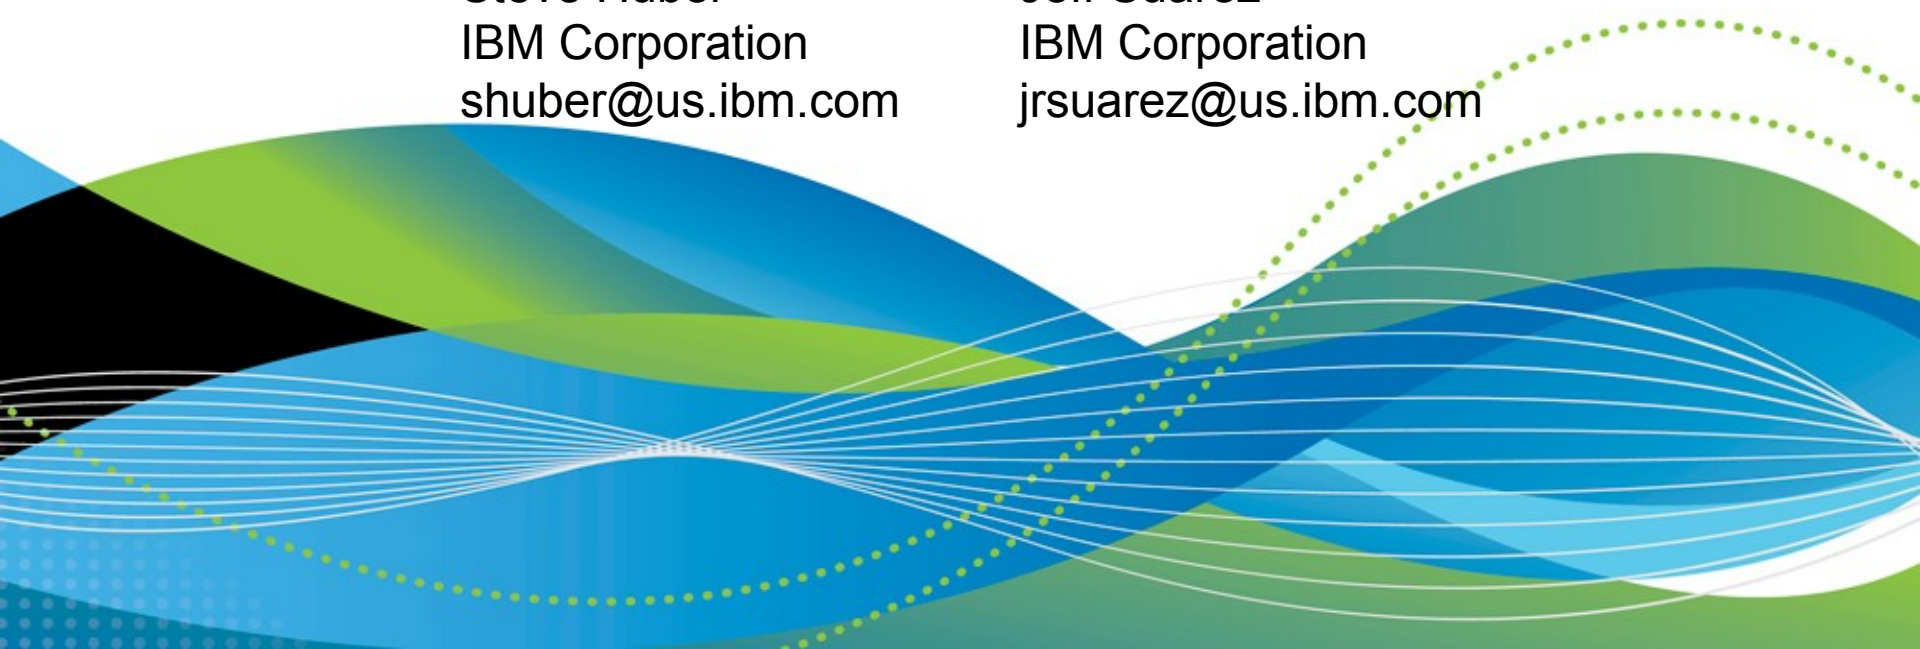


Making the Most of DFSMSdss and SMS: Hints, Tips, and Best Practices in your z/OS Environment

