





Moving JES2 SPOOL Data Offload, Migrate, and non-JES2 Methods


Adam Nadel
anadel@us.ibm.com
IBM - Poughkeepsie, NY


Thursday, March 5, 2015
Session Number 16895



#SHAREorg


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



Copyright (c) 2014 by SHARE Inc.  script where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/4.0/>

Acknowledgments: Thank you to Miguel Perez (SDM & XRC Level 2) from Tuscon, AZ



Trademarks

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

- IBM®
- MVS™
- Redbooks®
- RETAIN®
- z/OS®
- zSeries®

The following are trademarks or registered trademarks of other companies.

- Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- 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.
- SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.
- 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.

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

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



Session 16895

SPOOL Offload - Overview



- Offload (transmit) data from the spool to a dataset
- Reload (receive) data from the data set to the spool
 - The offload medium can reside on tape or DASD
- Preserve jobs and SYSOUT across a COLD start scenario
- Provide temporary relief during spool space full/shortage condition
- Simple spool dataset “back up”
- Converting to another DASD type

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



SPOOL offload has existed in the product for 25+ years and is the most “tried and true” method of managing SPOOL data from solely within JES2. However, we will find newer methods also can be used to achieve the desired results (and in some cases be more dynamic). SPOOL offload follows the same protocol / behavior as other NJE. Specifically, NJE records and BSAM processing performs the I/O to the specified offload medium. Transmitters and receivers are logical devices defined to pass the data.

SPOOL Offload - Overview



- OFFLOAD(n)
 - Defines the offload dataset DSN=
 - Associates the logical offload device
 - Corresponds to the respective transmitter statements
- OFF(n).JT
- OFF(n).ST
- OFF(n).JR
- OFF(n).SR

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



A maximum of eight offload devices can be defined via OFFLOAD(n) statements; however, at any given time only eight total offload devices can be active within the JES2 MAS.

SPOOL offload facility has the ability to transmit (or receive) all jobs and SYSOUT data or specific subset based on WS= selection criteria rules – these are set on the respective JT/ST and JR/SR definitions. For instance, you can select to transmit (or receive) jobs based on job class, jobnumber range, routecode, etc.

SPOOL Offload – Offload Setup



```
$DOFFLOAD1
```

```
$HASP882 OFFLOAD1 DSN=OFFLOAD.DATASET1, STATUS=DRAINED,  
$HASP882 ARCHIVE=ONE, CRTIME=RESET, LABEL=SL, PROTECT=NO,  
$HASP882 RETPD=30, TRACE=NO, UNIT=(161,1), VALIDATE=YES,  
$HASP882 VOLS=255
```

offload dataset name

Job timestamp associated with original submission prior to reload or reload time

days tape is retained

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



If offloading to DASD, the offload dataset should be pre-allocated since you cannot specify SPACE= on the OFFLOAD statement. If not pre-allocated and pre-cataloged, then you can define the tape on which the data set will reside by specifying UNIT, LABEL, RETPD, PROTECT, etc.

STATUS = status of the offload device

ARCHIVE = specifies which offload devices JES2 prevents from scheduling jobs for offload

LABEL = for non-cataloged dataset, specifies the type of label processing that is required for the tape


PROTECT = specifies whether offload dataset requires SAF protection

TRACE= specifies whether JES2 trace records should be produced

UNIT= device number/type

VALIDATE = specifies whether JES2 should validate the logical record length of first (header) record

VOLS = volume count to be used for the offload



SPOOL Offload – Offload Setup

```

$HASP884 OFF1.JT
$HASP884 OFF1.JT STATUS=STARTABLE,CLASS=B,CREATOR=,
$HASP884 DISP=DELETE,HOLD=,JOBNAME=,NOTIFY=NO,RANGE=(J
$HASP884 1,65534),ROUTECD=( ),START=YES,SYSAFF=( ),
$HASP884 VOLUME=(,,),WS=(CL/)

