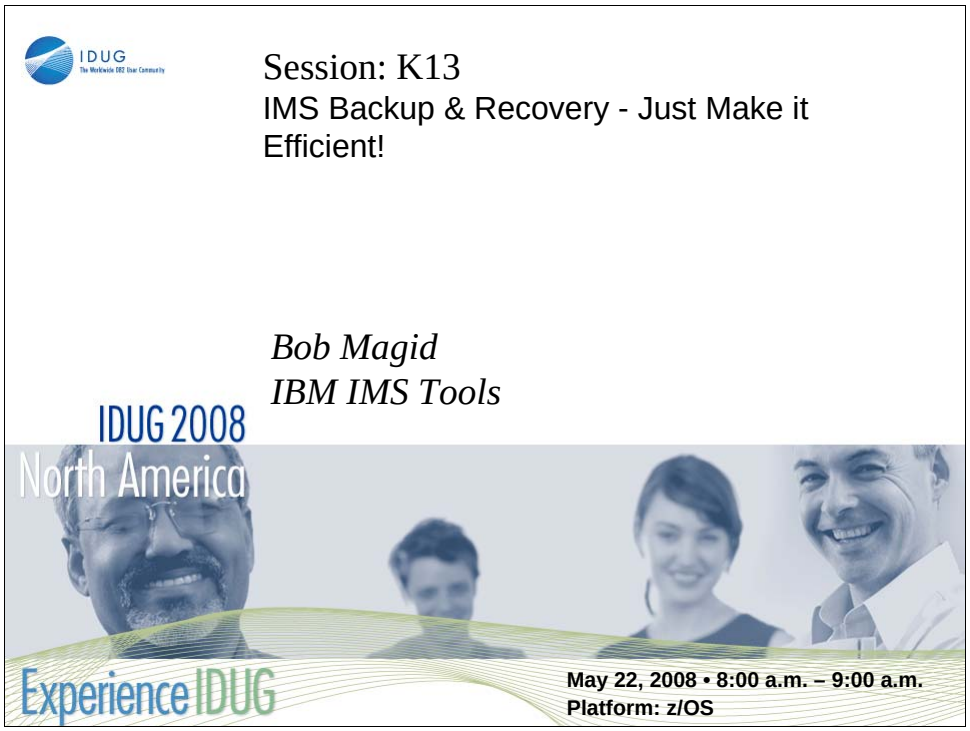

You need the highest quality of services for your mission-critical IMS databases. IBM provides high-quality IMS Tools to give you the best support to reach the goal of reducing the data unavailability period by also providing constant integrity. This session puts into perspective the different solutions for doing database backup and recovery and helps you choose the best one!

Session objectives:

Define different solutions available to do IMS database backup and recovery using IMS Tools.

Provide information for technical personnel to make informed decisions.

Provide technical detail on backup and recovery.

Provide detailed performance information on backup and recovery tools.

Provide information on fast replication technology.

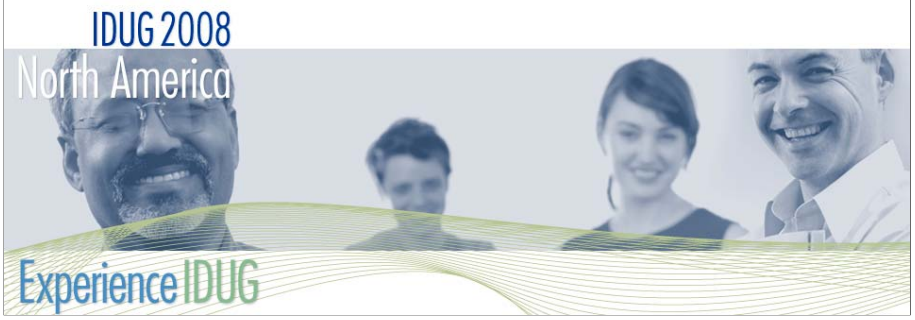

Agenda

- Basics on Backup and Recovery
- Using Change Accumulation – Why?
- Using IMS HP Image Copy Tool - Fast Replication option
- Performance figures

IMS Backup & Recovery - Just Make it Efficient!

You need the highest quality of services for your mission critical databases hosted in IMS DB. IBM does provide high quality tools to give you the best support to reach the goal of reducing the data unavailability period by also providing constant integrity. This session will put into perspective the different solutions to do database backup and recovery and will position the BEST one!

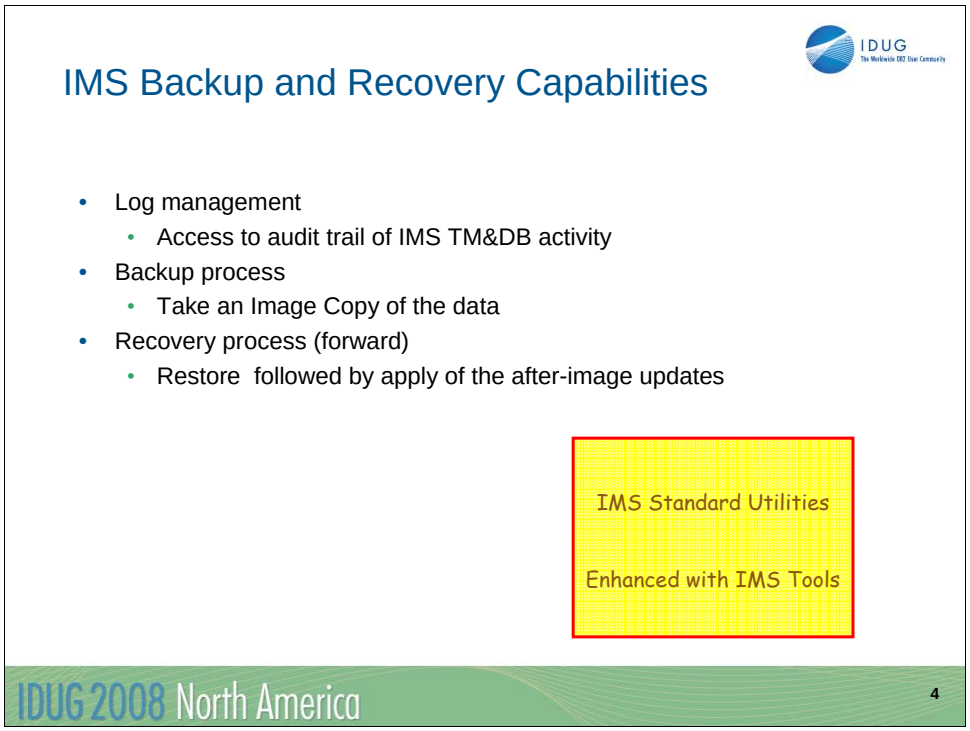

Straightforward concept – Image Copy + changes as input to recovery process brings the DB up to date.

IMS Tools Backup & Recovery Solution

- The IBM IMS Tools Backup & Recovery Solution is an effective and efficient way to ensure your IMS data is recovered quickly and accurately.
- These 5 tools combine to form an integrated and automated solution to meet your recovery needs
  - IMS High Performance Image Copy
  - IMS High Performance Change Accumulation Utility
  - IMS Database Recovery Facility
  - IMS Index Builder
  - IMS High Performance Pointer Checker & IMS HP Fast Path Utilities

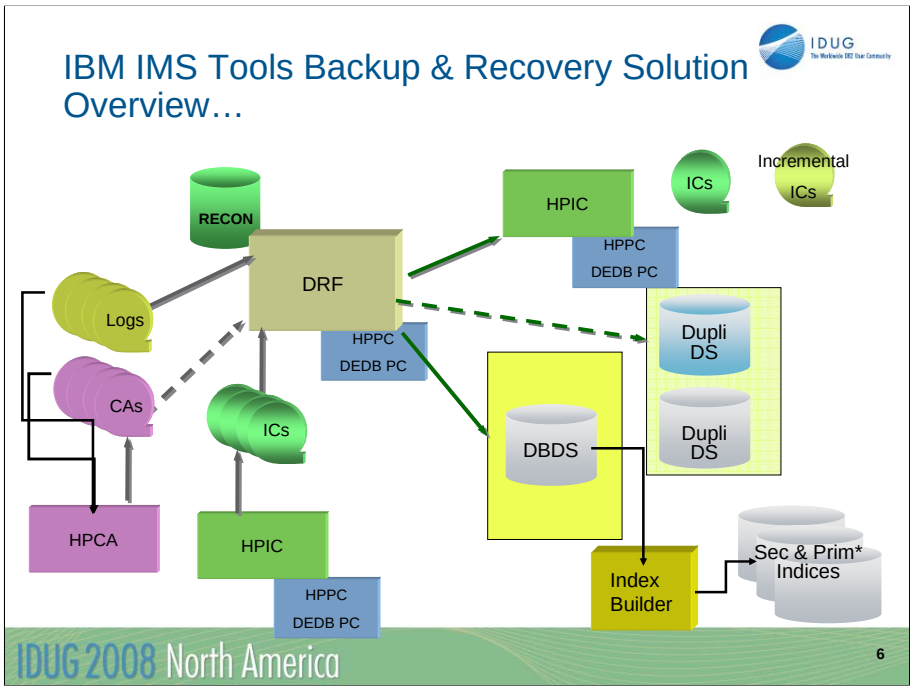

To use the Full-Function database HASH check option with image copy job, you need one of the following:

– IBM IMS High Performance Image Copy for z/OS, Version 3 Release 2 (5655-K96)

– IBM IMS High Performance Image Copy for z/OS, Version 4 Release 1 (5655-N45)

To use the HASH Check option in the IMS Parallel Reorganization job, you need IMS Parallel Reorganization for z/OS, Version 3 Release 1 (5655-M28) or later.

