

IBM Responses to Requirements Submitted by the SHARE MVSS Project

Barbara McDonald
IBM DFSMS Product Management
bawhite@us.ibm.com

March 13, 2012
Session 10944



Disclaimer

The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.

Agenda


- IBM Responses to SHARE Requirements
 - Requirements Waiting for Responses
 - Requirements Open for Discussion
 - Existing requirements with updated responses
 - Requirements Available in R13

SHARE Requirements Summary




- Waiting for Response

Requirement #	Title	Status
SSMVSS11001	Force System Determined Blocksize (SDB)	Waiting for response (Atlanta) Note: Cannot find in FITS – may need to be resubmitted. Open for discussion (Anaheim 2011)



SSMVSS11001

- **Force System Determined Blocksize (SDB)**
 - **Open for discussion; waiting for response**
- **Description**
 - The way the Force System Determined Blocksize option of the SMS Data Class as introduced in z/OS Release 10, when set to Y, creates the belief that SDB will be forced for a given data set. This is true during the creation of the data set only. Since most user's create and then use (open) the data set, the Force System Determined Blocksize option does not work when an application program provides a blocksize value in the DCB when issuing the Open macro.
- **Benefit:**
 - Since most files are allocated and opened within a step, this will provide a SDB as requested in the dataclas regardless of the DCB parm
- **Solution:**
 - This REQUEST is being submitted as a requirement to DFSMSdfp to provide the function to actually force SBD for data sets within a Data Class for the life of the data set regardless of the presence of any coded blksize value in the DCB JCL parm or program.




SHARE Requirements Summary



- Open for Discussion

Requirement #	Title	Status
SSMVSS11003	DFSMS: Change How System Managed Buffering Handles Empty Datasets	More discussion needed (Atlanta) Note: Cannot find in FITS – may need to be resubmitted. Open for discussion (Anaheim 2011)
SSMVSS10002	Excessive/invalid HSM RECALL commands generated by IKJEFT01	More discussion needed (Atlanta) Returned for user clarification AK – Acknowledged (Anaheim 2011)
SSMVSS11005	Increase the maximum VSAM Control Interval size (CISZ)	More discussion needed (Atlanta) Returned for user clarification Open for discussion (Orlando 2011)
SSMVSS11013	Enhance DFSMSdss To Process An Offline DASD Volume	Open for discussion (Atlanta 2012)
SSMVSS11014	ICKDSF: Automated Control Statement Generation for 1,000s of Vols	Open for discussion (Atlanta 2012)



SSMVSS11003

▪ DFSMS: Change How System Managed Buffering Handles Empty Datasets

• Open for discussion


- VSAM is really WAD, however, its not the best design! There is a difference in the way NSR and LSR handle SEQ processing, LSR requires a POINT and NSR does not. COBOL is not issuing a POINT since its a NSR user. When SMB is used to convert it to use LSR under the covers, a POINT it now required.

▪ Description

- Currently, if a program opens a VSAM data set with an ACB that specifies MACRF=DIR (which is what COBOL will specify if ACCESS IS RANDOM is coded on the File Definition) and the HURBA is zero, the open will fail if System Managed Buffering (SMB) is invoked. SMB selects ACCBIAS=DO, and that does not work for an empty data set. SMB should check the HURBA, and if it is zero ACCBIAS=DW should be selected instead.

▪ Benefit:


- The application can open a data set that has a HURBA=0 if 1) the program is changed to specify MACRF=SEQ (or any combination other than MACRF=DIR) in Assembler or ACCESS IS DYNAMIC in COBOL, or 2) codes ACCBIAS=DW in the JCL on the DD statement for the data set. The first option is actually the correct fix, but in either case requires the application to make a change to a program or job that possibly has not changed in a number of years. If there are a lot of programs coded this way, the number of changes could be substantial. Making this change in System Managed Buffering would fix this coding deficiency without any adverse affects, and it would allow customers to continue rolling out SMB without fear of causing job failures.



