#SHAREorg

SHARE

# IBM Responses to Requirements Submitted by the SHARE MVSS Project

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

August 12, 2013
Session 13774

SHARE in Boston

# Disclaimer

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

# Agenda

- IBM Responses to SHARE Requirements
  - Requirements Waiting for Responses
  - Requirements Open for Discussion
  - Existing requirements with updated responses
  - Requirements Available in 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)**<br>SUG – Suggestion (San Fran 2013)<br>Open for discussion (Anaheim 2012) |

# 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 (previously reviewed)

| Requirement # | Title | Status |
|---|---|---|
| SSMVSS13001 | DFHSM should provide a REXX interface | **Open for discussion (San Fran 2013)**<br>Cannot find in FITS – was this submitted? |
| SSMVSS10002 | Excessive/invalid HSM RECALL commands generated by IKJEFT01 | **More discussion needed (Anaheim 2012)**<br>Returned for clarification (Atlanta 2012)<br>AK – Acknowledged (Anaheim 2011) |
| SSMVSS11003 | DFSMS: Change How System Managed Buffering Handles Empty Datasets | **More discussion needed (Anaheim 2012)**<br>**Note:** Cannot find in FITS – need to resubmit.<br>Open for discussion (Anaheim 2011) |
| SSMVSS11005 | Increase the maximum VSAM Control Interval size (CISZ) | **More discussion needed (Anaheim 2012)**<br>**Note:** Cannot find in FITS – may need to resubmit.<br>Returned for clarification (Atlanta 2012)<br>Open for discussion (Orlando 2011) |

# SSMVSS13001

- **DFHSM should provide a REXX interface**
  - **Open for Discussion (San Francisco 2013)**
  - **Note: Cannot find in FITS. Was this submitted from San Francisco?**
- **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

# SSMVSS10002

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

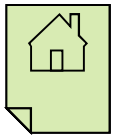

# SSMVSS11003

- **DFSMS: Change How System Managed Buffering Handles Empty Datasets**
  - **Open for discussion**
    - VSAM is really WAD, however, its not the best design! There is a difference in the way NSR and LSR handle SEQ processing, LSR requires a POINT and NSR does not. COBOL is not issuing a POINT since its a NSR user. When SMB is used to convert it to use LSR under the covers, a POINT it now required.
    - **Note:** Cannot find in FITS – may need to be resubmitted. IBM recommendation is to resubmit the requirement to DFSMS, and we'll reassign the FITS requirement to COBOL.
- **Description**
  - Currently, if a program opens a VSAM data set with an ACB that specifies MACRF=DIR (which is what COBOL will specify if ACCESS IS RANDOM is coded on the File Definition) and the HURBA is zero, the open will fail if System Managed Buffering (SMB) is invoked. SMB selects ACCBIAS=DO, and that does not work for an empty data set. SMB should check the HURBA, and if it is zero ACCBIAS=DW should be selected instead.
- **Benefit:**
  - The application can open a data set that has a HURBA=0 if 1) the program is changed to specify MACRF=SEQ (or any combination other than MACRF=DIR) in Assembler or ACCESS IS DYNAMIC in COBOL, or 2) codes ACCBIAS=DW in the JCL on the DD statement for the data set. The first option is actually the correct fix, but in either case requires the application to make a change to a program or job that possibly has not changed in a number of years. If there are a lot of programs coded this way, the number of changes could be substantial. Making this change in System Managed Buffering would fix this coding deficiency without any adverse affects, and it would allow customers to continue rolling out SMB without fear of causing job failures.

