

IBM Responses to Requirements Submitted by the SHARE MVSS Project

Barbara McDonald, Glenn Wilcock
IBM - bawhite@us.ibm.com
IBM - wilcock@us.ibm.com

*August 4, 2014
Session 16134*



#SHAREorg



SHARE is an independent volunteer-run information technology association
that provides **education, professional networking and industry influence.**

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.

Connect with IBM System z on social media!



Subscribe to the new [IBM Mainframe Weekly](#) digital newsletter to get the latest updates on the IBM Mainframe!



[System z Advocates](#) **
[IBM Mainframe- Unofficial Group](#)
[IBM System z Events](#)
[Mainframe Experts Network](#)
[SHARE](#)



[IBM System z](#) **
[IBM Master the Mainframe Contest](#)
[IBM Destination z](#)
[SHARE Inc.](#)



[IBM System z](#) **
[IBM System z Events](#)
[Destination z](#)
[SHARE](#)

System z SMEs and Executives:
Deon Newman - [@deonnewm](#)
Steven Dickens - [@StevenDickens3](#)
Michael Desens - [@MikeDesens](#)
Patrick Toole - [@Pat_Toole_II](#)
Kelly Ryan - [@KellykmRyan](#)
Richard Gamblin - [@RichGx](#)

Blogs

[IBM Mainframe Insights](#) **
[Millennial Mainframer](#)
[#MainframeDebate](#) blog
[SHARE](#) blog
[IBM Destination z](#)



[IBM System z](#) **
[Destination z](#)



[IBM Mainframe50](#)

Include the hashtag [#mainframe](#) in your social media activity and [#mainframe50](#) in 50th anniversary activity

Agenda

- IBM Responses to SHARE Requirements
 - Requirements Waiting for Responses
 - Requirements Open for Discussion
 - Existing requirements with updated responses
 - DFSMSHsm Short List
- Backup
 - Requirements Available in V2.1

SHARE Requirements Summary

- Waiting for Response

Requirement #	Title	Status
SSMVSS12006	Tape Virtualization Engines should pre-cache during DR recovery	SUG – Suggestion (Boston 2013) SUG – Suggestion (San Fran 2013) Open for discussion (Anaheim 2012)
SSMVSS13010	Provide 'MOVE' option to IEBCOPY	Waiting for response (Pittsburgh 2014) Open for discussion (Anaheim 2014)
SSMVSE14003	Provide SET IGGCAT=xx to dynamically change catalog options	Waiting for response (Pittsburgh 2014) Open for discussion (Anaheim 2014)



SSMVSS12006



- **Tape Virtualization Engines should pre-cache during DR recovery**
 - **Response: SUG – Suggestion (Boston 2013) – responded SUG 11/20/12**
 - Comments from development: *Since the cache in a TS7740 is limited, there would probably need to be some selection criteria as to which stacked/logical volumes are brought back to cache. There is a pre-stage tool that can bring volumes back efficiently, so it may be possible to exploit this tool. Timeframe for delivery is TBD.*
- **Description**
 - When recovering a TSxxxx Virtualization Engine which uses physical cartridges, the customer should have the option of having all logical volumes on inserted physical cartridges recalled to the disk cache - to the limit of the disk cache threshold.
- **Benefit:**
 - Staging to disk cache at recovery time will eliminate a lot of physical cartridge contention during z/OS (or other OS) recovery. Doing recalls in parallel as described in the suggested solution will speed the necessary recalls as well.
- **Solution**
 - The interfaces pertaining to I/O station interface needs a UI item, button, or switch which acts in concert with the "Copy/Export Recovery" function to toggle "DR Insert" vs. "Normal Insert". During DR insert, all inserted physical tapes ought to be queued for mount by all drives but one, and each valid logical volume on the physical tapes recalled to disk cache until a threshold is reached. When the threshold is reached, recalls should quiesce (assuming those in progress had already allocated the necessary space for recall, the event logged, and cartridge insert return to "normal". The one drive omitted is there in case of a hardware error which might be correctable via swapping to a different drive.



SSMVSS13010



- **Provide 'MOVE' option to IEBCOPY**

- **Response: Waiting for Response (Pittsburgh 2014); cannot find in FITS or RFE**
 - Data Set Commander for z/OS Version 8 Release 1 (formally ISPF Productivity Tool for z/OS) does provide this capability.

- **Description**

- Provide an option for IEBCOPY that deletes the members that are successfully copied to another PDS(E). "MOVE" is my suggested keyword as that is the same function used to do this with ISPF.