To use the HASH Check option in the IMS Database Recovery Facility, you need IMS Database Recovery Facility for z/OS Version 3 Release 1 (5655-N47).

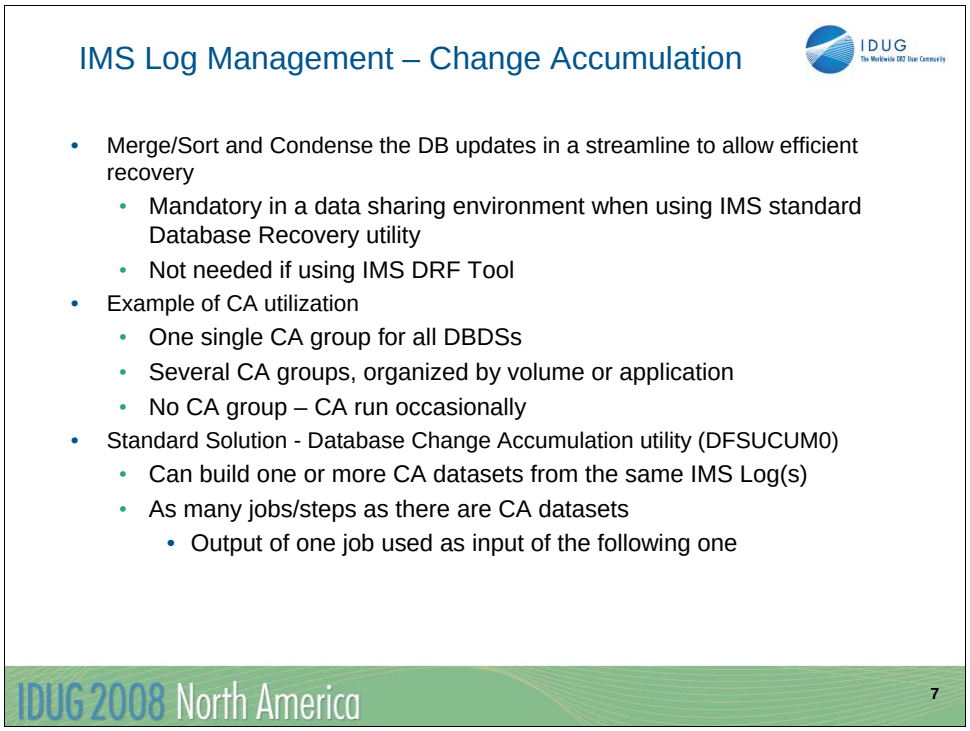

Use the **Database Change Accumulation utility (DFSUCUM0)** to streamline the recovery information you provide to the Database Recovery utility. This utility takes information from log data sets; output is in the form of a sequential data set. Utility processing involves:

1. Eliminating all non database change records
2. Specifying a purge date (or dates) to eliminate all database records before that date
3. Sorting the acceptable database change records
4. Combining all database change records that update the same database physical record

The resulting records are sequenced by data set within the database. This utility invokes the Sort/Merge program, which is an execution prerequisite. This utility also sorts as a minor field RBA or key.

This utility can be executed several times over a period of time to incorporate additional database changes and to delete changes that are no longer useful.

The following are some CA utilizations:

•Define a single change accumulation group for all DBDSs. In this case, the Database Change Accumulation utility reads the log data set only once, but the Database Recovery utility must read records in the output change accumulation data set for potentially unnecessary databases. This technique improves the performance of the Database Change Accumulation utility, but degrades performance of the database recovery.

•Define several change accumulation groups. In this case, you divide your databases into groups of several databases, often by volume or by application. If you take image copies frequently, you can keep the number of records scanned by the Database Recovery utility to a minimum.

•Define no change accumulation groups. Run the Database Change Accumulation utility, but not as a regular job.

When you need to recover a database, specify that the utility accumulate changes for the affected DBDS only.

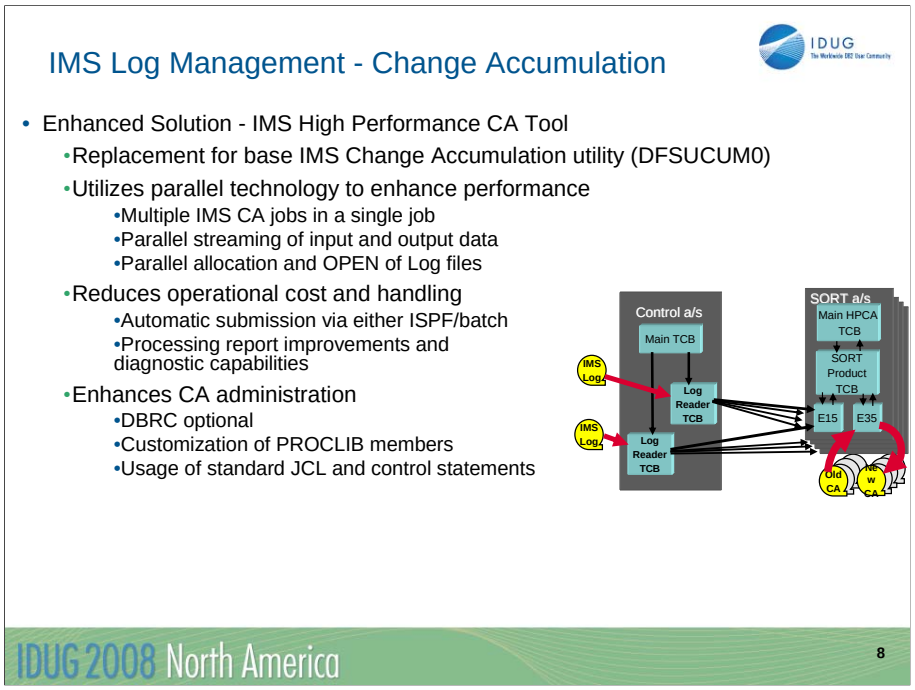

As part of the process to create control information, the ISPF/UI interfaces with DBRC to run the GENJCL procedure for CAGROUPS for which change accumulation will run. This mode of operation is considered as *ISPF mod*e.Optionally, you can use the ISPF/UI to generate only a part of the control information required by HPCA and generate the remaining information prior to running HPCA as a first step of a two-step job stream. This mode of operation is called *batch mod*e. In batch mode, the first step runs the GENJCL procedure of DBRC utility and merges the output of GENJCL with the ISPF/UI-generated control information. For this step, you must provide input to the GENJCL procedure. The second step is the processing of HPCA.

The extended batch interface uses a one-step process that can perform most functions of the ISPF/UI and the two-step batch process. By performing only one job-step, the extended batch interface can interface with DBRC, create HPCA procedures, and process multiple address space SORTS. The one-step process does not require ISPF or DBRC.