```

Status must be STARTABLE to indicate it is ready

Work selection criteria


Disposition of job following processing by offload

```

$HASP886 OFF1.ST
$HASP886 OFF1.ST STATUS=STARTABLE,CREATOR=,DISP=DELETE,
$HASP886 OUTDISP=(WRITE,KEEP),HOLD=,JOBNAME=,
$HASP886 NOTIFY=NO,RANGE=(J1,65534),ROUTECD=( ),
$HASP886 START=YES,VOLUME=(,,),WS=(Q/),BURST=,FCB=,
$HASP886 FLASH=,FORMS=,LIMIT=(0,*),PLIM=(0,*),
$HASP886 PRMODE=( ),QUEUE=X,UCS=,WRITER=

```

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

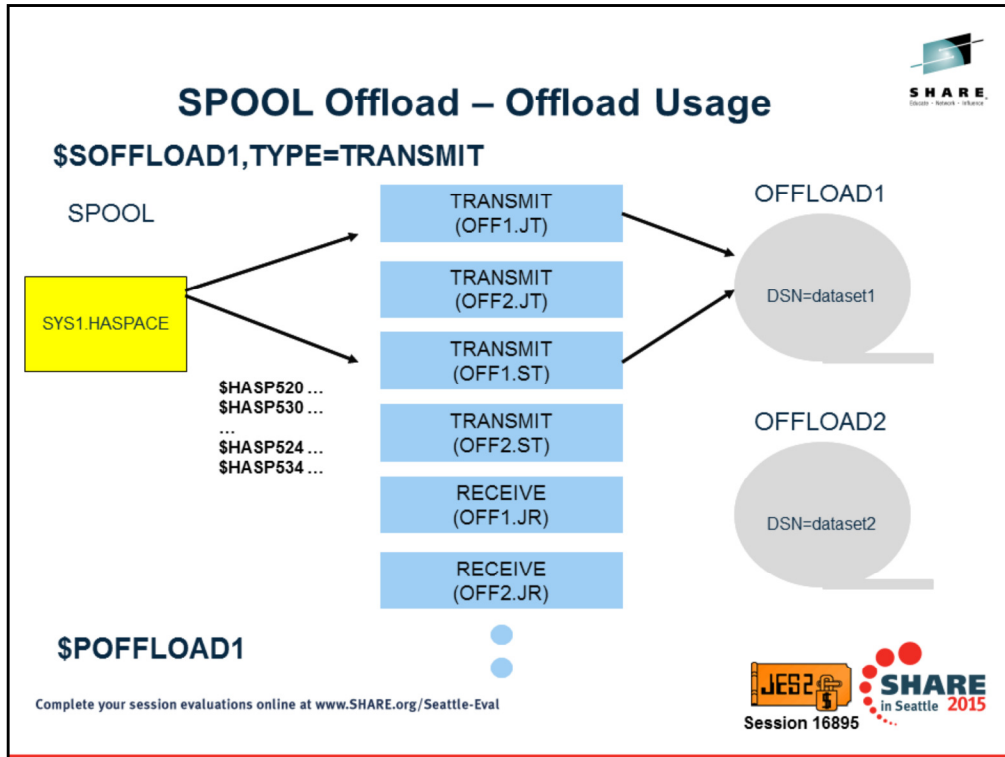


Session 16895

Most of the parameters on the JT/ST are honored only if specified as part of the workload selection criteria WS= . For additional information regarding how to specify WS criteria, please refer to the JES2 Initialization and Tuning Guide (specifically Tables 60 and 61).

If want to OFFLOAD everything, then you would specify WS=(/). If you want to speed up the process you can use multiple offload devices, but be sure to specify WS criteria so that they do not compete with one another. Such as one device selecting the HOLD queue, while another selecting the OUTPUT queue –or- one device selecting class A thru L, while another selecting M thru Z.

If the available WS criteria options are insufficient, it is also possible for an installation to create and define their own USER-defined selection criteria via the \$WSTAB macro. There is a limit of 19 selection criteria for any device (JES2 and USER criteria combined)



As each respective job is processed by the offload transmitter, \$HASP520 messages will be produced. SYSOUT transmitters will produce \$HASP530 message. Once transmission is complete and devices go inactive, JES2 issues \$HASP524 for job transmitter and \$HASP534 for SYSOUT transmitter. After the devices are inactive and transmissions complete, command \$POFFLOAD1 will close and deallocate the offload dataset.