- **Benefit:**

- Being able to move PDS(E) members from one PDS(E) to another in a single execution of a program. Speed and ease of use over having to code ISPF in batch, which isn't trivial for most users or using another job step with IEHPROGM or IDCAMS etc. to delete the PDS members from the original library they were copied from.

- **Solution**


- an IEBCOPY 'MOVE' option with the similar syntax to the COPY option.

- **Impact:**

- Continue to use multiple utilities / job steps to move members from one PDS(E) to another or develop / use ISPF in batch code with ISPF services to move the members.

- **Discussion:**

- None listed



SSMVSE14003

- **Provide SET IGGCAT=xx to dynamically change catalog options**
 - **Response: Waiting for Response (Pittsburgh 2014); cannot find in FITS or RFE**
- **Description**
 - Allowing a SET IGGCAT=xx would allow a user to dynamically change catalog options without having to restart the catalog address space (CAS).
- **Benefit:**
 - Restarting CAS can delay processing of all work on the system. Providing a dynamic change would reduce this interruption.
- **Solution**
 - None listed
- **Impact:**
 - Performance issues might occur when trying to change a catalog parameter.
- **Discussion:**
 - None listed



SHARE Requirements Summary

- Open for Discussion (new)

Requirement #	Title	Status
SSMVSS14001	HSM - Prioritize On Demand Migrations	Open for discussion (Pittsburgh 2014)
SSMVSS14002	Change DEFrag TYPRUN=NORUN into separate ADRDSSU command	Open for discussion (Pittsburgh 2014)
SSMVSS14005	All FRBACKUP and FRRECOV Messages Should Be Sent Back to the User	Open for discussion (Pittsburgh 2014)
SSMVSS14006	Prevent VTOC corruption	Open for discussion (Pittsburgh 2014)
SSMVSS14007	Track VTOC and VTOC index updates	Open for discussion (Pittsburgh 2014)



SSMVSE14001


- **HSM - Prioritize On Demand Migrations**
 - **Open for Discussion (Pittsburgh 2014)**
- **Description**
 - When a volume exceeds the high threshold the ENF72 message will trigger an ON Demand Migration for that volume. In our environment we have many PROD volumes that exceed the high threshold, but at the same time we also have many TEST volumes that exceed their threshold. What we would like is the possibility to (somehow) prioritize the migration of the PROD volumes over the migration of the TEST volumes (or any other non-PROD volume for that matter).
- **Benefit:**
 - With the requested prioritization possibility we would ensure that PROD volumes would have enough space available. TEST volumes should not take priority over PROD.
- **Solution**
 - We're running on 1.13 but will be migrating to 2.1 within a couple of months and I believe that there is some kind of priority setting for the storage group that these PROD- or TEST-volumes belong to but I understood that it is not applicable for On Demand Migration. The priority setting comes only into play with Automatic Space Management and Backup processing.
- **Impact:**
 - Possible PROD errors because there's not enough space available. Manual volume additions to the Storage Group and/or manual migrations to make enough space available.



SSMVSE14002



- **Change DEFrag TYPRUN=NORUN into separate ADRDSSU command**
 - **Open for Discussion (Pittsburgh 2014)**
- **Description**
 - The ADRDSSU DEFrag command with PARM='TYPRUN=NORUN' is a very handy way to produce a VTOC extents report, similar to the old DISKMAP program on the CBT tape. However, if you forget to specify the correct PARM= keyword on the EXEC statement, then DSS will actually attempt to perform the DEFrag on the specified volume, which can be an unpleasant surprise. Regardless, ADRDSSU controls usage of the DEFrag command via READ access to the same FACILITY class profile, with or without the TYPRUN=NORUN parm. There is no way to authorize a user to execute the DEFrag command only when PARM='TYPRUN=NORUN' is also specified.
- **Benefit:**
 - The benefit is primarily for reducing the possibility of an unintended DEFrag attempt. Using the DEFrag command for a DISKMAP function can also be somewhat intimidating for something that is read-only in nature. By clearly separating the two functions into respective FACILITY resources will not only provide better granularity of control, but will also benefit auditing concerns.
- **Solution**
 - One approach might be to clone this specific function and make it optionally available under a different DSS command name such as MAP, (sans the PARM= requirement).
- **Impact:**
 - None, except that many will be less inclined to exploit a DSS feature that can be very useful.