Point-In-Time CA in V1.3 - for disaster recovery scenario

Only includes committed data

Can be used as input to DRF but cannot combined with subsequent log input

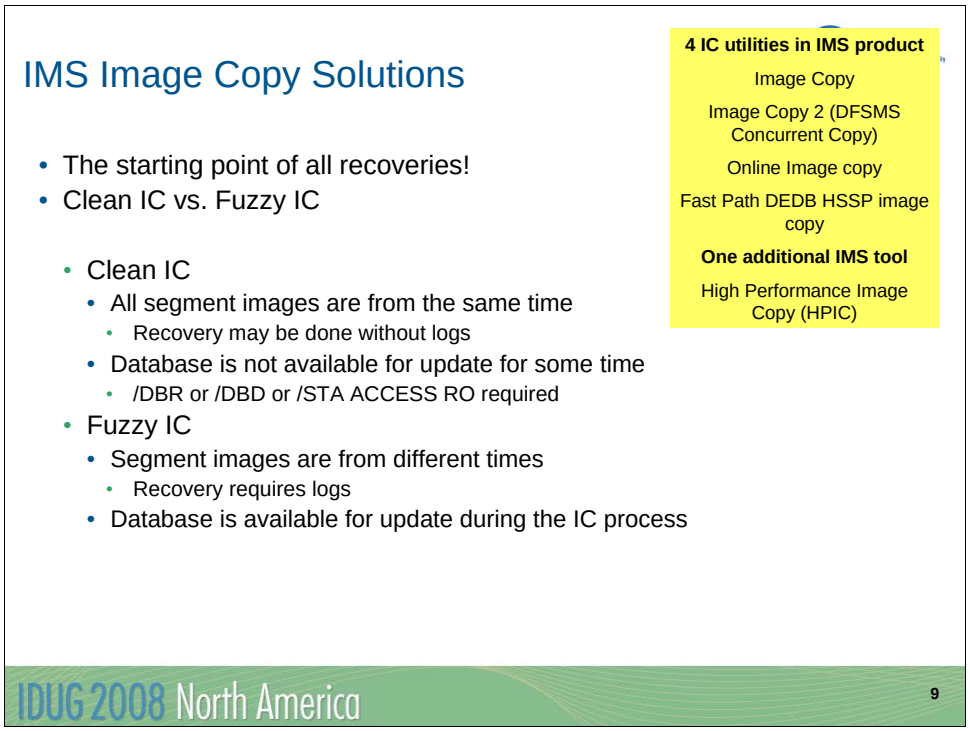

The starting point of all recoveries process! Frequency depends on expected recovery outage length.

Recommendation:

It is almost impossible to copy all data sets or areas for a database **at the same time**.

When you recover a database to a prior state, you must recover all data sets belonging to the database, as well as all logically related databases (including those related by application processing) to the same point to avoid data integrity problems.

Recommendation: Take an image copy immediately after running batch jobs that do not use IMS logging.

IMS Image Copy Solutions – From IMS Product

- Database Image Copy (DFSUDMP0)
  - Clean and fuzzy image copies
  - Fuzzy IC also called “Concurrent Image Copy”
    - Fuzzy IC not valid for VSAM KSDSs
- Online Database Image Copy (DFSUICP0)
  - Fuzzy image copies
    - Updates allowed only in the same online system
      - Limited data sharing support
  - Runs in online system – special BMP
    - Uses online buffer pools – performance implications
  - Valid for all data set types: OSAM, VSAM ESDS and KSDS
  - Does not support Fast Path databases

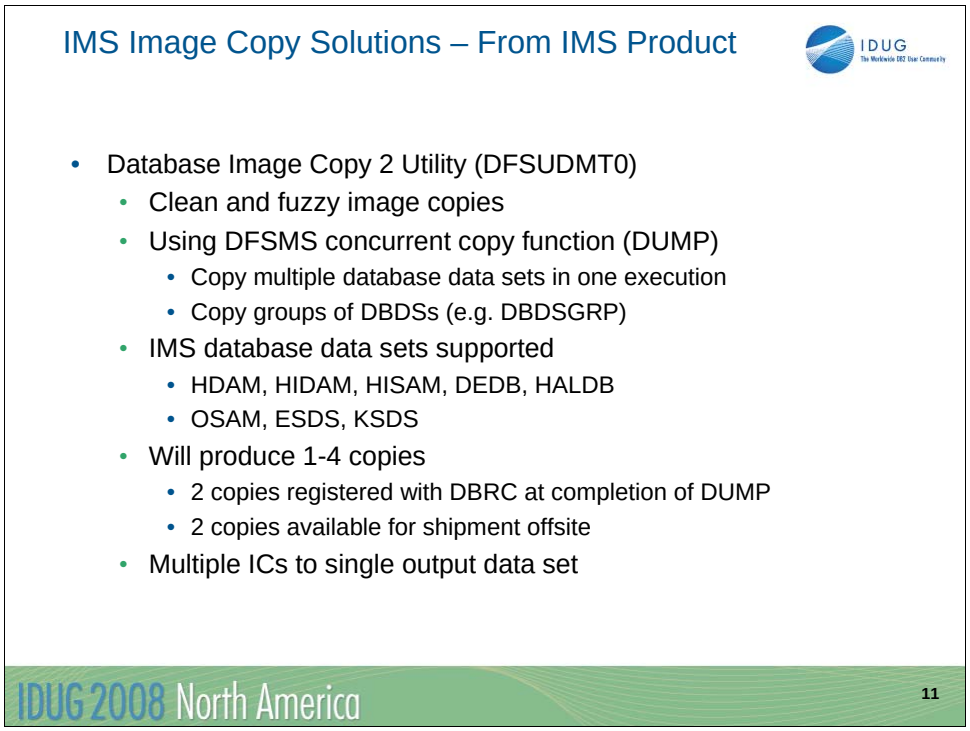

The concurrent copy function of DFSMS is a hardware and software solution that allows you to back up a database or any collection of data at a point in time and with minimum down time for the database. The database is unavailable only long enough for DFSMS to initialize a concurrent copy session for the data, which is a very small fraction of the time that the complete backup will take.

Once the concurrent copy session has been established, the copy is said to be logically complete. This concurrent copy process creates what is called the **logical copy**. After the concurrent copy initialization, updates may be resumed while DFSMS is reading the data and creating an output copy. The copy that is made will not include any of the update activity; it will be as if the backup were made instantaneously when it was requested.

A **physical copy** is made when DFSMS copies the logical copy to the output medium. This copy is registered with DBRC as either a SMSNOCIC or SMSCIC image copy. Depending upon the size of the data set, the physical copy may take time to produce. The physical copy is used by DBRC for recovery.

By using the DFSMS concurrent copy function the Database Image Copy 2 utility increases database availability. The utility can copy a database that is either stopped or active. If the database is stopped, it can be restarted after the logical copy is complete, and database updating can continue. The database is available to IMS without waiting for the physical copy to be complete. Note that nonrecoverable databases must be stopped before the utility is run.

You also have the option to wait until the physical copy is complete before releasing the database for update. This is useful in cases where an image copy is required for a specific purpose (like end-of-month processing). In this case, it is safer to wait for the

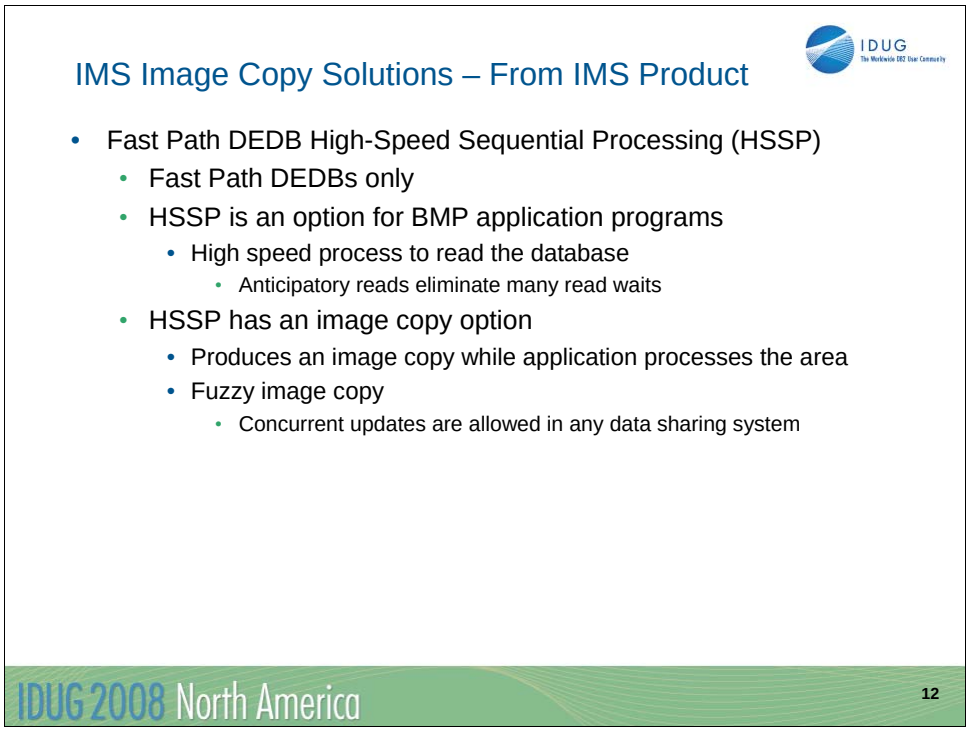

High-Speed Sequential Processing (HSSP) is a function of Fast Path that handles sequential processing of DEDBs.

Some reasons you may choose to use it are that, HSSP:

•Generally has a faster response time than regular batch processing.

•Optimizes sequential processing of DEDBs.

•Reduces program execution time.

•Typically produces less output than regular batch processing.

•Reduces DEDB updates and image copy operation times.

•Image copies can assist in database recovery.

•Locks at UOW level to ease "bottle-necking" of cross IRLM communication.

•Uses private buffer pools reducing impact on NBA/OBA buffers.

•Allows for execution in both a mixed mode environment, concurrently with other programs, and in an IRLM-using global sharing environment.

