

IBM Information

>>> On Demand

2007



Quicker DB2 Recovery - Tips From the Trenches

Robert Goodman



Certified DB2 Professional

Certified Business Continuity Planner

Data Servers – DB2 for z/OS



Act.Right.Now.

Quicker DB2 Recovery – Tips from the Trenches

Major Points

- Overview of the DB2 backup and recovery process
- Periodic tasks to streamline DB2 recoveries
- Comparison of backup and recovery strategies
- Recovery parameters that make a BIG difference!
- The latest recovery techniques and tricks of the trade



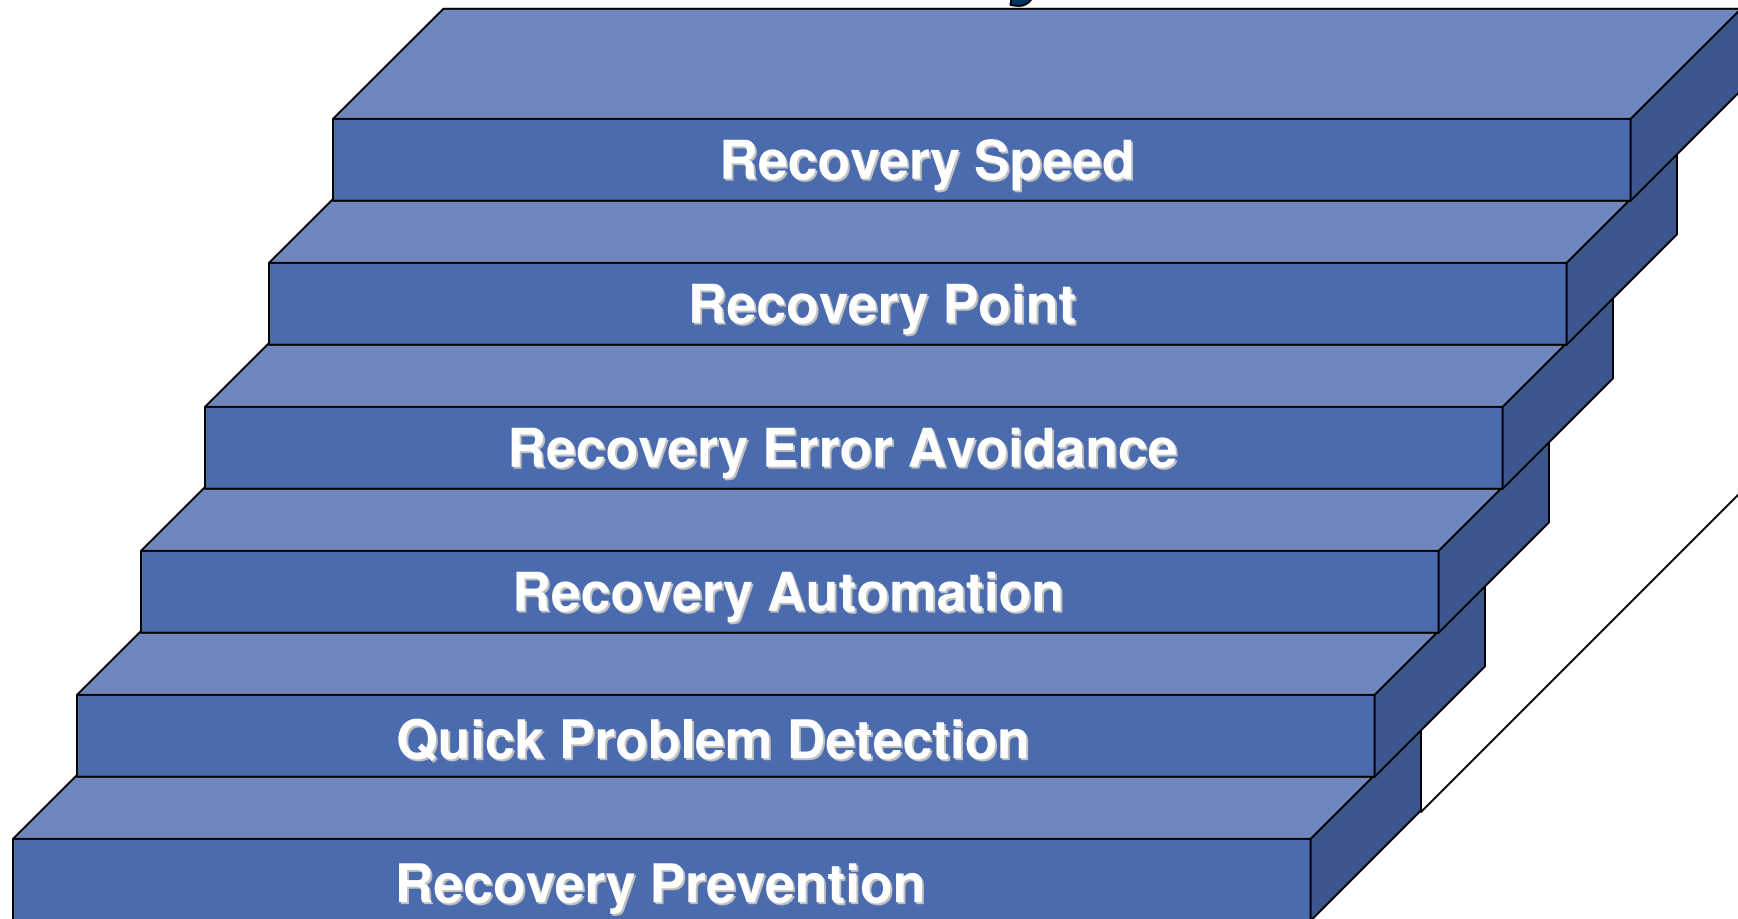
Act. Right. Now.



A Holistic Approach to Recovery



Recovery Success!



3



How Are Your Recovery Skills?



They Need to Be Rock Solid!

21st Century Databases Require Precise Skills!

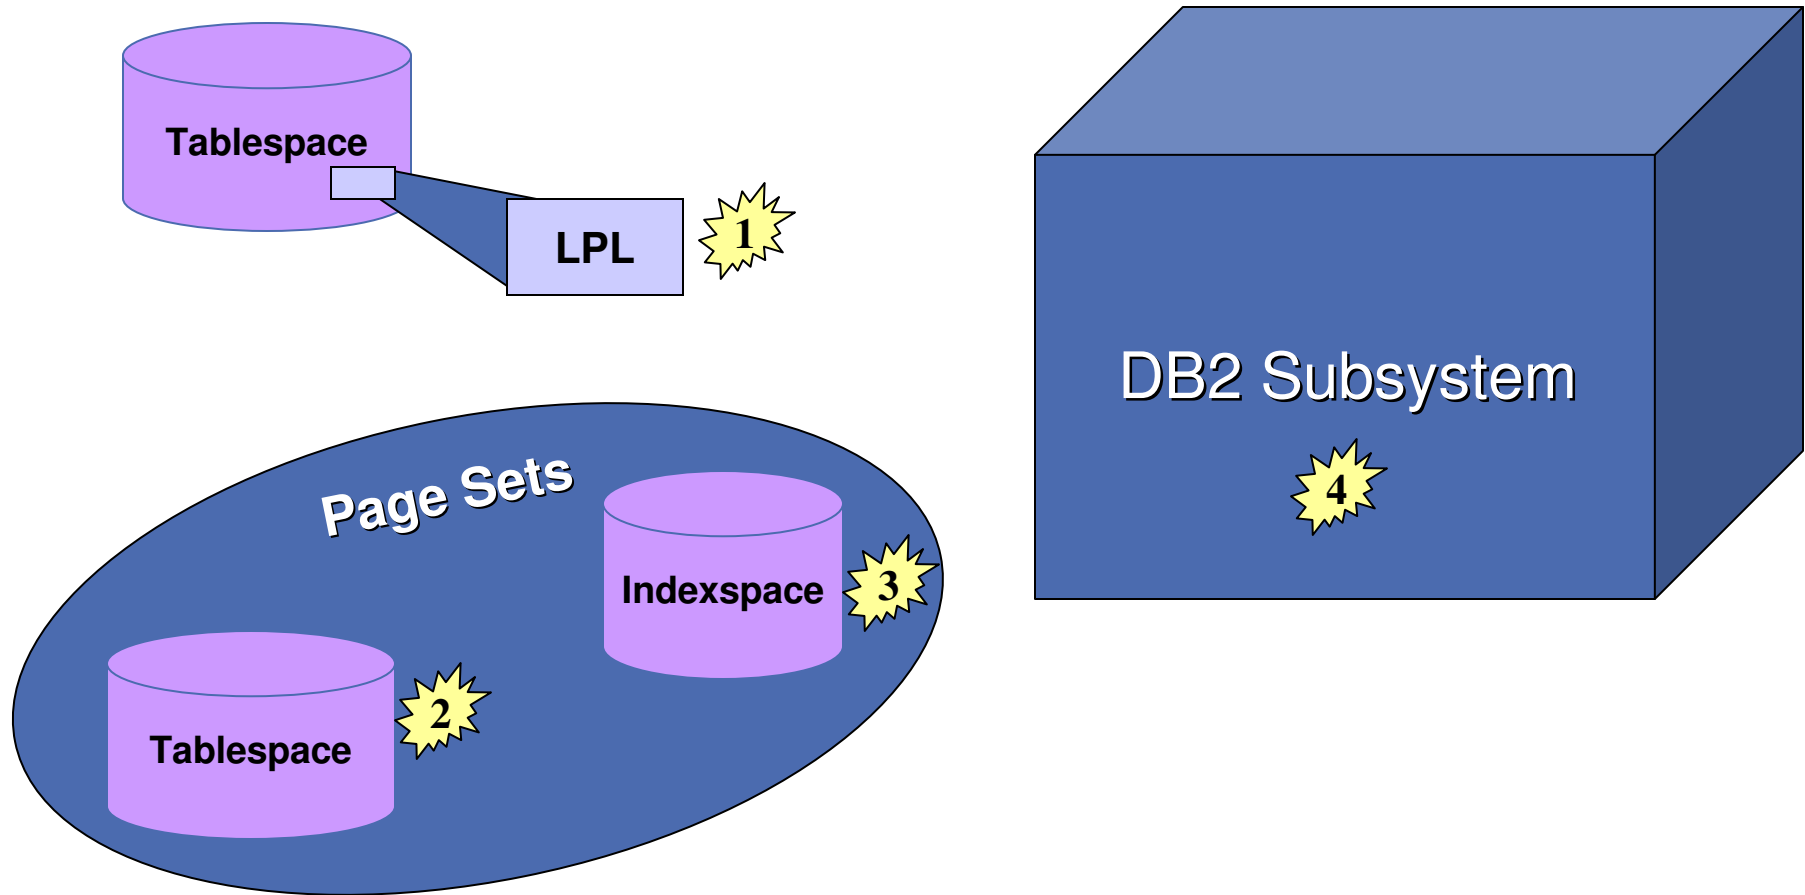
4



IBM INFORMATION ON DEMAND 2007

Act Right. Now.

