IBM Software Group

Gros plan sur…

florence_dubois@fr.ibm.com

DB2 Information Management Software

© 2004 IBM Corporation

IBM Software Group

Storage management before and after DB2 for z/OS V8

DB2 Information Management Software

© 2004 IBM Corporation

Before DB2 V8 – What is the problem?

- Background
  - Each address space has an addressing range of 2GB based on 31 bit addressing
    Maximum of 16MB available "below the line"
    Maximum of 2032MB available "above the line"
  - Practical maximum available to DB2 below and above the line much less

- GETMAIN processing by DB2
  - Requests may be conditional or unconditional
  - "Short on Storage" condition can occur for both
  - DB2 recovery routines may be able to clean up
  - Individual DB2 threads may abend with 04E/RC=00E200xx when
    Insufficient storage available e.g., 00E20003 & 00E20016
  - Eventually DB2 subsystem may abend with abend S878 or S80A when critical task and no toleration of error

© 2004 IBM Corporation

What are the drivers?

- Workload growth
- DB2 subsystems recycled less frequently
- Growth in thread related storage over successive DB2 releases
- Slow downs and CTHREAD / MAXDBAT throttle set too high
- New workload types: ERP & CRM e.g., use of local DSC
- Long running persistent threads
- Dynamic SQL
- Widespread use of Compression plus wide partitioning
- Size of ECSA to support WebSphere, IRLM PC=NO, etc
- Over allocation of ECSA and other extended common areas
- Query parallelism with high degree
- Heavy concurrent SQL prepare activity without DSC turned on
- Over allocation of bufferpools, EDM Pool, etc

© 2004 IBM Corporation

What is allocated in DBM1?

- Storage allocated for both subsystem and threads
- Majority of storage allocated « above the line »

- DB2 Code
- Virtual (Buffer) Pools
- Controls blocks for VP and HP buffers
- Lookaside buffers for data space BP
- EDM Pool including Global Dynamic Statement Cache
- RID Pool
- Compression dictionaries for open datasets
- RDS OP Pool
- Agent Local Storage
- Stack Storage
- Internal Trace Pool used by Buffer Mgr. and Data Mgr.
- Local Dynamic Statement Cache
- Storage required by various service tasks
- Fast log apply buffers
- Real Time Statistics blocks

© 2004 IBM Corporation

How much VSTOR is available to DBM1?

- RMF Virtual Storage Private Area Report
  - Interval data collected in SMF Type 78-2
  - RMF Monitor I session option: *VSTOR(D,xxxxDBM1)*
  - RMF Post Processor option: *REPORTS(VSTOR(D,xxxxDBM1))*
- Calculate amount of storage available above the line

  REGION ASSIGNED – (MAX LSQA/SWA/229/230 PAGES ALLOCATED + MAX USER REGION PAGES ALLOCATED)
- How much is enough?
  - Greater than 500MB spare is AOK (**GREEN**)
  - Between 200-500MB spare is boundary condition (**AMBER**)
  - Less than 200MB action is required (**RED**)

© 2004 IBM Corporation

How much VSTOR is used in DBM1?

- DB2PM Statistics Report|Trace Layout Long
  - ZPARM SMFSTAT=(....,6) to generate IFCID 225
  - ZPARM STATIME=5 (mins)
  - ZPARM SYNCVAL=0
- Develop and set virtual storage budget
  - Determine how much non-thread related storage is required
  - Develop how much storage is used per active thread
  - Plan on keeping at least 200 MB spare for tuning, growth, recovery, etc.
  - Determine how many active threads can be supported
- Set CTHREAD and MAXDBAT defensively for robustness to protect system

© 2004 IBM Corporation

64-bit REAL storage support

- Prerequisites: zSeries, OS/390 R10 or z/OS, ESAME 64-bit LPAR, DB2 V6+
- Data spaces provided a good short term solution by exploiting 64-bit REAL
  - Buffer pools and statement caching in data spaces
    Frees up space for other work in the DBM1 address space
    Performance penalty when not 100% backed by real storage
- Advantages of data spaces over hiperpools
  - Read and write cache with direct I/O to data space
  - Byte addressability
  - Very large buffer pool sizes
    32GB for 4K page size
    256GB for 32K page size
- Excellent performance experienced with z900 and large processor storage
  - Performance dependent upon being in 64-bit REAL mode

© 2004 IBM Corporation

DB2 V8 – 64-bit VIRTUAL storage support

- DB2 Version 8 is 64-bit exclusive
  - Buffer pools always allocated above 2GB bar
  - Eliminates need for hiperpools and data space pools
  - "Data access" modules enhanced to access 64-bit addressable buffers "in place"
  - NO internal data movement as per data space buffer pools today
- Sizing and placement
  - Buffer pool max size is 1TB

    The actual maximum = the REAL storage available
  - Total buffer pool max size is 1TB
  - Page manipulation blocks (PMBs) moved above the bar
  - Castout buffers (data sharing) above the bar

© 2004 IBM Corporation

Limit changes

- Max number of read, write, castout engines increased to 600
- z/OS V1.2 provides a new MEMLIMIT JCL keyword
  - Controls how much VSTOR above 2GB bar is available in address space
  - DB2 sets MEMLIMIT value to 4TB (minimum) ensuring sufficient memory for operation
- VPSIZE default values:

BP0	20000	(Minimum 2000 (from 56))
BP8K0	1000	
BP16K0	500	
BP32K	250	

© 2004 IBM Corporation

What else is moving above the 2-GB bar?