SPOOL Offload – Reload Setup



Status must be STARTABLE to indicate it is ready

```
$HASP883 OFF1.JR  
$HASP883 OFF1.JR STATUS=STARTABLE,CLASS=,CREATOR=,HOLD=,  
$HASP883 JOBNAME=,MOD=(CLASS=,HOLD=,ROUTECD=,  
$HASP883 SYSAFF=),NOTIFY=NO,RANGE=(J1,2147483647),  
$HASP883 ROUTECDE=(),START=YES,SYSAFF=(),WS=(CL/)
```

Work selection criteria

Characteristics of pre-execution jobs (or output) upon reload

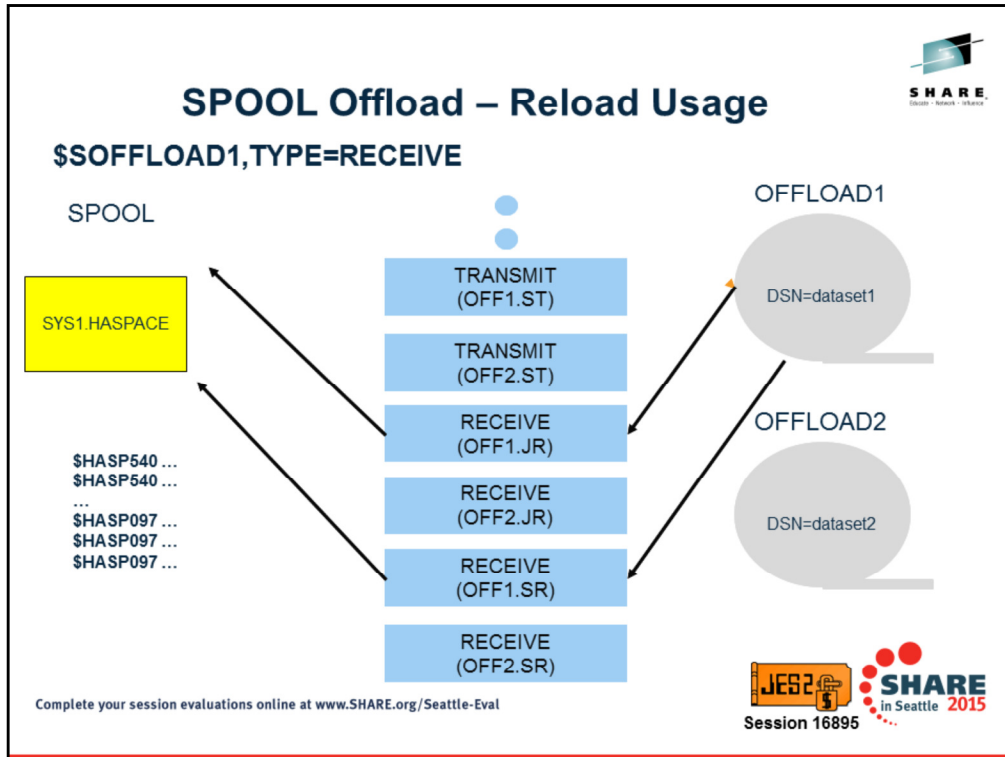
```
$HASP885 OFF1.SR  
$HASP885 OFF1.ST STATUS=STARTABLE,  
$HASP885 OUTDISP=(WRITE,HOLD,KEEP,LEAVE),CREATOR=,  
$HASP885 HOLD=,JOBNAME=,MOD=(BURST=,OUTDISP=,FCB=,  
$HASP885 FLASH=,FORMS=,HOLD=,PRMODE=,QUEUE=,ROUTECD=,  
$HASP885 UCS=,WRITER=),NOTIFY=NO,RANGE=(J1,214783647),  
$HASP885 ROUTECDE=(),START=YES,WS=(Q/),BURST=,FCB=,  
$HASP885 FLASH=,FORMS=,PRMODE=(),QUEUE=X,UCS=,WRITER=
```

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



Session 16895

Most of the parameters on the JR/SR are honored only if specified as part of the workload selection criteria WS= . For additional information regarding how to specify WS criteria, please refer to the JES2 Initialization and Tuning Guide (specifically Tables 60 and 61).



