



DFSMSHsm DB2 Backup Enhancements

Session 17805

Glenn Wilcock
DFSMSHsm Architect
IBM



#SHAREorg



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

Permission is granted to SHARE Inc. to publish this presentation paper in the SHARE Inc. proceedings; IBM retains the right to distribute copies of this presentation to whomever it chooses.
Except where otherwise noted, this work is licensed under
<http://creativecommons.org/licenses/by-nc-sa/3.0/>



Agenda

- **Solution Overview**
- **Common Dump Queue**
- **Dump Stacking**
- **Update Dump Class Expiration**
- **Fast Replication Messaging**

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

Business Continuity Overview

Business Continuity


Maintaining business operations in the event of an outage – with processes and infrastructure that are responsive, highly available and scalable

Three key characteristics

- ✓ Recovery Time Objective (RTO)
 - The *acceptable* amount of time you can afford to be without your data
- ✓ Recovery Point Objective (RPO)
 - The amount of data that can be *acceptably* recreated
- ✓ Backup Window Objective (BWO)
 - The *acceptable* amount of time operations can be quiesced to create a copy

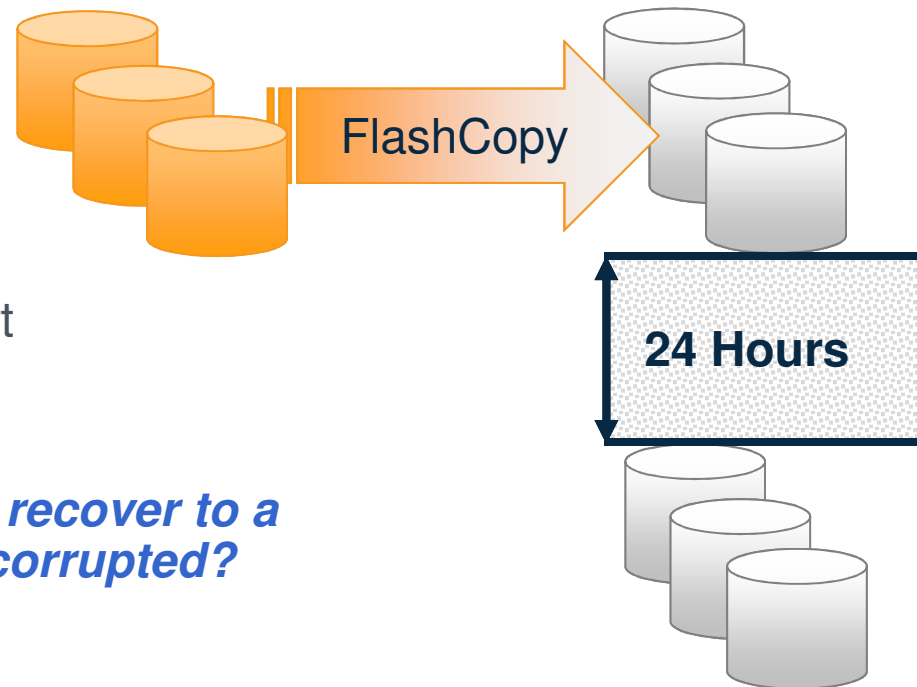
Business Continuity Overview *(continued)*

- Responsiveness to *Logical* data loss is sometimes overshadowed
 - Data corruption, User errors, Application errors, Localized data loss, etc
 - What are your RTO / RPO / BWO for these types of data loss?
 - Data-mirroring solutions do not help
 - ! *The corruption is just instantly copied from the local site to the remote site*
 - Not a question of *if* it will happen, but rather *when* it will happen



Introduction to Continuous Data Protection

- **Traditional Point-in-Time Backup**
 - Taken at specific time or data points
 - Only captures data at the point of the backup
 - Low RPO
 - RTO requirements vary
 - Disk – Short
 - Tape - Long
 - BWO requirements vary
 - Point-in-Time Copy – Short
 - Standard I/O – Long
- *Wouldn't it be nice to be able to recover to a point right before the data was corrupted?*

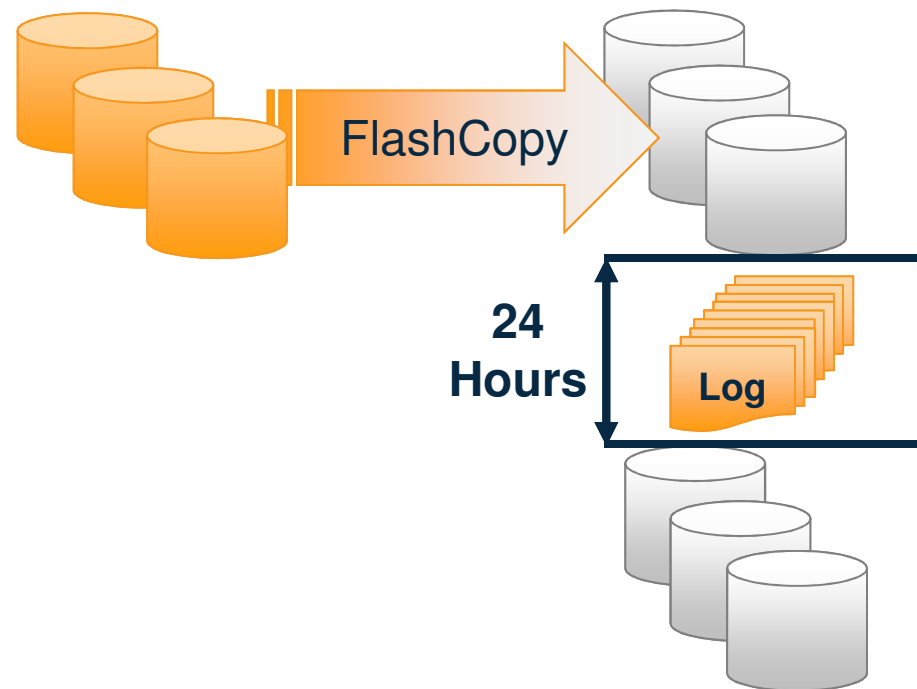


Introduction to Continuous Data Protection

(continued)



- **Continuous Data Protection (CDP):**
 - ★ **Continuously captures all changes**
 - Journaling combined with Point-in-Time copies
 - ★ **Eliminates backup window**
 - Short/Transparent BWO
 - ★ **High RPO**
 - ★ **Generally short RTO**
 - Long from tape



