


SHARE
Technology • Connections • Results

IBM Responses to Requirements Submitted by the SHARE MVSS Project

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

Tuesday, March 1, 2011
Session 9015




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 Open for Discussion
 - Requirements Waiting for Response
 - Existing requirements with updated responses
 - Requirements Available in R12
 - Re-Review of Existing Requirements


SHARE Requirements Summary



SHARE
Technology • Connections • Results

- Open for Discussion


Requirement #	Title	Status
<u>SSMVSS11001</u>	Force System Determined Blocksize (SDB)	Open for discussion (Anaheim 2011)
<u>SSMVSS11002</u>	Provide option on EDGHSKP to WAIT for datasets in use	Open for discussion (Anaheim 2011)
<u>SSMVSS11003</u>	DFSMS: Change How System Managed Buffering Handles Empty Datasets	Open for discussion (Anaheim 2011)
<u>SSMVSS11004</u>	DFSMS: DATACLAS Needs Additional SMB Information	Open for discussion (Anaheim 2011)



SHARE
in Anaheim
2011


SSMVSS11001

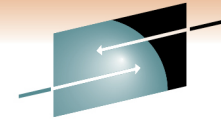
- **Force System Determined Blocksize (SDB)**
 - **Open for discussion**
- **Description**
 - The way the Force System Determined Blocksize option of the SMS Data Class as introduced in z/OS Release 10, when set to Y, creates the belief that SDB will be forced for a given data set. This is true during 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.



SSMVSS11002

- **Provide option on EDGHSKP to WAIT for datasets in use**
 - **Open for discussion**
- **Description**
 - Currently EDGHSKP generates an RC12 if any of the datasets specified to it (extract, message, activity, VRSELRPT, etc) are in use. EDGHSKP should provide an option to WAIT for the datasets to become available.
- **Benefit:**
 - Time lost in unnecessary reruns will be eliminated. In order for the restart to be successful, someone has to manually monitor DFRMM to determine when to rerun EDGHSKP. Regular batch jobs wait for datasets, EDGHSKP should also have that ability
- **Solution:**
 - Add WAIT as an option to the PARM string for EDGHSKP. The RMM started task should also WAIT if the datasets somehow become ENQ'd after EDGHSKP but before DFRMM gets control.






S H A R E

Technology • Connections • Results

SSMVSS11003


- **DFSMS: Change How System Managed Buffering Handles Empty Datasets**
 - **Open for discussion**
- **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 it 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.




SHARE
in Anaheim
2011

SSMVSS11004

- **DFSMS: DATACLAS Needs Additional SMB Information**
 - **Open for discussion**
- **Description**
 - The DATACLAS definition needs to be updated so that customers can specify ACCBIAS=SO, ACCBIAS=SW, ACCBIAS=DW, ACCBIAS=DO, RMODE31=ALL, RMODE31=BUFF, RMODE31=CB, and RMODE31=NONE.
- **Benefit:**
 - Having this parameters in the DATACLAS definition would prevent users from having to code them in their JCL when the defaults of ACCBIAS=SYSTEM and RMODE31=BUFF do not work for a particular job or set of jobs. It is much easier to make one DATACLAS change than potentially hundreds of JCL changes.




SHARE Requirements Summary



SHARE
Technology • Connections • Results

- Waiting for Response


Requirement #	Title	Status
<u>SSMVSS10002</u>	Excessive/invalid HSM RECALL commands generated by IKJEFT01	AK – Acknowledged (Anaheim 2011)
<u>SSMVSS10008</u>	DFSMSdss: Re-architect DSS to begin processing sooner	AK – Acknowledged (Anaheim 2011)




SHARE
in Anaheim
2011

SSMVSS10002

- **Excessive/invalid HSM RECALL commands generated by IKJEFT01**
 - **Response: Acknowledged**
- **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.





