



B72

IMS Database Image Copy 2 - the Complete Story

Pat Schroeck

IMS
Technical Conference

Sept. 27-30, 2004

Orlando, FL

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States or other countries or both:

DFSMS

DFSMSdss

Enterprise Storage Server

IBM®

IMS

IMS/ESA

RAMAC®

IMS Database Image Copy 2 -- Topics

- Database Image Copy 2 overview
 - ▶ What is Database Image Copy 2
 - ▶ Key benefits
 - ▶ Capability
 - ▶ DFSMS interface
- Invoking the Database Image Copy 2 utility
- IMS Version 8 enhancements
- GENJCL support

Database Image Copy 2 Overview

- The Database Image Copy 2 utility (DFSUDMT0) invokes DFSMS Concurrent Copy to copy IMS databases
 - ▶ Utility was introduced in IMS/ESA Version 6
- Takes advantage of DFSMS Concurrent Copy and Virtual Concurrent Copy (Snapshot)
 - ▶ For databases that reside on supported hardware
 - Enterprise Storage Server (ESS) -- Shark
 - 3990 Storage Control with Concurrent Copy feature
 - RAMAC Virtual Array (RVA) with Snapshot capability
- Consistent ('clean') or 'fuzzy' copies are produced, per user specification

Why use Image Copy 2?

- For 'clean' copies, reduces the time that the database or area must be offline (or not available for update)
- 'Fuzzy' KSDS image copies can be taken
 - ▶ The Database Image Copy utility (DFSUDMP0) does not support fuzzy KSDS copies
- Supports data sharing environment
 - ▶ Online Database Image Copy (DFSUICP0) does not support data sharing
- Up to 4 output copies can be created
 - ▶ Primary and secondary copies are registered with DBRC
 - ▶ Additional copies may be used for 'second site'/disaster recovery purposes