JES2 will call SAF to verify that a userid is valid and has access to the data with specific security labels. The receiving system will assign a new date/time to the SYSOUT receiver upon reload unless CRTIME=RESTORE specified (in which case pre-offload date/time is preserved).

As each respective SYSOUT is received, message \$HASP540 is produced. Once transmission is complete and devices go inactive, JES2 issues \$HASP097 for job receiver, SYSOUT receiver, and overall offload device. JES2 also automatically drains the devices.

SPOOL Offload – Common Problems



- Offloaded job disposition not what you expect?
 - OFF(n).ST or OFF(n).JT DISP = parm controls what happens to the original job/output following offload
- Offloaded everything with WS=(/) but some jobs were not offloaded?
 - Issue \$DJnnn to find its status. It is possible they were HELD in execution phase and now AWAITING OUTPUT. May require \$AJnnn to release the job
- Attempting to offload the same job (or SYSOUT) multiple times?
 - When a job (or SYSOUT) is offloaded, it is marked to indicate which offloader has processed it via the OFFS= keyword on \$DJ (or \$DOJ)
 - Can be processed by different offloaders, but not by the same one unless the OFFS= demarcation is removed via \$TJ (or \$TOJ)

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



SPOOL Offload – Common Problems



- Reloaded job's timestamp not what you expect?
 - OFFLOAD(n) CRTIME= parm influences the job creation timestamp following reload
- Reloaded job number not what you expect?
 - JOBDEF RASSIGN=YES/NO controls whether new job ID is assigned or original preserved across NJE/offload
 - New job ID assigned if original job number is already in use or outside range

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



Session 16895

SPOOL Migration (\$MSPL) - Overview



- **\$M SPOOL command to move data off volume**
 - Faster than \$P SPOOL (minutes not days)
 - Enabled in JES2 release 1.13 via OA36158 (UA64366)
- **Works with active address spaces using volume**
 - Less activity is better/faster but no need to IPL to stop any active jobs
- **Goal of SPOOL migration is to stop using a SPOOL data set**
 - It is NOT to eliminate the internal representation of the volume
 - Old data set can be deleted and SPOOL volume take offline
- **After successful SPOOL migration**
 - \$DSPOOL still shows volume
 - \$DJQ,SPOOL= still displays volume
 - Volume status is MAPPED

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



SPOOL migration allows an installation a way to quickly move data off of a SPOOL volume in a period of minutes instead of the days that a drain command would take. The processing can be done with active address spaces still accessing the volume. The goal of the command is to get the source data set moved to either a new volume or merged onto an existing SPOOL volume. The internal representation of the volume will remain after it is merged onto an existing volume and will persist until all jobs that were using the volume have been purged. This implies that the volume will still be displayed in \$D SPOOL commands and in the volume list of a \$DJQ,SPOOL command. The status of the “remnant” volume will be MAPPED.

SPOOL Migration (\$MSPL) - Terminology



- **Source Volume** – the SPOOL to be migrated
- **Target Volume** – the SPOOL to receive the migrated data

- **MERGE migration** – copy a source volume to free space on an existing/active target volume in the JES2 configuration
- **MOVE migration** – copy an inactive source volume to a new target volume that is not in the JES2 configuration

- **Active migration** – a migration that is currently being processed
- **Migrator** – the JES2 member that coordinates the migration
- **Migration phase** – the current 'step' of the migration process

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



Here are some key terms used in a SPOOL migration process. Understanding these will help understand the explanations that follow and the explanations in the publications.

