



IBM Responses to Requirements Submitted by the SHARE MVSS Project


Barbara McDonald
IBM - bawhite@us.ibm.com

March 10, 2014
Session 15079



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
 - Requirements Available in V2.1



SHARE Requirements Summary



- Waiting for Response

Requirement #	Title	Status
SSMVSS13005	PDSE V1 to V2 Bulk Conversion Utility	ACK – Acknowledged (Anaheim 2014)
SSMVSS13008	PDSE V2 health checks	ACK – Acknowledged (Anaheim 2014)
SSMVSS13009	Externalize indicator for PDSE V1/V2	ACK – Acknowledged (Anaheim 2014)
SSMVSS12006	Tape Virtualization Engines should pre-cache during DR recovery	SUG – Suggestion (Boston 2013) SUG – Suggestion (San Fran 2013) Open for discussion (Anaheim 2012)



SSMVSS13005

▪ PDSE V1 to V2 Bulk Conversion Utility

- **ACK – Acknowledged (Anaheim 2014)**

▪ 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**
 - **ACK – Acknowledged (Anaheim 2014)**
- **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

- **ACK – Acknowledged (Anaheim 2014)**

▪ 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:

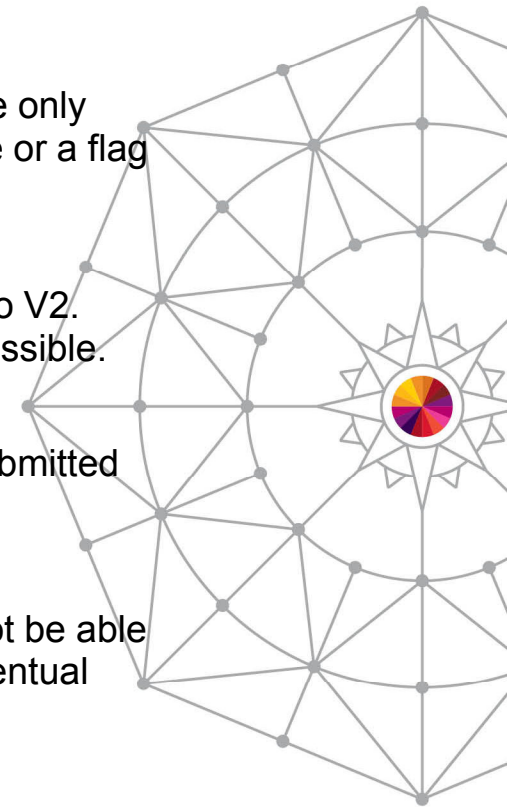
- Users benefit from being able to potentially expedite converting PDSE V1 to V2. IBM benefits by from expediting since V1 can be deprecated as soon as possible.

▪ 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.



SSMVSS12006



■ Tape Virtualization Engines should pre-cache during DR recovery

• **SUG – Suggestion (Boston 2013)**

- 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.




SHARE Requirements Summary



- Open for Discussion (new)

Requirement #	Title	Status
SSMVSS13010	Provide 'MOVE' option to IEBCOPY	Open for discussion (Anaheim 2014)
SSMVSE14003	Provide SET IGGCAT=xx to dynamically change catalog options	Open for discussion (Anaheim 2014)