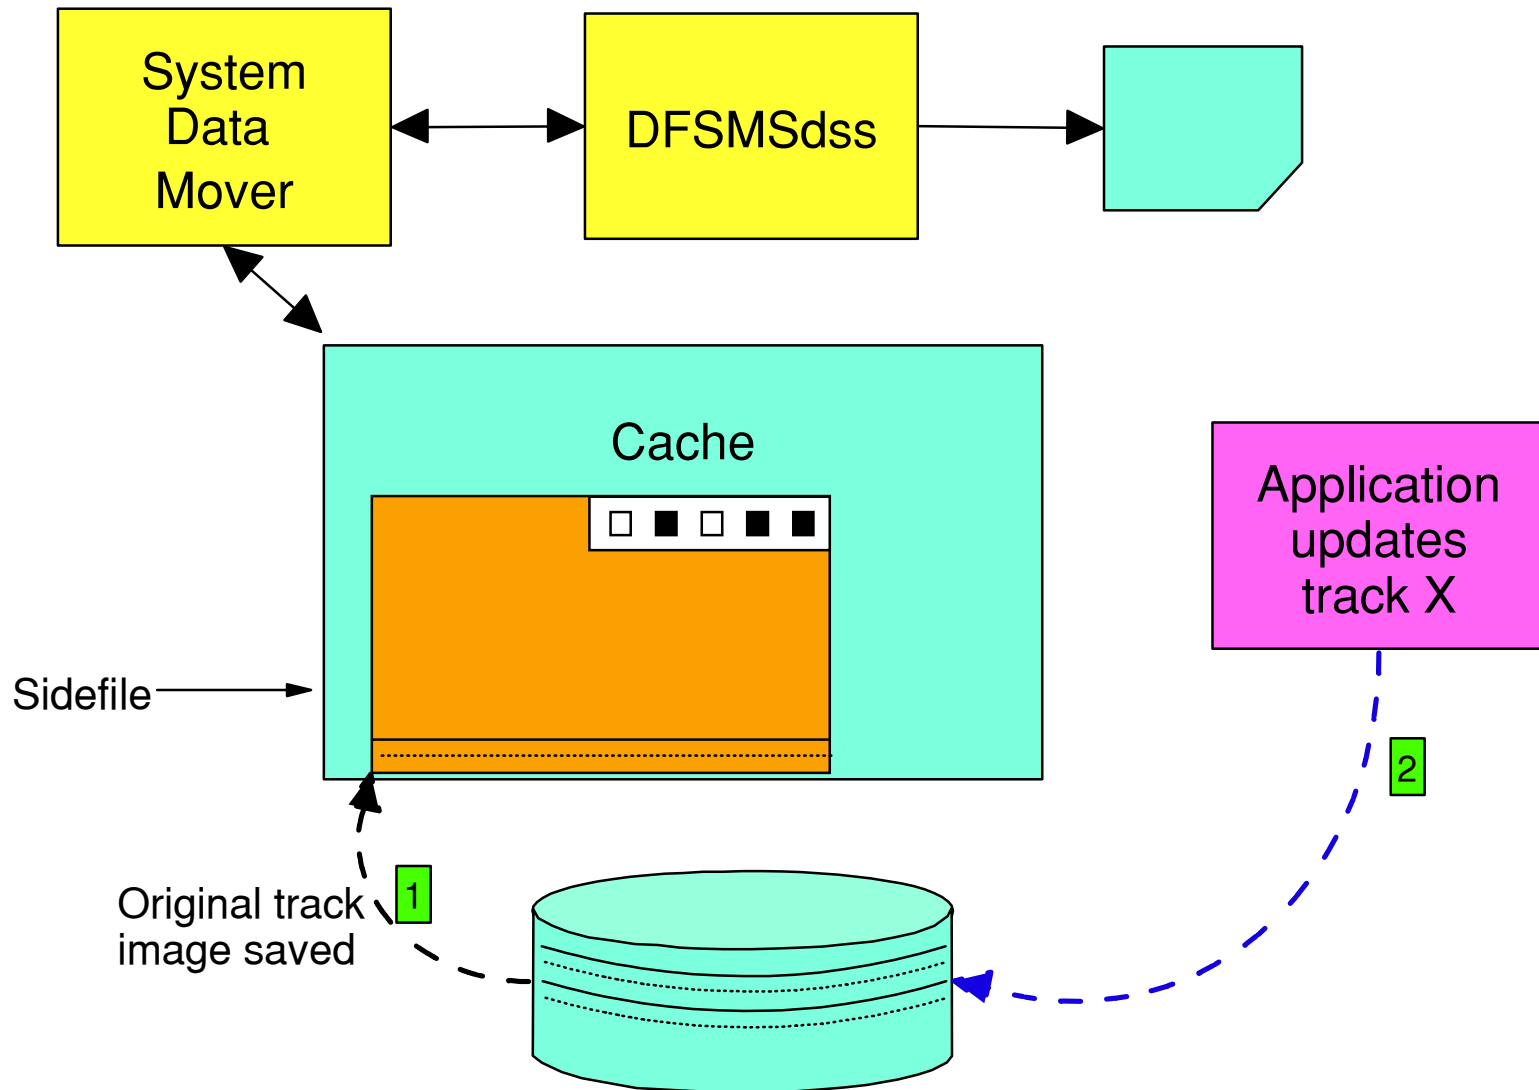
IMS Image Copy utilities

	Database Image Copy (DFSUDMP0)	Online Database Image Copy (DFSUICP0)	Database Image Copy 2 (DFSUDMT0)	High Performance Database Image Copy for z/OS (Tool)
Environment	batch	online (BMP)	batch	batch
Full function or FP	both	Full function only	both	both
Fuzzy KSDS copies	N	Y	Y	N
Reduced DB outage for clean copies	N	N	Y	N
Supports datasharing for fuzzy copies	Y	N	Y	Y
Hardware requirement	N	N	Y	N
Multiple data sets copied in parallel	N	N	Y (V8) N (V7)	Y
Compressed output	N	N	Y	Y
Integrated pointer checking	N	N	N	Y
Stacked output on tape	N	N	N	Y
Dynamic allocation of input and output	N (Y- for HALDB input in V8)	N	N (Y - for HALDB input in V8)	Y

Reduced DB outage for clean copies

- DFSMS Concurrent Copy processes in two phases
 - ▶ Logical copy phase (concurrent copy initialization) -- takes a very quick 'snapshot' of the data
 - ▶ Physical copy phase -- produces the physical dump data set(s)
- Database/area can be made available for update after logical copy phase has completed
 - ▶ Consistent copy is produced even though updates occur during the physical copy phase
- Logical copy complete notification is provided by the DFS3121A message
- The option exists to disallow update access until physical copy completion
 - ▶ Use when a clean copy taken at a specific point in time is required

DFSMS Concurrent Copy Write Interception



Reduced DB outage for clean copies...