SPOOL Migration (\$MSPL) - Terminology



- **Mapped Volume** –
 - When a Merge Migration completes, the Source Volume becomes Mapped
 - Deleted when all jobs with space on it have been purged
 - No longer allocated to the SPOOL data set
 - SPOOL data set on the volume can be deleted
 - Physical device can be removed
- **Mapped Target** -
 - A volume with at least one Mapped Volume mapped onto it
- **Reserved** –
 - Attribute of any SPOOL set via \$TSPOOL(volser),RESERVED=YES|NO
 - Indicates if the SPOOL volume is selectable but not allocatable
 - Can be used to Reserve a volume for future Merge migration(s)
 - Reserved volumes have no entries in the BLOB

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



Reserved is a new attribute of a SPOOL volume. Similar to the draining state, a reserved volume can have work on it selected for processing. But unlike drain processing, there is no attempt by JES2 to get the data off the volume or to make it go away. It is useful for a standby volume during normal processing or to have a volume ready to receive data that is being merged onto it. The reserved attribute can be set via a \$T SPOOL,RESERVED=YES|NO.

SPOOL Migration (\$MSPL) - Phases



- **PENDING** – Command issued and queued for processing
 - **INITIALIZING** – Create data structures and subtasks
 - **SETUP** – Prepare source and target data set
 - **COPY** – First pass copy of all data from source to target
 - **CATCHUP** – Second pass copy of tracks updated by active applications
 - **CANCEL** – Error phase that synchronizes stopping migration
 - **BACKOUT** – Error phase to undo any work done in migration
 - **CLEANUP** – Delete data structures and end active migration
-
- Cancel can be requested up until start of **CATCHUP** phase
 - Internal cancel can occur later in error recovery cases
 - Phase start/end messages issued to SYSLOG
 - Can be sent to console with `DEBUG VERBOSE=YES`
 - Some source volume state changes occur before the **INITIALIZING** phase and after the **CLEANUP** phase

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



A normal SPOOL migration goes through 6 phase from start to finish. Each phase is indicated by a message in SYSLOG (or on the console if `DEBUG VERBOSE=YES` is specified). In the event of a migration being canceled (either by operator command or error processing) the 2 phases in blue are triggered. These phases can occur at any time during the migration process. All migration (including cancel processing) end with a cleanup phase.

Operator initiated cancels can be done up to the start of the **CATCHUP** phase. Internal CANCELS can be performed without loss of data in the early part of the **CATCHUP** phase. Internal error driven cancels after the start of the catchup phase can result in some loss of data.

There is some status changes that occur using normal SPOOL command processing that are part of the migration but actually occur outside these phases.

SPOOL Migration (\$MSPL) – Move vs Merge



- **MOVE** takes all data on an existing volume and moves it to a new one
 - Source must be INACTIVE (\$Z SPOOL done)
 - No active jobs on the volume
 - Target cannot be currently an active SPOOL volume
 - Can specify space to use to create data set on target
 - At the end of move, old (source) volume does not exist
 - Target after a move is active
- **MERGE** takes all data on one volume and merges it onto free space on another volume
 - Most flexible migration option
 - Source can be in any state with active jobs/address spaces
 - Less activity is good
 - Results is a mapped volume that goes away when all jobs using it are deleted
 - Similar to \$PSPL but device no longer in use

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



There are 2 forms of SPOOL migration, MOVE and MERGE. In a move migration, you take one existing, INACTIVE SPOOL volume and move it to a new volume that is not currently part of the SPOOL configuration. If you have 3 SPOOL volumes before a move, then you will have 3 SPOOL volumes after the move. As stated, for a move, the source volume must be INACTIVE (HALTED). There are other restrictions listed later for a move.

A merge migration takes the data on an existing SPOOL volume (in any state) and merges in into contiguous space on a target volume. With a merge, if you start off with 3 volumes before the merge you end up with 2 volumes after the merge. The 3rd volume will display but it is not being used. It is considered mapped.

Merge is the least restrictive process. Any source volume can be merged to an appropriate target volume.