SSMVSE14005



- **All FRBACKUP and FRRECOV Messages Should Be Sent Back to the User**

- **Open for Discussion (Pittsburgh 2014)**

- **Description**

- All messages associated with a particular FRBACKUP or FRRECOV command should be sent back to the user who issues these commands. Alternatively, these messages should be sent to an activity log created specifically for the FRBACKUP or FRRECOV command similar to the way activity logs are created for ABACKUP and ARECOVER commands.
 - Here is an example of what the problem is. I issued the following command from a batch job.
 - HSEND WAIT FRBACKUP COPYPOOL(TESTDSMX) EXECUTE - DUMP RETURNCONTROL(DUMPEND) DUMPCLASS(FRDUMP)
 - The following are the messages returned to the batch job.
 - ARC1001I FRBACKUP COPYPOOL(TESTDSMX) EXECUTE DUMP RETURNCONTROL (DUMPE COMMAND DUMPCLASS(FRDUMP)
 - ARC1808E ONE OR MORE FAILURES OCCURRED DURING FAST REPLICATION ***** OF ***** .
 - An extremely long list of messages went to SYSLOG. As you can see in the use case, messages associated with that particular FBACKUP command are scattered all over the system. The batch job that issued the command gets an indication that something did not work, but there is no information about why it did not work. SYSLOG had a little more information, but it did not provide the real reason for the failure. I had to issue a RELEASE HARDCOPY command and go look in the activity log to find the problem.

- **Benefit:**

- Most users of the FRBACKUP command do not have access to issue the RELEASE HARDCOPY command or access to the actual activity log. Therefore, when problems are encountered, a storage administrator has to be contacted to help diagnose the problem. This is time consuming, and it could cause delays in the batch cycle. It would be much easier and faster to diagnose problems with the FRBACKUP and FRRECOV commands if all of the messages were in a single place that would be easy for the issuer of these commands to get to.

- **Solution**

- A solution similar to the way the messages are handled for the ABACKUP and ARECOVER commands would be satisfactory.



SSMVSE14006



▪ Prevent VTOC corruption

- **Open for Discussion (Pittsburgh 2014)**
- Already under consideration by development. Initial thoughts include:
 - 1) ICKDSF invalidate the old VTOC location when initializing a volume or moving the VTOC.
 - 2) OCEOV compare the UCB VTOC location with the VTOC location in the volume label and failing the dataset open if there is a mismatch.
 - 3) DADSM/CVAF failing allocation if the VTOC pointed to by the UCB is invalidated
 - 4) Cutting SMF records for updates to VTOC and Index
 - 5) Health Check to periodically validate the VTOC and index.
 - 6) Copy services I.E. PPRC, Flash Copy, HSM and DSS functions all need to be looked at also.

▪ Description

- Currently, it is possible for the VTOC and VTOC index to get out of synch. When this condition occurs, the results can be catastrophic since data loss can occur. We need IBM to evaluate its ability to mitigate this corruption by performing some sort of (intent) integrity checking prior to making an update to the VTOC or the VTOC Index to preserve the integrity of these resources.

▪ Benefit:


- This is a data loss issue. It caused significant problems for one of our applications. Six weeks later, we still have been unable to pinpoint the cause.

▪ Solution

- None provided.

▪ Impact:

- This issue caused an in-use part of the volume to show as available. A subsequent allocation overlaid a production data set, causing a 16-hour outage.



SSMVSE14007



▪ Track VTOC and VTOC index updates

– **Open for Discussion (Pittsburgh 2014)**

- Adding trace records for DADSM and CVAF may help to understand when an overlay has started.
- Also need to consider trace records for ICKDSF actions (e.g., Init, Move, Expand VTOC) as well as require other software, such as DFSMSdss and ISVs, that manipulate the VTOC to also add SMF Trace records.

▪ Description

- The results of updates to the VTOC and VTOC Index when improperly done can be catastrophic. However, there is no facility available to record these requests and who they were made by. Some type of facility is required to warehouse VTOC and Index update requests so its potentially possible to determine what product/component made the request so they can be notified of their error.

▪ Benefit:


- VTOC corruption issues lead to data loss. We experienced a significant data loss issue recently when two different volumes had corrupted VTOCs (VTOC and index were out of synch). We have been unable to determine who or what made the improper updates to the VTOCs. Having the ability to determine the cause would help prevent further data loss issues.

▪ Solution

- This item could also cut SMF Type x (DADSM) and Subtype x' (DADSM Subfunction) and SMF Type y (CVAF) and Subtype y' (CVAF Subfunction) records for all requests made to DADSM and CVAF.