- Virtual Concurrent Copy is used for databases that reside on RAMAC Virtual Array (RVA) DASD
 - ▶ Virtual Concurrent Copy is also called 'Concurrent Copy-compatible Snapshot'

- Virtual Concurrent Copy uses Snapshot capability
 - ▶ Logical copy phase consists of taking a Snapshot copy of the data set
 - ▶ Snapshot copy is offloaded to the image copy output data set(s) during the physical copy phase
 - ▶ Use of Snapshot vs. Concurrent Copy is transparent to the IMS utility
 - handled by DFSMSdss
 - ▶ Snapshot copy is not visible to the Image Copy 2 utility
 - is not externalized or recorded in the RECON

Image Copy 2 copy options

- Summary of copy options
 - ▶ **Clean copy -- allow updates after logical copy**
 - reduced database downtime
 - 'XL' option (eXclusive, release after Logical copy)
 - ▶ **Clean copy -- disallow updates until after physical copy**
 - like Database Image Copy (DFSUDMP0) without CIC (concurrent image copy) option
 - 'XP' option (eXclusive, release after Physical copy)
 - ▶ **Fuzzy copy -- no impact on database availability**
 - 'S' option (Shared)
- 'SMS' image copy types
 - ▶ **SMSNOCIC** -- image copy created by the Image Copy 2 utility while update access to the database was not allowed
 - ▶ **SMSCIC** -- image copy created by Database Image Copy 2 concurrent with update access

Database Image Copy 2 -- general

- Provides data set level processing
 - ▶ Does not support volume dump/restore
- Copies a single data set per execution [V7]
- Copies multiple data sets per execution [V8]
- Up to 4 output copies can be created
 - ▶ Primary and secondary copies are registered with DBRC
 - ▶ Additional copies may be used for 'second site'/disaster recovery purposes
- Image copies are in DFSMS dump format
- Image copies can be in compressed format [V7 and above]

Database Image Copy 2 -- general...

- DFSUDMT0 invokes DFSMSdss for DUMP specifying the CONCURRENT (CC) option
- DFSMSdss is invoked through the DFSMSdss cross memory API
 - ▶ via ADRXMAIA
 - ▶ DFSMSdss processing runs in a separate 'server' address space (IEESYSAS)

Image Copy 2 setup requirements

- Database data sets must reside on
 - ▶ Enterprise Storage Server (Shark), or
 - ▶ 3990 Storage Control with Concurrent Copy capability, or
 - ▶ RAMAC Virtual Array subsystem with Snapshot capability

- Databases and areas must be registered with DBRC

- For fuzzy KSDS image copy
 - ▶ KSDS must be SMS-managed
 - ▶ BWO(TYPEIMS) must be specified for the KSDS via DEFINE or ALTER CLUSTER
 - BWO = backup while open

Database Image Copy 2 -- invocation

- Database Image Copy 2 (DFSUDMT0) executes as an offline IMS utility
- DFSUDMT0 executes under the IMS region controller -- stand-alone mode is not supported

```
//TOCOPY1 EXEC PGM=DFSRRRC00, PARM='ULU', DFSUDMT0, . . .
```

- DBRC is required for execution
- Image copy options are specified on the utility control (SYSIN) statement(s)
 - ▶ For V7, only one control statement can be specified
 - ▶ For V8, multiple control statements can be specified

Image Copy 2 -- invocation (cont'd.)

Format of the 'DBDS select' control statement:

position	description
-----	-----
1	
2	Number of output copies
3	
4 - 11	Input database name
12	
13 - 20	Input DD name or area name
21	
22 - 29	Copy 1 (primary output copy) DD name
30	
31 - 38	Copy 2 (secondary output copy) DD name
39	
40 - 47	Copy 3 DD name
48	
49 - 56	Copy 4 DD name

Image Copy 2 -- invocation...