SPOOL Migration (\$MSPL) – Move vs Merge



- **General restriction (for move and merge migrations):**
 - The *Source Volume* cannot be a *Mapped Target*
 - Cannot merge A to B and then move or merge B to C
 - Once A no longer exists then B can be migrated to C
 - The *Source Volume* cannot be actively migrating or extending
 - The track size of the *Target Volume* cannot be less than the *Source Volume*
 - The *Source Volume* cannot be stunted
 - All MAS members must be JES2 z/OS 1.13 & z11 checkpoint mode
- **Each SPOOL migration requires a separate XCF group**
 - Used to manage messages for each unique migration
 - JES2 limits migration to 5 concurrent migrations per MAS
 - Use D XCF,COUPLE to display MAXGROUP formatted in CDS
 - Group name is SYSMGxxx
 - xxx is the decimal source SPOOL extent

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



These are general restrictions for MOVE and MERGE. The source and target cannot be being processed by another SPOOL command. The target track size cannot be smaller than source track size (merge is on a track by track basis). You must be in z11 CKPT mode (and cannot go to z2 mode once a migration has been requested). All members must be z/OS 1.13.

Also, there can be only one level of mapped volumes. You cannot merge A to B and then try to merge B to C with A still mapped to B.

Each SPOOL migration uses a separate XCF group during the active part of the migration. You must have sufficient space in your XCF couple data set to create the new group. JES2 will only do 5 active migration at a time in the MAS to limit the number of groups JES2 uses (there is no programmatic way to query how many free groups are left).

SPOOL Migration (\$MSPL) – Move vs Merge



- **Move migration moves an INACTIVE volume to a new volume**
- **Upon successful completion**
 - The *Source Volume* no longer exists
 - The *Target Volume* exists and is active
 - Could be RESERVED if requested on \$M SPOOL command
- **Source Volume STATUS= values**
 - INACTIVE ->MIGRATING ->does not exist
- **Target Volume STATUS= values**
 - Does not exist ->ACTIVE
- **Additional move migration restrictions**
 - The *Source Volume* must be INACTIVE
 - *Source Volume* cannot be in Absolute format (instead, do a merge)
 - The Target Volume will inherit the Source Volume Tracks per Track Group value

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



A move migration is described here. Move replaces one volume with another volume. At the beginning of the move, the source volume must be inactive and the target does not exist in the SPOOL configuration. At the end of the move, the source volume no longer exists in the SPOOL configuration and the target volume is active with all the data from the source volume. Most of the properties of the target are inherited from the source.

The restrictions specific to a move migration are listed.

SPOOL Migration (\$MSPL) – Move vs Merge



- **Merge migration moves a *Source Volume* to an free space on an active *Target Volume***
- **Upon successful completion**
 - The *Source Volume* still exists but is STATUS=MAPPED
 - Still displays in \$DSPOOL and in \$DJQ,SPOOL lists
 - The *Target Volume* is a mapped on volume
- **Source Volume STATUS=values**
 - INACTIVE ->MIGRATING ->MAPPED
- **Additional merge migration restrictions**
 - The *Target Volume* must be *Active* (can be *Reserved*).
 - The *Target Volume* cannot be stunted.
 - The *Target Volume* must use relative addressing

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



A merge migration merges data from a source volume (in any state) to available free space on an existing target volume. At the end of the merge, the source is mapped and is selectable if the target is selectable. The source still appears in \$D SPOOL and \$DJQ,SPOOL command until all jobs that have space on the volume have purged. However, JES2 is no longer allocated to the source SPOOL data set and it can be deleted.

There are no additional restrictions on the source volume, but there are some restrictions on the target.

SPOOL Migration (\$MSPL) – Usage



- **\$M SPOOL command syntax (merge)**
\$M SPOOL(*volser*),TARGET=*target*
- **\$M SPOOL command syntax (move)**
\$M SPOOL(*volser*),TARGET=*target*
[,SPACE=(CYL|TRK|MAX,*size*)]
[,DSNAME=*dsname*]
[,RESERVED]
- **\$M SPOOL cancel command**
\$M SPOOL(*volser*),CANCEL
- **Multi-source move command is also supported**
\$M SPOOL(*volser1*,*volser2*,*volser3*...),TARGET=*volser*
- 1st volume can be a move or a merge, remainder are merges
- Migration happens 1 volume at a time (one per target)

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



Session 16895

The command syntaxes are listed here. A merge supports a source and target volume. Move supports specifying the target size, data set name, and reserved status. The move can create the SPOOL data set on the target in the same way a \$S SPOOL command can.