▪ Impact:

- We cannot diagnose what caused our VTOC corruptions. This leaves us vulnerable to further data corruption issues and loss of reputation.



SHARE Requirements Summary

- Existing requirements with updated responses

Requirement #	Title	Status
SSMVSS13003	Comprehensive Programming Environment Documentation	Under Consideration (Pittsburgh 2014) Resolver Assigned (Anaheim 2014) Open for discussion (Boston 2013)
SSMVSS13005	PDSE V1 to V2 Bulk Conversion Utility	Uncommitted Candidate (Pittsburgh 2014) ACK – Acknowledged (Anaheim 2014)
SSMVSS13008	PDSE V2 health checks	Uncommitted Candidate (Pittsburgh 2014) ACK – Acknowledged (Anaheim 2014)
SSMVSS13009	Externalize indicator for PDSE V1/V2	Return for addtl info (Pittsburgh 2014) ACK – Acknowledged (Anaheim 2014)



SSMVSS13003



- **Comprehensive Programming Environment Documentation**

- **Response: Under Consideration**

- **Description**

- Most z/OS programming interface manuals contain an 'Environment' section for each system service being described. This section documents the allowable run-time environment at the time the service is called. For example:

- <Sample Environment Section>

- Minimum authorization: Supervisor state

- Dispatchable unit mode: Task or SRB

- Cross memory mode: PASN=HASN or PASN^=HASN

- AMODE: 24-bit or 31-bit. To reference the copy of the parameter string, the user must be in 31-bit addressing mode.

- ASC mode: Primary or AR

- Interrupt status: Enabled for I/O and external interrupts

- Locks: No locks held

- Control parameters: Must be in the primary address space

- </Sample Environment Section>

- This information should be provided in a similar format for the DFSMS programming interfaces described in z/OS DFSMS Macro Instructions for Data Sets.

- **Benefit:**

- Not all of this information is documented in the existing publication. "Trial and error" approaches are time consuming and error prone and can lead to bugs or worse (e.g., integrity issues in authorized code). Comprehensive documentation will clearly spell out the intended environment and any restrictions the programmer(s) should be aware of.

- **Discussion:**

- None listed



SSMVSS13005

▪ PDSE V1 to V2 Bulk Conversion Utility

– **Response: Uncommitted candidate (RC in FITS)**

- DFSMSdss COPY would be the preferred solution as it already supports generic qualifiers (ie wildcards), the RENUNC keyword to avoid duplicate catalog entries, and the DELETE keyword if customers wanted to “convert in place”.

▪ Description

- The upcoming PDSE V1 to V2 conversion did not provide a utility to convert PDSE's in bulk. While IEBCOPY can convert a single PDSE, this is not feasible for site with hundreds or thousands of PDSE's.

▪ Benefit:

- Users benefit from labor savings converting PDSE V1 to V2. IBM benefits by expediting the conversion from V1 to V2 PDSE's so that V1 can be deprecated as soon as possible.

▪ Solution

- DFDSS seems the most logical utility to use. It already provides the CONVERT PDSE syntax, so a CONVERT PDSEV2 or something similar seems logical. If DFDSS is used, you should also provide selection keywords, such as PDSEVER,EQ,1, PDSEVER,EQ,2, etc.

▪ Impact

- If IEBCOPY remains the only means to convert PDSE's from V1 to V2, conversions to V2 will be significantly delayed due to the labor-intensive effort required to run IEBCOPY for each individual V1 PDSE.



SSMVSS13008



▪ PDSE V2 health checks

– **Response: Uncommitted candidate (RC in FITS)**

- Agree this is potential candidate for a future release once the adoption of PDSE V2 is more widespread. What timeframe is required?

▪ Description

- The release of PDSE V2 did not come with any health checks. Health checks for IGDSMSxx, the SMS dataclass parms, etc. should be created with the recommendation to allocate PDSE V2 by default.

▪ Benefit:

- Users benefit from labor savings converting PDSE V1 to V2. IBM benefits by expediting the conversion from V1 to V2 PDSE's so that V1 can be deprecated as soon as possible.

▪ Solution

- Left to the developers, but a health check for IGDSMSxx, and for the SMS dataclasses seems like a minimum requirement.


▪ Impact

- Customer may not take the actions required to convert to PDSE V2, delaying the eventual deprecation of V1.



