



IBM Responses to Requirements Submitted by the SHARE MVSS Project

Barbara McDonald
IBM - bawhite@us.ibm.com

Ed Petka, MVSS Requirements Coordinator
edward.petka@fiserv.com

August 10, 2015
Session 17815



#SHAREorg



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



© Copyright IBM Corp. 2015

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 z Systems 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](#)

[Google+](#)



Key Websites

- [Mainframe Insights](#)
- [Engines of Progress](#)
- [z Systems Software](#)
- [IBM z Systems Redbooks](#)
- [DFSMS](#)
- [z/OS Storage Management Tools](#)

Blogs

[IBM Mainframe Insights](#) **
[Millennial Mainframer](#)
[#MainframeDebate blog](#)
[SHARE blog](#)
[IBM Destination z](#)
[Dancing Dinosaur](#)
[Bob's Mainframe Blog](#)



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



[IBM Mainframe50](#)

Include the hashtag #Mainframe, and also see #z13 and #IBMz

IBM Redbooks



- New look for [IBM z Systems Redbooks](#)
 - Connect with them on Facebook, Twitter, LinkedIn
 - [5 Things to Know About IBM Redbook for IBM z13](#)
- Check out their BLOGS
 - [zEDC Compression](#)
 - [DS8870 and z13 Synergy](#)
- Stay tuned for a new DFSMS V2.2 Redbook
 - **z/OS V2R2: Storage Management and Utilities**
- Visit their Booth at the SHARE Expo
 - IBM Redbooks Booth #101

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



Agenda

- IBM Responses to SHARE Requirements
 - IBM RFE Community
 - DFSMSHsm “Short List” Survey Results
 - PDSE “Short List” Survey Results from Seattle and Redo
 - SHARE MVSS Top “XX” List
 - Requirements Waiting for Responses
 - Existing requirements with updated responses
 - Requirements Available in V2.2
- Backup

What is the IBM RFE Community?



