



## What's New in DFSMSdss

Lisa Gundy / Corey Smeaton  
IBM

August 15, 2013  
Session #14138

[lisat@us.ibm.com](mailto:lisat@us.ibm.com) / [csmeato@us.ibm.com](mailto:csmeato@us.ibm.com)



# Legal Disclaimer

## NOTICES AND DISCLAIMERS

Copyright © 2013 by International Business Machines Corporation.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) described herein at any time without notice.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Consult your local IBM representative or IBM Business Partner for information about the product and services available in your area.

Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

# Legal Disclaimer

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not necessarily tested those products in connection with this publication and cannot confirm the accuracy of 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.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

# Trademarks



The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

BookManager*	Enterprise Storage Server*	IP PrintWay	RMF
CICS*	ES/9000*	Language Environment*	S/370
DB2*	FlashCopy*	Lotus*	S/390*
DB2 Universal Database	GDPS*	Multiprise*	Tivoli*
developerWorks*	HiperSockets	MVS	TotalStorage*
DFSMSdfp	IBM*	Notes*	WebSphere*
DFSMSdss	IBM eServer	OS/390*	z/Architecture
DFSMShsm	IBM e(logo)server*	Parallel Sysplex*	z/OS*
DFSMSrmm	IBM logo*	RACF*	zSeries*
DFSORT	IMS	RAMAC*	
Domino	InfoPrint*		

Intel is a trademark of the Intel Corporation in the United States and other countries.

Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group in the United States and other countries.

\* All other products may be trademarks or registered trademarks of their respective companies.

## **Notes:**

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

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.

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 subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

# Agenda


- **z/OS V2.1**
  - Storage Tiers
  - RLS for Catalogs
  - Reset during Restore
  - Reset for mounted zFSs
  - Physical DS Rename and Replace
  - Physical DS Alternate SMS Volume
  - FlashCopy for Extended Format Sequential DS
  - FlashCopy Consistency Group Verification
  - GDG Support for PDSE
  - ACS Variable for EAV
- **Recent DFSMSdss APARs**
  - EAV:
    - OA42848
  - FlashCopy Scalability:
    - OA42214 OA42513  
OA42210

# z/OS V2.1

## Storage Tiers




- A storage tier is a **class** of devices that has a defined set of performance, availability, accessibility and capacity characteristics
- HSM Primary Space Management will drive class transitions within Ln tiers
  - Example of tier characteristics:
    - L0 consists of mirrored SSD
    - L1 consists of non-mirrored SSD
    - L2 consists of mirrored Fibre Channel
    - L3 consists of non-mirrored Fibre Channel
    - L4 consists of SATA
    - Ln devices are directly accessible by the user (DASD)



# z/OS V2.1

## Storage Tiers

- Each tier would be assigned a storage class and storage group through the ACS routines
  - ACS routines are used to direct data to a particular tier
  - ACS routines do not ‘understand’ the tier characteristics
    - Left to the user to assign accordingly
- Class transition attributes will consist of
  - Time since creation
  - Time since last used
  - Periodic




# z/OS V2.1

## Storage Tiers



- DSS will drive moving the data between tiers
  - COPY with DELETE
  - Will serialize the data set
  - If serialization cannot be obtained (e.g., data set open to an application), DSS will close the data set, move it, and reopen it
    - Done for DB2, CICS, and zFS data sets
    - Management class attribute will determine which applications DSS is to drive to close/open the data set
    - A new user exit option also added
    - Coexistence PTFs available
- See session #14125 for details





