number of CONTROL INTERVAL ADDS
ENDREQ	– Number of LSR ENDREQ requests
WRTBFR	– Number of LSR WRTBFR requests
IAM STATISTICS	– Number of IAM statistic requests
IAM FLUSH BUFFER	– Number of IAM flush buffer requests
CLOSE	– Number of CLOSE requests
OPEN	– Number of OPEN requests
CLOSE TYPE=T	– Number of TEMPORARY CLOSE requests
VERIFY	– Number of VERIFYs
INVALID REQUESTS	– Number of requests with RPL error's
RECORD LENGTH CHANGES	– Number of records with length changes
SEQ CHAINED BLOCKS READ	– Number of chained buffers during READ
SEQ CHAIN WRITTEN	– Number of chained buffers during WRITE

IAM367 . INSUFFICIENT STORAGE AVAILABLE TO ACQUIRE ADDITIONAL BUFFERS

This is an informational message provided by IAM's Real Time Tuning buffer management. It is provided when an attempt to acquire additional data buffers fails because of insufficient virtual storage. IAM acquires additional buffers when it had determined that this would improve performance. Increasing the region size will result in fewer disk I/O's to process this file.

IAM368 ... BUFNO VALUE GREATER THAN nn MAY IMPROVE PERFORMANCE

This informational message is provided whenever IAM determines that a greater number of buffers than provided would result in fewer disk I/O's to process this file. The MAXBUFNO= parameter on the IAM Override statement can be used to increase the number of buffers IAM will use. If the actual number of DISK BLOCKS READ and DISK BLOCKS WRITTEN is small there may be little benefit to be gained.

IAM371 INCREASING THE NUMBER OF PRIME EXTENSION BLOCKS MAY IMPROVE PERFORMANCE.

This informational message is provided whenever IAM determines that a record has been added to the INDEPENDENT OVERFLOW instead of being added to the PRIME EXTENSION because the PRIME EXTENSION is full.

SECTION 64: IAM ESDS TRACE FACILITY

64.00 IAM ESDS TRACE FACILITY

64.10 IAM 6.2 TRACE FACILITY FOR ESDS FILES.

In order to trace IAM 6.2 activity, you must specify, at a minimum, the TRACEDDNAME=ddname override. If nothing else is specified in the override, IAM will perform an I/O End (TRIOE) trace. In order to request additional or other trace information, you must use the TRACEREQUEST= override.

One or more of the following may be entered in any order:

- TRIOS** – TRace I/O Start
- TRIOE** – TRace I/O End (default if TRACEREQUEST not specified)
- TRBFR** – TRace Buffer manager calls
- TREXCP** – TRace EXCP (physical I/O requests)
- TRXTND** – TRace eXTeND (new extent requests)

For example, to request tracing of I/O Start and I/O End requests, enter the following on an ACCESS OVERRIDE card:

```
TRACEREQUEST=(TRIOS,TRIOE)
```

Tracing will only take place for the specified DDNAME that the IAM override refers to.

IAM 6.2 TRACE RECORD SELECTION UTILITY

Since the IAM trace facility could generate many lines of information, IAM supplies the IAMATSEL utility to locate and report on only the information you may be interested in without having to scan through (perhaps) hundreds or thousands of lines of trace data.

IAMATSEL is a program that reads an input file containing trace records generated by IAM 6.2 trace requests, and sends its output to an output file. The following is an example JOB:

```
//IAMTSEL JOB 'IAMATSEL',CLASS=M,MSGCLASS=X,NOTIFY=USERID
//STEPLIB DD ...
//SYSPRINT DD SYSOUT=*
//IAMTRPRC DD DISP=SHR,DSN=trace.output.dataset <== INPUT FILE
//IAMTRPRO DD SYSOUT=* <== OUTPUT DD
//SYSIN DD *
SELECT ...
/*
```

The SELECT statement above, has many options that can help you locate only the trace information you desire. The parameters of the SELECT statement are described below.

The first set of SELECT parameters are strictly character matching tests, and are broken down by fieldnames found in the IAM 6.2 Trace output. The TYPE= and REQUEST= fields have a limited number of possible values, where the RPL=, RBN=, PLH=, OPTCD=, RTNCD=, RINFO=, RECLen=, KEY=, and TIME= fields will contain addresses, numeric values, etc.

TYPE= IOS|IOE|BFR|EXCP|XTND|ALL

Select the TYPE of trace records you want to see. One or more of the above types can be requested. If more than one type is requested, they must be placed within parens (IOS,IOE,BFR).

REQUEST= GET|PUT|POINT|ERASE|CHECK|ENDREQ|READ|RELEASE|WRITE|FLUSH|ALL

Select the type of VSAM, IAM, or IAM Buffer Manager requests to report on. GET, PUT, POINT, ERASE, CHECK, and ENDREQ are VSAM requests that will appear on IOS and IOE trace type records. READ, RELEASE, WRITE and FLUSH are IAM Buffer Manager requests that will appear on BFR trace type records. There are no specific request types for EXCP or XTND records.

RPL= (up to 8 hex 'characters')

64.00 IAM ESDS TRACE FACILITY

RBN=	Select request using a specific RPL address. All characters entered (up to 8) will be used to determine if a matching RPL address is valid. Note: RPL= and RBN= reference the same field in trace output.
PLH=	Up to 8 hex 'characters'.
OPTCD=	Up to 4 hex 'characters'.
RTNCD=	Up to 8 hex 'characters'. Note: RTNCD= and RINFO=
RINFO=	reference the same field in trace output.
RECLN=	Up to 4 hex 'characters'.
KEY=	Up to 8 hex 'characters'.
TIME=	Time values with the format 'HH.MM.SS'.

The following SELECT parameters are used to locate a specific condition that has occurred without regard to specific addresses and return codes:

COND=	GETFAIL PUTFAIL Locate a GET or PUT failure in the trace output.
PRERECS=	A decimal number of up to 6 digits which specifies how many trace records to print out before the GET or PUT failure trace record. This will let you see (perhaps) what led up to the failure.

The following SELECT parameters are used to control the overall trace record selection process:

INDD=	DDNAME of trace records to scan. The default is IAMTRPRC.
OUTDD=	DDNAME of trace record selection output. the default is IAMTRPRO.
SKIP=	A decimal number of up to 6 digits which indicates how many trace records should be skipped over before the selection process begins. This is useful for large trace files.
COUNT=	A decimal number of up to 6 digits which indicates a maximum number of trace selection records to output.

IAMM

INNOVATION ACCESS METHOD

MESSAGES & CODES

THIS PAGE WAS INTENTIONALLY LEFT BLANK

SECTION 80: IAM V6.0 IAM MESSAGES AND CODES

80.01 IAM MESSAGES AND CODES INTRODUCTION

80.01 INTRODUCTION

INTRODUCTION This section describes the various printer, console and TSO messages which may be output by INNOVATION's IAM system and the various ABEND codes with which it may terminate.

The general format of IAM messages and WTORs is as follows:

MESSAGES FORMAT

IAMnnn Message-text

or

IAMWnn WTO/LOG Message-text

IAM UTILITY RETURN CODES IAM passes a return code at the end of the step, unless it ABENDS. A return code of zero indicates that IAM has performed all functions successfully. A return code of four (004) indicates that the testing period for a trial version of the product has expired. The production version of the product library does not contain a date routine. Any other return code is accompanied by error messages and indicates that errors of some kind have occurred during the execution. Check the listing for the error messages.

VSAM RETURN CODES IAM's VSAM Interface (VIF) will pass return codes comparable to the return codes set by VSAM. The return codes are passed in the RPL or by IDCAMS. If the IAM return code is different than VSAM, it will be documented in [Sections 80.20 through 80.22 \(VSAM Return Codes\)](#).

ABEND CODES IAM ABEND codes range from U0100 to U0999. Most user ABENDS are preceded by an IAM message.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

The following write-to-operator, write-to-operator-with-reply and TSO messages are issued by IAM:

IAMW01 DD= 'ddname' I/O ERROR CODE=X'decb iobecb csw cccchhhrr'

Reason: This message is provided by the IAM processor when a file access fails with an I/O error. The message contains the following diagnostic information:

DECB – displays the two error bytes of the BDAM DECB(+1).
IOBECB – displays the first four bytes of the IOB.
CSW – displays the channel status word.
ccccchhhrr – displays the cylinder, head and record head number the error was detected on. The head and record number may not be accurate. IBM's message IEA000I/IOS000I, if present, contains the real track address.

VIF RETURN CODES: A Return Code of 12, X'0C', and a reason code of X'04' or X'10' is stored in the RPL. The SYNAD exit will be invoked, if so specified by the program.

IFI BATCH & ONLINE RETURN CODES: Reason code of C 'Y' is stored in the DCB at displacement decimal +138 and the DECB used for the I/O will contain the completion code returned. IFI batch processing terminates with a U0101. IFI online processing continues. IAM native terminates with a U0101.

Action: If an IBM IEA000I/IOS000I message appears on the JCL LOG, a hardware error has occurred. The IBM message gives the sense information (ex: data check, equipment check). If this message is not displayed, contact INNOVATION.

Examine the error information provided in the IBM message to determine the cause of the error. The format of this information is documented in the IBM data management SRL for the operating system in use. The file in question is unreadable in its present state and must be reestablished. This may be the result of a hardware error and if possible the new file should be allocated to a different physical location or volume.

The utility IAMRECVR may be used to recover a file that is no longer usable due to I/O or logical error conditions. This utility can be used to off load records from those portions of the file that have not been physically damaged. The user program should make a decision on the action to take when this occurs, for example to continue processing without this file or to terminate processing until this file is made available. The appropriate action will depend entirely on the application and the user program's evaluation of the diagnostic information that is returned.

IAMW02 REPLY RETRY WAIT OR CANCEL FOR ENQ. WAIT ON 'dsname'

Reason: The IAM file is not available to this job because some other job is currently accessing the file. A load process is always protected against concurrent access (read or update) or from another load. The Share Options specified when the file was DEFINEd determine which types of access can concurrently share the file.

[Sections 10 and 11](#) of this manual document IAM's Share Options.

Action: The operator reply determines the course of action IAM will take. The following responses are allowed:

RETRY – IAM retries the ENQUEUE for the data set. If the file is still unavailable, the message will be re-issued.
WAIT – IAM waits for the data set to become available. Caution: The job could time out.
CANCEL – IAM fails the OPEN request with a return code indicating the file is not currently available.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW03 DD= 'ddname' FILE FULL DUE TO INSUFFICIENT CORE

Reason: This message is displayed by IAM when it finds an insert has failed because of insufficient storage to expand the incore overflow index. The record that was to be added was not placed in the file. Overflow may still be available at other locations within the file. Consequently, subsequent inserts may or may not fail depending on where in the file they are placed. Additionally, changes in the storage available within a region may allow a subsequent GETMAIN to succeed.

This message is displayed a maximum of 10 times per execution.

VIF RETURN CODES: A return code of eight (8) and a reason code of X'1C' is stored in the feedback field of the RPL. The LERAD exit will be invoked, if so specified in the program.

IFI BATCH & ONLINE RETURN CODES: A reason code of C'F' is stored in the DCB at displacement decimal +138 and the DECB used for the I/O will contain a completion code of file full.

Action: The definition of the file's OVERFLOW parameters as well as the OCOREO% and OCOREX% parameters should be reviewed. Increasing the OCOREO% value to acquire a larger overflow index area when the file is OPENed should correct the problem. The job step region size or, for CICS users, the CICS OSCORE parameter may also need adjustment so more storage will be available for use by this particular task.

IAMW04 DD= 'DDNAME' OPEN ERROR - DATA SET NOT AVAILABLE, IN USE BY JOB JOBN

Reason: IAM was unable to open the specified file because it was already opened by another job. For ESDS files, if IAM is able to determine which JOB or user has the file, it will be so indicated in the message.

VIF RETURN CODES: A reason code of 168(X'A8') is stored in the ACB error flags field (ACBERFLG) and the OPEN is failed with a return code of 8.

Action: Rerun the job when the indicated job has terminated or closed the IAM file that was in use.

IAMW05 DD= 'ddname' OPEN ERROR - MORE THAN 2 OPENS FOR UPDATE IN TASK

Reason: This message is displayed by IAM when it finds a task has issued multiple OPENS for update against the same dataset.

VIF RETURN CODES: A reason code of X'F8' is stored in the ACB error field at displacement decimal +49 and the OPEN CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10').

IFI BATCH & ONLINE RETURN CODES: Reason code of C'M' is stored in the DCB at displacement decimal +138 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10').

Action: Change the program to insure that the same file is never OPENed for update more than once within the processing task without first closing it. Concurrent multiple Opens for read-only processing are permitted in the IFI interface.

IAMW06 DD= 'ddname' OPEN ERROR - DD STATEMENT MISSING

Reason: This message is displayed when a task has issued an OPEN, but the 'ddname' the OPEN was issued against did not exist.

VIF RETURN CODES: Reason code of X'80' is stored in the ACB error field at displacement decimal +49 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10').

IFI BATCH & ONLINE RETURN CODES (without IAMOPEN-IAMCLOSE): Reason code of C'O' is stored in the DCB at displacement decimal +138 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10'). IFI batch processing terminates with a U0101. IFI online processing continues.

Action: Check to see if the missing DD statement is not the result of a misspelling. If not, add the DD statement to the JCL.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW07 DD= 'ddname' I/O ERROR CODE=X'decb iobecb csw cccchhhrr'

Reason: This message is displayed when a file OPEN fails with an I/O error. The message contains the following diagnostic information:

- DECB – displays the two error bytes of the BDAM DECB (+1).
- IOBECB – displays the first four bytes of the IOB.
- CSW – displays the channel status word.
- cccchhhrr – displays the cylinder, head and record number the error was detected on. The head and record number may not be accurate. IBM's message IEA000I/IOS000I, if present, contains the real track address.

NOTE: This message may also occur when an attempt is made to OPEN a non-IAM or non-VSAM data set as VSAM.

VIF RETURN CODES: A Return Code of 12, X'0C', and a Reason code of X'B8' is stored in the ACB error field at displacement decimal +49 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10').

IFI BATCH & ONLINE RETURN CODES: Reason code of C'R' is stored in the DCB at displacement decimal +138 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10'). IFI batch processing terminates with a U0101. IFI online continues.

Action: If an IBM IEA000I/IOS000I messages appears on the JCL LOG, a hardware error has occurred. The IBM message gives the sense information (ex: data check, equipment check). If this message is not displayed, check to see if the file is a properly formatted IAM or VSAM file. If the file is IAM or VSAM, contact INNOVATION.

Examine the error information provided in the IBM message to determine the cause of the error. The format of this information is documented in the IBM data management SRL for the operating system in use. The file in question is unreadable in its present state and must be reestablished. This may be the result of a hardware error and if possible the new file should be allocated to a different physical location or volume.

The utility IAMRECVR may be used to recover a file that is no longer usable due to I/O or logical error conditions. This utility can be used to off load records from those portions of the file that have not been physically damaged. The user program should make a decision on the action to take when this occurs. For example, to continue processing without this file or to terminate processing until this file is made available. The appropriate action will depend entirely on the application and the user program's evaluation of the diagnostic information that is returned.

IAMW08 DD= 'ddname' OPEN ERROR - INSUFFICIENT STORAGE IN REGION

Reason: This message is displayed by IAM when a file OPEN fails because storage is not available for data buffers, the index or control information.

VIF RETURN CODES: Reason code of X'88' is stored in the ACB error field at displacement decimal +49 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10').

IFI BATCH & ONLINE RETURN CODES: Reason code of C'G' is stored in the DCB at displacement decimal +138 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10'). IFI batch processing terminates with a U0130. IFI online processing continues. IAM native terminates with a U0130.

Action: The job step region size must be increased. For CICS users the CICS OSCORE parameter may also need adjustment so more storage will be available for use by this particular task.

IAMW09 OPEN ERROR - FILE ALLOCATED BUT NOT CREATED DD= 'ddname'

Reason: This message is displayed by IAM when a task has issued an OPEN against an IAM file but the dataset (referenced by 'ddname') was an allocation that had never been loaded as an IAM file.

VIF RETURN CODES: Reason code of X'F8' is stored in the ACB error field at displacement decimal +49 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10').

IFI BATCH & ONLINE RETURN CODES: Reason code of C'B' is stored in the DCB at displacement decimal +138 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10'). IFI batch processing terminates with a U0118. IFI online processing continues. IAM native terminates with a U0118.

Action: The file must be created/loaded before it can be used or the file was not successfully loaded.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW10 OPEN ERROR - NOT AN IAM/VSAM FILE OR CREATE FAILED DD= 'ddname'

Reason: This message is displayed by IAM when a task has issued an OPEN for the dataset (referenced by 'ddname') but the file did not contain valid IAM or VSAM data. The dataset may be the result of a load that failed, or some other processor overwriting the file.

VIF RETURN CODES: Reason code of X'F8' is stored in the ACB error field at displacement decimal +49 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10').

IFI BATCH & ONLINE RETURN CODES: Reason code of C'1' is stored in the DCB at displacement decimal +138 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10'). IFI batch processing terminates with a U0302. IFI online processing continues. IAM native terminates with a U0302.

Action: First, determine if the file being OPENed is supposed to be an IAM or a VSAM file. If it is supposed to be an IAM file, review the output from the job which created/loaded the file for errors.

IAMW11 DYNCORE DISABLED DUE TO INSUFFICIENT CORE DD= 'ddname'

Reason: This message is displayed by IAM when a user requested storage for IAM's Dynamic Table option and the storage is not available. Processing continues as if the request for Dynamic Tabling had not been requested.

Action: The user should increase the region size or decrease the value requested for DYNCORE.

IAMW12 DATA COMPRESS ERROR xxxxxxxx xxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxx

Reason: During a file access of a data compressed IAM file, IAM decompression failed. The first four bytes are the internal RDW, the remaining data is the key.

Action: Run an IAMRECVR DIAGNOSE function to validate that the file is still valid. Contact INNOVATION for assistance.

IAMW13 DD= 'DDNAME' FILE FULL, REORGANIZATION REQUIRED

Reason: IAM was unable to add a record into a file, or unable to accept a larger size updated record. For KSDS type of files, the Independent Overflow area has been filled. For ESDS files, IAM was not able to obtain additional DASD space to expand the size of the file.

VIF RETURN CODES: A reason code of 28(X'1C') is stored in the RPL error flags field (RPLERRCD) and the PUT is failed with a return code of 8. If the application has provided a LERAD exit, the exit will be invoked.

Action: Reorganize the file using an IDCAMS REPRO. For KSDS files, it may be necessary to make the size of the overflow areas larger. For ESDS files, a larger space allocation may be necessary.

IAMW16 DD= 'ddname' OPEN ERROR - IAM DOES NOT CURRENTLY SUPPORT LSR

Reason: This message is displayed by IAM when a task has issued an OPEN for a dataset and the Local Shared Resources processing option was specified in the ACB. The OPEN fails.

VIF RETURN CODES: Reason code of X'E4' is stored in the ACB error field at displacement decimal +49 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10').

Action: If the task opening the dataset has retried the OPEN after removing LSR from the ACB, the task will continue processing normally after the IAMW16 message is issued. In this case the task itself is prepared to process normally when LSR support is not available. If a task abnormally terminates following the message the files in question must have the request for LSR support removed. Contact INNOVATION for further technical assistance.

NOTE: In CICS change the FCT entry to specify the parameter LSRPOOL=NONE. (See Section 90.10).

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW17 DD= 'ddname' ADD FAILED - MORE THAN 1 ACB CONCURRENTLY UPDATING FILE.

Reason: IAM does not support concurrent file update. This message will be displayed when IAM detects a concurrent attempt by more than one user (ACB) to ADD, update VL or update compressed format records. IAM expects proper procedures to be used to prevent concurrent update as it will eventually result in damage to the file. This message is displayed a maximum of 10 times per execution.

VIF RETURN CODES: A return code of eight (8) and a reason code of X'1C' (file full) is stored in the feedback field of the RPL. The LERAD exit will be invoked, if so specified in the program.

Action: In a multi-processor environment IAM files must be protected, by a global enqueue facility or manual scheduling, from concurrent update by tasks running on different processors. Within a single processor multiple tasks will be prevented from concurrently OPENing a file for update by the proper use of SHAREOPTIONS. IAM defaults to a cross region SHAREOPTION 1, multiple read access users or 1 update user. The use of SHAREOPTIONS 3 and 4 disable IAM's enqueue protection for a file as does the IAMOVRID parameter UPDATENQ=NONE. Use these facilities with great care. Within a single task use of multiple ACB's against an IAM file will defeat IAM's enqueue protection for that file. To preserve the integrity of your IAM files any application program that accesses an IAM file with multiple ACB's must be changed to ensure only one ACB is used to access the file for update. (For further assistance contact INNOVATION).

IAMW18 DD='ddname' USER NOT AUTHORIZED FOR UPDATE

Reason: For Enhanced Format files, IAM issued a RACROUTE to verify that the user had authority to access the data set for the desired mode, and received a return code indicating that the user is not authorized for the requested function.

Return Codes: A reason code of x'98' is stored in the ACB error flags field, and the open request is failed with a return code of 8.

Action: Correct the error and rerun the job. Error could be caused by an incorrect data set name on the DD card, or this most likely is a legitimate error.

IAMW19 THIS VERSION OF THE IAM PROGRAM DOES NOT SUPPORT THE VSAM INTERFACE

Reason: The IAM system level VSAM Interface (VIF) loaded an IAM access/create module but found it to be at a level previous to the ones that support IAM's VSAM Interface (VIF).

VIF RETURN CODES: Reason code of X'F1' is stored in the ACB error field at displacement decimal +49 and the OPEN/CLOSE flags at displacement decimal +48 will not show the file as OPENed (X'10').

Action: If STEPLIB/JOBLIB was not specified, check the LINKLIST libraries to see if a prior version of IAM is ahead of the current version. If STEPLIB/JOBLIB was specified, verify that the library specified on that DD statement contains IAM Version 5.0 or higher.

IAMW20 DD= 'ddname' OPEN ERROR -- error description

Reason: The file load process has detected an invalid, unsupported, or inconsistent parameter for the indicated file attribute. The particular attribute that is invalid is indicated in the error description text. The possibilities are:

- LRECL Specification – Must be at least as long as key offset (RKP) plus the key length, and no larger than 32,755.
- Key Length Specification – Must be greater than 1, and less than 250.
- Blocking Factor Specification – Must be between 1 and 15.
- Blocksize Specification – Must exceed the maximum record length by at least five bytes, and be at least 300 bytes.
- RKP Specification – Must be less than 4092 bytes.
- Independent Overflow Specification – For compatible format files, cannot be a value that will cause the Overflow area to exceed 64,000 blocks.
- Integrated Overflow Specification – Must be between 0 and 99.
- Prime Extension Specification – For Compatible Format files, cannot exceed 32,767.
- Delete Processing Request – Should not occur with the IAM / VSAM Interface.
- Data Control Area (DCA) Validation – Should not occur with the IAM / VSAM Interface.

Return Codes: A reason code of 192(X'C0') is stored in the ACB error flags field (ACBERFLG) and the open is failed with a return code of 8.

Action: Correct the invalid specification, and rerun. This error may occur when using VIF if the IAM file was allocated through some means other than an IDCAMS DEFINE and the user failed to correctly provide file specifications through a CREATE Override Control Statement. Additionally, conflicts between the file itself and the user program's internal file definition must be resolved.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW21 DD='ddname' LOAD ERROR -- EXCEEDED 255 EXTENTS

Reason: During the loading of a multivolume file, the file required more than 255 extents, which is more than the number of extents allowed.

Return Codes: If failure occurred on a WRITE or PUT request, the RPL is failed with a Logical Error code of x'1C', file full logical error, and the request will receive a return code of 8. If the error occurred during CLOSE, then the CLOSE will fail with a return code of 4, and the ACB error flags set to x'90'. The file is left marked in an unloaded state.

Action: Delete and redefine the data set, increasing the primary and/or secondary space allocation values so that the file will fit within 255 extents.

IAMW22 DD='ddname' FILE REORGANIZATION RECOMMENDED -- reason

Reason: A file reorganization is being recommended for one of the following reasons, as indicated in the message:

1. OVERFLOW INDEX EXCEEDS 4 MEG
2. OVERFLOW EXCEEDS 1000 CYLINDERS
3. EXCEEDED 13 EXTENTS ON SINGLE VOLUME

Return Codes: None.

Action: For reasons 1 and 2, performance on the indicated file may be adversely affected due to the size of the overflow area. In particular, it may take several minutes to open the file, and sequential processing may be detrimentally affected. It is, therefore, recommended that the file be reorganized at the earliest convenient time to prevent further performance deterioration. The file can be quickly reorganized with FDRREORG®, or if that product is not available, then use IDCAMS REPRO.

For reason 3, because IAM files have a non-VSAM file structure, they are limited to 16 extents per volume. The file indicated currently has 14 or more extents, so future growth will be restricted. To prevent an out of space condition, action should be taken at the earliest possible time. If there is sufficient space for the file to expand on the volume that it currently resides, either use FDR/CPK® to merge extents, or reorganize the file doing a DELETE and DEFINE of the data set, specifying a larger space allocation. The current space allocation values can be determined by performing a LISTCAT ALL on the data set. If the current volume has insufficient space, then the file should be moved to a different volume, where more DASD space is available.

IAMW30 DD= 'ddname' OPEN ERROR -- FILE IN USE

Reason: The IAM file is currently in use by some other job or an attempt was made to simultaneously load multiple IAM files to the same data set.

VIF RETURN CODES: A reason code of 192(X'C0') is stored in the ACB error flags field (ACBERFLG) at displacement 49(X'31') and the open indicator, ACBOPEN(X'10'), in the ACB OPEN/CLOSE flags field (ACBOFLGS) at displacement 48(X'30') is not set.

NATIVE OR IFI RETURN CODES: A U0303 abend is issued.

Action: An IAM load must be executed without any other job accessing the file at the same time.

IAMW31 DD= 'ddname' OPEN ERROR -- NOT DASD DEVICE

Reason: The device the IAM file was being loaded to is not a disk or a supported DASD device.

VIF RETURN CODES: A reason code of 192(X'C0') is stored in the ACB error flags field (ACBERFLG) at displacement 49(X'31') and the open indicator, ACBOPEN(X'10'), in the ACB OPEN/CLOSE flags field (ACBOFLGS) at displacement 48(X'30') is not set.

Action: Check to see if the DD statement is pointing to the IAM file. An IAM file must reside on a DASD device. The version of IAM you are using may not support this DASD device. If necessary, call INNOVATION for further assistance.

IAMW32 DD= 'ddname' OPEN ERROR -- DD STATEMENT MISSING

Reason: The user failed to supply a 'ddname' statement for the file or used a DDNAME on the 'DD' statement that was different from the one coded in the program (or specified on a control statement).

VIF RETURN CODES: A reason code of 192(X'C0') is stored in the ACB error flags field (ACBERFLG) at displacement 49(X'31') and the open indicator, ACBOPEN(X'10'), in the ACB OPEN/CLOSE flags field (ACBOFLGS) at displacement 48(X'30') is not set.

Action: Check the IAM file's DD statement to ensure it is not misspelled.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW33 DD= 'ddname' OPEN ERROR -- FILE IN USE

Reason: The 'IAMWKDD' file is already in use. A program attempted to use the same work file data set name for the concurrent load of multiple IAM files. Unique work files must be specified for each IAM file when multiple files are to be created concurrently.

NOTE: This message will not appear when creating an IAM file under MVS, MVS/XA and MVS/ESA. A temporary work file will be dynamically allocated if the work file called for by the user is in use.

VIF RETURN CODES: A reason code of 192(X'C0') is stored in the ACB error flags field (ACBERFLG) at displacement 49(X'31') and the open indicator, ACBOPEN(X'10'), in the ACB OPEN/CLOSE flags field (ACBOFLGS) at displacement 48(X'30') is not set.

Action: Use the IAM CREATE Override Control Statement to specify unique WORK file ddnames for individual IAM files that will be loaded concurrently.

NOTE: If the user allows IAM to dynamically allocate the work file, this message will not appear.

IAMW34 DD= 'ddname' OPEN ERROR -- NOT DASD DEVICE

Reason: The 'IAMWKDD' device type is invalid. The DDNAME named in the message was not allocated to disk or a supported DASD device.

VIF RETURN CODES: A reason code of 192(X'C0') is stored in the ACB error flags field (ACBERFLG) at displacement 49(X'31') and the open indicator, ACBOPEN(X'10'), in the ACB OPEN/CLOSE flags field (ACBOFLGS) at displacement 48(X'30') is not set.

Action: Change the 'IAMWKDD' DD statement to point to a disk or supported DASD device.

NOTE: If the user allows IAM to dynamically allocate the work file, this message will not appear.

IAMW35 DD= 'ddname' OPEN ERROR -- DD STATEMENT MISSING

Reason: The user failed to supply a 'ddname' statement or used a DDNAME on the DD statement that was different from the one coded in the program (or specified on a control statement) for the 'IAMWKDD' file.

VIF RETURN CODES: A reason code of 192(X'C0') is stored in the ACB error flags field (ACBERFLG) at displacement 49(X'31') and the open indicator, ACBOPEN(X'10'), in the ACB OPEN/CLOSE flags field (ACBOFLGS) at displacement 48(X'30') is not set.

Action: Check the spelling of the IAMWKDD ddname used. Use the IAM CREATE Override Control Statement to specify a specific WORK file ddname.

NOTE: If the user allows IAM to dynamically allocate the work file, this message will not appear.

IAMW36 DD= 'ddname' OPEN ERROR -- BLOCKSIZE CALCULATION

Reason: The BLOCKSIZE that has resulted from IAM's calculation is not appropriate for the device in use.

VIF RETURN CODES: A reason code of 192(X'C0') is stored in the ACB error flags field (ACBERFLG) at displacement 49(X'31') and the open indicator, ACBOPEN(X'10'), in the ACB OPEN/CLOSE flags field (ACBOFLGS) at displacement 48(X'30') is not set.

Action: Possible logic error call INNOVATION for assistance.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW37 DD=DDNAME I/O ERROR ECB=XX CSW=XXXX SENSE=XXXX [OP RBN=BLOCK]

Reason: This message is displayed when an I/O error has occurred processing an Enhanced Format IAM file, or during a load of an IAM file. This message may also be accompanied by an IEA0001/IOS0001 error message. The message contains the following information:

ECB – displays high order byte of ECB as posted. Possible values include:

- 41 – I/O terminated with error. CSW and/or Sense bytes are useful.
- 42 – Block is not within DASD extents for data set.

CSW – contains the UNIT/CHANNEL status bytes from the CSW. Possible values include:

- 0C00 – Normal Status
- 0E40 – Wrong length record
- 0Dxx – End of file

SENSE – contains the first two sense bytes from the IAM internal IOB. Possible values include:

- 8000 – Command reject, the device or control unit did not recognize the command
- 4000 – Device requires operator Intervention
- 1000 – Equipment check
- 0800 – Data check
- 0020 – End of cylinder
- 0008 – Block not found

For file loads, the job step will be abended with a U0233 abend code.

For file access, additional information includes:

- op = RD error occurred on input operation or
- op = WR error occurred on output operation
- RBN = the relative block number being read or written

The job will be abended with a U0233 abend code.

Action: If an IBM IEA0001/IOS0001 message appears on the JCL LOG, a hardware error has occurred. The IBM message gives the sense information (ex: data check, equipment check). Examine the error information provided in the IBM message to determine the cause of the error. The format of this information is documented in the IBM data management SRL, for the operating system in use. The file in question is unreadable in its present state and must be reestablished. This may be the result of a hardware error and if possible the new file should be allocated to a different physical location or volume.

For other I/O errors, those not accompanied by a hardware failure error message, determine the cause of the error. It is recommended that an IAMRECVR DIAGNOSE function be executed on the problem file, to see if there are any problems with the file integrity. Some common reasons for these types of I/O errors include:

- File has been improperly moved or restored to a device type different than it was originally loaded on. Frequently, this will fail with block not found I/O error, with a SENSE error code of x'0008'.
- Multivolume file has been improperly moved, or improperly cataloged. Frequently, this will result in an I/O error of block not within extent, DECB error code of x'0010'.
- Storage overlay of IAM I/O control blocks. When this type of error occurs, most jobs and application programs do successfully process the IAM file, as does IAMRECVR. The failure is typically limited to one or a few jobs. A SYSUDUMP will be needed to determine the cause of this type of error.

The utility IAMRECVR may be used to recover a file that is no longer usable due to I/O or logical error conditions. This utility can be used to off load records from those portions of the file that have not been physically damaged. The user program should make a decision on the action to take when this occurs, for example, to continue processing without this file or to terminate processing until this file is made available. The appropriate action will depend entirely on the application and the user program's evaluation of the diagnostic information that is returned. If further assistance is needed, contact Innovation Data Processing.

IAMW38 DD= 'ddname' DSPSERV CREATE FAILED, RC=xx REAS=nnnnnnnn

Reason: IAM attempted to create a data space for holding the index structure during the file load, however, the request was rejected by MVS for the indicated return code and reason code. This is an informational message only, IAM will attempt to continue processing, and utilize a dynamically allocated temporary work file on DASD.

Action: Contact INNOVATION for assistance to resolve the problem.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW39 DD= 'ddname' ##### RECORDS ACCEPTED PRIOR TO ABEND

Reason: IAM has determined that the task loading the specified file has abended after loading the indicated number of records. This information may be useful for determining how much to adjust the space parameters if some type of Sx37 abend has occurred. The number of records that were actually written to the file may be slightly less than indicated due to buffering. Please note that the IAM file MUST BE RELOADED SUCCESSFULLY before attempting to otherwise access the file.

Action: Correct the error condition as indicated by the abend, and rerun the job. Depending on the cause of the abend, the file may have to be deleted and redefined, for example if a larger space requirement is necessary.

IAMW40 CARD IMAGE -- *cc....cc*

Reason: The input control statement(s) read from the 'IAMOVRID' DD statement is displayed when the control statement(s) contains an error or if requested by the user via the 'LOG=YES' operand.

Action: None.

IAMW41 CONTROL STATEMENT OPEN FAILED -- DDNAME= 'ddname'

Reason: The DDNAME listed was required as control statement input to the override processor, IAMOVRID. An OPEN was attempted, but failed. Processing of the override service is terminated. Normal processing continues.

Action: Review the execution job log messages for more detail on cause of the failure. Correct the 'ddname' statement and, if necessary, rerun the job.

IAMW42 INVALID CONTINUATION CARD

Reason: User coded a delimiting comma following the last keyword on a control statement input to the override processor, IAMOVRID, and neglected to provide the next logical record. Processing of the override service is terminated. Normal processing continues.

Action: Correct the control statement. The job will continue to run, but may fail if the Override data is critical.

IAMW43 I/O ERROR READING CONTROL STATEMENTS -- DDNAME= 'ddname'

Reason: An I/O error occurred reading the data set referenced by 'ddname'. IAMOVRID is terminated, but normal processing continues.

Action: Examine any system message(s) to determine the cause of the error. The format of system messages is documented in the IBM MESSAGE SRL for the operating system in use. The job will continue to run, but may fail if the Override data is critical.

IAMW44 CONTROL STATEMENT BYPASSED -- 'error description'

Reason: An error was encountered by the override processor, IAMOVRID, during the processing of user supplied Override Control statements. The error description will be from the following list:

1. DDNAME NOT SPECIFIED -- The DDNAME operand was missing or misspelled. DDNAME is required to relate the override to a specific IAM file.
2. MAXIMUM GLOBAL OVERRIDES EXCEEDED -- The in-storage table which holds the Overrides is full. A maximum of 200 control statements may be specified.

Action: Correct the Override statements as follows:

1. Add a corrected DDNAME operand to the control statement and, if necessary, rerun the job.
2. Reduce the number of global overrides to 200 control statements or less. If more Override statements are required, contact INNOVATION technical support for a modification to IAMOVRID to expand the in-storage table.

IAMW46 'ddname' OPEN FAILURE -- IAMNINFO PROCESSING TERMINATED

Reason: The output report 'ddname' statement could not be OPENed by IAMNINFO. This DDNAME is usually 'IAMINFO', but may have been overridden by the user. Processing continues without interruption with the IAM INFO report bypassed.

Action: Correct the allocation of DDNAME 'ddname' so the next execution of the job will produce the IAM INFO report.

IAMW47 I/O ERROR MONITORING DSN - 'dsname' - IAM MONITOR TERMINATED

Reason: An IAM monitor facility processor encountered an I/O error while writing monitor/trace data. Monitoring has been discontinued. Normal IAM processing continues.