A partnership to provide transparency in product support

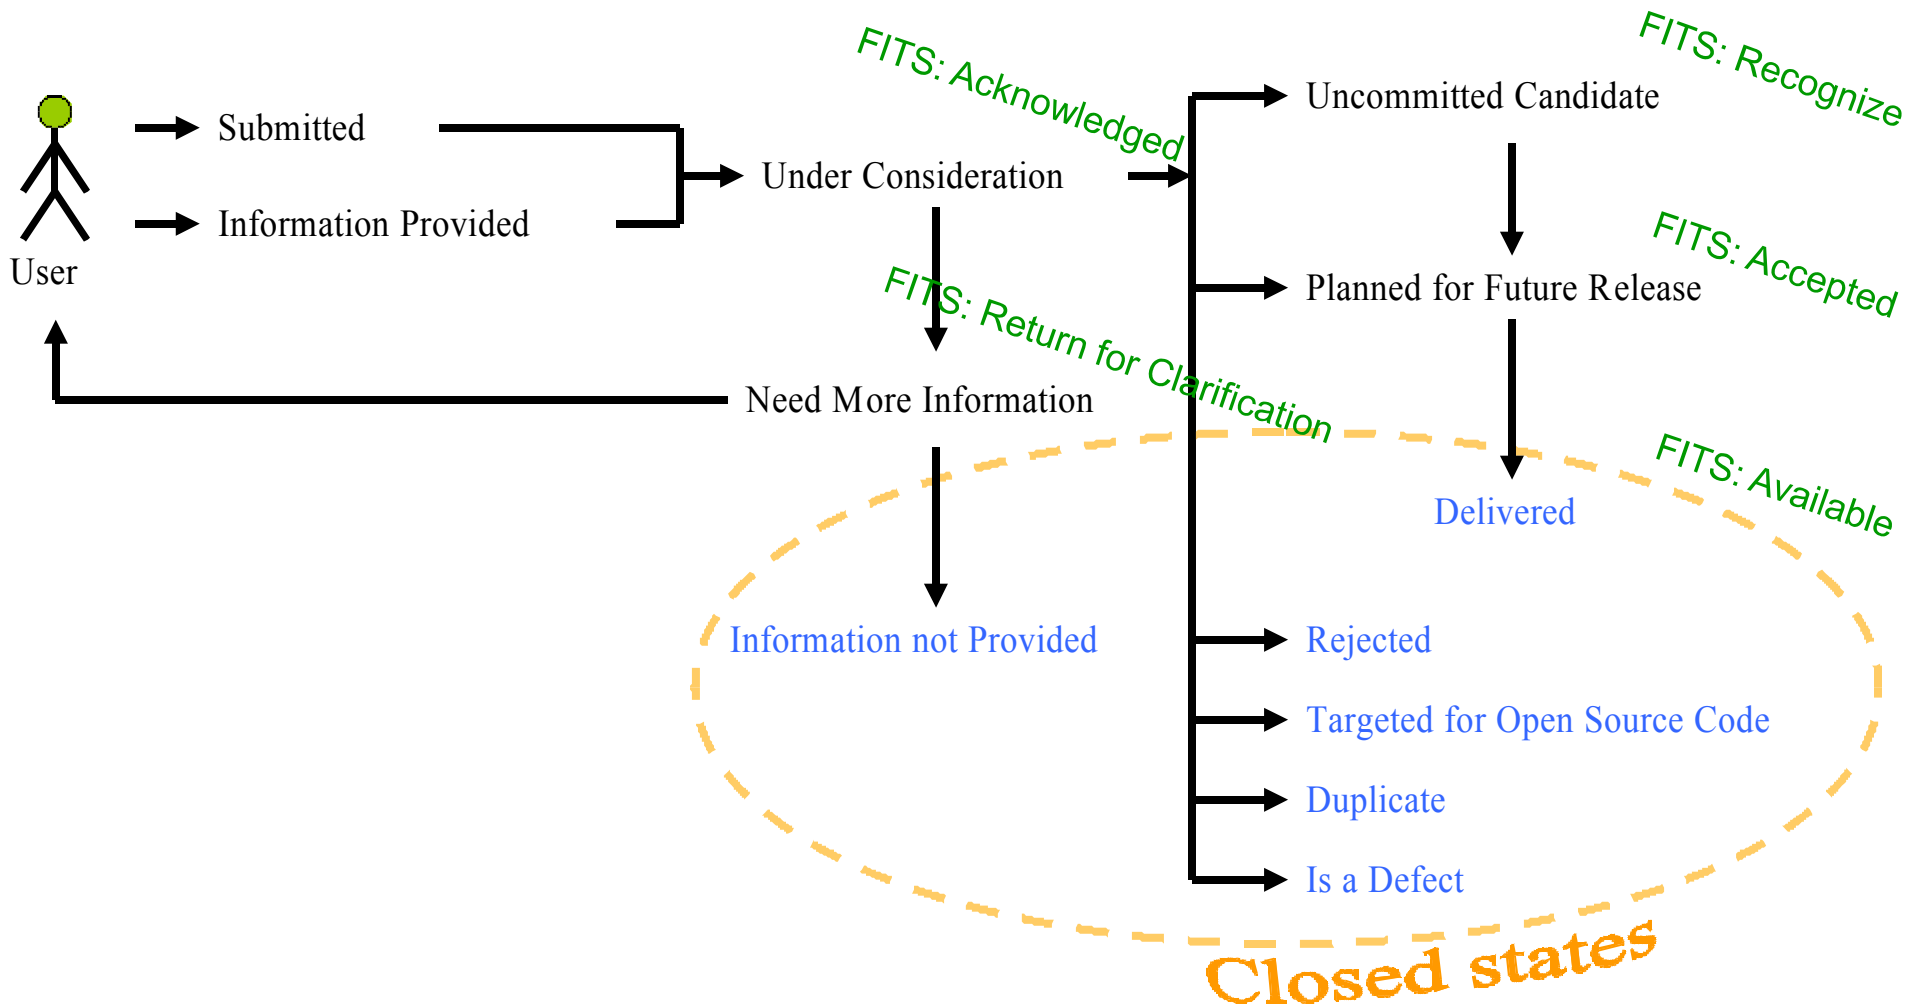
- A community of IBMers and IBM Customers to interact and collaborate on Requests for Enhancements (RFEs) in IBM products
- Strategic replacement for the IBM Internal only FITS MO database, process and tooling
- Hosted on the IBM developerWorks website
 - RFE developerWorks: <http://www.ibm.com/developerworks/rfe>
- SHARE is in the process of converting from FITS to RFE.
 - The Requirements team has submitted to the SHARE Board a proposal to ensure that the alliance and advocacy continues between SHARE and IBM.



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



RFE Requirement State Flow



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



RFE States



Open states

- ▶ **Submitted:** IBM has not yet evaluated this RFE and plan to provide an update **within 30 days of submission.**
- ▶ **Under consideration:** The IBM team is evaluating this RFE. A decision or request for more information will be provided **within 90 days.**
- ▶ **Uncommitted Candidate:** This request will not be delivered within the release currently under development, but the theme is aligned with our multi-year strategy. IBM is soliciting RFE Community feedback for this request through activities such as voting. IBM will update this request in the future.
- ▶ **Planned for Future Release:** This request is a candidate for a future generally available (GA) release. IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. IBM will update this request to reflect any changes.
- ▶ **Need More Information:** IBM is requesting more information from the submitter before the request can be evaluated. This request will remain open in this status for 30 days, but if the submitter doesn't respond, the request will be closed to status Information Not Provided (see below).
 - Note that the move to Not Provided is NOT automatic. It is left to the discretion of each product team.
- ▶ **Information Provided:** The submitter has provided additional information, to assist IBM in evaluating the request.
 - Note that this move is for ANY comment by the customer.