S H A R E

Technology • Connections • Results

SSMVSS10008

- **DFSMSdss: Re-architect DSS to begin processing sooner**

- **Response: Acknowledged**

- **Description:**

- When running the logical COMPRESS, COPY, and RELEASE commands in particular, DSS locates all of the data sets that meet the selection criteria and stores the necessary information to act on these data sets in storage. It then processes the list of data sets. If a job is acting on selected data sets in a large group of volumes or in a large storage group, DSS may run for hours before it actually does any work. The amount of storage used by the job can also a large amount.


- Here is an example.

```
RELEASE ADMINISTRATOR –  
  STORGRP(PS3) –  
  INCLUDE(**) –  
  EXCLUDE(SYS1.**)  
  BY(MGMTCLAS,EQ,M3NLB4CB) –  
  WAIT(0,0)
```

- This job ran for 26 hours before it released 1 track of unused space, and it used over 251MB of storage. In addition, it got lots of messages either indicating the data set was no longer on the volume (it had been deleted) or that it was migrated. Of course, DSS did a locate nearly 2 days earlier.


- **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.



SHARE
in Anaheim
2011

SHARE Requirements Summary




SHARE
Technology • Connections • Results

- Existing requirements with updated responses

Requirement #	Title	Status
<u>SSMVSS09001</u>	DFSMSHsm – Protect an old single backup version. MR0703094736 (Denver, Boston) MR0205111833 (Anaheim)	SG – Suggestion (Denver 2009) SG – Suggestion (Boston 2010) RC – Recognized (Anaheim 2011)
<u>SSMVSS10003</u>	HSM: Make TAPECOPY restartable after a tape takeaway	RC – Recognized (Anaheim 2011)
<u>SSMVSS10004</u>	HSM: RECYCLE should automatically restart after tape takeaway	RC – Recognized (Anaheim 2011)
<u>SSMVSS10005</u>	HSM: Provide the ability to have more than 2 copies of HSM tapes	RC – Recognized (Anaheim 2011)
<u>SSMVSS10006</u>	HSM: Segment HSM Migration and Backup Tapes	RC – Recognized (Anaheim 2011)
<u>SSMVSS10007</u>	HSM: Change RECYCLE to not require tapes in connected sets	RC – Recognized (Anaheim 2011)


SHARE Requirements Summary



SHARE
Technology • Connections • Results

- Existing requirements with updated responses


Requirement #	Title	Status
<u>SSMVSS07007</u>	FCQUERY command should show whether a volume is FlashCopy capable	RC – Recognized (2007) AC – Accepted (Anaheim 2011)
<u>SSMVSS053122</u>	SDM: FlashCopy Should Cut an SMF Record	RC – Recognized (2007) RC – Recognized (Anaheim 2011)
<u>SSMVSS09004</u>	SMS volume selection should also use physical disk response time	AK – Acknowledged (Denver 2009) RJ – Rejected (Anaheim 2011)
<u>SSMVSS07013</u>	Ability to list failed audit tapes that are eligible to resume	AC – Accepted (2008) AV – Available (Anaheim 2011)
<u>SSMVSS09008</u>	DFSMSdfp - Multi volume PDSE	AK – Acknowledged (Anaheim 2011)

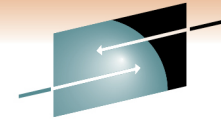


SHARE
in Anaheim
2011

SSMVSS09001

- **DFSMSHsm – Protect an old single backup version**
 - **Response: RC - Recognized**
- **Description**
 - When an old dataset (e.g. a source PDS) has not been used for a while and later on is brought back to life by a new project, then the existing old single backup version (BV) has too low protection. As the management class in this case need to have frequency 1 the old previously single BV disappears already after two days of updating because of its age.
The old single BV need to protected by the "Retain days extra backup versions" to secure expected logical recovery needs.
- **Benefit**
 - Decreased possibility for loss of data when you are able to recover to the end point of the previous dataset "life time" during the entire "Retain days extra backup versions". I.e. more often relief from the indescribable huge amount of work to recreate old code.