DB2 Recoverable Objects

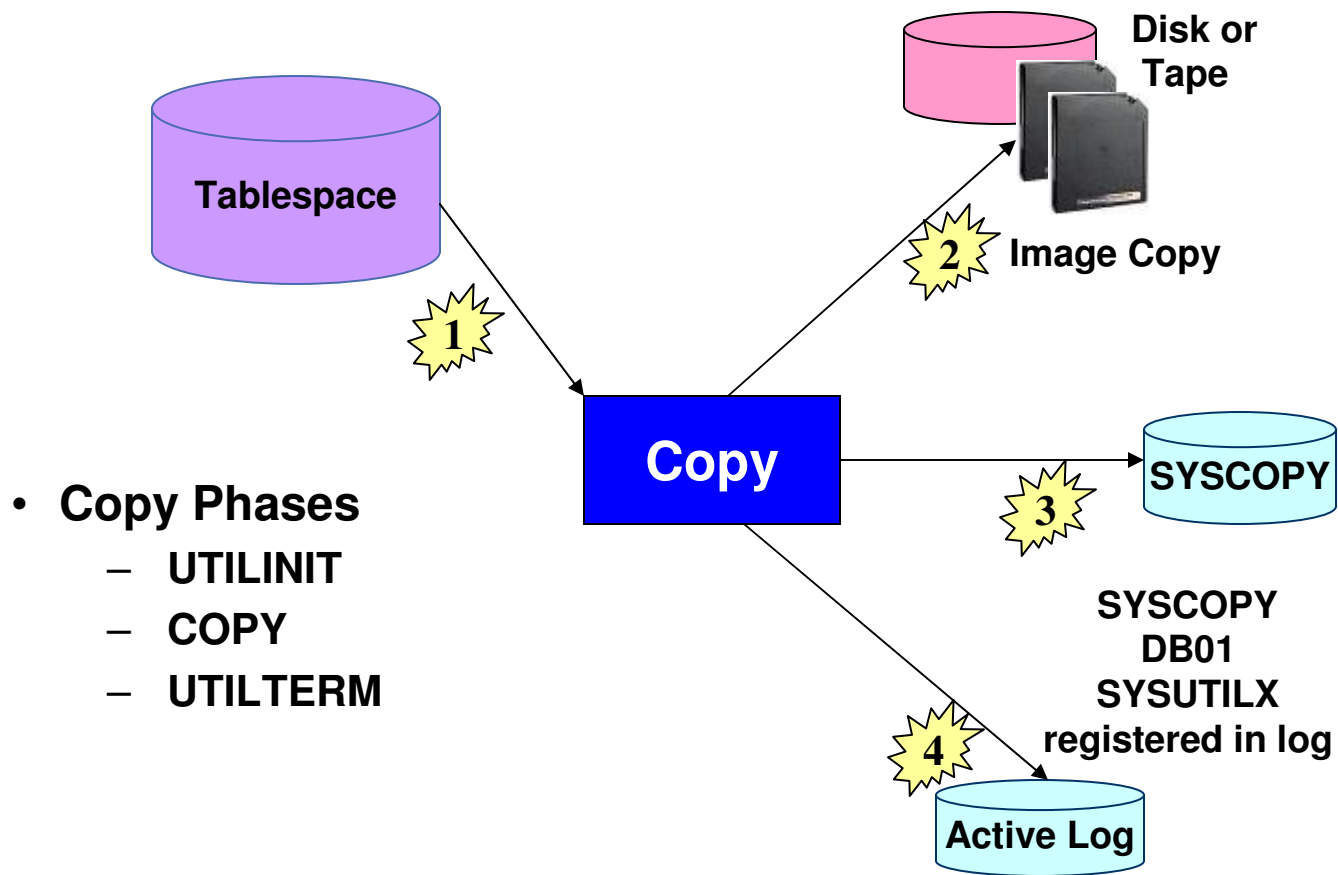


DB2 Does Not Backup or Recover Tables!

5



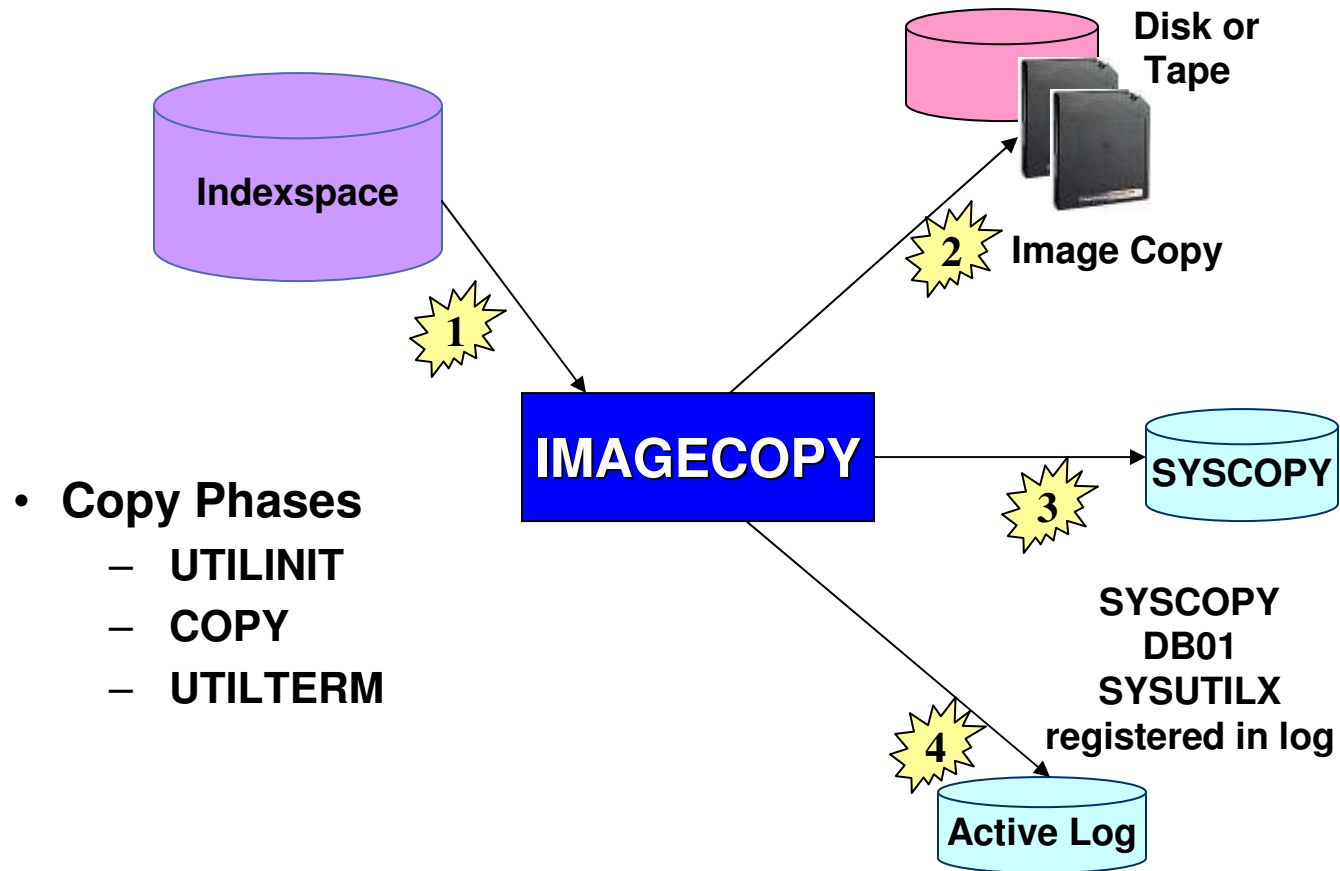
DB2 Tablespace Copy Concept



REORGs, LOADs & Rotates Require Full Image Copies to be Recoverable!