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



RFE States



Closed states

- ▶ **Delivered:** This request has been delivered into a previously released version of the product, or imminently in the next available release.
- ▶ **Rejected:** IBM has considered the request, and it will not be delivered in the future. If this request receives continued customer interest, it may be resubmitted for consideration after 18 months from the submit date.
- ▶ **Information not Provided:** IBM had previously requested more clarifying information. Because the additional information was not provided within 30 days, the request has been closed.
 - [Note: This is NOT automatic.](#)
- ▶ **Is a Defect:** IBM has identified this as a defect requiring a code change to resolve. The RFE has been closed and will be tracked via APAR management. The APAR number will be provided on the RFE, and made available when closed via via Lotus Support Search, Rational Support Search or Tivoli Support Search.
- ▶ **Targeted for Open Source Code:** IBM has determined this feature would need to be developed within the Open Source Community and further tracking of the request should be done in the applicable system (i.e. Bugzilla or other).
- ▶ **Duplicate:** IBM has identified the RFE as a duplicate of another RFE already posted to this community. Please refer to the linked "duplicate of" RFE record for information on this request.

DFSMSHsm “Short List” Survey Results

- Thanks to everyone who provided input for prioritizing the 'short list' of small HSM requirements.
 - Note: Survey results based on both average and mode. Refer to [Pittsburgh Requirements](#) session for details.

- **D-Type APARs**
 1. Create FSRs for skipped data sets *(In Development now)*
 2. Use Storage Group priority for On Demand Migration and Interval Migration
 - *Previously listed as a release level change, but due to the high voting, will try to deliver this as a D-Type APAR instead*
 3. Store the value of CA Reclaim and PDSE V2 in the MCD record so that they can be included with DCOLLECT
 4. ALTER CA Reclaim w/o Recall

- **Release Enhancements**
 1. Provide a SETSYS for the minimize size that a data set must be in order to be eligible for migration
 2. Have QUERY run under a separate task as opposed to locking up the main HSM task
 3. Eliminate ARC1102I messages when using the common recall queue
 4. Provide an option to send QUERY output to SYSOUT or a data set
 5. Enable Fast Subsequent Migration for spanned data sets
 6. Close CDSes after CDS Backup so that the space statistics are correctly reflected in the catalog (only necessary in a nonRLS environment)
 7. Limit the number of concurrent recalls for per user

PDSE “Short List” Survey Results

- Thanks to everyone who provided input for prioritizing the 'short list' of small PDSE requirements.
 - Refer to [Seattle Requirements](#) session for details.

Small Enhancements	Very High	High	Medium	Low	Very Low	Score	# Responses
PDSE Full Cross Sysplex Sharing	4	1	1		1	28	7
Externalize indicator for PDSE V1/V2	2	2	2			24	6
PDSE Sharing (avoid corruption)	2	1	2		1	21	6
PDSE V2 Health Checks	2	2	1			21	5
Multi-volume PDSE	1	2	1	1		18	5
PDSE REORG Function (non-disruptively)	1	3			1	18	5
PDSE V1 to V2 Bulk Conversion Utility	2	1		1	1	17	5
Allow PDSE Member Level Security	1		2	2	1	16	6
PDSE data compression (member level)		1	1	3	1	14	6

Small Enhancements	Very High	High	Medium	Low	Very Low	Score	# Responses
PDSE Full Cross Sysplex Sharing	4	1	1		1	22	7
Externalize indicator for PDSE V1/V2	2	2	2			16	6
PDSE Sharing (avoid corruption)	2	1	2		1	15	6
PDSE V2 Health Checks	2	2	1			13	5
PDSE REORG Function (non-disruptively)	1	3			1	12	5
Multi-volume PDSE	1	2	1	1		11	5
PDSE V1 to V2 Bulk Conversion Utility	2	1		1	1	11	5
Allow PDSE Member Level Security	1		2	2	1	10	6
PDSE data compression (member level)		1	1	3	1	9	6

PDSE “Short List” Survey Results


- Thanks to everyone who provided input for prioritizing the 'short list' of small PDSE requirements.
 - Refer to [Seattle Requirements](#) session for additional details.