To cancel a SPOOL migration, specify the source volser whose migration you want to cancel



SPOOL Migration (\$MSPL) – Usage

- **\$D SPOOL...,MIGDATA** helps determine migration requirements

- SPACE_USED is high water mark of used space on volume
- LARGEST_FREE is largest contiguous free space on the volume

\$D SPOOL,MIGDATA

```
$HASP893 VOLUME (SPOOL2)  MIGDATA=(SPACE_USED=433410,  
$HASP893                    LARGEST_FREE=16590)  
$HASP893 VOLUME (SPOOL5)  MIGDATA=(SPACE_USED=418215,  
$HASP893                    LARGEST_FREE=31785)
```

- Display all volumes having contiguous free space greater than 17000 tracks:

\$D SPOOL,MIGDATA=LARGEST_FREE>17000,MIGDATA

```
$HASP893 VOLUME (SPOOL5)  MIGDATA=(SPACE_USED=418215,  
$HASP893                    LARGEST_FREE=31785)
```

- Note: Track groups in BLOB are considered to be used (not free)

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



Session 16895

The \$D SPOOL,MIGDATA command will provide information you need to determine what volumes can be merged or moved where. It lists the current high water mark on each volume and the largest contiguous free space. This can be used to determine where to move a SPOOL volume to. One note, SPOOL space in the BLOB is considered used when listing used and available space. To get space out of the BLOB, set the volume to reserved (but do not reserve all volumes because then SPOOL will be considered full).



SPOOL Migration (\$MSPL) – Usage

- **Use \$D SPOOL to monitor a migration and check results**
- **Display all spool volumes that are currently migrating**

```
$D SPOOL(*),STATUS=MIGRATING
$HASP893 VOLUME (SPOL7) STATUS=MIGRATING-MOVE, TARGET=SPOL2
$HASP893 VOLUME (SPOL8) STATUS=MIGRATING-MOVE, TARGET=SPOL2
$HASP646 75.0000 PERCENT SPOOL UTILIZATION
```
- **Display all spool volumes that have a Target (are Mapped)**

```
$D SPOOL, TARGET^=' ', TARGET
$HASP893 VOLUME (SPOL4) TARGET=SPOL11
$HASP646 80.0000 PERCENT SPOOL UTILIZATION
```
- **Display all spool volumes that are reserved**

```
$DSPL, RESERVED=YES
$HASP893 VOLUME (SPOL2) STATUS=RESERVED, PERCENT=20
$HASP893 VOLUME (SPOL3) STATUS=RESERVED, PERCENT=40
$HASP646 30.0000 PERCENT SPOOL UTILIZATION
```
- **\$D SPOOL,MPERCENT** – Displays percent of migration that is complete
- **\$D SPOOL,PHASE** – Displays current migration phase

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



Session 16895

Here are some useful commands that can be used to display information about SPOOL volumes.



SPOOL Migration (\$MSPL) – Common Problems

- What is MOVE vs MERGE?
 - The primary difference relates to the state of the respective source and target volumes. MOVE copies an inactive source volume to a target not currently in the configuration. MERGE copies a source volume to a target that exists in the configuration.
- What does volume STATUS=MAPPED mean?
 - The source SPOOL volume has been migrated. *Logically*, it remains until all jobs residing on it are purged. *Physically*, the data set and volume are not being used
- When can I delete the source SPOOL data set?
 - After successful completion of migration (STATUS=MAPPED)
- When can I disconnect/repurpose my source SPOOL volume?
 - After successful completion of migration (STATUS=MAPPED)

- **REFERENCE**

http://www-03.ibm.com/systems/z/os/zos/jes2_spoolmigration.html

SPOOL Migration FAQ is really useful!