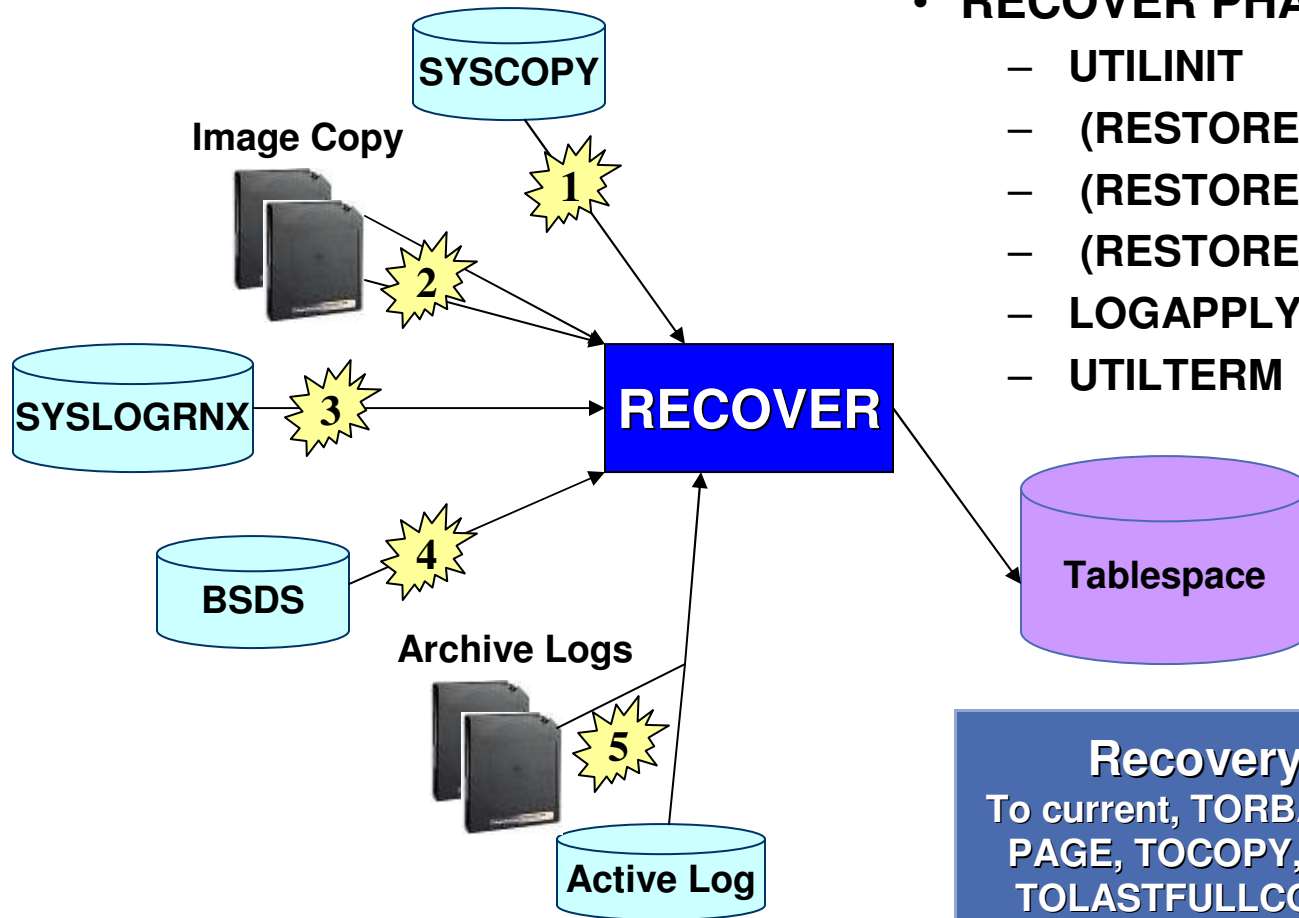
DB2 Indexspace Copy Concept



Make Sure That Copied Indexes Have the COPY YES attribute!



DB2 Tablespace Recovery Concept



• RECOVER PHASES

- UTILINIT
- (RESTORE) or
- (RESTORER and
- (RESTOREW)
- LOGAPPLY
- UTILTERM

Recovery Options
To current, TORBA, TOLOGPOINT,
PAGE, TOCOPY, TOLASTCOPY,
TOLASTFULLCOPY, LOGONLY


REPORT RECOVERY Tells What Will Be Used



REBUILD INDEX vs. RECOVER INDEX

- REBUILD INDEX

- All indexes can be rebuilt



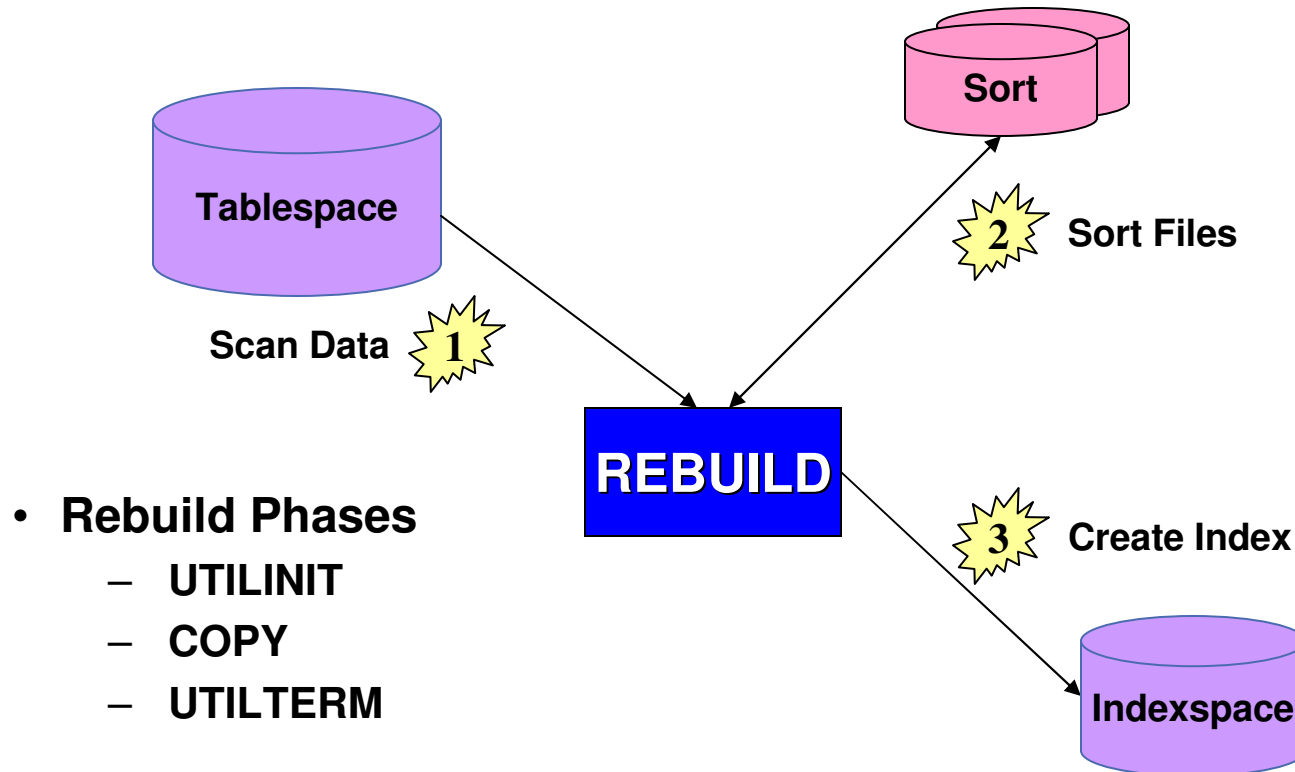
- RECOVER INDEX

- Prerequisites
 - ✓ ALTER INDEX with COPY YES option
 - ✓ IMAGECOPY index
- RECOVER INDEX available
- Can always use REBUILD INDEX if needed

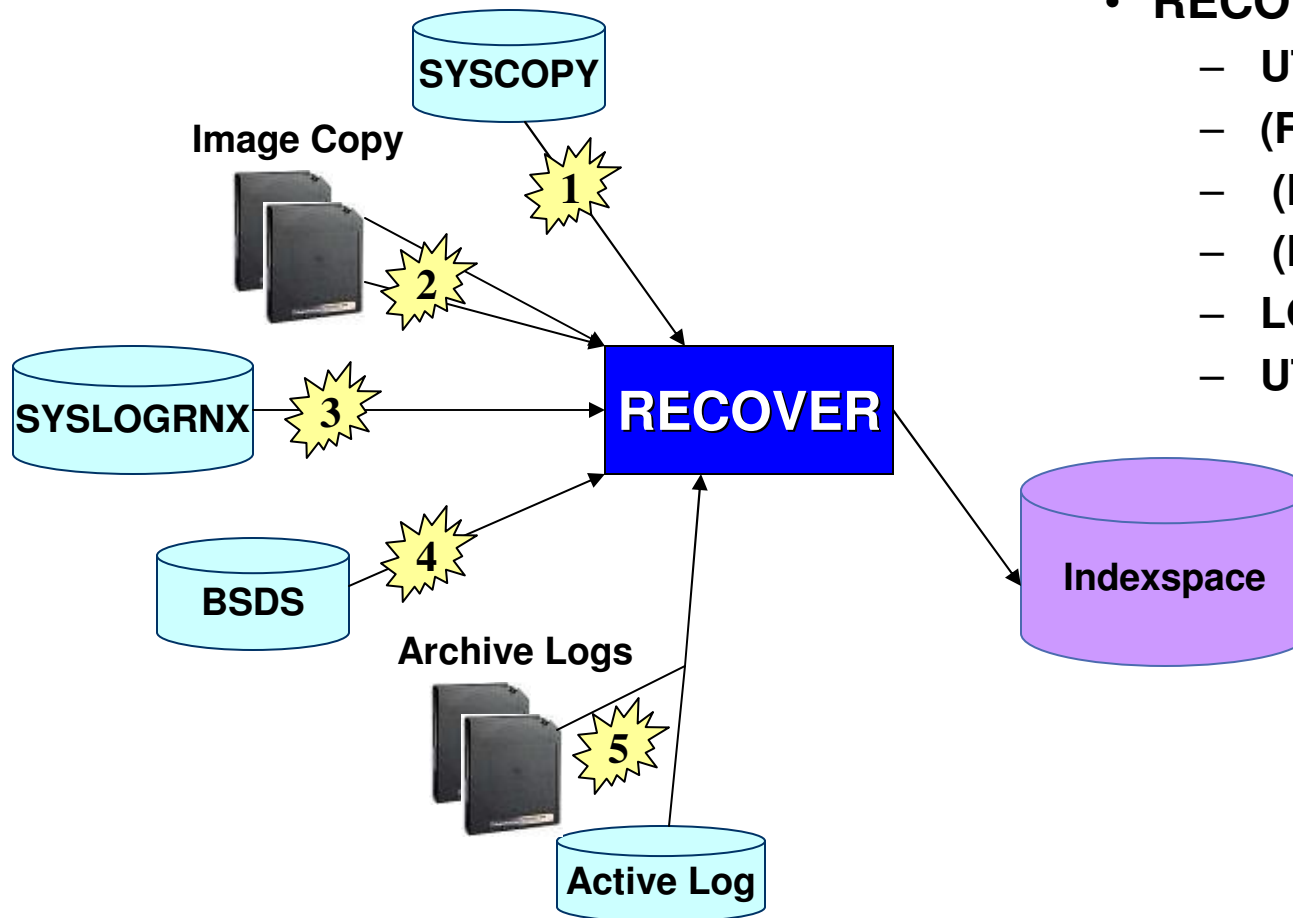
COPY NO Indexes Must Be Rebuilt



DB2 REBUILD INDEX Concept



RECOVER Index Concept




- **RECOVER PHASES**

- UTILINIT
- (RESTORE) or
- (RESTORER and
- (RESTOREW)
- LOGAPPLY
- UTILTERM



Critical Recovery Resources



- ✓ **Enough tape drives**
- ✓ **Adequate DB2 DASD**
- ✓ **Adequate work DASD**
- ✓ **Adequate Size Bufferpools**
- ✓ **Type DASD & Attach**
- ✓ **Job Initiators**
- ✓ **Adequate logs on DASD**

- **System Constraints**
 - **tape**
 - **CPU**
 - **I/O**
 - **memory**
 - **buffers**
 - **extents**



Things That Make a BIG Difference!