Small Enhancements	Value
PDSE Full Cross Sysplex Sharing	usability
Externalize indicator for PDSE V1/V2	usability, space utilization
PDSE Sharing (avoid corruption)	usability
PDSE V2 Health Checks	usability, scalability
Multi-volume PDSE	usability, scalability
PDSE REORG Function (non-disruptively)	scalability, space utilization, performance
PDSE V1 to V2 Bulk Conversion Utility	usability, scalability, performance
Allow PDSE Member Level Security	usability, performance
PDSE data compression (member level)	usability, scalability, space utilization

PDSE 'Short List'

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

PDSE “Short List”



Requirement #	Title	Status
SSMVSS13008 / RFE51884	PDSE V2 health checks	Uncommitted Candidate (Pittsburgh 2014) ACK – Acknowledged (Anaheim 2014)
SSMVSS13009 / RFE51900	Externalize indicator for PDSE V1/V2	Planned for a Future Release (Orlando ‘15) Uncommitted Candidate (Seattle 2015) Return for addtl info (Pittsburgh 2014) ACK – Acknowledged (Anaheim 2014)
SSMVSS13005 / RFE51885	PDSE V1 to V2 Bulk Conversion Utility	Uncommitted Candidate (Pittsburgh 2014) ACK – Acknowledged (Anaheim 2014)
SSMVSS09008 / RFE51892	DFSMSdfp - Multi volume PDSE	Uncommitted Candidate (Seattle 2015) RC – Recognized (Orlando 2011) AK – Acknowledged (Anaheim 2011)
Numerous	Allow PDSE Member Level Security with RACF Data Set Profiles	Uncommitted Candidate
RFE51899	PDSE data compression of the individual member in the data set	Under Consideration
RFE51897	PDSE REORG Function	Uncommitted Candidate
Needed	PDSE Cross Sysplex Sharing to Avoid Corruption	Under Consideration
Needed	PDSE Full Cross Sysplex Sharing	Under Consideration

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



SSMVSS13008

■ PDSE V2 health checks

– **Response: Uncommitted candidate**

- Comments from development: *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

Updated Response: Planned for a Future Release

- Re-discussed in Pittsburgh and the majority agreed what IBM should provide is support in DCOLLECT. Need to ensure the last reference date is not getting updated in this case.
- **Comments from Development:** Since we will have to allocate and open the PDSE, which could present performance problems, propose a new DCOLLECT parameter PDSEINFO that would be used to determine if the user wants the PDSE version number and is willing to suffer the performance impact. Should other PDSE attributes require opening the PDSE in the future, the PDSEINFO parameter could also be reused.
- Also, need to add support in HSM to have an indicator in the migration record for V2, so that DCOLLECT can retrieve the information about migrated PDSEs without a RECALL.
- **Comments from Development:** Any issues if migrate support comes later than the initial support for “active” PDSE’s?

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.



SSMVSS13005

■ PDSE V1 to V2 Bulk Conversion Utility

– **Response: Uncommitted candidate**

- Comments from Development: *DFSMSdss COPY would be the preferred solution as it already supports generic qualifiers (ie wildcards), the RENUMC keyword to avoid duplicate catalog entries, and the DELETE keyword if customers wanted to “convert in place”. Also would like input on MIGRATE / RECALL options.*

■ 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.



SSMVSS09008

- **DFSMSdfp - Multi volume PDSE**
 - **Response: Uncommitted Candidate**
- **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 (QUINNEW) 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.



Member Security

- **Allow Permissioning at the PDSE Member Level of RACF Data Set Profiles**
 - **Response: Uncommitted candidate**
- **Description / Use Case**
 - Several business areas current utilize the same partitioned data set for their requested universal command(ed) user ID/passwords for FTP. Would like to have the ability to segregate access at the member level. I don't want to administer multiple DSNs and would like to keep the profile as it is with one data set.
 - Member security for PDSs and PDSEs to control security separately for each member. The customer acknowledges that for a PDS this will be advisory security meaning that there will be simple ways for an assembler programmer to bypass it.
 - Client requests the option of protecting individual members of a PDSE. It would allow RACF administrators more control of PDSE members. Also instead of telling the customer they need to set up a new library and move all their members over if they wish to make security tighter for certain members a new RACF data set profile could be added instead which would be much less work.



RFE51899: PDSE Data Member Compression

- **PDSE compaction or compression of the data in the data set (ie via zEDC)**
 - **Response: Under Consideration**
- **Description / Use Case**
 - Customer wants to have an enhancement. PDSEs should be able to handle compressed (compact) data; especially when most of them are for source code or reports where over 50% data compression could be achieved and save disk space.
- **Benefits**
 - To manage and save more efficiently the space on dasd.



RFE51897: PDSE REORG Function

- **PDSE REORG Function**
 - **Response: Uncommitted candidate**
- **Description / Use Case**
 - A REORG function/utility is required for a PDSE which provides the capability to reorg a PDSE directory non-disruptively.
 - Ability to release the over-allocated, unused space in a PDSE that is no longer required, after members have been deleted. DFSMS should be able to reduce the size of the pdse after the space have been used and then released.