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



Types of Continuous Data Protection

- **Block Based**
 - Capture done at *storage level*
 - Time-ordered capture of every block write
 - Capturing process does not '*understand*' the data
 - Post processing may be required for a *data consistent* recovery
 - **True CDP**
- **Software Based**
 - Specific *software* journals every update
 - Recovery is tightly integrated with the software
 - Enables data consistent recovery
 - **True CDP**
- **File Based**
 - Runs on application host (Linux, AIX, Windows, etc)
 - Backup created when file is written to disk
 - Policies can be based on needs of various file types
 - **Near CDP**

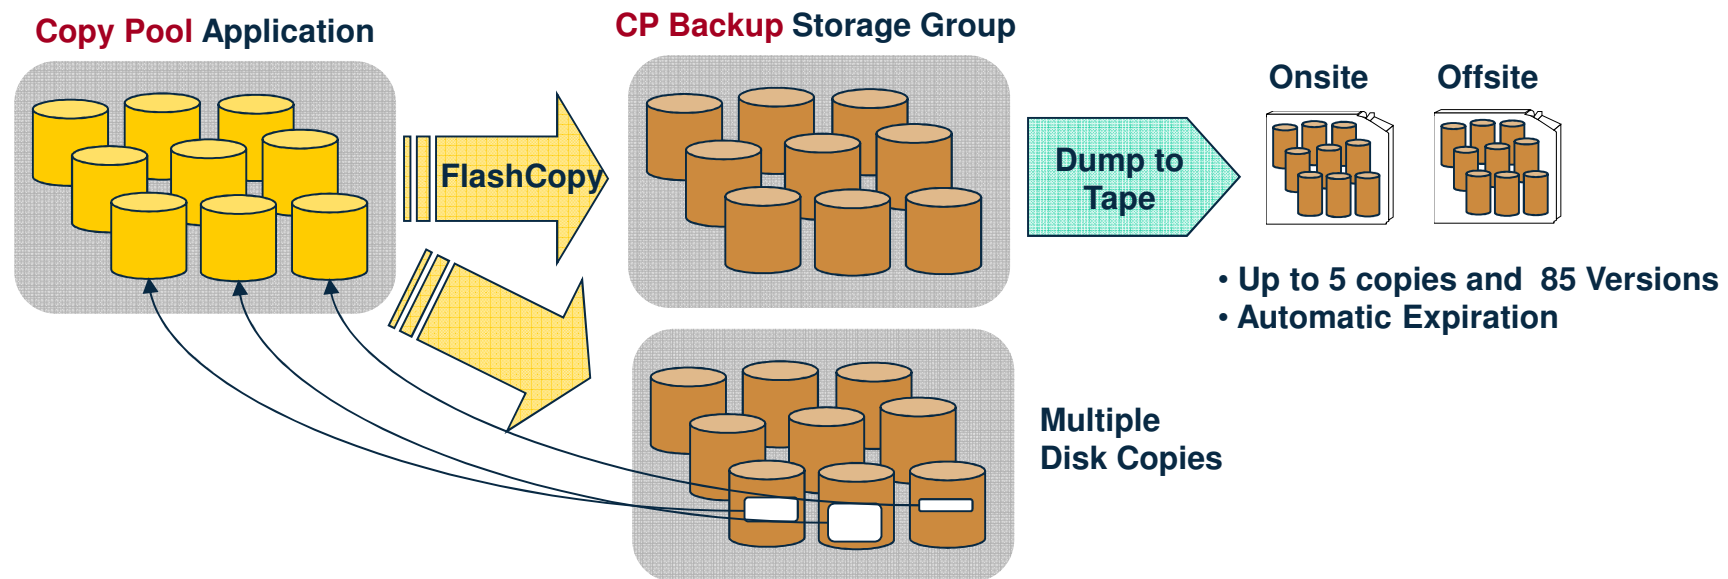
zCDP for DB2

- Software based CDP for DB2 on System z
 - Joint solution between DFSMS and DB2
- Solution based on Point-in-Time (PIT) backups combined with DB2 logging
 - ★ Eliminates need for DB2 Log Suspend
 - Only Object-level creates, extends, renames and deletes are suspended
 - Hundreds of volumes backed up in a matter of minutes
 - ★ Managed tape copies created from PIT copies
 - ★ Recovery at the System or Tablespace level
- Base Support: DB2 V8, z/OS V1R5 (2003)

zCDP for DB2

HSM function that manages Point-in-Time copies

- Combined with DB2 BACKUP SYSTEM, provides non-disruptive backup and recovery to any point in time for DB2 databases and subsystems (SAP)



★ Recovery at all levels from either disk or tape!