Action: If a Monitor report is needed, correct the cause of the I/O error and rerun the job.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW48 IAMNINFO PARAMETER LIST ABSENT OR IN ERROR -- PROCESSING TERMINATED

Reason: The parameter list required by the dynamic file status display processor, IAMNINFO, was missing, was overlaid or is in error. This maybe an internal error. Processing continues without the IAM INFO reported printed.

Action: If you are unable to determine the reason for the message, call INNOVATION for further assistance.

IAMW49 IAMREORG PARAMETER LIST ABSENT OR IN ERROR -- PROCESSING TERMINATED

Reason: The parameter list required by the dynamic reorganization processor, IAMREORG, was missing or in error. The common errors are:

1. The pointer to the parameter list is missing.
2. The parameter list is not fullword aligned.
3. The parameter list contains invalid data.
4. Reserved bytes in the parameter list are not set to hexadecimal zeros. A U0590 ABEND is forced.

Action: Correct the processing program calling 'IAMREORG' for dynamic reorganization services. If you are unable to determine the reason for the message, call INNOVATION for further assistance AFTER obtaining a storage dump.

IAMW50 IAM VTOC ACCESS FAILED COMP=xxxx CODE=xxxx CLUSTER=clustername

Reason: During an IDCAMS DEFINE or RECATALOG of an IAM file, an attempt to access or update a VTOC failed. The completion code and return codes from CVAF are displayed in the message.

Action: Refer to the IBM manual 'Common VTOC Access Facility Diagnosis Reference' for the meaning of the codes. Correct the problem and re-submit. If unable to correct the problem, contact INNOVATION for assistance.

IAMW52 IAM SHOWCAT INTERCEPT FAILED

Reason: The IAM SHOWCAT intercept function within VIF failed.

Action: Obtain a core dump and call INNOVATION for assistance.

IAMW54 PROBABLE IAM FILE HAS NOT BEEN DEFINED, DSN=dsname

Reason: The processing program issued a SHOWCAT catalog request for a non-VSAM file that is cataloged, but does not have the IAM information in the OWNER ID field (as established by DEFINE). This message was issued by IAM's catalog request intercept facility.

Action: The catalog requests being issued by this program require that an IAM file used under VIF be DEFINEd as an INDEXED VSAM CLUSTER via IDCAMS in an ICF or VSAM catalog. Check to ensure that this IAM file was not DEFINE'd as NON-VSAM or that it was not cataloged via JCL.

IAMW55 PROBABLE IAM FILE IS CATALOGED IN CVOL, DSN=dsname

Reason: The processing program issued a SHOWCAT catalog request for a non-VSAM file that is cataloged in a CVOL. This message was issued by IAM's catalog request intercept facility.

Action: The catalog requests being issued by this program require that an IAM file used under VIF be DEFINEd as an INDEXED VSAM CLUSTER via IDCAMS in an VSAM or ICF catalog. Check to ensure that this IAM file was not DEFINE'd into a CVOL or that it was not cataloged via JCL into a CVOL.

IAMW56 IAMDEFINE OF NON-SUPPORTED FILE TYPE CLUSTER=dsname

Reason: An IDCAMS DEFINE was issued with an indication that the file should be an IAM file, but the file type cannot be converted to IAM. IAM supports single index KSDS files, and ESDS files.

Action: Correct the IDCAMS DEFINE to either change the file type to one supported by IAM, or remove the indication that the file is to be an IAM file.

IAMW57 IAM ALLOCATION FAILED COMP=xxxx CODE=xxxx CLUSTER=clustername

Reason: An error occurred during an IDCAMS DEFINE of an IAM file. The codes displayed correspond to the return code and reason codes of the IDC3009I message from IDCAMS, and there will also be a IDC3009I message on SYSPRINT with the same codes. There may also be additional IDC or IAMW messages.

Action: Refer to IDCAMS error message IDC3009I for meaning of the codes. Correct the problem, and resubmit. It may be necessary to issue an IDCAMS DELETE command before attempting to resubmit the DEFINE.

IAMW58 LISTC INTERCEPT FAILED COMP=xxxx CODE=xxxx CLUSTER=UNKNOWN

Reason: A LISTC catalog request failed and received xxxx as the completion and return codes.

Action: Refer to VSAM errors message IDC3009I for meaning of the codes. Correct the problem and resubmit.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW59 LOCATE FOR AN IAM FILE FAILED CODE=(cc)xxx (DSN=.....)

Reason: Locate for an IAM file failed for one of the following reasons:

1. CODE=S1xxx The IAM SHOWCAT intercept issued a locate which failed with return code xxxx on the specified data set.
2. CODE=S2xxx The IAM SHOWCAT intercept issued a locate which failed with return code xxxx. The locate was issued by Cnumber, so the data set name is unknown.
3. CODE=L1xxx The IAM LOCATE intercept issued a locate which failed with return code xxxx. The locate was issued by Cnumber, so the data set name is unknown.

Action: Make sure that the data set is still properly cataloged. If not, an IDCAMS DEFINE RECATALOG must be done. If further assistance is required, contact INNOVATION.

IAMW60 IAM DYNALLOC FAILED 'description'

Reason: During the processing of an IDCAMS DEFINE for an IAM file, IAM'S attempted use of Dynamic Allocation failed for the specified reason. This message presents a brief English description of the error code returned by Dynamic Allocation, which is supplied in the IAMW61 error message. Both messages are printed on the system log. There will also be an IDC3009I message on SYSPRINT, with an appropriate error code.

Action: Correct the error situation, as described with the matching text below, and rerun the DEFINE.

TEXT: DATA SET NAME IN USE BY ANOTHER JOB/USER

Reason: The dataset name has been enqueued on by another job/user.

Action: Through whatever software facilities available, determine which job and/or users are enqueued on the data set, and rerun the DEFINE upon the completion of the other job/user.

TEXT: VOLUME NOT MOUNTED ON SPECIFIED UNIT

Reason: The specified volume was either not mounted, or was mounted but not on the unit specified by the UNIT= keyword on the IAM override control statement for this file. For non-specific volume requests, (i.e., with VOL(ANYVOL) coded), there were no volumes mounted as storage for the unit name specified on the IAM Override Control statement, or SYSDA.

Action: Mount the required volume, or change the volume and/or unit specification.

TEXT: SPECIFIED UNIT NAME IS UNDEFINED

Reason: The unit name specified on the IAM Override Control statement for this file does not exist on the system that the define was attempted.

Action: Correct the unit name specification, or run on the proper operating system.

TEXT: REQUIRED CATALOG NOT MOUNTED

Reason: The catalog required for the definition of the IAM data set is on a volume that is not currently mounted.

Action: Insure that the volume containing the user catalog is mounted, and rerun the DEFINE.

TEXT: DATA SET ALREADY EXISTS

Reason: The data set being DEFINEd already exists in the catalog, and may or may not be on the volume it is cataloged to.

Action: Make sure that the cluster name is correct and if not correct it. If it is correct, delete the data set from the catalog (and volume if applicable) and rerun the DEFINE.

TEXT: DUPLICATE DATA SET NAME ON VOLUME

Reason: The data set already exists on the specified volume, and is not cataloged.

Action: Delete the data set from the volume, and rerun the DEFINE.

TEXT: NO SPACE IN VTOC

Reason: There was no space in the VTOC (Volume Table of Contents) for the new data set on the specified or selected volume.

Action: Either correct the error by increasing the size of the VTOC on the volume (this can be done by use of IDP's COMPAKTOR), delete unwanted data sets from the volume, or select a different volume.

TEXT: VTOC I/O ERROR OR CVAF ERROR

Reason: An I/O error occurred on the VTOC during file allocation.

Action: Review SYSLOG for other messages indicating a more precise cause of error. Correct the problem and rerun DEFINE.

TEXT: REQUESTED SPACE NOT AVAILABLE ON VOLUME

Reason: The volume specified or selected did not have sufficient space to satisfy the request.

Action: Ensure that the space requested is actually needed, and adjust if possible. (**NOTE:** IAM files generally require less space than VSAM files.) Otherwise, select a different volume, remove unneeded data sets from the volume, or run COMPAKTOR to consolidate free space.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

TEXT: USER NOT AUTHORIZED TO ALLOCATE DATA SET

Reason: The job lacks RACF authorization to DEFINE the data set.

Action: Contact the Security Administrator for assistance.

TEXT: INSTALLATION EXIT REJECTED ALLOCATION REQUEST

Reason: A dynamic allocation exit routine in the system did not allow the allocation request to be processed.

Action: Correct the DEFINE to the installation requirements.

TEXT: REQUIRED CATALOG NOT AVAILABLE

Reason: The user catalog required may have been DELETED or disconnected from the system master catalog, or may have been damaged and is being recovered.

Action: Correct the error with the user catalog, and rerun the DEFINE command.

TEXT: DUPLICATE DATA SET NAME IN CATALOG

Reason: The data set name already exists in the catalog, and may or may not exist on disk.

Action: Make sure the cluster name is correctly specified. If it is delete the current entry from catalog (and disk if applicable).

TEXT: NO SPACE IN CATALOG

Reason: Insufficient space in the catalog to contain the record for the new data set.

Action: Enlarge the catalog, and rerun the DEFINE.

TEXT: SMS FAILED REQUEST. REFER TO PRIOR MESSAGE(S)

Reason: The allocation request was failed by SMS. There should be preceding messages from SMS indicating the reason for the error.

Action: Correct the problem indicated by the SMS error messages, and try request again.

IAMW61 IAM DYNALLOC FAILED COMP=nnnn CODE=nnnn CLUSTER=clustname

Reason: The dynamic allocation requested by IAM to perform the DEFINE operation failed with the printed error codes. This message may be accompanied by an IAMW60 message.

Action: Refer to message IAMW60, if printed, and/or the IDC3009I error message on SYSPRINT. The error codes from Dynamic Allocation are documented in the MVS/XA and MVS/ESA System Macro and Facilities manual, the MVS Job Management SPL, and under the ISPF tutorial. Correct the error condition as indicated by the error codes, and rerun the DEFINE.

IAMW62 IAM OPEN FAILED FOR DDNAME=ddname CLUSTER=clustname

Reason: During DEFINE processing of an IAM file, IAM attempted to OPEN the DEFINED file, however the OPEN failed. Additional IBM messages may appear on the system log.

Action: Determine the cause of the OPEN failure, correct the error, and rerun the job. For a new DEFINE (as opposed to RECATALOG), DELETE and re-DEFINE the data set. NOTE: At this point, the data set has been allocated and cataloged, but is not yet usable by IAM.

IAMW63 IAM I/O ERROR: 'synad message'

Reason: During the processing of a DEFINE command for an IAM file, an I/O error occurred when reading or writing the IAM control information.

Action: Using the standard SYNAD message and other messages that may appear on SYSLOG, determine the cause of I/O error and correct it. If this was not a RECATALOG operation, DELETE and DEFINE the IAM data set again.

NOTE: At this point, the data set has been allocated and cataloged, but is not yet usable by IAM.

For RECATALOG operations, the file is either not a previously DEFINED or loaded IAM file, or there is an error with the data set requiring recovery. The recovery can be done by restoring the data set from a good copy or possibly by using program 'IAMRECVR'. Use of the recovery program may result in data loss.

IAMW64 UNEXPECTED END OF FILE READING AN IAM FILE FOR RECATALOG REQUEST -- NOT VALID IAM DATA SET

Reason: During the processing of a DEFINE RECATALOG command for an IAM file, an end of file occurred while attempting to read the file characteristics.

Action: The data set is empty. The recatalog request was not performed. Either the file was never an IAM file, in which case no corrective action is required, or the data set has been clobbered. To recover the data set, it can be restored from a good backup, or a recovery attempted with program 'IAMRECVR'. Recover the file, then retry the recatalog processing.

IAMW65 IAM SCRATCH FAILED COMP=nnnn CODE=nnnn VOLSER=vvvvvv

Reason: After an error attempting to catalog an IAM file being DEFINED, an attempt to DELETE the data set from the indicated volser failed. The codes are returned from SCRATCH, which are documented in the SYSTEM DATA ADMINISTRATION manual.

Return Codes: The return code for the DEFINE is based on the original condition that caused the error.

Action: The data set is still on the specified volume. Refer to an immediately preceding IAMWnn message for the data set name. The data set must be manually scratched from the indicated volume.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

- IAMW66 IAM REALLOC FAILED CODE=nnnn INFO=nnnn CLUSTER=clustname**
Reason: After successfully defining an IAM file, IAM had determined that the job step had DD cards which were allocated to the file, but were allocated to the wrong volume. The attempt to reallocate the file with dynamic allocation failed, with the indicated error codes.
Return Codes: The DEFINE completes with a return code of 0.
Action: The IAM file has been successfully defined, but attempts to REPRO into the IAM file within the same step may fail. A subsequent REPRO into the IAM file can be done.
- IAMW67 IAM SMS ALLOC FAILED RC=X'xx' REAS=X'xxxxxxxx' CLUSTER=clustname**
Reason: The define of an IAM file failed using DADSM allocation with the specified return code and reason code. Refer to the DADSM Create (ALLOCATE) Function Return Codes section of the IBM MVS/ESA DADSM/CVAF Diagnostic Aids for a description of the error codes.
Action: Take the appropriate corrective action based on the error codes indicated, and retry the define request.
- IAMW68 IAM UNIT NAME SEARCH FAILED, RC=xx CLUSTER=clustname**
Reason: During the define of a multi-volume non-specific allocation, an IAM call to the MVS Unit Name look up service failed as indicated in the message. The return code, if provided, is documented in the IBM MVS System Modifications Manual.
Action: If the condition indicated by the return code cannot be corrected, contact INNOVATION for support. As a circumvention, try a different UNIT override, or switch to specific volume allocation.
- IAMW69 IAM xxx ELIGIBLE VOLUMES, nnn VOLUMES REQUIRED CLUSTER=clustname**
Reason: During the define of a multi-volume non-specific allocation, IAM found xxx storage volumes in the specified UNIT name pool, however, nnn volumes were needed to satisfy the allocation request.
Action: Change the UNIT override to indicate a UNIT name that has sufficient storage volumes, or reduce the number of volumes requested.
- IAMW70 DD= 'ddname' PUT ERROR -- WORK FILE DATASPACE OVERFLOW**
Reason: The size of the data space used to hold the index structure during a file load was insufficient for the file indicated by the ddname.
VIF RETURN CODES: A reason code of 244(X'F4') is stored in the RPL error flags field (RPLERRCD) at displacement 15(X'0F'). If the error was detected during close, the program is abended with a U0246 abend.
Action: Increase the size of the data space by using the IAM CREATE override, specifying the DATASPACE keyword. As an alternative, specify a DATASPACE=0, which will force the use of a temporary work file.
- IAMW71 TRACE DEACTIVATED - text indicating reason**
Reason: The IAM ESDS trace facility has detected an error during activation. The possible reasons include A) the DDNAME IAMATR31 is not available, or B) there is insufficient storage available to obtain the trace work area.
Action: If the IAMATR31 DD was not specified, add it to the job. If storage was not available for the trace work area, a REORG of the IAM file may be required.
- IAMW79 DD= 'ddname' OPEN ERROR - text indicating reason**
Reason: IAM was unable to open the specified file due to damaged control or index information on the file.
VIF RETURN CODES: A reason code of 188 (X'BC') is stored in the ACB error flags field (ACBERFLG) and the OPEN is failed with a return code of 8. If an IAMDEBUG DD DUMMY is in the job step, then the job will be abended with a U0184.
Action: Use the IAMRECVR to recreate the file. For diagnosis, please do the following:
1. Rerun job with an //IAMDEBUG DD DUMMY card added and a //SYSUDUMP or //SYSABEND. If you have the ABENDAID, please insure that the standard IBM dump will be taken.
2. Save the damaged file, or back it up using FDR/DSF, or DFDSS.
3. Contact Innovation Data Processing for assistance.
- IAMW80 MODULE 'modname' INSTALLED AT 'address' - VER nn**
Reason: The message is issued in response to a VIF status request when the module 'modname' is in place and is ready to provide IAM services to programs using ACB's to access VSAM files. The 'address' given is the virtual storage location of the named IAM VIF module. The version level number is also given for the module.
Action: None, information message only.
- IAMW81 THE IAM SYSTEM MODULES ARE 'status'**
Reason: This message is provided by IAMSTART when the IAM VIF modules are installed, or in response to VIF action commands when the IAM system level VSAM interface is already in place. The 'status' of the modules may be, ACTIVE, REACTIVATED, or ALREADY INSTALLED. This message includes the version and release level of the VIF modules installed.
Action: None, information message only.

80.02 WTO/WTOR CONSOLE AND TSO TERMINAL MESSAGES

IAMW82 THE IAM SYSTEM MODULES ARE 'status'

Reason: This message is provided by IAMSTART in response to VIF action commands when the IAM system level VSAM interface is already in place. The 'status' of the modules may be, NOT ACTIVE (the IAM system level VSAM interface is not in place), INACTIVE (the IAM system level VSAM interface is in place but is not active).

Action: None, information message only.

IAMW84 IDPSTART ESTAE RECOVERY ENTERED FOR ABEND Sxxx Uxxxx AT OFFSET xxxx

Reason: This message is provided by IAMSTART when VIF ABENDS. An attempt is made to provide diagnostic information.

Action: If the problem persists call INNOVATION for assistance.

IAMW86 IDPSTART FAILURE - REASON= 'number' 'reason'

Reason: The activation of the IAM VSAM Interface failed for one of the following reasons:

- 01 OPERATING SYSTEM NOT MVS OR SP 1.2 OR HIGHER
- 02 INVALID INPUT PARAMETERS
- 03 UNABLE TO OBTAIN AUTHORIZATION
- 04 'vector table name' VECTOR TABLE IS INVALID
- 05 SYSLIB DD MISSING OR OPEN ERROR
- 06 'modname' NOT FOUND IN SYSLIB
- 07 'modname' MODULE LOAD ERROR
- 08 ERROR MODIFYING THE SVC TABLE
- 09 'modname' NOT FOUND
- 10 'modname' NOT FOUND IN THE LINKLIST
- 11 'modname' UNABLE TO DE-INSTALL
- 12 'modname' CDE OR LPDE ABOVE 16M ERROR
- 13 'modname' MODULE ABOVE 16M ERROR
- 14 GETMAIN ERROR
- 15 'modname' HAS SMP INSTALLED IDP MODULE
- 16 UNABLE TO OBTAIN DISPATCHER LOCK
- 17 FREEMAIN ERROR
- 18 UNABLE TO OBTAIN LOCAL LOCK
- 19 'modname' CDE NOT FOUND
- 20 ENQUEUE/DEQUEUE ERROR ON IAM RESOURCE
- 21 ENQUEUE/DEQUEUE ERROR ON SYSZSVC
- 22 INVALID VECTOR TABLE STATUS
- 23 INVALID VECTOR TABLE CDE ADDRESS
- 24 RC=xxxx FROM SVCUPDTE
- 27 'modname' and IAMVECTB ARE OUT OF SYNCH
- 28 VECTOR TABLE CANNOT BE STOPPED
- 31 HIGHER VERSION OF VIF ALREADY STARTED

Action: If the problem persists call INNOVATION for assistance.

IAMW89 IAM - TRIAL VERSION FROM INNOVATION DATA PROCESSING EXPIRES IN 'nnn' DAYS (PLEASE CONTACT INNOVATION)

Reason: This is a trial version of the IAM system. The number of days the trial will remain active is displayed.

Action: When there are 10 or fewer days before the trial is due to expire this message will become non-deletable. To prevent the trial from expiring call INNOVATION for an extension PARM= value and use the JCL shown below to extend your trial's expiration date.

The JCL to extend the trial is as follows:

```
//EXTEND EXEC PGM=IAMEXTND, PARM=xxxxx
//STEPLIB DD DISP=SHR, DSN=your.user.lib
//SYSLIB DD DISP=SHR, DSN=your.user.lib
//SYSDIAG DD SYSOUT=A
```

IAMW99 INTERNAL LOGIC ERROR -- JOB TERMINATED

Reason: IAMOVRID has encountered an illogical condition.

Action: Obtain a core dump and call INNOVATION for assistance.

80.03 IAM SYSTEM SYSPRINT MESSAGES

80.03 IAM SYSTEM SYSPRINT MESSAGES

The following general messages are used by the IAM system programs. These messages are written to a SYSOUT print data set. The DD statement the message is written to depends upon the particular IAM program issuing the message.

IAM001 function prog VER nn INNOVATION DATA PROCESSING DATE yyddd PAGE nn.

Reason: Generalized page header for the named processing functions which are provided by the programs named in the printed title.

Action: None, information message only.

IAM002 CARD IMAGE

Reason: An image of the first eighty (80) positions of each input statement read by the IAM utility program.

Action: None, information message only.

IAM003 IAM FILE SUCCESSFULLY CREATED

Reason: A completion message noting that an IAM file has been successfully created by the utility program.

Action: None, information message only.

IAM004 DUMP SUCCESSFULLY COMPLETED

Reason: A completion message noting that the dump of an IAM file has ended normally and a sequential copy of the IAM file was successfully created. This dataset can be kept as a backup of the IAM file and can be used to recreate the IAM file.

Action: None, information message only.

IAM005 REORGANIZATION SUCCESSFULLY COMPLETED

Reason: A completion message noting that the RESTORE phase of a reorganization has ended and the IAM file has been successfully reorganized. The sequential dataset created during the DUMP phase can be kept as a backup of the file for recovery at a later time.

Action: None, information message only.

IAM006 COPY SUCCESSFULLY COMPLETED

Reason: A completion message noting that an IAM file has been successfully created by the utility program as a copy of an existing IAM file.

Action: None, information message only.

IAM007 RECREATE SUCCESSFULLY COMPLETED

Reason: A completion message noting that a VSAM (or ISAM) file has been successfully created by the utility program using an existing IAM file as the source of the data.

Action: None, information message only.

IAM090 DIAG MINI DUMP IAMDMPRT VER nn - DATE/TIME - yyddd/hh.mm.ss PAGE nn

Reason: General page heading for the MINI DUMP processor listing the version, level, date, time and page number.

Action: None, information message only.

IAM091 REGS general registers (storage)

Reason: Data formatted by the MINI DUMP processor, printing from 1 to 16 general registers and, optionally, related storage. All register oriented data is identified by this message number.

Action: Take corrective action as needed.

80.03 IAM SYSTEM SYSPRINT MESSAGES

- IAM092 type LEN=nnnn(hhhh) LOC=xxxx**
Reason: Control block formatted by the MINI DUMP processor. 'type' is the name of the control block, 'nnnn' is the length in decimal, 'hhh' is the length in HEX, and 'xxxx' is the HEX storage location.
Action: Take corrective action as needed.
- IAM093 CCWS LEN=nnnn(hhhh)**
Reason: CCWS formatted by the MINI DUMP processor.
Action: None, information message only.
- IAM094 reason - CCWS NOT PRINTED(TERMINATED)**
Reason: An IAM related program requested a MINI DUMP with CCWS formatting. However, for the reason printed, CCWS printing was terminated.
Action: Contact INNOVATION for further assistance.
- IAM099 NEAR RELATIVE DATA LOCATION nnnn 'error description'**
Reason: The common parsing routine encountered an error in parsing the user specified control statements. The approximate location of the error was position nnnn, counting the first position as 000.
Action: Correct the error and re-submit job.
- IAM100 IAM File analysis information**
Reason: Lists criteria describing the IAM file and the processing performed. The following information is provided:
- **Items common to all IAM file types**
 - FILE FORMAT:* ENHANCED or COMPATIBLE. Beginning with IAM version 6.3, a new ENHANCED file format is available. The old pre-allocated file structure is still available, and is described as COMPATIBLE.
 - FILE STATUS:* Describes if the IAM file has been LOADED with data, or has been DEFINED but not loaded.
 - RECORD SIZE:* The logical fixed or maximum variable record length in the IAM file, given in bytes.
 - BLOCK SIZE:* The actual block size of the IAM file, given in bytes.
 - BLOCK FACTOR:* The blocking factor specified when this file was last DEFINEd, loaded or reorganized. If the file has been DEFINEd, but not yet loaded, this value will reflect the CISIZE, if any, specified on the DEFINE, which will be converted to a blocking factor when the file is loaded.
 - KEY SIZE:* The length of the record key field, in bytes.
 - REL KEY POS or KEY OFFSET:* The location of the key relative to the start of the record (RKP), in bytes. The first byte in a record is RKP=0. This amount will be 4 higher than the value DEFINEd for files with variable length or compressed records, due to IAMs internal RDW.
 - FILE TYPE:* KSDS or ESDS. Describes the type of IAM file.
 - DEVICE TYPE:* Indicates the disk device type that the IAM file was originally DEFINEd or loaded on.
 - VOLSER:* This is the first volume on which the IAM file currently resides on. For multi-volume files, additional VOLSERS (up to 4) are printed on subsequent lines. The additional VOLSERS may represent candidate volumes, and may not yet contain any extents of the IAM file.
 - NO OF TRACKS USED:* The actual number of tracks used to accommodate the IAM file in its entirety, including overflow areas.
 - RELEASE:* Indicates whether or not IAM will automatically release space the next time the file is loaded or reorganized. Once an IAM file is loaded and the unused portion of its allocation is released, IAM will automatically change RELEASE=YES to RELEASE=NO, unless overridden by an IAM CREATE override statement.
 - SHARE OPTIONS:* Indicates the first VSAM share option that was specified when the IAM file was DEFINEd.
 - DATA COMPRESS:* Indicates whether or not the IAM file can contain compressed data. Values are either NO or ENABLED. ENABLED indicates that the DATA COMPRESS option was either specified as YES or defaulted to YES based on the global options table, and the IAM file met all other required criteria for data compression. NO indicates the IAM file does not contain any compressed data records.
 - TOTAL RECORDS:* Indicates the total number of user data records currently residing in the file.
 - INSERTS:* Indicates the total number of records added to the file since it was last loaded or reorganized.
 - UPDATES:* Indicates the total number of records updated in the file since it was last loaded or reorganized.
 - DELETES:* Indicates the total number of records deleted from the file since it was last loaded or reorganized.
 - NOTE:* The above four statistics may not be accurate if an update job failed to successfully CLOSE the IAM file. This would only occur if the address space failed, an update job was removed from the system with the force command or a system failure occurred.
 - FILE DEFINED:* Indicates the date and time the file was DEFINEd.

80.03 IAM SYSTEM SYSPRINT MESSAGES

FILE LOADED: Indicates the date and time the file was LOADED.

STORAGE REQUIRED FOR COMPRESSED KEY STRUCTURE: The amount of virtual storage required for this IAM file's index. This message only appears if the IAM file has a compressed index.

STORAGE REQUIRED FOR KEY STRUCTURE: The amount of virtual storage required for this IAM file's index. This message only appears if the IAM file does not have a compressed index.

NUMBER OF IAM DATA BLOCKS: The number of data blocks contained within the IAM file.

- **Items specific to COMPATIBLE format KSDS IAM files**

INTEGRATED OVERFLOW: The Integrated Overflow (CI%) percentage specified when the file was last defined, loaded or reorganized.

ACTUAL INTEGRATED OVFL: Indicates the percentage used in establishing the Integrated Overflow (CI%) for the IAM file. This figure may be lower than the requested Integrated Overflow due to rounding based on the record length and actual blocksize being used.

INDEPENDENT OVFL REQ: The number of Independent Overflow records specified when this IAM file was last defined, loaded or reorganized.

INDEPENDENT OVERFLOW: The number of physical blocks IAM reserved for Independent Overflow records. This number is developed by IAM during file loading using the number of Independent Overflow records requested, the maximum record size, and the actual blocksize being used.

INDEPENDENT OPEN: Specifies the percentage of Independent Overflow keys which will be reserved when the file is OPENed for update operations, as specified by the OCOREO% keyword.

NOTE: This field is reported as 0 (zero) until the file has been loaded, unless overridden at DEFINE time.

PRIME EXTENSION: The number of blocks reserved by IAM for resume load type of operations when this file was last defined, loaded or reorganized.

NOTE: This field is reported as 0 (zero) until the file has been loaded, unless overridden at DEFINE time.

PRIME EXTENSION USED: Indicates the number of Prime Extension blocks that are currently being used.

PERCENT AVAILABLE IN INDEPENDENT OVERFLOW: Indicates the percentage of the total Independent Overflow records that are currently available. A low value, in the range of 10 to 25, provides an indication that the file needs to be reorganized.

NUMBER OF RECORDS IN INDEPENDENT OVERFLOW: Indicates the number of records currently residing in the Independent Overflow area of the IAM file.

NO OF UNUSED RECORDS IN INDEPENDENT OVERFLOW: The number of additional records the Independent Overflow area can accommodate. This value added to the value in NUMBER OF RECORDS IN INDEPENDENT OVERFLOW may be slightly higher than the number of Independent Overflow Records Requested, because IAM uses a fully blocked overflow area.

- **Items specific to ENHANCED format KSDS and ESDS IAM files**

CI SIZE: The CONTROLINTERVALSIZE used that last time the IAM file was defined. This value can be specified or can be calculated by IAM.

HIGH USED RBA: Indicates the highest Relative Byte Address (RBA) in use within the IAM file.

HIGH ALLOCATED RBA: Indicates the highest allocated Relative Byte Address (RBA) within the IAM file.

FREESPACE - CI% and FREESPACE - CA%: The values specified in the IDCAMS FREESPACE(CI% CA%) parameter the last time the IAM file was defined.

EXTENDED OVERFLOW: The number of Extended Overflow blocks that are currently in use.

EXTENDED ALLOCATED: The total number of Extended Overflow currently allocated in the IAM file.

EXTENDED AVAILABLE: The number of Extended Overflow blocks currently available for use within the IAM file.

EXTENDED PE: The number of Extended Prime Extension (PE) blocks currently in use within the IAM file.

VOLUME COUNT: The number of volumes currently in use by the IAM file.

TOTAL EXTENTS: The total number of extents that the IAM file contains.

PRIMARY SPACE: The primary space value used the last time the IAM file was defined.

SECONDARY SPACE: The secondary space value used the last time the IAM file was defined.

MULTIVOLUME: Specifies the allocation type (PRIMARY, SECONDARY) that will be used when a multivolume IAM file requires a new volume. Normally, an IAM file will only obtain the primary allocation on the first volume of a multivolume file. This option allows IAM to modify the value to be the same as the primary value.

MAX SECONDARY: Indicates the largest secondary that IAM will request when obtaining a new extent. The MAXSECONDARY global option and IAM override values are specified as a multiplication factor to be used to calculate a larger secondary value if the IAM file has greater than five (5) extents. The value shown here is the actual largest value that IAM will use after its calculations.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM102 CREATE OF AN IAM FILE FAILED REASON= 'error description'

Reason: A record loaded from the input data set failed for one of the following reasons:

- A. OUT OF SPACE – The IAM file does not contain sufficient space for the number of records being loaded.
- B. DUPLICATE RECORD PASSED – The current record has the same key as the previous record.
- C. OUT OF SEQUENCE – The current record key is out of sequence.

Action: Correct the error condition. This conflict must be resolved before IAMUTIL will create the IAM file.

IAM103 TAPE I/O ERROR OR FAILED TO OPEN - IAMUTIL TERMINATED

Reason: An error condition most likely resulting from one of the following reasons:

- 1. The TAPE1 DD statement was missing, misspelled, or there was incorrect DCB information on the TAPE1 DD.
- 2. A physical I/O error was detected.

Action: Check the job log for system error information and, if possible, correct the error condition.

IAM104 INVALID CONTROL CARD - REASON 'number' - JOB TERMINATED

Reason: The control statement listed above contains at least one error, the first of which is identified by the reason code 'number'.

- 1. Control statement was blank up to column 60.
- 2. The operation keyword was invalid (ex: CREATE, DUMP, etc.).
- 3. An option keyword was invalid (ex: L=..., K=..., B=..., etc.)
- 4. FROM=C' " or TO=C' " was invalid. The keyword did not end in a quote or the number of characters specified was greater than 40.
- 5. A keyword specified invalid numerics after the equal (=) sign.
- 6. A numeric value was greater than 15 digits in length.
- 7. An option keyword did not end with a blank or a comma.
- 8. The FROM key value was greater than the TO key.
- 9. No input control statements were read. One or more control statements must be present.
- A. Expected continuation statement was not received. A comma followed by a blank was found on the preceding statement.
- B. A required option keyword was not found.
- C. Keyword FROM=X or TO=X was invalid. The keyword did not end in a quote or the data within the quotes contained invalid Hex Data.
- D. The DEXIT or REXIT module name specified on the control statement is invalid or contains greater than eight (8) characters.

Action: A U0502 ABEND is forced. Correct the input statement shown in the preceding message and resubmit the job.

IAM105 SIMULATION REPORT - STARTING TIME nnnn - ENDING TIME nnnn

Reason: Summary heading line for a report indicating the number of inquires against a SIM file in the specified time period. The number of keys found is also reported.

Action: None, information message only.

IAM106 KEY xxxx WAS ADDED (EXISTS) DELETED (NOT DELETED)

Reason: Detail line for a report indicating the number of ADD or DELETE commands against a SIM file and the number that failed.

Action: None, information message only.

IAM107 FILE IS FULL ADD TERMINATED

Reason: Detail line for a report indicating the user exceeded the capacity of the overflow areas on this add to a SIM file. This add operation is terminated.

Action: None, information message only.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM108 OUTPUT TAPE BLKSIZE IS nnnnn BYTES.

Reason: This message is issued by DUMP and REORG. The BLKSIZE specified is either what was specified by the user on the TAPE1 JCL or calculated by IAMUTIL. IAMUTIL will always round the requested blocksize down to a multiple of the LRECL for fixed format files.

Action: If a tape is not labeled, this blocksize must be specified on the JCL for any subsequent restore operation.

IAM109 TAPE RECFM BLKSIZE OR LRECL WAS NOT specified - IAMUTIL TERMINATED

Reason: The TAPE1 DD statement did not specify required JCL DCB parameters.

Action: If the tape is not labeled, the blocksize and LRECL must be specified on the JCL for any input operation.

IAM110 DISKIN OPEN ERROR OR I/O ERROR - JOB TERMINATED

Reason: An IAM create was requested with a VSAM (or ISAM) file specified as input but the DISKIN DD was missing, failed to open or experienced an I/O error during processing.

Action: Check the job log for system error information and if possible correct the error condition.

IAM111 DISKOUT OPEN ERROR OR I/O ERROR - JOB TERMINATED

Reason: An IAM recreate was requested with a VSAM (or ISAM) file specified as output but the DISKOUT DD statement was missing, failed to open or experienced an I/O error during processing.

Action: Check the job log for system error information and if possible correct the error condition.

IAM211 'function' PROCEEDING 'to/from' DDNAME 'ddname'

Reason: The named 'function' is being executed 'to' or 'from' the DDNAME 'ddname' listed.

Action: None, information message only.

IAM212 FILE CHARACTERISTICS

Reason: Describes the file characteristics of the output data set referenced by the DDNAME 'ddname' in the preceding IAM211 message.

BLKSIZE – The blocksize of the output file.

LRECL – The logical record length of the file.

RECFM – The record format of the output file, either FIXED or VARIABLE blocked.

Action: None, information message only.

IAM213 'file type' RECOVERY FILE CREATED DDNAME= 'ddname'

Reason: IAMRECVR has completed the output of a file to the DDNAME 'ddname'. The 'file type' is as follows:

SEQUENTIAL – Sequential file of data records. The file can be used as input to a file create.

IAM – An IAM file created by the RECOVER operation.

DUPLICATE – A key/log file containing duplicate records found in the IAM file that is being recovered by IAM-RECVR.

Action: Refer to the documentation of IAMRECVR for the uses of the different output files it creates.

IAM214 'function' COMPLETED 'to/from' DDNAME 'ddname'

Reason: The named 'function' has been successfully executed 'to' or 'from' the DDNAME 'ddname' listed.

Action: The DUMP function completes with the creation of a sequential copy of the records read from the IAM file. This is the file addressed by the 'ddname' statement named in the message. It can be kept as a backup of the IAM file and can be used to recreate the IAM file if the RESTORE function does not successfully complete.

IAM260 LISTC PROCESSING FAILED FOR DSN= 'dsname'

Reason: An error occurred during IAM processing of an IDCAMS LISTC command for a possible IAM file. This message will be preceded by one or more error messages describing the error in more detail.

Action: Review other messages for source of error, and correct as necessary. IDCAMS LISTC processing will continue normally, however no information will appear in the IAMPRINT file for the named file.