SHARE
Technology • Connections • Results

SSMVSS10003

- **HSM: Make TAPECOPY restartable after a tape takeaway**
 - **Response: Recognized**
- **Description:**
 - If tape take away takes a tape away from a TAPECOPY task, keep the target tapes mounted and continue copying the primary tape after the recall or ABARS backup finishes with the tape.
- **Benefit:**
 - If tape takeaway takes a tape away from a TAPECOPY task, the target alternate tape is dismounted and returned to scratch status. The TAPECOPY command has to be restarted from the beginning, even if the TAPECOPY was 90+% complete. This is a huge waste of resources (MIPs, tape drives, and time), and it elongates the time that a primary tape does not have a copy.
- **Solution:**
 - None listed




SHARE
in Anaheim
2011




SHARE
Technology • Connections • Results

SSMVSS10004

- **HSM: Make TAPECOPY restartable after a tape takeaway**
 - **Response: Recognized**
- **Description:**
 - When recall or ABARS takes a tape away from recall, once the recall or ABARS backup finishes the recycle should automatically resume where it left off.
- **Benefit:**
 - Very large tapes can take a long time to recycle. If tape takeaway causes a recycle task to terminate soon after it starts, it may be several hours before the RECYCLE command is issued again. That is time that could be used to copy more data off of the tape.
- **Solution:**
 - None listed




SHARE
in Anaheim
2011




SHARE
The Power of Connecting Results

SSMVSS10005

- **HSM: Provide the ability to have more than 2 copies of HSM tapes**
 - **Response: Recognized**
- **Description:**
 - HSM currently provides the ability to have 1 alternate copy. Since tapes are getting larger and can contain so much data, it may not be sufficient to just have 2 copies of the data. HSM should provide the ability to create more than 2 copies. The ability to create all of the copies concurrently or some copies at a later time using the TAPECOPY command should be provided.
- **Benefit:**
 - This functionality would provide additional protection for critical business data.
- **Solution:**
 - None listed




SSMVSS10006




SHARE
Technology • Connections • Results

- **HSM: Segment HSM Migration and Backup Tapes**
 - **Response: Recognized**
- **Description:**
 - Tapes today are very large. If HSM tapes are segmented, HSM could be changed to reuse segments that become empty. This would greatly reduce the need to recycle full tapes. In fact, RECYCLE should be changed to allow segments to be recycled.
- **Benefit:**
 - Recycling large tapes uses a lot of MIPs and tape drives. Segmenting tapes is viewed as a potential way to reduce the overhead associated with running HSM.
- **Solution:**
 - None listed




SSMVSS10007

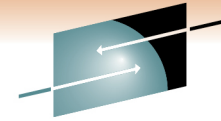


SHARE
Technology • Connections • Results

- **HSM: Change RECYCLE to not require tapes in connected sets**
 - **Response: Recognized**
- **Description:**
 - When a RECYCLE command is issued, whether for a single volume or several volumes, volumes that meet the recycle criteria but are part of a connected set will require that all of the volumes in the connected set also be recycled. This can cause lots of unnecessary movement of data and unnecessary use of MIPs, and this could also take a very long time (depending on how many volumes are in the connected set) to recycle the actual volume or volumes that need to be recycled.
 - The RECYCLE function should be changed to allow a volume in the middle of a connected set to be recycled without having to recycle all of the other volumes in the connected set. Ideally, HSM should move the data set that spans tapes first to the new tape, then the rest of the data sets, and finally the last data set on the tape that spans to another tape.
- **Benefit:**
 - For large connected sets on very large tapes, it could take days to recycle the data off a tape in the middle of the connected set. If the tape in the middle has a low percentage of utilization, it can cause HSM to use more tapes than necessary. If the tape in the middle has an I/O error on it, taking days greatly increases the probability of the tape being requested and getting the I/O errors again. This enhancement will not only solve these issues, but it will also reduce MIPs and tape drive usage, both expensive resources.
- **Solution:**
 - None listed