Session # 10954

Steve Huber
IBM Corporation
shuber@us.ibm.com

Jeff Suarez
IBM Corporation
jrswarez@us.ibm.com



Legal Disclaimer

NOTICES AND DISCLAIMERS

Copyright © 2012 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
DFSMSshsm	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

- **System Managed Storage Overview**
 - What is SMS
 - Benefits of SMS
 - SMS Configurations
 - ACS Routines
 - Best Practices
- **DFSMSdss Best Practices**
 - Converting nonSMS data to SMS without data movement
 - Keywords related to SMS processing
 - SG ACS and renaming data sets
 - Volume Selection
 - VSAM CA Size Requirements
- **FlashCopy Hints and Tips**

What is system managed storage?

- Lets the operating system take over storage management tasks
 - Data set allocation
 - Backup management
 - Space management
 - Availability management
- Reduces number of people needed to manage storage

What is SMS?

- SMS is an MVS subsystem
 - Manages the current storage management policy (active configuration)
 - Reduces end user data set creation & allocation complexity
 - Increases installation control of DASD, tape, and optical storage
- There is one SMS and one SMS address space per instance of MVS (z/OS)
- SMS runs in both the user's and the SMS address space

SMS Design Considerations

- Clearly separate the domains of users, data and storage media
- Introduce the role of storage administrator
- Preserve customer investment in JCL and other structures

Why should I use it?

- Reduce out of space abends (X37)
- Reduce device fragmentation
- Balance allocations across a pool of devices
- Improve storage utilization
- Help achieve device independence

Some Basic Terms

- SMS configuration
- Minimal configuration
- Base configuration information
- Storage group
- Storage class
- Management class
- Data class
- ACS routine

Some Basic Terms...

- SMS configuration
 - Minimal configuration
 - Base configuration information
 - Storage group
 - Storage class
 - Management class
 - Data class
 - ACS routine
- **A configuration is a storage management policy**
 - **It contains elements which define that policy:**
 - **Storage groups & volumes**
 - **Storage classes**
 - **Management classes**
 - **Data Classes**
 - **Automatic Class Selection (ACS) routines**
 - **Optical and tape libraries and drives**
 - **Aggregate groups**

Some Basic Terms...

- SMS configuration
- Minimal configuration
- Base configuration information
- Storage group
- Storage class
- Management class
- Data class
- ACS routine
- **Base configuration information**
- **One storage class definition**
- **One storage group with at least one volume**
- **A storage class ACS routine**
- **A storage group ACS routine**

Some Basic Terms...

- SMS configuration
- Minimal configuration
- Base Configuration Definition
- Storage group
- Storage class
- Management class
- Data class
- ACS routine
- **System & system group names**
- **Default management class**
- **Default unit**
- **Default device geometry**

Some Basic Terms...

- SMS configuration
- Minimal configuration
- Base configuration information
- Storage group
- Storage class
- Management class
- Data class
- ACS routine
- **Physical storage managed by SMS**
 - **Collection of DASD volumes**
 - **Volumes in tape libraries**
 - **Volumes in optical libraries**
 - **Virtual I/O storage**
- **Can be enabled, quiesced, quiesced new, disabled or disabled new**
- **Can be set to auto migrate, auto backup and/or auto dump**

Some Basic Terms...

- SMS configuration
- Minimal configuration
- Base configuration information
- Storage group
- Storage class
- Management class
- Data class
- ACS routine
- **Performance attributes**
 - **Direct & sequential millisecond response**
 - **Direct & sequential bias**
 - **Initial access response time**
- **Availability**
- **Accessibility**
- **Guaranteed space**
- **Guaranteed synchronous write**

Some Basic Terms...