80.03 IAM SYSTEM SYSPRINT MESSAGES

- IAM261 DYNAMIC ALLOCATION FAILED,RETURN CODE= REASON CODE= (DSN= 'dsname')**
Reason: IAM attempted to dynamically allocate a file, however the allocation failed with the specified codes. The dataset name if not displayed in this message is presented in the accompanying IAM260 message.
Action: Look up the Dynamic Allocation error codes in the IBM System Macros and Facilities manual or the ISPF Tutorial and take appropriate corrective action, for assistance contact INNOVATION.
- IAM262 DATA SET IS NOT AN IAM FILE**
Reason: During an IDCAMS LISTC IAM attempted to process a data set as an IAM file, however it was determined that the data set was not an IAM file. The data set name is displayed in the IAM260 message.
Action: None, information message only.
- IAM263 UNABLE TO OBTAIN VIRTUAL STORAGE FOR IAM PROCESSING**
Reason: During an IDCAMS LISTC command, a GETMAIN issued by IAM failed. The accompanying IAM260 message names the data set.
Action: If the IAM report is required, increase the region specification for IDCAMS.
- IAM264 'operation' FAILED FOR IAM FILE, DD=ddname**
Reason: During an IDCAMS LISTC the specified operation, OPEN, RDJFCB, or DEVTYPE issued by IAM failed. The data set name is presented in the accompanying IAM260 message.
Action: If other error messages are not present indicating the problem, contact INNOVATION for assistance.
- IAM265 DATA SET IS EMPTY, HAS NOT BEEN DEFINED OR LOADED**
Reason: During an IDCAMS LISTC IAM encountered an empty data set attempting to read the IAM control information. The data set is probably not an IAM file. The data set name is presented in the accompanying IAM260 message.
Action: None, information message only.
- IAM266 LOAD OF MODULE 'modulename' FAILED**
Reason: During an IDCAMS LISTC command, IAM encountered an error loading the named support module. There should also be an accompanying message in SYSLOG indicating the cause of the problem (i.e., a S106, S306, or S806 error). The data set name is presented in the accompanying IAM260 message.
Action: Find the IBM error message, and take corrective action. Possibilities are insufficient virtual storage to load the module, or IAM is not in the system LINKLIST and no STEPLIB has been provided.
- IAM267 IAM 'command' COMMAND FAILED, COMPLETION CODE IS 'cc'**
Reason: During processing of an IAM file for IDCAMS LISTC command, the indicated IAM command failed with the indicated completion code. The data set name is presented in the accompanying IAM260 message. No information about the file will be listed.
Action: Contact INNOVATION for assistance.
- IAM268 DYNAMIC DEALLOC FAILED,RETURN CODE= REASON CODE= (DSN= 'dsname')**
Reason: IAM attempted to dynamically deallocate the previously dynamically allocated file. The deallocation failed with the specified codes. The dataset name if not displayed in this message is presented in the accompanying IAM260 message.
Action: Look up the Dynamic Allocation error codes in the IBM System Macros and Facilities manual or the ISPF Tutorial and take appropriate corrective action, for assistance contact INNOVATION.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM269 IAM CPL PROCESSING FAILED

Reason: IAM was attempting to process a CATALOG parameter list to determine data set name and volume information, however the CPL did not contain the expected information. There will be no information listed about the IAM file.

Action: Contact INNOVATION for assistance.

IAM303 CARD IMAGE - * control statement image *

Reason: A display of the SYSIN data set input control statements.

Action: None. Information message only.

IAM307 DSN ID. 'nn' DATA SET NAME 'dsn' IAM MODE 'mode' IAM DPC 'xx'

Reason: Summary report heading line indicating the number 'nn' of the monitored IAM data set in respect to 'nn' value shown in message IAM339, the 'dsn' name of the IAM data set monitored, the IAM mode, and the IAM Data Protection Code currently in effect. Individual summary lines in this report show:

REQUESTS PROCESSED	– Total IAM requests processed.
I/O OPERATIONS	– Physical EXCP used for requests.
INACTIVE PERIODS	– Number of times IAM waited on requests.
IAM ACTIVATED	– Monitor starting time stamp.
IAM DEACTIVATED	– Monitor ending time stamp.
TOTAL TRANS TIME	– Cumulative IAM processing time.
TOTAL I/O TIME	– Cumulative I/O processing time.
TOT INACTIVE TIME	– Cumulative time IAM spent waiting for requests.
AVG TRANS TIME	– Average time for IAM to satisfy a request.
AVG I/O TIME	– Average time from start of I/O to user CHECK.
AVG INACTIVE TIME	– Average time IAM was idle waiting for a request.

Action: None. Information message only.

80.03 IAM SYSTEM SYSPRINT MESSAGES

- IAM310 NO USABLE MONITOR OR TRACE DATA SETS FOUND**
Reason: User requested analysis of monitor and/or trace data sets but no valid input data could be found.
Action: Check for IAM312 messages detailing the type(s) of data sets identified by the monitor facility. Correct as required and resubmit the job.
- IAM312 DATA SET QUALIFIED BY DDNAME 'ddname' 'dataset type'**
Reason: IAMONRPT interrogates all data sets qualified by ddnames beginning with 'IAMN'. As each data set is examined, the 'dataset type' is listed, i.e.: contains monitor data, trace data, or is empty and does not contain monitor or trace data. Only those data sets which have attributes corresponding to the operation keywords specified will take part in the processing.
Action: None, unless the 'dataset type' listed does not match what you would expect. If errors occurred or empty data sets were encountered, check the JOB LOG of the job(s) that created the data sets for any abnormalities.
- IAM313 NO DDNAMES WITH PREFIX'IAMN' FOUND**
Reason: User neglected to provide any suitable ddnames for monitor/trace report processing. All input monitor or trace data sets must have a DDNAME prefix of 'IAMN' in order to be processed. Execution is terminated without any additional processing.
Action: Check the spelling of the ddnames used in the execution of IAMONRPT. If spelling errors are present or if the DD statements were left out, insert them and rerun the job.
- IAM316 RECOVERY CAN BE FORCED BY SPECIFYING VALUES FOR THE FOLLOWING -**
Reason: An attempt was made to recover an IAM file that has damaged control records. The unreadable control records requires the user to specify key data normally extracted from the file.
Action: The next line(s) displayed will detail the fields required. Use this information to continue the recovery.
- IAM317 I/O ERROR ON DISK TRACK X'cccchhhh'**
Reason: IAM encountered an I/O error while reading or writing the VTOC. A MINI DUMP will be printed.
Action: IAM will continue processing the VTOC bypassing this track. Contact INNOVATION for further assistance.
- IAM318 * WARNING * DATASET- 'dsn' DEFINED FIXED CONTAINS VARIABLE LEN RECORDS**
Reason: The DEFINE for this cluster shows the average and maximum record lengths to be equal. IAMSIMVS detected one or more records that were not equal to the average record length DEFINEd for this cluster.
Action: IAMSIMVS continues processing as if the file contained fixed length records. Blocking, overflow and other values established for a file however vary depending upon whether record lengths are fixed or variable. To obtain a more accurate estimate of IAM's space savings for this file, include the 'VARIABLE' keyword with the 'SELECT' option.
NOTE: To properly identify this file as containing variable length records the DEFINE for this file should be changed so the average RECORDSIZE value is less than the maximum. During processing, if the DEFINEd average RECORDSIZE value is not changed, IAM will return a record length error for this file.
- IAM319 PREMATURE END OF FILE - AFTER BLOCK nnnnnnnn**
Reason: IAMRECVR detected a premature end of file at block nnnnnnnn. Data blocks may have been lost.
Action: Review the contents of the recovered file. A section of the file being recovered may not have been readable.
- IAM320 MAXIMUM BLOCKS LOST DUE TO END OF FILE - nnnnnnnn**
Reason: An end of file error erases the remainder of the track. The reported number of blocks could have existed on the track but were not yet read when the end of file was encountered.
Action: Review the contents of the recovered file. A section of the file being recovered may not have been readable.
- IAM320 TRACING 'dsn' DSNAME IAM MODE 'mode' DPC 'cc'**
Reason: Summary report heading line indicating the 'dsn' name of the data set being traced, the IAM mode, and the IAM Data Protection Code 'cc' in effect.
Action: None, information message only.
- IAM321 INVALID VARIABLE LENGTH FIELD - BLOCK nnnnnnnn**
Reason: IAMRECVR detected an invalid RDW length field for a given record (i.e.: RDW exceeds maximum LRECL). The block number is printed and the remainder of block is bypassed.
Action: Review the contents of the recovered file. A section of the file being recovered may not have been readable.
- IAM322 DROPPED DUE TO I/O ERROR - BLOCK nnnnnnnn**
Reason: IAMRECVR has dropped block nnnnnnnn. The first 24 bytes of the block at the location of the errors is printed in hexadecimal. Processing continues.
Action: Review the contents of the recovered file. A section of the file being recovered may not have been readable.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM323 SEQUENCE CHECK - BLOCK nnnnnnnn

Reason: An out of sequence record was encountered. The IAM block number is printed along with 24 bytes of the key in hexadecimal. Processing continues.

Action: Review the contents of the recovered file. A section of the file being recovered may not have been readable.

IAM324 UNABLE TO CALCULATE LOST BLOCKS - ASSUMING nn MISSING

Reason: An end of file erases the remainder of the track. Any blocks that were on the track at the time the EOF record was created have been destroyed.

Action: Review the contents of the recovered file. A section of the file being recovered may not have been readable.

IAM325 FILE NOT COMPRESSABLE - DATA LEN PAST KEY > 11 BYTES - DSN=

Reason: IAMSIMVS determined that the records contained within the file were not eligible for IAM record compression. IAM only compresses the data in a record located past the end of the key and only if the length of that data is more than 10 bytes. (i.e., MAXLRECL-(RKP+KEYLEN) > 10).

Action: No action is required. IAMSIMVS will continue the simulation for the file as if it were converted to an IAM file in non-compressed format.

IAM326 IAMCRTSM ERROR - 'error description' - DSN=dsn

Reason: IAMCRTSM was invoked to simulate the conversion of a VSAM file and encountered a logical error. The error description will be in the format:

INVALID BLOCKSIZE/BLOCKING FACTOR
CODE=nnn

The CODE=nnn translates directly into an IAM create abend code (see IAM user abend, error, and completion codes later in this section).

Action: If the message was issued for invalid blocking ensure that the blocksize specified is both larger than the record length and compatible with the device. If the CODE= message is issued, refer to the actions for the associated IAM create abend code.

IAM330 DATA SET MONITORED: 'dsname'

Reason: MONITOR Report header for;

A. The User Report showing for each request to IAM the following:

COMMAND	- The IAM command issued	TIME ON QUEUE	- The time the request was on a queue
COMP COD	- The IAM request completion code	I/O TIME USED	- The time to complete the physical I/O if any
REQUEST ISSUED	- The TOD IAM received the request	I/O IAM SNS	- The IAM sense returned if any
REQUEST COMPLETED	- The TOD IAM completed the request	REQUEST NUMBER	- A transaction ID number
REQUEST TIME	- The time to complete the logical request	REQ ACT	- Number of concurrent requests to a file

B. The Graphic Report depicting the following distribution curves:

1. The percentages of logical requests and I/O operations completing within a specified time span.
2. A summary of the percentage of all logical requests and I/O operations completing within a specified time span.

Action: None, information message only.

IAM331 TRANSACTION SAMPLE FROM ddyyy-hhmm TO hhmm HRS ..

Reason: Header for a detail transaction analysis report showing:

TRANS NUM	- An internal IAM trace point ID		
RESULTED FROM	- The name for the trace point	(003)I/O CMLPTD	(007)NO I/O REQ
	(000)INITIALIZE	(004)RETURN	(255)TERMINATE
	(001)MAINLINE		
	(002)I/O START		
DATASET UTILIZED	- Unique ID for this file	DECB CMD/CMP	- The I/O CMD completed in the DECB
IAM CMD	- The IAM command this request was for	TIME STAMPED	- A time stamp from the trace point
IAM COMP	- The IAM completion code for this request	TRANS TIME	- Time to complete the logical request
DECB ADDR	- The virtual address of the I/O DECB	I/O TIME	- Time from start of I/O to user check
		CNCR TRANS	- Concurrent requests queued to this file

Action: None, information message only.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM332 DATASET TRACED 'dsn'

Reason: Report header for the Monitor Trace Report showing for each request made to IAM the following:

COMMAND	- The IAM command issued
COMP COD	- The IAM request completion code
REQUEST COMPLETED	- The TOD IAM completed the request
I/O TIME USED	- The time to complete the logical I/O
IAM SNS	- The IAM sense returned if any
REQUEST NUMBER	- A transaction ID number
KEY IN USE	- The key the command was issued against.

Action: None, information message only.

IAM338 I/O ERROR PROCESSING MONITOR DATA FOR DSN 'dsname' - EOF FORCED

Reason: The IAM monitor program, IAMONRPT, encountered an I/O error attempting to read the data output from a monitored use of the IAM data set 'dsname' name. Rather than abort the interpretation of the data, an END OF FILE was forced and processing continues.

Action: None, information message only.

IAM339 SUMMARY REPORT FOR nn MONITORED DATA SETS

Reason: Summary report heading line indicating the number of monitored IAM data sets reported on during this execution.

Action: None, information message only.

IAM340 'recovery operation' - DSN= 'dsname'

Reason: Describes the user specified IAM file 'recovery operation' and the name of the IAM file being processed.

Action: None, information message only.

IAM341 CONTROL BREAKS IN OVERFLOW BLOCKS - = nnnnnnnn

Reason: Issued by the DIAGNOSE and RECOVER operations of the IAM file recovery program IAMRECVR, indicating that records within the overflow blocks of the file being processed are out of sequence.

Action: When issued during RECOVER, it indicates that the RECOVER output is out of sequence and must be sorted, either by IAMRECVR or externally prior to an IAM create. When issued during DIAGNOSE, no additional processing is necessary.

IAM342 IAM FILE CONTAINS NO DETECTABLE ERRORS

Reason: Issued by the DIAGNOSE and RECOVER operations of the IAM file recovery program IAMRECVR, indicating that the IAM file is not damaged and is acceptable for IAM processing.

Action: None, information message only.

IAM343 SEQUENCE ERRORS IN DATA BLOCKS - = nnnnnnnn

Reason: Issued by the DIAGNOSE and RECOVER operations of the IAM file recovery program IAMRECVR, indicating physical damage to the IAM file.

Action: The file must be created again from a suitable backup or recovered and reconstructed. Prior to create, sort the data set output from the RECOVER operation.

IAM344 NUMBER OF DROPPED BLOCKS - = nnnnnnnn

Reason: Issued by the DIAGNOSE and RECOVER operations of the IAM file recovery program IAMRECVR, indicating physical damage to the IAM file. IAMRECVR encountered I/O errors during the processing of the IAM file and nnnnnnn blocks were dropped from the file.

Action: The file must be created again from a suitable backup or recovered and reconstructed. Prior to create, sort the data set output from the RECOVER operation.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM345 NUMBER OF DUPLICATE RECORDS - = nnnnnnnn

Reason: Issued by the DIAGNOSE and RECOVER operations of the IAM file recovery program IAMRECVR, indicating that duplicate keys were found in the IAM file.

During a RECOVER operation, duplicate records can be ignored, printed, logged for subsequent application to the file, or applied directly if an IAM file is being created as the output of the recovery program.

Action: See the documentation for the RECOVER program, DUPLICATES operand. The file must be created again from a suitable backup or recovered and reconstructed. Prior to create, sort the data set output from the RECOVER operation.

IAM350 'ddname' OPEN FAILED - mmmmmmmm PROCESSING TERMINATED

Reason: A file required by the processing module mmmmmmmm failed to OPEN. This message will always be preceded by MSG IAM465 detailing the reason for the OPEN failure. A user ABEND is forced.

Action: Correct the problem(s) with the library addressed by the 'ddname' statement and rerun the job. If you are not able to identify the problem, call INNOVATION for further assistance AFTER obtaining a storage dump.

IAM351 IAM PROCESSING OF USER DCA FAILED - mmmmmmmm PROCESSING TERMINATED

Reason: IAM was unable to process a user specified IAM DCA (Data Control Area) passed to module mmmmmmmm for dynamic service. The DCA was in error in that it was:

1. Not present
2. Not fullword aligned
3. Contained invalid data
4. Reserved bytes were not set to hexadecimal zero.

Processing is terminated and a user ABEND is forced.

Action: Correct the problem(s) in the program that called module mmmmmmmm for dynamic service and rerun the job. If you are not able to identify the problem, call INNOVATION for further assistance AFTER obtaining a storage dump.

IAM352 DEVICE ALLOCATED TO 'ddname' IS NOT DISK OR TAPE DEVICE - mmmmmmmm PROCESSING

Reason: A file required by the processing module mmmmmmmm can only reside on a disk or tape device. A user ABEND is forced.

Action: Allocate the 'ddname' statement to either a disk or tape device and rerun the job.

IAM353 DEVICE ALLOCATED TO 'ddname' WILL NOT ACCEPT LRECL OF IAM DATASET

Reason: The device allocated to the 'ddname' statement does not have a capacity record as large as the LRECL of an IAM file. This is usually caused by allocating the 'ddname' statement to a disk device with a smaller track capacity than the disk device used for the IAM file. A user ABEND is forced.

Action: Allocate the 'ddname' statement to either the same disk device type as the IAM file or to a tape device and rerun the job.

IAM360 STEP - ssssssss DDNAME - 'ddname' DATA SET MONITORED - 'dsname'

Reason: Identifies the IAM job step name, the DDNAME and, optionally, the data set name of the IAM file being processed by the information service routines of the IAM monitor facility.

Action: None, information message only.

IAM361 INFO REQUESTED BY PROGRAM 'program' PERFORMING 'description' PROCESSING

Reason: Identifies the IAM processing program that requested the printing of the information block, the type of processing being performed, and the time the information block was printed by the information service routines of the IAM INFO report.

Action: None, information message only.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM362 DATA CHARACTERISTICS

Reason: Heading line which always precedes the IAM data characteristics when listed by the IAMINFO report.

- LOGICAL RECORD LENGTH – The maximum length of the record. IAM adds an internal 4 bytes to variable length and compressed records.
- The simulated CI SIZE that was specified.
- KEY SIZE – The length of the KEY (KSDS).
- KEY OFFSET – The position of the key in the record relative to zero (0). IAM adds an internal 4 bytes to variable length and compressed records.
- TOTAL RECORDS – The total number of records in the IAM file.
- TOTAL RECORDS DELETED – The number of records deleted while the IAM file was OPEN.
- TOTAL RECORDS UPDATED – The number of records modified while the IAM file was OPEN.
- TOTAL RECORDS INSERTED – The number of records inserted while the IAM file was OPEN.
- AVERAGE RECORD LENGTH – During create (OPEN OUTPUT) IAM calculates an average length for the records that were loaded.
- LARGEST RECORD PROCESSED – During create (OPEN OUTPUT) IAM reports the length of the largest record that was written.

NOTE: The last two fields only appear for a file being loaded.

IAM363 IAM FILE CHARACTERISTICS

Reason: Heading line which always precedes the IAM file characteristics when listed by the IAMINFO report.

- BLOCKING FACTOR – This is the original Blocking Factor requested. A value of 1–9 indicates number of blocks per track. A larger value is a specified blocksize.
- BLOCK SIZE – Physical block size of the IAM file.
- TRACKS IN USE – Number of tracks used.
- DELETE BYTE SUPPORT – Optional ISAM delete byte support.
- DATASET TYPE – Type of data set: KSDS or ESDS.
- SHARE OPTIONS – SHAREOPTIONS as defined or defaulted to.
- NUMBER OF IAM DATA BLOCKS – Total number of blocks in the IAM dataset.
- HIGH ALLOCATED RBN – Highest allocated Relative Block Number (Enhanced KSDS and ESDS files only).
- INTEGRATED OVERFLOW (CI%) – The percentage of each IAM data block left for new or modified records.
- DASD RESERVE (CA%) – The CA FREESPACE value which IAM used to reserve space for future file expansion (Enhanced KSDS and ESDS) or used by IAM to help calculate the number of Independent Overflow records to allocate (Compatible KSDS).
- FILE DEFINED DATE/TIME – The date and time that the file was defined, as taken from the TOD stamp in the file header.
- FILE LOADED DATE/TIME – The date and time of the completion of the last file load or reorganization for this file.
- FILE UPDATE DATE/TIME – The date and time the file was last closed by a program that updated the file.

IAM364 IAM OVERFLOW CHARACTERISTICS (Compatible KSDS files only)

Reason: Heading line which always precedes the IAM overflow characteristics when listed by the IAMINFO report.

- INDEPENDENT RECORDS REQ – Number of Independent Overflow records requested or calculated from the DEFINE.
- REQUESTED INTEGRATED PER INDEPENDENT RECORDS ALLOC – Percentage of Integrated Overflow (CI%) requested.
- ACTUAL INTEGRATED PER INDEPENDENT RECORDS USED – Actual number of Independent Overflow records allocated by IAM.
- PE BLOCKS REQUESTED – Actual percentage of Integrated Overflow (CI%) allocated by IAM.
- Number of records currently in use in Independent Overflow.
- INDEPENDENT RECORDS EMPTY PE BLOCKS – ALLOCATED,USED – Number of Prime Extension (PE) blocks requested or calculated from the DEFINE.
- Number of records currently available in Independent Overflow.
- Actual number of Prime Extension (PE) blocks allocated by IAM and the number of PE blocks in use.
- INDEPENDENT BLOCKS ALLOC – Number of blocks used for Independent Overflow.
- INDEPENDENT PCT-OPEN,FULL – The percentage of Overflow keys to be reserved at OPEN for inserts, and the keys to be acquired when the incore tables are filled.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM365 IAM EXECUTION STATISTICS

Reason: Heading line which always precedes the IAM execution statistics when listed by the IAMINFO report.

TOTAL STORAGE REQUIRED	– Total storage acquired by IAM during this execution.
COMPRESSED KEY STRUCTURE	– Indicates whether index compressed was used for this file. (Compatible KSDS file only).
PRIME INDEX	– Indicates the amount of virtual storage required for the index to the prime data area of this file. Also if the index is in a compressed format, the word COMPRESSED appears. (Enhanced KSDS and ESDS files only).
STORAGE ABOVE THE LINE	– Total amount of storage acquired above the 16MB line.
COMPRESSED DATA STRUCTURE	– Indicates whether data compression was used for this file.
REQUESTS PROCESSED	– Number of requests made to the IAM file since last OPEN.
REQUESTS FAILED	– Number of requests made to IAM which did not complete normally (i.e. EOF on sequential READ, RECORD NOT FOUND, etc.).
DISK BLOCKS READ	– Number of physical I/Os used to read blocks from the file.
DISK BLOCKS WRITTEN	– Number of physical I/Os used to write blocks in the file.
DYNAMIC BUFFER RETRIEVALS	– Number of EXCPs saved (not executed). IAM's Real Time Tuning Buffer management was able to satisfy this number of I/O requests from buffers that it maintains in storage.
MAXIMUM BUFFERS USED	– The maximum number of data buffers used by IAM's Real Time Tuning during this execution. Under MVS/XA and MVS/ESA all buffers are acquired above the 16MB line.
MINIMUM BUFFERS USED	– The initial number of data buffers acquired by IAM's Real Time Tuning during this execution. Under MVS/XA and MVS/ESA all buffers are acquired above the 16MB line.
MAXIMUM BUFFERS AVAILABLE	– The maximum number of data buffers that could have been acquired by IAM's Real Time Tuning if necessary.
DYNAMIC TABLE RETRIEVALS	– Number of random read requests satisfied by IAM's Dynamic Tabling of data records in virtual storage. Under MVS/XA and MVS/ESA IAM's Dynamic Table (DYNCORE) is maintained in virtual storage above the 16MB line.
DYNAMIC TABLE RECORDS	– The total number of data records IAM placed into the Dynamic Table.
DATA SPACE USED (M BYTES)	– The amount of area actually used in the data space during file load.
DATA SPACE SIZE (M BYTES)	– The size of the data space created during file load.

IAM366 IAM COMMAND EXECUTION SUMMARY (Compatible format KSDS files only)

Reason: Heading line which always precedes the IAM command execution summary when listed by the IAMINFO report.

R.(READ)	– Number of random READ requests with an exact key specified.
W.(WRITE)	– Number of update requests processed.
F.(GETNEXT)	– Number of IAM GETNEXT commands processed.
A.(ADD)	– Number of insert requests processed.
D.(DELETE)	– Number of delete requests processed.
N.(READ OR GETNEXT)	– Number of IAM Read or GETNEXT commands processed.
G.(GET)	– Number of sequential get commands processed.
S.(START KEY EQUAL)	– Number of IAM Start Key Equal commands processed.
K.(START KEY CLASS)	– Number of point requests processed.
P.(GET PREVIOUS)	– Number of Get Previous requests processed.
U.(READ FOR UPDATE)	– Number of Read for Update commands processed.
RESERVED	– Reserved
I.(INFORMATION)	– Number of IAM Information commands issued.
E.(STATISTICS)	– Number of IAM statistic commands issued.
C.(CLOSE)	– Number of CLOSE requests processed.
O.(OPEN)	– Number of IAM OPEN requests processed.
X.(FLUSH)	– Number of requests with RPL errors.
VARIABLE LENGTH CHANGES	– Number of records with their record length changed.
DELETED RECORDS BYPASSED	– Number of records bypassed because they contained X'FF' in first byte.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM366 IAM COMMAND EXECUTION SUMMARY (ESDS/Enhanced KSDS files only)

Reason: Heading line which always precedes the IAM command execution summary when listed by the IAMINFO report.

GET RANDOM	– Number of random READ requests with an exact key specified.
PUT UPDATE	– Number of UPDATE requests.
GET SEQUENTIAL	– Number of sequential GET requests.
PUT ADD	– Number of INSERT requests.
GET PREVIOUS	– Number of GET PREVIOUS requests.
POINT (START BROWSE)	– Number of POINT requests.
GET CI (SEQUENTIAL)	– Number of sequential control interval GETS (ESDS only).
GET KGE/GENERIC	– Number of random or skip sequential reads, indicating search for key greater or equal, and/or partial key search (KSDS only).
PUT CI (UPDATE)	– Number of CONTROL INTERVAL UPDATES (ESDS only).
POINT KGE/GENERIC	– Number of POINT or START BROWSE requests indicating a search for key greater or equal, or partial key search (KSDS only).
GET CI (RANDOM)	– Number of CONTROL INTERVAL GETS (ESDS only).
GET (SKIP SEQUENTIAL)	– Number of GET's issued in skip sequential mode, with an exact key specified (KSDS only)
PUT CI (ADD)	– Number of CONTROL INTERVAL ADDs (ESDS only).
ERASE	– Number of DELETE record requests (KSDS only).
ENDREQ	– Number of LSR ENDREQ requests.
WRTBFR	– Number of LSR WRTBFR requests.
IAM STATISTICS	– Number of IAM statistic requests.
IAM FLUSH BUFFER	– Number of IAM flush buffer requests.
CLOSE	– Number of CLOSE requests.
OPEN	– Number of OPEN requests.
CLOSE TYPE=T	– Number of TEMPORARY CLOSE requests.
VERIFY	– Number of VERIFY requests.
INVALID REQUESTS	– Number of requests with RPL errors.
RECORD LENGTH CHANGES	– Number of records whose lengths changed.
SEQ CHAINED BLOCKS READ	– Number of chained buffers during READ.
SEQ CHAINED BLOCKS WRITTEN	– Number of chained buffers during WRITE.

IAM367 THERE WAS INSUFFICIENT STORAGE AVAILABLE TO ACQUIRE ADDITIONAL BUFFERS

Reason: IAM's Real Time Tuning tried to acquire additional data buffers but there was insufficient virtual storage available.

Action: Increase the REGION value on the JOB statement or the EXEC statement so the next execution of the job will have storage available to acquire additional buffers.

IAM368 SPECIFYING A BUFNO VALUE GREATER THAN nn MAY IMPROVE PERFORMANCE

Reason: IAM's Real Time Tuning buffer management found that, for this mix of data and file processing commands, if additional buffers had been available they would have been acquired. Additional buffers were not acquired because it would have exceeded the maximum buffers allowed for this job.

Action: If you wish to increase the maximum buffers for this job, specify the MAXBUFNO=nn parameter on the IAM Override Control statement ([Section 11](#)) for this IAM file.

NOTE: If the number of I/O's (EXCPs) for this file is relatively small, there is no real need to increase the maximum number of buffers.

IAM369 ISAM SIMULATOR EXECUTED SUMMARY

Reason: Heading line which always precedes the ISAM simulation command summary when listed by the IAM Monitor facility.

Action: None, information only.

IAM370 JOB CHARACTERISTICS

Reason: Heading line which describes the job printed by the IAMINFO command of IAMSMF.

Action: None, information only.

IAM371 INCREASING THE NUMBER OF PRIME EXTENSION BLOCKS MAY IMPROVE PERFORMANCE

Reason: An attempt to insert a record into prime extension failed because the prime extension was full. The record was inserted into independent overflow instead.

Action: Define the IAM file with a larger prime extension.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM372 IAM EXTENDED AREA CHARACTERISTICS (Enhanced KSDS/ESDS files only)

Reason: Heading line which always precedes the IAM Extended area when listed by the IAMINFO report (ESDS and Enhanced KSDS format files only).

EXT. OVERFLOW RECORDS	– Number of records currently contained in extended overflow blocks.
EXT. OVERFLOW BLOCKS	– Number of extended overflow blocks.
EXTENDED BLOCKS ALLOCATED	– Number of extended overflow blocks that will fit within the current amount of DASD space allocated to the file.
EXTENDED PE BLOCKS	– Number of extended blocks assigned as Prime Extension (PE) blocks.
EXTENDED BLOCKS USED	– Actual number of extended blocks in use. This includes Overflow, PE, and Extended index blocks.
EXTENDED BLOCKS AVAILABLE	– Number of extended blocks that are available for use within the currently allocated DASD space.

IAM400 'processing function' - 'progname' - VER v.r. -

INNOVATION DATA PROCESSING - DATE: 'yyddd' PAGE: 'nn'

Reason: Generalized page header for the named 'processing functions' which are provided by the 'program' named in the printed title.

Action: None, information message only.

IAM401 PARM DATA - * parm-field-data *

Reason: Displays the program control information specified in the "PARM=" field of the EXEC statement. This data will not be displayed if program is invoked under TSO.

Action: None, information only.

IAM402 INVALID CONTINUATION

Reason: User coded a delimiting comma following the last keyword on a control statement and did not provide the next logical record.

Action: Correct and resubmit job.

IAM403 REQUIRED OPERAND(S) NOT SPECIFIED - cc...cc

Reason: The operand cc...cc is required for the execution of the command. It must be specified; no defaults are available.

Action: Correct and resubmit job.

IAM404 WORKING STORAGE AREA SIZE OF nn BYTES EXCEEDED - SUBCOMMAND IGNORED

Reason: The maximum number of working storage bytes available to the subcommands has been exceeded.

Action: The number of bytes available for subcommand working storage is derived from the 'MAXCORE' operand, which has a lower limit of 1000 bytes and an upper limit of 120000 bytes. If the upper limit has been reached, call INNOVATION for technical support. If the upper limit has not been reached, resubmit the job specifying a value for 'MAXCORE' greater than the number displayed.

IAM405 MAXIMUM CONTINUATION COUNT OF nnnn EXCEEDED - COMMAND FLUSHED

Reason: The user control statement used too many continuations.

Action: Reduce the number of continuations to the value nnnn. Resubmit the job.

IAM407 CONTROL STATEMENT ERROR - 'action taken'

Reason: An error was encountered during the processing of user supplied control statements. Always preceded by one or more messages which define and delimit the error(s). The 'action taken' is one of the following:

1. JOB TERMINATED – Processing will stop after the first error has been encountered.
2. SKIPPING FOR COMMAND – Processing will continue for all command statements within the SYSIN data set.
3. RE-ENTER COMMAND OR END – Message when the SYSIN data set is assigned to a TSO terminal. Re-enter command in error or 'END' to complete the processing.

Action: Correct and resubmit job.

IAM408 NO CONTROL STATEMENTS WERE FOUND JOB TERMINATED

Reason: SYSIN data set is empty or contained only comment statements ('*' in column 1).

Action: Correct and resubmit job.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM410 DSNNAME="dsname" -- 'error description' -- PROCESSING BYPASSED

Reason: The 'data set name' printed encountered the 'error' described and subsequent processing was bypassed.

The error description can be one of the following:

- NOT RECOGNIZABLE IAM FILE
- TRKCALC FUNCTION FAILED
- NOT KSDS VSAM
- FILE HAS ALTERNATE INDEX
- NAMEDS DATA NOT RETURNED
- UNABLE TO LOCATE DATA/INDEX
- FILE HAS ALTERNATE INDEX
- KEYLEN GREATER THAN 250 BYTES
- RKP GREATER THAN 4096
- FILE WAS NEVER LOADED
- IS NOT VSAM OR IAM

Action: Check the data set name for correctness. If the wrong data set name was specified, specify the correct data set name and resubmit the job. If the message indicates the data set is not an IAM file it is because the IAM control record was not present. Use program IAMRECVR specifying the IAM file characteristics on the 'RECOVER' subcommand to facilitate data recovery.

IAM413 DSN/DSG INDEX ERROR - 'error description'

Reason: When using the index level option to select data sets or data set groups the user either:

1. Specified too many index levels. (The maximum is 22).
2. Failed to provide significant data. (The leading periods were followed by a comma or blank).

Action: Correct and resubmit.

IAM418 COMMAND PROCESSING DETECTED ERROR -- 'action taken'

Reason: An error was encountered during the processing of the subcommand specified on a user supplied control statement. Always preceded by one or more messages that define and delimit the error(s). The action taken is one of the following:

1. RE-ENTER COMMAND OR END – message for user when the SYSIN data set is assigned to a TSO terminal. Re-enter command in error or 'END' to complete the processing.
2. SKIPPING FOR COMMAND – processing will continue for all command statements found within the SYSIN data set.
3. REVERTING TO SYSIN – the error occurred when reading from an alternate command input source. Processing will continue for commands in the SYSIN data set.

Action: Correct and resubmit job.

IAM421 LOCATE ERROR -- 'error description' - DSN= 'dsname'

Reason: A LOCATE SVC was issued requesting identification of the component named by DSN= 'dsname'. The LOCATE either failed or returned a component:cluster name type code that is not currently supported.

1. If the 'error description' is in the form: RETURN CODE rc - REASON IGGOCLAA - 'reason number' The LOCATE failed. Error codes are documented in the IBM System Messages Manual; under message number IDC3009I.
2. If the 'error description' is in the form: ENTYPE -- C(X'xx') The program encountered a type of component entry that it does not presently support.
3. If the 'error description' is in the form: NAMEDS -- C(X'xx') The program encountered a type of cluster entry that it does not presently support.
4. If the 'error description' is in any other form, an error exists in the catalog.

NOTE: This message only appears if DFEFERRPT=YES had been setup earlier with IAMZAPOP or specified on the control statement. This is only a warning message. The program does not associate the component shown by DSN= 'dsname' with a cluster name, processing continues.

Action: Do one of the following:

1. For the RETURN CODE type of error, look up the codes under message IDC3009I. The most likely cause is a STEPCAT or JOBCAT DD statement was not supplied for a user catalog that has entries in it for components on the volume(s) being processed, and those entries are not ALIAS'ed.
2. For any other type of errors, contact INNOVATION technical support for assistance.

IAM457 ENQ FAILED -- RC= 'nn' -- QNAME=c...c - RNAME=c...c

Reason: An attempt was made to reserve or enqueue on the resource listed. Control was not returned. An unexpected failure occurred.

Action: Usually followed by a U0659 ABEND. Call INNOVATION for additional assistance AFTER obtaining a storage dump.