- RID pool
- Sort pool
- Compression dictionaries
- EDM pool - DBDs and OBDs

© 2004 IBM Corporation

RID pool

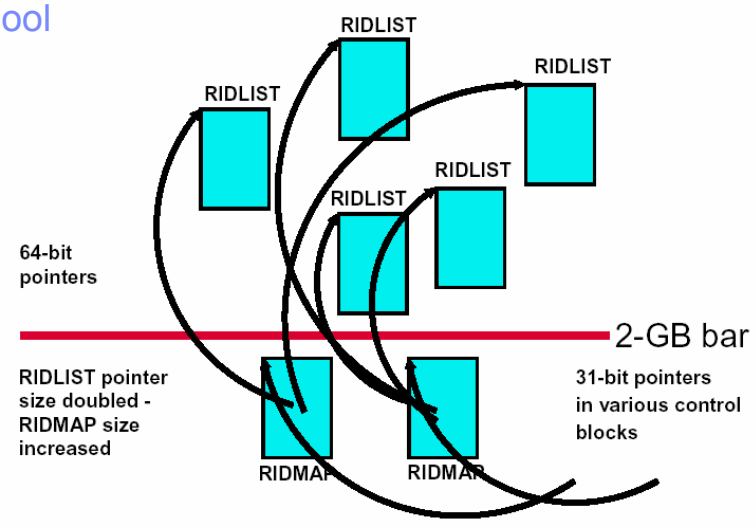

© 2004 IBM Corporation

Compression dictionaries

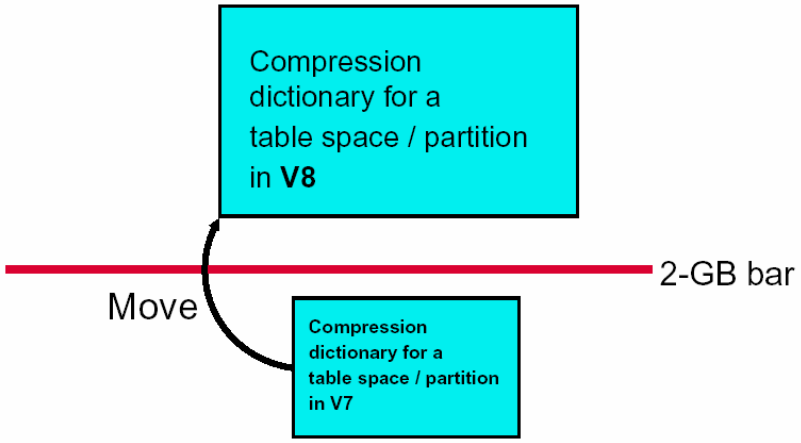

© 2004 IBM Corporation

Sort pool

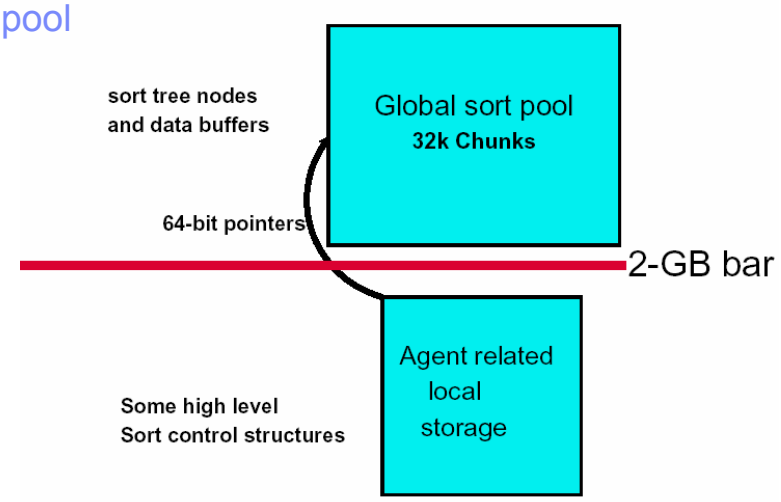

© 2004 IBM Corporation

EDM pool

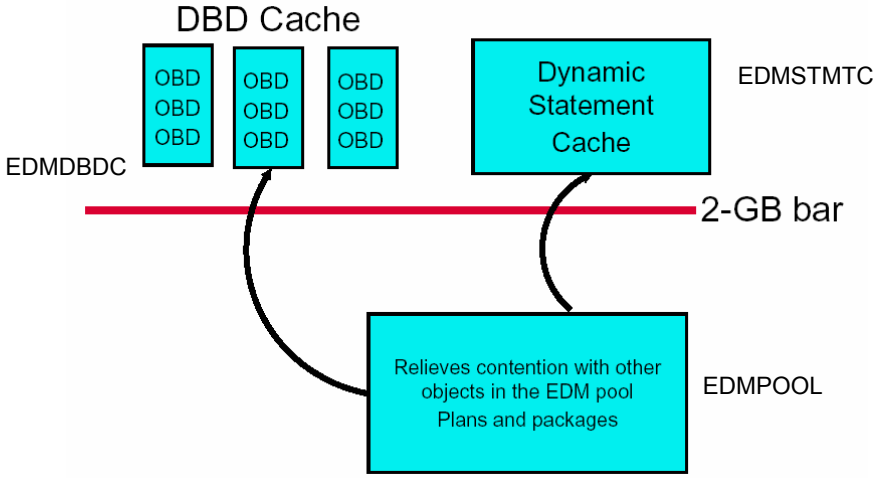

© 2004 IBM Corporation

How much is 64-bit virtual support worth for you? ...

- At first sight, re-engineering of DBM1 to exploit 64-bit should make a very significant difference in terms of providing significant VSCR
  - More concurrent threads
  - Higher transaction throughput
- Now consider V6 or V7 installation achieving significant VSCR
  - Maximum use of Dataspace bufferpools (400-800MB)
  - Using dataspace extension to EDM Pool for GDSC (80-160MB)
- What is the 'net' benefit in V8
  - Dataspace Lookaside Pool (80-120MB) is eliminated
  - Buffer Manager control blocks (70-120MB) going above 2GB bar
  - Other items going above the 2GB bar (compression dictionaries, certain EDM/RID pool components, thread sort pool, etc.)

© 2004 IBM Corporation

How much is 64-bit virtual support worth for you? …

- Most of the thread storage stays below the 2GB bar
  - Agent Local, Stack Storage
  - Local Dynamic Statement Cache (0-400MB)
  - Expect some regression
- How many additional threads can be supported will depend on 'Thread Footprint' which will vary by workload depending on
  - Duration of thread
  - SQL workload
  - RELEASE parameter setting on plan/package bind
  - Effectiveness of thread storage contraction (CONTSTOR=YES)
- Net benefit might only work out at 300-500MB for some installations
- Still need to monitor and track VSTOR usage with IFCID 225

© 2004 IBM Corporation

DB2 PE Statistics Report V8 - Extract

DBM1 AND MVS STORAGE BELOW 2 GB		QUANTITY
TOTAL DBM1 STORAGE BELOW 2 GB	(MB)	81.70
TOTAL GETMAINED STORAGE	(MB)	34.06
VIRTUAL BUFFER POOLS	(MB)	N/A
VIRTUAL POOL CONTROL BLOCKS	(MB)	N/A
EDM POOL	(MB)	32.12
COMPRESSION DICTIONARY	(MB)	N/A
CASTOUT BUFFERS	(MB)	N/A
DATA SPACE LOOKASIDE BUFFER	(MB)	N/A
HIPERPOOL CONTROL BLOCKS	(MB)	N/A
DATA SPACE BP CONTROL BLOCKS	(MB)	N/A
TOTAL VARIABLE STORAGE	(MB)	33.74
TOTAL AGENT LOCAL STORAGE	(MB)	29.23
TOTAL AGENT SYSTEM STORAGE	(MB)	20.36
NUMBER OF PREFETCH ENGINES		4.00
NUMBER OF DEFERRED WRITE ENGINES		207.00
NUMBER OF CASTOUT ENGINES		0.00
NUMBER OF GBP WRITE ENGINES		0.00
NUMBER OF P-LOCK/NOTIFY EXIT ENGINES		0.00
TOTAL AGENT NON-SYSTEM STORAGE	(MB)	8.86
TOTAL NUMBER OF ACTIVE USER THREADS		14.35
RDS OP POOL	(MB)	0.00
RID POOL	(MB)	0.00
PIPE MANAGER SUB POOL	(MB)	0.00
LOCAL DYNAMIC STMT CACHE CNTL BLKS	(MB)	0.99

...