Format of the 'DBDS select' control statement
(cont'd.):

position	description
-----	-----
57	
58	<u>S</u> X -- S (for "Shared") specifies a fuzzy copy (SMSCIC); X (for "eXclusive") specifies a clean copy (SMSNOCIC)
59	<u>L</u> P -- for clean copy, updates are to be allowed after Logical copy (L) or after Physical copy (P); ignored for fuzzy copy
60	C -- compress output copies [V7 and above]
61	1 2 3 4 -- DFSMSdss OPTimize value [V8]
62 - 72	Reserved

Recovery with 'SMS' image copies

- The Database Recovery utility (DFSURDB0) provides recovery from SMSCIC and SMSNOCIC image copies
 - ▶ Must be run with DBRC to process 'SMS' image copies
 - ▶ Existing JCL and/or skeletal JCL for GENJCL.RECOV can be used
- IMS Online Recovery Service (ORS) also provides recovery from 'SMS' image copies
- DFSMSdss is invoked to RESTORE the data set
 - ▶ Change accumulation data is not merged with image copy data during restore -- change accum is applied after restore

Image Copy 2 -- V8 enhancements

- Provide better coordination and management
 - ▶ User can /DBR the databases, run one image copy job, /STA the databases
- Multiple database data sets can be copied in a single utility execution
- Group name support
- Single output data set for multiple image copies
- User can specify the DFSMSdss OPTIMIZE option

Multiple utility control statements

- Multiple database data sets can be copied in a single utility execution
 - ▶ All DBDSs specified on the control statements are passed to DFSMSdss on a single invocation
 - ▶ DFSMSdss starts multiple dump processes in parallel
 - ▶ Logical completion is achieved in a brief period of time
 - *However, there is a limitation*
- Current DFSMSdss limitation:
 - ▶ One image copy step cannot stack multiple output image copy data sets on a tape volume!!
 - ▶ The Same Dataset option (see later charts) can provide an alternative

Multiple utility control statements...

- Notification of 'logical copy complete' is given for the database/area
 - ▶ DFS3121A message is issued when all DBDSs for a database or HALDB partition are logically complete
 - *(unless group name statement supplied -- see later charts)*
 - message is changed from Version 7
 - ▶ DFS3121I messages for the DBDSs that are logically complete are written to SYSPRINT
 - ▶ DFS3121A is only issued for 'XL' option, not for 'XP' or 'S' options
 - XL -- clean IC, DBs can be started after logical copy complete
 - XP -- clean IC, DBs cannot be started until physical copy complete
 - S -- fuzzy IC

Multiple utility control statements...

- Logical completion notification -- example

```
DFS3121A LOGICAL COPY COMPLETE FOR DB/AREA dbname1; n OF m  
DATA SETS FAILED
```

```
DFS3121I COPIED DB/AREA dbname1 DDN ddname1 DSN dsname1  
DFS3121I COPIED DB/AREA dbname1 DDN ddname2 DSN dsname2  
DFS3121I COPIED DB/AREA dbname1 DDN ddname3 DSN dsname3  
DFS3121I COPIED DB/AREA dbname1 DDN ddname4 DSN dsname4
```

Multiple utility control statements...

- Notification of 'physical copy complete' is given for the database/area
 - ▶ New message, DFS3141A, is issued when all DBDSs for a database or HALDB partition are physically complete
 - (*unless group name statement supplied -- see later charts*)
 - ▶ DFS3141I messages for the DBDSs that were successfully copied are written to SYSPRINT
 - ▶ DFS3141A is only issued for 'XP' option, not issued for 'XL' or 'S'