- Entire copy pool, individual volumes and ...
- Individual data sets

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

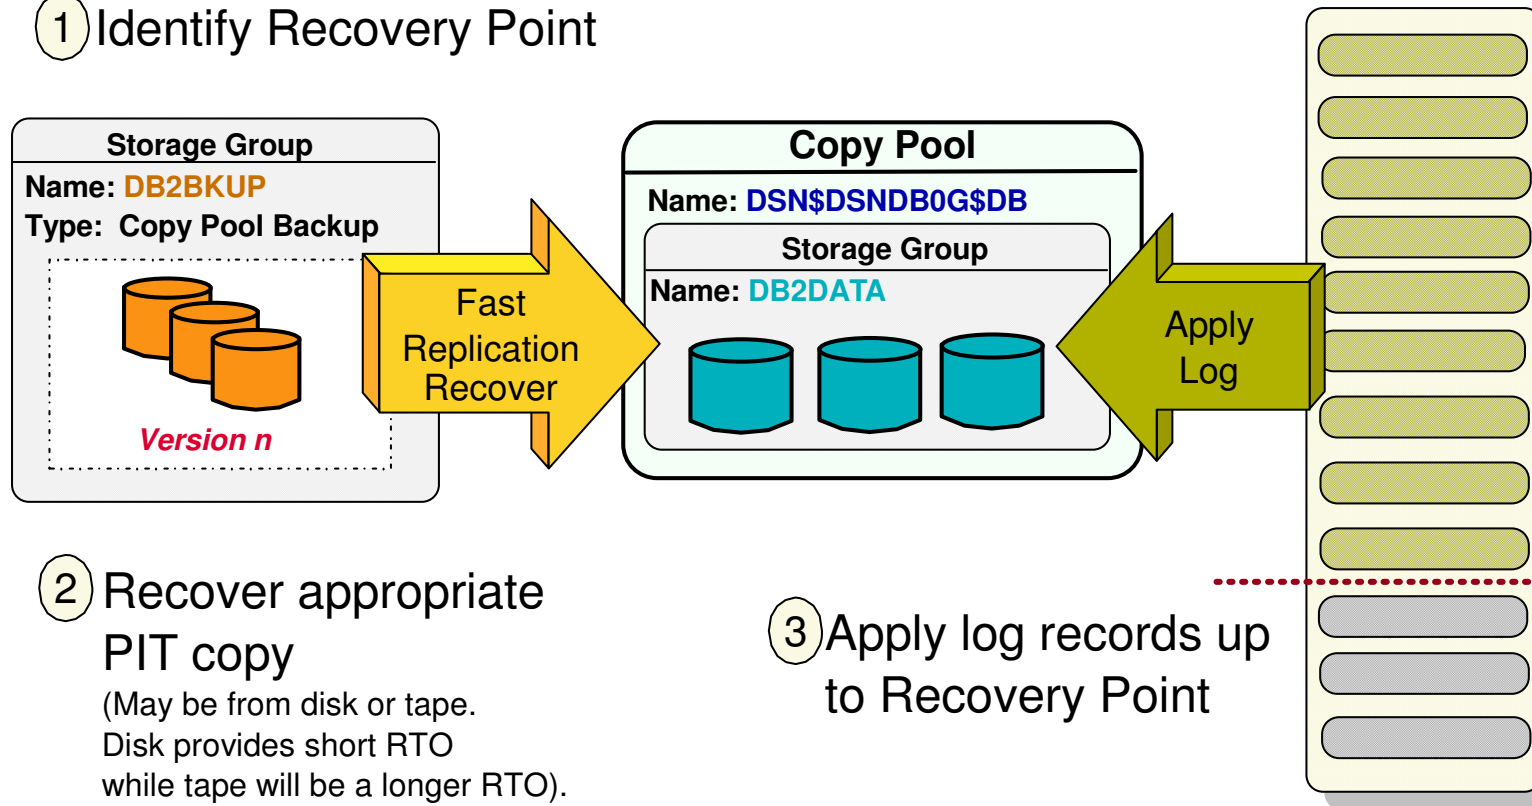
All Flavors

- **Backup**
 - ✓ Full
 - ✓ Incremental
 - ✓ No background copy
 - ✓ Space Efficient
 - ✓ Preserve Mirror
 - ✓ Consistency Groups
- **Recover**
 - ✓ Fast Reverse Restore
 - ✓ Preserve Mirror
 - ✓ Data Set Level
 - ✓ Database Level
 - ✓ Deleted data sets
 - ✓ Data sets to a new volume
 - ✓ Data set Rename

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

DB2 RESTORE SYSTEM

1 Identify Recovery Point



2 Recover appropriate PIT copy
(May be from disk or tape.
Disk provides short RTO
while tape will be a longer RTO).

3 Apply log records up to Recovery Point

DFSMSHsm V2R2


DFSMSHsm Fast Replication Enhancements

- Dump (Command, Auto, Fast Replication) commands today must be processed on the same DFSMSHsm host that initiated the request.
 - The resources of a single DFSMSHsm host do not scale to the large copy pool sizes that need to be processed.
- **New enhancement:** Distribute dump work across all hosts in a sysplex-wide Common Dump Queue (CDQ).
 - Introduces an architecture to distribute workload across host members in an HSMPLex, and the new architecture allows for flexible configurations.
 - Invoked and controlled via SETSYS commands.

? Why it Matters: Improves overall throughput by distributing the workload across the sysplex instead of concentrating it on a single LPAR.