# z/OS V2.1

## Storage Tiers



- DSS Closing/Opening CICS data set - ADR568I

```
COPY DS(INC(C9SDSS.CICSTS.V4R1.ENT1)) -
  STORCLAS(SC9TCICS) -
  DELETE -
  BYPASSACS(*) -
  TOLERATE(ENQF) -
  ALldata(*) -
  ALLEXCP -
  SPHERE -
  ADMIN
DR101I (R/I)-RI01 (01), TASKID 001 HAS BEEN ASSIGNED TO COMMAND 'COPY '
DR109I (R/I)-RI01 (01), 2011.255 15:04:58 INITIAL SCAN OF USER CONTROL STATEMENTS COMPLETED
DR050I (001)-PRIME(01), DFSMSDSS INVOKED VIA APPLICATION INTERFACE
DR016I (001)-PRIME(01), RACF LOGGING OPTION IN EFFECT FOR THIS TASK
DR006I (001)-STEND(01), 2011.255 15:04:59 EXECUTION BEGINS
DR568I (001)-EXCIC(01), INVOCATION OF CICS INTERFACES BEGIN
      ***** DFSMSdss EXCI Client Program (ADREXCIC) *****
      *   Input Parameters: CICS CLOS APPLID=CICSIVP1 DSN=C9SDSS.CICSTS.V4R1.ENT1
      *   Linking to CICS EXCI server routine ADREXCIS
      *   The EXEC CICS server request completed, data set is CLOSED
      ***** End of DFSMSdss EXCI Client Program (ADREXCIC) *****
DR711I (001)-NEWDS(01), DATA SET C9SDSS.CICSTS.V4R1.ENT1 HAS BEEN ALLOCATED USING STORCLAS SC9TCICS, NO DATACLAS, AND NO MGMTCLAS
DR806I (001)-TOMI (03), DATA SET C9SDSS.CICSTS.V4R1.ENT1 COPIED USING A FAST REPLICATION FUNCTION
DR431I (001)-XVSAM(01), DATA SET C9SDSS.CICSTS.V4R1.ENT1 IN CATALOG TEST.CAT.C9SDSS HAS BEEN DELETED
DR568I (001)-EXCIC(01), INVOCATION OF CICS INTERFACES BEGIN
      ***** DFSMSdss EXCI Client Program (ADREXCIC) *****
      *   Input Parameters: CICS OPEN APPLID=CICSIVP1 DSN=C9SDSS.CICSTS.V4R1.ENT1
      *   Linking to CICS EXCI server routine ADREXCIS
      *   The EXEC CICS server request completed, data set is OPENED
      ***** End of DFSMSdss EXCI Client Program (ADREXCIC) *****
DR801I (001)-DDDS (01), 2011.255 15:05:04 DATA SET FILTERING IS COMPLETE. 1 OF 1 DATA SETS WERE SELECTED: 0 FAILED SERIALIZATION
AND 0 FAILED FOR OTHER REASONS
DR454I (001)-DDDS (01), THE FOLLOWING DATA SETS WERE SUCCESSFULLY PROCESSED
      C9SDSS.CICSTS.V4R1.ENT1
```

# z/OS V2.1

## RLS for Catalogs



- Why RLS for User Catalogs?
  - Performance
    - Contention on SYSIGGV2 bcsname when updating catalogs
  - Availability
    - Need to split for contention issues
    - Lack of availability when splitting and recovering
  - Integrity
    - Catalog damage due to utility updates while CAS is updating
    - Lack of sysplex control of serializing catalogs
  - Recovery
    - Difficult and error prone process

# z/OS V2.1

## RLS for Catalogs



- Implementation
  - New parameters added for DEFINE, ALTER
    - LOG(NONE) parameter allows the catalog to become eligible for RLS access
      - DEFINE
      - ALTER
  - Many new MODIFY commands introduced

# z/OS V2.1

## RLS for Catalogs



- DSS serialization change for catalogs
  - SYSIGGV2 resource not obtained for catalogs open in RLS
    - SMSVSAM will obtain SYSIGGV2 bcsname SHARED while catalog is opened for RLS access
      - Ensures integrity from programs relying on SYSIGGV2
  - When restoring catalogs
    - If replacing existing catalog it must first be locked or suspended
    - BCSRECOVER(LOCK|SUSPEND) keyword
      - LOCK will invoke a sysplex wide close of the catalog and lock the catalog (failing new unauthorized requests), if the catalog is not already locked or suspended. If DSS locks the catalog, it will unlock it at the end of the restore
      - SUSPEND will invoke a sysplex wide close of the catalog and suspend new unauthorized requests in the client space if the catalog is not already locked or suspended. If DSS suspends requests, it will resume the requests at the end of the restore

# z/OS V2.1

## RLS for Catalogs



- When restoring catalogs (continued)
  - **Migration Action:** Existing catalog must be locked or suspended prior to invoking Restore (ADR439E issued)
    - Must specify either
      - BCSRECOVER(LOCK|SUSPEND) keyword on RESTORE
      - F CATALOG,RECOVER,LOCK|SUSPEND(bcsname) before initiating the RESTORE command
      - IGG.CATLOCK Facility Class Profile must be defined
- When copying or dumping catalogs
  - DSS will quiesce (QUICOPY) the catalog to ensure any updates being made are held until DSS ends the quiesce (QUIEND)
  - Ensures catalog is not being updated across the sysplex while moving or backing up the catalog

# z/OS V2.1

## RLS for Catalogs



- Coexistence PTFs
  - When attempting to COPY or DUMP a catalog open in RLS in V2.1 on a lower release error will be issued during filtering
  - ADR724E with return and reason codes indicating catalog is open in RLS on another system
    - VSAM open errors will also be issued to SYSLOG
- See session #14145 for details