© 2004 IBM Corporation

DB2 PE Statistics Report V8 – Extract …

```
...
   THREAD COPIES OF CACHED SQL STATEMENTS          0.00
   BUFFER & DATA MANAGER TRACE TBL       (MB)      0.00
  TOTAL FIXED STORAGE                    (MB)      0.21
  TOTAL GETMAINED STACK STORAGE          (MB)     13.69
 STORAGE CUSHION                         (MB)     18.16

 24 BIT LOW PRIVATE                      (MB)      0.14
 24 BIT HIGH PRIVATE                     (MB)      0.48
 31 BIT EXTENDED LOW PRIVATE             (MB)     33.57
 31 BIT EXTENDED HIGH PRIVATE            (MB)    100.61
 EXTENDED REGION SIZE (MAX)              (MB)   1610.00
 EXTENDED CSA SIZE                       (MB)    280.86

DBM1 STORAGE ABOVE 2 GB                                 QUANTITY
 ------------------------------------------  ------------------
 FIXED STORAGE                           (MB)              2.54
 GETMAINED STORAGE                       (MB)            270.57
    COMPRESSION DICTIONARY               (MB)              0.00
    CACHED DYNAMIC SQL STATEMENTS (MAX)  (MB)            100.38
    DBD CACHE (MAX)                      (MB)            100.38
 VARIABLE STORAGE                        (MB)             67.29
 VIRTUAL BUFFER POOLS                    (MB)              0.00
 VIRTUAL POOL CONTROL BLOCKS             (MB)              0.00
 CASTOUT BUFFERS                         (MB)              0.00
```


© 2004 IBM Corporation

Real Storage Use

- Important subsystems such as DB2 should not be paging to auxiliary storage
- RMF Monitor II Address Space Data (ASD) Report
  - Real storage usage by DB2 DBM1 address space
- RMF Monitor III Option 3 Storage Frames (STORF) resource
  - Number of REAL+AUX frames used by DB2 DBM1 address space
- UIC value highlighted in the ASD report is a good indicator if you are running out of real frames
  - If this value starts to dive it is an indication that you are starting to run out of real storage

    > 200 (**GREEN**)

    100 - 200 (**AMBER**)

    < 100 (**RED**)

© 2004 IBM Corporation

RMF Mon II ASD Report

```
                    RMF - ASD Address Space State Data          Line 1 of 44

                          CPU=  2/  2 UIC=2540 PR=   0          System= NB01 Total
16:54:35            S C  R  DP   CS   ESF   CS   TAR  X  PIN   ES  TX  SWAP   WSM
JOBNAME  SRVCLASS   P L  LS PR    F         TAR   WSS  M   RT   RT  SC   RV     RV

*MASTER* SYSTEM     1 NS    FF 4649               0     ----      0     0
PCAUTH   VEL30STC   1 NS    F7  101               0  X  ----      0     0
RASP     SYSTEM     1 NS    FF  155               0  X  ----      0     0
TRACE    VEL30STC   1 NS    F7   67               0  X  ----      0     0
DUMPSRV  SYSTEM     1 NS    FF   73               0     ----      0     0
XCFAS    SYSTEM     1 NS    FF 1220               0  X  ----      0     0
GRS      SYSTEM     1 NS    FF 1440               0  X  ----      0     0
SMSPDSE  SYSTEM     1 NS    FF 1776             432  X  ----      0     0
CONSOLE  SYSTEM     1 NS    FF  381               0  X  ----      0     0
WLM      SYSTEM     1 NS    FF  898               0  X  ----      0     0
ANTMAIN  SYSTEM     1 NS    FF 2009               0  X  ----      1   998
ANTAS000 VEL30STC   1 NS    F7 2293               0  X  ----      1   998
OMVS     SYSTEM     1 NS    FF 23.2K           20.5K X  ----      0     0
IEFSCHAS SYSTEM     1 NS    FF   33               0  X  ----      1     0
JESXCF   SYSTEM     1 NS    FF   86               0  X  ----      0   998
ALLOCAS  SYSTEM     1 NS    FF 1628               0  X  ----      1     0
```


© 2004 IBM Corporation

RMF III Option 3 STORFRMF

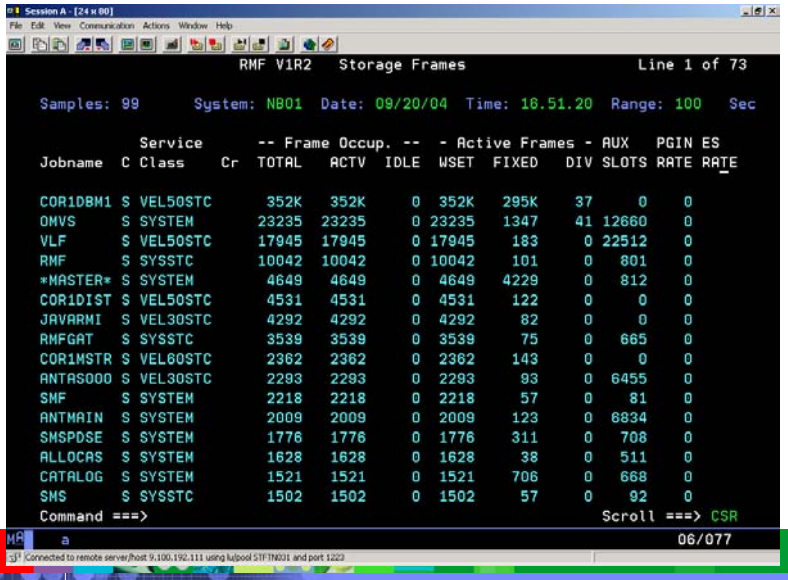

```
                          RMF V1R2  Storage Frames                 Line 1 of 73

 Samples: 99      System: NB01  Date: 09/20/04  Time: 18.51.20  Range: 100   Sec

            Service       -- Frame Occup. --   - Active Frames -  AUX   PGIN ES
 Jobname  C Class     Cr  TOTAL   ACTV  IDLE   WSET  FIXED   DIV SLOTS RATE RATE

 COR1DBM1 S VEL50STC      352K   352K     0  352K   295K    37     0    0
 OMVS     S SYSTEM      23235  23235     0 23235   1347    41 12660    0
 VLF      S VEL50STC    17945  17945     0 17945    183     0 22512    0
 RMF      S SYSSTC      10042  10042     0 10042    101     0   801    0
 *MASTER* S SYSTEM       4649   4649     0  4649   4229     0   812    0
 COR1DIST S VEL50STC     4531   4531     0  4531    122     0     0    0
 JAVARMI  S VEL30STC     4292   4292     0  4292     82     0     0    0
 RMFGAT   S SYSSTC       3539   3539     0  3539     75     0   665    0
 COR1MSTR S VEL60STC     2362   2362     0  2362    143     0     0    0
 ANTAS000 S VEL30STC     2293   2293     0  2293     93     0  6455    0
 SMF      S SYSTEM       2218   2218     0  2218     57     0    81    0
 ANTMAIN  S SYSTEM       2009   2009     0  2009    123     0  6834    0
 SMSPDSE  S SYSTEM       1776   1776     0  1776    311     0   708    0
 ALLOCAS  S SYSTEM       1628   1628     0  1628     38     0   511    0
 CATALOG  S SYSTEM       1521   1521     0  1521    706     0   668    0
 SMS      S SYSSTC       1502   1502     0  1502     57     0    92    0
 Command ===>                                                Scroll ===> CSR
```