SSMVSS13010



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

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

- **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**

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

- **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

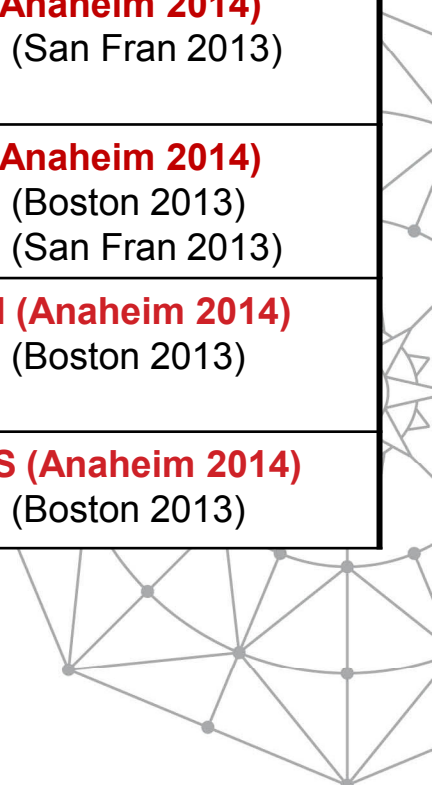


SHARE Requirements Summary



- Existing requirements with updated responses

Requirement #	Title	Status
<u>SSMVSS13001</u>	DFHSM should provide a REXX interface	RC – Recognized (Anaheim 2014) Open for discussion (San Fran 2013)
<u>SSMVSS13002</u>	Allow Override of Messages to TSO Session	RC – Recognized (Anaheim 2014) Open for discussion (Boston 2013) Open for discussion (San Fran 2013)
<u>SSMVSS13003</u>	Comprehensive Programming Environment Documentation	Resolver Assigned (Anaheim 2014) Open for discussion (Boston 2013)
<u>SSMVSS13004</u>	Add the ability to add a date and/or timestamp to a dataset name	Retired – per MVSS (Anaheim 2014) Open for discussion (Boston 2013)



SSMVSS13001



■ DFHSM should provide a REXX interface

- **RC – Recognized (Anaheim 2014)**

■ Description

- DFHSM could supply a REXX interface that would return stem variables like SDSF does.

■ Benefit:

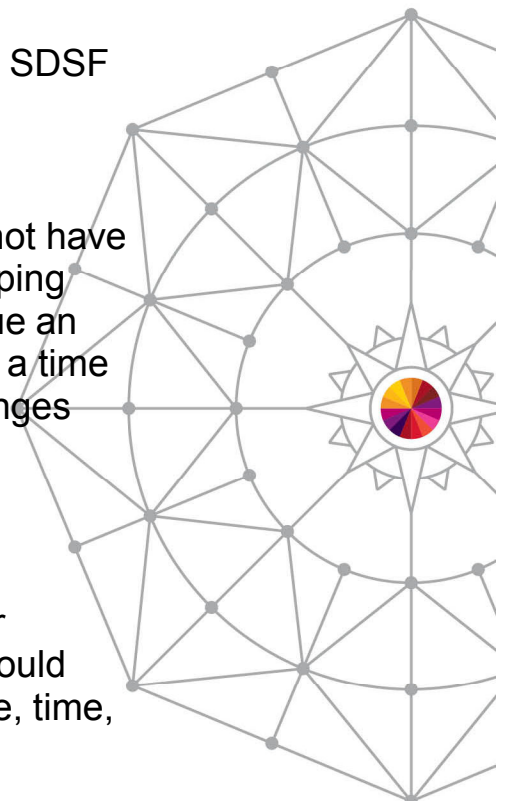
- By providing an HSMEXEC or ARCEXEC function for REXX, a user would not have to parse through lines and lines of output, but could use the PARSE and looping functions in REXX to produce reports they want. Currently a user has to issue an HSM command in REXX trap it if possible, and then extract the data. This is a time consuming process. And prone to errors if the messages or information changes with a release of the software.

■ Solution

- Provide a function like ISFEXEC or ISPEXEC for HSM. It could be called ARCEXEC that would pass back variables that contain the info required. For example, if you wanted to HLIST a dataset, ARCEXEC HLIST dsn MCDS would return stem vars containing any Informational messages, dataset name, date, time, etc as is currently produced in the HLIST command.