- SMS configuration
- Minimal configuration
- Base configuration information
- Storage group
- Storage class
- Management class
- Data class
- ACS routine
- **Space management attributes**
 - Expiration & retention attributes
 - Migration attributes
 - GDG management attributes
- **Backup attributes**
 - Backup frequency
 - Backup versions
 - Backup retention
- **Class transition attributes**
- **Aggregate backup attributes**

Some Basic Terms...

- SMS configuration
- Minimal configuration
- Base configuration information
- Storage group
- Storage class
- Management class
- Data class
- ACS routine
- **RECORD or RECFM**
- **LRECL**
- **Space**
- **DSNTYPE**
- **Volume count**
- **VSAM attributes**
- **RETPD or EXPDT**
- **Compaction**

Some Basic Terms...

- SMS configuration
- Minimal configuration
- Base configuration information
- Storage group
- Storage class
- Management class
- Data class
- ACS routine
- **Used to determine SMS classes and storage groups**
- **Used for both data sets and objects**
- **Can override specifications of SMS classes and groups on:**
 - **JCL DD statements**
 - **Dynamic allocation requests**
 - **DFSMSdss COPY, RESTORE & CONVERTV**
 - **DFSMShsm RECALL & RECOVER**
 - **IDCAMS DEFINE, ALTER & IMPORT**
 - **OAM STORE, CHANGE & class transition**

SMS Best Practices

- ACS Routines
 - REMINDER - ACS processing is uncaptured time
 - Use FILTLISTs with wildcards
 - Don't use long FILTLISTs with complete DSN's
 - Use SELECT statements whenever possible
 - Allows you to exit the routine a little faster
- CDS's
 - Ensure CDSs have REUSE specified
 - Ensure CDSs have plenty of space
 - Separate the ACDS and COMMDS

SMS Best Practices (cont)

- Configuration
 - See OA33127 and apply the PTF
 - After R11, null volume entries are not dealt with well
 - II14602 describes how to ‘prevent’ OA33127’s loop

DFSMSDSS Keywords related to SMS

- **STORGRP** keyword
 - Specifies that all online volumes in the Storage Groups be dynamically allocated
 - Can specify up to 255 storage group names
 - Do not need to specify SELECTMULTI keyword
 - Catalog filtering is used to find data sets but DSS will only select data sets for volumes in the specified storage groups
 - May cause an increase in job run time because DSS has to read VTOCs (similar behavior if specifying LOGINDD and LOGINDY)
 - Can be used for COPY, DUMP, and RELEASE commands
 - Customers have used this to determine which volumes in a storage group data sets have been copied to using TYPRUN=NORUN

DFSMSDss Keywords related to SMS

- **STORCLAS** keyword
 - Specifies a storage class name for DFSMSDss to pass to ACS
 - Does not guarantee that the storage class specified will be assigned to the data set
 - Must specify **BYPASSACS** to guarantee specified **STORCLASS** is assigned to target
 - You can also use a guaranteed-space defined storage class to placed data sets on specified **OUTDD/OUTDYNAM** volumes
 - **SG ACS** must use **SC** to determine storage group
 - Example shows moving data sets from all storage classes, except two, into one storage class using **BY FILTERING**

```
COPY DATASET(INCLUDE(**) -  
BY(STORCLAS,NE,(SCNAME1,SCNAME2))) -  
STORCLAS(SCNAME3) BYPASSACS(**) DELETE
```

SG ACS and renaming data sets using DFSMSdss

- Using RENAMEU to have ACS select storage group without STORCLAS keyword using FILTLIST (name masking)

```
PROC STORCLASACS  
FILTLIST JOBC    INCLUDE ('** .DSTGT.**')
```

```
/*  Assigns storage class      */  
  WHEN (&DSN = &JOBC)  
  DO  
  SET &STORCLAS = 'DSTARGET'  
  EXIT CODE(0)
```

ADDRSSU job :

```
COPY DS(INC(PROD.DSSRC.DS01.XX1900)) -  
  RENAMEU((* .DSSRC.** , * .DSTGT.**))    <----- name assigned
```