1. DB2 Subsystem 

2. Image Copies 

3. Housekeeping 

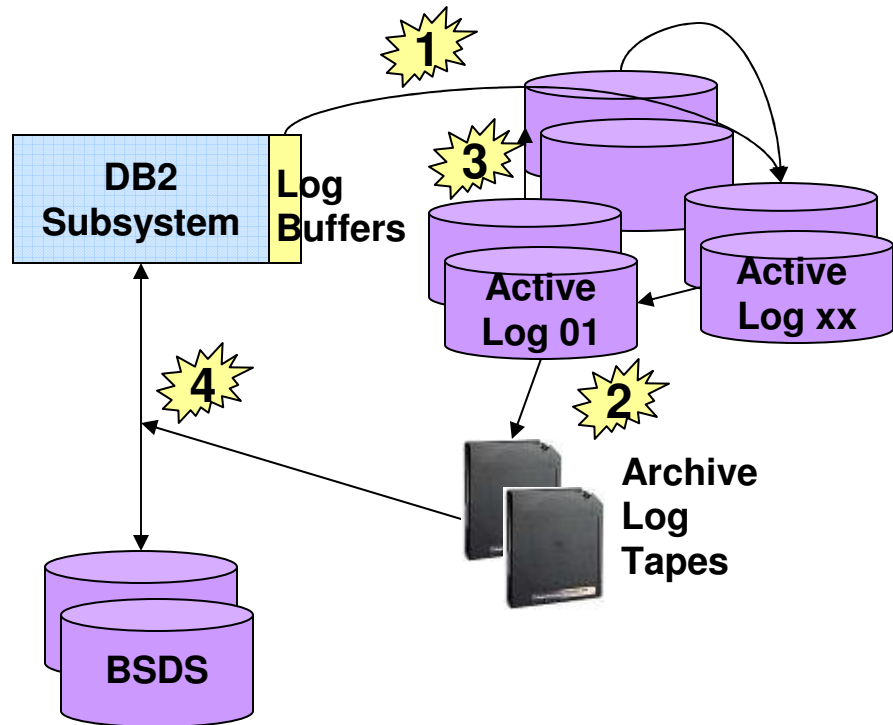

4. RECOVER Utility 

5. Index REBUILD/RECOVER 

6. Disaster Recovery Strategies 



Tip #1 Properly Size DB2 Logs



- Dual Logging
- Quantity & Span
 - on-site
 - off-site
 - DASD / tape / ATL
- Archives
 - clean up old logs
 - Change Log Inventory BSDS - DSNJU003
- Retention

Undersized Active Logs Can Force Recovery to Use Archive Tapes!



Tip #2 Enable Fast Log Apply



- Requires DB2 V6 or higher
- Must enable feature
- Log Apply Storage 0-100 mb (V8 default 100 mb)
- Up to 10 concurrent recoveries will use fast log apply
- Automatically invoked if storage is allocated
- After recovery storage is reused
- Cuts recovery time by 33% average


***Log records
are read,
sorted and
applied in
parallel!***



Fast Log Apply is Always ON for DB2 Restart!



Tip #3 Manage Deferred Restart



■ Deferred Restart Parameters

□ Options

- LIMIT BACKOUT
 - YES or AUTO
- BACKOUT DURATION


These Postpone Backout Processing on DB2 Restart!



Tip #4 Eliminate Long Running UOWs



- Long Running Units of Work
 - Detecting
 - See the DB2 master log DB2 log message - *DSNR035I UNCOMMITTED UR AFTER ### CHECKPOINTS*
 - Resolving
 - Assign problem tickets
 - Catch in quality assurance
 - Preventing
 - Use Netview to automatically cancel long running UOWs



Enable the "UR CHECK FREQ" option to Detect LR UOWs!



Tip #5 System Wide Recovery



- BACKUP SYSTEM
- RESTORE SYSTEM
- Prerequisites
 - z/OS 1.5 or above
 - ESS FlashCopy API
 - SMS copy pools
 - SMS copy pool backup storage groups
- Recovery is Simple!
 - Only shallow recovery skill sets are needed
 - Does DB2 subsystem & all applications
 - In DB2 9 - Can recover individual tablespaces
 - Makes recovery simple!



This is Ideal for ERPs!

This Is Available Starting in V8, Works Even Better in DB2 9



Tip #6 Copy Frequency



- **Factors Affecting Frequency**
 1. Amount of log apply
 2. Archive log tape mount time
 3. Log read time
 4. Time to apply DO and UNDO pages from log
- Adjust to meet recovery time objective (RTO)

The More Frequent Copies, Mean Less Recovery Time!




Tip #7 Backup Strategies



- Backup Design Parameters
 - Recovery granularity
 - **R**ecovery **P**oint **O**bjective (RPO)
 - **R**ecovery **T**ime **O**bjective (RTO)
 - Indexes to image copy
 - On-site / Off-site strategy
 - Frequency of off-site transport
 - Disaster recovery strategy

***Does Your Recovery Strategy
Match Your Recovery Objectives?***



20



Tip #8 Stacking Copies on Tape



Copy Utility JCL (stack copies & retain mounts)

```
//RS54009N DD DSN=&PREFIX..TIMGCPYD.SRS9754.A009(&GDG),  
//          DISP=(NEW,CATLG,CATLG),  
//          UNIT=T3590,  
//          VOL=(,RETAIN,,25),  
//          LABEL=(001,SL,EXPDT=99000),  
//          DCB=(SYS1.GDGMODEL,TRTCH=COMP)  
//RS54P01X DD DSN=&PREFIX..TIMGCPYD.XRS9754P.A009(&GDG),  
//          DISP=(NEW,CATLG,CATLG),  
//          UNIT=AFF=RS54009N,  
//          VOL=(,RETAIN,,25,REF=* .DA97185.S010.RS54009N),  
//          LABEL=(002,SL,EXPDT=99000),  
//          DCB=(SYS1.GDGMODEL,TRTCH=COMP)  
//RS54002X DD DSN=&PREFIX..TIMGCPYD.XRS97540.A001(&GDG),  
//          DISP=(NEW,CATLG,CATLG),  
//          UNIT=AFF=RS54009N,  
//          VOL=(,RETAIN,,25,REF=* .DA97185.S010.RS54P01X),  
//          LABEL=(003,SL,EXPDT=99000),  
//          DCB=(SYS1.GDGMODEL,TRTCH=COMP)
```


- Factors
 - # tablespaces
 - number of tapes
 - # pages on tape
 - restart time
- Automated generation
- Frequency & Urgency
 - I/O characteristics
 - Minimize log apply
- Granularity



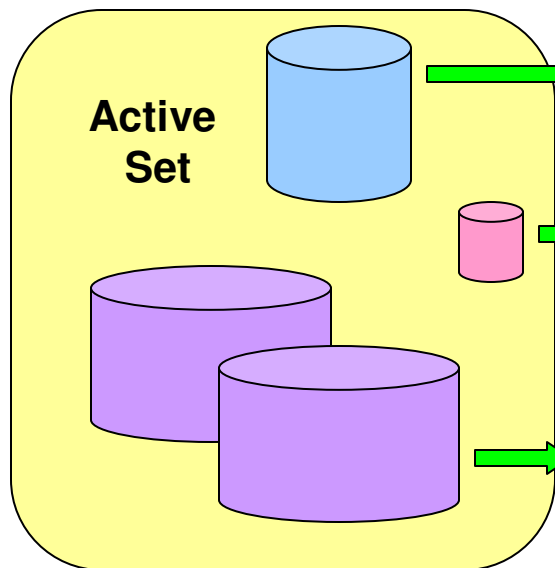
Tip #9 Copy Active Part & Indexes Together



Partitions Part Indexes



➤ Log Apply is faster because index & data updates match



Active Part

Active Part Index

Non Part Indexes

```

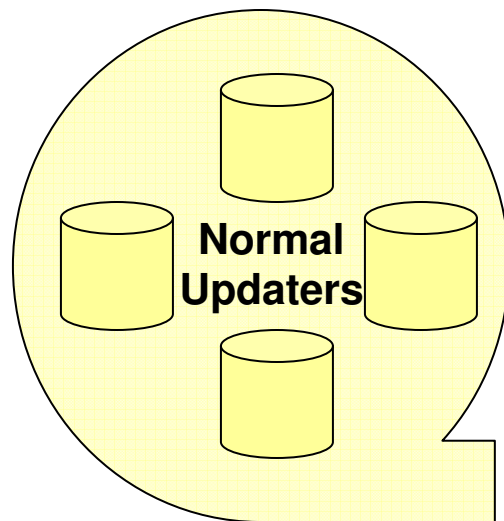
COPY TABLESPACE DRS97PRD.SRS9754
  DSNUM          009
  COPYDDN        (RS54013N)
  SHRLEVEL       CHANGE
  INDEXSPACE     DRS97PRD.XRS9754P
  DSNUM          009
  COPYDDN        (RS54P01X)
  SHRLEVEL       CHANGE
  INDEXSPACE     DRS97PRD.XRS97540
  COPYDDN        (RS54002X)
  SHRLEVEL       CHANGE
  INDEXSPACE     DRS97PRD.XRS97541
  COPYDDN        (RS54103X)
  SHRLEVEL       CHANGE
    
```



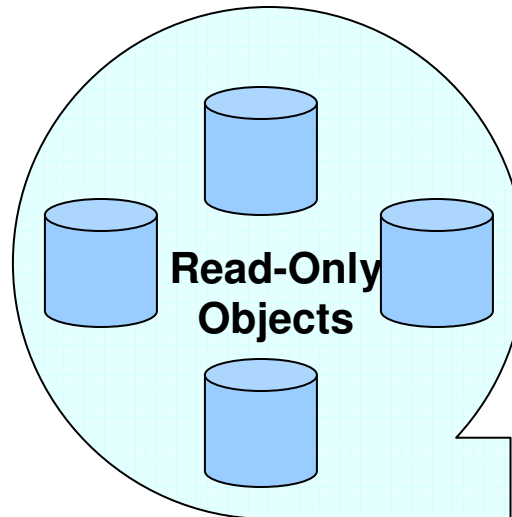


Tip #10 Group Copies By I/O Characteristics

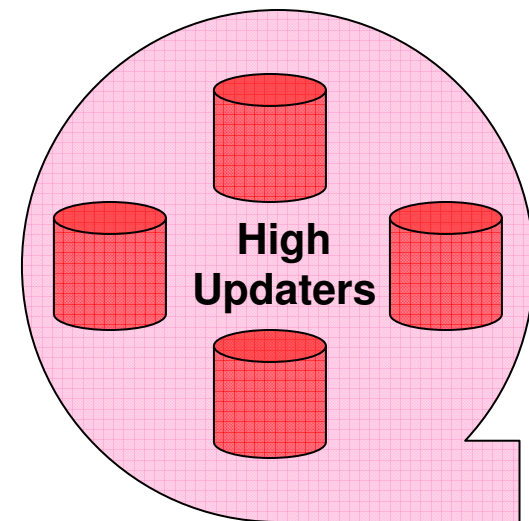
- Match I/O Characteristics in Log Apply of copy pagesets



Once A Day



Once A Week



Twice A Day



This Speeds Disaster Recovery



Tip #11 Copy to Disk Rather Than Tape



Tablespaces Indexes



It's Fast & Easy if You Can Afford the DASD!