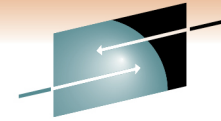
SHARE
in Anaheim
2011



SSMVSS07007

- **FCQUERY command should show whether a volume is FlashCopy capable**
 - **Response: Accepted**
 - Original response: Recognized
- **Description:**
 - Currently, the only ways to determine whether a source volume and a target volume are FlashCopy capable is to log into the box (using the ESS Specialist or the Storage Manager) or to use the ANTRQST API and the QFRVOLS query. The first method is a multi-step process that is cumbersome and takes several minutes to get an answer. The second requires that the customer write a program.
 - The FCQUERY command should be enhanced with a parameter that will tell the command to return the FlashCopy capability (FlashCopy V1 or V2) of the volumes provided (possibly allow for multiple volumes to be specified with a single command, but not required).
- **Benefit:**
 - This would provide significant reduction in time and complexity when added volumes to SMS storage groups that require FlashCopy capable volumes.
- **Solution:**
 - Something like "FCQUERY DEVN(4000) FCLEVEL" or "FCQUERY DEVN(4000) FCSTATUS" could provide V1, V2, or FNI (for Feature Not Installed) for the FCLEVEL or FCSTATUS field.






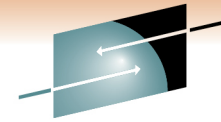
SHARE
Technology • Connections • Results

SSMVSS053122

- **SDM: FlashCopy Should Cut an SMF Record**
 - **Response: Recognized**
 - Original response: Recognized
- **Description:**
 - Whenever a FlashCopy session is established from a z/OS system, the System Data Mover (SDM) should cut an SMF record to indicate what time the request was made, what system the request was made on, what job and userid made the request, what volumes data was copied from, and what volumes data was copied to. If it is known, the amount of data being flashed should also be included.
- **Benefit:**
 - Currently, the storage administrators, performance analysts, and capacity planners have no way of knowing when FlashCopy is used, what volumes are affected (and what subsystem they are on), and who is using it. This information is vital when resolving performance problems, moving volumes from FlashCopy capable devices to devices that do not support FlashCopy, knowing when to increase backend capacity or order FlashCopy for more subsystems, and many other reasons.
- **Solution:**
 - SDM currently cuts a type 42 subtype 4 for Concurrent Copy sessions. Perhaps this same record could be used for FlashCopy with a different request type..




SHARE
in Anaheim
2011



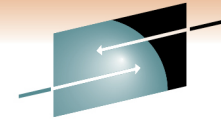
SHARE
Technology • Connections • Results

SSMVSS09004

- **SMS volume selection should also use physical disk response time**
 - **Response: RJ - Rejected**
 - *SMS selects from a list of logical volumes. Each logical volume may be mapped onto several different physical volumes each having a different response time. It would be illogical to attempt to come up with a response time for the logical volume.*
 - Original Response – Acknowledged
- **Description**
 - We need the ability to direct data set allocation based on the response time of the physical back end device(s).
- **Benefit**
 - An easier migration to newer technology and the ability to specify performance goals for data sets. The storage administrator could direct low access density (I/O per second per GB) data sets to cheaper but slower devices, and high access density data sets to faster disk drives. Currently, if different disk drives are used (with different response time capabilities), they must be defined in different SMS Storage Groups and the SMS ACS routines changed to manually direct allocation.
- **Solution**
 - The MSR (Millisecond response) time is used to distinguish between different device types. This could be expanded to also incorporate the response time of the back end disk drive. The RPM value of the hard disk drive is available to z/OS, and could be used to calculate a new MSR value. The latest MSR is for 2105. Unfortunately, this doesn't leave room for unlimited future expansion (since the values are already quite low), but would be adequate for the near future until a more robust solution is found.




SHARE
in Anaheim
2011



SHARE
Technology • Connections • Results

SSMVSS07013

- **Ability to list failed audit tapes that are eligible to resume**
 - **Response: AV – Available**
 - *IBM believes that the request described has been solved with a current product / service / policy / etc. Available on z/OS V1R10.*
 - Original Response – Accepted
- **Description**
 - Would like a command to list all failed HSM audits that are eligible to be resumed.
- **Benefit**
 - This will give the HSM user the ability to easily identify tapes that have started an audit process but did not complete.
- **Solution**
 - A command like LIST TTOC SELECT(RESUME) that lists the volumes eligible to resume would be great! Another option would be to handle like the FAILEDRECYCLE and FAILEDCREATE in the LIST TTOC - Perhaps a FAILEDAUDIT?