SSMVSS13009

- **Externalize indicator for PDSE V1/V2**
 - **Response: Returned for more information**
 - Adding an indicator in the VTOC or catalog will create an incompatibility with existing PDSE V2 data sets which would require explicit migration action on behalf of customers.
- **Description**
 - There is no indicator available to show if a PDSE is V1 or V2. Currently, the only way to get this data is to use the FAMS OCO interface. A non-OCO service or a flag in the catalog or VTOC is necessary to provide this information.
- **Benefit:**
 - Users benefit from being able to potentially expedite converting PDSE V1 to V2. IBM benefits by from expediting since V1 can be deprecated as soon as possible.
- **Solution**
 - Perhaps a catalog or VTOC indicator. A separate requirement has been submitted for the LISTDSI service.
- **Impact**
 - With inability to easily identify PDSE V1 and V2 datasets, customer may not be able to take the actions required to convert to PDSE V2, which will delay the eventual deprecation of V1.



DFSMSHsm 'Short List'

D-Type APARS

- **Update CDS migration record w/o recall for ALTER CARECLAIM**
 - Do not recall a data set that is altered for the CA Reclaim value
- **CA Reclaim status in migration records**
 - When a VSAM dataset is migrated to HSM, CA Reclaim information is not stored. Customer is looking for a way of knowing if the VSAM dsn that is migrated had CA RECLAIM(YES) or (NO). Currently, they have to recall the data set and then do LISTCAT.
- **FSRs for skipped data sets**
 - During Primary Space Management, DFSMSHsm will only create an FSR record for files that pass eligibility checking. When a file is failed without attempting activity on it, no FSR is created but if it fails after this point, DFSMSHsm will cut an FSR record. The client would like to see the failures included regardless of when the failure occurs as they use FSR records to monitor system activity.

Small Release Enhancements

- **Target Migration Tier by Data Set Size**
 - Provide a SETSYS command for Minimum/Maximum data set size that should be processed

SETSYS MIGRATION(DATASETSIZE(ML1MIN(nn) ML1MAX(nn) ML2MIN(nn)))

- **Update CRQ to handle multiple concurrent recall requests for the same data set. Eliminate the ARC1102I message.**
- **Use RACF to authorize specified users to unlimited recall requests. For other users, limit to a specified number.**

Small Release Enhancements

- **Write QUERY output to a file or SYSOUT**
- **QUERY should be its own task**
 - Query executes under the main HSM task, quiescing all other work
- **Use SG Priority for ODM and Interval Migration**
- **Support Fast Subsequent Migration for spanned data sets**
- **Close the CDSs after backup so that catalog stats are correct**
 - Not an issue when using RLS

Backup

SHARE Requirements Summary



- Requirements To Be Available in V2.1

Requirement #	Title	Status
<u>SSMVSE10018</u>	Provide creation date for alias	AC – Accepted (Anaheim 2012) RC – Recognized (Anaheim 2011)
<u>SSMVSS10004</u>	HSM: RECYCLE should automatically restart after tape takeaway	AV – Available (Anaheim 2014) AC – Accepted (San Francisco 2013) RC – Recognized (Anaheim 2011)
<u>SSMVSS10003</u>	HSM: Make TAPECOPY restartable after a tape takeaway	AV – Available (Anaheim 2014) AC – Accepted (San Francisco 2013) RC – Recognized (Anaheim 2011)
<u>SSMVSS01007</u>	DFSMS: IDCAMS needs to support RLS	AC – Accepted (Pittsburgh 2014) RC – Recognized (2007)
<u>SSMVSS11010</u>	Remove line limit for PDSE members	AV – Available (Boston 2013) AC – Accepted (Atlanta 2012) Open for discussion (Orlando 2011)
<u>SSMVSS08002</u>	Space Release will not release over-allocated space for MVS files	AC – Accepted (Boston 2010) AC – Accepted (Denver 2009)
<u>SSMVSS07002</u>	DFSMSdss: Do not reset DS1DSCHA bits on RESTORE FULL	AV – Available (Anaheim 2014) SUG – Suggestion (San Diego 2007)
<u>GO6SMG90004</u>	DFDSS allow renaming of VSAM data sets on a physical restore	AV – Available (Anaheim 2014)



SSMVSE10018

- **Provide creation date for alias**
 - **Response: AC – Accepted (Anaheim 2012)**
- **Description**
 - We are attempting to cleanup obsolete HLQs (High Level Qualifiers). If an alias has no associated datasets, there is no easy way to determine whether this is a new alias and no datasets have been created or this is an obsolete alias that should be deleted. The alias creation date would also be useful for aliases that are related to datasets.
- **Benefit:**
 - Identifying the creation date would simplify alias cleanup.
- **Solution:**
 - LISTCAT needs to show the creation date. Creation date is currently not stored.