80.03 IAM SYSTEM SYSPRINT MESSAGES

- IAM458 DEQ FAILED -- RC= 'nn' -- QNAME=c...c - RNAME=c...c**
Reason: An attempt was made to dequeue the controlled resource listed. Control would have been released but an unexpected failure occurred.
Action: Usually followed by a U0659 ABEND. Call INNOVATION for additional assistance AFTER obtaining a storage dump.
- IAM465 UNABLE TO OPEN (DDNAME=DSN=) 'name' - 'reason'**
Reason: The 'ddname' named in the message failed to OPEN for the 'reason' shown in the text. The recognized reasons are:
-- ABEND exit taken.
-- VOL= 'volser'
-- denied by OPEN exit.
-- DD= 'ddname' statement missing/ misspelled or incorrectly specified.
-- RC=nn ERROR CODE=nnn
Action:
If the named 'ddname' is required for the requested operation correct the error and resubmit the job, otherwise the message may be ignored.
- IAM471 DDNAME= 'ddname' I/O ERROR -- SYNAD= 'message'**
Reason: A permanent I/O error was detected on the data set referenced by the DDNAME 'ddname'.
Action: Examine the SYNADAF message to determine the cause of the error. The format of this message is documented in IBM SRL publications. Call INNOVATION for additional assistance.
- IAM476 DDNAME= 'ddname' REFERENCES A DEVICE TYPE OTHER THAN DISK**
Reason: The DDNAME 'ddname' listed can only be allocated to a disk device.
Action: Check the unit specification in the JCL for errors, correct and resubmit.
- IAM484 INTERNAL LOGIC ERROR -- JOB TERMINATED**
Reason: The program has encountered an illogical condition.
Action: Call INNOVATION for additional assistance AFTER obtaining a storage dump.
- IAM485 SORT FAILURE HAS OCCURRED RC= 'nn' - 'action description'**
Reason: Your installation's sort product has set a return code other than zero. The return code problem description can be found in the sort manual supplied with the sort program product. In many cases the sort program will provide an error message on DDNAME SYSOUT and or on the system console.
Action: If you cannot correct the problem from this information, call INNOVATION for additional assistance.
- IAM486 VTOC READ ERROR - VOLSER=vvvvvv - SEEK ADDRESS 'cccchhhh' - 'action description'**
Reason: An error was encountered while reading the VTOC of a disk volume. Depending on user options, processing may continue bypassing the remainder of the VTOC or terminate with a storage dump.
Action: Obtain a core dump and call INNOVATION for assistance.
- IAM487 VTOC WRITE ERROR - VOLSER=vvvvvv - SEEK ADDRESS 'cccchhhh' - 'action description'**
Reason: An error occurred while rewriting the VTOC of a disk volume. Depending on user options, processing may continue bypassing the remainder of the VTOC or terminate with a storage dump.
Action: Obtain a core dump and call INNOVATION for assistance.
- IAM488 VTOC OUT OF SEQUENCE - VOLSER=vvvvvv SEEK ADDRESS 'cccchhhh' - 'action description'**
Reason: IAM VTOC read routines check the cchhr portion of the count field for continuity. During this check, one of the records was found to be in error. This message is usually followed by a diagnostic MINI DUMP. Depending on user options, processing may continue, bypassing the remainder of the VTOC, or may terminate with a storage dump.
Action: Check the MINI DUMP for the record in error. Check the volume for hardware problems. If all else fails, call INNOVATION for additional assistance.
- IAM491 'function' FUNCTION STARTED TIME=hh.mm.ss**
Reason: Identifies type of function and time the function started.
Action: None, information message only.
- IAM492 'function' FUNCTION ENDED -- hh.mm.ss CONDITION CODE - nnnn**
Reason: Identifies type of function and time the function ended. The return code is printed if it is non-zero.
Action: Review the completion code. If the completion code is greater than zero, check the output for the preceding messages that describe the reason for a non-zero completion.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM493 'function' - BYPASSED -- PRIOR FUNCTION TERMINATED WITH KEYWORD OR COMMAND DETECTED ERROR

Reason: A previous command upon which this function may be dependent encountered serious errors.

Action: Correct the condition and resubmit.

IAM495 NO RECORDS MATCHED SELECTION CRITERIA

Reason: The selection criteria specified did not cause any records to be selected for processing. If program IAMSMFVS, this message will also appear if the required SMF record types are not being collected.

Action: Make sure the selection criteria is correct.

If IAMSMFVS: Make sure you are collecting SMF records types 4 or 30 subtype 4 and type 64. If you are running against an SMF history tape, make sure that the required record types are being copied to the history tapes.

If IAMSMF: Make sure you are collecting SMF records types 4 or 30 subtype 4 and type 14,15, 64. If you are running against an SMF history tape, make sure that the required record types are being copied to the history tapes.

IAM496 MODULE 'modname' -- NOT USABLE WITH RELEASE v.r. PROGRAMS -- EXECUTION TERMINATED

Reason: Module 'modname' is a release/version that is incompatible with the load module being executed.

Action: Check for a STEPLIB/JOBLIB DD statement pointing to a library other than the correct library for the product level you expect to use. If present, correct the library name and resubmit the job. If there is no STEPLIB/JOBLIB present, the module 'modname' is being obtained from a LINKLIST library. Add a STEPLIB or JOBLIB DD statement specifying the correct load module library and resubmit the job.

IAM497 cc..cc ABNORMALLY TERMINATED DUE TO KEYWORD/COMMAND DETECTED ERRORS

Reason: The common parsing routine encountered errors in parsing the user specified control statements. In addition, the user has set the option KWDC=ABEND via program IAMZAPOP. Always preceded by one or more error description messages.

Action: Previous message(s) describe the error(s); see those messages for further details.

IAM498 cc...cc PROCESSING COMPLETED WITH ERRORS

Reason: The named program completed the requested processing but encountered abnormalities in the process.

Action: Check the output for preceding messages that may describe the errors in detail.

IAM499 cc...cc PROCESSING COMPLETED

Reason: The named program has completed processing as requested.

Action: Check the output for messages that describe the results.

IAM516 CAMLST REGISTERS R0=nnnnnnnn R1=nnnnnnnn R15=nnnnnnnn

Reason: A CAMLST request failed. The type of CAMLST function (shown in the immediately preceding message) and the registers make diagnosing the problem relatively simple. The return code(s) from catalog management is(are) documented in the IBM SRL SYSTEM PROGRAMMING LIBRARY: DATA MANAGEMENT (for MVS) or CATALOG ADMINISTRATION GUIDE (for MVS/XA) or DATA MANAGEMENT FOR SYSTEM PROGRAMMERS (for non-MVS).

Action: Check the return code(s) and take corrective action if the error is apparent or call INNOVATION for additional assistance.

IAM530 MODULE 'modname' NOT FOUND - 'ddname' - 'dsname'

Reason: A BLDL was issued for the module 'modname' in the dataset 'dsname' referenced by 'ddname' and the module was not found. The module is required to support a requested processing function.

Action: Verify that the library dataset name 'dsname' specified on the 'ddname' statement was correct. If incorrect, correct it and resubmit the job. If the dataset name is correct, check the IAM product install listing to see that all steps executed successfully. If required, rerun the installation steps missed. If necessary contact INNOVATION for additional assistance.

80.03 IAM SYSTEM SYSPRINT MESSAGES

- IAM531 MODULE 'modname' CONTAINS NO TEXT RECORDS - 'ddname' - 'dsname'**
Reason: When an attempt was made to read the module 'modname' in the dataset 'dsname' referenced by 'ddname' and no text records were found. The module is required to support a requested processing function.
Action: Verify that the library dataset name 'dsname' specified on the 'ddname' statement was the correct one. If incorrect, change and resubmit the job. If the dataset name is correct, check the IAM product install listing to see that all steps executed successfully. If required, rerun the installation steps missed. If necessary contact INNOVATION for additional assistance.
- IAM532 MODULE 'modname' I/O ERROR READING - 'ddname' - 'dsname'**
Reason: When an attempt was made to read the module 'modname' in the dataset 'dsname' referenced by 'ddname' an I/O error was encountered. The module is required to support a requested processing function.
Action: Check the SYNAD error message(s) and MINI DUMP produced to see if the cause of the error is an obvious one. If so, correct and resubmit the job. If necessary contact INNOVATION for additional assistance.
- IAM533 MODULE 'modname' I/O ERROR WRITING - 'ddname' - 'dsname'**
Reason: When an attempt was made to write the module 'modname' in the dataset 'dsname' referenced by 'ddname' an I/O error was encountered.
Action: This is a serious error which may result in the load module library referenced by 'dsname' being unusable. If necessary, restore or recreate the library. Check the SYNAD error message(s) and MINI DUMP produced to see if the cause of the error is an obvious one. If so, correct and resubmit the job. If necessary contact INNOVATION for additional assistance.
- IAM534 'ddname' OPEN FAILURE -- AVAILABLE COMMANDS: HELP,END**
Reason: The required 'ddname' statement could not be OPENed. Commands which access modules in the 'ddname' dataset cannot be executed.
Action: If the 'ddname' dataset is required for the operation being attempted, correct and resubmit the job. Otherwise, none, information only.
- IAM535 MODULE 'modname' READ UNSUCCESSFUL -- 'command' BYPASSED**
Reason: The user requested function 'command' be performed against module 'modname' but the module was not successfully read. This message will be preceded by message IAM530, IAM531, or IAM532 detailing the reason the read failed.
Action: See message IAM530, IAM531, or IAM532, as required.
- IAM536 MODULE IAMOPT CANNOT PERFORM COPY ON THIS VERSION --- SYSLIB= 'DSNAME'**
Reason: The user requested the COPY function on a version of the options table that is not supported. COPY is not supported with versions of the table prior to 6.1.
Action: Make sure you are only copying to and from a 6.1 or higher level of the IAM options table.
- IAM537 MODULE IAMOPT CANNOT PERFORM COPY ON THIS VERSION --- COPYTODD= 'DSNAME'**
Reason: The user requested the COPY function on a version of the options table that is not supported. COPY is not supported with versions of the table prior to 6.1.
Action: Make sure you are only copying to and from a 6.1 or higher level of the IAM options table.
- IAM538 MODULE 'modname' -- WRONG VER/LEVEL 'ddname' - 'dsname'**
Reason: Module 'modname' is a release/version that is incompatible with the load module being executed.
Action: Verify that the 'ddname' DD statement is pointing to the correct library for the product level you expect to use. If incorrect, change and resubmit the job. If the data set name is correct, check the IAM product install listing to see that all steps were executed successfully. If required, rerun the installation steps missed. If necessary, contact INNOVATION for additional assistance.
- IAM539 MODULE 'modname' 'function' - 'ddname' - 'dsname'**
Reason: The module 'modname' from the dataset 'dsname' referenced by 'ddname' has had function 'function' performed as requested by the user.
Action: None, information only.
- IAM540 ccccccc CONTAINS INVALID CHARACTERS -- ZAP REJECTED**
Reason: The value specified for the operand 'ccccccc' contained one or more characters that were not A-Z, 0-9, \$#@.
Action: Remove the invalid character(s) and resubmit the job.

80.03 IAM SYSTEM SYSPRINT MESSAGES

- IAM541 THE FOLLOWING OPTIONS HAVE BEEN CHANGED IN -- 'dsname'**
Reason: The AUDIT function has found changed options in the options table. The changed options will be listed.
Action: N/A
- IAM542 ccccccc CONTAINS INVALID INDEX STRUCTURE -- ZAP REJECTED**
Reason: The value specified for the operand ccccccc contains two (2) or more consecutive periods (..) in violation of operating system dataset naming conventions.
Action: Correct the error(s) and resubmit the job.
- IAM543 ccccccc DOES NOT CONTAIN CHARACTER STRING ccccccc -- ZAP REJECTED**
Reason: The value specified for the operand ccccccc does not contain the character string ccccccc as the required start of an index level.
Action: Correct the error(s) and resubmit the job.
- IAM544 ccccccc REJECTED -- EXCEEDS MODIFIABLE PORTION OF IAMOPT**
Reason: The length of the value specified for the operand ccccccc taken in conjunction with the offset operand exceeds the modifiable portion of the IAMOPT table.
Action: Correct the error(s) and resubmit the job.
- IAM545 VERIFY FAILED -- CHAR/HEX PRINT FORCED**
Reason: The VERIFY of existing contents failed. A character/hexadecimal print of the module IAMOPT is produced. Always preceded by message IAM544.
Action: See message IAM544.
- IAM546 AUDIT REQUEST COMPLETE ALL DEFAULTS SET IN -- 'dsname'**
Reason: The AUDIT function has found no changed options in the options table.
Action: N/A
- IAM547 AUDIT REQUEST COMPLETE FOR -- 'dsname'**
Reason: The AUDIT function has completed. Changes were found and listed.
Action: N/A
- IAM550 ERRORS ENCOUNTERED DURING EXECUTION -- REWRITE CANCELLED**
Reason: One or more modules from the SYSLIB data set were scheduled for rewrite at termination or by the execution of a REWRITE command. However, previous commands failed to complete successfully. The errors encountered have been documented by preceding error messages.
Action: Correct the error conditions documented by the preceding error messages and resubmit the job.
- IAM567 OBSOLETE OPERAND 'keyword' specified - 'action taken'**
Reason: User specified the obsolete operand 'keyword'. Rather than bypass the command, the operand was either converted internally to a currently supported operand that will perform the same or a similar function to the one that was replaced or it will be ignored.
- | | | | |
|-----------|-----------------------------|----------|----------------------------|
| 'AC' | – CONVERTED TO 'DSORG=' | 'HI' | – IGNORED |
| 'CORE=' | – CONVERTED TO 'SORTCORE=' | 'IN=' | – CONVERTED TO 'DATATYPE=' |
| 'DSN=NO' | – CONVERTED TO 'NODSNAMES' | 'MSG' | – CONVERTED TO 'SORTMSG=' |
| 'DSN=ALL' | – CONVERTED TO 'ALLDSNAMES' | 'MAXR' | – CONVERTED TO 'MAXSTACK=' |
| 'DSN=YES' | – IGNORED | 'NOAUTO' | – IGNORED |
| 'FD=' | – CONVERTED TO 'FROMDATE=' | 'TD' | – CONVERTED TO 'TODATE=' |
- Action:** When convenient, change the control statement to the current standard.
- IAM569 WARNING - ccccccc NOT CATALOGED - DSN= 'dsname'**
Reason: The dataset 'dsname' indicated is the new default for the file type indicated by ccccccc. A CAMLST LOCATE was issued to verify that the dataset is cataloged as required for IAM to dynamically allocate. The data set name was not found by LOCATE. Message IAM516 detailing the CAMLST return code(s) is always printed following this message.
Action: See message IAM516. This is only a warning message. The data set name listed has become the new name of the indicated default.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM574 'request' - DDNAME= 'ddname' - 'action'

Reason: Failed I/O 'request' resulting in the name 'action'.

** DD NOT OPEN – EXTRACT REQUEST DENIED
GET FAILED – EOF FORCED
CLOSE FAILED – JOB TERMINATED
OPEN FAILED – COPY REQUEST DENIED
OPEN FAILED – EXTRACT REQUEST DENIED
OPEN FAILED – REPORT REQUEST DENIED
** OPEN FAILED – REQUEST DENIED
OPEN FAILED – TO DDNAME COPY DENIED
RJFCB FAILED – REQUEST DENIED

The 'ddname' shown in the message is required as either INPUT TO or OUTPUT FROM a processing program. An unexpected failing request for an I/O related service against that DDNAME has forced the program to take the remedial action shown in the message text.

Action: Correct the reason for the error and resubmit the job.

IAM600 MAXSTACK VALUE OF nnnnnn EXCEEDED - SMF RCD PROCESSING TERMINATED

Reason: The SMS record processing subcommand being executed utilizes an in storage stack specified by the operand MAXSTACK. This storage stack is full. The specified size is insufficient to allow all of the selected SMF records to be processed. IAMSMF terminates without producing a report.

Action: Re-execute the program and specify the operand MAXSTACK= with a value greater than the value nnnnnn printed in this message.

IAM601 SMF RECORDS -- READ.nnnnnnnn USED.nnnnnnnn DROPPED.nnnnnnnn

Reason: Documents the SMF records processed as follows:

READ – number of SMF records read from the input data set. May not reflect the total number of records in the data set if an IAM600 message was issued.
USED – number of SMF records selected from the input data set by the criteria specified by the user.
DROPPED – number of SMF records dropped from the input data set because of length checking, or other user specified criteria.

Action: None, information only.

IAM602 FOLLOWING SMF RECORD DROPPED - LENGTH CHECK

Reason: User specified the CHECKLENGTH operand, causing all SMF records selected to be checked against a table of minimum lengths. This record failed the length check. The first 32 bytes of the record are printed in hexadecimal.

Action: Use the keyword RECSIZE=nn on the report control statement to specify an appropriate size.

IAM603 VSAM ERROR -- (DDNAME=DSN=) 'name' 'error description' -- R15=xxxxxxx -CODE=xxxxxxx

Reason: An error occurred during the processing of a VSAM data set. The error description will be one of the following:

CLOSE failed – close of VSAM data set. GET failed – get next record.
GENCB failed – generation of a control block. OPEN failed – OPEN of VSAM data set.

Action: Check the values of R15 and CODE against the return codes listed in the 'VSAM REFERENCE FOR MVS/370 DFP', 'VSAM ADMINISTRATION: MACRO INSTRUCTION REFERENCE FOR MVS/XA', or 'OS/VIRTUAL STORAGE ACCESS METHOD PROGRAMMER'S GUIDE', to determine the cause of the error. If possible, correct and re-execute the job. Call INNOVATION for additional assistance.

IAM619 MODULE 'modname' TOO LARGE TO PROCESS -- 'ddname' - 'dsname'

Reason: IAMZAPOP attempted to read the module 'modname' in the data set 'dsname' referenced by the statement 'ddname'. The module was required to support a processing function but, due to the size of the module (or previously read modules), not enough buffer storage was available to complete the read.

Action: Run IAMZAPOP specifying 'ZAP BUFSIZE=nnnn', where nnnn is the size of the buffer in bytes. Specify a value between 122,880 and 3,145,728. After the buffer size has been changed, rerun the job that failed.

IAM620 MAXJOB VALUE OF nnnnn EXCEEDED -- ADDITIONAL JOB NAMES BYPASSED

Reason: The maximum number of unique job names tabled by IAMSMFVS has been exceeded. SMF records that match the specified selection criteria but with job names other than those already tabled will be bypassed and not reported.

Action: Specify a MAXJOB value greater than nnnnn but less than or equal to 32,000 and, if required, rerun the job.

80.03 IAM SYSTEM SYSPRINT MESSAGES

- IAM621 MAXDSN VALUE OF nnnn EXCEEDED -- ADDITIONAL DATA SET NAMES BYPASSED**
Reason: The maximum number of unique data set names tabled by has been exceeded. Records that match the specified selection criteria but with data set names other than those already tabled will be bypassed and not reported or will be shown in a second report if present.
Action: Specify a MAXDSN value greater than nnnnn but less than or equal to 32,000 and, if required, rerun the job.
- IAM622 MAXDSN VALUE OF nnnnn EXCEEDED -- ADDITIONAL VSAM CLUSTER RELATIONSHIPS BYPASSED**
Reason: IAMSFMVS issues a LOCATE for each ICF/VSAM DATA and/or INDEX component and adds the resulting cluster name to the tabled data set names. Those DATA and/or INDEX components that remain unrelated (i.e., have not had the cluster name appended) when the value of MAXDSN is exceeded will be printed as unique data sets.
Action: ICF/VSAM clusters can take up to three (3) entries in the data set name table that is generated from the MAXDSN operand. If you only plan to extract information on ICF/VSAM clusters and you expect to have approximately 1000 clusters listed, specify a MAXDSN no less than 3000. If the error reoccurs, specify a MAXDSN value greater than nnnnn but less than or equal to 32,000 and, if required, rerun the job.
- IAM630 ARITHMETIC OVERFLOW ERROR -- GENERATION TERMINATED**
Reason: IAMKEYUN has experienced an internal processing error.
Action: Call INNOVATION for additional assistance.
- IAM631 GENERATED KEY NOT HIGHER THAN PREVIOUS -- GENERATION TERMINATED**
Reason: IAMKEYUN has experienced an internal processing error.
Action: Call INNOVATION for additional assistance.
- IAM632 'TOTREC' KEYS GENERATED -- DDNAME='ddname'**
Reason: IAMKEYUN completed processing, and based on the users request has generated the number of record keys 'TOTREC' shown in the message to the data set specified by the 'ddname' statement.
Action: None, information message only.
- IAM633 'command' PROCESSING TERMINATED**
Reason: IAMKEYUN has completed processing the 'command' shown in the message and will now scan the input for additional commands. The 'commands' IAMKEYUN expects to process are: WRITE COMMAND, UNLOAD COMMAND, GENERATE COMMAND, CREATE COMMAND, or LIST COMMAND.
Action: None, information message only.
- IAM634 SEQUENCE CHECKS FORCE SORTING OF DATA BEFORE IAM CREATION**
Reason: The file recovery program found the records it was processing were not in ascending sequential order. An IAM file creation expects the IAM file to be loaded in order. The set of keys the recovery program is processing cannot be used as is to load an IAM file. The Independent Overflow records in an IAM file may not be in collating sequence. Recovery of a file with Independent Overflow could result in this message.
Action: The recovered records must be sorted before they can be used to load an IAM file.
- IAM635 LOG AND IAM FILE ARE INCOMPATIBLE**
Reason: The file recovery program found the log file records it was applying were not compatible with the IAM file it was rebuilding.
Action: The log file being applied and the IAM file in question should be checked to see if they actually represent the same data before attempting to continue with a file recovery.
- IAM636 APPLY FAILED -- RECORD KEY FOLLOWS --**
Reason: The file recovery program found the record it was applying was not compatible with the IAM file it was rebuilding.
Action: The log file being applied and the IAM file in question should be checked to see if they actually represent the same data before attempting to continue with a file recovery.
- IAM637 IAM INTERNAL BLOCK --**
Reason: The block shown contains file control information.
Action: None, information message only.
- IAM638 IAM BLOCK NUMBER nnnnnn --**
Reason: The block shown is relative block number 'nnnnnn' in the file.
Action: None, information message only.

80.03 IAM SYSTEM SYSPRINT MESSAGES

IAM700 Innovation Access Method Trace VER nnnnnnnn Date: 99.999 Page: 999

IAM701 DDNAME:ddname DSN:dataset name

IAM702 Type Time Request RPL/RBN PLH OPTCD RC/RINFO RECLEN Key

IAM703 IAM700 – IAM703 are headings for IAM trace output.

IAM704 Trace Ended

IAM704 indicates the end of IAM tracing.

IAM705 There were trace lock failures

If any trace lock failures occurred during processing, this message will be issued.

IAM710 IOS 08.33.38.1900 PUT 00023A10 00040298 2040 01FE 00002D00

This is the I/O Start trace detail line. The contents of each field are as follows:

Type – IOS **PLH** – The address of the active PLH for this I/O.
Request – One of the following, depending on the type of I/O requested: GET, PUT, POINT, ERASE, CHECK or ENDREQ. **OPTCD** – The values of RPLOPTCD1 and RLPOPTCD2 at the time of the request.
Time – Time the trace record was generated. **RC/RINFO** – blank
RPL/RBN – The address of the RPL used to request the I/O. **RECLEN** – The requested record length (if applicable).
Key – The records RBA or key (if applicable).

IAM711 IOE 08.33.39.2200 PUT 00023A10 00040298 0000 01FE 0005A000

This is the I/O End trace detail line. The contents of each field are as follows:

Type – IOE **PLH** – The address of the active PLH for this I/O.
Request – One of the following, depending on the type of I/O requested: GET, PUT, POINT, ERASE, CHECK or ENDREQ. **OPTCD** – N/A
Time – Time the trace record was generated. **RC/RINFO** – The return code from the I/O.
RPL/RBN – The address of the RPL used for this request. **RECLEN** – The record length (if applicable).
Key – The records RBA or key (if applicable).

IAM712 BFR 08.33.39.2300 WRITE 00000026 00040298 1008 0005A1FE

This is the Buffer Manager trace detail line. The contents of each field are as follows:

Type – BFR **PLH** – The address of the active PLH for this I/O.
Request – One of the following, depending on the type of processing required: READ, WRITE, RELEASE, FLUSH. **OPTCD** – PLH option bytes 1/2.
Time – Time the trace record was generated. **RC/RINFO** – 'OV' if record was from/to Independent Overflow.
RPL/RBN – The Relative Block Number (RBN) of the requested data block. **RECLEN** – N/A.
Key – The records RBA or key (if applicable).

IAM713 EXCP 08.33.39.2400 READ 00000024 00040298 0001

This is the EXCP trace detail line. The contents of each field are as follows:

Type – EXCP **Time** – Time the trace record was generated.
Request – One of the following, depending on the type of processing required:
 PRFMT – Preformat additional blocks **RPL/RBN** – Relative Block Number (RBN) of first block processed.
 WREOF – Write EOF **PLH** – The address of the active PLH for this I/O.
 WRADD – Write new block **OPTCD** – Number of blocks to process.
 READ – Read single or multiple blocks **RC/RINFO** – N/A.
 WRITE – Write single or multiple blocks **RECLEN** – N/A.
Key – N/A.

IAM714 XTND 08.34.02.0100 EXTEND 00000028 00000000 0102 0000 20980001

This is the XTND trace detail line. The contents of each field are as follows:

Type – XTND **Time** – Time the trace record was generated.
Request – One of the following, depending on the type of EXTEND processing required:
 EXPAND – Expand into allocated but unused area. **RPL/RBN** – RBN of first new block after EXTEND.
 EXTEND – Obtain new extent. **PLH** – Number of new index blocks after EXTEND.
 EXP/EXT – Expand and obtain new extent. **OPTCD** – Volume number and extent number of new extent.
Key – EXTEND processing internal flags. **RC/RINFO** – Return code from EXTEND processing.
RECLEN – N/A.

IAM999 END OF JOB

Reason: IAMUTIL is terminated with the completion of the processing for the last control statement.

Action: None, information message only.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

The following user abend codes are issued by IAM programs.

U0000 – AMIGOS ERROR ENTERED VIA DUMP

Internal Logic Error. Issued by modules AMICB, AMIIO, AMIOP, AMIPL, and AMIRC. Contact INNOVATION for assistance.

U0101 – IAM I/O ERROR ON IAM FILE

1. The file could not be OPENed.
 - A. The user failed to supply a DD statement for the file.
 - B. The user specified an incorrect DDNAME in the control statement or DD statement.

Diagnosis: See log for an operating system message indicating the name of the missing DD statement (IEC1301 'ddname' DD statement missing).

2. An unrecoverable I/O error occurred on the file during OPEN or during processing.

Diagnosis: When the error occurs during processing, R10+X'04' points to the DDNAME of the file and R7 points to the DCB address of the IAM file.

U0114 – IAMACC PROGRAM CHECK

1. Since initial OPEN another concurrent task has relocated data blocks in the file.
2. A record present in Independent Overflow at OPEN has been deleted from the file by a concurrent task.

Diagnosis: R10+X'04' points to the DDNAME of the file. In the case of 1 above, an update may have resulted in a physical block being overlaid. After the other task has completed, use IAMUTIL specifying LIST STAT,RECCNT to check the file for damage, then resubmit the original job. In the second case simply resubmit the job.

U0115 – IAMACC PROGRAM CHECK

The file being accessed contains overwritten data blocks. This may have resulted when a CREATE or REORG was run concurrently against a file already OPEN and in use by an update task.

Diagnosis: R3 points to the last record key accessed. R10+X'04' points to the data buffer. Reload the file from a valid backup or use IAMRECVR to OFF-LOAD the data records to a sequential file for recovery processing.

U0116 – IAMACC PROTECTION MODE VIOLATION

1. A WRITE, ADD, or DELETE was attempted in READ-ONLY mode.
2. In ISAM interface mode the DCBMACRF did not specify WRITE or PUTX operations.

Diagnosis: R10+X'04' points to the DDNAME of the file.

U0117 – IAMACC PROGRAM CHECK

Internal Logic Error. Contact INNOVATION for assistance.

U0118 – IAMACC BLKSIZE ON DISK IS ZERO

The file being access was only allocated, it was never created.

Diagnosis: R10+X'04' points to the DDNAME of the file. A valid IAM file will have a FORMAT-1 DSCB which specifies DSORG-DA, LRECL=0 and a valid physical blocksize.

U0119 – IAMACC PROTECTION MODE ERROR

1. The protection mode must be U, R, E, or J.
2. The mode parameter specified in the ISAM interface was invalid.

Diagnosis: R10+X'04' points to the DDNAME of the file. R10+X'04' points to the DDNAME of the file. R10+X'14' points to the invalid mode set.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

U0120 – IAMACC COMMAND CODE ERROR

The command code in sequential mode must be I, R, W, G, D, E, or C; in random mode must be I, R, A, W, D, S, K, G, E, C, N, or F.

Diagnosis: R10+X'04' points to the DDNAME of the file. R10+X'15' points to the invalid command.

U0121 – IAMACC OVERFLOW ADD ERROR

The disk file was not consistent with the index in memory. Most likely another partition or machine is concurrently adding records to the file.

Diagnosis: R10+X'04' points to the DDNAME of the file. This may have resulted in a physical block being overlaid. After the other task has completed, use IAMUTIL specifying LIST STAT,RECCNT to check the file for damage, then resubmit the original job.

U0125 – IAMACC DATA LENGTH ERROR (VARIABLE LENGTH RECORD)

The data length field in the IAMACC Data Control Area which specifies the portion of the record to be read or written is greater than the maximum record length, or less than key length plus RKP.

Diagnosis: R3 at entry to ABEND contains the record's length. R14 at entry to ABEND points to the record area.

U0126 – IAMACC RDW LENGTH ERROR (VARIABLE LENGTH RECORD)

During sequential processing, the length of the record (the RDW value) was changed and a rewrite was attempted.

Diagnosis: R14 at entry to ABEND points to the record area.

U0127 – IAMACC DATA LENGTH ERROR

The record length given in the program's DCA which specifies the length of the record to be read or written is greater than the DEFINEd record length, or less than key length plus RKP.

Diagnosis: R10+X'04' points to the DDNAME of the file.

U0128 – IAMACC SEQUENTIAL UPDATE ERROR

In sequential processing mode, the user attempted to DELETE/REWRITE a record containing a key value which was different from the key of the last record read. In sequential mode, only the last record read may be deleted/rewritten.

Diagnosis: R15 points to the record before rewrite. R7 points to the key of the new record.

U0129 – IAMACC SEQUENTIAL COMMAND ERROR

In sequential mode, the first command a user issued was DELETE/REWRITE. In sequential mode, a record must be read before it can be DELETED/REWRITTEN.

Diagnosis: R10+X'04' points to the DDNAME of the file.

U0130 – IAMACC GETMAIN REQUEST FAILED

The user region did not contain sufficient core to OPEN the IAM file. The region size must be increased.

Diagnosis: R0 contains the amount of core requested.

U0184 – IAMACC VSAM INTERFACE ERROR

IAM encountered an unexpected error while processing a request, and issued an ABEND rather than an error code because the user provided an IAMDEBUG DD DUMMY JCL statement.

U0186 – IAMACC ISAM INTERFACE ERROR

User specified ISAM PUT LOCATE access to an IAM file with data compression. This function is not supported by IAM's ISAM interface.

U0187 – IAMACC ISAM INTERFACE ERROR

The ISAM LRECL specified for the file is not equal to the IAM LRECL.

U0188 – IAMACC ISAM INTERFACE ERROR

User attempted a concurrent rewrite to the same block of data from two DCBs (one random, the other sequential).

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

U0189 – IAMACC ISAM INTERFACE ERROR

1. User issued a START or SETL command against a sequential ISAM DCB which did not specify SETL processing the DCBMACRF.
2. User attempted to OPEN two ISAM DCBS (one random, one sequential) to the same IAM file at the same time and the sequential DCB was not set up for SETL processing in the DCBMACRF.

To compensate for this problem, the user should specify 'DUAL' on the PGM statement of IAM's ISAM Interface and rerun the job with 'DUAL' to override the internal DCB characteristics.

U0190 – IAMACC ISAM INTERFACE ERROR

In batch processing, a user issued a check MACRO after an operation that was in error, but the DCB did not specify a SYNAD routine address.

Diagnosis: R7 points to the ISAM DCB.

U0191 – IAMACC ISAM INTERFACE ERROR

User issued a PUTX command to rewrite a record but the DCBMACRF was not set up for PUTX.

U0192 – IAMACC ISAM INTERFACE ERROR

User issued a PUTX rewrite command for a key which was different than the key of the last record read. The PUTX must not change the key field of the record read.

Diagnosis: R15 points to the record being rewritten. R3 points to the last key read.

U0193 – IAMACC ISAM INTERFACE ERROR

1. User issued a PUTX rewrite command before a GET was issued.
2. A PUTX was issued after an unsuccessful SETL operation.

U0194 – IAMACC ISAM INTERFACE ERROR

A SETL or START command was issued and failed. The user failed to provide a SYNAD address routine.

Diagnosis: R7 points to the ISAM DCB.

U0195 – IAMACC ISAM INTERFACE ERROR

User attempted to open an ISAM file in which the DCBMACRF did not specify either READ or GET.

Diagnosis: R7 points to the ISAM DCB.

U0196 – IAMACC ISAM INTERFACE ERROR

Internal Error. Contact INNOVATION.

U0197 – IAMACC ISAM INTERFACE RANDOM

On an ADD, WRITE, or DELETE command, the nominal key was different than the actual key in the record.

Diagnosis: R15 points to the record area. R3 points to the nominal key. R11 points to the DECB. (+X'05' command, +X'08' DCBADDR, +X'10' RCDADDR, +X'14' NOMKADDR)

U0198 – IAMACC ISAM INTERFACE SEQUENTIAL

The user issued a GET after an end of file.

Diagnosis: R10+X'04' points to the DDNAME of the file.

U0199 – IAMACC ISAM INTERFACE SEQUENTIAL

1. EOF was detected and no EOF address was specified in the DCB.
2. MACRF did not specify either move or locate mode.
3. SETL specified incorrect type byte.
4. SETL for data only and the RKP was equal to zero.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

U0201 – IAMACC RECORD LENGTH SPECIFICATION ERROR

1. The value of LRECL is equal to or less than zero (0).
2. The value of LRECL is less than the RKP + KEYLEN.
3. For variable files, the value of LRECL is less than four (4).

Diagnosis: R10 points to CDCA.

U0202 – IAM KEY LENGTH SPECIFICATION ERROR

1. The value of KEYLEN is less than one (1) or greater than 250.
2. The value of KEYLEN is greater than the value of LRECL.

Diagnosis: R10 points to CDCA.

U0203 – IAMCRT BLOCKING FACTOR CODE SPECIFICATION ERROR

The blocking factor code value is equal to or less than zero (0). The blocking factor code, when used with the DASD device allocated to the IAM file, generates a BLKSIZE which is:

- A. For fixed length format files, less than LRECL.
- B. For variable length format files, less than LRECL+4.

Diagnosis: R10 points to CDCA.

U0204 – IAMCRT BLKSIZE SPECIFICATION ERROR

1. The BLKSIZE value is less than minimum (300+KEYLEN).
2. The BLKSIZE value is less than LRECL for fixed files or less than LRECL+4 for a variable length format file.

Diagnosis: R10 points to CDCA.

U0205 – IAMCRT RKP SPECIFICATION ERROR

1. The RKP value is less than zero (0) for a fixed file or less than four (4) for a variable length file.
2. The RKP value is greater than 4095.

Diagnosis: R10 points to CDCA.

U0206 – IAMCRT INDEPENDENT OVERFLOW SPECIFICATION ERROR

The number of Independent Overflow records requested is less than zero (0) or greater than 32,767.

Diagnosis: R10 points to CDCA.

U0207 – IAMCRT INTEGRATED OVERFLOW SPECIFICATION ERROR

The percentage value specified as Integrated Overflow is less than zero (0) or greater than 99 percent.

Diagnosis: R10 points to CDCA.

U0208 – IAMCRT PRIME EXTENSION SPECIFICATION ERROR

The number of blocks of Prime Extension requested is less than zero (0) or greater than 32,767.

Diagnosis: R10 points to CDCA.

U0209 – IAMCRT OVERRIDE CONTROL STATEMENT SPECIFICATION ERROR

An error has been encountered while processing the user supplied control statements.

Diagnosis: Message 'IAMKWDP' is issued as a WTO,(ROUTCDE=11), detailing the nature and location of the error.

U0210 – IAMCRT DELETE BYTE SUPPORT SPECIFICATION ERROR

The OPTCD value specified in the control block or the control statement is not 'I', 'L', 'IL', or 'M', 'X' or 'ND'.

Diagnosis: R10 points to CDCA, R15 points to OPTCD.

U0211 – IAMCRT ISAM INTERFACE ERROR

ISAM defined RKP=0 with a RECFM=F (unblocked) is not supported.