DFSMSDss Keywords related to SMS

- MGMTCLAS keyword
 - Specifies a management class to replace the source management class to pass to ACS
 - Specifying NULLMGMTCLAS provides a null management class to the input of the ACS routines
 - Also does not propagate source management class to target
 - You can also do BY FILTERING by MGMTCLAS name

```
COPY DATASET(INCLUDE(**) BY(MGMTCLAS,EQ,  
  (MGTNAME1,MGTNAME2))) -  
  STORCLAS(SCNAME3) DELETE
```


DFSMSDss Keywords related to SMS

- **BYPASSACS** keyword
 - Way to force the specified **STORCLAS** and **MGMTCLAS** to be assigned to target
 - ACS routines are not invoked

```
COPY DATASET(INCLUDE(**)  
MGMTCLAS(MGNAME1)  
STORCLAS(SCNAME3)  
BYPASSACS(**)
```

DFSMSDSS Keywords related to SMS

- Using NULLSTORCLAS and BYPASSACS(**) keywords
 - Way to force the COPY and RESTORE of the data sets to be nonSMS managed
 - ACS routines are not invoked

```
COPY DATASET(INCLUDE(**)  
  BYPASSACS(**)  
  NULLSTORCLAS  
  RENAMEU((SOURCE.** ,TARGET.**))  
  OUTDYNAM(TARGETV)
```

DFSMSDss and SMS volume selection

- SMS follows same sequence of steps as it does for normal allocations (outside of DSS)
 - Volumes must be defined in a storage groups selected for the data set by the storage group ACS routine
 - For data sets allocated for FlashCopy, SMS will attempt to select volumes in the same SFI as the source data set
 - Performance considerations:
 - When processing many (100's) of FC's, when accessing target, response time will be better when FC source and target reside in the same cluster on an DS8000
 - Keeping data sets from crossing extents pools is also a HW performance benefit
 - If you have storage groups with hundreds of volumes “consider” using SMS PARMLIB option FAST_VOLSEL(ON|OFF)

BY Filtering (should I or shouldn't I)