•Optimizes database maintenance by allowing the use of the image-copy option for an updated database.

Selecting the **image-copy option with HSSP** reduces the total elapsed times of DEDB updates and subsequent image-copy operations. As database administrator, you decide whether to make an image copy of a database using HSSP. If you specify image copying, HSSP creates an asynchronous copy that is similar to a concurrent image copy.

The image copy process can only be done if a database is registered with DBRC. In addition, image copy data sets must be initialized in DBRC. HSSP image copies can also be used for database recovery. However, the Database Recovery Utility must know that an HSSP image copy is supplied.

IMS Image Copy Solutions – From IMS Product …

	Image Copy (DFSUDMP0)	Image Copy 2 (DFSUDMT0)	Online IC (DFSUICP0)	DEDB HSSP IC
Part of base IMS product	Yes	Yes	Yes	Yes
Execution	Batch job	Batch job	Online (BMP)	Online (BMP)
Hardware requirement	No	Yes	No	No
Database data sets supported	Full function and DEDBs	Full function and DEDBs	Full function only	DEDBs only
Fuzzy copies	CIC option with restrictions	Yes, without restrictions	Yes, with restrictions	Yes always
- Fuzzy copies of KSDSs (HISAM, SHISAM, Index, PSINDEX)	No	Yes	Yes	
- Data sharing with fuzzy copies	Yes	Yes	No	
Minimal database unavailability for clean image copies	No	Yes	No	NA
Multiple data sets copied in one execution	Yes	Yes	Yes	no
Compressed output	No	Yes	No	No
Pointer checking	No	No	No	No
Stacked output	No	No	No	No
Dynamic allocation	No	No	No	Yes

Database Synchronization Point = Image Copy

You must consider two things when making backup copies of databases and areas:

•How frequently to make new copies

•How long to keep old (back-level) copies

There are no precise answers to these questions. Generally, the more frequently you copy, the less time recovery will take. The farther back in time your old copies go, the farther back in time you can recover; remember that program logic errors are sometimes not discovered for weeks. On the other hand, making each new copy requires work, and each old copy you save uses additional resources.

The only firm guidelines are these:

•If you create a new database, immediately make a backup copy of it.

•If a database is composed of several data sets, be sure to copy all data sets at the same time.

•If you reorganize a database, immediately make a new backup copy of it.

Additional Functionalities

- •Pointer checking: To validate the DB while the IC process is reading the DB
- •Dynamic allocation: To use the same JCL
- •Compressed Output: To reduce the size of the output dataset
- • …

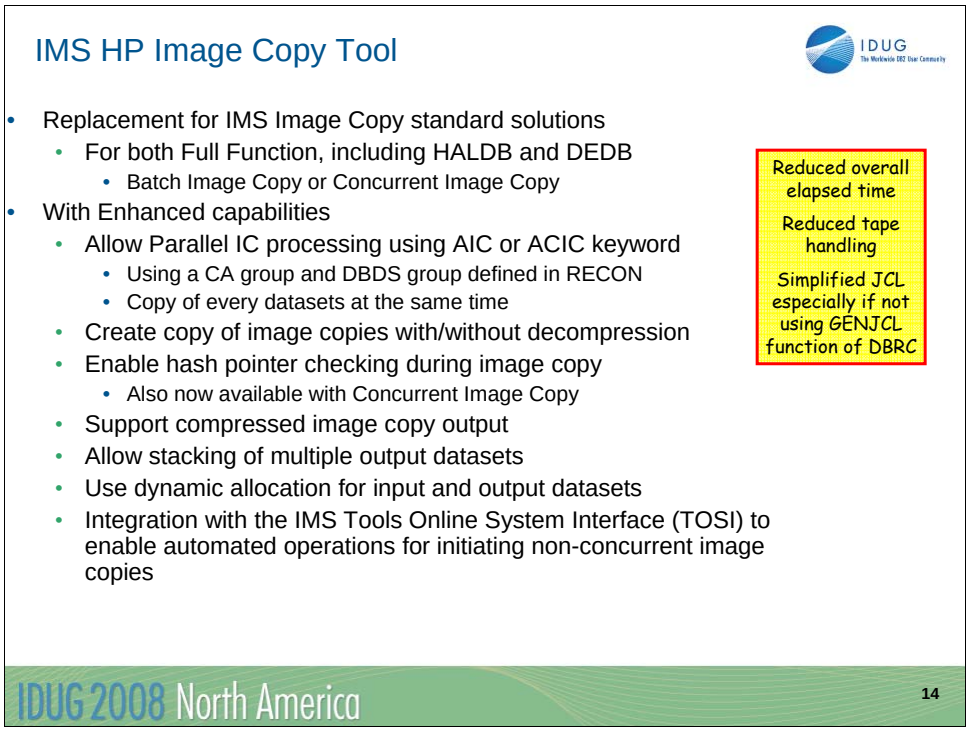

Parallel Processing: Image copy for more than one database data set in parallel within a single step. GRPLIM parameter to limit the number of parallel tasks within a group.

IC Creation: Provides capability to read an image copy and create up to seven copies of it in one execution

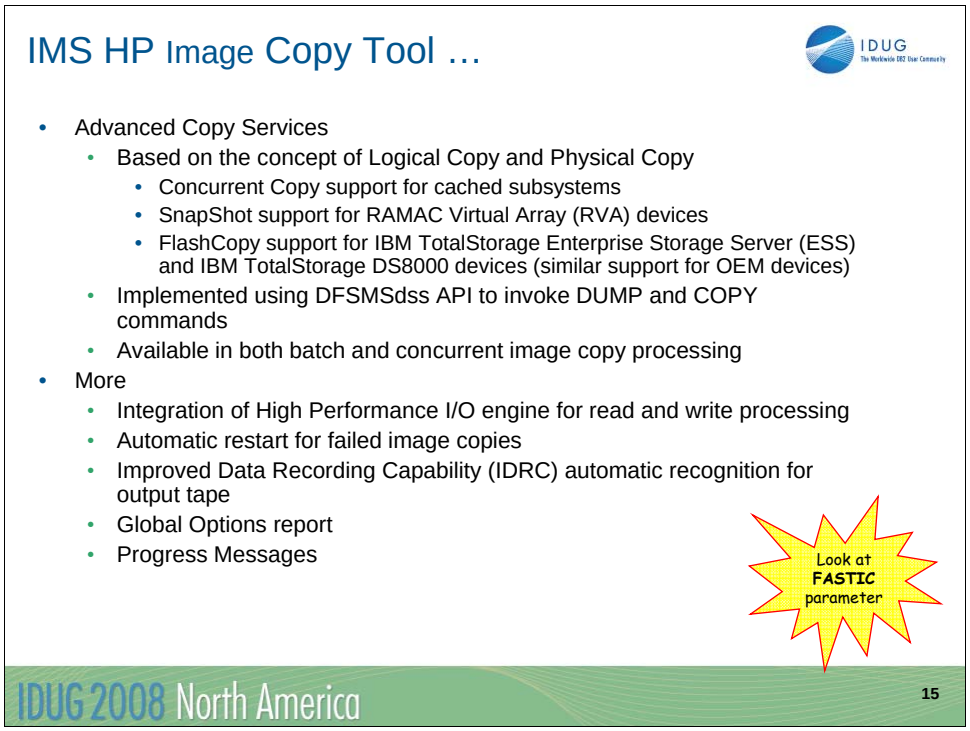

Advanced Copy services feature is initiated by using the keyword, **FASTIC=**. This keyword is available only on the GLOBAL statement and is, therefore, in effect for all image copies created by the job step. **FASTIC=** processing is valid for the batch image copy functions (IC/AIC) and concurrent image copy (CIC/ACIA). When you specify **FASTIC=**, the DFSMSdss API is used in creating the image copy. The default for **FASTIC=** is NO, which causes standard HPIC image copy processing to occur.

Concurrent Copy is NOT synonymous with Concurrent Image Copy: a software solution that is available on all cached storage subsystems, uses central storage and sidefiles which reside in the storage subsystem; Source data set must be on a device that supports Concurrent Copy; The target device need not be Concurrent Copy capable and may be on tape

SnapShot Copy is a feature on the RAMAC Virtual Array (RVA) device or similar OEM devices; provides fast data replication using a combination of hardware and software services; SnapShot copy functions only when the source and target locations reside within the same RVA subsystem; Output image copy cannot be on tape.