PDSE Sharing

- **PDSE Cross Sysplex Sharing to Avoid Corruption**
 - **Response: Uncommitted Candidate**
- **Description / Use Case**
 - PDSE was not designed to be shared outside a SYSPLEX and depends on the Sysplex Infrastructure. To enable cross-sysplex sharing, z/OS needs to provide cross sysplex messaging (XCF enhancement).
 - A PDSE will become corrupt when a system outside of a sysplex updates a PDSE while the PDSE is being updated by systems within the sysplex. The systems within the sysplex are abiding by PDSE extended sharing rules, and the system outside the sysplex is not.
 - Usability problem for only-readers outside the sysplex where members disappear and abend0F4s occur because of dirty caches. This later problem, referred to as in-storage corruption, is caused by legitimate changes to the PDSE within the sysplex where the outside the sysplex reader will not be signaled about the PDSE data set changes.
- **Benefits**
 - Improved reliability to avoid accidental sharing resulting in PDSE corruption.



PDSE Sharing

- **PDSE Full Cross Sysplex Sharing**
 - **Response: Uncommitted Candidate**
- **Description / Use Case**
 - PDSE was not designed to be shared outside a SYSPLEX and depends on the Sysplex Infrastructure. To enable cross-sysplex sharing, z/OS needs to provide cross sysplex messaging (XCF enhancement).
 - User wants to be able to create executable programs in PDS using COBOL V5.1; however, the program modules, which are compiled by Enterprise COBOL V5, must use PDSE. The customer faces a problem for the development and the test environment because they cannot share PDSE. Each test environment, the compiling system, some regression test systems and long-running test systems, is MONOPLEX or individual SYSPLEX separately. Additionally, to release the programs into the production systems, the libraries are shared between the compiling system and the production systems.
- **Benefits**
 - Improved flexibility, such as sharing PDSE between production and test environments.



SHARE MVSS Top 'XX' List

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

SHARE MVSS Top 'XX' List



Requirement #	Description
GO6SMG90047	<p>DFP %INCLUDE capability for predefined text in ACS language</p> <ul style="list-style-type: none"> Code (usually FILTLISTS) in the ACS routines must be duplicated by hand. This makes it difficult to maintain the ACS routines and can introduce errors.
SSSHARE016520	<p>DFSMS: Tape Label Architecture Update In Support Of Hi Cap Tape</p> <ul style="list-style-type: none"> The current architecture used to locate data sets on multi-file tape was architected in the era of 800 bpi round tape, and exists unchanged to this day. This technology is inadequate in the era of high capacity tape. The archaic method (used by DFSMSdfp) to locate data sets on a tape involves using the forward-space-file tape command, which causes the drive to read the data on the tape to find the next tape label. This has been an adequate method up to the era of real 3490E tape, but fails to scale in the era of tens of gigabytes of data recorded on a single tape.
SSMVSS98021	<p>Provide a Utility to Diagnose and Correct SCDS Errors</p> <ul style="list-style-type: none"> When the SCDS gets corrupted, IBM should have a utility that could diagnose and correct errors in the SCDS. Currently, the only way to handle this type of problem is the delete the corrupted SCDS and recover it from a backup.
SSMVSS02001	<p>DFSMS: Enhance Library Displays-Show Private and Error Tapes</p> <ul style="list-style-type: none"> Tape library displays should be enhanced to show private and error tapes in addition to the scratch count already shown.
SSMVSS01012	<p>DFSMSrmm: Provide ability to ignore ABEND retention for MOD tapes</p> <ul style="list-style-type: none"> If a user opens a tape for MOD processing, and then later experiences an ABEND, RMM will give the tape the ABEND retention, even if the tape existed with a permanent retention before the MOD operation. Since the ABEND retention is typically only a few days and would result in data loss for the MOD tape, RMM should provide the ability to ignore the ABEND retention for tapes opened MOD.
SSSHARE011543	<p>DF/DSS RELEASE: ALLOW PARTIAL RELEASE</p> <ul style="list-style-type: none"> The RELEASE command releases all free space in a data set. This may be desirable in some cases. However, it is also very desirable to be able to release just part of the free space. RELEASE should be able to release just free extents, or to release a specified number of tracks. It is also desirable to be able to specify the amount of free space to remain in the data set and release the rest.