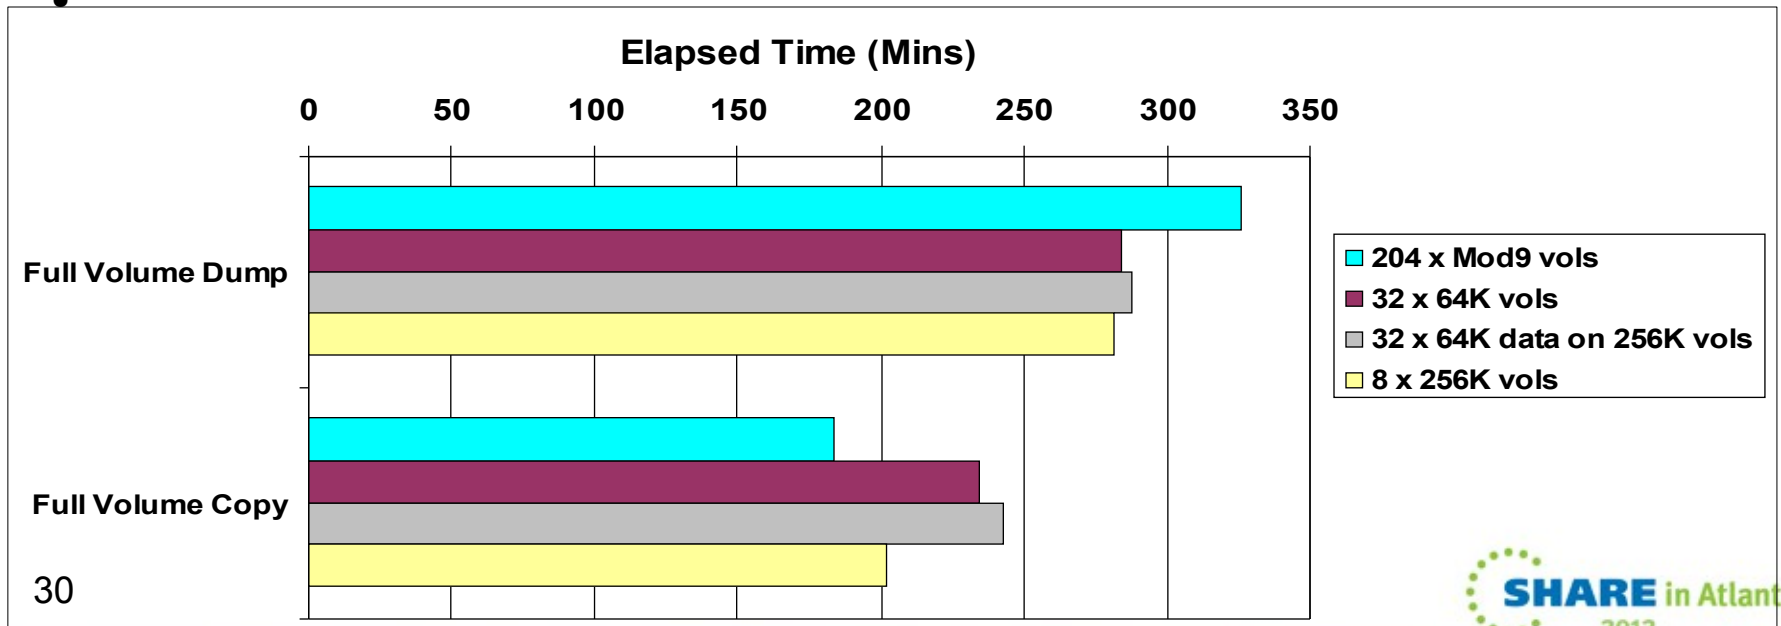
- Allows additional filtering based on data set characteristics
- Benefits
 - Limits the number of data sets processed when specifying generic include criteria
 - Done by reading VTOC to see if data set that meets include criteria should be selected
 - Many uses
 - *Finding uncataloged data sets- BY(CATLG EQ NO)*
 - *Useful if you want to delete them- (DUMP DELETE)*
 - *May get error if data set is uncataloged on multiple volumes*
 - *Processing data sets that have changed since last backup- BY(DSCHA EQ YES)*
 - *Allows for multiple BY criteria to be specified without additional reads*
- Drawbacks
 - Requires DSS to read VTOCs
 - Adds overhead
 - *Based on your include criteria this could turn out to be a lot more I/O and CPU*

VSAM CA Size Requirements

- New in V1R10
- New VSAM allocations must have a control area size of 1, 3, 5, 7, 9, or 15 tracks
- IDCAMS will automatically force any allocations to abide
- When copying data sets that don't follow the CA size requirements DFSMSDss must process it using IDCAMS
 - Prevents use of Fast Replication
 - Must be able to obtain exclusive enqueue
 - Keep in mind when an allocation is at or over the BPV the allocation will use chunk managed space and primary allocation will be in 21 cyl increments
 - May also prevent FR if source is full and target is allocated bigger (ADR918I RC26.....LDSs will be an exception soon with a recent APAR)
- Tool available to check for incompatible CA size
 - Please refer to the zOS V1R10 Migration publication for more information
- OA33531: Striped VSAM with spanned records not EAS eligible

Got EAVs?