FlashCopy is a feature on the IBM Total Storage Enterprise Storage Server (ESS) and IBM Total Storage DS8000 devices or similar OEM devices; provides fast data replication using hardware services found within supported devices FlashCopy functions only when the source and target locations reside within the same subsystem; Output image copy cannot be on tape

RESTART= The RESTART keyword specifies whether or not to use the checkpoint restart function for automatic restart of image copy processing. This keyword operates

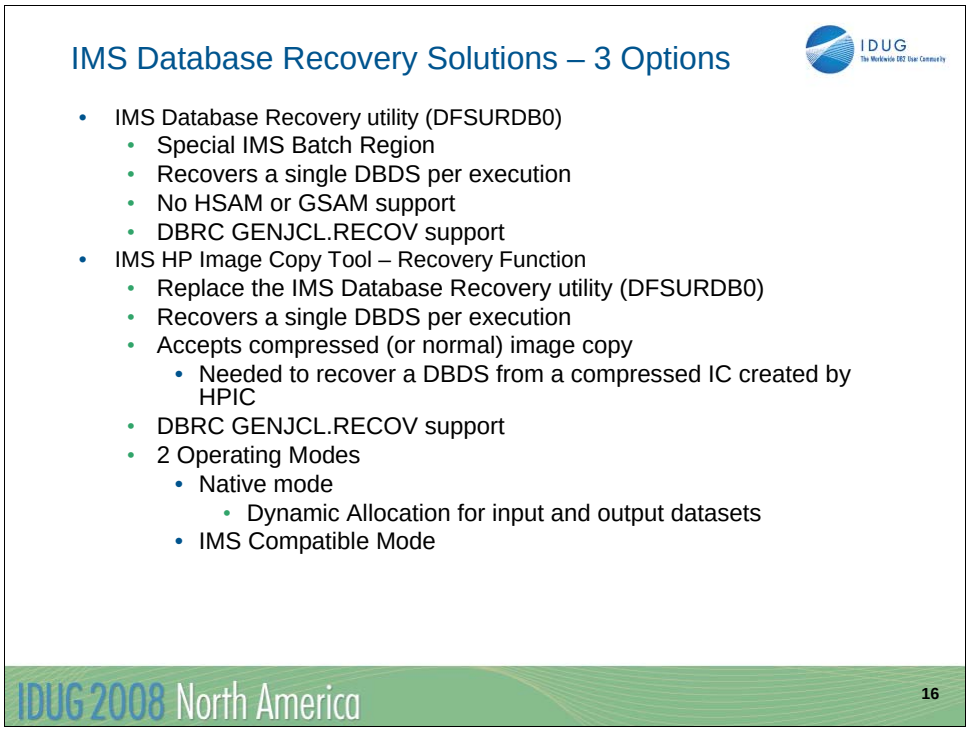

ICE Recovery Operating modes

Native Mode

•Dynamic allocation supported

•ICEIN dataset with adequate control statements

•EXEC PGM=FABJMAIN

IMS Compatible Mode

•Dynamic allocation NOT supported

•EXEC PGM=DFSRRC00,PARM='UDR,DFSURDB0,dbdname'

•JCL is almost same as for standard recovery utility

•Can use GENJCL.RECOV in standard manner

•ICE Recovery Utility automatically detects if input image copy needs de-compressing

•*information in IC header record*

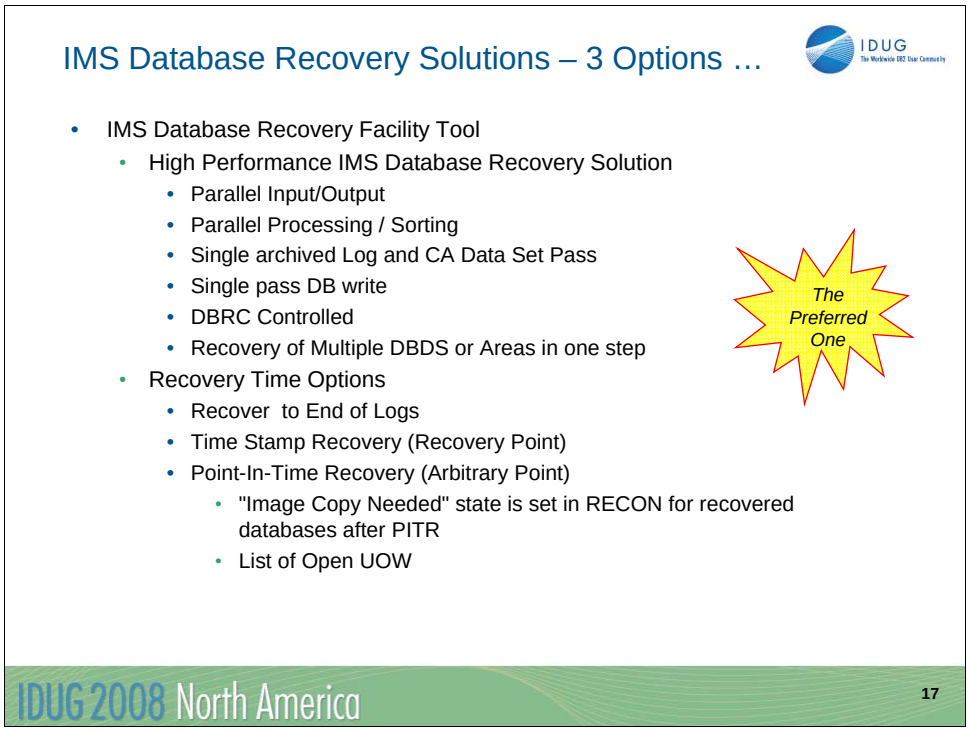

Database data sets and Fast Path areas being recovered that have image copies in a **DFSMSDSS SAMEDS format**, are restored at the beginning of recovery, before log and change accumulation data sets are read. The database data sets in the same SAMEDS are not necessarily recovered in the same subordinate address space.

All database data sets to be recovered, as well as the log, change accumulation, and image copy data sets used for recovery, must be registered with the **IMS Database Recovery Facility (DBRC).**

Perform time stamp forward recovery with time stamps that are not restricted to allocation boundaries. The ability to recover to any point-in-time allows you to recover IMS databases and areas to the same point in time as DB2 databases using time stamp recovery.

Supported IMS DB types:

Full Function: HDAM, HIDAM, HISAM, SHISAM, and INDEX

HALDB: PHIDAM, PHDAM, and PSINDEX

DEDB: ADS and MADS

Not supported IMS database types:

•GSAM, SHSAM, HSAM, and MSDB

•HALDB PHIDAM primary indexes*

•Indirect List Data Sets (ILDSs)*

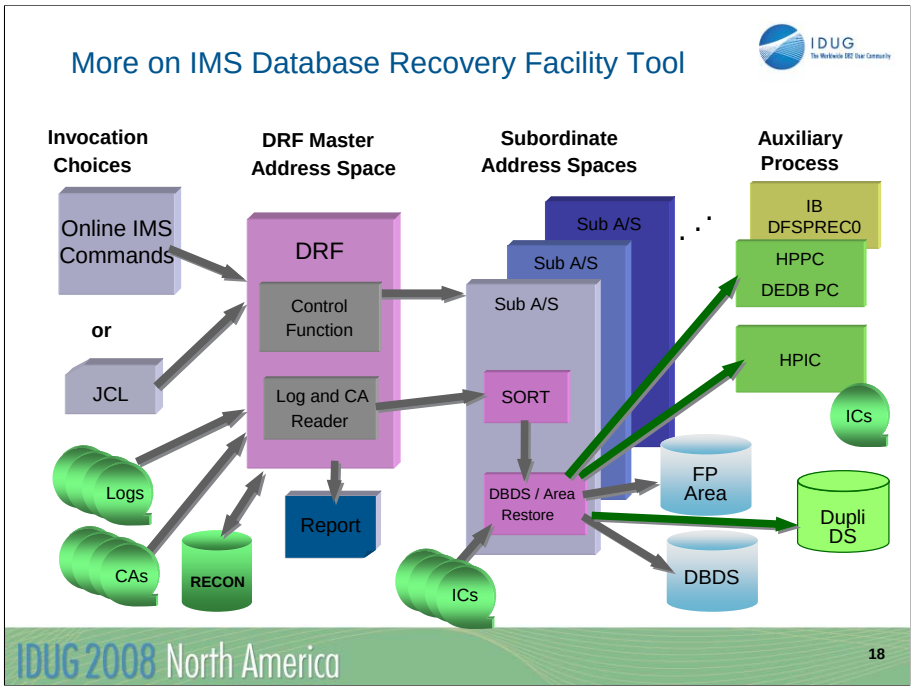

IMS with DRF only analyzes the syntax of the input /RECOVER and /DISPLAY RECOVERY commands.

The commands and accompanying parameters are translated into a control block, DFSRWAI, which is sent from the IMS command processor to the IMS Database Recovery Manager, DBRM, ITASK. DBRM checks to see if a recovery environment has been started. If not, it attempts to load the DRF communication modules.

If this is unsuccessful, it attempts to load the ORS communication modules. Assuming the DRF modules were loaded, it creates the DRF address space via ASCRE. It sends the RWAI to the DRF address space. The DBRM waits to process the results of recovery or another command. When the results of recovery are received by DBRM, the appropriate messages are issued. Everything else happens in the DRF address spaces.