## z/OS V2.1

# Reset during Restore



- Problem: DSS unconditionally resets the data-set-changed indicator (DS1DSCHA) for all data sets restored during full volume restore
  - Share Requirements SSMVSS07002, SSMVSS07009
  - 5 Marketing Requirements
    - MR0418071919, MR0302074136, MR0409076057, MR1119075914, MR0604021857
- Solution: New keyword RESET(YES|NO|DUMP) during full volume RESTORE
  - Allow users to specify how DSS should set the DS1DSCHA indicator for all data sets restored to the target volume

# z/OS V2.1

## Reset during Restore



- RESET(YES|NO|DUMP\*) keyword
  - RESET(DUMP) - default
    - If RESET is specified during DUMP FULL then the target will resemble volume **after** it was dumped
    - DS1DSCHA bits will be turned off for all data sets residing on volume
    - This applies only to volumes dumped with V2.1 support (if dumped prior, will default to RESET(NO))
  - RESET(YES)
    - DS1DSCHA bits will be turned off for all data sets residing on volume
    - This is behavior prior to V2.1
  - RESET(NO)
    - The target will return to the state **prior** to the dump
      - Original VTOC restored
      - DS1DSCHA bits for data sets on volume will remain untouched
    - \*If RESET is not specified during DUMP then this becomes default
- ★ New options also provided in ADRUFO to override keyword specification

HSM users: Please note new RECOVERRESET(YES|NO|DUMP) in V2.1



# z/OS V2.1

## Reset during Restore



- New RACF Facility Class for RESET keyword
  - Applies to both DUMP and RESTORE
  - The new RACF Facility Class profiles are
    - STGADMIN.ADR.DUMP.RESET for DUMP
    - STGADMIN.ADR.RESTORE.RESET.YES for RESTORE
  - When the corresponding profile is defined
    - Users without READ access to the corresponding RACF Facility Class profile will not be able to specify the RESET keyword
    - ADR707E will be issued

## z/OS V2.1

# Reset for mounted zFS data sets



- Problem: DSS RESET during logical DUMP is ignored for mounted zFS data sets
  - DSS only resets the DS1DSCHA bit if exclusive serialization can be obtained
  - For mounted zFS data sets, DSS invokes Unix System Services (USS) to quiesce the file system to take a consistent backup
  - Incremental backup was also not supported
  - New function will now allow DSS to reset DS1DSCHA bit for mounted zFS data sets while quiesced

## z/OS V2.1

# Physical DS Rename and Replace



- Problem: DSS does not allow renaming of VSAM data sets during physical DS COPY and RESTORE
  - Only supported for nonVSAM today
  - 7 Marketing Requirements
    - MR1011075120
    - MR030912475
    - MR0309124752
    - MR1011013327
    - MR1114053052
    - MR0819115958
    - MR0907116329
  - Gives users access to production data set and restored data set without having to restore entire volume
    - Ability to unload data from the backup version

# z/OS V2.1

## Physical DS Rename and Replace



- Solution:
  - New keyword RENAMEU will be supported for VSAM COPY and RESTORE
    - VSAM Alternate Indexes will continue to remain unsupported
  - RENAME will continue to remain unsupported for VSAM
  - Recataloging after DSS processing remains necessary
  - REPLACEUnconditional will also be supported for physical DS RESTORE
    - Physical DS COPY supports REPLACEU today

## z/OS V2.1

# Alternate Volume Selection for Physical DS



- Problem: Physical DS Copy and Restore only supports one output volume
  - Volume selection is not an option
  - If target does not have the space the operation fails
    - Scenarios:
      - COPY FULL from SRC to TGT, COPY DS from TGT to SRC
      - DUMP FULL from TGT to tape, RESTORE DS from tape to SRC
    - Space may no longer be available on SRC if data set size changed
      - Preallocated target may have been scratched/reallocated
- Solution: DSS adding a new user interaction module (UIM) exit
  - Physical DS Alternate SMS Volume (EIOPTION 30)
  - Added to ADREID0 mapping
  - Allows program to pass a list of volumes that DSS may use to attempt to allocate DS on, when necessary
    - Only called if original target volume could not be used

## z/OS V2.1

# FlashCopy Extended Format Sequential DS



- Problem: If you have a single-striped multivolume data set FlashCopy cannot be used
  - There is a suffix in each block that contains block number for volume
  - Same issue if source is single-striped single volume but target requires multiple volumes
  - DSS must use a utility to copy the data set (ADR918I RC 12)
- Solution: Support a new version for extended format sequential data sets called version 2
  - Volume boundary awareness is removed
  - A new PS\_EXT\_VERSION=1|2 keyword will be supported in IGDSMSxx member of PARMLIB
  - Newly created data sets will be the version type specified in PARMLIB
  - You can set it also via the SETSMS command with PS\_EXT\_VERSION(n)