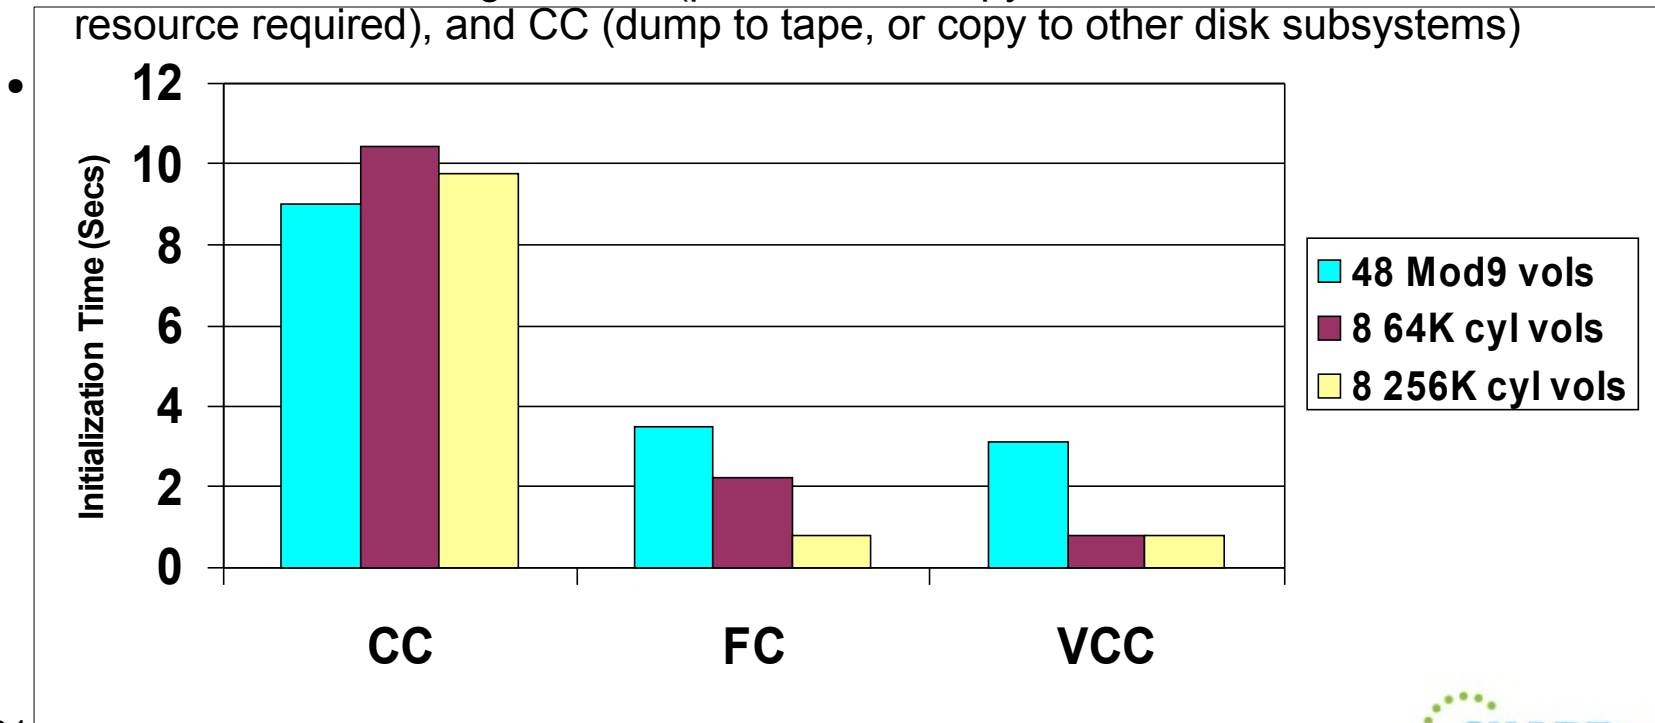
- Beginning in V1R10 volumes could be 223GB (262,668 cylinders)
- Support for 1 TB volumes (1,182,006 cyl) was made available in Sept 2011
- DFSMSdss Performance testing has found that there is no degradation when consolidating data onto EAVs (COPY, DUMP/RESTORE,DEFRAG) using traditional I/O



Got EAVs? (cont.)