Tip #12 Use MODIFY RECOVERY




■ Using MODIFY RECOVERY

- Deletes old entries in SYSCOPY & SYSLGRNX
- Speeds recovery process
- MODIFY RECOVERY TS automatically cleans copies & versions



SYSLGRNX



2/21/2005 Run

```
MODIFY RECOVERY
TABLESPACE DRS97PRD.SRS9754
DELETE AGE(030) DSNUM 009
```

Index Image Copies are Implicitly Pruned Until DB2 9



Tips #13 Drop Unneeded Stuff



Fewer Objects to Recover

- Tablespaces
- Indexes

Housekeeping

- Views
- Alias



Tips #14 Routine DB2 Audits



- Too many SYSADMs
- Compressed data
- Commit frequency
- RESTRICT ON DROP
- REORG & AREO*
- High table & index versions
- Extents
- Under allocations
- RESTRICT ON DROP on production tables
- LOG NO tablespaces



Routine Audits Can Prevent the Need to Recover

27



Tips #14 Routine DB2 Audits



Housekeeping



**Survey
Says**

- How many have lots of multi-dataset segmented tablespaces?
- How many have partitioned tablespaces getting close to dssize limits on partitions?
- How many have small partitioned (non-EA) tablespaces that are almost full?
- How many have audits to detect these conditions prior to death by hitting wall?
- How will you get out of these traps?

These Conditions Can Force You Into a Big Outage!

28



Tips #15 Image Copy Audits



- All tablespaces
- All partitions
- Unrecoverable tablespaces
 - ✓ Not copied since a reorg or load
- All indexes with COPY YES attribute
- Indexes without COPY YES and with copies
- Copy timeliness
- Onsite & offsite copies
- Old unusable copies

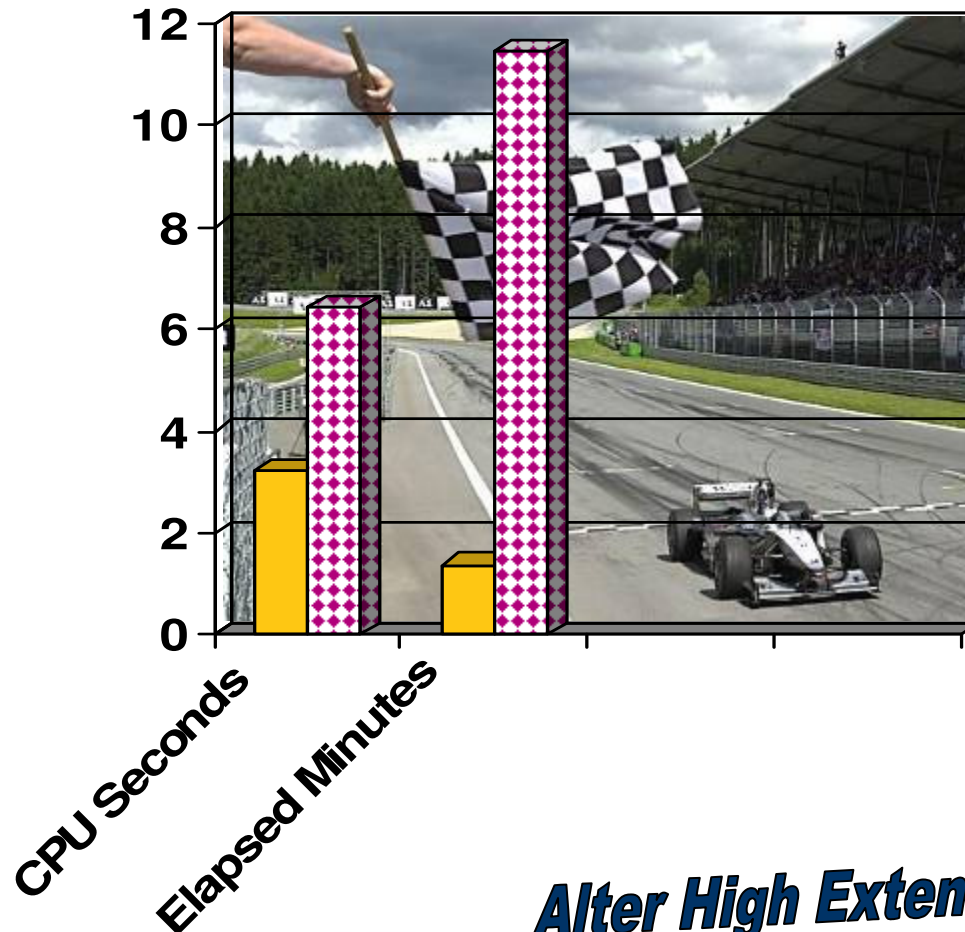
If You Don't Copy You Can't Recover!



Tips #16 Eliminate High Extents



High Extent Performance



Benchmark TOCOPY Recoveries

- 2 Extent Data & Index
- ◆ 200+ Extent Data & Index

Runs 8 Times Faster!
Uses 1/2 the CPU!

Alter High Extent Objects Prior to Recovery!



Tip #17 Minimal Recovery - LPL



- LPL Recovery
 - DB2 V7 & Earlier
 - Don't do normal recovery first
 - Do START DB to resolve LPL
 - If LPL recovery fails recover object
 - DB2 V8
 - LPL cleanup automatic
 - Message if automatic cleanup fails (DSNI005I)
 - Do START DB to resolve if automatic fails
 - Doesn't lock entire tablespace or partition



In V8 Catch LPL Recovery Failures with NetView!





Tip #18 Dynamic Allocations

```

TEMPLATE TSYSUT1 DSN(&USER..&TS.SYSUT1
  STOGROUP SCSSX
  REBUILD INDEX (PCTL.TDP97240)
  WORKDDN (TSYSUT1)

```

```

// IF RC LE 4 THEN
//RECOVIND EXEC DSNUPROC,
//*****
//* THIS STEP RECOVERS THE SDP9732 TABLESPACE
//*****
//          SYSTEM=DSN2,      * SUBSYSTEM ID
//          UID='XDP97240',    * UTILITY ID
//          UTPROC='',         * UTILITY PROC-PARMS
//*DSNUPROC.SORTWK01 DD SPACE=(CYL,(050,0100))
//*DSNUPROC.SORTWK02 DD SPACE=(CYL,(050,0100))
//*DSNUPROC.SORTWK03 DD SPACE=(CYL,(050,0100))
//*DSNUPROC.SORTWK04 DD SPACE=(CYL,(050,0100))
//*DSNUPROC.SYSUT1  DD SPACE=(CYL,(0150,0100))
//DSNUPROC.SYSIN DD *
  REBUILD INDEX (PCTL.XDP97240)
/*
// ENDIF

```

Comment
Out

Dynamic Dataset Allocations

- Omit the DDs
- Use TEMPLATES to specify dynamic dataset parameters
- Eliminates allocation errors
- Prevents over-allocation
- Only SYSIN DD required

Left Out DDs - Dynamic Alloc



Tip #19 Recover Multiple Items in 1 Pass



- Causes single pass of logs for log apply
- Especially important for disaster recovery
- Good for active partition and indexes

Recover Utility JCL

RECOVER

TABLESPACE DRS97PRD.SRS9754 DSNUM 009
 INDEXSPACE DRS97PRD.XRS9754P DSNUM 009
 INDEXSPACE DRS97PRD.SRS97540
 INDEXSPACE DRS97PRD.SRS97541

Single Log Apply

Recover Utility JCL

RECOVER

TABLESPACE DRS97PRD.SRS9754 DSNUM 009
 RECOVER
 INDEXSPACE DRS97PRD.XRS9754P DSNUM 009
 RECOVER
 INDEXSPACE DRS97PRD.SRS97540
 RECOVER
 INDEXSPACE DRS97PRD.SRS97541