## z/OS V2.1

# FlashCopy Extended Format Sequential DS



- DSS Changes
  - DSS will use FlashCopy only if input and output are version 2
    - V2.1 only
    - If source is V2 and target is V1 then FlashCopy will not be used
    - For new DSS allocations version type will be the version type of source
    - Preallocated version type will be preserved (for logical, not physical)
    - DSS allocations will not be controlled using PARMLIB
  - New **ADR918I** RC of **28** will be issued to indicate that the source extended format version type does not match that of the target
  - New **ADR918I** RC of **29** will be issued to indicate that the version 2 source extended format data set has an EOF on the non-last volume

## z/OS V2.1

# FlashCopy Extended Format Sequential DS



- The version type will reside in the NVR and will be displayed in IDCAMS LISTCAT and DCOLLECT
- Coexistence APARs are available for pre-V2.1 including:
  - COPY/DUMP/RESTORE of v2 on lower level releases are supported
  - Target COPY|RESTORE will be allocated as V1
  - Preallocated V2 will also be supported
  - If V2 is preallocated and must be scratched and reallocated, it will be reallocated as V1
  - FlashCopy of V2 single striped multi-volume will not be supported on lower level systems

★ Do not set version 2 until all sharing and backup systems are at this release or you install compatibility PTFs




## z/OS V2.1

# FlashCopy Consistency Group Verification



- Overview
  - FlashCopy S1 to T1 with FCCGFREEZE
    - Writes cannot proceed on S1
    - Any writes occurring on S2-S4 are not dependent writes
  - FlashCopy S2 to T2
    - Writes cannot proceed on S1 or S2
    - Any writes occurring on S3-S4 are not dependent writes
  - FlashCopy S3 to T3 and S4 to T4
  - T1-T4 contain a consistent copy
  - Issue CGCREATED to "UNFREEZE"
    - Specify FCCGVERFIY on S1
    - Writes may proceed on S1-S4



## z/OS V2.1

# FlashCopy Consistency Group Verification



- DSS will be modified to accept multiple volume serial numbers on the FCCGVERIFY keyword of the CGCREATED command
  - Today FCCGVERIFY allows one volsers
    - FCCGVERIFY during the CGCREATED command tells DSS to query the volume to see if it is still “FROZEN”
  - V2.1 will allow FCCGVERIFY to accept up to 255 volsers
    - Useful if the consistency group being formed spans many LSS's or subsystems that may not have the same timer
- DSS will verify all volumes specified within the FCCGVERIFY keyword whether any volume fails verification or not
  - A new bit will be available in ADRUFO to instruct DSS to stop verifying volumes specified within the FCCGVERIFY keyword once a volume has failed verification
    - Default is to verify all volumes specified


# z/OS V2.1

## GDG Support for PDSE



- Problem: Customers have expressed interest in being able to define a PDSE as a GDS
  - Sequential, Partitioned (PDS), and Direct Access data sets can be GDS's today
    - IDC3009I 48-152 when attempting to define a PDSE as a member of a GDG base
- Solution: PDSE can be defined as a GDS
  - Using DSNTYPE = LIBRARY or through DATACLAS
  - IDCAMS LISTCAT will display both GDG Status and DSNTYPE (SMS only)
  - DSS TGTGDS keyword for COPY and RESTORE is now supported for PDSE
    - ACTIVE, DEFERRED, ROLLEDOFF, SOURCE are options
    - Preallocated GDS status is preserved
    - Coexistence APAR available

## EAV Key Design Points



- Maintains 3390 track format
- **Track-managed space:** the area on an EAV located within the first 65,520 cylinders
  - Space is allocated in track and cylinder units
  - Storage for “small” data sets
- **Cylinder-managed space:** the area on an EAV located above the first 65,520 cylinders
  - Space is allocated in **multicylinder units (MCU)**
    - A fixed unit of space larger than a cylinder. (On an EAV MCU = 21 cylinders)
    - Requests rounded up to next MCU
  - Storage for “large” data sets
- Track-managed space comparable to same space on non-EAVs

## z/OS V2.1

### ACS Variable for EAV



- Problem: EATTR keyword is supported for JCL, dynamic allocations, AMS Define and DATACLAS, but not in ACS Routines
  - Does not allow SMS volume selection to select a volume for allocation based on EATTR specification
- Solution: DSS will pass EATTR specification to SMS for an allocation request if source EATTR=OPT
  - Will allow SMS to consider EATTR setting when deriving Storage Class and Storage Groups