SSMVSS10002




- **Excessive/invalid HSM RECALL commands generated by IKJEFT01**
 - **Response: Returned for User Clarification**
 - TSO/E can not prescreen commands before passing the command string to the command processor. However, it sounds like the customer simply wants HSM to screen the command before sending it to the coupling facility? If so this requirement is for DFSMSHsm.
- **Description:**
 - When IKJEFT01 passes HSEND RECALL commands to HSM using a dataset pattern, the catalog should be checked to ensure that HSM does not attempt to recall datasets that are already on L0. HSM currently issues a recall regardless of catalog status. Thousands of invalid recalls can flood the system, overwhelming the CRQ and its coupling facility structure.
- **Benefit:**
 - A dramatic decrease in overhead can be realized. Fewer MWEs in HSM, fewer HSM activity records will be written. More importantly, HSM recall users will not experience recall delays when excessive CRQ utilization forces HSM recalls to fall back to LOCAL mode.
- **Solution:**
 - IKJEFT01 should do a catalog check before passing the HSEND RECALL request to HSM. Any recalls for files already on L0 should be dropped before enlisting HSM for the recall.




SSMVSS11005

- **Increase the maximum VSAM Control Interval size (CISZ)**
 - **Response: Returned for User Clarification**
 - VSAM Development will need more information to determine if the requirement is feasible.
 - Possible to support for LDS only?
- **Description**
 - CISZ(s) are currently limited to 32K, which limits the record size to 32-7 bytes. Increase the CISZs to 64K, 128K, 256K, 512K, 1024k or 1Meg. The 1Meg size would work well in conjunction with the z/OS Large page feature.
- **Benefit:**
 - Removing the 32k Limit on Application development will enable newer applications to be developed on z/OS, that would otherwise have to be done on "open" platforms where the 32K limit does not exist. To keep new application development on the z Platform.
- **Solution:**
 - z/OS does support greater than 32K I/O requests, but even if VSAM can't use them, physical records size could be used instead.
 - VSAM uses a Physical Record size, so a 64K CISZ, could be 4 physical records of 16K as an example.



SSMVSS11013

- **Enhance DFSMSdss To Process An Offline DASD Volume**
 - **Open for discussion (Atlanta 2012)**
- **Description**
 - There are DASD volumes that need to be accessed by DFSMSdss that can not be brought online, such as duplicate VOLSER. These volumes need to be processed by functions such as DUMP, for backups, and COPY, for data promotion.
- **Benefit:**
 - To backup a DASD volume from a SYSPLEX that does not have tape drives.
 - To backup a DASD volume from a SYSPLEX with tape media that is stored offsite.
 - To copy a DASD volumes from a development SYSPLEX to promote data.
- **Impact:**
 - As the wall between production and development continues to get stronger, it becomes more difficult to provide work arounds.



SSMVSS11014

- **ICKDSF: Automated Control Statement Generation for 1,000s of Vols**
 - **Open for discussion (Atlanta 2012)**
- **Description**
 - A complementary tool, utility, or modification to ICKDSF is required to automate generation of its control statements. Such function shall allow for specification and incorporation of various formatting capabilities (e.g. creation of an indexed or non-indexed VTOC and OWNER) for various UCB ranges and consistently named volume serial numbers when executing ICKDSF with PARM=NOREPLY for very large numbers (in the thousands) of storage volumes.
- **Benefit:**
 - Given the historical context when ICKDSF was first written (1978), addressable numbers of storage volumes using high performance replication technology in the last 2 decades have transcended the function originally envisioned of ICKDSF - i.e. initializing and maintaining limited numbers of physically mounted storage volumes, typically on a volume-by-volume basis.
 - As an example, a business continuity requirement recently necessitated initializing a series of 4,000 discontinuous UCB addresses with consistently named volume serial numbers of the form "BCxxxx", where "xxxx" was the UCB address. In total, creating the ICKDSF INIT control statements and running the batch jobs to initialize these volumes consumed more than 3 man-hours of tedious, error prone, and time-consuming editing to complete.
 - Unquestionably, software automation of control statement generation is warranted. In light of this limitation, such lack of utilitarian function is unquestionably and financially wasteful and unproductive - an impediment that is laboriously time consuming and which may delay recovery of one or more systems by many hours.
 - In addition, this much needed function within ICKDSF will complement other IBM DFSMS storage product offerings such as FlashCopy Manager, PPRC Manager, and other strategic IBM Business Continuity automation tools such as Geographically Dispersed Parallel Sysplex(GDPS) and z/OS TotalStorage Productivity Center (TPC) for Replication.
- **Discussion:**
 - Given the ease of coding this within REXX, is it worth IBM's time and resources to provide this functionality?