DFS3141A PHYSICAL COPY COMPLETE FOR DB/AREA dbname1; n OF m
DATA SETS FAILED

DFS3141I COPIED DB/AREA dbname1 DDN ddname1 DSN dsname1

DFS3141I COPIED DB/AREA dbname1 DDN ddname2 DSN dsname2

DFS3141I COPIED DB/AREA dbname1 DDN ddname3 DSN dsname3

DFS3141I COPIED DB/AREA dbname1 DDN ddname4 DSN dsname4

Multiple utility control statements -- Example

```

|...+...1...+...2...+...3...+...4...+...5...+...6...
//SYSIN      DD      *
2 DBNAME1    DDNAME1A ICOUT1A1 ICOUT1A2          XLC
2 DBNAME1    DDNAME1B ICOUT1B1 ICOUT1B2          XLC
2 DBNAME1    DDNAME1C ICOUT1C1 ICOUT1C2          XLC
2 DBNAME1    DDNAME1D ICOUT1D1 ICOUT1D2          XLC
2 DBNAME1    DDNAME1E ICOUT1E1 ICOUT1E2          XLC
2 DBNAME2    DDNAME2A ICOUT2A1 ICOUT2A2          S C
2 DBNAME2    DDNAME2B ICOUT2B1 ICOUT2B2          S C
2 DBNAME3    DDNAME3A ICOUT3A1 ICOUT3A2          S C
2 DBNAME4    DDNAME4A ICOUT4A1 ICOUT4A2          XP
2 DBNAME4    DDNAME4B ICOUT4B1 ICOUT4B2          XP
2 DBNAME4    DDNAME4C ICOUT4C1 ICOUT4C2          XP
/*

```

The example shows utility control statements to copy DBDSs for 4 different databases. Different processing options (S|XL|XP) are specified for the DBDSs. When the image copies for the DBDSs for DBNAME1 are all logically complete, a DFS3121A message is issued for DBNAME1. When the image copies for the DBDSs for DBNAME4 are all physically complete, a DFS3141A message for DBNAME4 is issued. Completion messages are not issued for DBNAME2 or DBNAME3.

Multiple utility control statements...

- One processing option (XL|XP|S) is used for all DBDSs of a database
 - ▶ If 'XP' was specified for any DBDS, 'XP' processing occurs for the database (DFS3141A message is issued)
 - ▶ Else, if 'XL' was specified for any DBDS, 'XL' processing occurs for the database (DFS3121A message is issued)
 - ▶ Else, 'S' was specified (or defaulted) for all DBDSs (neither completion message is issued)

Multiple utility control statements...

- Fuzzy KSDS image copy processing can fail because of update activity
 - ▶ Processing fails if a CI or CA split occurs during the logical copy phase
 - DFS3145A message is issued
 - User response: Rerun image copy job at a later time
 - ▶ V7 --> Utility retries the DFSMSDss DUMP command up to 10 times before failing
 - ▶ V8 --> Dump command is retried up to 10 times if KSDS is the only data set being copied
 - no retry if multiple DBDSs are being copied in the Image Copy 2 execution

Group name support

- Group name control statement
 - ▶ A new control statement can be used to assign a group name to the collection of DBDSs that are to be copied
 - ▶ The group name statement is followed by control statements identifying the DBDSs
 - Group name statement must be the first control statement
 - ▶ Processing options specified (or defaulted) on the group statement override options specified for the DBDSs
 - S|XL|XP, compress option, optimize option
 - ▶ Completion notification is reported for the group, not for the DB/area
 - E.g., DFS3121A LOGICAL COPY COMPLETE FOR GROUP
grpname...

Group name support...

- **The group name specified for Image Copy 2 is not required to be a DBRC DB or DBDS group**
 - ▶ The utility does not check for or use DB or DBDS groups
- **Using DBRC groups, however, can simplify operations**
 - ▶ /DBR DATAGROUP AAA101
 - ▶ GENJCL.IC GROUP(AAA101)
 - the generated jcl includes a group name statement for AAA101
 - ▶ /STA DATAGROUP AAA101