The IMS Database Recovery Facility subordinate address space is responsible for the following tasks:

v sorting log data

v restoring image copy

v sorting and applying change accumulation data

v applying log data

v calling related utilities such as HPIC, HPPC, IB, DFSPREC0, and DEDB PC (with V3)

In the subordinate address space, the log data is received from the primary address space and passed to the appropriate sort tool for use in recovery processing. The

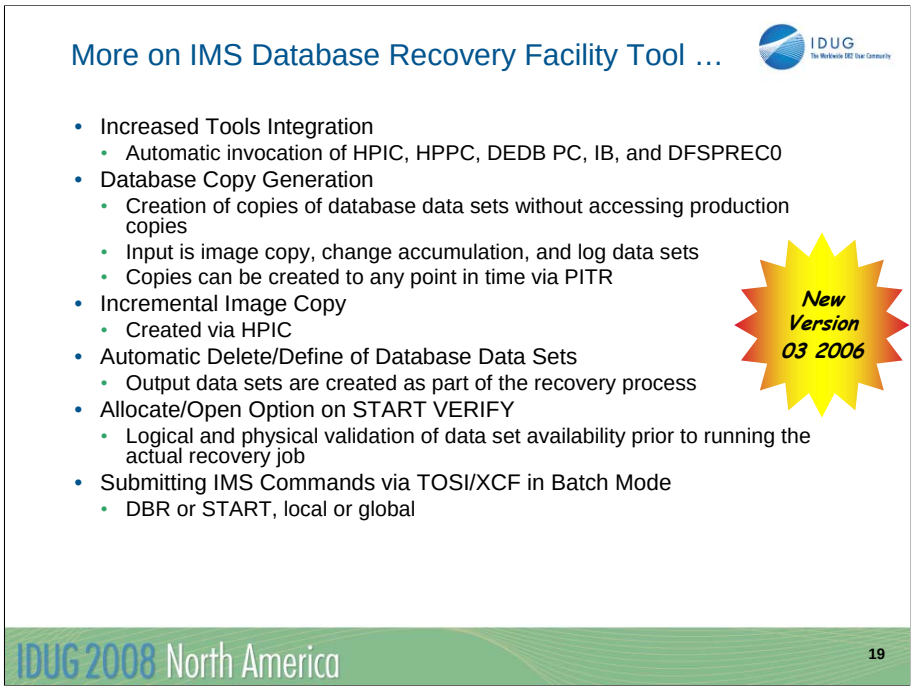

Beginning with IMS Database Recovery Facility V3R1, you can perform the following tasks:

•Integrate recovery tasks into a single job step, which can include: – creating image copies – recovery validation through pointer checker – rebuilding your index data sets – rebuilding your HALDB index list data sets

•Create image copy data sets during the recovery process to eliminate the need to follow recovery with a separate image copy invocation.

•Create index data sets during the recovery process, including primary, secondary, and HALDB index list data sets, in order to eliminate the need to follow recovery with a separate index build invocation.

•Create offline, cloned copies of the database which can be used for tasks such as application program testing, audit, and database recovery.

•Validate data during the recovery process, which eliminates the need to follow a recovery with a separate validation procedure invocation.

•Automatically delete and create database data sets during the recovery process, which eliminates the need to precede recovery by separately performing these tasks.

•Generate new image copy data sets by applying database updates from a log data set or change accumulation data set to existing image copy data.

•Verify data sets by specifying options such as allocate and/or open verification in order to verify availability and usability of recovery input data sets.

Beginning with IMS Database Recovery Facility V3R1, you do not need to issue the IMS **DBR DB** and **DBR AREA** commands before recovery and the IMS **STA DB** and **STA AREA** after recovery.

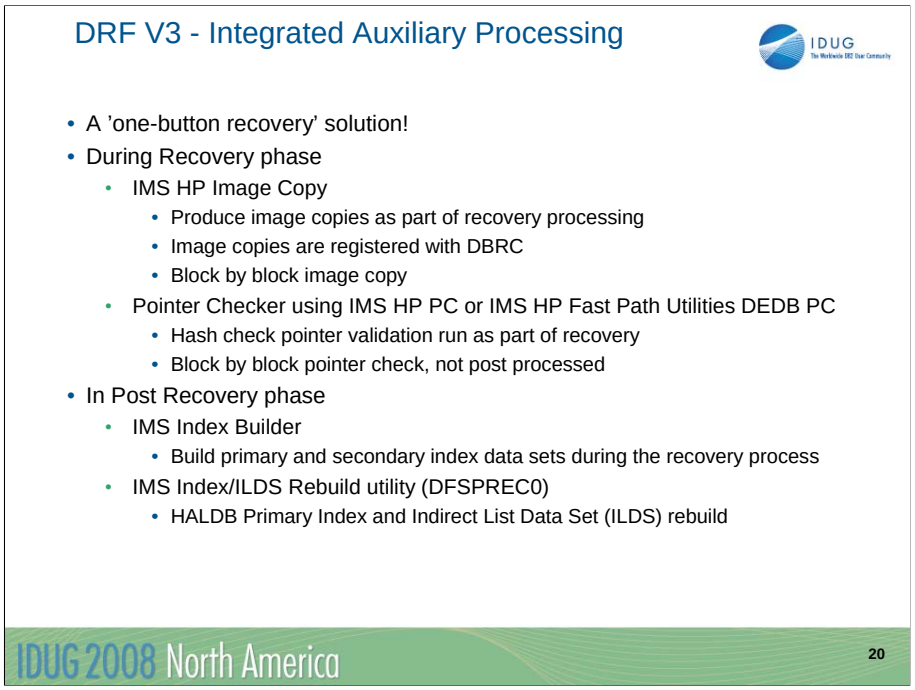

IMS Database Recovery Facility has a 'one-button recovery' solution that allows you to define recovery and post-recovery tasks and to submit one job to accomplish these tasks.

Integrated auxiliary processing allows you to use the DB Control Suite interface or batch JCL and control statements to tell IMS Database Recovery Facility that the following four processes (or a combination of them) are to be performed as part of the database recovery, rather than being performing in separate job steps after the recovery is complete:

v Image Copy generation using High Performance Image Copy (HPIC)

v Hash Pointer Checking for a Full Function DBDS using High Performance Pointer Checker (HPPC) or for a Fast Path DEDB using IMS Fast Path Basic Tools- DEDB Pointer Checker

v HALDB Primary Index rebuild (PHIDAM indexes are also partitioned) and HALDB Indirect List data set (ILDS) rebuild using DFSPREC0, a part of the IMS base product, after a time stamp recovery

v Non-HALDB primary index rebuild, and HALDB or non-HALDB secondary index rebuild using Index Builder (IB)

For integrated mode execution, there are additions to IMS Database Recovery Facility master and subordinate address space JCL, FRXDRF*xx* PROCLIB member, and control statements to communicate the recovery policy to the auxiliary utilities. For example, an IC(*options*) keyword is added to the IMS Database Recovery Facility **ADD** control statement indicating that an image copy of the database is to be performed as the database is being recovered. IMS Database Recovery Facility is responsible for communicating to HPIC all the data sets and options necessary to create image copies. The maximum completion code from the invocation of any utility or IMS Database Recovery Facility is returned as the final completion code for the recovery job.

Agenda

- Basics on Backup and Recovery
- Using Change Accumulation – Why?
- Using IMS HP Image Copy Tool - Fast Replication option
- Performance figures

Change Accumulation Process – Why ?

- Designed to speed-up Recovery
  - Accumulates changes to DB
    - Merging changes from online and/or batch subsystems
    - Required in a data sharing environment unless using DRF Tool
  - Minimizes log processing during Recovery
- Recovery Input
  - Image Copy
  - Change Accumulation
  - Logs created subsequent to most recent Change Accumulation
- Continuous Process
  - Post-process most log data for it to be effective
  - Resource intensive insurance policy
  - Operational complexities

IMS HP Change Accumulation Tool – V1R4 Benefits

- Simplified setup and configuration using a single address space
- Improved elapsed time performance
- Reduction or elimination of cross-memory data transfers
  - Single address space approach
- In-memory spill files using 64-bit addressing
  - Always limit the 64-bit storage usage by using MEMLIMIT= parameter
  - Need to monitor local paging

HPCA V1.3 vs Standard CA

Hardware	Software	Database
• CPU -- 4-cp 2064 • DASD -- ESS-F20, DS8000, RVA • 32 SLDS with 200 CYLs each on various DASD	• z/OS 1.6 • IMS V8 • HPCA 1.3	• 86 DEDB areas • 8 CAGRPs • Area size: 2,200 CYLs

- STD DFSUCUM0
  - 1 CA job a time, 8 CA serial runs
    - EXCPs: 1.538M
- HPCA 1.3 (Parallel Sorts=8)
  - 1 CA job on system
    - EXCPs: 265,398
    - 5.8 times less DASD IO