SHARE Requirements Summary



- Existing requirements with updated responses

Requirement #	Title	Status
SSMVSE00002	Request New Installation Option for GDG Management - PURGE	RC – Recognized (Atlanta 2012) Open for discussion (Orlando 2011)
SSMVSS11006	Package TDMF and LDMF with DFSMS	RJ – Rejected (Atlanta 2012) Open for discussion (Orlando 2011)
SSMVSS11007	ACS WRITE statement during RECALL/RECOVER	RC – Recognized (Atlanta 2012) Open for discussion (Orlando 2011)
SSMVSS11008	Update Reuse Capacity Periodically	AC – Accepted (Atlanta 2012) Open for discussion (Orlando 2011)
SSMVSS11009	Allow MAXEXTENTS in HSM to be used with Multi-Volume Non-VSAM	RC – Recognized (Atlanta 2012) Open for discussion (Orlando 2011)
SSMVSS11010	Remove line limit for PDSE members	AC – Accepted (Atlanta 2012) Open for discussion (Orlando 2011)
SSSTOR11002	Single utility for COPY/MOVE/RENAME/DELETE any dataset type	RJ – Rejected (Atlanta 2012) Open for discussion (Orlando 2011)
SSMVSE00006	NON-SMS HFS: Provide Same Flexibility as Other UNCATLG System DS	RJ – Rejected (Atlanta 2012) Open for discussion (Orlando 2011)


SSMVSE00002

- **Request New Installation Option for GDG Management - PURGE**
 - **Response: RC – Recognized (Atlanta 2012)**
- **Description**
 - With the advent of SMS managed storage, the SCRATCH and EMPTY options for GDGs are no longer sufficient to clean up old (rolled off) GDSs which were allocated with an expiration date. GDGs whose expiration date exceeds the rollover time (G0001V00 -> G9999V00), fail when attempting to rollover to G0001V00 again.
 - Prior to migrating to SMS managed storage, the rolled off generations were uncataloged and were deleted by home grown processes to delete uncataloged datasets even though their expiration date has not been reached. Now in an SMS environment, these jobs cause production work to fail at rollover (G9999v00 -> G0001V00). This occurs because an "unexpired" G0001V00 of a GDG exists on SMS managed storage (DASD or TAPE) when an attempt is made to allocate a new (+1) generation. SCRATCH/EMPTY options do not address this situation. To correct this situation, an IDCAMS job must be run to DELETE the dataset with the PURGE option.
- **Benefit:**
 - OW42558 would need to be re-worked to add this option for HSM migrated datasets.
- **Solution:**
 - The best solution would be an "installation configurable" parameter in SYS1.PARMLIB to request PURGE for GDSs which have rolled off. If not set, the system would continue to function as it now does.
- **Discussion**
 - I tried to open an APAR for this problem and was told by IBM that 'it is unlikely to happen so we won't accept the APAR'. So far, as far as I know, it has only happened to my installation once since converting to SMS, but I'm sure it will happen more. Posted by: Walt Sapp of State of California - DTS (HWC)




SSMVSS11006

- **Package TDMF and LDMF with DFSMS**
 - **Response: RJ – Rejected (Atlanta 2012)**
 - TDMF/LDMF are currently service offerings provided by IBM. There are no plans to integrate these products into DFSMS.
- **Description**
 - IBM now offers software previously marketed separately as TDMF and LDMF. This may be purchased by a customer or used as part of a services engagement. This software is not installed with z/OS ServerPac but must be installed separately and perhaps in haste. DFSMS ships a number of components DFSMSdss, DFSMSshm, DFSMSrmm, that are included in the common code base but separately entitled and enabled using SYS1.PARMLIB(IFAPRD00).
- **Benefit:**
 - Provide a way for customers to easily keep a current copy of TDMF and LDMF on the system service updated and ready to purchase, trial, or use as part of a services engagement.
- **Solution:**
 - Ship TDMF and LDMF code base in DFSMS to all customers or add function to DFSMS with some alterations perhaps incorporated into DFSMSdss.