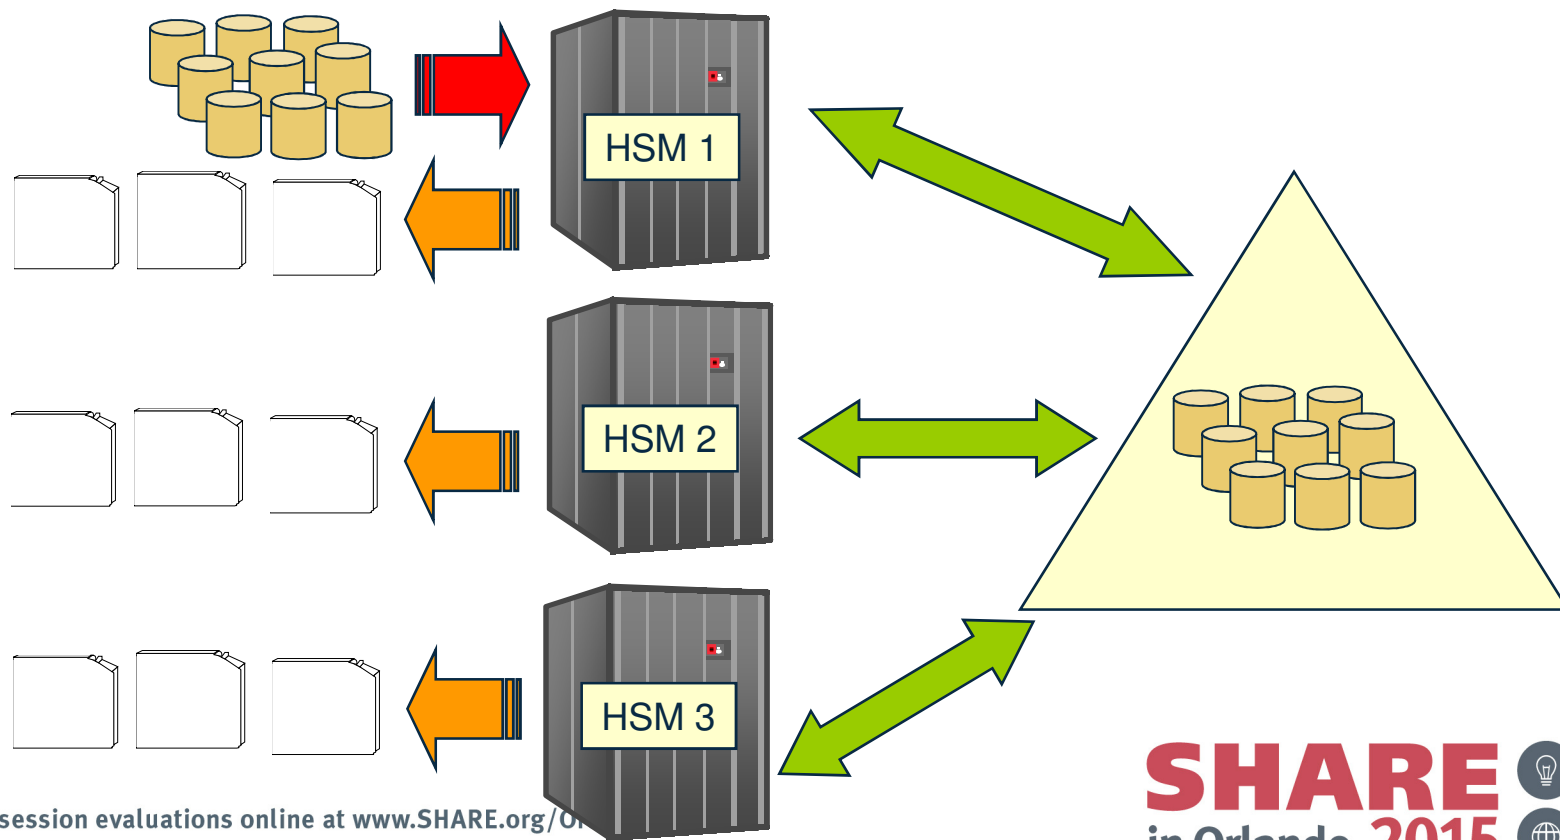
Overview – Common Dump Queue

- Today, all volumes associated with a copy pool must be dumped with a single command processed on a single host
 - Copy pools may be comprised of hundreds of volumes



Overview – Common Dump Queue

- Common Dump Queue introduces an architecture to distribute workload across host members in an HSMplex
 - HSMplex-wide queue for distributing work




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

Overview – Common Dump Queue

■ Architecture allows for flexible configurations

- Multiple queues in the same HSMplex, same and different functions
 - Hosts that both receive and process requests from the group
 - Hosts that just process requests
 - Hosts that receive requests, do but not process them
- No system affinity
- Supports priority

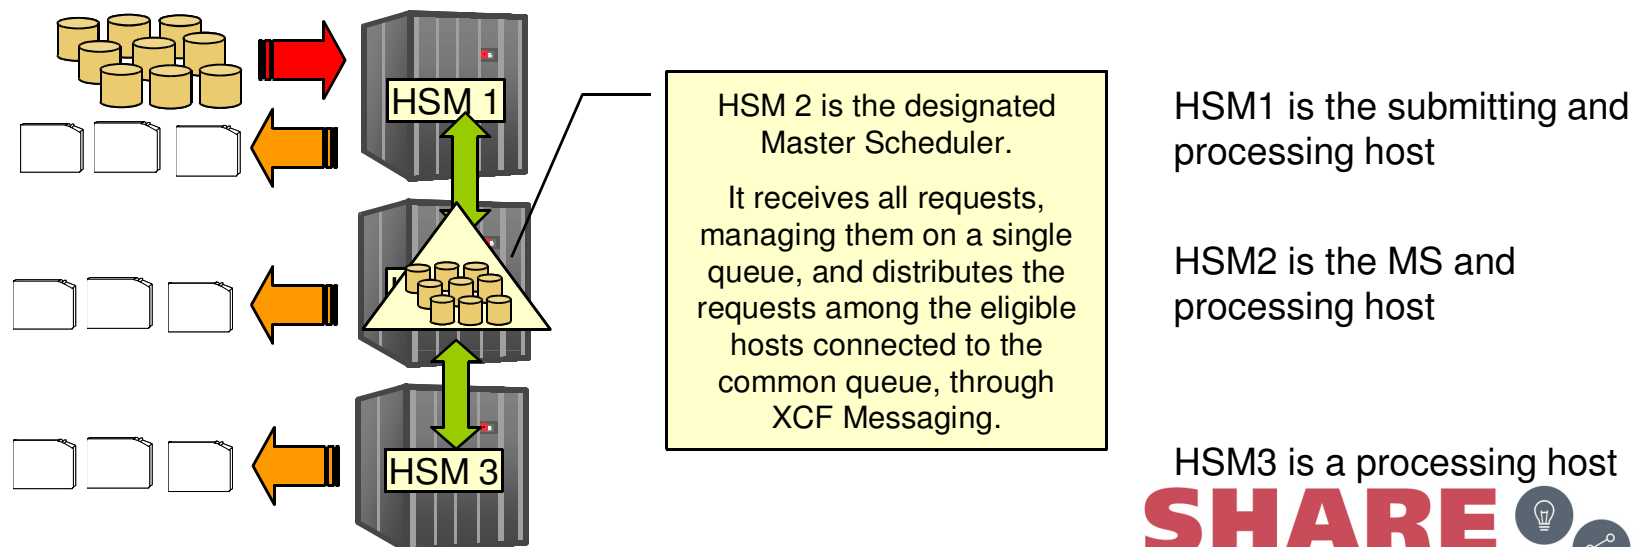


Complete your session evaluations online at www.SHARE.org/Orlando-Eval

Overview – Common Dump Queue

■ Implementation

- 'Common Queue' is managed by sending all requests to a 'Master Scheduler' that maintains a single, common queue
 - Unlike the Common Recall Queue that implements a persistent CF List Structure
 - A host is assigned as the Master Scheduler and is responsible for receiving all requests and distributing them among the available tasks across all of the eligible hosts connected to the common queue
 - Implemented via XCF Group Messaging



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

HSM3 is a processing host
SHARE
 in Orlando 2015



Overview – Common Dump Queue

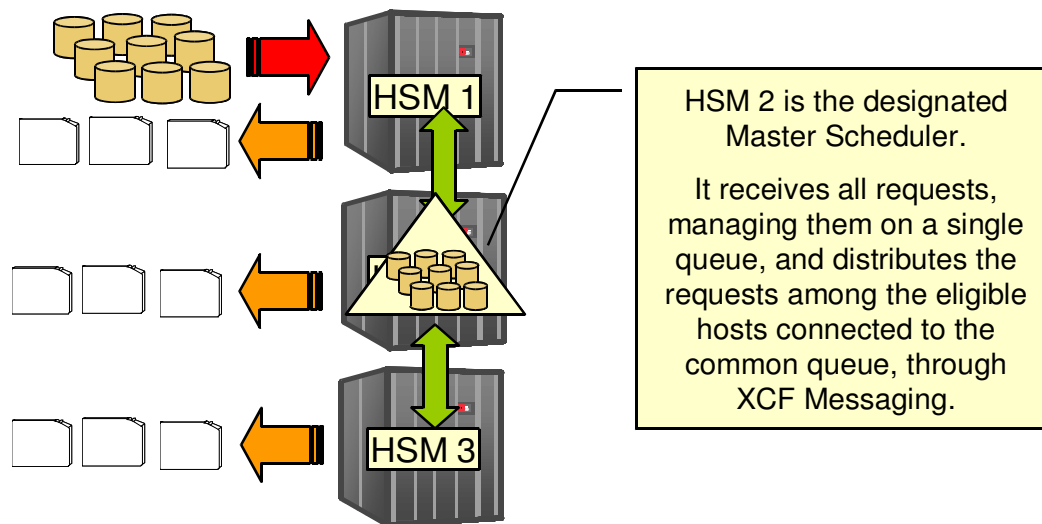
- Terminology and function
 - *Submitting Host*
 - Host that receives requests and using XCF messaging sends it to the *Master Scheduler (MS) host*
 - Processes command complete message from MS and posts user complete
 - *Master Scheduler*
 - Single Master Scheduler – race condition for all candidate hosts
 - The MS can be the submitting host - In this case the command is not sent
 - Accepts requests from *submitting host* members and locally submitted requests and maintains them as a single common queue of all requests for that grouped function
 - Assigns/distributes the requests among the available tasks across all of the eligible hosts (*processing hosts*) connected to the common queue
 - Manages work complete messages received from a *processing host* to support functions like STACKING
 - Sends command complete message back to *submitting host*