© 2004 IBM Corporation

Key messages

- V8 64-bit support will not eliminate VSC below the 2GB bar in DBM1
- Will provide valuable additional relief, but will vary by installation
- Will be able to exploit all available processor storage on latest processor models (currently 256GB, current DB2 limit of 1TB)
  - Some additional number of active threads and DBATs
  - XXL bufferpools to eliminate I/O and speed up remainder
  - Increase exploitation of ESA Compression
  - Larger thread Sort Pool
- Must have sufficient real storage to fully back increased usage
- Must continue to plan for, monitor and tune VSTOR usage below 2GB bar
- Additional exploitation of 64-bit virtual to move thread related storage above 2GB bar under consideration for future release

© 2004 IBM Corporation

BACKUP SYSTEM / RESTORE SYSTEM

DB2 Information Management Software

© 2004 IBM Corporation

Customer Requirements

- Need a fast non-disruptive way to backup/recovery all DB2 data
  - It is too slow and difficult to manage backups at table/index level
  - Need a way to ensure data and logs in backup are consistent for restart recovery
  - Consistent backup with no impact to applications
- Backups can be used to support
  - Point-In-Time Recovery on application errors
  - Disaster Recovery or Cloning Systems

© 2004 IBM Corporation

Prior to DB2 V8 – Volume-level backup

-SET LOG SUSPEND – Suspend writes

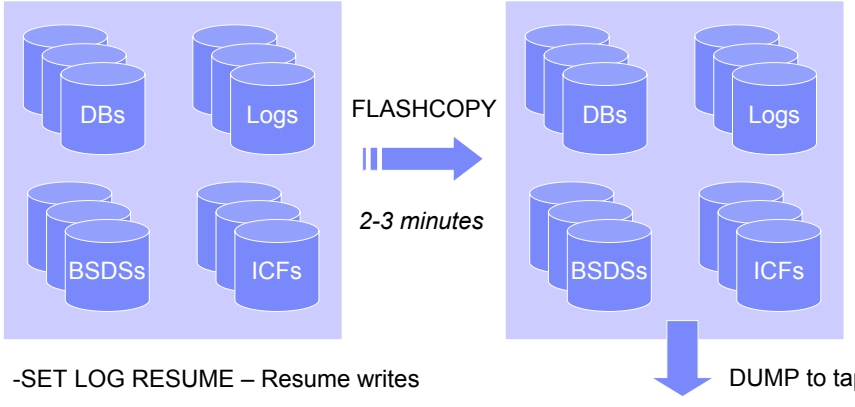

-SET LOG RESUME – Resume writes

© 2004 IBM Corporation

Prior to DB2 V8 – Recovery to the time of the backup

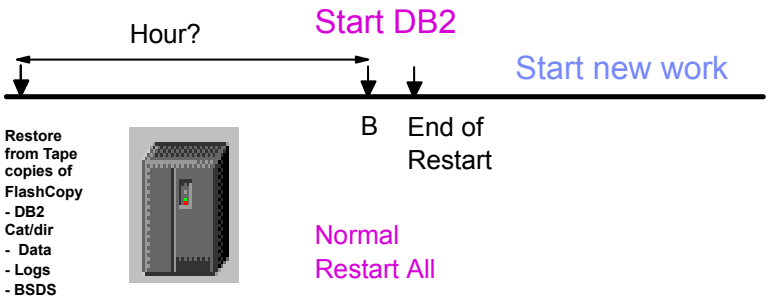

Analysis: Fast Restart – loss of data since FlashCopy

© 2004 IBM Corporation

Prior to DB2 V8 – Recovery to any point-in-time

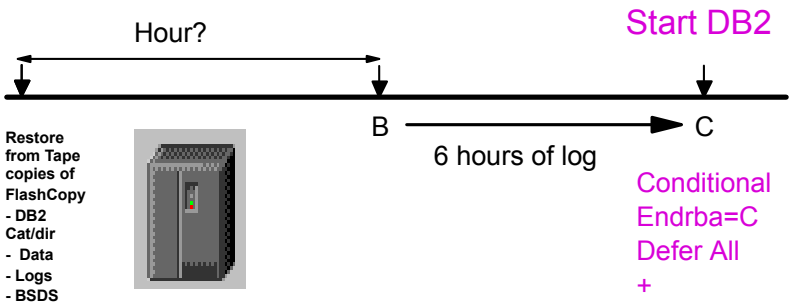

Analysis: Longer to come up – reduced data loss
Complex DB2 D/R procedures required

© 2004 IBM Corporation

With DB2 V8 – System PITR

- Provides an easier and less disruptive way for fast volume-level backup and recovery
  - Use FlashCopy API to backup DB2 data and logs
  - No longer need to suspend logs
  - Backup and recovery are managed by DB2 and DFSMShsm
- Two new utilities:
  - BACKUP SYSTEM
  - RESTORE SYSTEM
- Backup copies can also be used for:
  - Disaster recovery
  - System cloning

© 2004 IBM Corporation

DFSMS constructs and definitions

- COPYPOOL
  - New SMS construct
  - Set of SMS storage groups – maximum 256 storage groups
  - Has a VERSIONS attribute to specify the number of copy versions to be maintained on DASD – maximum is 85
  - Each DB2 system can have two SMS COPYPOOLs
    DATABASE COPYPOOL (DSN$location_name$DB)
    LOG COPYPOOL (DSN$location_name$LG)
- COPYPOOL BACKUP storage group
  - New storage group type
  - Used to hold volume copies of DASD defined in the COPYPOOL

© 2004 IBM Corporation

Requirements in a DB2 environment

- All volumes containing DB2 data sets, including the DB2 catalog and directory, the logs, and the BSDS, must be SMS managed.
- DB2 logs and BSDS should be separated from the rest of the DB2 data sets:
  - Own pool of volumes
  - Own ICF catalog
  - Defined in a separate SMS storage group
- Special care should be taken to ensure that the ICF catalogs stay synchronized with the data.
  - Must have separate ICF catalog for data and logs
  - Do not share ICF catalog with non-DB2 data