Group name support -- example

```
|...+...1...+...2...+...3...+...4...+...5...+...6...
//SYSIN      DD  *
G  GROUPXYZ                                     XL
  2 DBNAME1   DDNAME1A  ICOUT1A1  ICOUT1A2
  2 DBNAME1   DDNAME1B  ICOUT1B1  ICOUT1B2
  2 DBNAME1   DDNAME1C  ICOUT1C1  ICOUT1C2
  2 DBNAME1   DDNAME1D  ICOUT1D1  ICOUT1D2
  2 DBNAME1   DDNAME1E  ICOUT1E1  ICOUT1E2
  2 DBNAME2   DDNAME2A  ICOUT2A1  ICOUT2A2
  2 DBNAME2   DDNAME2B  ICOUT2B1  ICOUT2B2
  2 DBNAME3   DDNAME3A  ICOUT3A1  ICOUT3A2
  2 DBNAME4   DDNAME4A  ICOUT4A1  ICOUT4A2
  2 DBNAME4   DDNAME4B  ICOUT4B1  ICOUT4B2
  2 DBNAME4   DDNAME4C  ICOUT4C1  ICOUT4C2
/*
```

The example shows utility control statements to copy DBDSs for 4 different databases. The DBDSs to be copied are identified as members of group GROUPXYZ. All of the data sets will be copied using the same processing options (XL). When the image copies for all of the DBDSs are logically complete, one DFS3121A message will be issued for the group GROUPXYZ.

Single output data set for multiple image copies

- **Option is supported to create multiple image copies in one output data set**
 - ▶ Different than stacking multiple image copies on one output volume
- **Uses DFSMSdss capability to dump multiple input data sets into one output data set**
 - ▶ Multiple dumps are written serially to the output data set
 - ▶ The user specifies the **Same Dataset** (or **SameDS**) option on the DBDS control statement
 - 'S' in column 1
 - ▶ Up to 255 DBDS can be dumped to the same output data set
 - ▶ Not supported for DBDSs defined in the RECON as REUSE
 - ▶ IMS Online Recovery Service (ORS) recognizes this form of 'stacked' ICs and schedules a single restore operation

Single output data set for multiple ICs...


■ SameDS -- advantage

- ▶ Logical completion for all the data sets occurs quickly
 - Not overlapped with writing the physical copies

■ SameDS -- disadvantage

- ▶ Extended database recovery time
 - Restore of n th data set requires reading the image copies for all $n-1$ data sets dumped ahead of it

DUMP (DB01, DB02, DB03, ..., DBnn)



single output data set

Single output data set -- example

```
|...+...1...+...2...+...3...+...4...+...5...+...6...
//SYSIN      DD      *
  2 DBNAME1    DDNAME1A ICOUT101 ICOUT102          XL
S   DBNAME1    DDNAME1B                               XL
S   DBNAME1    DDNAME1C                               XL
S   DBNAME1    DDNAME1D                               XL
S   DBNAME1    DDNAME1E                               XL
  2 DBNAME2    DDNAME2A ICOUT201 ICOUT202          S
S   DBNAME2    DDNAME2B                               S
S   DBNAME3    DDNAME3A                               S
  2 DBNAME4    DDNAME4A ICOUT4A1 ICOUT4A2          XP
  2 DBNAME4    DDNAME4B ICOUT4B1 ICOUT4B2          XP
  2 DBNAME4    DDNAME4C ICOUT4C1 ICOUT4C2          XP
/*
```

The example shows utility control statements to copy DBDSs for 4 different databases. The five DBDSs for DBNAME1 are to be dumped to the ICOUT101 and ICOUT102 data sets. The two DBDSs for DBNAME2 and one DBDS for DBNAME3 are to be dumped to the ICOUT201/202 data sets. The DBDSs for DBNAME4 are each dumped into their own image copy output data sets.

Support DFSMSdss OPTIMIZE option

- DFSMSdss supports 4 levels of optimization
 - ▶ The four levels control the number of tracks of DASD which are transferred by one I/O request/command
 - OPTIMIZE(1) -- 1 track
 - OPT(2) -- 2 tracks
 - OPT(3) -- 5 tracks
 - OPT(4) -- 1 cylinder
 - ▶ The user can specify the optimization level, 1|2|3|4, on the group name or DBDS control statement (column 61)
 - ▶ The default is OPTIMIZE(1) for a fuzzy IC, OPTIMIZE(4) for clean ICs
 - These OPTIMIZE specifications are used in V6 and V7