Overview – Common Dump Queue

■ Terminology and function

– *Processing Host*

- Receives the assigned work request from the MS, places it on its local queue to process
- MAXDUMPTASKS determines the number of concurrent tasks (Max 64)
- Sends work complete messages back to MS with specific protocols to support functions like STACKING
- The processing host could be the submitting host and/or the MS host



HSM1 is the submitting and processing host

HSM2 is the MS and processing host

HSM3 is a processing host

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

Usage and Invocation

SETSYS COMMONQUEUE(DUMP(CONNECT(*basename*)))

- *basename* is a four character *suffix* for the XCF Group name
Dump Group name: *ARCDbasename*

Example - **SETSYS CQ(DUMP(CONN(QUE1)))**
Common Queue Name: **ARCDQUE1**

- A host may only connect to a single common queue for each function
 - Recall
 - Dump
- Member Name is **HOST#hsmid** in the XCF group
- Master Scheduler Candidate (MSC is defaulted to YES)

Usage and Invocation

■ D XCF,GROUP

```
IXC331I 08.21.32 DISPLAY XCF
          GROUPS (SIZE) : ARCDQUE1 (2)
```

■ D XCF,GROUP,ARCDQUE1

```
IXC332I 08.22.00 DISPLAY XCF
          GROUP ARCDQUE1:  HOST#A                HOST#B
```

■ D XCF,GROUP,ARCDQUE1,ALL

```
IXC333I 08.22.12 DISPLAY XCF
          INFORMATION FOR GROUP ARCDQUE1
          MEMBER NAME:      SYSTEM:      JOB ID:      STATUS:
          HOST#A           SYSTEM1      DFHSM       ACTIVE
          HOST#B           SYSTEM1      DFHSM       ACTIVE
```

```
INFO FOR GROUP ARCDQUE1 MEMBER HOST#A ON SYSTEM SYSTEM1
FUNCTION: HSM DUMP XCF GROUP
MEMTOKEN: 01000003 000C0001 ASID: 0038 SYSID: 01000001
```

Usage and Invocation

Connect Use Case

- Host A and B, not connected, group not defined

F H*,SETSYS CQ(DUMP(CON(QUE1)))

Response from Host A

ARC1566I CONNECTION TO THE DUMP GROUP ARCDQUE1 WAS
ARC1566I (CONT.) SUCCESSFUL
ARC1569I HOST=A HAS TAKEN OVER THE MASTER SCHEDULER
ARC1569I (CONT.) RESPONSIBILITIES FOR THE COMMON DUMP QUEUE GROUP
ARC1569I (CONT.) ARCDQUE1

Response from Host B

ARC1566I CONNECTION TO THE DUMP GROUP ARCDQUE1 WAS
ARC1566I (CONT.) SUCCESSFUL

- Host A 'won the race' to become the Master Scheduler
- The first connection, Host A, created the CDQ group, ARCDQUE1
- Host B connected to it when it joined the now already existing group

Usage and Invocation

SETSYS CQ(DUMP(MSC(Y|N)))

- Indicates whether or not a host is a master scheduler (MS) candidate
- Default is YES
- First host to CONNECT with a value of YES will become the MS
- When the MS shuts down, another eligible host becomes the MS
- NO may be specified on busy or inappropriate hosts
- Setting remembered for subsequent CONNECT
- More than one host should be eligible to be the Master Scheduler

SETSYS CQ(DUMP(CONNECT(*basename*) MSC(Y|N)))

- MSC setting can be combined with CONNECT

SETSYS MAXDUMPTASKS(0)

- Indicates that a host may submit dump requests, but not process them

Usage and Invocation

SETSYS COMMONQUEUE(DUMP(DISCONNECT))

SETSYS COMMONQUEUE(DUMP(DISCONNECT FORCE))

- The current connection to a common queue for this host is removed
- For a MS host, after all MS assigned work has completed
- For a Processing Host, after remote work running on the host has completed
- When **FORCE** is not specified, after this host's submitted work completes in the common queue
- When **FORCE** is specified, submitted work from this host will be failed locally but may still complete in the common queue
- If the host was the master scheduler, this role will be assigned to another candidate master scheduler host after this host disconnects
 - *If none available, local processing of work takes over*

Usage and Invocation

Disconnect Use case

- Host A and B connected to ARCDQUE1 group, Host A is MS
- Disconnect from Host A

F HA,SETSYS CQ(DUMP(DISC))

Response from Host A

ARC1568I HOST=A HAS LOST ITS MASTER SCHEDULER

ARC1568I (CONT.) RESPONSIBILITIES FOR THE COMMON DUMP QUEUE GROUP

ARC1568I (CONT.) ARCDQUE1

ARC1567I DISCONNECTION FROM THE DUMP GROUP ARCDQUE1

ARC1567I (CONT.) WAS SUCCESSFUL

Response from Host B

ARC1569I HOST=B HAS TAKEN OVER THE MASTER SCHEDULER

ARC1569I (CONT.) RESPONSIBILITIES FOR THE COMMON DUMP QUEUE GROUP

ARC1569I (CONT.) ARCDQUE1

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

Usage and Invocation

Master Scheduler Change Use Case

- Host A and B connected to ARCDWEWE group, Host B is MS
- Transition Host B out of being the MS

F HB,SETSYS CQ(DUMP(MSC(N)))

Response from Host B

ARC1568I HOST=B HAS LOST ITS MASTER SCHEDULER
ARC1568I (CONT.) RESPONSIBILITIES FOR THE COMMON DUMP QUEUE
GROUP ARCDQUE1

Response from Host A

ARC1569I HOST=A HAS TAKEN OVER THE MASTER SCHEDULER
ARC1569I (CONT.) RESPONSIBILITIES FOR THE COMMON DUMP QUEUE
GROUP ARCDQUE1