© 2004 IBM Corporation

Example of SMS definitions

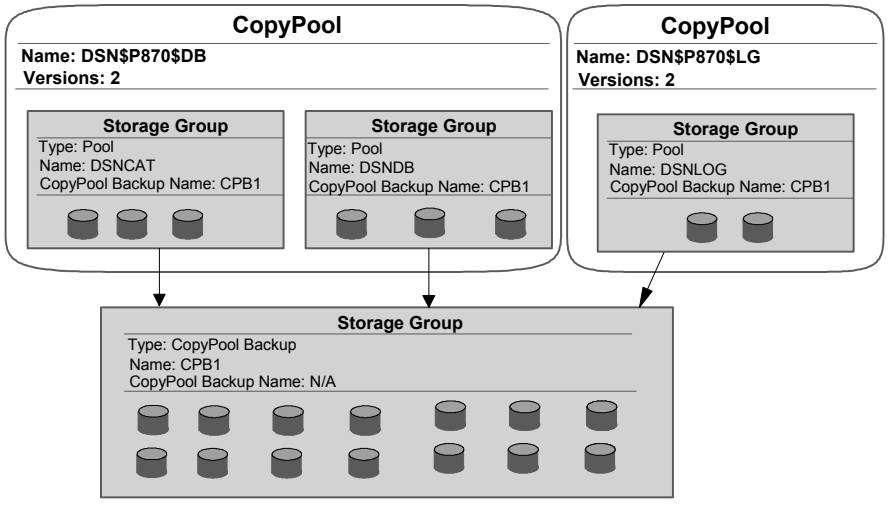

© 2004 IBM Corporation

```
                        ISMF PRIMARY OPTION MENU - z/OS DFSMS V1 R5
Enter Selection or Command ===> P

Select one of the following options and press Enter:
0  ISMF Profile              - Specify ISMF User Profile
1  Data Set                  - Perform Functions Against Data Sets
2  Volume                    - Perform Functions Against Volumes
3  Management Class          - Specify Data Set Backup and Migration Criteria
4  Data Class                - Specify Data Set Allocation Parameters
5  Storage Class             - Specify Data Set Performance and Availability
6  Storage Group             - Specify Volume Names and Free Space Thresholds
7  Automatic Class Selection - Specify ACS Routines and Test Criteria
8  Control Data Set          - Specify System Names and Default Criteria
9  Aggregate Group           - Specify Data Set Recovery Parameters
10 Library Management        - Specify Library and Drive Configurations
11 Enhanced ACS Management   - Perform Enhanced Test/Configuration Management
C  Data Collection           - Process Data Collection Function
L  List                      - Perform Functions Against Saved ISMF Lists
P  Copy Pool                 - Specify Pool Storage Groups for Copies
R  Removable Media Manager   - Perform Functions Against Removable Media
X  Exit                      - Terminate ISMF
Use HELP Command for Help; Use END Command or X to Exit.
```


© 2004 IBM Corporation


```
                              COPY POOL DEFINE

SCDS Name  . . : SMSCTL.SCDS
Copy Pool Name : DSN$P870$DB

To DEFINE Copy Pool, Specify:
  Description ==> COPY POOL FOR P870 DATA
              ==>
  Number of Recoverable DASD Fast
   Replicate Backup Versions  . . . . 2       (1 to 85 or blank)
  Storage Group Names:  (specify 1 to 256 names)
   ==> DSNCAT
   ==> DSNDB
   ==>
```

```
                              COPY POOL DEFINE

SCDS Name  . . : SMSCTL.SCDS
Copy Pool Name : DSN$P870$LG

To DEFINE Copy Pool, Specify:
  Description ==> COPY POOL FOR P870 LOGS
              ==>
  Number of Recoverable DASD Fast
   Replicate Backup Versions  . . . . 2       (1 to 85 or blank)
  Storage Group Names:  (specify 1 to 256 names)
   ==> DSNLOG
   ==>
```


© 2004 IBM Corporation

```
                    ISMF PRIMARY OPTION MENU - z/OS DFSMS V1 R5
Enter Selection or Command ===> P

Select one of the following options and press Enter:
0  ISMF Profile              - Specify ISMF User Profile
1  Data Set                  - Perform Functions Against Data Sets
2  Volume                    - Perform Functions Against Volumes
3  Management Class          - Specify Data Set Backup and Migration Criteria
4  Data Class                - Specify Data Set Allocation Parameters
5  Storage Class             - Specify Data Set Performance and Availability
6  Storage Group             - Specify Volume Names and Free Space Thresholds
7  Automatic Class Selection - Specify ACS Routines and Test Criteria
8  Control Data Set          - Specify System Names and Default Criteria
9  Aggregate Group           - Specify Data Set Recovery Parameters
10 Library Management        - Specify Library and Drive Configurations
11 Enhanced ACS Management   - Perform Enhanced Test/Configuration Management
C  Data Collection           - Process Data Collection Function
L  List                      - Perform Functions Against Saved ISMF Lists
P  Copy Pool                 - Specify Pool Storage Groups for Copies
R  Removable Media Manager   - Perform Functions Against Removable Media
X  Exit                      - Terminate ISMF
Use HELP Command for Help; Use END Command or X to Exit.
```


© 2004 IBM Corporation


```
                    COPY POOL BACKUP STORAGE GROUP DEFINE

SCDS Name . . . .  : SMSCTL.SCDS
Storage Group Name : CPB1

To DEFINE Storage Group, Specify:

Description ==> COPY POOL BACKUP STORAGE GROUP FOR P870 DATA + LOG
```

```
                          POOL STORAGE GROUP ALTER
Command ===>

SCDS Name . . . . . : SMSCTL.SCDS
Storage Group Name  : DSNCAT
To ALTER Storage Group, Specify:

 Description ==> FOR P870 CATALOG AND DIRECTORY
             ==>

 Auto Migrate . . N  (Y, N, I or P)   Migrate Sys/Sys Group Name . .
 Auto Backup  . . N  (Y or N)         Backup Sys/Sys Group Name  . .
 Auto Dump  . . . N  (Y or N)         Dump Sys/Sys Group Name  . . .
 Overflow . . . . N  (Y or N)         Extend SG Name . . . . . .
                                      Copy Pool Backup SG Name . . . CPB1
 Dump Class . . .                     (1 to 8 characters)
 ...
```


© 2004 IBM Corporation

DFSMShsm commands

- FRBACKUP – Create a Fast Replication backup
  - `FRBACKUP COPYPOOL(DSN$locn$DB) PREPARE`
  - `FRBACKUP COPYPOOL(DSN$locn$DB) NOVTOCENQ TOKEN(token)`
- FRRECOV – Recover a pool of volumes (or a target volume)
  - `FRRECOV COPYPOOL(DSN$locn$DB) VERIFY(YES) TOKEN(token)`
- QUERY – Determine status of physical backup
  - `QUERY COPYPOOL(DSN$locn$DB)`
- LIST – Provide information about each backup version of copy pool
  - `LIST COPYPOOL(DSN$locn$DB)`