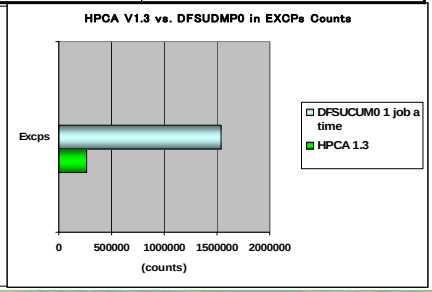

HPCA V1.3 vs Standard CA

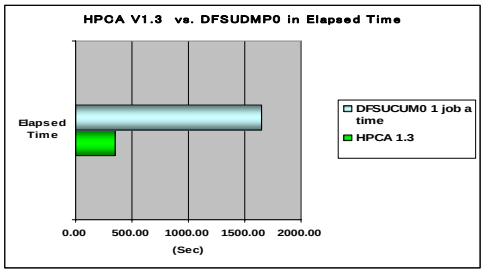

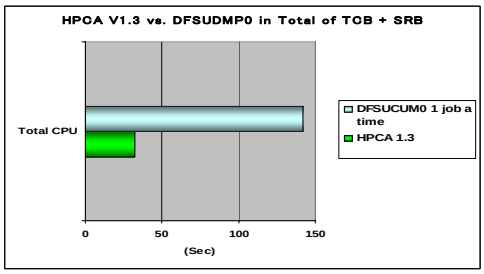

- Elapsed Time
  - STD DFSUCUM0
    - Elapsed Time (sec): 1,648
  - HPCA 1.3 (P.Sorts=8)
    - Elapsed Time (sec): 352
  - 4.7 times faster
- CPU TCB & SRB Time
  - STD DFSUCUM0
    - TCB: 16.45x8
    - SRB: 1.26x8
    - Total (TCB+SRB): 141.68
  - HPCA 1.3 (P.Sorts=8)
    - TCB: 29.92
    - SRB: 2.03
    - Total (TCB+SRB): 31.95
  - 4.4 times less CPU time

Agenda

- Basics on Backup and Recovery
- Using Change Accumulation – Why?
- Using IMS HP Image Copy Tool - Fast Replication option
- Performance figures

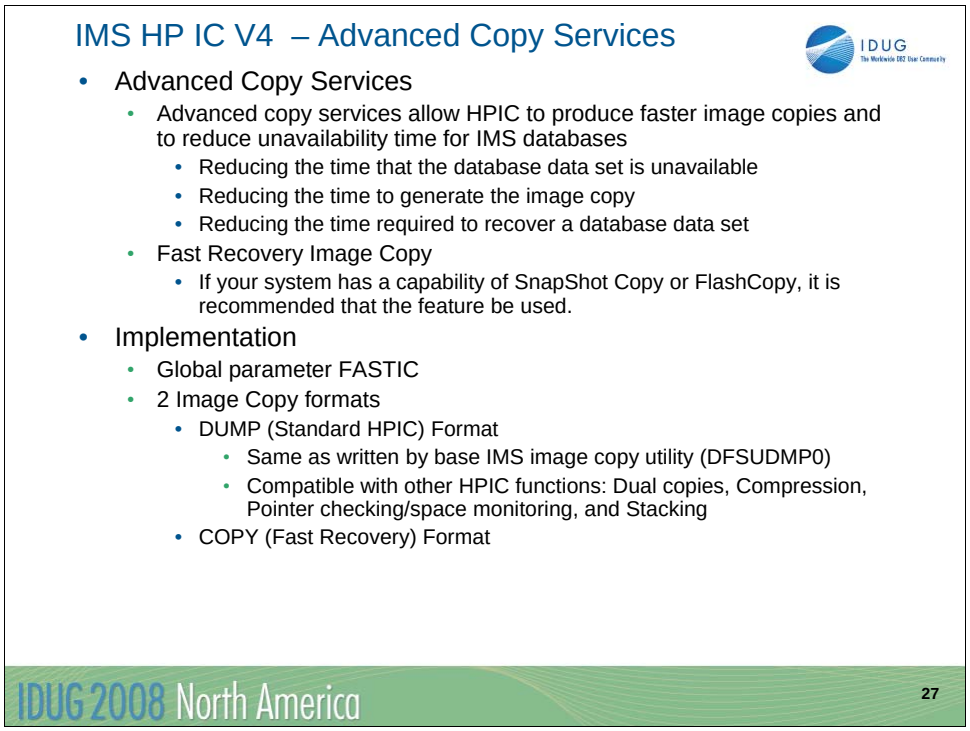

High availability backup solution

Supports both concurrent and non-concurrent copy

Almost no disruption for concurrent copy

Minimal disruption for non-concurrent copy

Allows for more frequent Image Copies

Provides a faster and more efficient recovery process

FASTIC=(REQ,)

Use of advanced copy services is required

Image Copy fails if advanced copy services are unavailable

Useful when database unavailable time must be limited

FASTIC=(PREF,)

Use of advanced copy services is preferred

Image Copy performed even if advanced copy services are unavailable

FASTIC=(PREF,) is the default

Advanced Copy Services - IMS image copy format

- Create IMS image copy format data set by less CPU time and less EXCPs.
  - Using DFSMSdss DUMP command

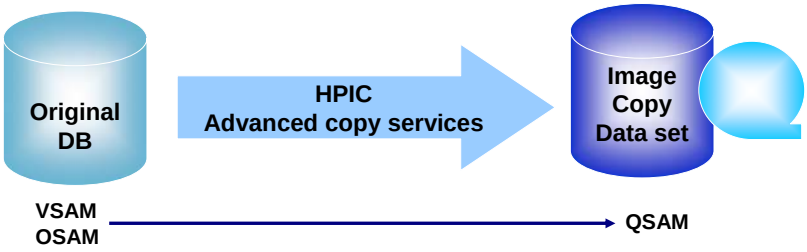

- DB unavailable time is shorter than HPIC V3.2 Image Copy function
- Enable HPIC various functions such as
  - Hash pointer checking
  - Compressed IC
  - Secondary IC

Advanced Copy Services – “Fast Recovery” image copy

- Fast Recovery image copy is exact replica of the source database
  - Using DFSMSdss COPY command

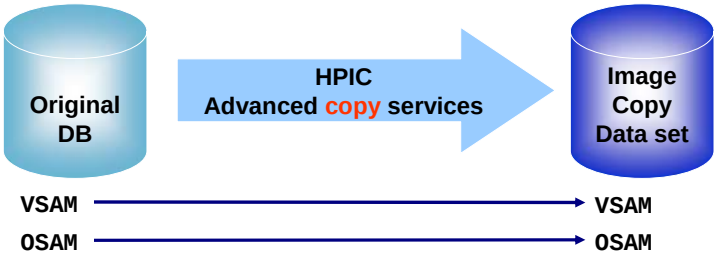

- Format of Output Image Copy is as same as IMS V10 New Copy ‘ fast replication function’ of IC2
  - IMS V10 : Registered as SMSOFFLC/SMSONLC
  - IMS V8 and 9: Registered as SMSNOCIC/SMSCIC
  - Note: Fast Recovery image copy of IMS V8 and 9 should use HPIC Recovery function or Database Recovery Facility (DRF) .

Differences of each Image Copy process

	HPIC V4			*IMS Utility*		
	FASTIC=				*IC2*	
	COPY	DUMP	NO	*Batch Image Copy*	*IC2*	*FlashCopy (IMS V10)*
Copy Method	FlashCopy SnapShot Concurrent (*)	Concurrent	OS I/O	*OS I/O*	*Concurrent*	*FlashCopy*
Performance	1	2	3	*4*	*2*	*1*
Tape output	N/A	OK	OK	*OK*	*OK*	*N/A*
Hash Check	N/A	OK	OK	*N/A*	*N/A*	*N/A*

* If FlashCopy or SnapShot Copy is available, these method are used preferentially .