Usage and Invocation

Shutdown

- Prior to shutdown, issue **SETSYS CQ(MSC(N))**
- Orderly quiesce of Master Scheduler responsibilities on this host to reassign them to a Master Scheduler candidate
- If not issued, the Shutdown will Fail queued requests and wait for currently processed requests to complete before shutting down

Master Scheduler loss of connectivity

- All hosts finish currently assigned work
- MS issues a message that it cannot complete work
- Submitting hosts take back submitted requests
- If a new MS is assigned, submitting hosts resubmit to new master scheduler
- Otherwise, they become local requests

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

Usage and Invocation

QUERY SETSYS

- The dump group common queue connection settings will be displayed in ARC1500I. Format of the changed message is:

```
ARC1500I PLEXNAME=HSMplex_name,PROMOTE PRIMARYHOST={YES
| NO},PROMOTE SSM={YES | NO},COMMON RECALL QUEUE BASE
NAME={basename | *****},COMMON RECALL QUEUE
TAPEDATASETORDER={PRIORITY | FBID | *****},COMMON DUMP QUEUE
BASE NAME={basename | ****} AND MSC={YES | NO | ***}
```

Example:

```
ARC1500I PLEXNAME=ARCPLEX0,PROMOTE PRIMARYHOST=NO,
PROMOTE SSM=NO,COMMON RECALL QUEUE BASE NAME=*****,
COMMON RECALL QUEUE TAPEDATASETORDER=***** ,COMMON DUMP
QUEUE BASE NAME=QUE1 AND MSC=***
```

- This Host requested to connect to CDQ group name, ARCDQUE1. The Master Scheduler Candidate setting was not specified but the default for it allows it to be eligible

Usage and Invocation

QUERY ACTIVE

- Display the current status on the system that issued the query command
- The dump group common queue active status will be displayed in new message ARC1560I:

ARC1560I COMMON type QUEUE FACTORS: GROUPNAME={groupname | *},
CONNECTION STATUS={CONNECTING | CONNECTED | DISCONNECTING |
UNCONNECTED | QUIESCED }, MASTERSCHEDULER HOSTID={hostid | NONE | *},
MASTER SCHEDULER STATUS={DISABLED | CANDIDATE | ENABLED | *}

Example:

ARC1560I COMMON DUMP QUEUE FACTORS: GROUPNAME=**ARCDQUE1**,
ARC1560I (CONT.) CONNECTION STATUS=**CONNECTED**, MASTERSCHEDULER
ARC1560I (CONT.) **HOSTID=A**, MASTER SCHEDULER STATUS=**ENABLED**

Usage and Invocation

QUERY COMMONQUEUE

- Display all the common queue requests queued or active in the group
- New DUMP option on QUERY CQ
- XCF messaging used to send and receive responses to be printed on the QUERY host
- New messages ARC1562I (active) and ARC1563I (waiting) may be issued

ARC1563I VOLUME {volser | volser, SGROUP=sg | volser, COPY POOL = cpname, }
{DUMPING | FRBACKUP DUMP | FRBACKUP DUMPNLY} FOR USER {userid |
AUTO | * } REQUEST {request-number [FROM HOST=hostid] | NONE }, WAITING TO BE
PROCESSED ON A COMMON QUEUE, nmwe MWES AHEAD OF THIS ONE,
GROUPNAME=groupname

ARC1562I VOLUME {volser | volser, SGROUP=sg | volser, COPY POOL = cpname, }
{DUMPING | FRBACKUP DUMP | FRBACKUP DUMPNLY} , FOR USER {userid |
AUTO | * } REQUEST {request-number [ON HOST=hostid] | NONE }, BEING
PROCESSED FROM A COMMON QUEUE, GROUPNAME=groupname

Usage and Invocation

QUERY WAITING

- Lists the number of requests *that originated from the host issuing the query* command that are waiting
- For a host connected to a CDQ (not the MS), the local queue will be searched and the master scheduler member's group queue will be searched
- XCF messaging used to send and receive responses to be printed on the QUERY host
- For a host connected to a CDQ (MS), the local queue will be searched
- Existing messages, ARC0168I and ARC1542I will be reported

ARC1542I WAITING MWES ON COMMON QUEUES: COMMON RECALL
QUEUE=type_mwes, **COMMON DUMP QUEUE=type_mwes**, TOTAL=total_mwes

Usage and Invocation

QUERY REQUEST or USER

- Display requests *that originated from the host issuing the query* command that match the criteria specified on the REQUEST or USER parameters
- For a host connected to a CDQ (not the MS), the local tasks and queue will be searched and the master scheduler member's group queue will be searched
 - XCF messaging used to send the command and receive responses to be printed on the QUERY host
- For a host connected to a CDQ (MS), the local tasks & queue will be searched
- Existing messages, ARC0161I and ARC0167I will be reported
 - QUERY ACTIVE also gets ARC0161I

Usage and Invocation

QUERY REQUEST or USER Use case

- Host A is the MS, Host B is a member
- Host B submitted 3 volumes for BACKVOL processing

F HB,Q REQUEST

ARC0101I QUERY REQUEST COMMAND STARTING ON HOST=B

ARC0101I QUERY REQUEST COMMAND STARTING ON HOST=A

ARC0101I QUERY REQUEST COMMAND COMPLETED ON HOST=A

ARC0161I DUMPING VOLUME **PRIM01** FOR USER **OPER**,
 ARC0161I (CONT.) REQUEST 00000037 ON **HOST=A**

ARC0167I DUMP MWE FOR VOLUME **PRIM03** FOR USER **OPER**,
 ARC0167I (CONT.) REQUEST 00000037, WAITING TO BE PROCESSED,00000

ARC0167I (CONT.) MWE(S) AHEAD OF THIS ONE

ARC0161I DUMPING VOLUME **PRIM02** FOR USER **OPER**,
 ARC0161I (CONT.) REQUEST 00000037 ON **HOST=B**

ARC0161I (CONT.) REQUEST 00000037 ON **HOST=B**

ARC0161I (CONT.) REQUEST 00000037 ON **HOST=B**

ARC0101I QUERY REQUEST COMMAND COMPLETED ON HOST=B

- Query sent to Host A, the MS. Host A reported messages get 'XCF messaged' back to Host B for printing. Host B reported messages are also printed

Usage and Invocation

HOLD DUMP | DUMP(AUTO) | DUMP(FRBACKUP)