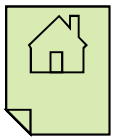

# SSMVSS11005

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

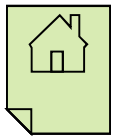

# SHARE Requirements Summary

- Open for Discussion (new)

| Requirement # | Title | Status |
|---|---|---|
| SSMVSS13002 | Allow Override of Messages to TSO Session | **Open for discussion (Boston 2013)** Open for discussion (San Fran 2013) |
| SSMVSS13003 | Comprehensive Programming Environment Documentation | **Open for discussion (Boston 2013)** |
| SSMVSS13004 | Add the ability to add a date and/or timestamp to a dataset name | **Open for discussion (Boston 2013)** |

# SSMVSS13002

- **Allow Override of Messages to TSO Session**
  - **Open for Discussion (Boston 2013)**
- **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

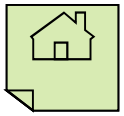

# SSMVSS13003

- **Comprehensive Programming Environment Documentation**
  - **Open for Discussion (Boston 2013)**
- **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**
  - **Open for Discussion (Boston 2013)**
- **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 generatation 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

# GO6SMG91037

- **DFP expand GDG limit beyond 255**
  - **Response: RC – Recognized (San Francisco 2013)**
- **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. Additionaly, 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.

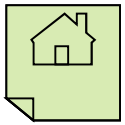

# SHARE Requirements Summary

- Existing requirements with updated responses

| Requirement # | Title | Status |
|---|---|---|
| SSMVSS12003 | REPRO MERGECAT Output should be condensed | **RC – Recognized (Boston 2013)**<br>AK – Acknowledged (San Fran 2013)<br>Open for discussion (Anaheim 2012) |
| SSMVSS12007 | VTS systems need to supply mount statistics per logical volume | **RC – Recognized (Boston 2013)**<br>RC – Recognized (San Fran 2013)<br>Open for discussion (Anaheim 2012) |
| SSMVSS11001 | Force System Determined Blocksize (SDB) | **RJ – Rejected (Boston 2013)**<br>RJ – Rejected (San Francisco 2013)<br>RJ – Rejected (Anaheim 2012)<br>Waiting for response (Atlanta 2012)<br>Open for discussion (Anaheim 2011) |

# SSMVSS12003

- **REPRO MERGECAT Output should be condensed**
  - **RC – Recognized (Boston 2013)**
- **Description**
  - REPRO MERGECAT currently produces up to 6 lines for each entry that is processed: one blank line, one or two for the IDC0639I message(Sphere conversion started), a blank line, and one or two for the IDC01402I (Sphere conversion ended). The one or two lines depend it seems on the length of the dataset name. This should be condensed to one line per entry, with indicators of what worked (or didn't).
- **Benefit:**
  - Improve the usability of the report when searching for exceptions.
- **Impact:**
  - If no solution is provided, customers and their businesses will continue to waste time when reviewing these listings.
- **Solution**
  - Suggestion: the one line should contain a message ID, a blank, a dataset name, and indicators for 'conversion started' 'conversion ended' and 'error encountered'. Error messages to follow the entry (if any), possibly indented.

# SSMVSS12007

- **VTS systems need to supply mount statistics per logical volume**
  - **RC – Recognized (Boston 2013)**
    - Comments from development: *Possible solution is to enhance the SMF21 record with additional statistics from the library. The SMF21 record is cut when we demount a volume. Requires changes in microcode and in host software, and timeframe for delivery is TBD.*
- **Description**
  - While the performance of logical volume mounts in a Virtualization Engine is not necessarily controllable, nor is it necessarily desirable to control it, we still need the statistics related to mounting logical volumes - including, but not limited to, the time to mount, the physical cartridge mounted if a recall was necessary, the mount time of the physical cartridge if necessary, indicators of whether the physical volume was already mounted requiring a wait for this logical volume mount request, or whether action was necessary to make space in the disk cache for a mount request. I'm sure there may be others.
- **Benefit:**
  - These statistics are necessary to answer performance / elapsed time questions for systems using the VE - for example, increase in elapsed time of a critical batch job may be questioned; we currently can say 'the tape mount took some time' but we cannot directly answer why.
- **Solution**
  - Enhance Allocation or O/C/EOV (whichever applies) to record tape mount time, have VE send the statistics record as part of the 'drive is ready, tape is mounted' completion and add that information to the mount record.