# z/OS V2.1

## NEW DEBUG KEYWORD



- New DEBUG(SMSMSG(MINIMAL|SUMMARIZED|DETAILED)) Keyword
  - Can be specified during COPY, RESTORE, and CONVERTV
  - Instructs DSS to output SMS messages and ACS Write requests to the DSS job log
  - SMS messages will be under the new ADR803I message for successful allocations
  - ADR709E will continue to be issued for failures
  - New bits representing DEBUG SMSMSG options will also be added to ADRUFO

# Agenda

- **z/OS V2 R1**
  - Storage Tiers
  - RLS for Catalogs
  - Reset during Restore
  - Reset for mounted zFSs
  - Physical DS Rename
  - Physical DS Alternate SMS Volume
  - FlashCopy for Extended Format Sequential DS
  - FlashCopy Consistency Group Verification
  - PDSE for GDGs
  - ACS Variable for EAV
- **Recent DFSMSdss APARs**
  - EAV:
    - OA42848
  - FlashCopy Scalability:
    - OA42214 OA42513  
OA42210

## APARs of interest

### OA42848: Migration of nonVSAM DS to EAS

- Problem: nonVSAM data sets cannot be copied to EAS on EAV if source is defined with EATTR = NS or NO (new allocations)
  - Logical data set COPY operations: DFSMSDss uses EATTR attribute of source data set to determine EATTR attribute of the target.
    - Example, source has an EATTR value of NO, when the target data set is newly allocated, the target data set will result in the EATTR value of NO.
- Target nonVSAM data set would have to be preallocated in EAS with EATTR=OPT



## APARs of interest

### OA42848: Migration of nonVSAM DS to EAS

- Solution: Provide a patch byte to allow you to override the EATTR attribute of the target data set to EATTR=OPT for logical data set copy
  - The new patch byte can only be enabled using the SET PATCH command (dynamic)
  - Does not apply if the target data set is preallocated
    - Even though a preallocated data set can be scratched and reallocated, the preallocated EATTR attribute is preserved
  - Dynamic patch option only applies to logical data set copy for nonVSAM data sets
  - Please contact IBM Support to inquire about PTF availability
    - GA tentatively set for Oct 2013 but ++APARs are available
- ❖ Note: With this patch, all selected data sets of the Logical non-VSAM data set command will have their EATTR value overridden with the EATTR=OPT. Users may want to initially run with TYPRUN=NORUN to ensure they want to include the listed data sets as part of the migration to EATTR=OPT. If not, then specific data sets may be excluded.

# APARs of interest

## FlashCopy Scalability

- Problem:
  - SMS volume selection for target(s) of DSS dataset copy with fast replication (FlashCopy) selection.
    - Resources consumed (CPU, virtual storage, LOCAL lock contention) for candidate selection increases as number of volumes in storage group (SG) increases.
- Volume selection today
  - SMS invokes ANTRQST QFRVOLS
    - passes SG volume list to System Data Mover (SDM)
  - SDM eventually invokes Asynchronous Operation Manager (AOM) to perform query to storage controller
  - Each candidate volume is examined for eligibility:
    - same Storage Facility Image (SFI) as source data set's volume?
    - FlashCopy enabled?
    - number of existing relationships or already a target?
    - ...

# APARs of interest

## FlashCopy Scalability

- How it works today (continued)
  - For each candidate
    - SRB scheduled
    - two I/O requests issued
      - *Each I/O request separately invokes a device driver resulting in Getmain/Freemain of working storage, thus a LOCAL lock (used for all private area Virtual Storage Manager (VSM) serialization)*
    - I/O scheduled for all volumes at once causing IOS to allocate resources to handle the spike
  - Results interpreted and returned to SDM and SMS
  - SMS allocates a data set on the target volume selected from the list
- How does this affect z/OS?
  - High CPU usage in ANTMAIN address space
    - Caused by having a large number of volumes in every storage group eligible for allocation of the target data set
  - LOCAL Lock contention
  - Excessive Virtual Storage consumption

# APARs of interest - FlashCopy Scalability



- How do we fix this?
  - Reduce the size of the Candidate Volume Pool...
  - SMS filters out those candidates that are ineligible to participate in a Fast Replicate relation with the source volume
  - From remaining list, SMS selects a subset of volumes, and invokes QFRVOLS
  - Sample size default = 100 candidates for single volume request from one Storage Group
  - Sample size increased by 20 for each additional target volume needed
  - Sample size increased by 20 for each additional Storage Group
  - If no viable candidate is found, SMS repeats QFRVOLS request with remaining SG volumes