Multiple Log Applies





Tip #20 Retain Tape Mounts

- Retaining Tape Mounts

- Prevents rewinds
- Retains Mounts

- Prior to DB2 V7

- You must manually code DDs

- With DB2V7 & UQ65183 **PTF**

- DB2 detects and automatically retains mounts & prevents rewinds

- With DB2V8 - not a problem

- Use PARALLEL to invoke

Prior to DB2 V7
Add DDs to
RECOVER Step
Look up &
code *vol-sers*

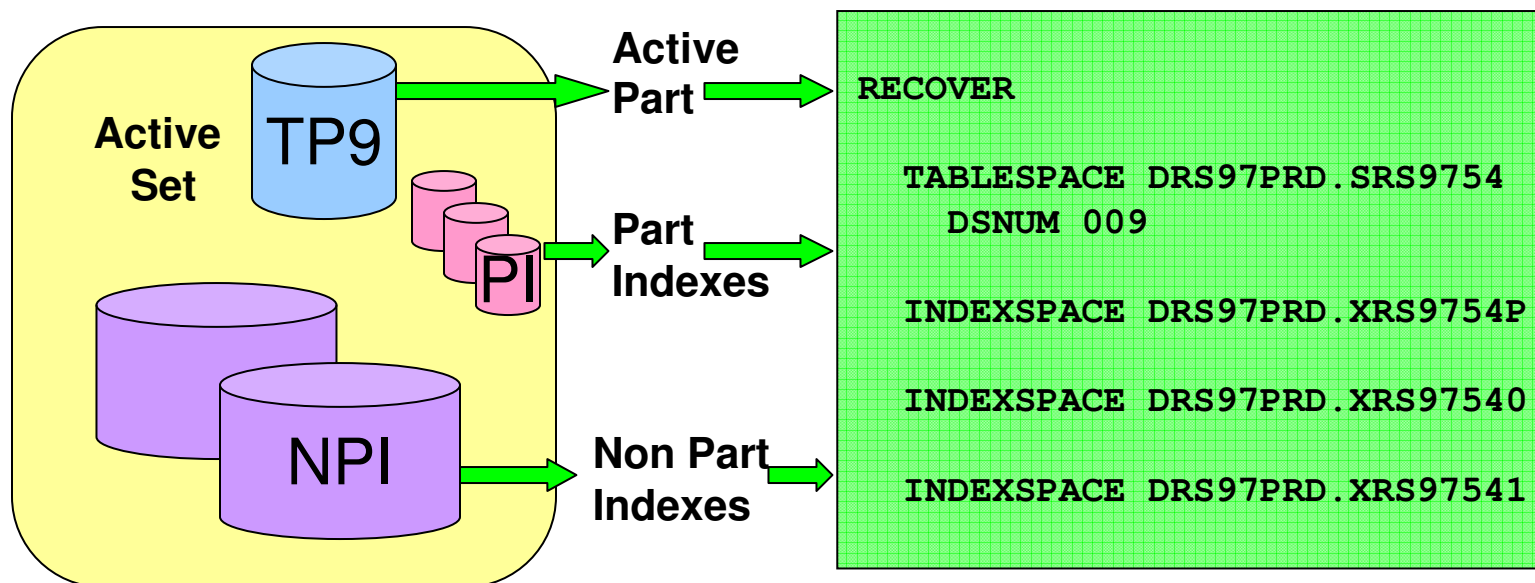
```
//ddname1 DD UNIT=TAPE,DSN=dsname1,DISP=(OLD,PASS),
//          LABEL=1,VOL=(,RETAIN,SER=vol-ser)
//ddname2 DD UNIT=TAPE,DSN=dsname2,DISP=(OLD,PASS),
//          LABEL=2,VOL=(,RETAIN,SER=vol-ser)
//ddname3 DD UNIT=TAPE,DSN=dsname3,DISP=(OLD,PASS),
//          LABEL=3,VOL=(,REF=* .ddname1)
```

If You Don't Use the Parallel Keyword, It's Not Invoked!



Tip #21 Recover Part & Indexes Together

- RECOVER Active Part & Indexes Together
 - Log Apply is similar for active sets

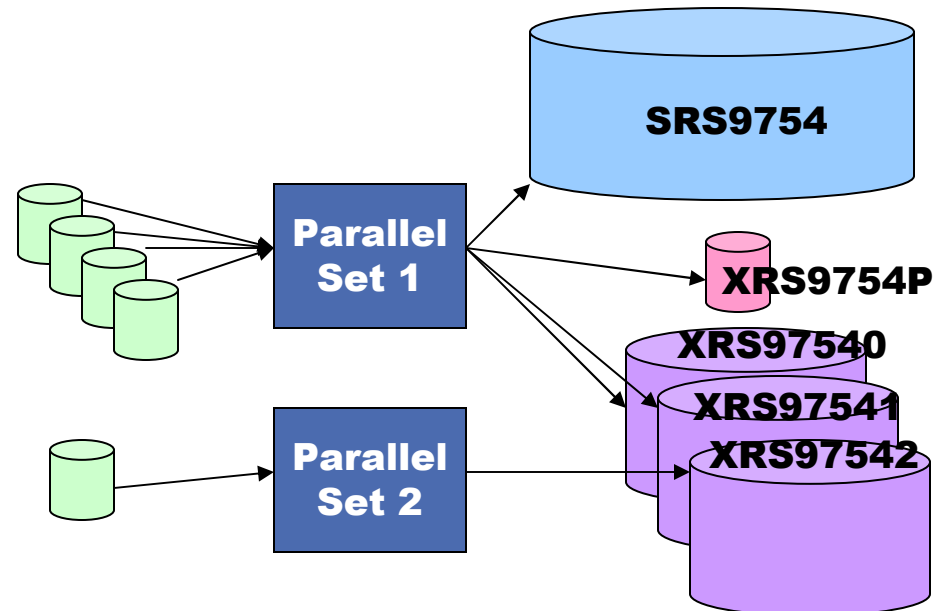


Tip #22 Recover Indexes in Parallel



Parallel Index RECOVER Option

- Image Copies must be on disk or tape
- Attempts to retain tape mounts when option used
- Lays down tablespaces in parallel
- Log Apply happens after all tablespaces are recovered
- Omit NUM-OBJECTs to let system determine # units



RECOVER **PARALLEL (4)**

```


TABLESPACE DRS97PRD.SRS9754
INDEXSPACE DRS97PRD.XRS9714P
INDEXSPACE DRS97PRD.XRS97140
INDEXSPACE DRS97PRD.XRS97141
INDEXSPACE DRS97PRD.XRS97142
  
```

Specifying Parallel(0)

Lets DB2 Figure Optimal # Objects!

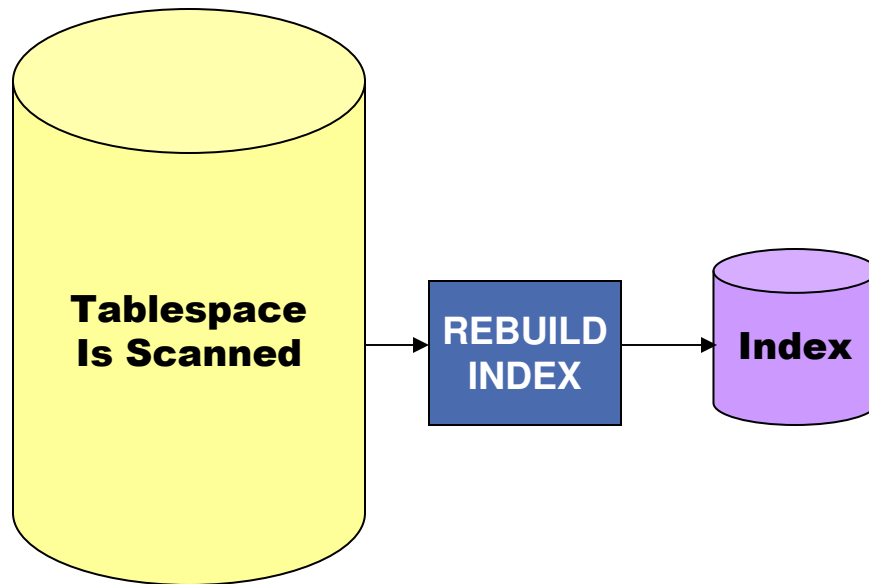


Tip #23 Rebuild vs. Recover Index



- **REBUILD INDEX**

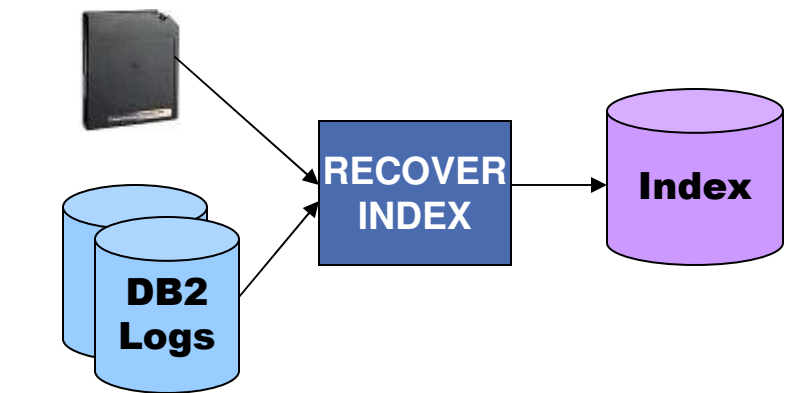
- Reads data rebuilds index




- **RECOVER INDEX**

- ALTER INDEX to **COPY**
YES to activate
- Image Copy index

Image Copy



Tip #23 Rebuild vs. Recover Index



■ When To Use REBUILD INDEX

- Small, medium & large tablespaces & indexes
- Large tablespaces & indexes with updates
- All indexes not Image Copied
- In most cases REBUILD faster than RECOVER


■ When To Use RECOVER INDEX

- Large R/O tablespaces & indexes
- Large periodic updated tablespaces & indexes

**Index Size & Log Apply
Are Critical Factors**

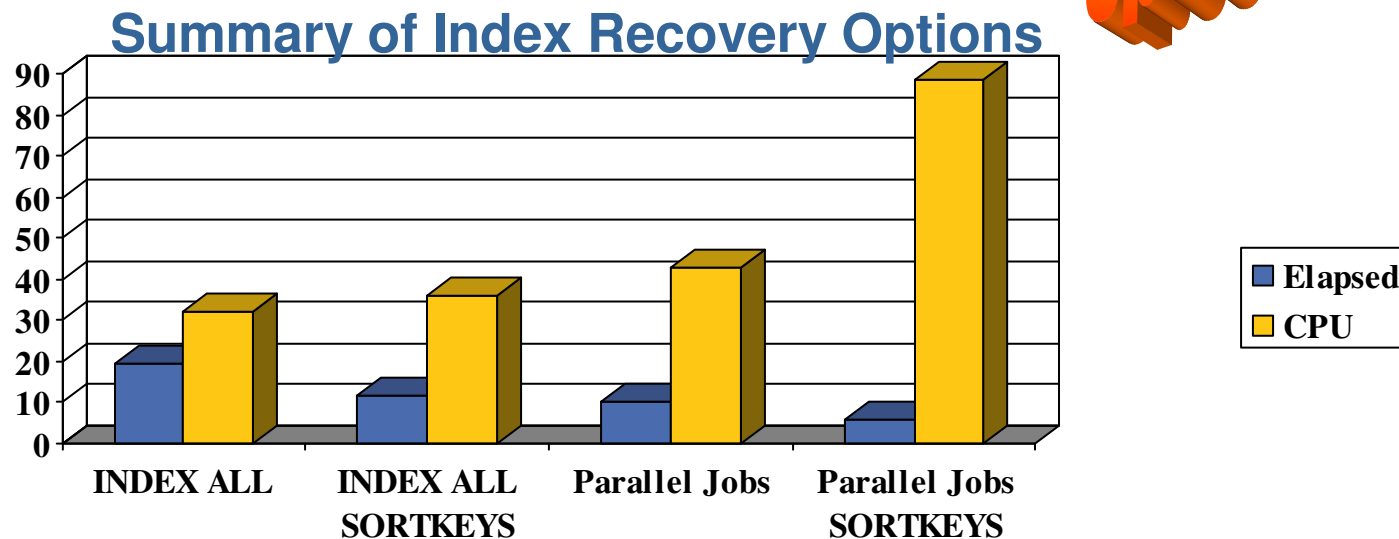


Tip #24 Use the SORTKEYS Option




- **SORTKEYS Option**
 - Use when elapsed time is critical
 - Sorts indexes in parallel disables SYSUTs
 - Dramatically Increases CPU
 - Default in V8