- HOLD continues to prevent dump processing based on its current options
 - Dump tasks not selected to process requests
 - WAIT types failed
- With CDQ 'any' specified HOLD
 - Prevents a host from being selected to process CDQ work
 - ! On a MS prevents MS duties from being performed
- Best reserved for stopping already active dump processing / scheduling
 - To configure a CDQ host that will not be selected for dump processing, it is preferable to use **SETSYS MAXDUMPTASKS(0)**
 - Host can still submit DUMP requests to CDQ
 - Requests are processed in other available dump hosts in the CDQ
 - If host is the MS, this will not impact MS responsibilities

Usage and Invocation

- If any variant of **RELEASE DUMP** is issued on a CDQ host:
 - That host will be eligible to be selected for CDQ dump processing if the following conditions are met:
 - No other HOLD for DUMP is in place for the host.
 - SETSYS MAXDUMPTASKS(x) is specified with a value greater than 0
 - If the host is the CDQ MS, it will also be eligible to perform MS duties for the particular dump function that has been released

Usage and Invocation

- **CANCEL**

- Cancels queued requests that **originated on the cancelling host** for a request number or userid that matches the CANCEL request parameters
- Applies to CDQ when the host is connected to a common queue
- *Canceling active requests with the TCBADDRESS must be done on the host that is processing the request*
- XCF messaging used to send the command and receive responses to be printed on the CANCEL host
 - For a MS host, the command is processed locally
 - For a non MS host, the command is processed locally and on the MS
- Existing and unchanged messages, ARC1008I and ARC0931I will be reported

Usage and Invocation

- **ALTERPRI**

- Changes the priority of queued requests waiting to be processed in DFSMSHsm for a request number or userid matching the alter request
 - Applies to CDQ when the host is connected to a common queue
- XCF messaging used to send the command and receive responses to be printed on the ALTERPRI host
 - For a MS host, the command is processed locally
 - For a non MS host, the command is processed locally and on the MS
- Existing messages, ARC0980I and ARC0981I will be reported

ARC0980I ALTERPRI {REQUEST | USER | DATASETNAME} COMMAND STARTING ON HOST=hostid

ARC0981I ALTERPRI {REQUEST | USER | DATASETNAME} COMMAND COMPLETED ON HOST=hostid,RC=retcode

Dump Stacking

DEFINE DUMPCLASS

- Added a new MINSTACK, minimum stack, parameter and an MAXSTACK alias for the STACK parameter

```
|_____|
|__ _STACK__ _(maxv)__| |__MINSTACK(minv)__|
|_MAXSTACK_|
```

- maxv* represents the preferred maximum number of dump copies that the system should use to place on a tape volume
 - minv* represents the preferred minimum number of dump copies that the system should use to place on a tape volume
 - Utilizes more dump tasks, but may have less dump copies on tape
 - MAXSTACK value must be equal to or larger than the MINSTACK value
 - If MINSTACK value not specified the MAXSTACK value used for it
- Example, `DEFINE DUMPCLASS(DCLASS1 MAXSTACK(100) MINSTACK(20))`

Dump Stacking

STACK and MINSTACK Use Case

- 400 volumes to be dumped
- Dump class defined with MAXSTACK(100)
- 5 dump tasks
- **Today**
 - 100 volumes are stacked onto 4 tapes, utilizing 4 dump tasks
- **Enhancement**
 - Add MINSTACK(20) to dump class
 - Volumes dumped in sets of 20 across all available dump tasks (5)
 - Additional volumes added to each tape in sets of 20 until MAX reached
 - 80 volumes will be stacked onto 5 tapes, utilizing all 5 dump tasks
- MINSTACK may also be specified on the BACKVOL command

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

Multiple Copy Pools

- Multiple Copy pools may now be stacked onto the same tape
- This will occur when the copy pools are dumped at the same time and use the same dump classes
- This only occurs with Automatic Dump prior to z/OS V2R2

Dump Expiration

- New **UPDTCDS** command will update the expiration date for a copy pool dump copy
 - Today, it can take *thousands* of FIXCDS commands to update the expiration date of a single copy pool dump version

**UPDTCDS COPYPOOL(*name*) VERSION(*ver*)
DUMPEXPIRATION(DCLASS(*name*) NEWDATE(*date*))**

- This new command updates all of the CDS records that need to be updated to reflect the new expiration date.
 - The Fast Replication Dump record
 - The dump record for each volume dumped
- This command will be extended to other functions in the future

Messaging

- **DFSMSHsm Fast Replication Messaging Enhancements**
 - Finding messages related to a Fast Replication command is very difficult because they are written to shared HSM logs.
 - You have to find the correct log and then search through all of the messages for all HSM activity.
 - **New enhancement:** Provide an option to dynamically allocate a unique message data set and write all DFSMSHsm and DFSMSdss messages for a specific FRBACKUP, FRRECOV or copy pool Automatic Dump to that data set.
 - **SHARE Requirement Addressed:** SSMVSS14005
- **? Why it Matters:** Greatly simplify problem determination.

Messaging

SETSYS FASTREPLICATION(MESSAGE DATASET(YES|NO HLQ(hlq)))

- Indicates that all Fast Replication messages associated with an FR command or copy pool automatic dump should go to a unique data set
- Message Data Sets must be SMS managed, so SMS and ACS routines must be setup prior to allocation
- DFSMSHsm dynamically allocates message data sets with the following characteristics:
 - Physical Sequential
 - FBA
 - LRECL of 121
 - Block size of 1210
- Naming Convention:
msgdsprefix.FB.Cccccccc.ippppppp.Dyyddd.Thhmmss.z
- Collisions
 - Time is incremented if a duplicate name already exists

Backup Messaging

msgdsprefix.FB.Ccccccc.ippppppp.Dyyddd.Thhmmss.z

- *msgdsprefix* is the value of the HLQ parameter of the SETSYS command. If the HLQ parameter is not specified, the default prefix of 'HSMMSG' will be used.
- *cccccc* is up to 7 characters to identify the copy pool that was processed. For a non-DB2 copy pool, this will be the first 7 characters of the copy pool name. For a DB2 copy pool, this will be the first 7 characters of the location qualifier in the DB2 copy pool name.
- *i* is an identifier, indicating whether the message data set contains messages for:
 - D – DB2 data L – DB2 log X – non-DB2 copy pool data
- *pppppp* is up to 7 characters to identify the copy pool that was processed. For a non-DB2 copy pool, this will be the next 7 characters of the copy pool name (the characters following *cccccc*). For a DB2 copy pool, this will be the next 7 characters of the location qualifier in the DB2 copy pool name.
- *yyddd* is the year and day when the FRBACKUP or copy pool auto dump operation started.
- *hhmmss* is the time in hours, minutes, and seconds when the FRBACKUP or copy pool auto dump operation started.
- *z* indicates the status of the FRBACKUP or copy pool auto dump operation, and will be set to one of the following values:
 - S – Success, F – Fail, I – In progress