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



Session 16895

Mirroring & DR Environments - Overview



- **Copying SPOOL and checkpoint data outside of JES2 management / jurisdiction**
- **SPOOL and checkpoint must be copied together asynchronously**
 - Same consistency group
 - All data must be copied (ie including record 0 etc)

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



Session 16895

There are many products available to asynchronously copy SPOOL and checkpoint data. The concepts discussed here are applicable to any product attempting to move JES2 data; however, some we will draw some specific notes to SDM and XRC

Mirroring & DR Environments - Terminology

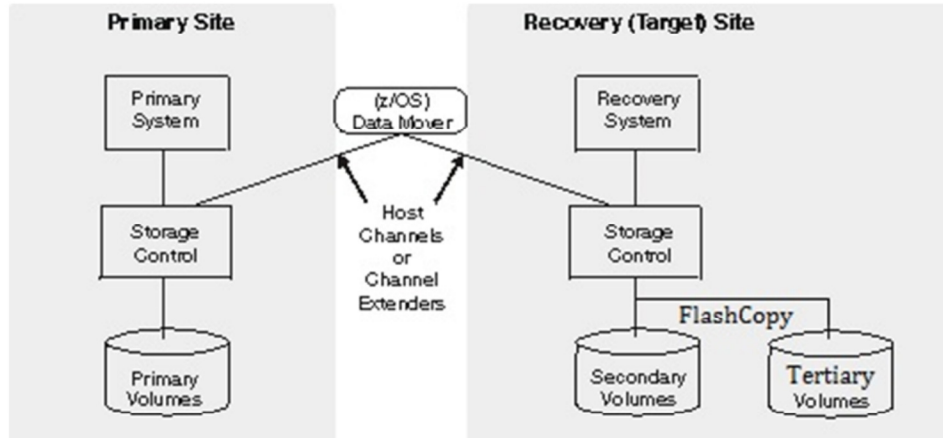


- **Primary volumes** – DASD volumes located at production site containing original production data.
- **Secondary volumes** – DASD volumes located at DR site that contain an asynchronous (delayed) copy of the production data.
- **Tertiary volumes** – DASD volumes located at DR site that contain a FlashCopy of the secondary volumes at a specific consistency timestamp.
- **Consistency groups** – Grouping of writes to the secondary volumes arranged together to ensure the order of updates are preserved within the same consistency timestamp.

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



Mirroring & DR Environments – XRC Setup



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



Mirroring & DR Environments – Common Problems



- Define SPOOL and checkpoint volumes within the same XRC session to ensure they will be in the same consistency group.
- Do not IPL the DR system from secondary volumes! IPL from the tertiary volumes to preserve the primary/secondary mirror relationship.
- Check ANTC8414W / ANTV8110I messages to verify XRC sessions are INTERLOCKED within the same consistency group when the FlashCopy is performed to the tertiary volumes.
- Check ANTR9102I messages to make sure XRECOVER was successfully performed after the FlashCopy.

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



Session 16895

Mirroring & DR Environments – JES2 symptoms



- **What is acceptable/reasonable to see?**
 - Running JES2 in duplex mode with primary CKPT on CF and secondary CKPT on DASD
 - The duplex copy (CKPT2) can be up to 10 writes behind primary CKPT
 - Mirroring from CKPT2, it is *acceptable (albeit still unlikely)* to encounter a small handful of **CBIMPLxx** errors
- **What is not acceptable/reasonable to see?**
 - \$HASP486 JES2 CKPTx DATA SET – DAMAGED RECORD DETECTED
 - \$HASP096 DISASTROUS ERROR AT SYMBOL BERTERR
 - JES2 not starting or immediately terminating following start

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



If JES2 is started using the duplexed CKPT, then it may be up to 10 writes behind (relative to primary CKPT on CF). Therefore, there could be a handful of differences between SPOOL and checkpoint that can result in CBIMPLxx errors – typically reflecting that job is in a different state than expected and/or trackgroup is assigned to a different job than expected.

Any of the *not* acceptable symptoms are indicative that there is a significant mismatch between SPOOL and checkpoint data – the copy was not complete. If the source JES2 system is operating normally, then there is nothing to suggest any actual damage to SPOOL or checkpoint data.

Questions?



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