IMS HP IC V4 – Compression Routines

- Compression consumes more CPU, tradeoff is between media and CPU cycles
- COMPRTN keyword specifies routine to compress FABJCMPn
  - "1" compresses repeating characters
  - "2" compresses free space
  - "3" compresses both repeating characters and free space
  - "4" is used for DEDB with SDEP segments only
- Recommendation
  - For disk - COMPRTN=FABJCMP3
  - For IDRC tape devices - none

HPIC V4 vs DFSUDMP0/DFSUDMT0 - PHIDAM/VSAM (no IC compression)

Hardware	Software	Database
• CPU -- 2-cp 2084 • DASD -- ESS-F20	• z/OS 1.6 • IMS V8 • HP IC V4.1 and V3.2	• PHIDAM/VSAM • 3,893,360 segments • 561,440 roots • Size: 2,620 CYLs

- Comparison
  - STD DFSUDMP0
    - Elapsed Time (sec): 312
    - DB Unavailable (sec): 312
  - HPIO = YES vs STD DFSUDMP0
    - Elapsed Time (sec): 123
      - 2.5 times faster
    - DB Unavailable (sec): 123
      - 2.5 times less CPU
  - FASTIO = COPY vs STD DFSUDMP0
    - Elapsed Time (sec): 43
      - 7.3 times faster
    - DB Unavailable (sec): 31
      - 10.1 times less CPU time

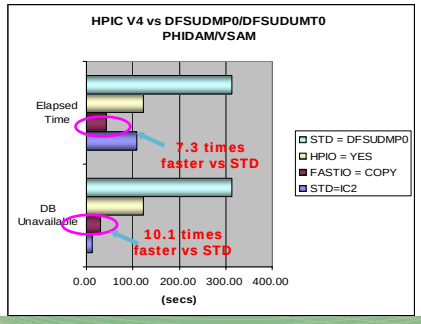

Agenda

- Basics on Backup and Recovery
- Using Change Accumulation – Why?
- Using IMS HP Image Copy Tool - Fast Replication option
- Performance figures

Study Environment & Methodology

- Hardware
  - 2-LPAR z990 (4-cp each) plus ICFs(1-cp each)
  - DS8000 for 24-AREA DEDBs and 4 SYSDAs; ESS for others
- Software
  - zOS 1.7
  - IMS V9
  - DRF V3 (GA)
  - HPIC 4.1(GA)
  - HPCA 1.4 (pre-GA)
- FP workload
  - 24-area TPCA (Banking Industry) application, Debit/Credit with SDEP insertion; running @ 1,000+ tx/sec (2-way IMSplex)

- **Methodology**
  - 24 2200-cyl DEDB areas with base crisp HPIC4 ICs
    - (**COMPRTN=FABJCOM4 to conserve DASD space)**
  - After 24+ hours of online stressing and with 64 SLDS/RLDS generated:
    - **Made sure IMS chkpt were taken on 3 consecutive OLDS prior to run FlashCopy HPICs for all 24 areas and observe FC's impact to online tran rates**
    - **Then /CHE FREEZE 2 IMS.**
  - Run DRF Recover-CURRENT with FC ICs and last 7 RLDS for IMZ1/IMZ2.
  - Deleted all FC ICs from the RECONs and ran HPCA 1.4 (max logs equal 64).
  - Ran DRF Recover-CURRENT with base CA Master and 7 remaining logs.

Study Observations

- FlashCopy IC for 24 areas
  - Average FlashCopy IC elapsed time was .7 second
  - The impact to online tran rates running at a combined 1,100 tx/sec was less than 3% decrease.
- DRF needs to go back only 2 prior chkpts when using CICs. (Only 1 prior chkpt for FF.) Thus eliminating the handling voluminous RLDS/SLDS.

FlashCopy/Change Accumulation/Recovery Performance

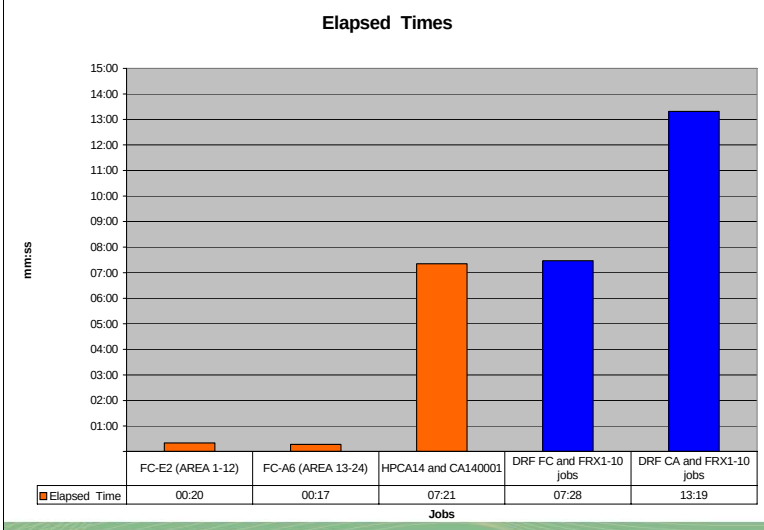

FlashCopy/Change Accumulation/Recovery Performance

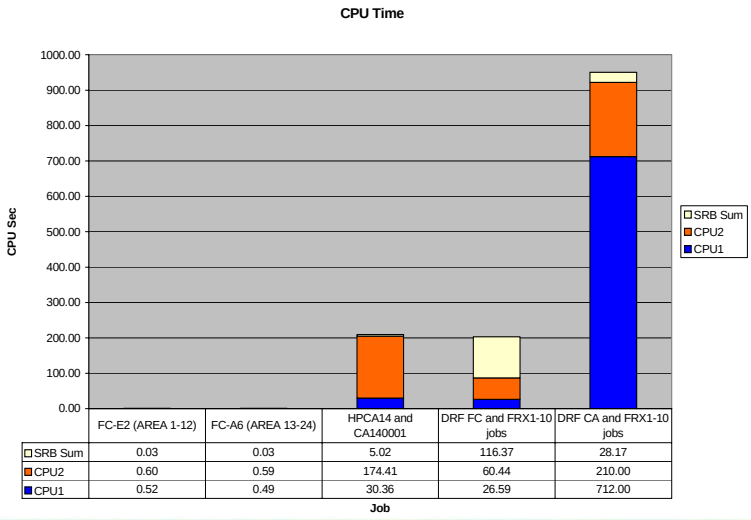

	FC-E2 (AREA 1-12)	FC-A6 (AREA 13-24)	HPCA14 and CA140001	DRF FC and FRX1-10 jobs	DRF CA and FRX1-10 jobs
SRB Sum	0.03	0.03	5.02	116.37	28.17
CPU2	0.60	0.59	174.41	60.44	210.00
CPU1	0.52	0.49	30.36	26.59	712.00

FlashCopy/Change Accumulation/Recovery Performance

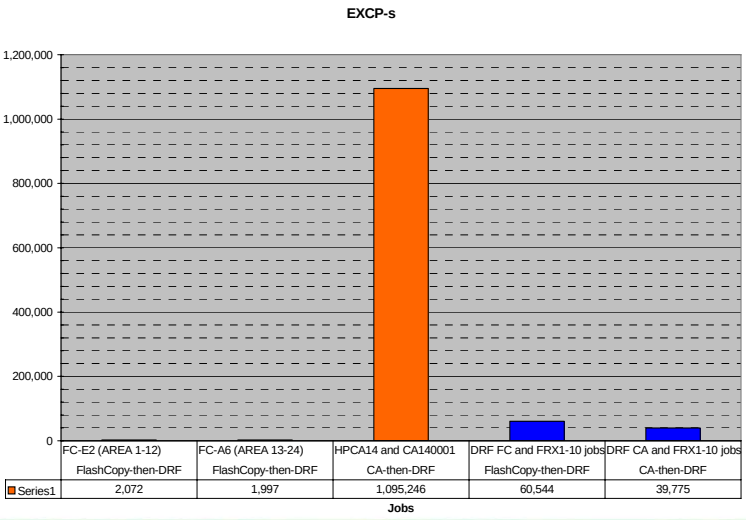

FlashCopy vs Change Accumulation Performance

FlashCopy (HPIC)	Elapsed (seconds)	CPU (seconds)	SRB (seconds)	Total EXCPs
AREA 1-12	20	.52+ .60 IEESYSAS	.02+ .01 IEESYSAS	2,072
AREA13-24	17	.49+ .59 IEESYSAS	.02+ .01 IEESYSAS	1,997
HPCA	**Elapsed Tm (mm:ss)**	**CPU (mm:ss)**	**SRB (seconds)**	**Total EXCPs**
HPCA14 and CA140001	07:21	00:30.36+ 02:54.41	2.81+ 2.21	1,095,246

Recovery Performance

Type	Elapsed Tm (mm:ss)	CPU (mm:ss)	SRB (seconds)	Total EXCPs
DRF with Flashcopy	07:28	00:26.59+ 00:60.44 (10 FRX addrspc)	.44+ 115.93 (10 FRX addrspc)	60,544
DRF with Change Accumulation	13:19	11:52+ 03:30 (FRX1-10)	.30+ 27.87 (FRX1-10)	39,775

Summary

- Study reveals compelling reasons to re-consider recovery strategy
  - Major gains in efficiency and speed using Fast Replication
  - Major gains in efficiency and speed during Recovery
  - Reduced operational complexity
- Consider impact of workload growth on future Change-Accum runs
  - Resources required
  - Ability to keep-up
- Doesn't have to be all-or-nothing
  - Start with high update activity databases
  - Phased migration
- Consider using Fast Replication even if not using Change-Accum

Session K13

IMS Backup & Recovery - Just Make it Efficient!

Bob Magid

IBM

magid@us.ibm.com