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

SHARE MVSS Top 'XX' List



Requirement #	Description
GO6SMG93003	DFHSM Tape Blocksize Improvement <ul style="list-style-type: none"> Currently, DFHSM uses a 16K blocksize for its tape data sets. This is not optimum for transfer performance or storage efficiency.
GO6HSM91001	SP MGT multiple 'windows' per day on each DFHSM host <ul style="list-style-type: none"> Allow DFHSM Space Management to execute multiple times per day on the same CPU. Allow the specification of different Storage Groups for each window.
SSMVSS053124	DSS: Allow More Data Sets to be Specified w/o 80A Abend <ul style="list-style-type: none"> DFSMSdss needed to be changed so that more data sets can be specified on data set level operations without experiencing an 80A abend. Currently, if the data set pattern mask or the volume mask causes a lot of data sets to be acted on within a single DSS step, the step can abend with an 80A abend.
SSMVSS053127	DSS - Improve Catalog Search Performance <ul style="list-style-type: none"> The performance of catalog search operations (DATASET INCLUDE and/or EXCLUDE) requires improvement.
SSMVSS97016	Optional parm on DELVOL to delete MCP rec & VTOC copies <ul style="list-style-type: none"> An optional parameter should be added to the HSM DELVOL command to delete the MCP record for this volume being deleted and the VTOC copies on the ML1 volumes. Currently, this is a manual process that the user has to perform.
GO6SMG90034	DFDSS option to suppress message level <ul style="list-style-type: none"> When copying many datasets with DFDSS, there are 4 messages displayed for each dataset copied (for 1000 datasets there are 4000 messages). If DFDSS has an error then all 4000 messages must be examined to find out where the error occurred. There should be an option (parm) to flag only messages that have a Wlevel or above message.
SSMVSS03008	DFSMS: Allow Scratch Tapes to be Read by Both OAM and RMM <ul style="list-style-type: none"> Both OAM and RMM prevent reading scratch tapes. OAM denies the allocation of a TCDB scratch tape for input and RMM denies a scratch tape for input during O/C/E. Both OAM and RMM should be modified to allow scratch tapes to be read.

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

SHARE MVSS Top 'XX' List



Requirement #	Description
SSMVSS053129	Improve Time Stamp Granularity For SMF-6x Records <ul style="list-style-type: none"> SMF 6X records need to carry a timestamp that is more precise and consistent than the Local time stamp.
SSMVSS97012	Remove Restrictions On ACS Rtn Use of &SIZE and &MAXSIZE <ul style="list-style-type: none"> When a JCL user codes LIKE on a DD statement, the &SIZE and &MAXSIZE read-only variables are not available to the ACS routines.
GO6SMG90014	DFHSM ISMF display DFHSM backup versions <ul style="list-style-type: none"> There is no provision with the ISMF panels to display information on backup copies of uncataloged datasets for the purposes of HALTER, HRECOVER, or HBDELETE processing. including datasets which have been deleted.
SSMVSS053115	DFSMSdss - TYPRUN=NORUN and DEFrag for SYSRES <ul style="list-style-type: none"> Allow the DSS DEFrag option to properly complete its function on a system residence volume (SYSRES/IPL volume) if and only if TYPRUN=NORUN is specified.
SSMVSS97018	Make Batch Navquest Functions Easier to Use under ISMF <ul style="list-style-type: none"> The old NaviQuest functions should be changed to also build the JCL dynamically. The user should be prompted for the information that changes each time, like SYSIN INFORMATION, by the use of panels. Then the user should be prompted about whether to SAVE the JCL, SUBMIT it, SAVE and SUBMIT it, or cancel out of it. This would make the NaviQuest functions consistent with the other functions within ISMF..
SSSHARE017174	ISMF: Add ability to save lists to PS or PDS files <ul style="list-style-type: none"> ISMF saved lists currently are saved in the ISPF profile dataset. These lists are difficult to manipulate or share with other users.
SSMVSS033105	DFSMS: Jobname & Date/Time That DASD D/S Was Last Modified <ul style="list-style-type: none"> Provide a command that gives the name of the JOB or TSO USERID that last modified a DASD dataset, and the date/time of that occurrence. If the file hasn't been modified, then the command should give the date/time it was created with the JOBNAME or TSO USERID..

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

SHARE MVSS Top 'XX' List



Requirement #	Description
SSMVSS033104	DFSMS: SMF60 Records for DELETE USERCATALOG <ul style="list-style-type: none">The IDCAMS DELETE USERCATALOG command, however, does not create SMF-60 DE records to indicate when the catalog's VVR components have been deleted from the VVDS. As a result, forward recovery of a VVDS can become inaccurate.
SSSHARE012292	ISMF: Select Criteria for DASD Based on SMS Status(ie.disnew) <ul style="list-style-type: none">ISMF panels do not have selection criteria for SMS status of volumes.
GO6SMG89022	DFP ENHANCEMENT FOR NEW ALLOCATIONS <ul style="list-style-type: none">DFP SHOULD TAKE FRAGMENTATION AND NUMBER OF DSCB'S (AND INDEX) ENTRIES INTO ACCOUNT WHEN DOING A NEW ALLOCATION.