© 2004 IBM Corporation

BACKUP SYSTEM utility

- Invokes DFSMShsm to take fast volume copies of the DB2 data and / or logs
- No DB2 quiesce point is required, nothing stops as in SET LOG SUSPEND
  - BACKUP SYSTEM FULL
    - Allow recovery of the entire system in later stage
    - Have to define the "database" and "log" COPYPOOLs
    - Backup both database and then log (active logs and BSDS)
  - BACKUP SYSTEM DATA ONLY
    - Only "database" COPYPOOL has to be defined for database backup

© 2004 IBM Corporation

BACKUP SYSTEM process

```
//BACKUP EXEC DSNUPROC,PARM='P870,BACKUP'
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
BACKUP SYSTEM DATA ONLY
```

- When BACKUP SYSTEM is issued, DB2
  - Suspends 32K page writes
  - Suspends data set creation, deletion, rename and extension operations
  - Prevents data sets from pseudo-close
  - *Records* the Recover Based Log Point (RBLP) in DBD01

    Starting point for RESTORE SYSTEM to apply log records

    Data sharing – Oldest system checkpoint across all members of group

© 2004 IBM Corporation

BACKUP SYSTEM process…

- During the backup . . .
  - Invokes DFSMShsm to take FlashCopy of 'DB' COPYPOOL

    `FRBACKUP COPYPOOL(DSN$locn$DB) NOVTOCENQ TOKEN(token)`
  - Invokes DFSMShsm to take FlashCopy of the 'LG' COPYPOOL, if for Backup System Full

    `FRBACKUP COPYPOOL(DSN$locn$LG) NOVTOCENQ TOKEN(token)`
  - Each member updates BSDS with the system backup information

    Up to 50 entries

    In data sharing only the submitting member logs BSDS information
  - Resume the quiesced activities

© 2004 IBM Corporation

```
DSNU000I DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = BACKUP
DSNU1044I DSNUGTIS - PROCESSING SYSIN AS EBCDIC
DSNU050I DSNUGUTC - BACKUP SYSTEM DATA ONLY
DSNU1600I DSNUVBBD - BACKUP SYSTEM UTILITY FOR DATA STARTING,
COPYPOOL = DSN$P870$DB
TOKEN = X'D7F8F7F0BB8E5142855EF907BB8E502C0AC4'.
DSNU1614I DSNUVBBD - BACKUP SYSTEM UTILITY FOR DATA COMPLETED
SUCCESSFULLY,
COPYPOOL = DSN$P870$DB
TOKEN = X'D7F8F7F0BB8E5142855EF907BB8E502C0AC4'
ELAPSED TIME = 00:00:16.
DSNU1602I DSNUVBBD - BACKUP SYSTEM UTILITY COMPLETED, ELAPSED TIME =
00:00:19.
```

```
                    BACKUP SYSTEM UTILITY HISTORY
                         SUBSYSTEM ID P870
                       20:59:59 JULY 22, 2004

 START STCK                                      DATA COMPLETE  DATA/LOG  COMPLETE
    DATA                 LOG            RBLP        LRSN         DATE     LTIME
---------------- ---------------- ------------ ------------ --------------------
BB8E5142855EF907 0000000000000000 BB8E502C0AC4 BB8E502C0AC4 2004/07/22 13:59:45
                TOKEN = D7F8F7F0BB8E5142855EF907BB8E502C0AC4
BA1CCAA9A3820B8A BA1CCAB1FF964845 00120CDAC090 00120CDAC090 2003/10/02 16:12:43
                TOKEN = D7F8F7F0BA1CCAA9A3820B8A00120CDAC090
```