Impact:

- It is difficult to determine if HLQs are obsolete or current. Since the information isn't stored, there isn't an alternative solution.



SSMVSS10004

- **HSM: RECYCLE should automatically restart after tape takeaway**
 - **Response: AV – Available (Anaheim 2014)**
- **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



SSMVSS10003

- **HSM: Make TAPECOPY restartable after a tape takeaway**
 - **Response: AV – Available (Anaheim 2014)**
- **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



SSMVSS01007

- **DFSMS: IDCAMS needs to support RLS**
 - **Response: AC – Accepted (Pittsburgh 2014)**
- **Description**
 - IDCAMS needs to support VSAM Record Level Sharing (RLS) when using the EXPORT, IMPORT, PRINT, and REPRO commands. Currently, IDCAMS only provides the capability to use these commands when the RLS data set is defined with SHAREOPTIONS(2,3), and this support does not provide any type of READ integrity.
 - At minimum, IDCAMS should support the use of RLS for data sets when they are opened for input and the JCL RLS parameter is specified in the DD statement for the data set.
 - Previously, customers have requested this functionality, but IBM rejected the request (APAR OW45478) stating the program was working as designed..

Impact:

- IBM develops and markets VSAM RLS, but they do not provide any utilities that support this feature. This makes the RLS function less attractive to applications and inhibits its exploitation..



SSMVSS11010

- **Remove line limit for PDSE members**
 - **Response: AV – Available (Boston)**
- **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.



SSMVSS08002

- **Space Release will not release over-allocated space for MVS files**
 - **Response: AC-Accepted (Boston 2010)**
 - This response will be tracked and updated.
 - Initial support for VSAM data sets is available in z/OS V1R12, and support for sequential data is targeted to follow at a later date. .
- **Description**
 - Normal space release functions such as RLSE coded in JCL, Immediate or conditional space release in the Management class, DSS or FDR space release during defrag of compactor and the TSO FREE command will not release the over-allocated space for multivolume datasets when the data extents residing on subsequent volumes are empty. This occurs when the primary (or secondary) space allocation is excessive and causes the allocation to span volumes, when the file is populated and does not use all the space requested, space release occurs to the extent on the last volume written to. The empty space occupied by extents on subsequent volumes is not freed. This results in wasted, unused, unavailable space in the storage pools.



SSMVSS07002

- **DFSMSdss: Do not reset DS1DSCHA bits on RESTORE FULL**
 - **Response: AV – Available (Anaheim 2014)**
- **Description**
 - This would require a Design Change for DFSMSdss. By current design, when DFSMSdss is invoked to perform a full-volume RESTORE from a full-volume DUMP, it unconditionally resets the "dataset changed" (DS1DSCHA) bit in every Format-1 DSCB in the VTOC at RESTORE time. This results in the restored volume NOT being a true and correct copy of the original source volume at the time of the DUMP FULL (except in two special cases), and can lead to loss of data in disaster recovery or similar situations by misleading other programs (e.g., DFSMSHsm) as to the status of individual datasets on the restored volume, inducing the other program to take a different action regarding individual datasets on the restored volume than it would on the original source volume. The current design asserts that a full-volume dump is "by definition, a backup" of a dataset on that volume, and that resetting the DS1DSCHA bit at RESTORE time is justified because the dataset has not changed since the "backup". However, this assertion ignores the fact that the entity restored is the entire volume, not a particular data set, and therefore the RESTORE FULL should deliver a volume that is an exact replica of the original source volume in all material respects, including all status bits in the DSCBs. .



GO6SMG90004



- **DFDSS allow renaming of VSAM data sets on a physical restore**
 - **Response: AV – Available (Anaheim 2014)**
- **Description**
 - Currently the only way to rename a VSAM dataset on a restore is if the dataset was dump logically.
- **Benefit:**
 - When converting from an OEM product of DFDSS where the OEM product supported this function, the only way that we could have the same capability with DFDSS is to back-up all the packs that VSAM datasets reside on using the logical technique, which according to the DFDSS Users Guide does not perform as good as the physical technique, and then back-up all other packs using the physical technique. It is our feeling that it would be much easier and more competitive if DFDSS changed the physical restore technique to allow renaming of VSAM datasets on a restore.
- **Solution:**
 - Change the physical DUMP/RESTORE technique to allow renaming of VSAM datasets on a restore.



Complete your session evaluations online at www.SHARE.org/Pittsburgh-Eval



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.