Diagnosis: Change the ISAM file model DSCB to specify RECFM=FB (blocked). Check the user program to ensure this will be logically correct.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

- U0220 – IAMCRT CREATE DATA CONTROL AREA/CALL PARM LIST ERROR**
1. The CDCA is not fullword aligned, does not begin with the constant 'IDPC', contains blanks in the file DDNAME field or the reserved fields are not binary zero filled.
 2. The last entry in the CALL PARM LIST is not marked with a X'80' in the high order byte or there are more than three entries in the list.
 3. The CDCA has been incorrectly modified.
- Diagnosis:** R10 points to CDCA.
- U0221 – IAMCRT COMMAND CODE ERROR**
- The CDCA command field is invalid.
- Diagnosis:** R10 points to CDCA.
- U0223 – IAMCRT VARIABLE LENGTH RDW ERROR**
- The length in the first two (2) bytes of a variable length record (RDW) is greater than maximum LRECL or less than the KEYLEN + RKP.
- Diagnosis:** R10 points to CDCA. R15 points to record.
- U0224 – IAMCRT ISAM INTERFACE KEY SEQUENCE CHECK**
- The key in the last record written was not higher than the preceding record's key and there is no ISAM SYNAD error routine.
- Diagnosis:** R14 points to current key. R15 points to prior key.
- U0225 – IAMCRT GETMAIN REQUEST FAILED**
- The user region did not contain sufficient core to open the IAM file. The region size must be increased.
- U0230 – IAMCRT IAM FILE IN USE**
- An attempt was made to simultaneously create multiple IAM files to the same DDNAME.
- Diagnosis:** R15 points to DDNAME.
- U0231 – IAMCRT IAM FILE INVALID DEVICE TYPE**
- The DD statement for the IAM file was not allocated to a supported DASD device.
- Diagnosis:** R15 points to DDNAME.
- U0232 – IAMCRT IAM FILE OPEN ERROR**
1. The user failed to supply a DD statement for the file.
 2. The user specified an incorrect DDNAME in the control statement or the DD statement.
- Diagnosis:** R15 points to DDNAME.
- U0233 – IAMCRT IAM FILE I/O ERROR**
- An unrecoverable error occurred on the file.
- Diagnosis:** R14 points to DCB. R15 points to DDNAME.
- U0234 – IAMCRT WORK FILE IN USE**
- The program attempted to use the same work file DDNAME for the concurrent creation of multiple IAM files. Unique work files must be specified for each IAM file when multiple files are to be created concurrently.
- Diagnosis:** R15 points to DDNAME.
- U0235 – IAMCRT WORK FILE INVALID DEVICE TYPE**
- The DD statement for the work file was not allocated to a DASD device.
- Diagnosis:** R15 points to DDNAME.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

U0236 – IAMCRT WORK FILE OPEN ERROR

1. The user failed to supply a DD statement for the file.
2. The user specified an incorrect DDNAME on the control statement or the DD statement.

Diagnosis: R15 points to DDNAME.

U0237 – IAMCRT WORK FILE I/O ERROR

An unrecoverable error occurred in the file.

Diagnosis: R14 points to DCB. R15 points to DDNAME.

U0238 – IAMCRT OVERRIDE CONTROL FILE OPEN ERROR

1. The user failed to supply a DD statement for the file.
2. The user specified an incorrect DDNAME in the control statement or the DD statement.

Diagnosis: R15 points to DDNAME.

U0239 – IAMCRT OVERRIDE CONTROL FILE I/O ERROR

An unrecoverable error occurred while reading control statements.

Diagnosis: R14 points to DCB. R15 points to DDNAME.

U0240 – TO U0249 -IAMCRT PROGRAM CHECK

Internal error. Contact INNOVATION.

U0250 – IAM INTERFACE ERROR

The user program issuing the I/O request is not a subtask to IAM.

Diagnosis: The user programs expecting to do I/O against IAM files through IAM's ISAM Interface must be named on the IAM PGM control statement or must be a subtask to that program.

U0251 – IAM INTERFACE ERROR

Internal control error. Contact INNOVATION.

U0252 – IAM INTERFACE ERROR

Internal control error. Contact INNOVATION.

U0253 – IAM INTERFACE ERROR

User specified a CYL and HEAD (CCHH) access to an ISAM file. This function is not supported by IAM's ISAM interface.

U0254 – IAM INTERFACE ERROR

User issued a PUTX or RELEASE command without a preceding READ.

U0259 – IAMOCICS ERROR

The CICS interface found an FCT entry specifying ISAM, but the DCB in the FCT did not agree.

Diagnosis: R5 should point to the DCB. R10 will point to the FCT. The DCBADDR is based on the offset used in the CICS Interface.

U0260 – IAM CONTROL STATEMENT ERROR

1. Missing IAMCNTL DD * statement.
2. I/O error occurred while processing the IAMCNTL DD statement.

U0261 – IAM CONTROL STATEMENT ERROR

IAM read a control statement from the IAMCNTL data set which was blank to column 60.

U0262 – IAM CONTROL STATEMENT ERROR

Incorrect option was specified on the PGM= control statement.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

- U0263 – IAM CONTROL STATEMENT ERROR**
Keyword was specified that did not end with blank or comma.
Diagnosis: The error occurred on the last statement printed on the log.
- U0264 – IAM CONTROL STATEMENT ERROR**
The first control statement did not specify the PGM= KEYWORD.
- U0265 – IAM CONTROL STATEMENT ERROR**
An incorrect option was specified on the DD control statements.
Diagnosis: The error occurred on the last log statement printed.
- U0266 – IAM CONTROL STATEMENT ERROR**
1. The keywords IAM= or ISAM= are required on a DD control statement for batch processing.
2. The keyword IAM= should be used alone only if one of the online options is chosen.
Diagnosis: The error occurred on the last log statement printed.
- U0267 – IAM CONTROL STATEMENT ERROR**
The user exceeded the maximum number of DD statements. The default is 100 DD control statements if one of the online options is specified on the PGM statement, otherwise 10 is the default. Use the MAXDD option to override this default.
- U0268 – IAM CONTROL STATEMENT ERROR**
User did not specify any IAM DD statements in the IAMCNTL data sets.
- U0269 – IAM CONTROL STATEMENT ERROR**
An optional keyword was specified that required a numeric field. That field was non-numeric (EX: MAXDD=50).
Diagnosis: The error occurred on the last log statement printed.
- U0270 – IAM DD STATEMENT ERROR**
1. User specified a DDNAME on an ISAM= control statement which was missing.
2. The ISAM file specified by the ISAM= keyword failed to open.
Diagnosis: The error occurred on the last log statement printed.
- U0271 – IAM CONTROL STATEMENT ERROR**
Name specified on the IAM control statement exceeds 8 bytes.
Diagnosis: R4 points to the control statement.
- U0272 – IAM CONTROL STATEMENT ERROR**
MAXDD was specified with a value of zero.
- U0273 – IAM DD STATEMENT ERROR**
I/O error occurred on ISAM file.
Diagnosis: The error occurred on the last log statement printed.
- U0274 – IAM INTERNAL ERROR**
Internal control error. Contact INNOVATION for assistance.
- U0275 – WAIT BIT ERROR**
Internal control error. Contact INNOVATION for assistance.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

U0276 – IAM INTERFACE ERROR

1. IAM encountered an OPEN for a ISAM file not specified on the IAM DD control statements and "MIX" was not specified on the PGM statement.
2. An ISAM file was accessed by create processing. This file must be converted to IAM format in order for this program to run.

U0277 – IAM ONLINE INTERFACE ERROR

User specified ONLINE or one of the other TP options on the PGM control statement and ...

1. IAM did not find the required IAMISAM DD statement.
2. I/O error occurred on the file specified by the IAMISAM statement.

The DD statement should specify:

```
// IAMISAM DD UNIT=SYSDA,SPACE=(CYL,1),DCB=DSORG=IS
```

U0278 – IAM ONLINE INTERFACE ERROR

User specified ONLINE or one of the other TP options on the PGM control statement. An error occurred in the execution of the IAMOPEN or IAMOCICS modules by IAM. Contact INNOVATION for assistance.

U0279 – IAM ONLINE INTERFACE ERROR

User specified CICS and one of the other TP Options (ex: ONLINE) on the same PGM statement. These options are mutually exclusive.

U0280 – IAM FUNCTIONAL ERROR

Program attempted a sequential OPEN of an ISAM file which was already OPENed sequentially by another DCB.

Diagnosis: R14 points to the DDNAME of the file being OPENed. R15 points to the DCB.

U0281 – IAM MODULE INTERFACE ERROR

IAM is not running as an authorized program.

U0282 – IAM VSAM INTERFACE (VIF) ERROR

Internal control error. Contact INNOVATION for assistance.

U0283 – IAM VSAM INTERFACE (VIF) ERROR

One or more of the following modules were not in the IAM Vector table (IAMVECTB): IAM0CLA0, IAM0192A, IAM0200T, IAM0231T. Contact INNOVATION for assistance.

U0295 – IAM INTERFACE ERROR

Issued by module IAMISAHF. Internal control error. Contact INNOVATION for assistance.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

U0301 – IAMACC DATA CONTROL AREA/CALL PARM LIST ERROR

1. The ADCA did not begin on a fullword boundary.
2. The first four (4) bytes ('IDPx') were incorrectly specified.
3. The ADCA was incorrectly modified during processing.
4. The reserved fields in the ADCA were not binary zero filled.
5. The last entry in the CALL parameter list was not marked with a X'80' in the high order byte; the VL bit in an ASSEMBLER call.
6. The user specified an entry after the record area in the parameter list.

Diagnosis: R10 points to the ADCA.

U0302 – IAMACC FILE ACCESS ERROR

1. An attempt was made to access a non-IAM file with IAM.
2. An attempt was made to access an IAM file which was not properly CLOSED during creation.
3. The IAM file's control information on the disk has been destroyed or improperly modified.

Diagnosis: R10+X'04' points to the DDNAME of the IAM file. A validly created IAM file will have a FORMAT-1 DSCB which specifies DSORG=DA, LRECL=0 and a valid physical blocksize.

U0303 – IAMACC FILE ACCESS ERROR

An attempt was made to access a data set which for some reason is not available.

U0304 – IAMACC FILE ACCESS ERROR

An attempt was made to OPEN the same data set for update multiple times within the same task.

U0350 – IAM COBOL VSAM INTERFACE ERROR

During OPEN, IAM found the COBOL VSAM file information block (FIB) did not contain valid VSAM identification.

Diagnosis: Most likely a user program error caused the FIB to be overlaid.

U0351 – IAM COBOL VSAM INTERFACE ERROR

During OPEN, IAM found the LRECL for a fixed length disk file was not equal to the 01 LEVEL specified in the program.

Diagnosis: The file is inconsistent with the processing program.

U0352 – IAM COBOL VSAM INTERFACE ERROR

During processing, IAM read a variable length record from the file which was larger than the maximum 01 LEVEL specified in the program.

Diagnosis: The file is inconsistent with the processing program.

U0353 – IAM COBOL VSAM INTERFACE ERROR

A user's sequential rewrite/delete or dynamic read failed because a record that was present when the buffer priming command was issued is no longer in the file.

Diagnosis: Most likely the file is being concurrently updated by another partition, another machine or by multiple FD's for the same file.

U0354 – IAM COBOL VSAM INTERFACE ERROR

During OPEN processing, IAM found the KEYLEN and RKP of the IAM file not equal to the description of the file in the program.

Diagnosis: The IAM file is inconsistent with the processing program.

U0360 – IAM COBOL VSAM INTERFACE ERROR

Internal logic error. Call INNOVATION for assistance.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

- U0401 – SYSIN ERROR**
The SYSIN DD statement failed to OPEN for IAMACC, IAMKEYUN, IAMONRPT, IAMRECVR, IAMSMF, IAMSMFIS, IAMUTILS, or IAMZAPOP. It may have been missing or incorrect.
- U0402 – SYSPRINT ERROR**
The SYSPRINT DD statement failed to OPEN for IAMUTIL or IAMMODEL. It may have been missing or incorrect.
- U0501 – SYSIN ERROR**
1. The SYSIN DD statement was missing or incorrect for IAMUTIL.
2. An I/O error occurred when reading the input statements.
- U0502 – STATEMENT ERROR**
1. The IAM processor (IAMKEYUN, IAMONRPT, IAMRECVR, IAMSMF, IAMSMFIS, IAMUTIL, IAMUTILS, or IAMZAPOP) terminated because it found a control statement error. An error message is printed on SYSPRINT.
2. The IAM processor (IAMMODEL) terminated because of missing DD statements. An error message is printed on SYSPRINT.
- U0503 – IAMUTIL EXIT PROCESSING ERROR**
IAMUTIL detected an invalid command from an exit invoked routine.
- U0504 – UTILITY PROCESSING ERROR**
1. IAMUTIL detected a serious error which caused it to terminate the job. An error message is printed on SYSPRINT. See message IAM102, IAM109, IAM110, IAM111.
2. IAMSMF invoked a sort that failed. An error message is printed on SYSPRINT. See message IAM485.
- U0505 – IAM ERROR DURING SIM, ADD, DELETE, OR TEST FUNCTION**
Internal logic error. Call INNOVATION for assistance.
- U0506 – ISAM ERROR DURING SIM, ADD, OR TEST FUNCTION**
Internal logic error. Call INNOVATION for assistance.
- U0507 – IAMUTIL REORG ERROR**
User specified REORG against an IAM file. The TAPE1 DD statement specified a dummy data set.
- U0590 – IAMREORG ERROR**
REORG processing terminated. An error message is printed on SYSPRINT. See message IAMW49.
- U0591 – IAMREORG ERROR**
SYSPRINT failed to OPEN.
- U0592 – IAMREORG ERROR**
REORG processing terminated. An error message is printed on SYSPRINT. See messages IAM350, IAM351, IAM352, or IAM353.
- U0600 – UTILITY PROCESSING ERROR**
1. OPEN failed.
2. I/O error. A message is printed on SYSPRINT. See messages IAM471, IAM532, IAM533, IAM574, or IAM603.
3. CLOSE error. A message is printed on SYSPRINT. See messages IAM532, or IAM533.
- U0658 – UTILITY PROCESSING ERROR**
IAMSMF/IAMSMFVS error. A message is printed on SYSPRINT. See message IAM496.

80.10 IAM USER ABEND, ERROR AND COMPLETION CODES

U0659 - UTILITY PROCESSING ERROR

1. IAMEXTND internal logic error.
2. IAMKEYUN internal logic error.
3. IAMONRPT internal logic error.
4. IAMRECVR internal logic error.
5. IAMSMT internal logic error.
6. IAMSMTFVS internal logic error.
7. IAMUTILS internal logic error.
8. IAMZAPOP internal logic error.

See messages IAM484 or IAM530 on SYSPRINT.

Call INNOVATION for assistance.

U0660 - UTILITY STORAGE STACK OVERFLOW

Internal Logic Error. Call INNOVATION for assistance.

U0777 - IAMSTART PROCESSING ERROR

IAMSTART SETLOCK failed. Call INNOVATION for assistance.

U0787 - MAXIMUM OF 20 ENTRIES EXCEEDED IN PARAMETER LIST. (IAMLODP only)

The parameter list contained more than 20 entries.

U0788 - NO PARAMETER LIST ADDRESS PASSED. (IAMLODP only)

No parameter list passed to IAMLODP.

U0789 - USER ATTEMPTED TO CALL MORE THAN TEN (10) DIFFERENT LOAD MODULES

The maximum number of load modules was exceeded.

U0900 - UTILITY PROCESSING ERROR

1. IAMEXTND correctable user error. A message is printed on SYSPRINT. See message EXTND001.
2. IAMLKD error. A message is printed on SYSPRINT. See messages IAMLK900(SYSPRINT OPEN failed), IAMLK902(MISSING DD), IAMLK903(parameter error), IAMLK940(compress failed).

U0901 - UTILITY PROCESSING ERROR

1. IAMEXTND serious error. A message is printed on SYSPRINT. See messages EXTND003, EXTND005, EXTND008.
2. IAMLKDUT logic error. A message is printed on SYSPRINT. See message IAMLK950.

U0902 - EXTEND PROCESSING ERROR

IAMEXTND serious error. A message is printed on SYSPRINT.

See messages EXTND002, EXTND006, EXTND007, EXTND009.

80.20 IAM'S VSAM INTERFACE (VIF) ERROR CODES ISSUED BY IAM:

80.20 IAM'S VSAM INTERFACE (VIF) ERROR CODES ISSUED BY IAM:

The following are Return/Reason Codes issued by VIF which correspond to VSAM MESSAGE IDC3009I.

<u>RETURN CODE</u>	<u>REASON CODE</u>	<u>DESCRIPTION</u>
8	38	DATA SET ALREADY CATALOGED
16	0	SMS FAILED ALLOCATION
22	8	DEFINE IS NOT FOR AN INDEXED FILE
42		MVS/ESA DASDM ALLOCATION FAILED -
42	nnnn	THE REASON CODE IS THE RETURN CODE FROM DASDM
44	12	CALLER WORK AREA TOO SMALL
48	0	NOT IN AN ICF CATALOG
54		DASDM ALLOCATION FAILED -
54	nnnn	THE REASON CODE IS THE RETURN CODE FROM DASDM
58	0	OBTAIN FAILED
58	4	VOLUME NOT MOUNTED
58	8	DSCB NOT FOUND IN VTOC
58	12	CVAF READ/WRITE ERROR
62	0	INITIALIZATION OF IAM FILE FAILED
68	2	NO SPACE ON VOLUME
86	4	RECATALOG - DATA SET NOT IAM OR VSAM
86	6	RECATALOG - DATA SET NOT ON VOLUME
72	4	VOLUME SERIAL NOT ONLINE
108	0	FIELD NAME NOT FOUND
132	2	NO VOLUME FVT FOUND IN PARAMETER LIST
132	4	NO AMDSB FVT FOUND IN PARAMETER LIST
132	8	NO AVERAGE LRECL FVT FOUND IN PARAMETER LIST
132	10	NO SPACE FVT FOUND IN PARAMETER LIST
132	26	NO SPACE IN FPL
132	48	NO AMDSB IN FPL
136	2	NO VOLUME LENGTH
136	6	CLUSTER NOT FOUND
136	18	NO AVERAGE LRECL IN FPL
168	2	NON-GENERATED DEVICE TYPE
176	0	NO SPACE IN VTOC
192	0	EXCEEDED MAX RECORD SIZE
240	4	DEVTYPE FAILED
240	22	MISSING DDNAME
224	0	EXCEEDED 5 VOLSERS ON DD STATEMENT

80.21 IAM ERROR CODES PASSED IN THE RPL

80.21 IAM ERROR CODES PASSED IN THE RPL

The following are Return/Reason Codes passed by IAM in the RPL.

<u>RETURN</u> <u>CODE</u>	<u>RPL</u> <u>ERRCODE</u>	<u>DESCRIPTION</u>
08		<u>LOGICAL ERROR OCCURRED</u>
08	X'04'	SET LOGICAL END OF FILE
08	X'08'	SET DUPLICATE RECORD CODE
08	X'0C'	SKIP SEQUENTIAL OUT OF SEQUENCE
08	X'0C'	KEY SEQUENCE ERROR ON SEQUENTIAL ADD
08	X'0C'	INDICATE SEQUENCE CHECK DURING LOAD
08	X'10'	SET NO RECORD FOUND
08	X'14'	SET RECORD UNDER EXCLUSIVE CONTROL
08	X'1C'	SET FILE FULL FOR RPL
08	X'28'	CORE NOT AVAILABLE FOR LOCATE
08	X'2C'	WORK AREA NOT LARGE ENOUGH FOR RECORD
08	X'40'	MAXIMUM NUMBER OF PLH USED
08	X'44'	FILE NOT OPENED FOR UPDATE
08	X'48'	KEYED ACCESS ATTEMPTED ON ESDS
08	X'50'	ATTEMPTED ERASE ON ESDS
08	X'54'	PUT LOCATE REQUEST IS INVALID
08	X'58'	SEQUENTIAL REQUEST WITHOUT BEING POSITIONED
08	X'5C'	PUT OR ERASE WITHOUT READ FOR UPDATE
08	X'60'	PUT KEY DOES NOT MATCH READ
08	X'64'	ATTEMPTED TO CHANGE RECORD LENGTH ON ESDS FILE
08	X'68'	RBA PROCESSING IS NOT SUPPORTED
08	X'68'	GET PREVIOUS IN SKIP SEQUENTIAL MODE
08	X'6C'	RECORD LENGTH INVALID
08	X'70'	KEY LENGTH IN RPL EXCEEDS KEY SIZE
08	X'74'	SET INVALID REQUEST DURING LOAD
08	X'F1'	INVALID REQUEST TYPE
08	X'F2'	ECB ADDRESS INVALID
08	X'F4'	DATA SPACE OVERFLOW
0C		<u>PHYSICAL I/O ERROR OCCURRED</u>
0C	X'04'	SET READ I/O ERROR
0C	X'10'	SET WRITE I/O ERROR

80.22 IAM ERROR CODES PASSED IN THE ACB

80.22 IAM ERROR CODES PASSED IN THE ACB

The following are error flags passed by IAM in the ACB if the file is not OPENed successfully.

<u>RETURN CODE</u>	<u>ACB ERFLG</u>	<u>DESCRIPTION</u>
04	X'04'	CLOSE ERROR – ACB NOT OPEN
08	X'80'	OPEN ERROR – DD STATEMENT MISSING
08	X'88'	OPEN ERROR – INSUFFICIENT STORAGE IN REGION
08	X'A0'	OPEN ERROR – FILE ALLOCATED BUT NOT CREATED
08	X'A8'	OPEN ERROR – DATA SET NOT AVAILABLE
08	X'A8'	OPEN ERROR – MORE THAN 2 OPENS FOR UPDATE IN TASK
08	X'B8'	OPEN ERROR – I/O ERROR
08	X'BC'	OPEN ERROR – NOT IAM OR VSAM FILE OR LOAD FAILED
08	X'C0'	OPEN ERROR – UNUSABLE DATA SET OPENED FOR LOAD
08	X'F1'	SET ERROR – BAD VERSION
08	X'F8'	SET ERROR – BAD RETURN CODE
08	X'FE'	SET ERROR – WRONG VERSION

IAMM

INNOVATION ACCESS METHOD

INSTALLATION

THIS PAGE WAS INTENTIONALLY LEFT BLANK

90.01 IAM V6.3 INSTALLATION INSTRUCTIONS

90.01 INNOVATION ACCESS METHOD (IAM) INSTALLATION INTRODUCTION

OVERVIEW IAM is an efficient high performance alternative to using VSAM KSDS clusters to manage your SMP/E CSI data sets and in other purchased or in house written applications that run as batch programs, interactive TSO processors or CICS transactions. Programs coded to access single index KSDS VSAM files can transparently use IAM files without program or JCL changes. IAM has been designed to provide especially high performance and reliability. In addition, IAM files because of their simplified file structure and optional Data Compression feature usually take 30 to 70% less disk space than VSAM equivalents.

The installation of IAM is a simple three (3) step procedure.

First, load the IAM Installation Control Library (Sec. 90.02). Second, load the IAM program load library (Sec. 90.03). Third, execute the procedure to activate IAM's system level VSAM Interface (VIF) in your system (90.04).

Section 90.05 outlines a testing plan for IAM.

You are now ready to use IAM files in place of VSAM clusters. Whenever you specify the \$IAM parameter in a DEFINE statement for a KSDS cluster, you will create an IAM file instead of a VSAM cluster. Any program referencing the file will be automatically switched to IAM instead of using VSAM.

If you are planning to use IAM in a CICS environment, review CICS Special Considerations (Section 90.10) before you start.

PRODUCT TAPE The product distribution tape is recorded at 6250 BPI or 3480 in SL format with a volume serial number of IAM63T(P). The tape volume IAM63T contains a trial version of the library which is date protected and will expire on the date shown on the external tape label. The tape volume IAM63P contains the production version of the product library. There are no date protected modules in the production library.

TRIAL CUSTOMER If you are a trial IAM customer, the trial copy will expire on the date specified on the installation tape, unless it is extended by INNOVATION. You will be supplied a production library when you purchase the product. Reinstall the IAM system from the production tape before your trial expiration date.

CURRENT CUSTOMER If you are a current customer of IAM, you can install and test this new version of IAM without effecting any of your existing jobs. Programs can continue to use prior versions until you choose to have them use the new version.

Install this version of IAM in an authorized library. If the library is in the LINKLIST, make sure it is not in front of the production version. Execute the VIFSTART procedure from the new Installation Control Library. You will now have two versions of IAM installed in your system. Insert a STEPLIB pointing to the new IAM load library in the jobs you wish to test. IAM modules contain a key matching them to the proper VIF interface. The VIF Interface modules in the system will recognize which version of IAM the STEPLIB contains.

NOTE: The newest version of IAM must be started after all prior versions of IAM are active in the system.

90.01 IAM V6.3 INSTALLATION INSTRUCTIONS

PROCEDURE FOR TESTING WITH PREVIOUS VERSIONS OF IAM ACTIVE

If you are a current customer of IAM, you can install and test this new version of IAM without effecting IAM 6.3 any of your existing jobs. Programs can continue to use prior versions until you choose to have them use the new version.

Install this version of IAM in an authorized library.

If the library is in the LINKLIST, make sure it is not in front of the production version. Execute the VIFSTART procedure from the new Installation Control Library. You will now have two versions of IAM installed in your system. Insert a STEPLIB pointing to the new IAM Load Library in the jobs you wish to test. IAM modules contain a key matching them to the proper VIF Interface. The VIF Interface modules in the system will recognize which version of IAM the STEPLIB contains.

NOTE: The newest version of IAM must be started after all prior versions of IAM are active in the system.

PROCEDURE FOR TESTING NEW IAM LEVELS WITH THE SAME VERSION ALREADY ACTIVE

Beginning with IAM Version 6.2, it is possible to test different levels of the same version of IAM concurrently. This ability is required when different levels of the same version utilize the same VIF key. Please follow these instructions which will limit the test level of IAM to interact with only certain jobs for testing purposes.

Dynamically install the test version of IAM by executing a batch job using the following JCL:

```
//IAMSTRTT JOB 'DYNAMICALLY INSTALL TEST VERSION
//          OF IAM',IAM
//IAMSTRTT EXEC PGM=IAMSTART,
//          PARM='TEST, JOBNAME=jobname
//          =jobgroup* See Note 1
//STEPLIB DD DSN=i am . test . library,DISP=SHR See Note 2
//SYSLIB DD DSN=i am . test . library,DISP=SHR See Note 2
//ABNLDUMP DD DUMMY FOR ABEND-AID
//SYSUDUMP DD SYSOUT=A
```

This JCL is supplied in member IAMSTRTT in the Installation Control Library ([loaded as described in Section 90.02](#)).

Note 1: On the EXEC statement, JOBNAME= specifies the job(s) that are to be processed by the test version of the exits. JOBNAME= may be abbreviated as JOB=. The value specified may be either a single jobname, or a group name specified as a jobname prefix followed by an asterisk (*). TSO userids and started task names will also be checked for a match and will be eligible for processing by the test version of IAM. Any job that does not match will be processed by the production version of IAM. If JOBNAME=* or JOB=* is specified, then all jobs will be processed by the test version; in effect, the production version will be temporarily replaced.

Note 2: The SYSLIB DD statement, and, if used, the STEPLIB DD statement must specify the dsname of the test library, created as discussed above in Step A.

90.01 IAM V6.3 INSTALLATION INSTRUCTIONS IAM VERSION 6.3

90.01 IAM V6.3 INSTALLATION INSTRUCTIONS

To aid in the implementation of IAM Version 6.3, the following procedures should be used to avoid system failures and outages.

IAM 6.3 has 2 file formats for the KSDS files, COMPATIBLE and ENHANCED. The GLOBAL OPTIONS TABLE has been set so the default is to create COMPATIBLE format files. This file structure is completely compatible with all prior releases of IAM. Any files created with the ENHANCED file structure will not be able to be processed by prior levels of IAM. For details on the ENHANCED file format please refer to [Section 01](#), Introduction and V6.3 Enhancements.

Interfaces between version 6.1 or 6.2 and 6.3 are not compatible and will require an IPL to safely cut over IAM 6.3 into production.

CONCURRENT TESTING of IAM 6.3 with any prior version of IAM may be done as long as the latest version of IAM is started last. Included in the ICL library are PTF's P610023 for IAM 6.1/05P and below. These PTF's will prevent a version of VIF from starting if a later version is already active. We recommend you apply the appropriate version of this PTF if you are planning concurrent testing of different versions of IAM.

Should you have any questions regarding concurrent testing of IAM Version 6.3 with prior levels of IAM, or questions about placing IAM Version 6.3 into production, please call IAM Technical Support at Innovation Data Processing. (201) 890-7300.

If you are running IAM 6.0 and plan concurrent testing with 6.3, please contact IAM Technical Support.

Please follow these guidelines:

1. DO NOT COPY IAM VERSION 6.3 INTO ANY LIBRARY CONTAINING a prior version or release of IAM unless an IPL is to be done immediately after the copy. Problems with CICS, and abends in batch jobs have occurred when improper procedures were followed.
2. The newest version of IAM must be started after any prior versions if concurrent testing is to be done.
3. DO NOT ACTIVATE THE LOADJCLALLOC FEATURE if running more than one level of VIF on your system.

INSTALLING AND TESTING IAM VERSION 6.3

1. Install IAM Version 6.3 into an authorized library.
2. Set the IAM Version 6.3 global options as desired. The IAM Global Options table has changed so do not copy a prior version of the options table into the 6.3 library. Run the IAMZAPOP PRINT function on both the prior and new versions of IAM to compare the setting of the Global Options. After comparing the options, run IAMZAPOP to change the necessary options in the new 6.3 option table.

Many custom zaps from IAM Version 6.1 and 6.2 are now options in the Global Option Table, and are not necessary for Version 6.3.

3. IAM Version 6.3 VIF can now be started out of the new library. If you are concurrently running prior versions of IAM, always start the VIF for the prior versions before starting the new VIF. To test IAM Version 6.3, after starting VIF place JOBLIB or STEPLIB dd statements to the new library for the jobs or steps which are testing the new version.

90.01 IAM V6.3 INSTALLATION INSTRUCTIONS IAM VERSION 6.3

PLACING IAM VERSION 6.3 IN PRODUCTION

You can either create a new LNKLSTxx member of SYS1.PARMLIB containing the new IAM 6.3 load library and IPL using the new PARMLIB member, or copy the new IAM load library over the old IAM load library just before an IPL is to be done. All open IAM files must be closed before stopping VIF and copying the new IAM load library over the old one or abends may occur when the IAM files are closed.

INSTALLATION CHECKLIST

- 1 – Using the JCL example provided on page 6 of the Installation section of the IAM manual, download the ICL library from the IAM tape.
- 2 – Make sure the IAM load library is allocated in a single extent to avoid problems when the library is placed in the linklist. This will require about 30 tracks of 3380 space and 16 directory blocks. Then install the IAM load library using the JCL provided in the IAMLOAD member of the ICL library.
- 3 – Set the global options table as desired. Run IAMZAPOPOP PRINT to list the options from any prior version of IAM and from the 6.3 version. Example 1 in [section 91.04](#) of the IAM manual provides sample JCL to do this. Compare the options and run IAMZAPOPOP to set the options in the 6.3 table to match the prior version with any desired changes.

If you have changed the DATACOMPRESS option in prior versions it is recommended that you change it in the new version also or users may experience unexpected changes in file size and format.

DO NOT COPY THE PRIOR VERSION OF THE TABLE INTO THE 6.3 LOAD LIBRARY, THE TABLE HAS CHANGED FROM PRIOR VERSIONS.

- 4 – Check your custom zaps from prior versions of IAM. Many of the IAM 6.1 level custom zaps identified by the C-61 prefix on the fixes are now GLOBAL OPTIONS and are no longer necessary. If you had installed any custom zaps and they do not appear to be a GLOBAL OPTION please contact Innovation for a revised zap.
- 5 – Activate the VSAM INTERFACE (VIF) using either the VIFSTART batch job or the IAMSTART procedure provided in the ICL library. Make sure that any STEPLIBS or SYSLIBS in these jobs point to the newly allocated IAM load library.

If concurrently testing different versions of IAM make sure the newest version is started last.
- 6 – You can test this IAM version with prior versions active by adding steplibs to any batch jobs or CICS regions that wish to test under the new version.
- 7 – Cut over new IAM version into production. You can either create a new LNKLSTxx member of SYS1.PARMLIB containing the new IAM 6.3 load library and IPL using the new PARMLIB member, or copy the new IAM load library over the old IAM load library just before an IPL is to be done.

All open IAM files must be closed before stopping VIF and copying the new IAM load library over the old one or abends may occur when the IAM files are closed.

90.02 INSTALLING THE IAM INSTALLATION CONTROL LIBRARY

90.02 INSTALLING THE IAM INSTALLATION CONTROL LIBRARY

The IAM product distribution tape contains an Installation Control Library. This library supplies the user with JCL streams to complete the installation of IAM, machine readable documentation, etc.. File two of the installation tape contains IEBCOPY input that will create the Installation Control Library.

You should review member '@INDEX' in this library to determine which members might be helpful to you.

DISK SPACE REQUIREMENT

The Installation Control Library must be loaded to a partitioned data set on disk. You may load it to an existing data set (if it has sufficient space and proper DCB attributes) or allocate and load a new one. The following table shows the allocation parameters for the Installation Control Library:

DATA SET	RECFM	LRECL	BLK SIZE	BLOCKS	PDS DIR BLOCKS
IDP.ICLIAM63	FB	80	3120	50	6

This data set is allocated in blocks so that the system will calculate the correct number of tracks for your device type. If your installation chooses a different blocksize for this data set, you should adjust the number of blocks accordingly.

INSTALLATION INSTRUCTIONS

The JCL below will allocate and load the Installation Control Library. You must make the following changes to reflect your environment:

- A) 'DSN=IDP.ICLIAM63' on the ICLOUT DD statement should be changed to the name you wish to use for the Installation Control Library.
- B) 'VOL=SER=vvvvv' on the ICLOUT DD statement must specify a disk volume where the Installation Control Library will be allocated.
- C) 'UNIT=TAPE' on the ICLIN DD statement must specify a tape drive capable of reading the 6250 BPI or 3480 installation tape.
- D) 'VOL=SER=IAM63T' on the ICLIN DD statement must be changed to 'VOL=SER=IAM63P' if you are loading from a production installation tape.

90.02 INSTALLING THE IAM INSTALLATION CONTROL LIBRARY

```
JOB CONTROL //ICLLOAD EXEC PGM=IEBCOPY,REGION=1024K
FOR LOADING //SYSUT3 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
LIBRARY //SYSUT4 DD UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSPRINT DD DUMMY SEE NOTE 1
//ICLOUT DD DSN=IDP.ICLIAM63, <==USER-CHANGE
// VOL=SER=VVVVVV, <==USER-CHANGE
// UNIT=SYSDA,DISP=(,CATLG), SEE NOTE 2
// DCB=(LRECL=80,BLKSIZE=3120,RECFM=FB),
// SPACE=(3120,(50,10,6),,,ROUND)
//ICLIN DD DSN=IAMICL,
// UNIT=TAPE, <==USER-CHANGE
// DISP=OLD,LABEL=2,
// VOL=SER=IAM63T CHANGE T TO P IF PRODUCTION TAPE
//*
//SYSIN DD *
COPY I=((ICLIN,R)),O=ICLOUT
/*
```

NOTE 1: To list the members being loaded, replace SYSPRINT above with:

```
//SYSPRINT DD SYSOUT=A
```

NOTE 2: To load the members to an existing library, replace ICLOUT above with:

```
//ICLOUT DD DISP=OLD,DSN=IDP.ICLIAM63 <==USER-CHANGE
```

JCL STREAMS MEMBER '@INDEX' is the table of contents for this library. Installation Control Library provides the user procedures for:

- ...Installing the program load library.
- ...Loading the user Documentation library.
- ...Activating the IAM system level VSAM Interface (VIF).
- ...Using a provided test procedure to verify the installation.
- ...Using IAM's VSAM ANALYSIS programs.
- ...And many more.

90.03 IAM PROGRAM LOAD LIBRARY

90.03 IAM PROGRAM LOAD LIBRARY

The IAM program library is distributed on tape in IEBCOPY unloaded format. On MVS , MVS/XA and MVS/ESA systems, it is required that an authorized library be used, as certain modules are linked with an authorization code of 1. A library is authorized if it is accessed via the LINKLIST or its name is in SYS1.PARMLIB member IEAAPFxx. Under MVS/XA and MVS/ESA, even a library accessed via the LINKLIST may have to be listed in IEAAPFxx to be authorized.

INSTALLATION

The IAM program library must be loaded to a partitioned data set on disk. You may load it to an existing data set (if it has sufficient space and proper DCB attributes) or allocate and load a new one. The following table shows the allocation parameters for the product program library:

DISK SPACE REQUIREMENT	PRODUCT	RECFM	LRECL	BLKSIZE	BLOCKS	PDS DIR BLOCKS
	IAM	U	n/a	6144	250	25

This data set is allocated in blocks so that the system will calculate the correct number of tracks for your device type. If your installation uses a different block size for this data set, you should adjust the number of blocks accordingly.

90.03 IAM PROGRAM LOAD LIBRARY

JOB CONTROL FOR LOADING LIBRARY The JCL listed below is supplied on the Installation Control Library ([Section 90.02](#)) with a member name of 'IAMLOAD'. The JCL in the Installation Control Library reflects the type of product tape (Trial or Production) from which it was loaded.