© 2004 IBM Corporation

BACKUP SYSTEM

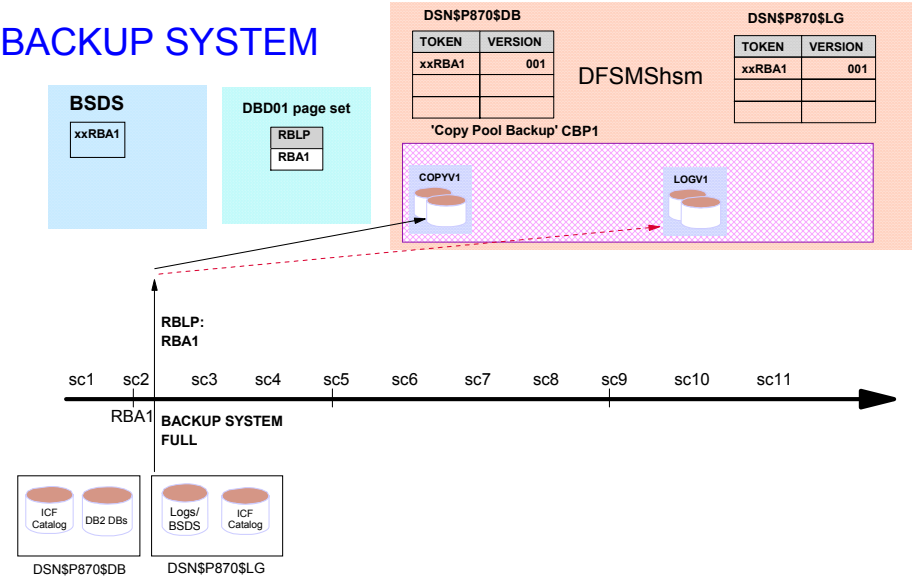

© 2004 IBM Corporation

BACKUP SYSTEM

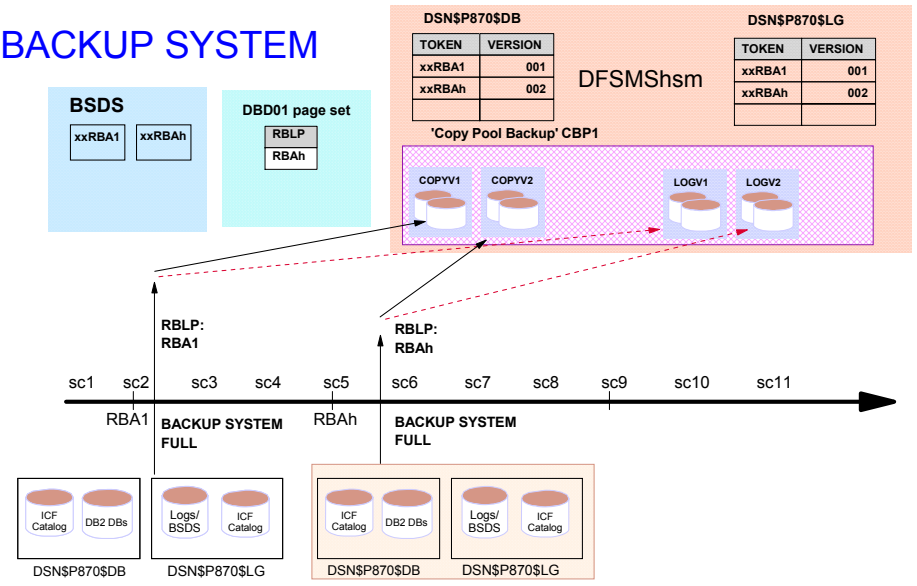

© 2004 IBM Corporation

BACKUP SYSTEM

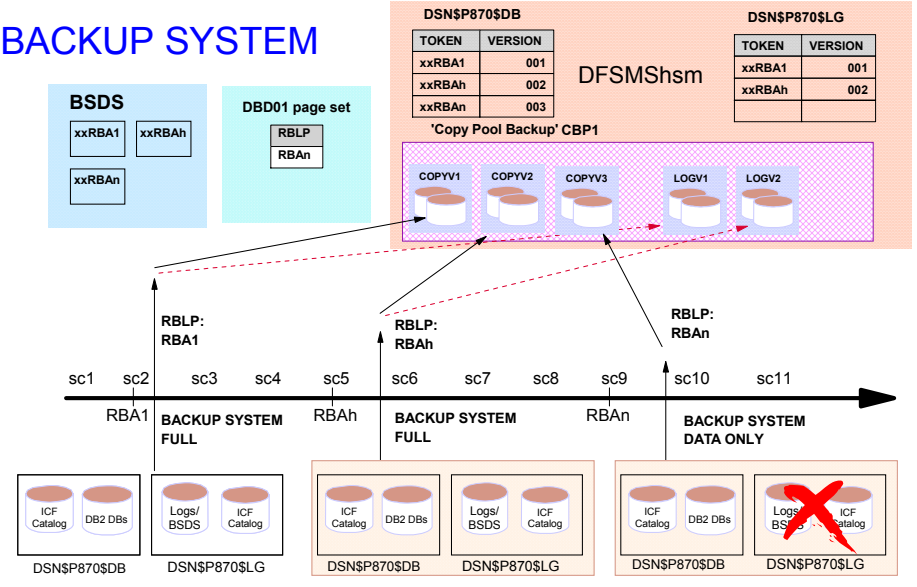

© 2004 IBM Corporation

RESTORE SYSTEM utility

- RESTORE SYSTEM utility is only needed to recover system to an arbitrary PIT
- To recover system only to the PIT at which the backup copy was taken
  - Use copies from BACKUP SYSTEM FULL to restore the database and log copypool

```
FRRECOV COPYPOOL(DSN$locn$DB) VERIFY(YES) TOKEN(token)
FRRECOV COPYPOOL(DSN$locn$LG) VERIFY(YES) TOKEN(token)
```

- Start DB2 and inflight URs are backed out

© 2004 IBM Corporation

System level restore to an arbitrary PIT

- RESTORE SYSTEM utility is needed
  - Use copies from BACKUP SYSTEM FULL or DATA ONLY

    RESTORE SYSTEM does not restore LOG backup copies, therefore copies from DATA ONLY is enough
- Two phases:
  - RESTORE phase: recover the database volumes from the latest BACKUP version prior to the arbitrary PIT
  - LOG APPLY phase: apply log records to recover DB objects to the PIT
- Need to unallocate the ICF Catalogs for data copy pool before RESTORE SYSTEM
  - F CATALOG,OPEN
  - F CATALOG,UNALLOCATE(...)

© 2004 IBM Corporation

System level restore to an arbitrary PIT

- Establish the 'PITR' conditional restart record
  - CRESTART CREATE SYSPITR=log-point (*truncation RBA value*)
  - MUST be in New Function Mode
- Start DB2 with a PITR CRCR
  - DB2 system enters into System Recover Pending mode
  - Implictly apply DEFER ALL, FORWARD = NO (except for in-doubt URs) and Access(Maint)
  - Write logs to rollback uncommitted changes
  - Reset database restrict status and utility job status

© 2004 IBM Corporation

System level restore to an arbitrary PIT

- RESTORE SYSTEM
  - Restore the "database" COPYPOOL version that was taken by BACKUP SYSTEM prior to the specified PIT recovery point
  - Perform log apply function
- RESTORE SYSTEM with LOGONLY specified
  - Performs log apply function only
  - Note: this option can run in z/OS 1.3 without BACKUP SYSTEM utility

    Using Log Suspend/Resume and backup volumes manually
- Stop DB2 --> resets the system recover-pending status
- Recover all objects that are marked in RECP or RBDP state

© 2004 IBM Corporation

System level restore to an arbitrary PIT – Data Sharing

- Establish the LRSN truncation point on all active members
  - CRESTART CREATE SYSPITR= end-lrsn
- Delete all CF structures
- Group restart each active member with the SYSPITR CRCR
  - All members MUST be restarted
- Restore system
  - Similar to the steps as in the non-data sharing environment