GENJCL.IC support

- One skeletal IC JCL member supports both utilities, DFSUDMP0 and DFSUDMT0
- IBM-supplied skeletal JCL member, ICJCL, was modified for Version 8
 - ▶ Version 8 users must use the new member or modify their existing skeletal JCL members to use the new format
- SMSCIC and SMSNOCIC keywords specify that JCL to invoke Image Copy 2 is to be generated
 - ▶ Mutually exclusive with existing specifications --
 - CIC | NOCIC | SMSCIC | SMSNOCIC
 - ▶ SMSNOCIC indicates that update access is not allowed while the copy is being taken
 - 'X' is generated on the Image Copy 2 control statement
 - ▶ SMSCIC indicates that the IC will be taken concurrent with updates
 - 'S' is generated on the control statement

GENJCL.IC support...

- SMSNOCIC and SMSCIC allow DFSMSdss options to be specified
 - ▶ SMSNOCIC(*smsopts*) or SMSCIC(*smsopts*) where *smsopts* are:
 - ▶ Compress [V7 and above]
 - SMSCIC(COMPRESS) | SMSNOCIC(COMPRESS)
 - ▶ Sameds and OPTIMIZE value (N=1|2|3|4) [V8 and above]
 - SMSCIC(COMPRESS,SAMEDS,3)
 - ▶ Options can be specified in any order
 - SMSNOCIC(4,S,C)
- For SMSNOCIC, the DBREL(L|P) parameter specifies the 'DB release' option for the control statement(s)
 - 'L' -- updates allowed after logical copy
 - 'P' -- updates allowed after physical copy

GENJCL.IC support...

- Support for 3 or 4 output copies is provided for database data sets defined with NOREUSE attribute
 - ▶ COPIES(3) and COPIES(4) can be specified along with
 - UNIT3 or UNIT4
 - VOLLIST3 or VOLLIST4
- **IMS V8 extensions**
 - ▶ GENJCL.IC for a DBDS group or database with ONEJOB and either SMSNOCIC or SMSCIC specified generates JCL to copy all of the DBDSs in one execution
 - one Image Copy 2 step (with multiple control statements) instead of multiple Image Copy 2 steps in one job
 - ▶ The group or database named on the command is specified in the utility control statements

GENJCL.IC support...

GENJCL.IC example [V8]

GENJCL command:

```
GENJCL.IC DBD(DBDVSAMB) SMSNOCIC(COMPRESS,3) DBREL(L) -  
ONEJOB COPIES(4)
```

Generated JCL:

```
//IC1      EXEC  PGM=DFSRR00,REGION=800K,  
//          PARM='ULU,DFSUDMT0,,,,,,,,,,,,,Y,,,,,,,,,'  
          .  
          .  
          .  
//SYSIN  
G  DBDVSAMB                                XLC3  
  4  DBDVSAMB DDVSAMB1 D1OU1      D2OU1      D3OU1      D4OU1  
  4  DBDVSAMB DDVSAMB2 D1OU2      D2OU2      D3OU2      D4OU2  
/*
```

GENJCL.IC support...

GENJCL.IC example [V7]

GENJCL command:

```
GENJCL.IC DBD(DBDVSAMB) DDN(DDVSAMB1) SMSNOCIC(COMPRESS) -  
          DBREL(L) COPIES(4) VOLLIST1(IC2001) -  
          VOLLIST2(IC2002) VOLLIST3(IC2003) VOLLIST4(IC2004)
```

Generated JCL:

```
//IC1      EXEC  PGM=DFSRRRC00,REGION=800K,  
//          PARM='ULU,DFSUDMT0,,,,,,,,,,,,,Y,,,,,,,,,'  
          .  
          .  
          .  
//SYSIN  
  4 DBDVSAMB DDVSAMB1 DATAOUT1 DATAOUT2 DATAOUT3 DATAOUT4 XLC  
/*
```

IMS Database Image Copy 2 -- Summary

- IMS Database Image Copy 2 exploits DFSMS Concurrent Copy and DFSMS Snapshot
- Improves DB availability when taking clean image copies
- Provides support that is missing from other IMS image copy utilities
 - ▶ Supports KSDS fuzzy image copies
 - ▶ Supports SYSPLEX data sharing
- Version 8 enhancements provide improved useability