SHARE
in Anaheim
2011


SSMVSS09008



- **DFSMSdfp - Multi volume PDSE**
 - **Response: AK – Acknowledged**
 - *Requirement being evaluated and response will be Accepted or Recognized.*
- **Description:**
 - Support multi volume PDSE.
- **Background:**
 - It's quite simple to keep the storage group utilization high by simply utilizing multi volume datasets where the allocations span to overflow (QUINEX) volumes when needed. We have seen no performance problems despite we use multi volume datasets quite much. But, when there is a PDS or a PDSE they will suffer in a setup like this. They often can't grow as wished within its single volume. I.e. we have to setup specific low utilized storage groups for PDS and PDSE because they can't grow by the multi volume model in a highly utilized storage group.
- **Benefit:**
 - Higher storage group utilization. Less PDSE expand outages.
- **Solution:**
 - Possibility to grow over several volumes like multi volume supported VSAM and PS. Details is up to the developer.




SHARE Requirements Summary



SHARE
Technology • Connections • Results

- Requirements Available in R12

Requirement #	Title	Status
SSMVSS08006	DFSMS – ICF Catalog Size Relief	AV – Available (Boston 2010)
SSMVSS07014	Provide a supported REXX interface for the System Data Mover	AV – Available (Boston 2010)
SSMVSS09006	Implement NOLIST option for RMM CLIST commands in dialog	AV – Available (Boston 2010)
SSMVSS09007	Prevent RMM impact when issuing WTOR in parallel	AV – Available (Boston 2010)
SSMVSS053125	DFSMS: DCOLLECT Needs to Include More Fields	AV – Available (Boston 2010)
SSMVSS064933	Catalog Support of HDELETE for DELETE GDG FORCE	AV – Available (Boston 2010)
SSMVSS063069 SSMVSS063070	DFSMS – Method to Empty a PDS/PDSE	AV – Available (Boston 2010)



SHARE
in Anaheim
2011

SSMVSS08006


■ DFSMS – ICF Catalog Size Relief

- **Response: AV-Available**

- This function is available with z/OS 1.12 (GA 9/10).


■ Description

- An ICF CATALOG is essentially a special form of a non-extended format VSAM KSDS. As such, it is limited to 4 gigabytes in size. Provide relief from the 4 gigabyte size restriction.




SSMVSS07014

- **Provide a supported REXX interface for the System Data Mover**
 - **Response: AV-Available**
 - This function is available with z/OS 1.12 (GA 9/10).
- **Description**
 - Currently the ANTRQST and ANTQFRVL macros are the supported ways to call the System Data Mover (SDM) from a program. Both of these macros must be called from an Assembler program. IBM should provide a supported method to call the System Data Mover from a REXX program.




SSMVSS09006



SHARE
Technology • Connections • Results

- **Implement NOLIST option for RMM CLIST commands in dialog**
 - **Response: AV-Available**
 - This function is available with z/OS 1.12 (GA 9/10).
- **Description**
 - RMM added the CLIST option to the dialog, but without NOLIST, the SEARCH command often terminates for lack of TSO storage. RMM should provide the NOLIST option to suppress the hit list and allow the SEARCH to complete.
- **Benefit**
 - A previous SHARE requirement added the CLIST option to the RMM ISPF dialogs. But that function is essentially useless. Since running the CLIST option in the dialog populates Rexx stem variables, many searches will not complete in the dialog due to "Machine storage exhausted" situations. By implementing NOLIST and not populating the Rexx stem variables, the SEARCH command will likely complete.
- **Solution**
 - Add NOLIST to the RMM ISPF dialog. If NOLIST is specified in the dialog, do not return the hits to Rexx stem variables and do not display the hit list. A message that the SEARCH completed will suffice. This should allow the SEARCH command to complete without exceeding storage. The previous requirement to add the CLIST operand to the dialog panels did not concentrate on this aspect of the solution, which, in hindsight, was a mistake.