© 2004 IBM Corporation

RESTORE SYSTEM

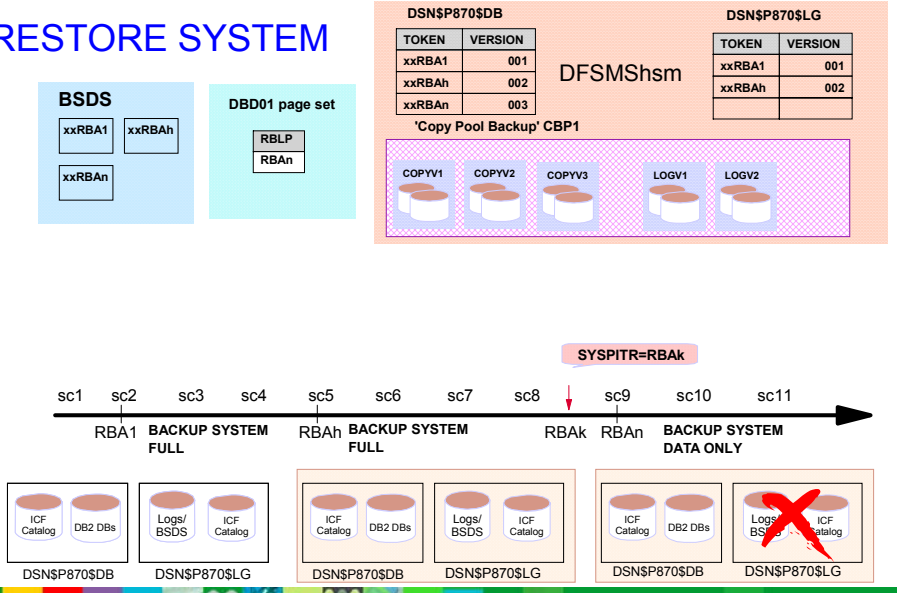

© 2004 IBM Corporation

RESTORE SYSTEM

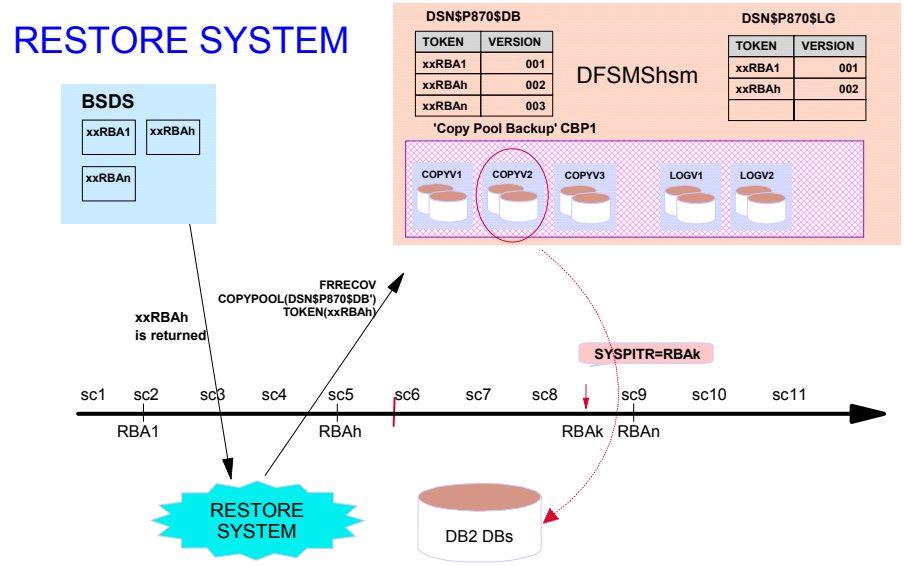

© 2004 IBM Corporation

RESTORE SYSTEM

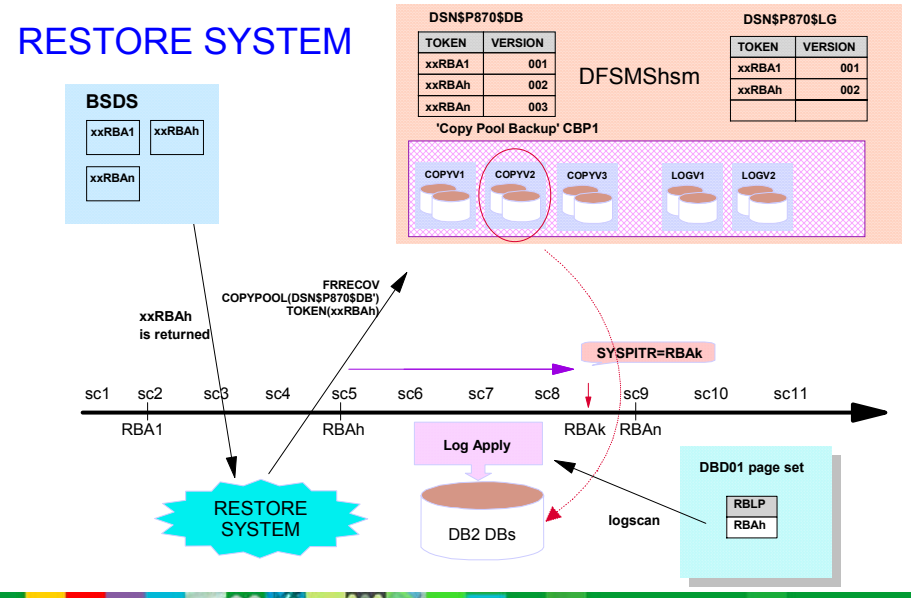

© 2004 IBM Corporation

System level restore – Notes

- When DB2 is in System Recover Pending state
  - Only RESTORE SYSTEM utility is allowed
  - START DATABASE command is not allowed
  - TERM UTIL command is not allowed
  - DISPLAY UTIL command will display only the status of RESTORE SYSTEM utility
  - SQL operation is not allowed

    Claim request on any DB2 objects will be rejected with a -904 SQL code (reason code of 00C20269)
- Restore of the database volumes is done in parallel

© 2004 IBM Corporation

System level restore – Log recovery

- DB2 reads the DBD01 header page to retrieve the recovery base log point (RBLP) → starting point for the log scan
- Handles table spaces and index spaces
  - CREATEs - will define data sets
  - DROPs
  - LOG NO events

    If LOG NO, the associated object is entered into RECP or RBDP state. Table spaces and indexes with COPY YES attribute will be marked RECP and the indexes with COPY NO will be marked RBDP
- Uses fast log apply (FLA) to recover objects in parallel

© 2004 IBM Corporation

System level restore – Log recovery (cont.)

- Log apply phase takes periodic checkpoints
  - Forces modified data pages to DASD and trigger system checkpoint
  - Updates the DBD01 header page with the new RBLP value
- At the end of log apply phase
  - Issues informational message if any object is marked RECP, RBDP or LPL during the log apply phase -> RESTORE SYSTEM RC=4
  - Reset the PITR state of each member
- RESTORE SYSTEM utility is restartable
  - In data sharing, only that member which issued the original RESTORE SYSTEM can issue the restart request

© 2004 IBM Corporation

Hardware and software prerequisites

- z/OS V1R5 and DFSMShsm
- DFSMShsm BCDS requires a block size of 6544
- DASD control units which support ESS Flashcopy APIs
- FlashCopy V2 is strongly recommended – allows 'source' and 'target' to be across LSS boundary, but must be within same ESS
- DB2 datasets must reside on SMS-managed volumes
- Must be in New Function Mode
- RESTORE SYSTEM LOGONLY can be executed under z/OS 1.3
  - Assumes you have used -Set Log Suspend
  - Manually dumped volumes (like today)
  - Manually restored volumes (like today)
  - Eliminates complex recovery procedures for Disaster Recovery

© 2004 IBM Corporation

Future directions

- *The FlashCopy source and target volumes can reside on different ESS*
- *DFSMShsm will automatically manage FlashCopy target volumes to tapes*
- *Use volume level backups as the source for DB2 object level recovery*
- *Manage data set level FlashCopy*
- *Support object level point-in-time recovery*
  - *Rollback uncommitted changes*

© 2004 IBM Corporation

Just published...

Disaster Recovery with DB2 UDB for z/OS
SG24-6370

Chapter 6 - SMS copy pools and DB2 point in time recovery

Chapter 19 - Local recovery: System PITR

Chapter 20 - Restart using tape dump of copy pools

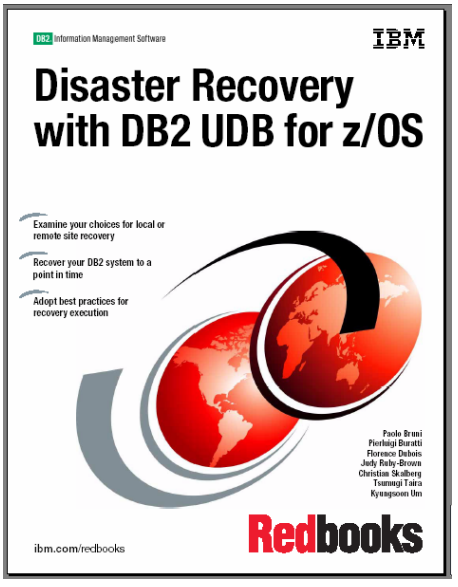

© 2004 IBM Corporation