SSMVSS11007

- **ACS WRITE statement during RECALL/RECOVER**
 - **Response: RC – Recognized (Atlanta 2012)**
- **Description**
 - Content of ACS WRITE statement during ACSENVIR=RECALL goes to HSM Command log, instead of the user/address space recalling data set.
- **Benefit:**
 - IGD01008I, 09I and 10I may contain important information directed to user/address space recalling data sets.
- **Solution:**
 - setsys parm to allow storage admin to choose what to do with IGD010XXI message (user, hsmlog, both, none).




SSMVSS11008

- **Update Reuse Capacity Periodically**
 - **Response: AC – Accepted (Atlanta 2012)**
- **Description**
 - When DFSMSHsm is started the Tape Reuse Capacity is calculated. As tape usage changes, this number is not updated causing many tapes to not be selected for RECYCLE processing. When DFSMSHsm is cycled, a new Reuse Capacity is calculated causing RECYCLE to now select those tapes that should have been previously recycled. As DFSMSHsm may stay running for many months, this can lead to thousands of tapes to be RECYCLED causing a peak in tape activity.
- **Benefit:**
 - Eliminate this peak tape activity following a DFSMSHsm restart.
- **Solution:**
 - Create a SETSYS parm to cause Tape Reuse Capacity to be periodically recalculated. Such as daily, weekly, monthly.




SSMVSS11009

- **Allow MAXEXTENTS in HSM to be used with Multi-Volume Non-VSAM**
 - **Response: RC – Recognized (Atlanta 2012)**
- **Description**
 - Allow MAXEXTENTS parameter to be used with Multi-Volume Non-VSAM Data sets.
- **Benefit:**
 - Our shop has most of our Non-VSAM data sets defined as Multi-Volume, so the great function of MAXEXTENTS used in DFSMSHSM to cause reallocation of the Non-VSAM data sets after hitting a defined threshold, can Not be used.
- **Solution:**
 - Desired solution: Allow Multi-Volume Non-VSAM data sets use of the MAXEXTENTS parameter of DFSMSHSM.




SSMVSS11010

- **Remove line limit for PDSE members**
 - **Response: AC – Accepted (Atlanta 2012)**
- **Description**
 - Message IEC036I 002-A8 is issued if a PDSE member exceeds 15,728,639 lines. This limit does not exist for PDS datasets. This creates situations where it is impossible to create partitioned dataset. For example, PDS is limited to 65K tracks. Once that limit is exceeded, PDSE is selected, but the 002-A8 abend occurs on a member that was successfully stored in a PDS. PDSEs were supposed to remove the limitations of PDS's, not introduce new limitations.
- **Benefit:**
 - Users requiring PDSE members with more than 16M lines have to resort to changing their applications to support physical sequential datasets. That is an unacceptable burden.
- **Solution:**
 - Remove the limit.



SSSTOR11002

- **Single utility for COPY/MOVE/RENAME/DELETE any dataset type**
 - **Response: RJ – Rejected (Atlanta 2012)**
 - From an ease of use and usability perspective, the requirement is certainly valid, as it would be very beneficial from a user perspective to modernize the various utilities into one utility. However, from a business case perspective, it's not feasible in the next 3-5 years.
- **Description**
 - Customer should not need to know which utility to use when copying or renaming different type of dataset, i.e. IEBCOPY for PDS, IEBGENER/IDCAMs for SEQ or VSAM/PDSE/zFS etc.. DF/DSS is NOT an utility for general users. We would like to have an intuitive tool, especially one supported under batch and ISPF, for user to COPY/MOVE/RENAME/DELETE any dataset, one single interface, with no external parameters besides the Source and/or Target dataset name. Same tool as on lower platform via CUT/COPY and PASTE, RENAME and DELETE.
- **Benefit:**
 - Ease of use. Standard and similar methodology across all platforms to provide user an intuitive tool without awareness of dataset type. Most user are using off the shelf application products and don't know what the data depository file type is. Asking user to remember which utility to use for specific type of dataset does not project z/OS as a user friendly environment. It also increase our company's cost to train user under z/OS, making it a very undesirable platform.
- **Solution:**
 - Provide batch and ISPF interface to a single tool that allows user to copy/backup/cloning/rename/delete any type of dataset just by providing source and/or target dataset name. If target dataset name does not exist, allocate one with source dataset attributes. Also, would like to use ISPF OPTION 3, one entry panel to handle any dataset type. Also, please allow utility to handle multiple datasets UNLIKE the single dataset restriction for IEBGENER.