WARNING – The JCL supplied in the Installation Control Library must be used to load the new version. DO NOT use JCL from a prior version's ICL.

```
//IAMLOAD      JOB      (IAM), 'IAM-LOAD LIBRARY INSTALL'
//*****
//*           THIS JOB INSTALLS THE INNOVATION ACCESS METHOD(IAM)
//*           TAPE VOLUME SERIAL MUST BE IAM63P
//*           USER CHANGES:
//*           CHANGE ONLY WHERE INDICATED BY |←-USER CHANGE|
//*           DSN=  USER LOAD LIBRARY
//*           UNIT= TAPE DRIVE UNIT NAME FOR 6250 BPI OR CART
//*           SER=  OUTPUT VOLSER FOR LOAD LIBRARY
//*****
//*
//COPY          EXEC   PGM=IEBCOPY,REGION=1024K
//SYSPRINT      DD     SYSOUT=*
//SYSUT3       DD     UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT4       DD     UNIT=SYSDA,SPACE=(CYL,(1,1))
//LOADOUT      DD     DSN=IDP.MODIAM63,                <-USER CHANGE
//              VOL=SER=VVVVVV,                       <-USER CHANGE
//              SPACE=(6144,(250,25,25),,,ROUND),
//              DCB=(LRECL=0,BLKSIZE=6144,RECFM=U),
//              DISP=(,CATLG),UNIT=SYSDA
//LOADIN       DD     DSN=IAMLOAD,
//              UNIT=TAPE,                              <-USER CHANGE
//              DISP=OLD,LABEL=1,
//              VOL=SER=IAM63T      CHANGE T TO P IF PRODUCTION TAPE
//SYSIN        DD     *
//              COPY INDD=((LOADIN,R)),OUTDD=LOADOUT
//*
```

INSTALLATION INSTRUCTIONS

Submit the member 'IAMLOAD' to load the IAM program library from your installation tape, after you have made the changes to the 'IAMLOAD' procedure to reflect your environment.

* * * **CAUTION** **CAUTION** **CAUTION** * * *

- 1) For current IAM customers, we recommend that you install this version in an authorized library separate from your current IAM library. Once you have tested this version, then copy the modules into your production library replacing the previous version, or replace the previous IAM library name with the new one in your system's LINKLIST.
- 2) Under MVS/XA and MVS/ESA, if you have installed the modules into a LINKLIST library, the operator may have to issue the command `MODIFY LLA,REFRESH` to inform the system of the update.
- 3) All new and existing IAM customers should review [Section 90.01](#) before installing the IAM system.

90.04 ACTIVATING THE IAM VSAM INTERFACE (VIF)

90.04 ACTIVATING THE IAM VSAM INTERFACE (VIF)

IAM's system level VSAM interface (VIF) allows CICS or other ONLINE systems, batch programs, TSO applications and IDCAMS to use IAM files in place of single indexed VSAM KSDS files without any change to programs or job step JCL. Both batch and online (synchronous and asynchronous) processing are supported. VIF supports any programming language which provides VSAM KSDS file creation, random, sequential and skip sequential capabilities.

VIF can be activated at any time after the IAM program load library is ready for use ([Sections 90.01, 90.02, 90.03](#)).

IDCAMS SUPPORT IAM supports most of the IDCAMS functions. The support for IDCAMS is documented in detail in [Section 10.01](#) of this manual. VIF provides support for the following IDCAMS functions:

- DEFINE** – Use the same parameters as you would to define a VSAM cluster, adding the parameter \$IAM. You must specify that a KSDS cluster is to be created. IAM will ignore DATA and INDEX names if they are specified. IAM will create a single sequential file on disk using the cluster name. The file will contain both the data and index information. IAM files on disk are marked with a DSORG of PS. IAM files can be in a VSAM or ICF Catalog.
- DELETE** – Will execute the same for a VSAM cluster or IAM file.
- LISTCAT** – LISTCAT for an IAM file will reflect a non-VSAM file. IDCAMS LISTCAT ALL will result in the dynamic allocation of an IAMPRINT DD to SYSOUT and the production of a file description report.
- PRINT** – Will execute the same for a VSAM cluster or IAM file.
- REPRO** – Will execute the same for a VSAM cluster or IAM file. Use REPRO to copy VSAM clusters to IAM files or back to VSAM.
- VERIFY** – IDCAMS will complete normally, however, since this function is not needed on a IAM file it will be ignored.
- IMPORT** – Is not currently supported by IAM. For a description of comparative functions [see Section 10](#) of this manual.
- EXPORT** – Is not currently supported by IAM. For a description of comparative functions [see Section 10](#) of this manual.

90.04 ACTIVATING THE IAM VSAM INTERFACE (VIF)

VIF INTERFACE VIF is a transparent interface that allows programs to use IAM files instead of VSAM clusters without program or JCL changes.

IAM IS NOT A REPLACEMENT FOR VSAM IN YOUR SYSTEM.

IAM DOES NOT EFFECT THE NORMAL USE OF VSAM IN THE SYSTEM.

VSAM CATALOG MANAGEMENT IS NOT CHANGED.

IAM is an alternative file structure which uses the same control block protocol as VSAM. IAM supports all of the VSAM control blocks that a normal program would use. Higher level languages (ex: COBOL) normally generate these VSAM control blocks for the user.

ACB The ACB is the control block used to OPEN a VSAM cluster. IAM will use the ACB just like VSAM. During OPEN, IAM will fill in the appropriate fields in the ACB and build system control blocks (ex: DEB) which are chained from it.

RPL The RPL is the control block used to request records or information from VSAM. IAM will process RPL commands just as VSAM would have, except it will execute much faster. IAM supports most of the commands that the RPL can pass to VSAM. CI and RBA processing commands are not supported. The RBA address is not returned in the RPL.

EXLST The EXLST is a control block used to specify exit names or addresses to be given control if an error occurs. IAM supports the exit list.

IAM also supports the following VSAM functions:

SHOWCB SHOWCB is used to return information concerning the file itself.

TESTCB TESTCB is used to test the contents of fields in the VSAM control blocks.

MODCB MODCB is used to modify the contents of fields in the VSAM control blocks.

SHOWCAT SHOWCAT is used to return file information contained in the file or Catalog.

IAM builds the VSAM control blocks used by these functions. IAM fills in the fields in these control blocks, just as VSAM would. If a field being requested is not applicable to IAM, a null or zero value will be passed back to the user.

HOW DOES VIF WORK

Once VIF is active in the system, IAM will inspect all VSAM OPEN, CLOSE and Catalog requests.

If the file being referenced is VSAM, the normal IBM routines are given control and VSAM processes the request. No further intercepts are made.

If the file being referenced is IAM, the OPEN of the ACB is simulated. All the appropriate VSAM control blocks are built and hooked up just like a VSAM OPEN. IAM OPENS the file internally as non-VSAM.

All RPL requests (GET, PUT, POINT, etc.) against an IAM file are passed to VIF. VIF deciphers the RPL command. IAM processes the request for the user just as VSAM would have, except much faster.

To use IAM, follow the same rules as you would for VSAM. In fact most end users of VSAM files never need to know that they are referencing an IAM file instead of a VSAM cluster.

IAM uses the values specified in the VSAM DEFINE statement (ex: KEYS, RECORDSIZE). However because the IAM file and overflow structure is different than VSAM's, IAM may change the CI size and overflow area values. [See Section 10.01](#) for details.

Just let IAM take its defaults. IAM's Real Time Tuning, buffer management and overflow structure take the work out of tuning VSAM files. To use IAM to your best advantage however, you should review this manual, especially the *OVERVIEW* section.

90.04 ACTIVATING THE IAM VSAM INTERFACE (VIF)

ACTIVATING VIF The IAM Installation Control Library contains five (5) members to control the IAM VSAM Interface in your system.

1. VIFSTART – A job stream to activate VIF for testing.
2. VIFSTATS – A job stream to check the status of VIF in the system.
3. VIFSTOP – A job stream to deactivate VIF.
4. IAMSTART – A procedure which activates VIF.
5. IAMCOMM – A command to activate VIF at IPL time.

ACTIVATING VIF FOR TESTING You can activate the IAM VSAM interface (VIF) at any time. Review the JCL in member 'VIFSTART' in the Installation Control Library. Change the library name specified in the 'LIB=' statement to reflect your IAM load library. Submit this procedure to activate the VIF interface. This can be done at anytime. This will not effect normal VSAM or system operations. The IAM VIF Interface will remain active until the next IPL of your operating system. VIF must be restarted after each IPL.

NOTE: The VIF Interface must be activated on each machine that IAM is used on. If the IAM program library is not in the LINKLIST, you must use a STEPLIB DD statement pointing to the IAM load library in the jobs processing IAM files.

PROCEDURE FOR TESTING IAM 6.3 WITH PREVIOUS VERSIONS OF IAM ACTIVE If you are a current customer of IAM, you can install and test this new version of IAM without affecting any of your existing jobs. Programs can continue to use prior versions until you choose to have them use the new version.

Install this version of IAM in an authorized library.

If the library is in the LINKLIST, make sure it is not in front of the production version. Execute the VIFSTART procedure from the new Installation Control Library. You will now have two versions of IAM installed in your system. Insert a STEPLIB pointing to the new IAM Load Library in the jobs you wish to test. IAM modules contain a key matching them to the proper VIF Interface. The VIF Interface modules in the system will recognize which version of IAM the STEPLIB contains.

NOTE: The newest version of IAM must be started after all prior versions of IAM are active in the system.

USING IAM AND OTHER SYSTEM SOFTWARE PRODUCTS IAM supports or co-exists with the majority of the many software products that make decisions based on a file's DSORG and or interface with the same MVS services as IAM's VIF. The main consideration when using IAM in conjunction with one of these decision making facilities, especially one that requires a system interface or started task, is the order in which the product interfaces are activated. As a general rule if a program provides a VSAM only service such as altering VSAM buffering or activating LSR it should be started before IAM's VIF. However, if a product is applicable to IAM files, for example, this would include but not be limited to CICS and batch VSAM journaling facilities like DRS™ from BMC as well as VSAM allocation control products like VAM™ from Sterling Software of California, then VIF should be activated first.

90.04 ACTIVATING THE IAM VSAM INTERFACE (VIF)

TESTING VIF STATUS You can test to see if the IAM VIF Interface is active in your system using the member 'VIFSTATS' in the Installation Control Library. Change the library name specified in the 'LIB=' statement to reflect your IAM load library. Submit this procedure to check the status of the VIF Interface. IAM will display a log message indicating the status of IAM (active or inactive). If the VIF Interface is active, IAM will display the IAM module names, version number and their locations in the system.

ACTIVATING VIF AT IPL TIME After you have completed your testing, or possibly before, you will want to start VIF automatically at each IPL. To have the IAM system level VSAM Interface (VIF) started automatically at each IPL, use the following procedure:

1. In member IAMSTART of the Installation Control Library, change the library name specified in the 'LIB=' statement to reflect your IAM load library.
2. Copy the member 'IAMSTART' into SYS1.PROCLIB (or the procedure library used by your installation for system level user procs). This is the start up procedure for the VIF Interface.
3. Copy the member IAMCOMM from the IAM Installation Control Library into the SYS1.PARMLIB member COMMND00 (or the COMMNDxx member used by your installation). This will create a record that will result in a system start command being issued for the IAM start up procedure at IPL time. The record in that member should look like the following statement.

```
COM= ' START IAMSTART '
```

You have now completed the installation of a dynamic IAM system level VSAM interface (VIF) that will automatically be activated each time the operating system is IPL'd.

DEACTIVATING VIF If you wish to deactivate the IAM VIF Interface in the system submit the member 'VIFSTOP' from the Installation Control Library. All VSAM requests will be directly routed to VSAM without any checking for IAM files.

NOTE: To reactivate VIF use the VIFSTART procedure.

WARNING: If you deactivate VIF in a production environment, all OPEN, CLOSE and IDCAMS processing against IAM files will fail.

90.05 TESTING IAM

90.05 TESTING IAM

Once you have completed the loading of the IAM Installation Control Library, the program load library and activating the IAM system level VSAM Interface (VIF), you are now ready to test IAM.

TESTING THE INSTALL

To make a quick test to see if IAM is properly installed in the system submit the job 'TESTVIF' from the IAM Installation Control Library. This job will execute IDCAMS to DEFINE, REPRO and DELETE an IAM file.

A LISTCAT is executed against the IAM file. If IAM is properly installed, LISTCAT will show the file as a non-VSAM file. The IAMPRINT DD statement will display the characteristics of the IAM file. If LISTCAT shows the file as VSAM, the IAM VIF Interface may not be properly installed. Refer to [Section 90.06](#) for helpful hints on diagnosing your situation.

NOTE: Before submitting the test job, review the JCL and control statements, making whatever changes are appropriate.

SMF ANALYSIS

You are now ready to test IAM in your system. The first thing you should do is identify your most active VSAM clusters. To aid you in this function, IAM supplies you with a SMF Analysis program (IAMSMFVS). This program is documented in [Section 40.01](#) of this manual. To use this program you must be collecting SMF type 64 and either SMF type 4 or 30 subtype 4 records.

A sample procedure is supplied as member 'IAMSMFVS' in the IAM Installation Control Library. Review this procedure and make any changes that are appropriate. You can select from your current SYS1.MANx data set or from history SMF tapes. It is recommended that you scan from one week to one month's worth of SMF history to get a true picture of your VSAM usage.

Submit member 'IAMSMFVS' to display the VSAM files with the most activity in your system. The first report will display the top 100 VSAM files in descending EXCP order. The second report will display the VSAM files in data set name sequence.

Converting just a few of your most active VSAM files can give noticeable and immediate benefits to your CICS and batch systems.

NOTE: If IAMSMFVS gives a 'IAM495 NO RECORDS MATCHED SELECTION CRITERIA' message, it probably means that you are not collecting SMF type 64 records. Check SYS1.PARMLIB member SMFPRMxx to see what records SMF is collecting. If type 64 records are not collected, you must change this member. You must issue a SET SMF=xx command to activate the SMF change or wait for your next IPL. After you have collected at least one days worth of data, run the IAMSMFVS program again.

90.05 TESTING IAM

GENERAL TESTING Once you choose which VSAM files you wish to test with, use IDCAMS to DEFINE the IAM file(s). Take your existing IDCAMS procedure and add the parameter OWNER(\$IAM) to the DEFINE statement.

Next, use IDCAMS REPRO to copy the VSAM cluster to the IAM file. Run a LISTCAT ALL against the IAM file with a IAMPRINT DD statement specified. The IAMPRINT report will give you the number of tracks used by IAM in addition to all of the IAM characteristics. The number of tracks used includes all of the prime records loaded into the file and the overflow areas. IAM will automatically release all of the unused space allocated, if secondary allocation has been specified.

Compare this to the VSAM cluster. Unfortunately, it is not easy to tell how much space VSAM really used. If the Data Component HI-USED-RBA is close to HI-ALLOC-RBA, the VSAM file is using most of its allocated space.

If you wish to calculate the used space do the following calculation. Use the values found in the DATA component.

Divide the HI-USED-RBA by CISIZE. Divide the result by CIs per CA. Multiply this value by the number of tracks per CA. This will yield the number of tracks used by the Data Component and imbedded index (if specified). Add in the tracks for the Index. This will give you the approximate number of tracks used by the VSAM cluster.

One of the easiest tests you can make is to compare an IDCAMS REPRO of a VSAM cluster and IAM file to tape. Many applications use REPRO to make backup copies of their VSAM files. When the jobs are completed, compare the wall clock time, CPU time (TCB and SRB) and EXCPs issued by each job. This will give you a comparison of the resources used to sequentially read an entire VSAM cluster versus using an IAM file.

PARALLEL TESTS You are now ready to run parallel tests of jobs using VSAM clusters compared to IAM file(s). Select an application to test, using the IAMSMFVS report as a guide.

Use IDCAMS to DEFINE the IAM file specifying a different cluster name(s) then REPRO the VSAM cluster(s) into the IAM file(s). Use a copy of any file(s) which may be changed by the application. Run the production job against the VSAM cluster(s). Use the same JCL for the parallel run, changing the JCL to point to the new IAM file(s). Insert a STEPLIB pointing to the IAM load library if it is not in the LINKLIST. Re-execute the production run using the IAM file(s).

Compare the results from each run. If you do not have the reporting tools necessary to get the SMF data you require, use the IAM supplied program IAMSMF ([Section 41.01](#)). Compare the statistics for each job. You will want to measure the wall clock time, CPU time (SRB and TCB), and EXCP counts (DATA and INDEX for VSAM). In addition, you will want to compare the disk space used by IAM versus the VSAM clusters.

TESTING WITH CICS Review the CICS Special Considerations ([Section 90.10](#)) before you start using files under CICS.

IAMINFO DD STATEMENT Add the following statement to each of the steps using IAM.

```
//IAMINFO DD SYSOUT=*
```

If this statement is present IAM will print a Run Time Statistics report each time an IAM file is OPENed and CLOSEd. The report will display the characteristics of the file, run time statistics (memory use, read and write I/Os, etc.), count of each command issued (GETS, PUTS, ADDS, DELETES, POINTS, etc.), number of buffers used and additional information. There is very little overhead associated (50 I/O's for the lines printed) with the report and it can be very valuable to you and INNOVATION in determining what each job is doing. ([See Section 12 IAM Run Time Statistics for additional details on using this report](#))

90.06 IAM INSTALLATION QUESTIONS

90.06 IAM INSTALLATION QUESTIONS

QUESTIONS REGARDING INSTALLATION

Here are the answers to some common questions that may be encountered after the installation of IAM.

Question: I have placed the parameter OWNER(\$IAM) in the VSAM DEFINE, but the file is created as a VSAM cluster:

Answer: There are several possibilities:

First, the IAM system level VSAM Interface (VIF) is not active in this system. Run the procedure VIFSTATS in the Installation Control Library to see if VIF is active. If VIF is not active, run VIFSTART to activate VIF.

Second, the IAM load library is not in the LINKLIST. The JCL LOG in this case may indicate a S806 abend for an IAM module name. If you do not wish to put IAM in the LINKLIST, insert a STEPLIB pointing to the IAM library in the IDCAMS DEFINE job. Third, the IAM library is not authorized. The IAM library must be authorized.

Question: I get VSAM error message IEC1611 pointing to an IAM file.

Answer: There are several possibilities:

First, the IAM system level VSAM Interface (VIF) is not active in this system. Run the procedure VIFSTATS in the Installation Control Library to see if VIF is active. If VIF is not active run VIFSTART to activate VIF.

Second, the IAM load library is not in the LINKLIST. If you do not wish to put IAM in the LINKLIST, insert a STEPLIB pointing to the IAM library in the job referencing the IAM file.

Third, the IAM library is not authorized. The IAM library must be authorized.

Question: I converted a file to IAM and all of a sudden it went back to a VSAM cluster.

Answer: There was a job executed against the IAM file which DELETED and re-DEFINED the file as VSAM without the OWNER(\$IAM) parameter specified. You must find this job and change the DEFINE procedure. Use the IAMSMF program copy and query commands to report on all the jobs which have defined the cluster so you can identify the IDCAMS job causing the problem. ([See Section 41.10 Example 3](#)).

Question: I get a RPL error code of X '6C' (followed by a COBOL file status 90) trying to load, add, or update a record on a IAM file. The same program works with VSAM.

Answer: This completion code indicates a record length error. Check the DEFINE for the IAM file. If the RECORDSIZE parameter specifies the same average record length as maximum, IAM will consider this to be a fixed length file. VSAM considers all files as capable of having variable length records. If the file is really variable, change the RECORDSIZE in the DEFINE to show a smaller value of the average record size than for the maximum record size. Other alternatives are to enable Data Compression on the file, to use the IAM CREATE override keyword VARIABLE or to change IAM's Global Options Table to default the VSAM RECFM to VARIABLE for all files.

90.06 IAM INSTALLATION QUESTIONS

Question: Can I test a new version of IAM if I am already using VIF in a product from an older version of IAM.

Answer: Yes, just insure the newest version of IAM's VIF is started after all prior versions of IAM are already active in the system.

Question: Can I use IAM and VIF with other software products that do VSAM allocation control or otherwise intercept VSAM.

Answer: Yes, IAM supports or co-exists with the majority of the many software products that make decisions based on a file's DSORG and or interface with the same MVS services as IAM's VIF. The main consideration when using IAM in conjunction with one of these decision making facilities, especially one that requires a system interface or started task, is the order in which the product interfaces are activated. As a general rule if a program provides a VSAM only service such as altering VSAM buffering or activating LSR it should be started before IAM's VIF. However, if a product is applicable to IAM files, for example this would include but not be limited to CICS and batch VSAM journaling facilities like DRST[™] from BMC as well as VSAM allocation control products like VAM[™] from Sterling Software of California, then VIF should be activated first.

90.07 INSTALLING THE USER DOCUMENTATION

90.07 INSTALLING THE USER DOCUMENTATION

Your IAM installation tape now contains a SOFTCOPY version of the documentation that can be downloaded and used with BOOKMANAGER to view the manual online. The SOFTLOAD member in the ICL library will download the PDS from the installation tape containing the SOFTCOPY doc. The BLDBOOKS member in the ICL library can then be used to create the sequential files needed to use the books with BOOKMANAGER. Follow the directions in the BLDBOOKS member to change the data set names to your installations specifications. BLDBOOKS will also create a file that contains the bookshelf index. We have also included a bookshelf in the ICL library, member BKSHELF, that can be copied into a sequential file with the appropriate file attributes for BOOKMANAGER. The bookshelf must be DSORG=PS,RECFM=VB,LRECL=255, BLKSIZE=8000 and can be allocated as just 1 track.

Your IAM installation tape contains IEBCOPY input which will create a documentation library. This library is a machine readable version of the IAM user documentation.

LIBRARY CONTENTS

The documentation library, when loaded, will contain members corresponding to the documentation sections listed in the manual.

<u>MEMBER</u>	<u>SECTIONS</u>	<u>DESCRIPTION</u>
IDXINTRO	00	TABLE OF CONTENTS *
MODIFICA	00	SUMMARY OF MODIFICATION
OVERVIEW	01	IAM PRODUCT OVERVIEW & IAM V6.3 ENHANCEMENTS
TECHDESC	02	IAM TECHNICAL DESCRIPTION
SMS	02	SMS SUPPORT IN IAM
STRUCTUR	03	IAM DATA SET STRUCTURE
IDCAMS	10	ACCESS METHOD SERVICES
IAMOVRID	11	IAM OVERRIDE STATEMENTS
IAMINFO	12	RUN TIME STATS AND SMF RECORDING
VSAMANAL	40	IAM SMF VSAM USAGE ANALYSIS
SMFANAL	41	IAM SMF ANALYSIS PROGRAM
SPACESAV	42	IAM SPACE SAVINGS
TRACE	43	IAM TRACE/DEBUG
MONITOR	44	IAM EVENT RECORDING MONITOR
RECOVERY	45	IAM RECOVERY PROGRAM
ESDS	60	ESDS SUPPORT
MESSAGES	80	MESSAGES AND CODES
INSTALL	90	INSTALLATION AND VIF ACTIVATION
GLOBALOP	91	IAM GLOBAL OPTION FACILITY

* Includes Summary of Modifications and Document Notation.

DISK SPACE REQUIREMENT

The user documentation must be loaded to a partitioned data set on disk. You may load it to an existing data set (if it has sufficient space and proper DCB attributes) or allocate and load a new one. The following table shows the allocation parameters for the documentation library:

<u>DATA SET</u>	<u>RECFM</u>	<u>LRECL</u>	<u>BLK SIZE</u>	<u>BLOCKS</u>	<u>PDS DIR BLOCKS</u>
IDP.DOCIAM63	FB	80	6160	720	3

This data set is allocated in blocks so that the system will calculate the correct number of tracks for your device type. If your installation uses a different block size for this data set, you should adjust the number of blocks accordingly.

90.07 INSTALLING THE USER DOCUMENTATION

INSTALLATION INSTRUCTIONS

The following JCL is supplied in the Installation Control Library. To load the user documentation make the following changes to reflect your environment:

- A) Change 'DSN=IDP.DOCIAM63' on the DOCOUT DD statement to the name of the library you wish to create for the documentation library.
- B) Change 'VOL=SER=vvvvvv' on the DOCOUT DD statement to specify a disk volume where the documentation library will be allocated.
- C) Change 'UNIT=TAPE' on the DOCIN DD statement to specify a tape drive capable of reading the 6250 BPI or 3480 installation tape.

JOB CONTROL FOR LOADING

The JCL below is supplied on the Installation Control Library (see Section 90.02) as member 'DOCLoad'. The JCL in the Installation Control Library reflects the type of installation tape, trial or production, from which it was loaded.

```
//IAMDOC      JOB      (IAM), 'IAM-DOC'
//*****
//*
//*          LOAD DOC LIBRARY FROM PRODUCTION DISTRIB
//*
//*****
//*          REQUIRED JCL CHANGES:
//*          VOL=SER=VVVVVV FOR DOCOUT
//*
//*****
//*
//COPY        EXEC    PGM=IEBCOPY,REGION=1024K
//SYSPRINT    DD      SYSOUT=*
//SYSUT3      DD      UNIT=SYSDA,SPACE=(CYL,(1,1))
//SYSUT4      DD      UNIT=SYSDA,SPACE=(CYL,(1,1))
//DOCOUT      DD      DSN=IDP.DOCIAM63,                <-USER-CHANGE
//              VOL=SER=VVVVVV,                        <-USER-CHANGE
//              SPACE=(6160,(600,25,2),,ROUND),
//              DCB=(LRECL=80,BLKSIZE=6160,RECFM=FB),
//              DISP=(,CATLG),UNIT=SYSDA
//DOCIN       DD      DSN=IAMDOC,
//              UNIT=TAPE,
//              DISP=OLD,LABEL=3,
//              VOL=SER=IAM63T      CHANGE T TO P IF PRODUCTION TAPE
//SYSIN       DD      *
//              COPY INDD=((DOCIN,R)),OUTDD=DOCOUT
//*
```

NOTE 1: To list the documentation members being loaded, replace SYSPRINT above with:

```
//SYSPRINT DD  SYSOUT=A
```

NOTE 2: To load the documentation members to an existing library, replace DOCOUT above with:

```
//DOCOUT DD  DISP=OLD,DSN=IDP.DOCIAM63                <==USER-CHANGE
```

BROWSING THE ISPF BROWSE DOCUMENT

Once loaded, the library may be browsed online using or any comparable online browse function under TSO or any other online system.

90.07 INSTALLING THE USER DOCUMENTATION

PRINTING THE DOCUMENT

The documentation library contains upper case, lower case and standard special characters. Most printers and terminals that do not support lower case will print or display the lower case characters in upper case. The documentation will never contain more than 80 characters on a line, and can be printed on narrow (8.5x11 inch) paper if standard 10 character/inch spacing is used. Sufficient margin is allowed on the left to 3-hole punch the documentation after printing (or to use pre-punched narrow paper). You must align the printer so that print column 80 is printed at the right edge of the 8.5x11 form.

JOB CONTROL TO PRINT THE DOCUMENT

The JCL below is supplied on the Installation Control Library as member 'DOCPRINT'. The program 'IAMDOCP' x is supplied in the IAM program load library. Make the changes to 'DOCPRINT' required to reflect your environment.

Submit member 'DOCPRINT' to print the entire document (Sections may be deleted by omitting the appropriate PRINT statement, enabling you to print a document customized to your needs). The following changes must be made to reflect your environment:

- A) Change 'DSN=IDP.MODIAM63' on the STEPLIB DD statement to the name of the IAM program load library.
- B) Change 'DSN=IDP.DOCIAM63' on the SYSUT1 DD statement to the name of the library you loaded the documentation library.
- C) Make any required changes to SYSUT2 for special SYSOUT class, special forms, multiple copies, character sets, etc., as required by your installation.

```
//IDPPRINT EXEC PGM=IAMDOCP
//STEPLIB DD DISP=SHR,DSN=IDP.MODIAM <==USER-CHANGE
//SYSPRINT DD SYSOUT=A
//SYSUT1 DD DISP=SHR,DSN=IDP.DOCIAM63 <==USER-CHANGE
//SYSUT2 DD SYSOUT=A <==USER-CHANGE
//SYSIN DD *
***
*** TO PRINT THE ENTIRE DOCUMENT
*** USE ALL OF THE PRINT STATEMENTS SUPPLIED
***
PRINT MEMBER=IDXINTRO
PRINT MEMBER=MODIFICA
PRINT MEMBER=OVERVIEW
PRINT MEMBER=TECHDESC
PRINT MEMBER=SMS
PRINT MEMBER=STRUCTUR
PRINT MEMBER=IDCAMS
PRINT MEMBER=IAMOVRID
PRINT MEMBER=IAMINFO
PRINT MEMBER=VSAMANAL
PRINT MEMBER=SMFANAL
PRINT MEMBER=SPACESAV
PRINT MEMBER=TRACE
PRINT MEMBER=MONITOR
PRINT MEMBER=RECOVERY
PRINT MEMBER=ESDS
PRINT MEMBER=MESSAGES
PRINT MEMBER=INSTALL
PRINT MEMBER=GLOBALOP
```

90.08 DFSMS SUPPORTED BY NEW FEATURES IN IAM V6.2

90.08 DFSMS SUPPORTED BY NEW FEATURES IN IAM V6.2

IAM V6.2 and higher contain several new features sure to be of interest to users of Data Facility Product (DFP) Version 3.1 and above. Primarily the definition and allocation of IAM data set's through the DFP's new Storage Management Subsystem (SMS) which also allows specification of availability, archive/migration and back-up criteria.

The principal IAM DFSMS features are:

- Integration of IAM's System Level VSAM Interface (VIF) and the DFSMS Storage Management Subsystem
- JCL definition and allocation of VIF managed IAM data sets
- TSO Allocate command definition and allocation of IAM data sets
- Allocation of temporary IAM data sets managed through VIF

INTEGRATION OF IAM WITH SMS

The Storage Management Subsystem (SMS) is a component of DFP 3.1 and above that can help storage administrators balance requirements for performance, availability, and space utilization. IBM explains that SMS and ISMF provide a means of integrating the operating system and any external software used to manage peripheral DASD storage so that site-wide standards can be set for all data sets. This is the concept of the DFSMS environment. With SMS Classes assigned to data sets software can direct storage management tasks previously handled by users and/or storage administrators. Three SMS classes, the Data Class, the Management Class, and the Storage Class establish defaults that set the allocation, performance and availability of data sets. When an IAM data set is defined with one or more of these classes, SMS provides default characteristics for the allocated data set. DFSMS was meant to provide a means of managing peripheral storage so that shop-wide standards can be set for all data sets. Storage administrators can also use SMS classes to standardize the allocation attributes of their IAM data sets. Attributes provided by SMS classes will override the standard defaults for IAM data sets. The order of precedence for IAM file attributes is now:

1. attributes established by SMS (ACS routines) for IAM data sets.
2. attributes specified for a particular SMS Data Class
3. default IAM attributes

DATACLAS

A \$IAM in the Data Class name (the keyword parameter is DATACLAS) can be used to define a data set as an IAM file. Default attributes provided by the Data Class for an IAM data set include:

- Data set organization (IAM INDEXED)
- record size
- keylength and offset
- primary and secondary space quantity
- control interval size/blocking factor
- freespace/overflow percentages
- share options

Data Classes can make it simpler to standardize attributes, such as CISIZE/BLOCKING FACTOR, which can be used to improve performance. Shop-wide implementation of SMS however will usually require a great deal of work. For example, to maintain currently established IAM file attributes through the use of Data Classes, it may be necessary to create hundreds of Data Classes.

90.08 DFSMS SUPPORTED BY NEW FEATURES IN IAM V6.2

STORCLAS The Storage Class (the keyword is STORCLAS) lets the storage administrator specify performance and availability characteristics for data sets. A Storage Class must be specified if the data set is to be SMS managed, Data Class or Management Class need not be specified at all. Specifying \$IAM in the Storage Class name will define a data set as an IAM file. The following are examples of SMS Storage Class attributes that can be specified for IAM data sets:

- direct millisecond response
- sequential millisecond response
- continuous availability
- guaranteed space

MGMTCLAS Management Class attributes (the keyword is MGMTCLAS) control the backup migration, deletion and archiving/auto-recall of data sets. Attributes that can be provided through Management Class for IAM data sets include:

- release of unused space
- expiration after nn days of non-use
- expiration after nn days regardless of use
- retention limit
- archiving after nn days of non-use
- backup frequency
- number of backups

If SMS is inactive, IAM data sets can still be established through the use of the Data Class parameter but Storage Class and Management Class are never valid unless SMS is active.

IAM DATA SETS THROUGH JCL

The creation of permanent IAM data sets with JCL, is accomplished with the use of new DATACLAS keyword. A number of other new keyword parameters have also been added to the JCL DD statement for SMS which are ignored if SMS is inactive. DFP's new SMS JCL parameters, however, have limited function and do not provide all the facilities of IDCAMS. Because of these limitations, to duplicate the attributes of an existing data set through JCL, it may be necessary to set up a unique data class for that data set. This limitation holds true for VSAM as well as for IAM data sets.

The primary means of defining a data set's allocation characteristics (RECORDG, KEYLEN, CISIZE, etc.) is with DATACLAS parameters. While some of these parameters — for example, the RECORG parameter — can be specified on the DD card the only way to specify other parameters in the JCL, such as the CISIZE for an IAM or VSAM data set is through a predefined Data Class. The OWNER parameter is another attribute that can only be specified through a predefined Data Class.

If a Data Class with the proper CISIZE, FREESPACE, and SHAREOPTIONS is available but has the wrong key length. The KEYLEN can be overridden on the JCL statement when the data set is created. Additionally, a LIKE parameter is available, but it is not the same as the VSAM MODEL parameter available in IDCAMS. The only attributes copied from a data set with the LIKE parameter are already available as JCL keyword parameters. The VSAM MODEL parameter is not available through JCL. For JCL to be used successfully to allocate IAM or VSAM data sets, a storage administrator may have to set up many appropriately predefined Data Classes.

90.08 DFSMS SUPPORTED BY NEW FEATURES IN IAM V6.2

IAM DATA SETS THROUGH THE TSO ALLOCATE COMMAND With DFP 3.1 and higher releases, the TSO Allocate command recognizes SMS parameters, such as DATACLAS, MGMTCLAS, and STORCLAS. In addition, it accepts the same keyword parameters for IAM and VSAM data sets as JCL. Some but not all attributes of an IAM data set can be overridden with keywords of the Allocate command. If SMS is used, the volume and unit parameter on the Allocate command are not recommended; these values are provided by the Storage Class which is required for all SMS-managed data sets. Until appropriate Data Classes have been set up, it may be easier to continue to use the TSO DEFINE command.

TEMPORARY IAM DATA SETS DFP 3.1 and later releases make it possible to allocate temporary IAM data sets managed through VIF within any job. They will only exist within a job and must reside on storage managed by SMS. To define a temporary IAM data set, you must specify the STORCLAS parameter, specify a data set name that begins with & or &&, such as DSN=&TEMP and specify \$IAM in the data set name, Data Class name or Storage Class name. Temporary data sets are deleted upon completion of the job in which they were defined.

DFSMS SUMMARY \$IAM in the SMS Data Class or Storage Class name can now be used to cause files to be defined as IAM. The use of SMS's new JCL allocation attributes may not be appropriate for all IAM and VSAM data sets. One way to accommodate users of IAM and VSAM data sets, is to specify a particular file's attributes in a unique Data Class. Additionally, when using the SMS JCL allocation feature, specifying \$IAM in the data set name will still define a file as IAM.

90.10 SPECIAL CONSIDERATIONS FOR CICS WITH COMPATIBLE FORMAT FILES

90.10 SPECIAL CONSIDERATIONS FOR CICS WITH COMPATIBLE FORMAT FILES

IAM supports VSAM processing in all current versions of CICS. There are no CICS usermods required when IAM's VSAM Interface (VIF) is used. No application program changes are required to use IAM files in place of VSAM clusters. There may be some changes to CICS parameters required, as described below, if certain features relating to VSAM are being used.

IAM files may be used in place of any single index KSDS VSAM clusters. IAM files do not support alternate index processing and can not be used in place of VSAM clusters which are alternate indexes or have alternate indexes unless other arrangements are made.

FILE CONTROL TABLE IAM files are defined in the File Control Table (FCT) as VSAM files. The only difference is that IAM files must not be defined as being part of an LSR pool.

LOCAL SHARED RESOURCES (LSRPOOL=) IAM does not support Local Shared Resources (LSR) the way VSAM does. The IAM concept of Real Time Tuning is far superior to VSAM's LSR. IAM manages its own buffer pools. IAM's buffer pool and VSAM's cannot be shared. CICS Ver 1.7 and above default to Local Shared Resources. FCT's that are to be used for IAM files in these CICS systems must specify the parameter LSRPOOL=NONE.