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

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)
SSMVSE14003	Provide SET IGGCAT=xx to dynamically change catalog options	Waiting for response (Orlando 2015) Waiting for response (Pittsburgh 2014) Open for discussion (Anaheim 2014)
SSMVSS14002	Change DEFrag TYPRUN=NORUN into separate ADRDSSU command	Waiting for response (Orlando 2015) Waiting for response (Seattle 2015) Open for discussion (Pittsburgh 2014)
SSMVSS14007	Track VTOC and VTOC index updates	Waiting for response (Orlando 2015) Waiting for response (Seattle 2015) Open for discussion (Pittsburgh 2014)



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



SHARE Requirements Summary

- Existing requirements with updated responses

Requirement #	Title	Status
SSMVSS13003 / RFE51803	Comprehensive Programming Environment Documentation	Uncommitted Candidate (Seattle 2015) Under Consideration (Pittsburgh 2014) Resolver Assigned (Anaheim 2014) Open for discussion (Boston 2013)
SSMVSS13010 / RFE70491	Provide 'MOVE' option to IEBCOPY	Uncommitted Candidate (Orlando 2015) Waiting for response (Seattle 2015) Waiting for response (Pittsburgh 2014) Open for discussion (Anaheim 2014)
SSMVSS14001 / RFE70492	HSM - Prioritize On Demand Migrations	Uncommitted Candidate (Orlando 2015) Waiting for response (Seattle 2015) Open for discussion (Pittsburgh 2014)
SSMVSS14005 / RFE55528	All FRBACKUP and FRRECOV Messages Should Be Sent Back to the User	Planned for Future Release (Orlando 2015) Waiting for response (Seattle 2015) Open for discussion (Pittsburgh 2014)
SSMVSS14006 / RFE58262	Prevent VTOC corruption	Planned for Future Release (Orlando 2015) Waiting for response (Seattle 2015) Open for discussion (Pittsburgh 2014)



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

SHARE Requirements Summary

- Requirements Available in V2.2

Requirement #	Title	Status
SSMVSE11033 / RFE51903	Add JCTJOBID and Sysplex to SMF types 14/15 (data set close)	Planned for Future Release (Orlando 2015) SHARE Top 75 Requirement
<u>GO6SMG91037</u> / RFE52495	DFP expand GDG limit beyond 255	Delivered (Orlando 2015) RC – Recognized (San Francisco 2013) Open for discussion (Anaheim 2012)
<u>SSMVSS14005</u> / RFE55528	All FRBACKUP and FRRECOV Messages Should Be Sent Back to the User	Planned for Future Release (Orlando 2015) Waiting for response (Seattle 2015) Open for discussion (Pittsburgh 2014)
<u>SSMVSS14006</u> / RFE58262	Prevent VTOC corruption	Planned for Future Release (Orlando 2015) Waiting for response (Seattle 2015) Open for discussion (Pittsburgh 2014)

SSMVSS12006



- **Tape Virtualization Engines should pre-cache during DR recovery**
 - **Response: SUG – Suggestion (Boston 2013) – responded SUG 11/20/12; no RFE**
 - *RFE - Uncommitted Candidate*
 - 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. Timeframe for delivery is TBD.*
 - **IBM Tape Analysis Tools** has a pre-stage tool that can bring volumes back efficiently, so it may be possible to exploit this tool. The tool is free and is available through the following link):
 - Tape Tools FTP site: <ftp://public.dhe.ibm.com/storage/tapetool/>
 - **ibmtools.txt** gives installation instructions; **updates.txt** lists updates and fixes
- **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.

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



SSMVSS13010

■ Provide 'MOVE' option to IEBCOPY

– **Response: Uncommitted Candidate (Orlando 2015)**

- *IBM Data Set Commander for z/OS Version 8 Release 1 (formally ISPF Productivity Tool for z/OS) does provide this capability for PDSs. See Data Set Commander for z/OS V8R1 Users' Guide ([SC19-4207-00](#)) for more details.*

■ 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



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

SSMVSE14003

- **Provide SET IGGCAT=xx to dynamically change catalog options**
 - **Waiting for Response (Orlando 2015); cannot find in RFE**
 - *RFE - Uncommitted Candidate*
 - *Comments from development: It is already on the list of things to look at for future releases. It might not be possible to do all parms do to control block structures, however, most could be changed. Each of the parms in IGGCATxx has a different scope and some could be implemented without a restart while others would be more difficult.*
- **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

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



SSMVSS14001

■ HSM - Prioritize On Demand Migrations

– **Uncommitted Candidate (Orlando 2015)**

- Comments from development: *Definitely doable and can go onto HSM backlog of small usability improvements.*

■ 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.



SSMVSS14002

- **Change DEFrag TYPRUN=NORUN into separate ADDRSSU command**
 - **Waiting for Response (Orlando 2015); cannot find in RFE**
 - *RFE - Uncommitted Candidate*
 - Comments from development: *The overhead of producing a new command is pretty large, however it is not complex code.*
- **Description**
 - The ADDRSSU 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, ADDRSSU 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.