■ Discussion:

- None listed



SSMVSS13002



■ Allow Override of Messages to TSO Session

- **RC – Recognized (Anaheim 2014)**

■ Description

- When an HLIST or other HSM command is done under TSO some of the messages cannot be suppressed to the TSO Terminal. It would be helpful to add a parm to the command to allow suppression of such messages. I will use HLIST as an example.
- In a REXX that does an HLIST I will sometimes get messages
 - ARC0138I NO MCDS INFORMATION FOUND FOR DATASET &dsn
- This appears to be a TPUT that cannot be suppressed. So for HLIST I would like to see HLIST DSN('dsn') BOTH NOTPUT which would suppress it or allow my suppression in REXX with MSG OFF functions. Or change the display to something I can suppress in REXX or Clist under TSO. This would also be true when running Batch TSO and HLIST command.

■ Benefit:



- Prevent irrelevant messages going back to a TSO Terminal that will either require a lot of enters, or cluttering up a TSO Terminal.

■ Solution

- HLIST DSN('dsn') BOTH NOMSG or
- HLIST DSN('dsn') BOTH SUPPR(arc?,ARC?,...) (allow specific message suppression) or
- HLIST DSN('dsn') BOTH NOTPUT

■ Discussion:

- None listed



SSMVSS13003



■ Comprehensive Programming Environment Documentation

- **Response: Resolver Assigned**

■ 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



SSMVSS13004



- **Add the ability to add a date and/or timestamp to a dataset name**
 - **Retired – per MVSS (Anaheim 2014)**
- **Description**
 - Instead of the longstanding "infinite GDG" requirement (see GO6SMG91037) add the ability to use the dynamic system symbols for date and time in dataset names, both in JCL and dynamic allocation (NOTE: does that make this an SCP requirement rather than a Storage Management requirement?).
- **Benefit:**
 - We have for a long time had the need for processes to create unique dataset names without changing JCL. That is the primary rationale for generation data groups. GDGs, however, probably due to ancient architectural issues, only allow 255 generations, which are often not enough for many applications which might need to retain daily data for a year or more. Allowing the use of date and time stamps would provide the same 'unique DSN' capability without artificial restriction of the number of unique datasets.
- **Solution**
 - Suggestion: dynamic system symbols are already available to started tasks, so adding them for batch jobs should not be hard (as of z/OS 1.13 I have not seen them available for batch?). Add a system symbol-substitution scan to the handling of the DSN key value in dynamic allocation.
- **Impact**
 - Awkward workarounds for providing more than 255 generations will continue to cost us time and money.
- **Discussion:**
 - Sounds like [GO6SMG91037](#)



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	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 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: RC – Recognized (2007)**

▪ 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.



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



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

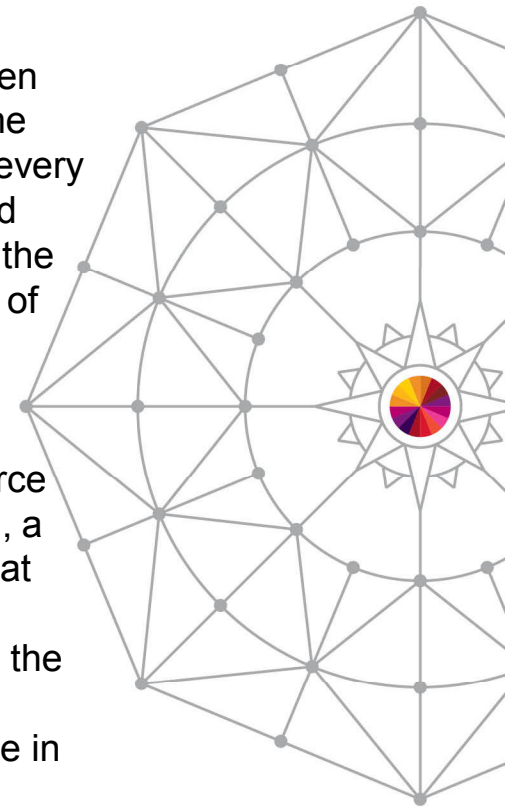
- Normal space release functions such as RLSE coded in JCL, Immediate or conditional space release in the Management class, DSS or FDR space release during defrag of compactor and the TSO FREE command will not release the over-allocated space for multivolume datasets when the data extents residing on 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/Anaheim-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 inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