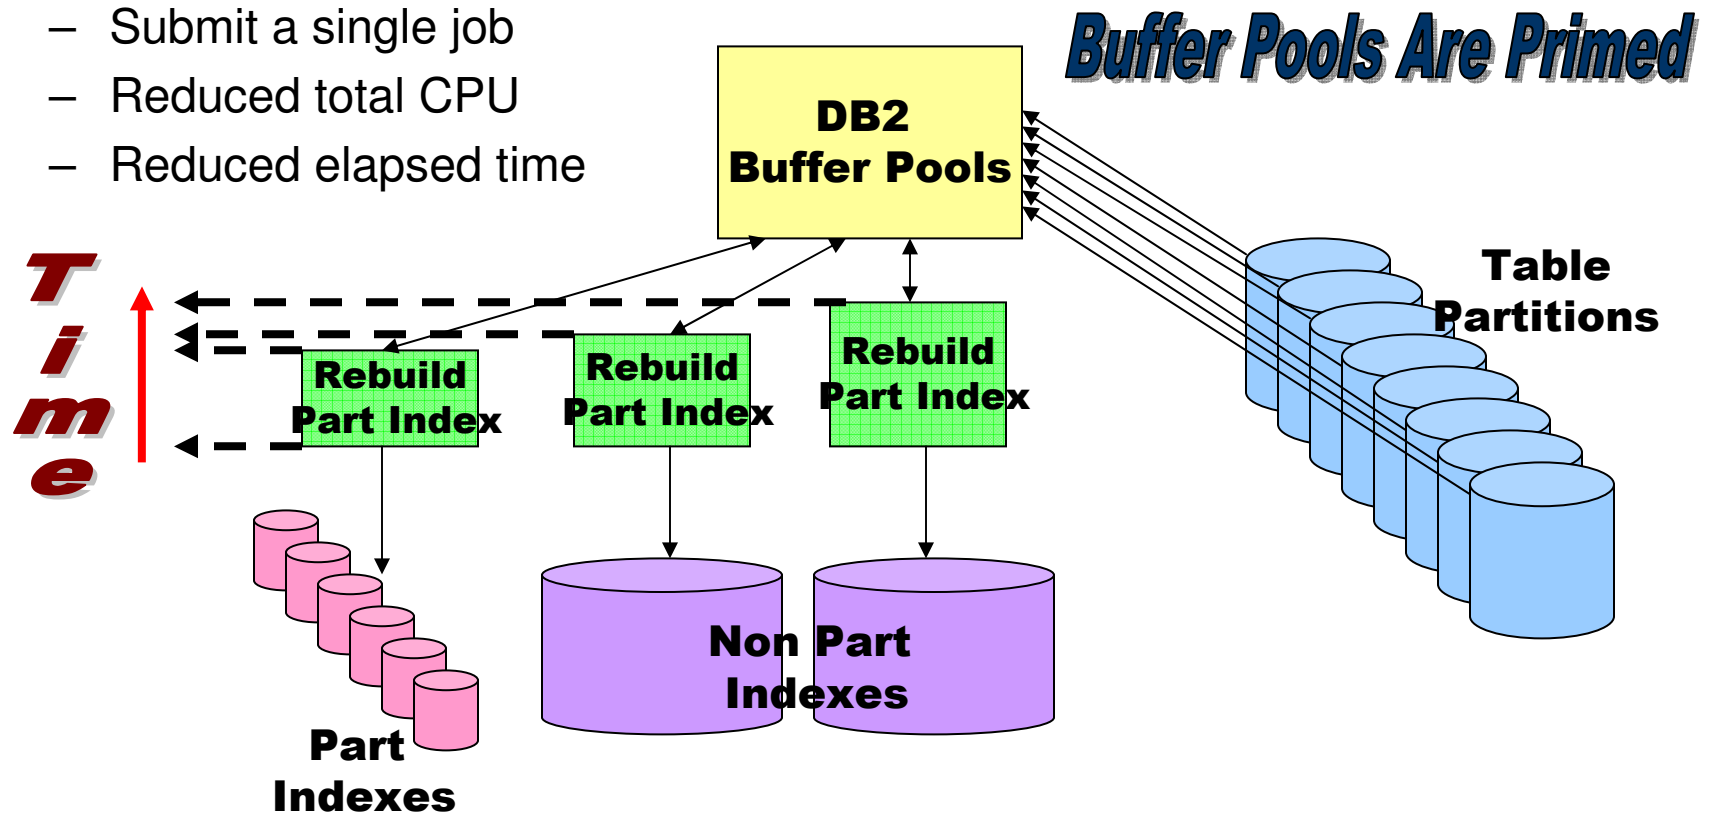
Speed vs. CPU




Tip #25 Rebuild Index All



- **SORTKEYS INDEX(ALL) to Build 2+ Indexes**
 - SORTKEYS with INDEX(ALL) faster doing many indexes
 - Submit a single job
 - Reduced total CPU
 - Reduced elapsed time

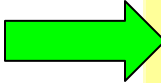


Tip #26 Omit SYSUT1 DDs



■ Omitting SYSUT1 DD

- Omit SYSUT1 from JCL
- REBUILD INDEX is not restartable if omit SYSUT1
- SORTKEYS option disables SYSUT1
- Eliminates intermediate work dataset



```
// IF RC LE 4 THEN
//RECOVIND EXEC DSNUPROC,
//*****
//* THIS STEP RECOVERS THE SDP9732 TABLESPACE
//*****
//                SYSTEM=DSN2,          * SUBSYSTEM ID
//                UID='XDP97240',        * UTILITY ID
//                UTPROC='',              * UTILITY PROC-PARMS
//DSNUPROC.SORTWK01 DD SPACE=(CYL,(050,0100))
//DSNUPROC.SORTWK02 DD SPACE=(CYL,(050,0100))
//DSNUPROC.SORTWK03 DD SPACE=(CYL,(050,0100))
//DSNUPROC.SORTWK04 DD SPACE=(CYL,(050,0100))
//*DSNUPROC.SYSUT1  DD SPACE=(CYL,(0150,0100))
//DSNUPROC.SYSIN DD *
//                REBUILD INDEX (PCTL.XDP97240)
/*
// ENDIF
```

In V8, Omitting SYSUT1 DDs is the Default!

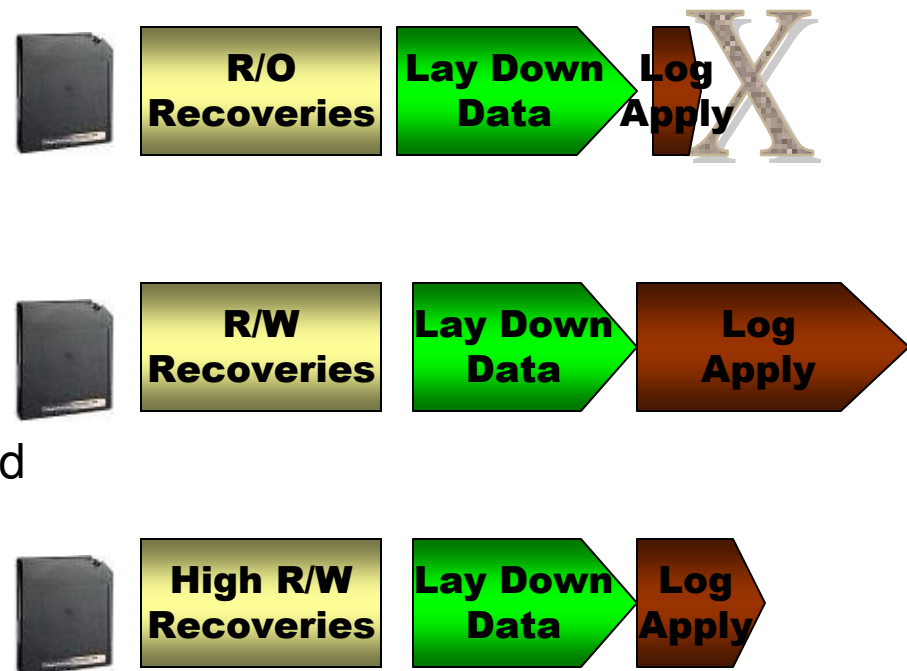



Tip #27 Balance By I/O Characteristics



■ Balance Copies By I/O Characteristics

- Longest Log Apply affects all objects in the recovery set
- Matching sets reduces elapsed recovery times
- R/W intensive objects are grouped for Log Apply





Tip #28 Write Archive Logs to Disk



Write Archive Tape Logs to Disk to Avoid Reading Tape Archives!