Recover Messaging

msgdsprefix.FR.Ccccccc.ippppppp.Dyyddd.Thhmmss.z

- *msgdsprefix* is the value of the HLQ parameter of the SETSYS command. If the HLQ parameter is not specified, the default prefix of 'HSMMSG' will be used.
- *cccccc* is up to 7 characters to identify the copy pool that was processed. For a non-DB2 copy pool, this will be the first 7 characters of the copy pool name. For a DB2 copy pool, this will be the first 7 characters of the location qualifier in the DB2 copy pool name.
- *i* is an identifier, indicating whether the message data set contains messages for:
 - D – DB2 data L – DB2 log X – non-DB2 copy pool data
- *pppppp* is up to 7 characters to identify the copy pool that was processed. For a non-DB2 copy pool, this will be the next 7 characters of the copy pool name (the characters following *cccccc*). For a DB2 copy pool, this will be the next 7 characters of the location qualifier in the DB2 copy pool name.
- *yyddd* is the year and day when the FRRECOV operation started.
- *hhmmss* is the time in hours, minutes, and seconds when the FRRECOV operation started.
- *z* indicates the status of the FRRECOV operation, and will be set to one of the following values:
 - S – Success, F – Fail, I – In progress
- Messaging datasets only supported for COPYPOOL operations at this time

Messaging

- A message data set is allocated after the request is determined eligible for processing
- Messages indicating parse errors or termination prior to the ARC1801I or ARC1841I start messages are not collected in message data sets
- Each message is prefixed with the date and time that the message was written to the data set, along with the ID of the host the wrote the message:

yy/mm/dd hh:mm:ss Hx message-text

- where H is a constant, and x is the host id

- *Some messages will appear multiple times in a message data set*
 - Each time a message is written to the operator or a log, the message is also logged in the message data set. Duplicates are not removed

Messaging

- **New Message**

ARC1870I MESSAGE DATA SET *msgdsname* CREATED FOR {FAST REPLICATION {BACKUP | BACKUP DUMP | BACKUP DUMPONLY | PREPARE | RECOVER } | AUTO DUMP} OF {COPY POOL *cpname* | VOLUME *volser* }

- **QUERY SETSYS**

ARC1823I MAXCOPYPOOL(FRBACKUP TASKS = *backup_tasks*, FRRECOV TASKS = *recover_tasks*, DSS TASKS = *dss_tasks*), FASTREPLICATION (DATASETRECOVERY = {PREFERRED | REQUIRED | NONE } FCRELATION = {EXTENT | FULL} VOLUMEPAIRMESSAGES = {YES | NO} **MESSAGEDATASET=({YES | NO} {HLQ=*hlq*})**)

14/11/21 10:09:10 H2 ARC1801I FAST REPLICATION BACKUP DUMP IS STARTING FOR COPY POOL DSN\$PAYROLLA123\$DB, AT 10:09:10 ON 2014/11/21

14/11/21 10:09:11 H2 ARC1805I THE FOLLOWING 00003 VOLUME(S) WERE SUCCESSFULLY PROCESSED BY FAST REPLICATION BACKUP OF COPY POOL DSN\$PAYROLLA123\$DB

14/11/21 10:09:11 H2 ARC1805I (CONT.) SRC001

14/11/21 10:09:11 H2 ARC1805I (CONT.) SRC003

14/11/21 10:09:11 H2 ARC1805I (CONT.) SRC002

14/11/21 10:09:11 H1 ARC1814I FAST REPLICATION BACKUP HAS COMPLETED SUCCESSFULLY AND DUMP MWES HAVE BEEN QUEUED FOR COPY POOL DSN\$PAYROLLA123\$DB, VERSION 3

14/11/21 10:09:11 H1 ARC0622I FULL VOLUME DUMP STARTING ON VOLUME SRC001(SMS) AT 10:09:11 ON 2014/11/21, SYSTEM 3090, TASK ID=ARCDV001, TO DUMP CLASS(ES)= DCLASS1

14/11/21 10:09:11 H1 ARC0640I ARCDV001 - PAGE 0001 5695-DF175
DFSMSDSS V2R02.0 DATA SET SERVICES 2014.325 10:09

14/11/21 10:09:11 H1 ARC0640I ARCDV001 - ADR035I (SCH)-PRIME(06),
INSTALLATION EXIT ALTERED BYPASS FAC CLASS CHK

14/11/21 10:09:11 H1 ARC0640I ARCDV001 - DUMP FULL
OUTDDNAME(SYS00070) 14/11/21 10:09:11 H1 ARC0640I ARCDV001 -
ALLEXCP ALLDATA(*) OPTIMIZE(1) TOLERATE(IOERROR) ...



DFSMSHsm DB2 Backup Enhancements

Session 17805

Glenn Wilcock
DFSMSHsm Architect
IBM



#SHAREorg



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

Permission is granted to SHARE Inc. to publish this presentation paper in the SHARE Inc. proceedings; IBM retains the right to distribute copies of this presentation to whomever it chooses.

Copyright (c) 2015 by SHARE Inc. Except where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/3.0/>



© Copyright IBM Corp. 2015

Sources for more information

- Information about [DFSMS](#) and components
 - http://publib.boulder.ibm.com/infocenter/zos/basics/index.jsp?topic=/com.ibm.zos.s.zdatamgmt/zsysprog_c_dfsmselements.htm
- Information about [DFSORT](#)
 - <http://www-01.ibm.com/support/docview.wss?rs=0&uid=isg3T7000077>
- Information about [z/OS Storage Management Tools](#)
 - <http://www-03.ibm.com/systems/storage/software/zos/index.html>
- Information about [IBM Tivoli Storage Productivity Center](#)
 - <http://www-03.ibm.com/systems/storage/software/center/index.html>
- Information about [IBM System Storage Disk](#) systems
 - <http://www-03.ibm.com/systems/storage/disk/ds8000/index.html>
- Information about [IBM System Storage Tape](#) systems
 - <http://www-03.ibm.com/systems/storage/tape/?lnk=mprST-tsys>

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*
DFSMSdftp	GDPS*	OMEGAMON*	Virtualization Engine
DFSMSdss	geoManager*	OS/390*	VisualAge*
DFSMSshsm	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.