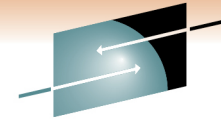
SHARE
in Anaheim
2011

SSMVSS09007



- **Prevent RMM impact when issuing WTOR in parallel**
 - **Response: AV-Available**
 - This function is available with z/OS 1.12 (GA 9/10).
- **Description**
 - When RMM is running in parallel with another tape management system and issues a WTOR, production tape mounts may wait for the RMM WTOR to be satisfied. This has caused production tape outages either due to the wait or an invalid operator reply. RMM should suppress these WTOR's by assuming replies which will allow production tape processing to continue unimpeded. This requirement is a scale back of a previous requirement to eliminate all warning mode WTOR's.
- **Benefit**
 - Both users and IBM will benefit by the elimination of production tape outages while RMM is running in parallel during a conversion.
- **Solution**
 - For example, the most common outage caused by RMM is EDG2103D when the journal fills up. RMM should simply assume a DISABLE reply, issue an informational reply, and continue to process. Halting production tape processing for RMM to offload the journal is unacceptable.





SHARE
Technology • Connections • Results


SSMVSS053125

■ DFSMS: DCOLLECT Needs to Include More Fields

- **Updated Response: AV-Available**
 - This function is available with z/OS 1.12 (GA 9/10).

■ Description


- DCOLLECT needs to be enhanced to include additional information in the DCOLLECT records. The following information should be included.
- For the "DC" records --
 - Data set name type
 - Extended Addressability
 - Recording technology
 - Performance Scaling
 - Reuse
 - BWO
 - Log Stream ID
 - Forward Recovery Log
 - RLS CF Cache
 - Space Constraint Relief
 - Reduce Space Upt
 - Dynamic volume count



SHARE
in Anaheim
2011

SSMVSS064933

- **Catalog Support of HDELETE for DELETE GDG FORCE**
 - **Updated Response: AV-Available**
 - This function is available with z/OS 1.12 (GA 9/10).
- **Description**
 - Need DELETE GDG FORCE to support the conversion of a DELETE request for a migrated GDS to an HDELETE. Currently this is only supported for individual dataset deletions. A DELETE GDG FORCE results in the recall of all migrated GDS's just to delete them. A DELETE GDG FORCE of a 255-entry GDG where each GDS is migrated to a different tape can take hours.



SSMVSS063069 / SSMVSS063070


■ DFSMS: Method to Empty a PDS / PDSE

- **Response: AV-Available**


- This function is available with z/OS 1.12 (GA 9/10).

■ Description

- An IBM supported single-step method is required to empty a PDS (Partitioned Data Set). The data set must be returned to the state at which it was originally allocated (such as no members, total number of directory blocks, and fully compressed status ... number of extents is not addressed by this requirement). This method should be available from the BATCH and TSO environments.



SHARE Requirements Summary




SHARE
Technology • Connections • Results

- Re-Review of Existing Requirements

Requirement #	Title	Status
<u>SSMVSS07009</u>	DFSMSdss RESOTRE FULL DSCHA bit based on DUMP RESET	RC – Recognized (Orlando 2008) RC – Recognized (Boston 2010)

SSMVSS07009

- **DFSMSdss RESTORE FULL DSCHA bit based on DUMP RESET**
 - **Response: RC – Recognized**
 - IBM agrees with the request and a solution appears to be a desirable objective. A solution, however, may not presently appear feasible or implementable.
- **Description**
 - RESTORE FULL should by default set the VTOC DS1DSCHA "changed" bits based on whether RESET was specified on the DUMP FULL. Only if RESET was specified should these bits be cleared. If RESET was not specified, DS1DSCHA bits should be restored without alteration. This is a modification of SHARE requirement SSMVSS07002, incorporating the intent of that requirement as well as what may have been the original intent of the DSS developers when the RESET option was designed.



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.