For CICS 1.6 users, if SERVREQ=(...,SHARE,...) is specified, it must be removed for IAM files to remove them from the LSR pool.

IAM's Real Time Tuning buffer management will provide a level of performance that is superior to the benefits offered by VSAM's LSR buffer look asides at a fraction of the CPU overhead.

STRING NUMBER (STRNO=) CICS FCT's should also be checked to insure the STRNO= parameter is specified and contains an adequate value. This parameter may have previously been specified as part of the LSR pool definition and consequently may not have been specified in the FCT. STRNO= is needed in the FCT entry for an IAM file because it controls the number of transactions allowed concurrent access to the file. STRNO= should always specify a minimum value of 2. If a file is expected to have a very high number of concurrent users a value of 10 or more may be needed. The value specified in STRNO= should actually be equal to or slightly higher than the maximum number of users who will concurrently be accessing the IAM file for BROWSE or GET-FOR-UPDATE operations. If the value specified is not adequate, you will get excessive 'waiting for strings' conditions from CICS. Each string specified results in an additional (keylength + 24) bytes of storage required by IAM in the VSAM control block area. Refer to the OSCORE section for complete information on storage requirements.

VSAM SUBTASKING VSAM subtasking is now supported in MVS/XA or MVS/ESA with IAM Version 6.1 Level 08 or higher. Previous releases of IAM required you to remove VSAM subtasking. This restriction has now been lifted.

IOCP PARAMETER The SIT IOCP parameter may need to be reduced or set to 0, if it has been specified. This parameter indicates, as a percentage, the number of I/O requests that must be posted before CICS will be dispatched. This can cause CICS to go into a permanent wait if there are multiple I/O requests outstanding to a single VSAM or IAM file, and more than one has to complete to satisfy the count. IAM can significantly reduce physical I/O operations the resulting redistribution of the physical I/O activity amongst the files may make the problem more probable than it had been previously, if the number of required posted requests is not lowered.

90.10 SPECIAL CONSIDERATIONS FOR CICS WITH COMPATIBLE FORMAT FILES

XA AND ESA OSCORE REQUIREMENTS For MVS/XA and MVS/ESA customers, IAM primarily uses storage above the line. However, there is need for some storage below the line to accommodate applications running in AMODE 24, and for data areas that are required to reside below the line by MVS. For CICS customers, all below the line storage come out of OSCORE. The use of below the line storage has been revised in Version 6.1 to reduce fragmentation, and provide additional storage relief for CICS 2.1 and applications which indicate that VSAM control blocks can be above the line. File load processing has also been enhanced to place buffers and other storage areas above the line.

STORAGE FOR EACH OPEN FILE As a general guideline, after taking into account the storage required for the IAM load modules, IAM files require about 2K to 4K more OSCORE than a VSAM file in an LSR pool. The following numbers provide approximations of OSCORE storage use for each open file, and reflect a typical 1/4 track blocked IAM file residing on a 3380, with index and data compression. Certain storage sizes vary, depending on file attributes, so actual storage required for any particular file may be slightly different. The example is for a STRNO value equal to 7 and assumes a keylength of 32 bytes. For additional strings, add (24 + keylength) bytes per string to the ACCESS VSAM control block storage requirement. These storage numbers reflect what is required during processing, after a file is opened. Open and Close processing may require some additional storage. It is recommended that at least 32K of additional OSCORE be available to support Open and Close.

<u>APPROXIMATE</u>	<u>ACCESS</u>	<u>CREATE</u>
Base requirement:	5K	6K
VSAM CB (above line in CICS 2.1):	2K	2K
Blocksize, rounded to 4k for Create:	12K	
	7K	20K

CICS users who frequently CLOSE and OPEN IAM files, without bringing CICS down, should set core rounding options in the IAM global options table, to help prevent fragmentation. This will round the getmain for the VSAM control blocks to a 4K boundary. This option is of particular benefit when the size of the VSAM control block area approaches or exceeds 4K due to large keylengths and/or STRNO values. Other large getmains, such as for the buffers, default to rounding. Smaller getmains, those less than 4K, do not require the rounding.

LOAD MODULE STORAGE Storage is also required for the IAM modules. Only the required module(s) are loaded. For example, to access a file without data compression, and with no overrides, requires modules IAMACC and IAMOPT. The modules are reentrant, so only one copy is required, regardless of the number of files OPENed. The modules, IAMACC, IAMCRT and IAMCOMPO can all be placed in LPA (Link Pack Area), to avoid occupying OSCORE storage. Placing modules in LPA will not reduce virtual storage. It will however appropriate storage from LPA and reduce OSCORE.

<u>MODULE</u>	<u>STORAGE</u>	<u>FUNCTION</u>
IAMACC	32K	File access processor
IAMCRT	32K	File load processor
IAMCOMPO	8K	Data compression
IAMOPT	2K	Global Options Table
IAMOVRIID	12K	IAM Override Processing
	86K	Maximum for all IAM routines

NOTE: Modules placed in LPA will require special consideration when maintenance is applied.

90.10 SPECIAL CONSIDERATIONS FOR CICS WITH COMPATIBLE FORMAT FILES

CSA STORAGE REQUIREMENTS IAM'S only use of CSA is for the VIF interface screen modules, such as Open/Close/Locate. CSA storage required is approximately 20K. There is an optional additional requirement of 24K to support SMS JCL allocation when IAM is not in the link list. This optional support is set by the LOADJCLALLOC parameter in the global options table. For a more exact amount, review the modules listed by the IAMSTATS program, and refer to the IAM load module library for the exact sizes of these modules.

CONVERTING VERSION 4.0 USERS TO VIF Most command level applications can be easily converted to use IAM V6.3 files through the VIF interface simply by changing the file's FCT entry from an ISAM FCT to a VSAM FCT. Files may be migrated to the VIF interface one file at a time while others are left under the control of ISAM FCT's.

There are considerable benefits to accessing IAM files through the VIF interface, for example, data buffers managed through ISAM FCT's remain below the 16MB line and are never chained.

NOTE: IAM files to be accessed by CICS applications using VSAM FCT's must have been DEFINE'd (or DEFINE RECATALOG'ed) using IDCAMS in a system with VIF active. Likewise, IAM's VSAM Interface, VIF must be active in the system in which CICS is run. This is explained in detail in [Section 10](#) of the IAM V6.3 Users Manual.

ISAM SUPPORT If you are currently using real ISAM files in the CICS system, it is recommended that the application be converted to use IAM files by changing the FCT's entry to VSAM. In CICS Ver 1.7 and above IBM has dropped ISAM support from CICS so this conversion is inevitable.

90.30 TECHNICAL OVERVIEW AND EXCEPTION CONSIDERATIONS

90.30 TECHNICAL OVERVIEW AND EXCEPTION CONSIDERATIONS

IAM supports single index (KSDS) VSAM files referenced sequentially or by key. The following are not KSDS files: Entry Sequenced (ESDS) VSAM clusters (sequential files like BSAM/QSAM), Relative Record (RRDS) VSAM clusters (like BDAM files), and LINEAR VSAM clusters. IAM does not support the creation of or processing of Alternate Indexes. It is our intention to support Alternate Indexes in a future release.

IMS files are usually ESDS and RRDS VSAM clusters. DB2 files are special format VSAM files. Records are accessed by internal tables built by DB2. IMS and DB2 cannot use IAM files.

IAM creates a physical file structure on disk which is simpler and far superior than VSAM. IAM builds and supports all of the VSAM control blocks used by most programs accessing VSAM clusters. If a field in a VSAM control block does not apply to IAM, zero or some other dummy value will be inserted.

IAM messages and return codes are comparable to VSAM. IAM logs its own special error messages to the JCL LOG. IAM's log messages may have a different meaning than VSAM. Check Messages and Codes ([Section 80](#)) for an explanation of the IAM message. The OPEN/CLOSE and I/O message areas may not contain the VSAM error message.

USING IAM AND OTHER SYSTEM SOFTWARE PRODUCTS

IAM supports or co-exists with the majority of the many software products that make decisions based on a file's DSORG and or interface with the same MVS services as IAM's VIF. The main consideration when using IAM in conjunction with one of these decision making facilities, especially one that requires a system interface or started task, is the order in which the product interfaces are activated. As a general rule, if a program provides a VSAM only service such as altering VSAM buffering or activating LSR it should be started before IAM's VIF. However, if a product is applicable to IAM files, for example, this would include but not be limited to CICS and batch VSAM journaling facilities like DRST[™] from BMC as well as VSAM allocation control products like VAM[™] from Sterling Software of California, then VIF should be activated first.

PROGRAMS CHECKING FOR VSAM

If a DD statement references an IAM file, the DSORG field in the JFCB will be marked as VSAM. If an OBTAIN is issued, the DSORG in the VTOC will indicate DA. Since IAM files are non-VSAM, they do not exist in the VVDS. Locate requests for an IAM file will indicate it is non-VSAM. Superlocate requests for an IAM file will indicate VSAM.

SECURITY

IAM supports the same level of security as VSAM. Support for RACF (or RACF compatible), is the same as for any OS file. OS PASSWORD is not supported from the DEFINE statement. If it is needed, it must be coded the same as any non-VSAM file.

ADDITIONAL DEFINE PARAMETERS

Special Considerations – A VSAM DEFINE for the most part will have the same meaning to IAM as it does to VSAM. However, because IAM's file structure is different than VSAM some of the VSAM parameters are changed or ignored. In addition, there are additional parameters available to the user of an IAM file. IAM gives expanded capabilities to the end user. The IAM parameters are coded using the IAM Control Override DD statement (IAMOVRID). Check [Section 11.01](#) for the Override documentation.

90.30 TECHNICAL OVERVIEW AND EXCEPTION CONSIDERATIONS

PROCEDURE FOR TESTING IAM 6.3 WITH PREVIOUS VERSIONS OF IAM ACTIVE

If you are a current customer of IAM, you can install and test this new version of IAM without effecting any of your existing jobs. Programs can continue to use prior versions until you choose to have them use the new version.

Install this version of IAM in an authorized library. If the library is in the LINKLIST, make sure it is not in front of the production version. Execute the VIFSTART procedure from the new Installation Control Library. You will now have two versions of IAM installed in your system. Insert a STEPLIB pointing to the new IAM Load Library in the jobs you wish to test. IAM modules contain a key matching them to the proper VIF Interface. The VIF Interface modules in the system will recognize which version of IAM the STEPLIB contains.

When you have completed testing the new version, load the IAM program library over the prior version or place this library ahead of the prior version in the LINKLIST. If necessary, change the IPL start procedure (IAMSTART) to point to the new library.

NOTE: The newest version of IAM must be started after all prior versions of IAM are active in the system.

90.40 INSTALLING THE IAM ISPF DIALOG (IAM 6.3)

90.40 INSTALLING THE IAM ISPF DIALOG (IAM 6.3)

INSTALLING THE PANELS AND MESSAGES LIBRARIES

The IAM ISPF Dialog Panels and messages are contained in 2 libraries on the installation tape. To install these libraries, submit IAMSPFIN from the Installation Control Library after making the following changes as required for your environment:

- A) Update the job card.
- B) Change ISPLIB='ISPF.PANEL.LIBRARY' to the name of the library you wish to contain the IAM panels.
- C) Change ISPLIB='ISPF.MESSAGES.LIBRARY' to the name of the library you wish to contain the IAM messages.
- D) Change UNIT=TAPE to specify a tape drive capable of reading the 6250 BPI or 3480 installation tape.

If you have installed the IAM panels and messages in libraries that are not specified in the ISPLIB and ISPLIB DD statements in the appropriate TSO logon procs, then you must add the dataset name specified for the panel library to the ISPLIB concatenation, and add the dataset name specified for the message library to the ISPLIB concatenation.

ADD THE IAM ISPF DIALOG TO AN ISPF OPTIONS MENU

You must add an option to the ISPF/PDF primary option menu (ISR@PRIM), or an option menu of your choice, to invoke the IAM ISPF DIALOG program.

Add a line that describes the IAM ISPF DIALOG to the PANEL BODY.

Example:

```
% I +IAM – IAM Dataset Utilities
```

Add a line that selects the IAMISPF program to the lines that translate the user entered options into the appropriate panel or program name in the PROC section of the panel definition.

Example:

```
&ZSEL = TRANS( TRUNC (&ZCMD,')
0,'PANEL(ISPOPTA)'
1,'PGM(ISRBRO)'
.
other panel options
.
I,'PGM(IAMISPF)' <=== Add
.
other panel options
:
:
X,'EXIT'
*,'?')
```

ADD THE IAM LOAD LIBRARY TO THE TSO LOGON PROCS IF REQUIRED

If IAM has not been installed in a linklist library, you must make the IAMISPF load module, and other required IAM system modules, available under ISPF by adding the IAM load library to the STEPLIB DD statement in the appropriate logon procs. Please note that the use of the ISPLIB DD for this purpose is strongly discouraged. If any of the IAM utility programs are executed from within IAMISPF via option U, they will NOT be loaded from the ISPLIB DD. Additionally, some of the IAM utility functions require APF authorization which will require that all libraries specified in the STEPLIB DD be authorized. Please refer to [Section 90.41](#) for additional information on APF authorization.

90.41 APF AUTHORIZATION FOR THE IAM ISPF DIALOG

90.41 APF AUTHORIZATION FOR THE IAM ISPF DIALOG

FOREGROUND APF AUTHORIZATION Some of the utility functions available under option U utilize system facilities that require APF (Authorized Program Facility) authorization. If any of these functions are used from within the IAM ISPF DIALOG, you must provide authorization for those utility programs as discussed below.

FUNCTIONS THAT NEED AUTHORIZATION The DIAGNOSE and PRINT functions of IAMRECVR require APF authorization. If you do not wish to execute these functions under the IAM ISPF DIALOG, you will not have to provide APF authorization for IAMRECVR.

PROVIDING AUTHORIZATION UNDER TSO If you have TSO/E release 2 or higher, you can provide APF authorization for IAMRECVR to run under the IAM ISPF DIALOG by updating the appropriate TSO authorized program table, as discussed below.

IEAAPFXX The IAM load modules must be available to the IAM ISPF DIALOG via an APF authorized library. The IAM load module library must be either:

- (a) a linklist library that is APF authorized, or
- (b) allocated to the STEPLIB DD in the LOGON PROC. Additionally, the IAM load module library and any other libraries that may be concatenated under the STEPLIB DD MUST be listed in the IEAAPFxx member of SYS1.PARMLIB even if they are also in the LNKLSTxx member. Please note that you can NOT use the ISPLLIB DD for this purpose.

TSO AUTHORIZED PROGRAM TABLE Programs that are allowed to run with APF authorization under TSO must be listed in the TSO Authorized Program Table. As discussed below, you must add IAMRECVR to the appropriate table. Starting with TSO/E release 4, the list of authorized programs may be specified in member IKJT000 in SYS1.PARMLIB. The old method of changing the table within a load module in the Link Pack Area is still supported. The authorized program lists are documented in the IBM manual "TSO/E Customization".

90.41 APF AUTHORIZATION FOR THE IAM ISPF DIALOG

UPDATING IKJTSO00 If your installation uses the IKJTSO00 member, then add IAMRECVR to the AUTHTSF NAMES list. If not, then use the table below to determine the CSECT(s) and LOAD module to change. An IPL or PARMLIB command is required in order for the changes to IKJTSO00 to take effect. The CSECT name(s) and load module of the table that must be updated for various levels of MVS and TSO/E are as follows:

Operating System	TSO	Load Module	CSECT(s)
MVS/XA+	TSO/E R3+	IKJTABLS	IKJEFTE8 and IKJEFTAP
MVS/XA+	TSO/E R2.1	IKJTABLS	IKJEFTE8
MVS/XA+	TSO/E R2.0	IKJEFT02	IKJEFTE8
MVS/370	TSO/E R3+	IKJEFT02	IKJEFTAP
MVS/370	TSO/E R2+	IKJEFT02	IKJEFTE8

“+” means “or higher”

An IPL with CLPA is required for the updated TSO authorized program table to take effect.

The authorized program tables are documented in the IBM manuals “System Programming Library: TSO” and “System Programming Library: TSO/E User Exits and Modifications Volume 2”

INSTALLING IAM CICS MONITOR IAMXMON consists of a single program and a single mapset:

Program: IAMXMON language (Assembler)

Mapset: IAMXMAP

These modules must be copied from the IAM LOADLIB into a LOADLIB that is part of the DFHRPL concatenation in order to work properly. IAMXMON is a CICS application, not part of the IAM access method code. A transaction ID must be assigned to invoke the IAMXMON program, such as IAMX or IMON.

SECTION 91: GLOBAL OPTION CHANGE FACILITY

91.01 GLOBAL OPTION CHANGE FACILITY - OVERVIEW AND JCL

91.01 GLOBAL OPTION CHANGE FACILITY - OVERVIEW AND JCL

The Global Option Change Facility gives the user a simple method of modifying installation options for executing the IAM system. The installation options include defaults for DEFINE, special features in IAM, and other processing options. This facility supplies the user with the ability to PRINT present global option values, change values and RESET options to their original values as supplied on the installation tape. You should carefully review the options available for the IAM system.

In order for changes made to the IAM options to take effect, you must run IAMSTART on all systems to re-load the IAM option table. It is not necessary to deactivate IAM's VSAM interface before running IAMSTART. If IAMSTART detects that IAM is already active, it will refresh IAM's global option table.

BATCH EXECUTION

To execute IAMZAPOP as a BATCH job use the following JCL:

JOB CONTROL REQUIREMENTS

The JOB Statement is user specified and depends upon installation standards.

EXEC STATEMENT

Must specify the program name of the Global Option Change Facility - IAMZAPOP.

STEPLIB OR JOBLIB DD STATEMENT

If required, must specify the load module library in which IAM resides.

SYSPRINT DD STATEMENT

Specifies the output message data set. This is a required statement and usually is a SYSOUT data set.

SYSLIB DD STATEMENT

Must specify the load module library in which IAM resides.

COPYTODD DD STATEMENT

Specifies the load library containing the target option table for the COPY operation. The changed options in the SYSLIB option table will be copied to the target option table.

SYSIN DD STATEMENT

Specifies the control statement data set required for all functions. Usually an input stream or DD * data set.

TSO EXECUTION

The GLOBAL OPTION CHANGE FACILITY program (IAMZAPOP) can be executed under TSO. The program will prompt the user for the commands. 'END' will save the new options, if any, and terminate the program. The allocations required to execute IAMZAPOP in the TSO Foreground are as follows:

```
ALLOC F(SYSLIB) DA('iam.library') SHR
ALLOC F(SYSPRINT) DA(*)
ALLOC F(SYSIN) DA(*)
IAMZAPOP
```

----- or if the IAM library is not in LINKLIST-----

```
CALL 'iam.library(IAMZAPOP)'
```

91.02 GLOBAL OPTION CHANGE FACILITY - FORMAT AND OPTIONS

91.02 GLOBAL OPTION CHANGE FACILITY - FORMAT AND OPTIONS

AVAILABLE FUNCTIONS

The Global Option Change Facility has the following commands:

- AUDIT** – Lists the Global Options that have been changed from their distributed defaults. This function is only valid on Version 6.1 levels and above of the Global Options table.
The AUDIT command format is: **AUDIT**
- CANCEL** – Terminates IAMZAPOP without updating the Option Table.
The CANCEL command format is: **CANCEL**
- COPY** – Copies the changed Global Options from the Option table in the SYSLIB data set to the COPYTODD data set. COPYTODD is a required DD statement for this operation. This function is only valid on Version 6.1 levels and above of the Global Options table.
The COPY command format is: **COPY**
- END** – Terminates IAMZAPOP processing and rewrites the Option Table if any option was changed. This command is intended for TSO users.
The END command format is: **END**
- HELP** – The HELP command will print or display a menu of the IAMZAPOP options and related documentation.
The HELP command format is: **HELP ALL**
- PRINT** – If PRINT is specified, IAM will print or display the current values of the Global Options table.
The PRINT command format is: **PRINT**
- RESET** – If RESET is specified, IAM will reset the Global Options to the original values on the installation tape.
The RESET command format is: **RESET**
- ZAP** – Modify the Global Options. This command can be used to enable or disable specified processing options and set DEFINE defaults for the IAM system. The operands for this command are documented by function in the following sections.
The ZAP command format is:
ZAP operand=value, ...operand=value

91.03 GLOBAL OPTION CHANGE FACILITY - ZAP COMMAND

91.03 GLOBAL OPTION CHANGE FACILITY - ZAP COMMAND

This section discusses options that apply to most of the programs within the IAM system.

ZAP

BLKFACT=n	,IAMDDNAME=c...c	,RECFM=c
,BUFNO=nn	,INDEPENDENT=nnnnn	,RECTYPE=nnn
,BUFOPNO=nn	,INTEGRATED=nn	,RELEASE=nnnn
,BUFSP=nnnnnnnn	,KWDC=xxxxx	,ROUTECD=nn
,CORELIMIT=nnnnnn	,LIMITKEYS=nn	,SMF=xxx
,CRBUFOPT=c..c	,LINECNT=nnn	,SORTCORE=nnnnnnnn
,DATACOMPRESS=nnnnnnnn	,LOGDDNAME=c.....c	,SORTMSG=xx
,DATASPACE=nnnn	,MAXBUFNO=nn	,SORTPF=xxxx
,DESCRIPTCODE=nn	,MAXOVERFLOW=nnnnnn	,STORCLASS=c.....c
,DISABLE=(opt1,..optn)	,MAXREGION=nnn	,VSAMBLOCKF=n
,DSORG=cc	,MAXSECONDARY=(x,y)	,VSAMDDNAME=c.....c
,DYNCAT=c.c	,MINCOMPRESS=nn	,VSAMWTO=c.c
,ENABLE=(opt1,..optn)	,MULTIVOLUME=cccccccc	,WORKDDNAME=xxxxxxxxx
,FROMDDNAME=c...c	,NATWTO=c.c	,WORKPRIMARY=nnn
,GETRND=c.c	,OCOREO%=nnn	,WORKSECONDARY=nnn
,ESDSINTEGRATED=nn	,OCOREX%=nnn	,WORKUNIT=c.....c
,IAMCC=xxxxx	,PE=nnnnn	

OPERANDS Following are the operands for the ZAP command.

BLKFACT= Specifies the blocking factor to be used when creating an IAM file without IDCAMS DEFINE. BLKFACT= may specify from 1 to 15 inclusive.

NOTE: This keyword does not apply if the file is DEFINEd with IDCAMS.

The default is 2.

BUFNO= Specifies the number of BSAM buffers to be used when loading an IAM file. BUFNO may be any number from 2 to 80 but will be rounded down to a multiple of 2. This option only has meaning when CRBUFOPT=BSAM.

The default is 12.

91.03 GLOBAL OPTION CHANGE FACILITY - ZAP COMMAND

- BUFOPNO=** Specifies a default value for the initial number of buffers IAM will acquire when OPENING a file for processing. IAM will acquire a number of buffers equal to this value or to the number of blocks contained on one track whichever is greater. The initial number of buffers IAM will acquire for a file can be overridden at execution time using the IAM Override Statement keyword MINBUFNO=.
- IAM's Real Time Tuning starts with this initial number of buffers and during file processing dynamically adjusts the number of buffers actually used for a file up or down based on demand.
- You may specify a value from 1 to 32 buffers.
- Under NON-MVS/XA and NON-MVS/ESA systems BUFOPNO defaults to 1 buffer.
- The default is 4 under MVS/XA and MVS/ESA systems.
- BUFSP=** Specifies a default, in bytes, for the maximum amount of storage that IAM is to use for buffers when accessing a file. IAM divides this value by the file's block-size to determine the number of buffers that will fit.
- For example: When a file with 1/4 track blocking (11,476) is on a 3380 IAM can fit 22 buffers in 256K of storage.
- IAM compares the number of buffers requested in the BUFMAXNO= parameter to this value. When the number of buffers requested by BUFMAXNO= can not fit in the amount of storage specified by BUFSP, IAM internally adjusts BUFSP upwards to an amount of storage that will hold the number of buffers requested by BUFMAXNO.
- BUFSP can not be specified less than 65,536 bytes (64K).
- The default is 262,144 (256K).
- CORELIMIT=** Specifies the smallest size IAM file that will be considered for compression. Used by IAM when a file is loaded to determine if the file's index should be compressed. Any number from 0 to 999999, inclusive, may be specified.
- The default is 8000 bytes.
- CRBUFOPT=** Specifies the EXCP options to be used during create processing.
- CYL** – Acquire enough buffers for one full cylinder. Each I/O is for one half of a cylinder.
- MCYL** – Acquire enough buffers for two full cylinders. Each I/O is for a full cylinder.
- TRK** – Acquire enough buffers for two tracks. Each I/O if for one track.
- MTRK** – Acquire enough buffers for ten tracks. Each I/O is for five tracks.
- BSAM** – Disables the use of EXCP by create. The BSAM access method will be used and will use the number of buffers specified by BUFNO.
- DATACOMPRESS=**
DATAAC= Specifies the smallest size IAM file that will be considered for automatic data compression. This value is the minimum number of tracks that a DEFINE can specify as a file's primary allocation and still qualify for automatic data compression. Any number from 0 to 99999999, inclusive may be specified.
- The default is that records in an IAM file will not be compressed unless DATACOMPRESS=YES is specified on an IAM Override Control statement for the file.

91.03 GLOBAL OPTION CHANGE FACILITY - ZAP COMMAND

DATASPACE= Specifies the size, in megabytes, of the Data Space to be used for the temporary storage of the index to the IAM file that is being loaded. This parameter is only valid for MVS systems which support the use of Data Spaces. Valid values are from 0 to 2048. A value of 0 results in the use of a dynamically allocated temporary data set.

Default is 128 megabytes.

DESCRIPTCODE= Specifies the descriptor code(s) to be used when issuing Write-To-Operator or Write-To-Operator-With-Reply messages. Any number from 1 to 16 inclusive may be specified. Multiple descriptor codes can be entered if specified as **DESCRIPTCODE=(nn,....,nn)**.

The default is 0 (X'0000').

DISABLE= Specifies the option(s) coded for this operand shall be deactivated. See the **ENABLE** operand for the options and their implications.

DSORG= Specifies the DSORG to be used when creating an IAM file.

DA – Sets a DSORG of DA (direct access)

PS – Sets a DSORG of PS (physical sequential)

The default is PS.

DYNCAT= Specifies whether IAM should allow dynamic allocation to catalog the IAM file. This option is available to enhance IAM's support of POOLDASD.

YES – Let dynamic allocation catalog file

NO – IAM will catalog the file

The default is NO.

ENABLE= Specifies that the option(s) coded for this operand shall be activated.

BIM – Enables IAM support for the BIM product. Default value is disabled. Requires a restart VIF to take effect.

DYNAL – Permits the use of Dynamic Allocation in the support of requested IAM services.

ENHANCED – Specifies that IAM files will default to the Enhanced file format when they are defined. Default is that this option is disabled, which means that IAM files will default to the Compatible format.

EURODATE – Changes the format of date fields on IAMINFO and IAMPRINT reports to a European format. Default is disabled.

NOOWNID – Disables IAM's use of the ownerid in the catalog record that describes an IAM dataset.

THIS OPTION MUST NOT BE SPECIFIED IF IAM VERSION 6.0 OR EARLIER WILL ACCESS FILES CREATED WITH THIS OPTION ENABLED.

NOREUSE – Causes IAM to honor the **NOREUSE** option if coded in the IDCAMS define statements.

VAM™ – Enables IAM's enhanced **VAM™** support.

91.03 GLOBAL OPTION CHANGE FACILITY - ZAP COMMAND

- ESDSINTEGRATED=** Sets the INTEGRATED OVERFLOW percent for ESDS files. Reserves space in each block for record length changes.
The default is 0.
- FROMDDNAME=** Specifies the ddname to be used by IAM utility programs when reading an IAM file.
The default is DISKIN.
- GETRND=** Specifies if getmains should be rounded to 4K.
The default is NO.
- IAMCC=** Specifies the return code an IAM utility abend will issue when an error is detected, nn may be any number from 8 to 255, inclusive or the word ABEND. ABEND will cause a U0888 abend to be issued.
The default is 12.
- IAMDDNAME=** Specifies the ddname to be used by IAM utility programs when creating or updating an IAM file.
The default is DISKOUT.
- INDEPENDENT=** Specifies the default number of records of Independent Overflow to be reserved when creating an IAM file and the user does not specify an Independent Overflow value. You may specify a value from 0 to 32,767, inclusive.
NOTE: This keyword does not apply if the file is DEFINEd with IDCAMS, see Section 10.04 FREESPACE(CI% CA%) for IDCAMS defaults. Under VM/CMS VMINDEP is applied not INDEPENDENT.
The default is 500.
- INTEGRATED=** Specifies the default percentage of each Prime Data block that is to be reserved when creating an IAM file for record insertions during later processing. You may specify any number from 0 to 99, inclusive.
NOTE: This keyword does not apply if the file is DEFINEd with IDCAMS.
The default is 10.
- KWDCC=** Specifies the return code IAM will issue when a command/operand related error is found when editing a command. You may specify any number from 8 to 255, inclusive or the word ABEND. ABEND will cause a U0502 abend to be issued.
The default is 12.
- LIMITKEYS=** Specifies the number of keys taken in a set when creating an IAM file with a compressed index. You may specify any number from 3 to 64, inclusive.
The default is 32.
- LINECNT=** Specifies the maximum number of lines to be printed on any report page. You may specify any number from 10 to 200, inclusive.
The default is 58.

91.03 GLOBAL OPTION CHANGE FACILITY - ZAP COMMAND

LOGDDNAME= Specifies the ddname of the sequential log file to be used by IAM utility programs when applying the log ([See Section 45 on IAMRECVR](#)).

The default is LOG.

MAXBUFNO= Specifies the default maximum number of buffers IAM is permitted to acquire during file processing. IAM's Real Time Tuning will dynamically adjust the number of buffers used for the file as demand warrants up to this maximum.

Under MVS/XA and MVS/ESA MAXBUFNO will cause IAM to internally adjust the BUFSP parameter value upwards to a size which can accommodate the number of buffers specified by MAXBUFNO. [See Section 02.50 IAM's Real Time Tuning for details](#).

Under NON-MVS/XA and NON-MVS/ESA the same as for XA systems with the exception that for online systems (programs using asynchronous processing) MAXBUFNO is always made equal to MINBUFNO for the file.

You may specify a value from 1 to 32 buffers.

The default is 5.

MAXOVERFLOW= Sets the maximum amount of overflow that will be allocated based on the CA% freespace value provided in the IDCAMS define.

The default is 50,000.

MAXREGION= Specifies the default maximum value, in megabytes, that IAM will dynamically adjust the above the line Region value to. To disable the feature set value to 0.

Default is 128.

MAXSECONDARY=(x,y) Default multiplication factors for IAM Dynamic Secondary Space Adjustment feature, where x is the value for file loads and y is the value during file updates. To disable this feature, set both values to 0.

Default values are (10,5).

MINCOMPRESS= Specifies the minimum acceptable percentage of compression achieved when creating an IAM file to determine if the file qualifies for a compressed index. 'nn' may be any number from 8 to 40, inclusive.

The default is 10.

MULTIVOLUME= PRIMARY/SECONDARY
Specifies which space allocation value IAM will use when it appears that a data set will take the next extent on the next volume.

Default is PRIMARY.

NATWTO= Under the NATIVE or ISAM interface, specifies the action to be taken if the IAM file is not available.

YES – An IAMW02 message is generated and the operator must reply Retry, Wait, or Cancel.

NO – The open is failed.

The default is YES.

91.03 GLOBAL OPTION CHANGE FACILITY - ZAP COMMAND

- OCOREO%=** Specifies the Independent Overflow Core Open percent used when an IAM file is OPENed for update processing. 'nnn' may be any number from 1 to 100, inclusive.
The default is 10.
- OCOREX%=** Specifies the Independent Overflow Core Full percent used when the memory used for overflow record processing of an IAM file opened for update processing is exhausted. 'nnn' may be any number from 10 to 100, inclusive.
The default is 10.
- PE=** Specifies the number of blocks of Prime Extension area to be reserved when creating an IAM file. 'nnnn' may be any number from 0 to 32767, inclusive.
The default is 3.
- RECFM=** Specifies the internal record format IAM is to use for VSAM defines of files that appear to have fixed length records.
F – Define the file as fixed.
V – Define the file as variable.
The default is F.
- RECTYPE=** Specifies the 'user' record type to be written if SMF recording is requested for IAM files, nnn may be a number from 128 to 255, inclusive. There is no default value.
This is a required field for SMF recording to be requested. Member IAMUSMF in the IAM Installation Control Library is a DSECT of the IAM SMF user record format.
- RELEASE=** Specifies the default value for automatic release.
YES – Unused disk space in an IAM file is to be released if Secondary allocation value is specified.
NO – Unused disk space is not to be released.
The default is YES.
- ROUTE CODE=** Specifies the route code(s) to be used when issuing WTO of WTOR messages. Any number from 1 to 16, inclusive, may be specified. Multiple route codes can be entered as ROUTE CODE=(nn,...,nn).
The default is 11 (X'0020').

91.03 GLOBAL OPTION CHANGE FACILITY - ZAP COMMAND

- SMF=** Specifies whether IAM is to write an IAM SMF user record when an IAM file is CLOSEd. The IAM SMF user record will contain all of the same information displayed in an IAM INFO Run Time Statistics Report. Member IAMUSMF in the IAM Installation Control Library is a DSECT of the IAM SMF user record's format.
- YES** – If a RECTYPE value is specified in IAM's GLOBAL OPTION TABLE an IAM SMF user record will be written whenever an IAM file is CLOSEd. This can be overridden by the SMF=NO parameter on an IAM Override Control Statement.
- NO** – IAM will not write IAM SMF user records. This can be overridden by the SMF=YES parameter on an IAM Override Control Statement, if a RECTYPE value was specified.
- The default is NO, but can be overridden at execution time if a RECTYPE value was specified.
- SORTCORE=** Specifies the amount of storage 'SORT' is to use. You may specify any number from 10000 to 8000000, inclusive.
- The default is 100000.
- STORCLASS=** Specifies the SMS storage class to be used if SMS is active and a storage class was not specified on the define request, or the installation ACS routines did not assign an SMS storage class. If this option is blanks, no SMS storage group will be assigned. To reset this option to blanks, specify STORCLASS=' '.
- The default is blanks.
- SORTMSG=** Specifies the message option to be used by the program 'SORT' if external sorting is required.
- AC** – all messages to the console
AP – all messages to the printer (SYSOUT)
CC – critical messages to the console
CP – critical messages to the printer
NO – no messages to be produced
PC – critical messages to both console and printer
- The default is CC.
- SORTPFX=** Specifies the ddname prefix to be used by 'SORT'. If the string specified is less than 4 characters, a dollar sign (\$) fill character will be used.
- The default is SORT.

91.03 GLOBAL OPTION CHANGE FACILITY - ZAP COMMAND

- VSAMBLOCKF=** Specifies the default blocking factor (number of blocks per track) to be used when an IDCAMS DEFINE does not specify CISIZE. Any value from 1 to 15 inclusive may be specified.
The default value is 4 (four blocks per track).
- VSAMDDNAME=** Specifies the default ddname IAMRECVR is to use when creating a VSAM file.
The default is VSAMOUT.
- VSAMWTO=** Specifies the action to be taken under the VSAM interface, when an IAM file is not available.
YES – An IAMW02 message is generated and the operator must reply Retry, Wait, or Cancel.
NO – The Open is failed.
The default is NO.
- WORKDDNAME=** Specifies the ddname of the work file used during an IAM load function.
The default is IAMWKDD.
- WORKPRIMARY=** Specifies the primary allocation in tracks of the work file dynamically allocated during an IAM load function. Any value from 1 to 65535, inclusive, may be specified.
The default is 30 (tracks).
- WORKSECONDARY=** Specifies the secondary allocation in tracks for the work file dynamically allocated during an IAM load function. Any value from 1 to 65535, inclusive, may be specified.
The default is 30 (tracks).
- WORKUNIT=** Specifies the unit name to be used when dynamically allocating the work file used during an IAM load function.
The default is SYSDA.

91.04 IAMZAPOP JCL EXAMPLES

91.04 IAMZAPOP JCL EXAMPLES

The following examples illustrate some of the ways of executing the GLOBAL OPTION CHANGE FACILITY.