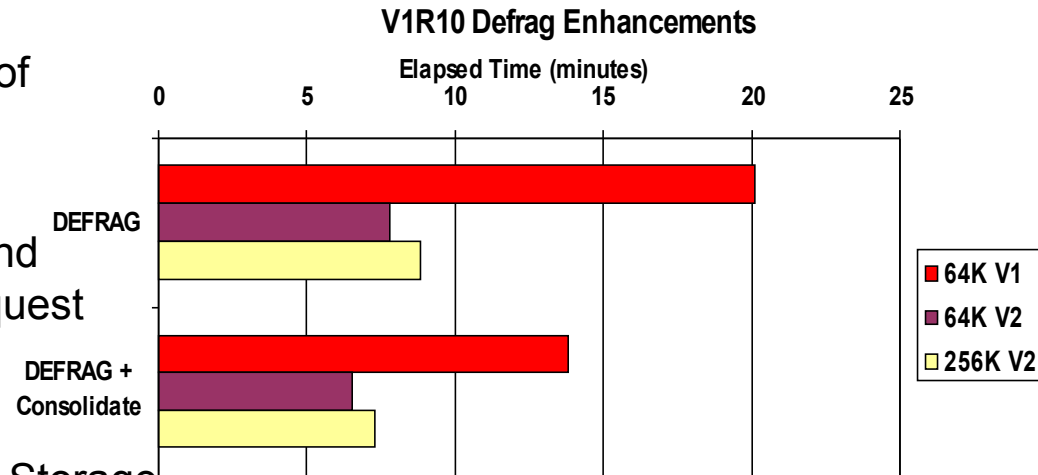
What about Copy Services?

- Initialization time required for Concurrent Copy (CC), FlashCopy (FC), and Virtual CC (VCC) initialization can be reduced by consolidating many smaller volumes onto fewer big ones
- VCC has the advantages of FC (point-in-time copy contained within DS8000, no host resource required), and CC (dump to tape, or copy to other disk subsystems)



DEFRAG and CONSOLIDATE

- Big performance improvement implemented in V1R10
- CONSOLIDATE can be run outside of DEFRAG
- Why DEFRAG?
 - Reduces time spent in SW to find enough space for allocation request
 - When do I DEFRAG
 - Monitoring FRAGINDEX
 - Tivoli OMEGAMON XE for Storage
 - Good practice is to DEFRAG regularly to prevent longer run times
 - SG Threshold value (85% default)
- Why CONSOLIDATE?
 - Reduces number of extents
 - Eliminates both SW and HW overhead
 - Keep in mind a read or write I/O specifies a range of sequential tracks (more CCWs in a CP)



DFSMSdss COPY/DUMP/RESTORE FULL

- Bypass authorization checking, expiration date checking, or unused space checking
- ADMINISTRATOR keyword
- PURGE keyword
- ALLDATA(*) ALLEXCP keywords
- If all of the options are not specified DSS will read every VTOC track and find each F1 DSCB

DFSMSdss DUMPCONDITIONING

- With Dump Conditioning, when a full volume copy is done, the source volume is copied to the target volume, but the target volume label does not change
 - Allows both the source and target to remain online
- Example:
 - Copy Volume A to Volume B with Dumpconditioning
 - Dump Volume B to Tape
 - Restore Tape to Volume C or back to Volume A
 - Volume C now looks exactly like Volume A
- You can also physically copy or restore individual data sets
 - To Volume C or A
 - If copying or restoring to C VSAM and multi volume data sets require recataloging

DFSMSdss TRACKS Processing

- Must be requested when processing CP VOLUMES with DSS
- Ending range specification
 - Specifying x'0FFFFFFFF' for cyl DSS will process to last cyl
 - 268,435,455 cylinders
 - Recommended so your batch jobs don't miss copying or dumping data
 - Example

DUMP TRACKS -

((X'1',X'0',X'0FFFFFFFF',X'0F')) -

INDD (DSK79D0) OPT(4) ADMIN CPVOLUME -
OUTDDNAME (DD9XWRK)

DFSMSDSS Parallel Processing

- DSS limits parallel processing to 80 sub tasks
- 24 subtasks in parallel per ASID gives you the maximum benefit
- DSS default limit is 80 subtasks per ASID
 - Can be overwritten via ADRPATCH (offset x'45')
 - Anything over 80 parallel subtasks could result in 878 abends

Enable or Disable DFSMSdss CSI usage?