# SSMVSS11001

- **Force System Determined Blocksize (SDB)**
  - **Response: RJ – Rejected (San Francisco 2013)**
    - If the application has hard coded the dcb blksize in the dcb then we shouldn't be ignoring it. The application could be dependent upon that and it could cause problems if we changed this 45+ year precedent. About the only time this would be acceptable is if the customer was writing a new application that understood this was happening or understood what their existing application required. But in the case of an existing application, most customers probably couldn't tell you what the application required and may not even have the source.
    - Concern is system programmers will get into trouble trying to force a blksize for those apps that hard code blksize in the DCB, and the step would be to find those apps.
    - Would it be useful to customer is IBM provided a z/OS utility program to analyze SMF type 15 records to determine which ones have a value for BLKSIZE (block size) that differs from what the system would have calculated as the optimal value?
- **Description**
  - The way the Force System Determined Blocksize option of the SMS Data Class as introduced in z/OS Release 10, when set to Y, creates the belief that SDB will be forced for a given data set. This is true during the creation of the data set only. Since most user's create and then use (open) the data set, the Force System Determined Blocksize option does not work when an application program provides a blocksize value in the DCB when issuing the Open macro.
- **Benefit:**
  - Since most files are allocated and opened within a step, this will provide a SDB as requested in the dataclas regardless of the DCB parm
- **Solution:**
  - This REQUEST is being submitted as a requirement to DFSMSdfp to provide the function to actually force SBD for data sets within a Data Class for the life of the data set regardless of the presence of any coded blksize value in the DCB JCL parm or program.

# SHARE Requirements Summary

- **Requirements To Be Available in V2.1**

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

# 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: AC – Accepted (San Francisco 2013)**
- **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

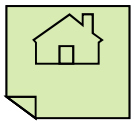

# SSMVSS10003

- **HSM: Make TAPECOPY restartable after a tape takeaway**
  - **Response: AC – Accepted (San Francisco 2013)**
- **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 (MIPâs, tape drives, and time), and it elongates the time that a primary tape does not have a copy.
- **Solution:**
  - None listed

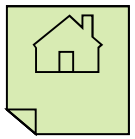

# SSMVSS01007

- **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 imput 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.

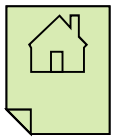

# SSMVSS08002

- **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.

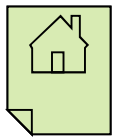

# SSMVSS07002

- DFSMSdss: Do not reset DS1DSCHA bits on RESTORE FULL
  - **Response: SUG – Suggestion (San Diego 2007)**
- **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. .

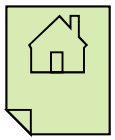

# System z Social Media Channels

- Top Facebook pages related to System z:
  - IBM System z
  - IBM Academic Initiative System z
  - IBM Master the Mainframe Contest
  - IBM Destination z
  - Millennial Mainframer
  - IBM Smarter Computing
- Top LinkedIn groups related to System z:
  - System z Advocates
  - SAP on System z
  - IBM Mainframe- Unofficial Group
  - IBM System z Events
  - Mainframe Experts Network
  - System z Linux
  - Enterprise Systems
  - Mainframe Security Gurus
- Twitter profiles related to System z:
  - IBM System z
  - IBM System z Events
  - IBM DB2 on System z
  - Millennial Mainframer
  - Destination z
  - IBM Smarter Computing
- YouTube accounts related to System z:
  - IBM System z
  - Destination z
  - IBM Smarter Computing
- Top System z blogs to check out:
  - Mainframe Insights
  - Smarter Computing
  - Millennial Mainframer
  - Mainframe & Hybrid Computing
  - The Mainframe Blog
  - Mainframe Watch Belgium
  - Mainframe Update
  - Enterprise Systems Media Blog
  - Dancing Dinosaur
  - DB2 for z/OS
  - IBM Destination z
  - DB2utor

# Trademarks and Disclaimers

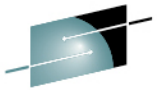

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

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