SSMVSE00006



- **NON-SMS HFS: Provide Same Flexibility as Other UNCATLG System DS**
 - **Response: RJ – Rejected (Atlanta 2012)**
 - HFS has been functionally stabilized for years. zFS is the strategic file system for z/OS UNIX System Services. zFS should be used in its place. Requirements for USS file systems should be directed to the zFS component.
 - Similar requirement opened on the MVS "Top 39" list: SSSHARE016559 - Duplicate Data Set Names On Different Volumes.
- **Description**
 - With the ability to have non-SMS managed HFSs, it should be possible to have "duplicate" HFSs in the same manner as an installation would have any System Target Data Set (i.e., duplicate names on different volumes, possibly cataloged using extended indirect volume serial numbers) and to utilize them simultaneously. Specifically, this includes:
 - Having the System recognize that same name HFSs on different volumes are unique and not have them prevented from being mounted simultaneously on the same or different System.
 - Being able to MOUNT an HFS Data Set with Volume and Unit specifications, bypassing the Catalog, so that a specific instance of a duplicate named HFS can be explicitly mounted (typically for service).
 - Allowing same named HFSs on different volumes to be mounted simultaneously at different mount points within the same file system.
 - Basically, an installation needs the ability to manipulate and manage System HFSs in the same manner as they do the other Target Data Sets that are built and maintained (by SMP/E) on the "logical" SYSRES.
- **Benefit:**
 - The ability of the (forever shrinking) Systems Programming Staffs to maintain OS/390's HFS components is a substantial impediment to migrating new releases and service in a timely manner, costing additional time and manpower that could be spent more productively.
- **Solution:**
 - Left to the developers.



SHARE Requirements Summary




- Requirements Available in R13

Requirement #	Title	Status
SSMVSS09005	DFHSM command to release DASD recalls	AV – Available (Orlando 2011) AC- Accepted (2009)

SSMVSS09005



- **DFHSM command to release DASD recalls**
 - **Response: AV – Available (Orlando 2011)**
- **Description:**
 - Once a HOLD RECALL command has been issued, there is no way to release DASD recalls without also releasing tape recalls. The requirement asks for a way to specifically release DASD recalls.
- **Benefit:**
 - If an operator experiences a major problem with a tape subsystem s/he often responds with HOLD RECALL to avoid recall failures. However, once this is done, there is no way to release DASD recalls without issuing a general RELEASE RECALL. This reintroduces the tape issues (which can sometimes cause problems beyond a simple recall failure). This can be a serious problem if the tape subsystem will be down for an extended period of time due to maintenance (which is also more likely than having a DASD subsystem down for maintenance). Work that would otherwise be able to run is held up despite the fact all its datasets are only migrated to DASD.
- **Impact:**
 - Work is unintentionally held up for no reason. Once the error is made, it can be extremely difficult to undo if hardware work has begun in the meantime.
- **Solution:**
 - Introduce a RELEASE RECALL(DASD) command.



Trademarks and Disclaimers



The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

CICS*	FICON*	Lotus*	System Storage
DB2*	FlashCopy*	MQSeries*	Tivoli*
DFSMS	GDDM*	Multiprise*	TotalStorage*
DFSMSdfp	GDPS*	OMEGAMON*	Virtualization Engine
DFSMSdss	geoManager*	OS/390*	VisualAge*
DFSMShsm	HiperSockets	Parallel Sysplex*	VM/ESA*
DFSMSrmm	HyperSwap	PR/SM	VSE/ESA
DFSORT	IBM*	QMF	VTAM*
DFSMS	IBM logo*	RACF*	WebSphere*
DS4000	ImagePlus*	Rational*	z/Architecture*
DS6000	IMS	RMF	z/OS*
DS8000	Intelligent Miner	System i	z/VM*
Enterprise Storage Server*	Language Environment*	System z	z/VSE
ESCON*		System z9	zSeries*
			zSeries Entry License Charge

The following are trademarks or registered trademarks of other companies:

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

NOTES:

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Information is provided "AS IS" without warranty of any kind.

Trademarks and Disclaimers (continued)



NOTES:

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven, ClusterProven or BladeCenter Interoperability Program products. Support for these third-party (non-IBM) products is provided by non-IBM Manufacturers.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.