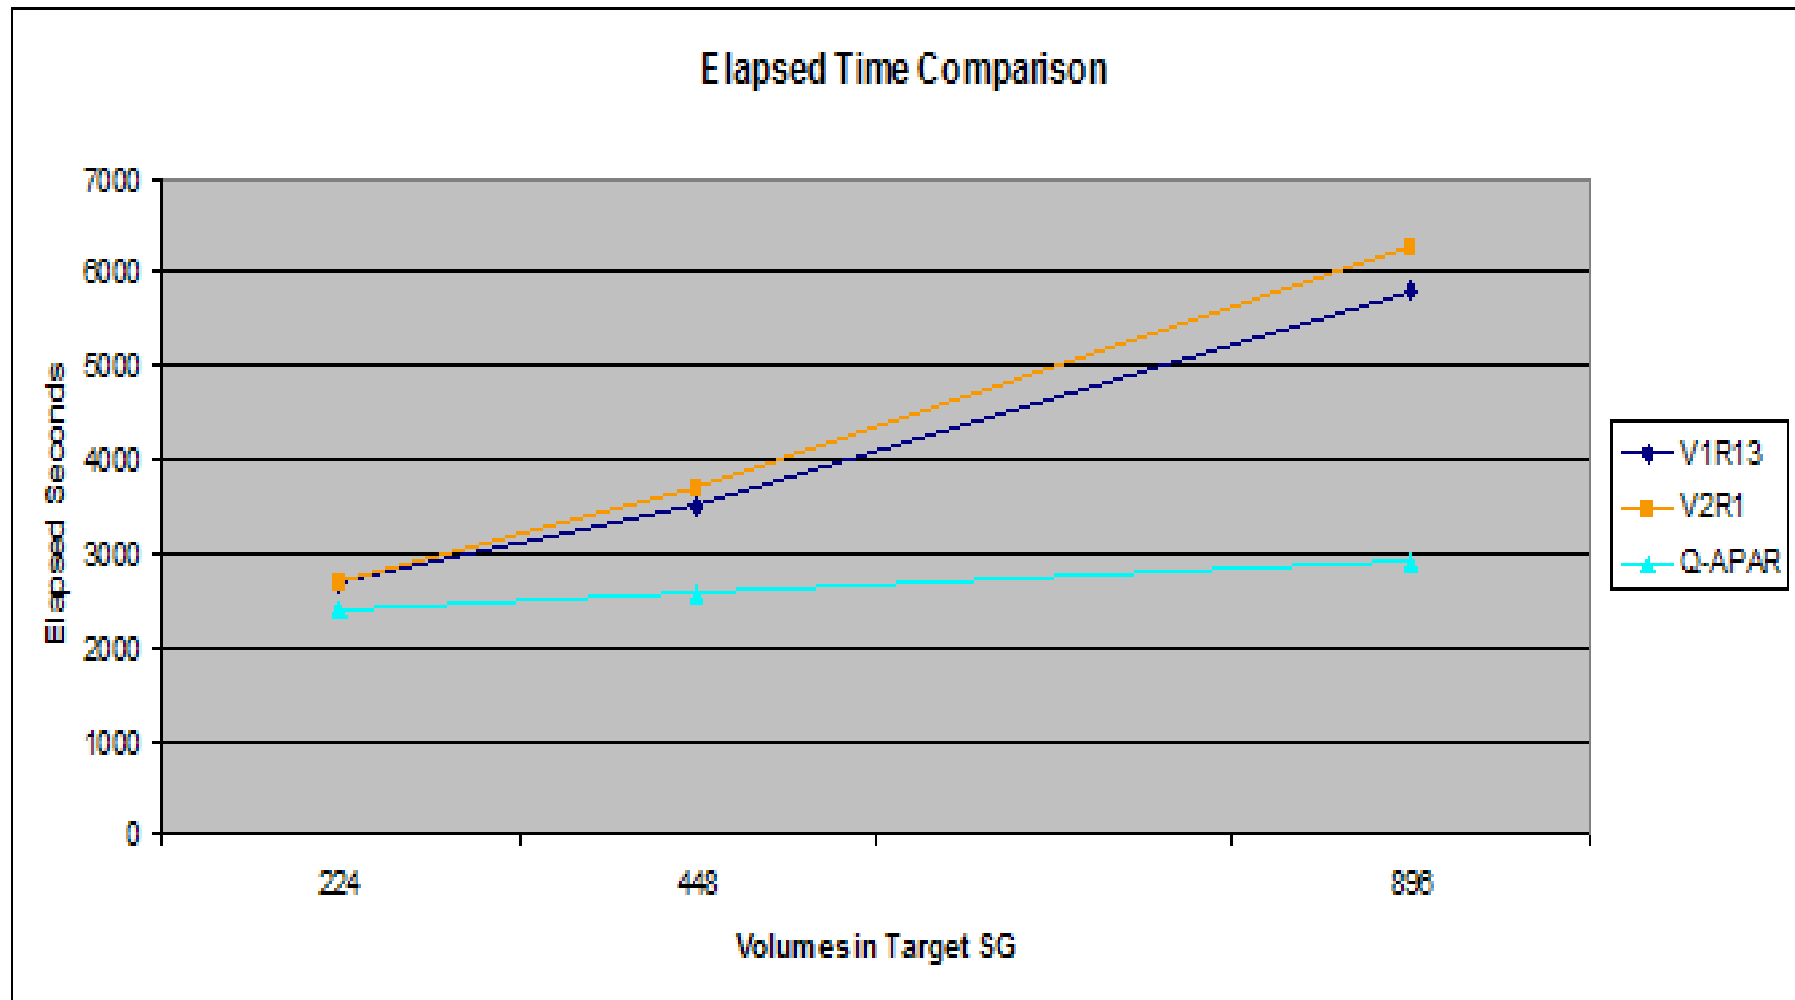
# APARs of interest - FlashCopy Scalability