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



SSMVSS14005

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

- **Planned for Future Release (Orlando 2015)**
 - Add SHARE as an interested party to the existing RFE55528
 - Comments from development: Available in V2.2

■ 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.

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



SSMVSS14006

■ Prevent VTOC corruption

– **Planned for Future Release (Orlando 2015)**

- *Add SHARE as an interested party to the existing RFE58262*
- *Comments from development: Already under consideration. IBM development understands that if the VTOC and INDEX structures get out of sync, this could cause significant problems including data loss. That is the reason why all DADSM/CVAF functions are designed to specifically execute in a sequence as to not cause this type of error, even if execution is interrupted. It is also not possible without significant performance loss, to compare the VTOC to the INDEX for all datasets if an update is being done by CVAF.*
- *Considering a few options:*
 1. *Produce a better diagnostic footprint by expanding and improving our SMF record for DADSM functions. This would help with first time data capture without requiring the DADSM CTRACE facility to be running.*
 2. *Design a new CVAFDIR interface that allows writers to only update certain fields in the VTOC DSCB records. Note: this would require vendors to rewrite their programs to use such a new interface.*
 3. *Cutting SMF records for updates to VTOC and Index*

■ 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.



SSMVSS14007

■ Track VTOC and VTOC index updates

- **Waiting for Response (Seattle 2015); cannot find in RFE**
- *RFE - Planned for Future Release*
- Comments from development: *Already under consideration*
 - *Initial thoughts include:*
 - *Cutting SMF records for updates to VTOC and Index*
 - *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.

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



SSMVSS13003

■ Comprehensive Programming Environment Documentation

- **Updated Response: Uncommitted Candidate**
- Comments from development: *Already under consideration by Development and ID and looking to focus on the DFSMS Macro Instructions for Data Sets. If there are other manuals which need attention, it would be good to start building that list for future updates.*

■ 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



GO6SMG91037

- **DFP expand GDG limit beyond 255**
 - **Delivered (Orlando 2015)**
- **Description**
 - Many applications require GDS(s) for a year or more. The current limitation of 255 for GDG does not support this need. This causes confusion when attempting to go back more than the 255 generations to decide which GDS(s) would be required to input into a process.
- **Benefit:**
 - Considerable time and effort on the part of analysts and schedulers could be saved by reducing confusion and time in the attempt to decide which GDG's are required for processing. Additionally, cost associated to reruns, related to this process would be eliminated.
- **Solution**
 - Raise the GDG limit to the max possible for the size of the field. This will provide the desired flexibility.
- **Discussion:**
 - I agree with this requirement, and will vote for it, but as for the 'year or more' I'd rather see implementation of something that allowed a date or time to be automatically substituted instead of a generalized generation. When I get some time I will try to write up that requirement.



Backup

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

IBM Tape Analysis Tools

IBM Tape Analysis Tools Overview

Nov 14, 2012

<ftp://ftp.software.ibm.com/storage/tapetool>

<=== as IBMTOOLS.EXE

Tool	Major Use	Benefit	Inputs	Outputs	Product pre-reqs
GETVOLS	Get volsers from list of dsns	Automate input to PRESTAGE	CA1, TLMS, RMM, ZARA, CTLT	volsers for requested dsns	none
GRPDSN	Genericize dsname lists	Speeds up the process of genericizing dsname lists which makes tape study more accurate	Dsname list	Genericized dsname list	none
IOSTATS	Report job elapsed times	Show runtime improvements	SMF 30 records	Job step detail reporting	none
LASTLIST	Report volsers last referenced between SDATE/EDATE	Determine how many volsers were last referenced during a period	CA1, TLMS, RMM, ZARA, CTLT	List of volsers and dsnames	none
LIBMANGR*	3494 library contents management	Allows user to identify and eject non mounted volumes	TCDB & CA1, TLMS, RMM, ZARA, CTLT	List of volumes not ref'd in N days, eject list, detail contents rpt	none
MOUNTMON*	Monitor mount pending and volume allocations	Determine accurate mount times and concurrent drive allocations	Samples tape UCBs	detail, summary, distribution, hourly, TGROUP, system reporting	none
OFFSITE	Identifies data sets sent off-site	Creates filter list to separate off-site workload	CA1, TLMS, RMM, ZARA, CTLT	Report plus disk file filter list of off-site dsnames	none
ORPHANS	Identify orphan data sets in Tape Management Catalog	Clean up tool	CA1, TLMS, RMM, ZARA, CTLT	Listing file showing all multi occurrence GDGs that have not been created in the last nn days.	none
PRESTAGE	Recall lvols to VTS	Ordered and efficient	BVIR VOLUME MAP	Jobs submitted to recall lvols	none

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

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

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.