EXAMPLE 1 The user wishes to display the present Global values.

```
//PRINT      EXEC  PGM=IAMZAPOP
//SYSPRINT   DD    SYSOUT=A
//SYSLIB     DD    DSN=iam.library,DISP=SHR
//SYSIN      DD    *
          PRINT
```

EXAMPLE 2 The user enables Automatic Data Compression for files 150 tracks or larger. The ZAP command will modify the IAM Option Table.

```
//ZAP        EXEC  PGM=IAMZAPOP
//SYSPRINT   DD    SYSOUT=A
//SYSLIB     DD    DSN=iam.library,DISP=SHR
//SYSIN      DD    *
          ZAP      DATACOMPRESS=150
```

EXAMPLE 3 Reset all of the Global Option values to their original values supplied on installation tape.

```
//RESET      EXEC  PGM=IAMZAPOP
//SYSPRINT   DD    SYSOUT=A
//SYSLIB     DD    DSN=iam.library,DISP=SHR
//SYSIN      DD    *
          RESET
```

EXAMPLE 4 Enable the IAM SMF recording option so an SMF user record is written every time an IAM file is CLOSED. IAM's SMF user record type is to be type 201.

```
//SETSMF     EXEC  PGM=IAMZAPOP
//SYSPRINT   DD    SYSOUT=A
//SYSLIB     DD    DSN=iam.library,DISP=SHR
//SYSIN      DD    *
          ZAP      RECTYPE=201,SMF=YES
```

EXAMPLE 5 COPY all of the changed Global Option values to a new IAM load library.

```
//COPY       EXEC  PGM=IAMZAPOP
//SYSPRINT   DD    SYSOUT=A
//SYSLIB     DD    DSN=iam.library,DISP=SHR
//COPYTODD   DD    DSN=iam.new.library,DISP=SHR
//SYSIN      DD    *
          COPY
```

EXAMPLE 6 AUDIT the Global Option table to list any values that have been changed from their distributed defaults.

```
//AUDIT      EXEC  PGM=IAMZAPOP
//SYSPRINT   DD    SYSOUT=A
//SYSLIB     DD    DSN=iam.library,DISP=SHR
//SYSIN      DD    *
          AUDIT
```

THIS PAGE WAS INTENTIONALLY LEFT BLANK

INDEX

Section numbers are followed by page numbers in parenthesis.

A

&ALLDD 10.03(4), 10.04(8,9), 10.05(15)
ACB 02.10(47), 02.50(53), 10.04(7), 60.20(2), 90.04(10)
ACS 00.02(14), 02.70(56), 10.01(1), 10.04(11), 90.08(18)
ACCESS 00.02(14,17), 10.01(1), 10.03(4), 11.01(1), 11.02(5),
11.03(9,13,14), 11.04(16), 90.08(20)
ALLOCATION 00.02(11,14), 02.70(58), 03.20(64), 10.01(2), 10.04(8-9),
10.05(14), 11.02(4,6-8), 11.04(16)
ALTER 10.02(3), 11.03(14), 60.20(4)
AMAXTASK 01.40(15)
APF 10.03(4), 90.40(28), 90.41(29)
AUDIT command 91.02(32)
AUTHORIZATION 90.40(28), 90.41(29)
Automatic Space Release 01.30(8)
Automatic Space Reserve 01.30(8)

B

BACKUPCOMPRESSED 01.50(20), 01.60(28)
Backup While Open 01.50(24)
BIM 01.50(24), 01.60(27), 91.03(35)
BLKFACT 91.03(31-33)
BLOCKSIZE 10.04(7,11), 10.07(20), 10.10(26), 10.11(27), 11.02(3),
11.03(11), 42.04(5), 44.03(3), 60.20(4), 62.30(2), 63.20(4), 90.10(24)
BOOKMANAGER 00.02.00(12)
BUFFERS 01.40(13), 02.50(53-54), 10.03(4), 10.04(7,11), 10.07(21),
11.01(1,3,5), 11.02(5), 11.03(11,12), 11.04(16), 12.04(24,25),
41.10(18), 41.11(25), 42.04(5), 50.01(3), 60.10(1), 60.20(2,4),
61.60(10-12), 62.20(1), 62.30(2), 63.10(1), 63.20(5,6), 90.05(14),
90.10(24-25)
BUFND 01.40(13), 42.04(5), 60.20(2)
BUFNO 00.02(12), 10.11(28), 11.02(2,3), 12.04(25), 41.10(18), 41.11(25),
63.20(6), 91.03(33)
BUFOPNO 91.03(34)
BUFSP 91.03(34)
BWO 01.50(24)

C

Cached Control Units 01.50(20)
CANCEL command 91.02(32)
CA% 00.02(11), 01.30(8), 02.20(49), 02.40(51), 02.90(60), 10.04(8,9),
10.09(25), 11.02(6), 12.04(23), 40.08(18), 42.04(6)
CI% 02.20(49), 02.90(60), 10.04(8), 10.07(20), 10.09(25), 11.02(5),
12.04(23), 42.04(6), 50.01(3), 61.10(1), 61.35(5), 62.10(1)
CICS 00.02(12,17), 01.40(13), 01.40(15), 02.10(47), 02.50(52-53),
10.01(1), 11.01(1), 11.03(11), 11.04(16), 43.01(1), 44.01(1),
44.05(7), 62.20(1), 90.01(1), 90.01(3-4), 90.04(9), 90.04(11),
90.05(13-14), 90.06(16), 90.10(23-25), 90.30(26), 90.41(30)
AMAXTASK 01.40(15)
Data Table Facility 01.40(15)
FCT 01.40(15)
IAM Execution Monitor 01.40(17)

CI freespace 01.20(4)
CISIZE 02.50(53), 10.04(7,11), 11.02(3), 40.01(2), 40.08(18,19),
90.05(14), 90.08(20-21)
COBOL 00.02(17), 02.10(47), 10.01(1), 11.04(16), 90.04(10), 90.06(15)
COMPATIBLE 01.20(6-7)
COMPRESSION 00.02(17), 02.30(50), 10.01(2), 10.04(6,8,10),
10.07(21), 11.02(4,5,8), 12.04(24), 40.01(2), 40.08(19),
42.01(1,2), 60.10(1), 60.20(2), 61.35(6), 61.60(11), 62.10(1),
90.01(1), 90.06(15), 90.10(24)
CONCURRENT I/O 01.40(14)
CONTROLINTERVAL 42.04(5), 60.20(2), 62.30(2)
CONVERTING 00.02(14), 10.01(1), 10.04(6), 10.05(13), 10.09(24,25),
42.01(2), 61.10(1), 90.10(25)
COPY command 91.02(32)
COPYING 00.02(13), 10.04(10), 10.09(24), 61.40(7), 90.01(4)
CORELIMIT 11.02(2,3), 91.03(34)
CPK 01.30(8), 10.11(27)
CRBUFOPT 02.50(54), 11.02(2,3), 61.40(8), 16.60(10), 91.03(34)
CREATE Override 01.30(8)

D

DATACLAS 02.70(56-57), 90.08(20-22)
DATACOMPRESS 00.04(21), 02.30(50), 10.04(6-8), 10.07(20,21),
11.02(2,4), 11.04(15), 61.60(10), 90.01(4), 91.03(34)
DATA COMPRESSION 02.30(50), 10.01(2), 10.04(6-10), 11.02(4,8),
40.01(2), 40.08(19), 42.01(1), 60.20(2), 61.35(5), 61.60(10),
62.10(1), 90.01(1), 90.06(15), 90.10(24)
DATACLAS 10.01(1), 10.04(5,7), 10.05(15), 11.02(4,8)
Data Integrity 01.20(5)
DATASPACE 00.02(13), 11.02(2,4), 91.03(35)
Data Table Facility 01.40(15)
Date and Time Stamps 01.50(25)
DECOMPRESS 01.50(22)
DEFINE 00.02(14,15), 02.10(47), 02.40(51), 02.50(53), 02.70(56-58),
10.01(1,2), 10.02(3), 10.04(5-12), 10.05(13-15), 10.07(19-21),
10.09(25), 10.11(27,28), 11.01(1), 11.02(3-8), 11.03(9,11,14),
11.04(15-16), 12.04(23), 41.01(1), 41.11(22,24), 42.01(2),
42.04(6), 43.04(5), 45.10(14), 50.01(2,3,4,6), 60.01(1), 61.10(2),
61.35(5,6), 61.60(10), 62.10(1), 62.30(2), 90.01(1), 90.04(9-10),
90.05(13-14), 90.06(15), 90.08(22), 90.10(25), 90.30(26)
DELETE 00.02(14), 02.10(47), 10.01(1,3), 10.03(4), 10.04(10,11),
10.05(15), 10.06(16), 10.09(25), 10.10(26), 10.11(27,28),
11.02(7), 12.03(19,20), 12.04(22,25), 40.09(20), 41.01(1),
41.10(21), 43.04(4), 43.05(8), 44.04(4), 50.01(2), 61.20(3),
61.30(4), 63.10(1,2), 90.04(9), 90.05(13-14), 90.06(15)
Mass 01.50(25)
DESCRIPTCODE 91.03(35)
DFDSS 10.05(15)
DFHSM 00.02(11)
DFSMS 1.3 01.50(24)
DIAGNOSE 00.02(15), 10.04(5), 40.08(18), 41.01(1), 41.11(23),
90.41(29)
DISABLE 91.03(35)

REVISED MARCH, 1996

INDEX

DSORG= 01.60(27), 91.03(35)

Dynamic

Secondary Space Adjustment 01.30(9)

Dynamic File Expansion 01.30(8)

Dynamic Region Size Adjustment 01.40(15)

DYNAL 91.03(35)

DYNAMIC TABLING 02.60(55), 12.04(24), 41.11(25), 60.20(3), 61.60(12)

DYNCAT 91.03(35)

DYNCORE 01.40(15), 02.60(55), 10.03(4), 11.03(9,10), 11.04(16),
12.04(24), 60.20(4), 61.60(12)

E

ECKD 01.50(20)

Empty data blocks 01.50(25)

ENABLE= 91.03(35)

BIM 91.03(35)

DYNAL 91.03(35)

ENHANCED 01.60(27), 91.03(35)

EURODATE 01.60(27), 91.03(35)

NOOWNID 91.03(35)

NOREUSE 91.03(35)

VAM™ 91.03(35)

END command 91.02(32)

ENHANCED 01.20(6-7), 01.60(27), 91.03(35)

Enhanced File Structure 01.20(3)

ENQ 00.02(12)

ENQUEUE 11.03(14)

EOV 01.20(4)

ESCON 01.50(20), 60.20(2)

ESDS 00.02(11,12,14), 02.10(47), 02.30(50), 10.04(8-10), 11.02(6),
11.03(14), 42.04(5), 50.01(2,3,4,6), 60.10(1), 60.20(2,3), 61.10(1),
61.20(3), 61.30(4), 61.35(5), 61.40(7,8), 61.60(10-12), 62.10(1),
62.30(2), 63.10(1,2), 63.20(4), 64.10(1,2), 90.30(26)

ESDSINTEGRATED 00.02(11), 10.04(8), 62.10(1), 91.03(36)

EURODATE 01.50(25), 01.60(27), 91.03(35)

EXLST 90.04(10)

EXPORT 10.02(3), 10.11(27,28), 90.04(9)

EXTENTS 00.02(11), 01.20(4), 01.30(8), 10.07(19,20), 10.11(27),
40.01(2), 40.08(19), 50.01(5), 60.10(1), 60.20(3), 61.30(4), 61.35(5)

Extended File Area 01.20(4)

Extended Overflow 01.20(4-5), 01.50.1(23)

Extended PE 01.50.1(23)

F

Factor

Secondary Space 01.30(9)

FCT 01.40(15)

FDR 10.05(15), 10.09(24), 10.11(27)

FDRREORG 01.50(21), 61.40(7,8)

File integrity 01.20(5)

File Structure 01.20(3)

FIXED 10.04(7,8,10), 10.07(19,20), 11.02(2,4,5,8), 12.03(20),
12.04(21,23), 42.04(6), 60.10(1)

Fixed length records 01.20(3)

FREESPACE 00.02(11,12), 01.20(4), 02.20(49), 02.90(60), 10.04(8),
10.05(14,15), 11.02(5,6), 11.04(15,16), 12.04(23), 42.01(1),
42.04(6), 50.01(4,5), 60.20(4), 61.10(1), 61.35(5), 62.10(1),
62.20(1), 90.08(21)

FROMDDNAME 91.03(36)

G

GETRND 91.03(36)

GIGABYTES 11.02(6), 60.10(1), 60.20(3), 61.10(2), 61.60(11)

GLOBAL OPTIONS 02.100(61), 01.20(6-7), 01.60(27), 10.04(5,10),
10.07(21), 11.02(3,4,7,8), 40.01(2), 50.01(8), 60.20(2), 61.40(1),
61.60(10), 90.01(3-4), 90.06(15), 90.10(24-25)

BIM 01.50(24)

COMPATIBLE 01.20(6-7)

ENHANCED 01.20(6-7)

MAXREGION 01.40(16)

MAXSECONDARY 01.30(11)

MULTIVOLUME 01.30(12)

Global Option Change Facility

AUDIT command 91.02(32)

CANCEL command 91.02(32)

COPY command 91.02(32)

END command 91.02(32)

HELP command 91.02(32)

PRINT command 91.02(32)

RESET command 91.02(32)

ZAP command 91.02(32)

BLKFACT= 91.03(33)

BUFNO= 91.03(33)

BUFOPNO= 91.03(34)

BUFSP= 91.03(34)

CORELIMIT= 91.03(34)

CRBUFOPT= 91.03(34)

DATACOMPRESS= 91.03(34)

DATASPACE= 91.03(35)

DESCRIPTCODE= 91.03(35)

DISABLE= 91.03(35)

DSORG= 91.03(35)

DYNCAT= 91.03(35)

ENABLE= 91.03(35)

ESDSINTEGRATED= 91.03(36)

FROMDDNAME= 91.03(36)

GETRND= 91.03(36)

IAMCC= 91.03(36)

IAMDDNAME= 91.03(36)

INDEPENDENT= 91.03(36)

INTEGRATED= 91.03(36)

KWDCC= 91.03(36)

LIMITKEYS= 91.03(36)

LINECNT= 91.03(36)

LOGDDNAME= 91.03(37)

MAXBUFNO= 91.03(37)

MAXOVERFLOW= 91.03(37)

MAXREGION= 91.03(37)

MAXSECONDARY=(x,y) 91.03(37)

MINCOMPRESS= 91.03(37)

MULTIVOLUME= 91.03(37)

NATWTO= 91.03(37)

OCOREO%= 91.03(38)

OCOREX%= 91.03(38)

PE= 91.03(38)

RECFM= 91.03(38)

RECTYPE= 91.03(38)

RELEASE= 91.03(38)

ROUTECD= 91.03(38)

SMF= 91.03(39)

SORTCORE= 91.03(39)

SORTMSG= 91.03(39)

SORTPFX= 91.03(39)

STORCLASS= 91.03(39)

VSAMBLOCKF= 91.03(40)

VSAMDDNAME= 91.03(40)

VSAMWTO= 91.03(40)

WORKDDNAME= 91.03(40)

WORKPRIMARY= 91.03(40)

WORKSECONDARY= 91.03(40)

WORKUNIT= 91.03(40)

Guaranteed Space 01.30(11)

INDEX

H

HELP command 91.02(32)

High values

Key 01.50.1(23)

HIPERBATCH 62.30(2)

I

IAM

Backup Compressed Feature 01.50(20)
 CICS Execution Monitor 01.40(17)
 Dynamic Region Size Adjustment Feature 01.40(15)
 Dyncore Feature 01.40(15)
 Enhancements for Version 6.3 01.10(1)
 Listcat Output 01.60(29)
 Overrides 01.60(28)
 Sequential I/O Capabilities 01.50(20)
 Trace Output 01.60(43)

IAMCC 91.03(36)

IAMDDNAME 91.03(36)

IAMINFO 00.02(13,15), 01.60(34), 02.100(61), 10.03(4), 10.09(25),
 10.10(26), 11.02(4), 11.03(10), 11.04(16), 12.01(17), 12.02(18),
 12.03(19), 12.04(21), 41.01(1), 41.10(17,18,20,21), 41.11(25),
 43.05(8), 61.40(8), 61.50(9), 61.60(12), 63.10(1), 90.05(14)

IAMNOLIC 61.50(9)

IAMOVRID 00.02(12), 02.70(57), 10.03(4), 10.05(15), 11.01(1), 11.03(13),
 11.04(15,16), 43.02(2), 43.05(7,8), 44.02(2), 44.05(7), 60.20(4),
 61.10(1,2), 61.35(5), 61.40(8), 61.50(9), 90.10(24), 90.30(26)

IAMPRINT 00.02(13), 01.60(29), 02.10(47), 02.20(48), 10.03(4),
 10.04(6,8), 10.07(17-20), 10.10(26), 10.11(28), 45.01(1), 50.01(7),
 61.30(4), 61.50(9), 90.04(9), 90.05(13-14)

IAMRECVR 00.02(14,15), 02.100(61), 45.01(1), 45.02(2), 45.06(6),
 45.09(11), 45.10(13,14), 50.01(6,7), 61.35(6), 90.41(29-30)
 DECOMPRESS 01.50(22)

IAMSIMVS 40.02(3), 40.09(21), 42.01(1,2), 42.02(3), 42.03(4), 42.04(5,6),
 42.05(7), 42.06(8,9)

IAMSMF 00.02(15), 01.60(34), 10.03(4), 12.01(17), 40.02(3), 40.03(5),
 40.09(22), 41.01(1), 41.02(2), 41.03(3), 41.04(5,7), 41.05(9),
 41.06(10), 41.07(11), 41.08(13,14), 41.09(15,16), 41.10(17-21),
 41.11(22-25), 61.20(4), 61.50(9), 90.05(14), 90.06(15)

IAMSMFVS 00.02(15), 12.01(17), 40.01(1,2), 40.02(3), 40.03(4,5),
 40.04(6), 40.05(7,10), 40.06(11), 40.07(12), 40.08(16-19),
 40.09(20-22), 42.02(3), 42.06(9), 60.20(3), 60.20(3,4), 62.10(1),
 90.05(13-14)

IAMSTART 90.01(4), 90.04(11-12), 90.30(27)

IAMXMON 01.40(17)

IAMZAPOP 01.20(7), 01.60(27)
 JCL Examples 91.04(41)

IAM360 01.60(37)

IAM361 01.60(37)

IAM362 01.60(37)

IAM363 01.60(38)

IAM365 01.60(40)

IAM366 01.60(41)

IAM367 01.60(42)

IAM368 01.60(42)

IAM372 01.60(39)

IDCAMS 00.02(17), 02.10(47), 02.20(48), 02.40(51), 02.70(56-58),
 10.01(1,2), 10.02(3), 10.03(4), 10.04(5-12), 10.05(13-15),
 10.06(16),
 10.07(17-22), 10.08(23), 10.09(24,25), 10.10(26), 10.11(27,28),
 11.01(1,3,5-8), 11.02(8), 11.03(14), 11.04(15,16), 12.04(23),
 41.11(22), 42.04(6), 45.01(1), 42.02(2), 45.10(14), 50.01(7),
 60.10(1), 61.10(1,2), 61.20(3), 61.30(4), 61.35(5,6), 61.40(7,8),
 61.50(9), 61.60(10), 62.10(1), 62.30(2), 90.04(9), 90.04(12),
 90.05(13-14), 90.06(15), 90.08(21), 90.10(25)
 REPRO 01.50(21)

IEFUSI 01.40(15)

IMPORT 10.02(3), 10.11(27,28), 90.04(9)

INDEPENDENT= 91.03(36)

INDEPENDENT OVERFLOW 00.02(15), 01.20(3), 02.40(51), 02.80(59),
 02.90(60), 03.30(65), 10.01(2), 10.04(6,8,9,11), 10.05(15),
 10.07(19-22), 10.10(26), 11.01(1), 11.02(6), 11.03(12), 11.04(15),
 12.04(23,25), 40.08(19), 41.10(18), 41.11(25), 43.04(5), 45.08(9),
 60.20(3)

INDEX 01.20(4)

INDEX COMPRESSION 02.20(48), 11.02(4)

INSTALLATION 00.02(15,16), 00.05(22), 10.01(1,5), 11.02(7),
 12.01(17), 40.09(20), 90.01(1-4), 90.02(5-6), 90.03(7-8),
 90.04(12), 90.06(15), 90.07(17-19), 90.40(28), 90.41(30)

INTEGRATED= 91.03(36)

INTEGRATED OVERFLOW 01.20(4), 10.04(8), 10.07(21), 12.04(23),
 61.10(1,2), 61.30(4), 61.35(5)

Integrity 01.20(5)

IRB 01.40(15)

ISAM 10.01(2), 11.03(9,13), 12.04(22), 40.05(8), 63.20(4), 90.10(25)

ISPF 00.02(14), 50.01(1-9), 90.07(18), 90.40(28), 90.41(29)

I/O Read Ahead 01.50(20)

J,K

JRNAD 01.50(24)

KEYS 10.04(10), 10.15(13-15), 10.07(21), 11.02(5-7), 11.03(12),
 11.04(15), 41.10(18), 43.04(5), 45.08(8,9), 50.01(8)

KWDCC 91.03(36)

L

LIMITKEYS 91.03(36)

LINECNT 91.03(36)

LISTCAT 00.02(14,17), 01.60(29), 02.10(47), 02.20(48), 10.02(3),
 10.03(4), 10.04(5,6), 10.07(17-20), 10.10(26), 45.01(1),
 50.01(5,7), 60.20(3), 61.30(4), 61.35(5), 61.50(9), 90.04(9),
 90.05(13-14)

LOGDDNAME 91.03(37)

Locate mode 01.20(5)

LSR 00.02(11), 01.40(15), 02.50(53), 11.03(9,11), 60.10(1), 60.20(2),
 62.30(2), 63.20(6), 90.04(11), 90.06(16), 90.10(23-24), 90.30(26)

LSTAR 01.30(8)

INDEX

M

MAXBUFNO 00.02(12), 01.40(13), 10.04(7), 10.07(21), 11.01(1),
11.02(2,5), 11.03(9,11), 11.04(16), 12.04(25), 50.01(3), 60.20(2),
61.40(8), 61.60(11,12), 91.03(37)

Maximum record size 01.50(26)

MAXOVERFLOW 00.02(11), 91.03(37)

MAXREGION= 01.40(16), 01.60(27-28), 91.03(37)

MAXSECONDARY 01.30(10-11), 01.60(27-28)

MAXSECONDARY=(x,y) 91.03(37)

MAXTASK 01.40(15)

MGMTCLAS 10.04(10), 10.05(15), 90.08(21-22)

MINBUFNO 10.04(7), 10.07(21), 11.01(1), 11.02(2,5), 11.03(9,11),
50.01(3), 60.20(2), 61.40(8), 61.60(11,12), 63.20(6)

MINCOMPRESS 91.03(37)

MODCB 62.30(2), 90.04(10)

MODEL 00.02(14), 10.04(6,10,12), 10.09(24), 11.03(13), 11.04(15,16),
50.01(2,6), 90.08(21)

MONITOR 01.40(17), 11.03(14), 40.01(1), 43.05(8), 44.01(1), 44.02(2),
44.03(3), 44.04(4,5,6), 44.05(7), 44.06(8,9,10), 90.41(30)

Multiple ACB 01.50(24)

MULTIVOLUME 01.30(11-12), 01.60(27-28), 91.03(37)
Parallel Sequential Read 01.50(21)

MVS/ESA 5.2 01.50(24)

N

NATWTO 91.03(37)

NOOWNID 91.03(35)

NOREUSE 00.02(11), 10.04(10), 91.03(35)

O

OCOREO% 10.07(21), 11.02(2,6,9,12), 91.03(37)

OCOREX% 00.02(11), 10.07(21), 11.02(2,6,9,12), 91.03(37)

Online systems 01.40(13), 01.40(15)

OPTCD=LOC 01.20(5)

OPTIONS 00.02(11,13), 10.03(4), 10.04(5,10,11), 10.07(17,19,21),
11.01(1), 11.02(2,4,7,8), 11.03(13,14), 12.03(19), 12.04(22),
40.01(2), 40.03(4,5), 41.04(5-7), 43.01(1), 43.04(4), 44.04(4),
50.01(2,6-8), 60.20(2,3), 61.10(1), 61.30(4), 61.35(5), 61.60(10),
63.10(1,2), 63.20(4,5)

OSCORE 01.40(17)

Overallocated 01.30(8)

OVERFLOW 00.02(11,12,15), 10.01(2), 10.03(4), 10.04(5,8,9,11),
10.05(15), 10.10(26), 11.02(1), 11.02(2,6), 11.03(12), 11.04(15),
12.03(19), 12.04(23,25), 40.01(1), 40.07(12), 40.08(19),
41.10(18,20), 41.11(25), 42.04(5,6), 42.06(8), 43.04(5), 45.08(9),
45.09(10,12), 45.10(13), 50.01(1,3,8), 60.10(1), 60.20(3), 61.10(1,2),
61.30(4), 61.35(5), 61.60(11), 62.10(1), 62.20(1), 63.10(1), 63.20(6)
Extended 01.50.1(23)
Independent 01.50.01(23)

Override 01.60(28)
BACKUPCOMPRESSED 01.50(20)
COMPATIBLE 01.20(6)
ENHANCED 01.20(6)
MAXBUFNO 01.40(13)
MAXREGION 01.40(16)
MAXSECONDARY 01.30(10)
MULTIVOLUME 01.30(11)
RELEASE 01.30(8)
secondary space adjustment factor 01.30(10)
TRACEDD 01.60(43)
TRACEREQ 01.60(43)

OVERRIDES 00.02(12), 10.01(2), 10.04(7), 11.01(1), 11.03(14),
11.04(15), 43.01(1), 44.02(2), 50.01(3), 60.20(2,4), 61.40(8),
61.50(9), 61.60(10)

OWNER 02.10(47), 10.01(1), 10.04(5), 10.05(14,15), 10.07(17,18),
10.09(25), 50.01(4,5), 60.10(1), 61.10(1,2), 90.05(14), 90.06(15),
90.08(21)

P

Paging 01.40(14)

Parallel Sequential Reads 01.50(21)

PE 00.02(15), 01.20(3-4), 01.50.1(23), 11.02(2,6), 12.03(19,20),
12.04(23,25), 42.04(5,6), 45.08(8,9), 50.01(3,8), 91.03(38)

PRIMARY 00.02(15), 10.01(2), 10.04(5,6,8,9), 10.09(25), 11.02(6),
43.02(2), 50.01(2-5,7)

Prime Data Area 01.20(4)

PRIME EXTENSION 01.20(3-4), 03.30(66), 10.01(2), 10.07(19,21),
11.02(6), 12.04(23,25), 40.08(19), 41.10(18), 45.08(9), 50.01(3,8),
60.20(3), 63.20(6)

PRINT 02.10(47), 10.02(3), 10.08(23), 11.02(7), 11.03(14), 11.04(16),
12.01(17), 12.03(19), 40.03(5), 41.01(1), 41.04(6-8), 41.07(11-12),
41.09(15,16), 43.02(2), 43.04(5), 43.05(7), 44.04(6), 44.05(7),
45.01(1), 45.03(3), 45.08(8,9), 45.09(10,11), 50.01(6-8), 63.10(1),
64.10(2), 90.01(3-4), 90.04(9), 90.41(29), 91.02(32)

Pseudo Logical Record Length Support 01.50(26)

PSEUDOLRECL 01.50(26), 01.60(28)

PSEUDORBA 00.02(12), 11.02(6), 50.01(3), 60.10(1), 60.20(3),
61.10(2), 61.35(5), 61.60(11)

R

RANDOM 02.50(52), 11.03(10,13), 12.03(20), 12.04(21,24,25),
41.11(25), 43.04(4), 44.04(4), 60.20(2,3), 63.10(2), 63.20(3,5,6)

RBA 00.02(12), 50.01(5), 60.10(1), 60.20(3), 61.30(4), 61.35(6),
61.60(11), 63.10(1), 90.04(10), 90.05(14)

Read Ahead 01.50(20)

Real storage 01.40(14)

REAL TIME TUNING 00.02(11), 01.40(13), 02.50(52), 02.50(53),
10.04(7), 11.02(5), 12.04(24), 60.10(1), 60.20(2), 63.20(5,6),
90.04(10), 90.10(23)

RECATALOG 00.02(17), 10.04(10), 10.05(15), 90.10(25)

RECFM 91.03(38)

RECORDSIZE 10.04(10), 10.015(13-15), 10.09(25), 11.02(5),
11.04(15,16), 42.04(6), 50.01(3-5), 61.10(1,2), 90.04(10),
90.06(15)

INDEX

RECOVER 45.01(1), 45.02(2), 45.03(3), 45.06(6), 45.09(10-12),
45.10(13,14), 61.35(6), 61.40(8)

RECTYPE 91.03(38)

Region 01.40(15)

RELEASE 01.30(8), 02.40(51), 10.01(2), 10.03(4), 10.04(6), 10.07(19,20),
11.02(2,7), 11.04(15), 12.04(22), 42.01(2), 42.06(9), 50.01(3),
60.20(2), 61.30(4), 61.35(5), 61.60(11), 63.20(4), 64.10(1), 91.03(38)

RESET command 91.02(32)

REUSE 00.02(11), 10.04(10), 10.05(13), 10.09(25), 10.10(26), 50.01(4,5),
61.40(8)

REORGANIZATION 01.50(21), 10.01(2), 10.10(26), 12.04(23), 61.40(7,8),
62.20(1)

REORGANIZE 10.10(26)

REPRO 00.02(17), 02.10(47), 10.01(1,2), 10.02(3), 10.014(10), 10.09(24),
10.10(26), 10.11(27), 11.02(3), 61.40(7,8), 90.04(9), 90.05(13-14)

RESERVE 01.30(8)

ROUTE CODE 91.03(38)

RPL 00.02(16), 02.10(47), 12.04(25), 41.11(25), 62.30(2), 63.20(6),
64.10(1), 90.04(10), 90.06(15)

S

Secondary space 01.30(9)

Secondary space adjustment 01.30(9)
Rules 01.30(9)

SECONDARY EXTENTS 60.10(1), 60.20(3), 61.35(5)

SEQUENTIAL 02.50(52), 10.01(1), 10.04(9), 10.09(25), 10.10(26),
11.03(13), 12.04(25), 40.01(1,2), 40.02(3), 40.07(13), 41.08(13),
41.09(15), 41.11(24), 43.04(4), 44.06(10), 45.01(1), 60.20(2),
61.35(5), 61.40(7,8), 61.50(9), 62.10(1), 63.10(2), 63.20(3,5,6)

Sequential I/O 01.50(20)

Sequential update 01.50(20)

SHARE OPTIONS 10.04(11), 10.07(19,21), 11.03(12,14), 12.03(19),
12.04(22), 50.01(3), 61.30(4), 61.35(5), 63.10(1,2), 63.20(4),
90.08(20)

SHOWCAT 02.10(47), 90.04(10)

SHOWCB 02.10(47), 62.30(2), 90.04(10)

Single record load 01.50.1(23)

SMF= 91.03(39)

SMF RECORDS 01.60(34), 02.100(61), 10.03(4), 11.03(13), 12.01(17,18),
40.01(1,2), 40.03(5), 40.05(8,10), 40.07(13-15), 40.08(16),
40.09(21), 41.01(1), 41.03(3,4), 41.04(8), 41.07(11,12), 41.08(14),
41.09(15,16), 41.10(17,19,20), 41.11(24,25), 60.20(4), 61.50(9)

SMS 00.02(12,14), 02.70(56-58), 10.04(5,7,10,11), 11.02(4,8), 50.01(1,3),
61.10(1), 90.08(20-22), 90.10(25)

SORTCORE 91.03(39)

SORTMSG 91.03(39)

SORTPFX 91.03(39)

SPACE SAVINGS 11.02(4), 42.01(1), 42.02(3), 42.03(4), 42.06(8)

Space release 01.30(8)

Space reserve 01.30(8)

SPANNED 10.04(11), 50.01(4,5), 62.30(2)

STATISTICS 00.02(14,16), 02.50(54), 02.60(55), 10.07(18,22),
11.02(7), 11.03(13), 11.04(16), 12.01(17), 12.02(18),
12.03(19,20), 12.04(21,24,25), 40.01(1), 40.05(7), 40.07(13),
40.09(22), 41.01(1), 43.04(5), 45.09(10-12), 50.01(1,5,9),
60.20(4), 61.35(6), 63.00(1,2), 63.20(3,5,6), 90.05(14)

STORAGE 00.02(2,14), 10.01(1), 10.04(5,11), 10.07(19,21,22),
11.01(1,4,5,6,8), 11.03(10,12), 12.03(19,20), 12.04(24,25),
40.05(9), 40.07(15), 41.04(8), 41.08(14), 41.10(19), 45.04(4),
50.01(3-5), 60.20(3,4), 61.30(4), 61.35(5,6), 61.60(12), 62.10(1),
62.20(1), 63.10(1), 63.20(5,6)

STORCLAS 00.02(14), 02.70(56-57), 10.01(1), 10.04(5,11), 10.05(13),
61.10(1), 90.08(21-22), 91.03(39)

STRNO 02.50(53), 90.10(23-24)

T

TABLING 02.60(55), 11.01(1), 11.03(10), 41.11(25), 60.20(3,4),
61.60(11), 63.20(5)

TESTCB 02.10(47), 62.30(2), 90.04(10)

TESTING 00.02(12,15), 00.05(22), 50.01(2), 61.60(11), 90.01(2-3),
90.04(11-12), 90.05(13-14), 90.30(27)

TOD stamps 01.50(25)

TRACE 01.60(43), 11.02(7), 11.03(14), 11.04(16), 43.01(1), 43.02(2),
43.03(3), 43.04(4,5,6), 43.05(7,8), 44.02(2), 44.03(3), 44.04(6),
44.05(7), 64.10(1,2)

TRACEDD 01.60(43)

TRACEREQ 01.60(43)

TRBFR 01.60(43), 01.60(45)

TREXCP 01.60(43), 01.60(46)

TRIOE 01.60(43), 01.60(45)

TRIOS 01.60(43-44)

TRXTND 01.60(43), 01.60(46)

TUNING 00.02(11), 02.50(52), 02.50(54), 10.04(7), 11.01(1), 11.02(5),
12.04(24,25), 44.01(1), 60.10(1), 60.20(2), 63.20(5,6)

U

UCB 01.50(24)

UPAD 01.40(15)

UPDATENQ 00.02(12), 11.03(14)

V

VAM™ 91.03(35)

VARIABLE 00.03(20), 10.04(7,10), 10.05(14), 10.07(19,20), 11.01(1),
11.02(2,4,8), 11.03(10), 12.03(19), 12.04(21,25), 41.10(19),
42.04(5,6), 42.06(8), 45.09(10,12), 50.01(3), 60.10(1), 61.10(1),
63.10(1,2), 63.20(3), 90.06(15)

Variable length records 01.20(3)

VERIFY 00.02(14), 02.10(47), 10.02(3), 63.10(2), 90.04(9)

INDEX

VIF 00.02(12,15-17), 01(2,8,9,16), 02.10(47), 10.01(1,2), 10.02(3),
10.04(5,10), 10.05(15), 10.06(16), 10.07(17), 10.08(23), 10.09(24),
10.11(27), 11.03(13), 12.04(22), 40.01(1), 42.01(2), 45.03(3),
50.01(9), 61.20(3), 61.30(4), 90.01(1-4), 90.02(6), 90.04(9-12),
90.05(13), 90.06(15-16), 90.08(20), 90.08(22), 90.10(23), 90.10(25),
90.30(26-27)

VIFSTART 00.05(22), 02.10(47), 90.01(1-2), 90.01(4), 90.04(11-12),
90.06(15), 90.30(27)

Virtual storage 01.40(15)

VSAM

JRNAD Exit 01.50(24)

UPAD Exit 01.40(15)

VSAMBLOCKF 91.03(40)

VSAMDDNAME 91.03(40)

VSAMWTO 91.03(40)

VSAM INTERFACE 00.02(15,17), 02.10(47), 10.01(1,2), 10.02(3),
10.03(4), 10.04(5), 10.06(16), 10.07(17), 10.08(23), 10.09(24),
42.01(2), 45.03(3), 45.09(11), 50.01(9), 90.01(1), 90.01(4), 90.02(6),
90.04(9), 90.04(11-12), 90.05(13), 90.06(15), 90.08(20), 90.10(23),
90.10(25)

W,X,Y,Z

WORKDDNAME 91.03(40)

WORKPRIMARY 91.03(40)

WORKSECONDARY 91.03(40)

WORKUNIT 91.03(40)

X37 abend 01.20(4)

ZAP command 91.02(32)

Numerics

3990-3 01.50(20)

3990-6 01.50(20)