- Good question, it depends. First a history.....
- Prior to V1R11, OA25644 introduced using CSI to locate data sets
 - Enabled via ADRPATCH
- Became the default in V1R11, but OA32120 changed it back
 - Reason: run times increased when specifying INCAT
- You should see a benefit when CSI is enabled when INCAT is not specified
- No benefit if input volumes or STORGRP is specified
 - CSI is not used by DFSMSdss
- If using CSI be sure to have fix provided in OA32165
 - Filter containing * in first character of qualifier will not select data sets
 - INCLUDE(DS(*.FILT or *P.FILT))

DFSMSdss DUMP/RESTORE on V1R12

- In V1R12 DSS began using BSAM to
 - Write bigger tape blocks
 - Use an EFSAM as DUMP output/RESTORE input on DASD
 - If backup is on EF SAM lower level releases cannot restore backup
 - One way around this is to COPYDUMP from an EFSAM to a Large SAM
 - Target blocksize has to match source
 - If backup is on tape DSS can restore it on lower releases
 - No matter what the blocksize is
 - We saw a performance improvement up to 36%
 - Average 20-30%

To COPY or to NOCOPY? That is the question!

- **BACKGROUND NOCOPY is typically the best choice to minimize activity within the physical box**
 - But.... You must ask why are you making a copy? And.... What type of application workload do I have?
 - For example:
 - Is the copy only going to be used for creating a tape backup?
 - *BACKGROUND NOCOPY should be used and the relationship withdrawn after the tape backup is complete*
 - Is the copy going to be used for testing or development?
 - *NOCOPY again is typically the best choice*
 - Will you need a copy of the copy?
 - *BACKGROUND COPY must be used so that the target will be withdrawn from its relationship after all of the tracks are copied thereby allowing it to be a source in a new relationship*
 - *Possibly use NOCOPY to COPY option*
 - *Most efficient when the source and target volumes are within the same cluster*
 - Is the workload OLTP (NOCOPY typically is the choice) or are there a large number of random writes and are not cache friendly (COPY may be the better choice)

DFSMSdss FlashCopy Batch Protection using SMS

- Forces data sets to be copied to a specific storage group when FlashCopy is used as the data mover
- Must create a new storage group to contain the FlashCopy target eligible volumes (example nonXRC volumes)
- Populate new SG with the FlashCopy target volumes
- Modify SG ACS routine to direct allocations to the new SG when the value of the new ACS variable &ACSENV2 is equal to 'FLASHCPY'
- Revalidate the ACS routine and re-activate the SMS configuration
- APARs OA32101 and OA32103
- Example

```
PROC &STORGRP
```

```
IF &ACSENV2 = 'FLASHCPY' THEN SET &STORGRP = 'FCTARGET'
```


Red Alert

- OA38942: Possible data loss

Abstract:

Possible data loss for EAV Users on releases z/OS 1.11, 1.12 and 1.13 using DSS Logical Dump or HSM Migrate and Backup processing

Description: DSS may skip some of the user's data set during dump processing and result in a failure when attempting to Restore, Recall or Recover these data sets. This failure will prevent any of the data set from being restored (ADR351E). The problem can occur when using EAV volumes with data sets in the EAS region (Extended Addressability Space - cylinder managed region) where the addresses of the data set extents are not in ascending cylinder order, ie. the CCCH address of a later extent is less than a previous extent.

Please see APAR OA38942 for additional information and actions to determine exposure.

Recommended Actions:

Apply ++APAR for OA38942

Reference Materials

- Publications:
 - SC35-0428: z/OS V1R10.0 Migration
 - GA22-7499: z/OS V1R12 Migration
 - GA22-7499 : DFSMSdss Storage Administration

Summary

- SMS Overview
- DFSMSdss Best Practices
- Questions?