- How do we fix this? (continued)
  - Throttle AOM QFRVOL Threads
    - Default to 64 SRBs per QFRVOLS request. As one SRB completes, reuse its storage for next candidate. Repeat until all candidates have been checked.
  - Reduce LOCAL lock contention
    - Replace Getmain/Freemain in QFRVOLS with CellPools.
      - *One CPOOL for SRB I/O, one for AOM Driver.*
      - *Old design obtained Local Lock 8 times per candidate volume*
      - *New design obtains Local lock 1 time per candidate volume*
  - Eliminate Read Device Characteristics I/O
  - Obtain device chars from internal control blocks



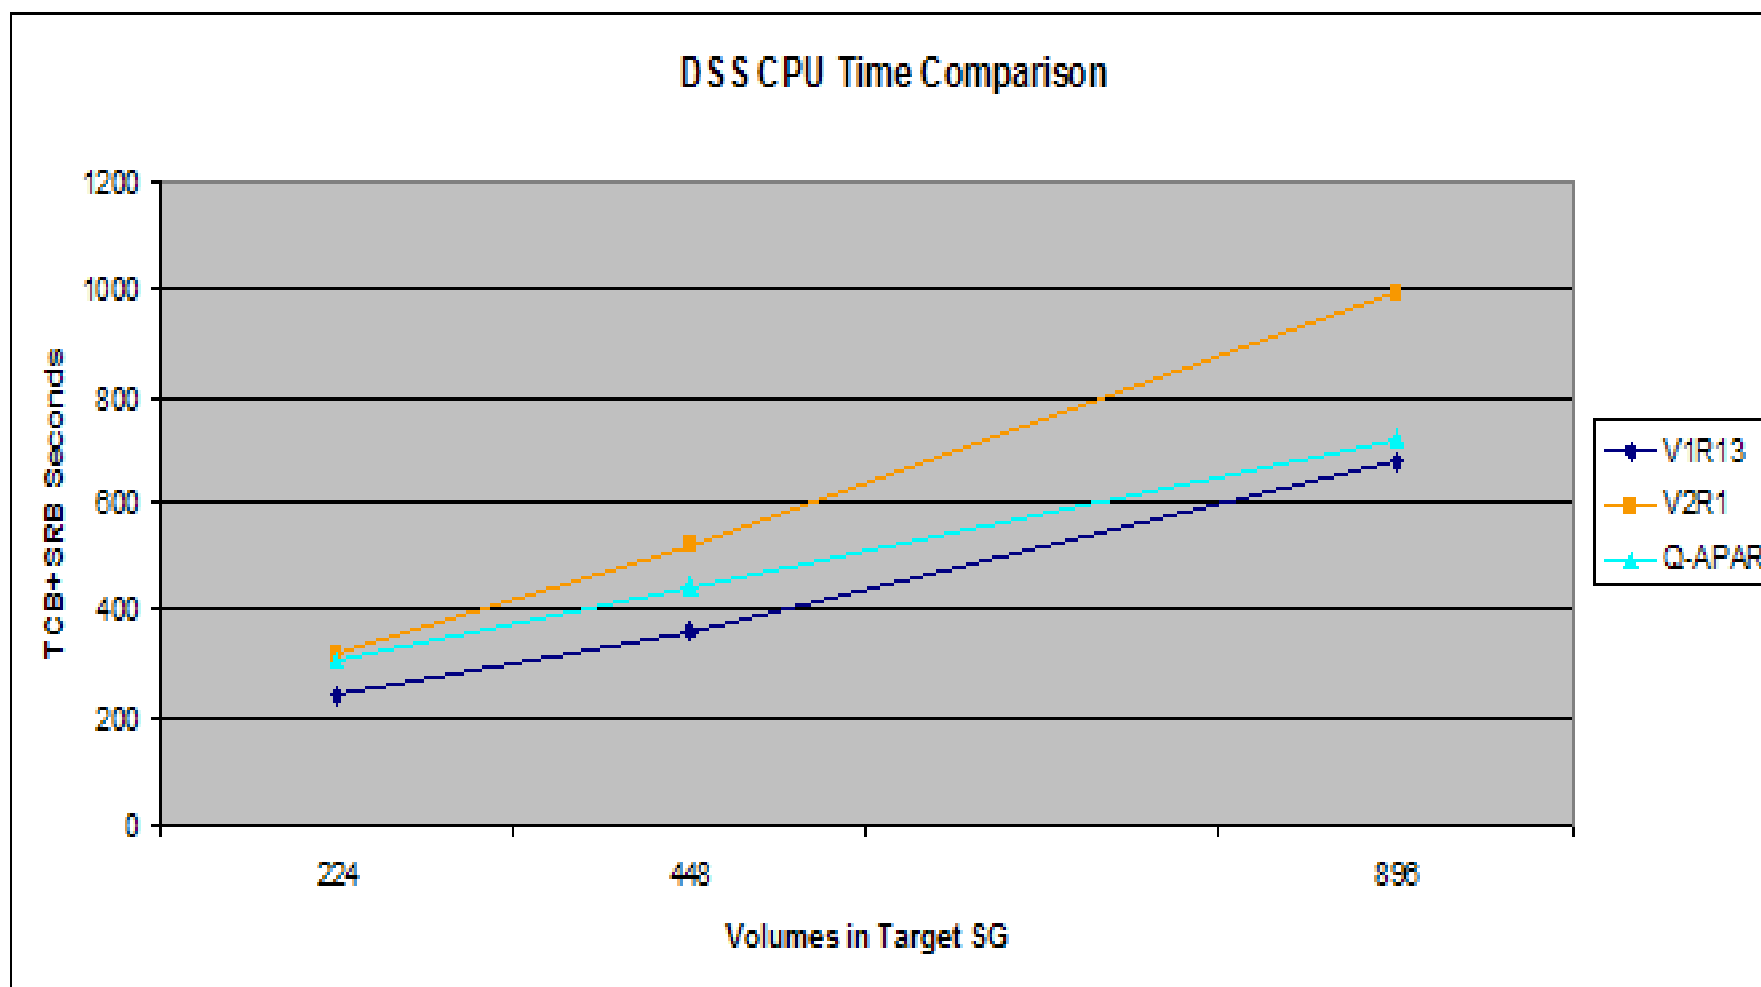
# APARs of interest - FlashCopy Scalability



Complete your sessions evaluation online at [SHARE.org/BostonEval](http://SHARE.org/BostonEval)




# APARs of interest -FlashCopy Scalability



Complete your sessions evaluation online at [SHARE.org/BostonEval](http://SHARE.org/BostonEval)



# APARs of interest - FlashCopy Scalability




Complete your sessions evaluation online at [SHARE.org/BostonEval](http://SHARE.org/BostonEval)





# APARs of interest - FlashCopy Scalability



Complete your sessions evaluation online at [SHARE.org/BostonEval](http://SHARE.org/BostonEval)



# APARs of interest

## FlashCopy Scalability

- APARs/PTFs
  - z/OS V2R1 only
  - AOM OA42214/OA42513 PTF UA69389
  - SMS OA42210 PTF UA69396

## Summary

- Significant number of new and enhanced functions
  - Emphasis on addressing customer requirements
  - Expanding the use of FlashCopy
  - New functions being delivered in z/OS V2.1 result in the largest DFSMSdss delivery ever in a single release
- Demonstrates IBM's commitment to continual improvement of the DFSMSdss product

# 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](#)



# THANK YOU!

## What's New in DFMSdss

Lisa Gundy / Corey Smeaton  
IBM

August 15, 2013  
Session #14138

[lisat@us.ibm.com](mailto:lisat@us.ibm.com) / [csmeato@us.ibm.com](mailto:csmeato@us.ibm.com)