Tip #29 Run Recoveries 10 Wide






- Fast Log Apply Must Be 100 mb
- 10 recoveries in log apply concurrently can exploit this
- This will speed log apply by 33%



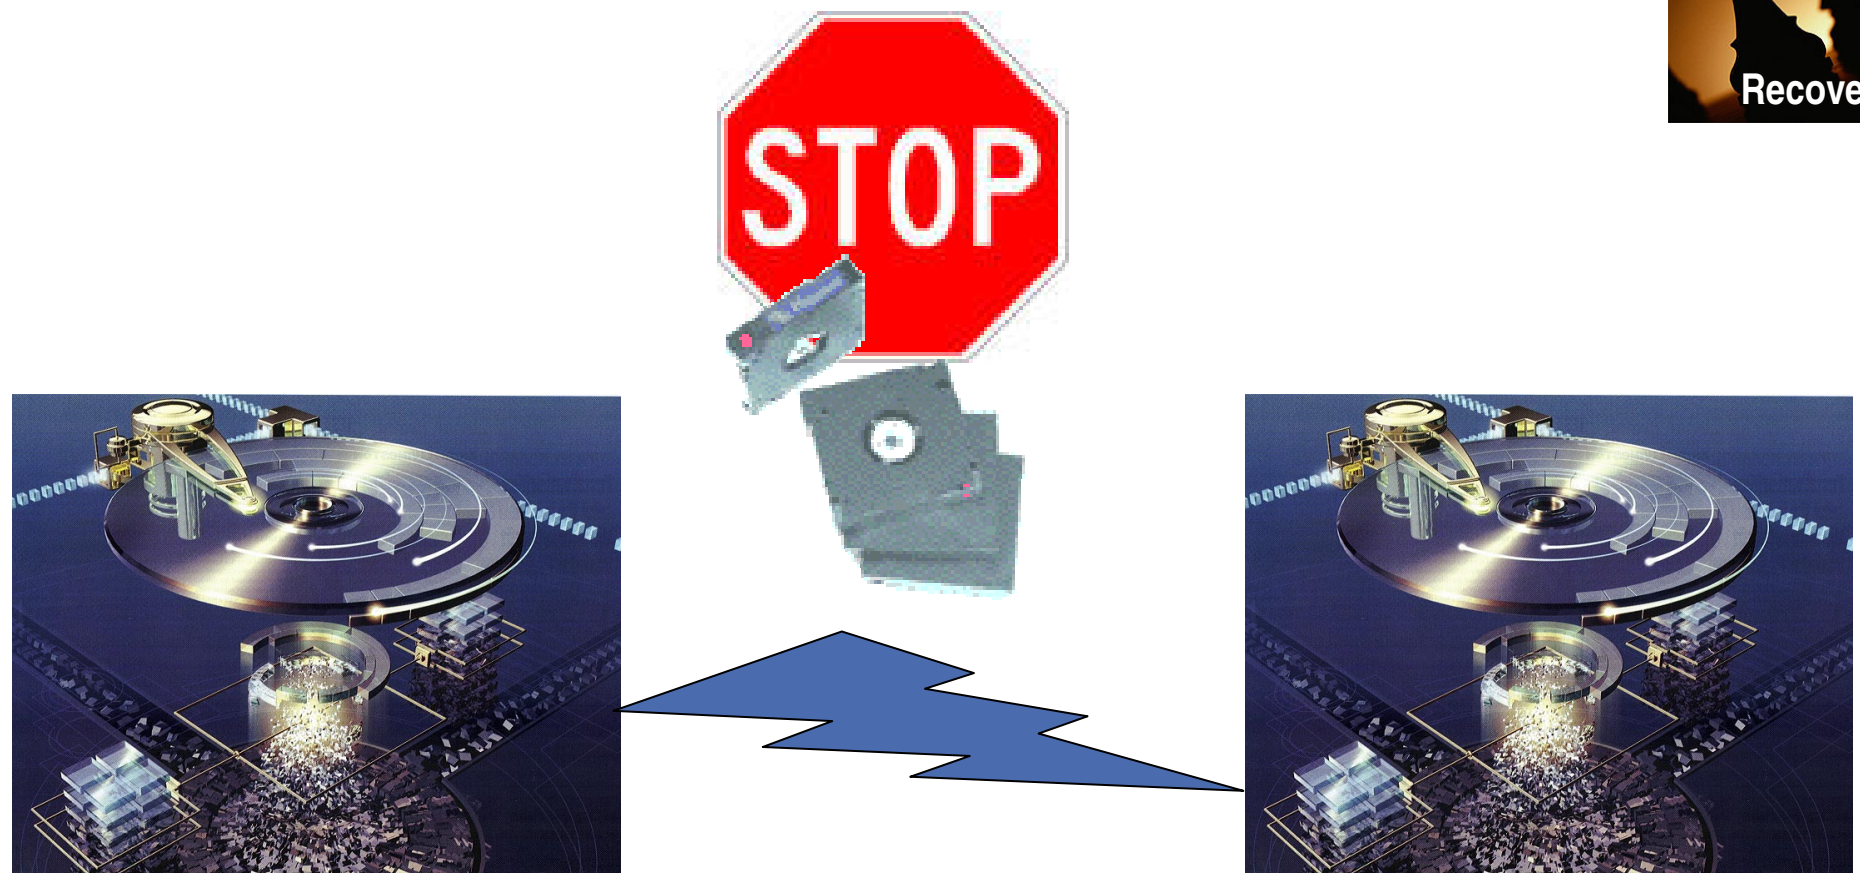


Tip #30 RECOVER All then REBUILD





Tip #31 Remote Mirror - Simplify DR



Continuous Data Protection Simplifies Disaster Recovery



Things That Make a BIG Difference!

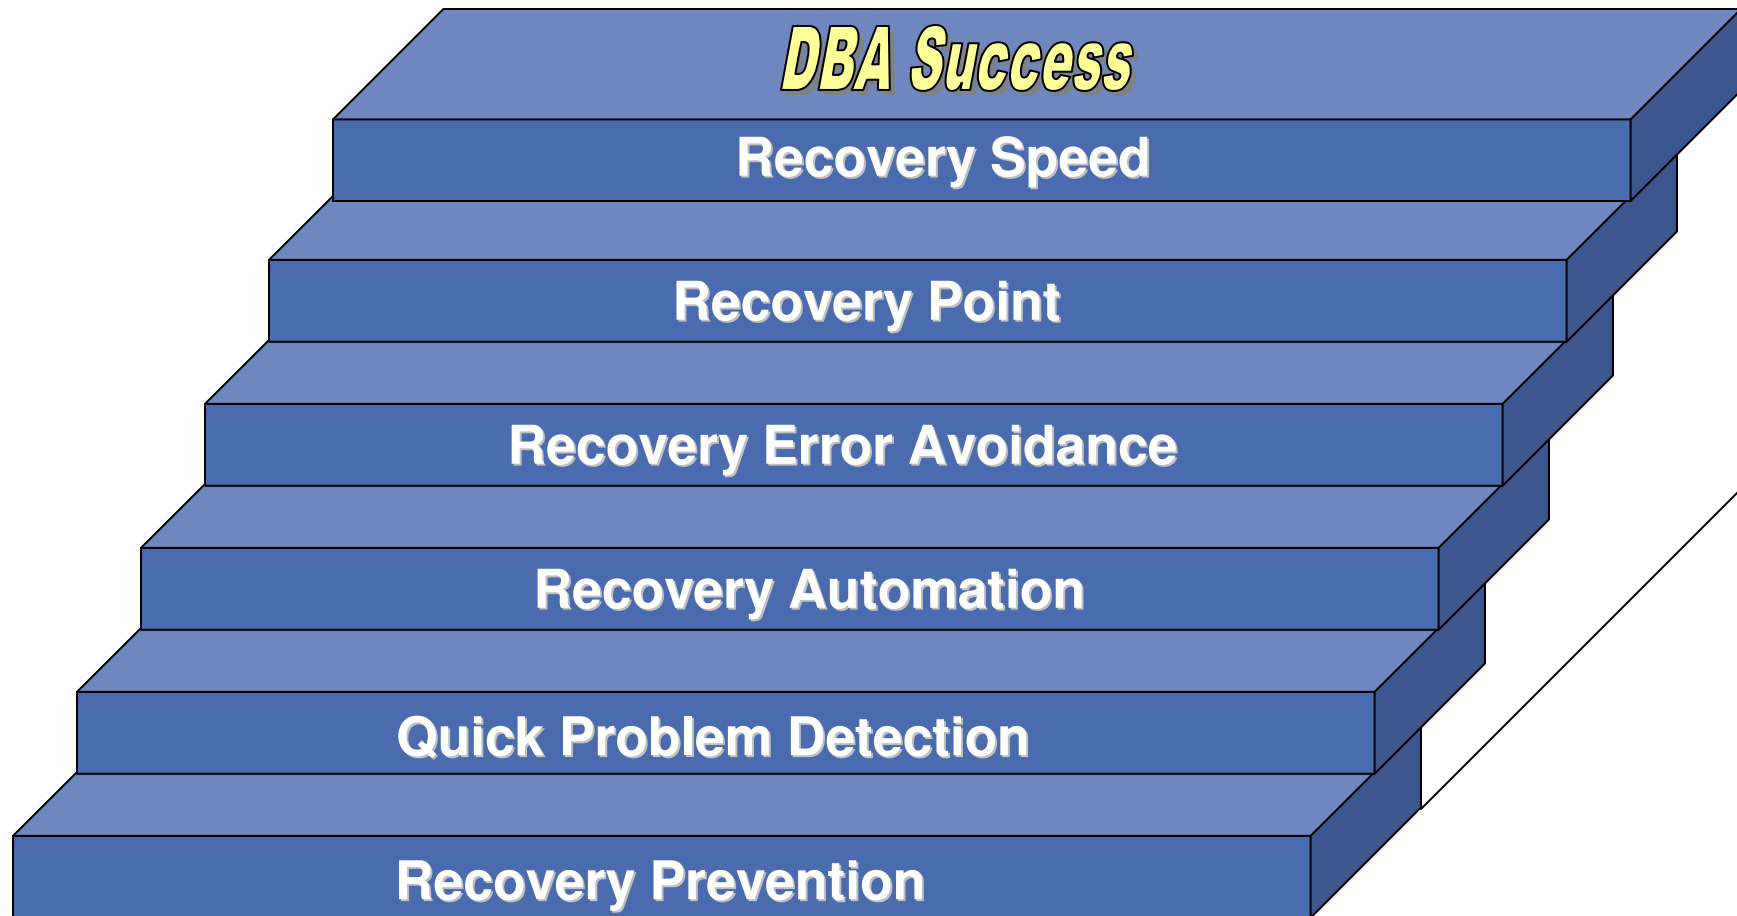


- **DB2 Subsystem**
 1. Properly Sized DB2 Logs
 2. Enable Fast Log Apply
- **Image Copies**
 1. Stacking Copies on Tape
 2. Copy active partition and indexes together
 3. Group Image Copies By Update Characteristics
- **Housekeeping**
 1. Using MODIFY RECOVER
 2. Routine Audits
- **RECOVER Utility**
 1. Dynamic Dataset Allocation
 2. Multiple Objects per RECOVER statement
- **RECOVER Utility (cont.)**
 3. Retaining Tape Mounts
 4. Recovering Active Part & NP Indexes Together
 5. Parallel RECOVER
- **Index REBUILD/RECOVER**
 1. REBUILD vs. RECOVER
 2. REBUILD INDEX(ALL)
 3. SORTKEYS Option
 4. Omitting SYSUT DDs
- **Disaster Recovery Strategies**
 1. Balance copies by write characteristics
 2. Put archive logs on disk
 3. Run recoveries 10 wide



A Holistic Approach to Recovery

Educate, Practice, Automate & Execute!





Questions?



IBM INFORMATION ON DEMAND 2007

Act Right. Now.