
LUT
University
LAND
OF THE
CURIOUS

TABLE OF CONTENTS

- Database backup
- Database recovery
- Disaster plan

CT60A7650 – DATABASE SYSTEMS MANAGEMENT

DATABASE BACKUP

Lecture

Jiri Musto, D.Sc.

PREPARING FOR PROBLEMS

- Instance failures
  - DBMS or operating system failures, software-related database failure
- Application or transaction failures
  - Scripts or programs using wrong input, run in the wrong order
- Media failures
  - Disk storage and file system failures, deleted data files, damaged hardware
- Recoverability should be one of the top priorities
  - Fast access makes no difference if you cannot recover a database

BACKUP PLAN

- Backup plan in case of a failure
- How to backup
  - Backup type, incremental, full?
- When to backup
  - What time, how often?
- What to backup
  - All, partial?
- How to recover
  - How many backups stored? What backup to use?

IMAGE COPIES (PHYSICAL BACKUP)

- Backup copy of data
- Can be used as a basis to recover or restore the database
- DBMS offer image copies as the go-to backup method
  - BACKUP, COPY, DUMP, EXPORT (depends on the DBMS)
- Can use tools outside of the DBMS
- Can have multiple generations of backups
  - Depends on the backup plan

GENERAL IMAGE COPY GUIDELINES

- Have at least two local copies of each image copy in the case of media failure
- Keep at least two generations of backups
- Consider creating backups on disk or external media
- Ensure the backup process can be restarted
- Use DBMS facilities to verify the correctness of the backup
- Data not stored in the database but used by database applications should be backed up
- May need to re-create indices

FULL VS. INCREMENTAL BACKUP

- Full image copy backup copies all data in the database object at the time of copy
- Incremental image copy backup (or differential backup) copies only the data that was changed since the last image copy backup
  - Incremental bacup is quicker and requires less space
  - Recovering from incremental backup may take longer
- Favor full backups with small database objects
- Some scenarios are not compatible with incremental backups

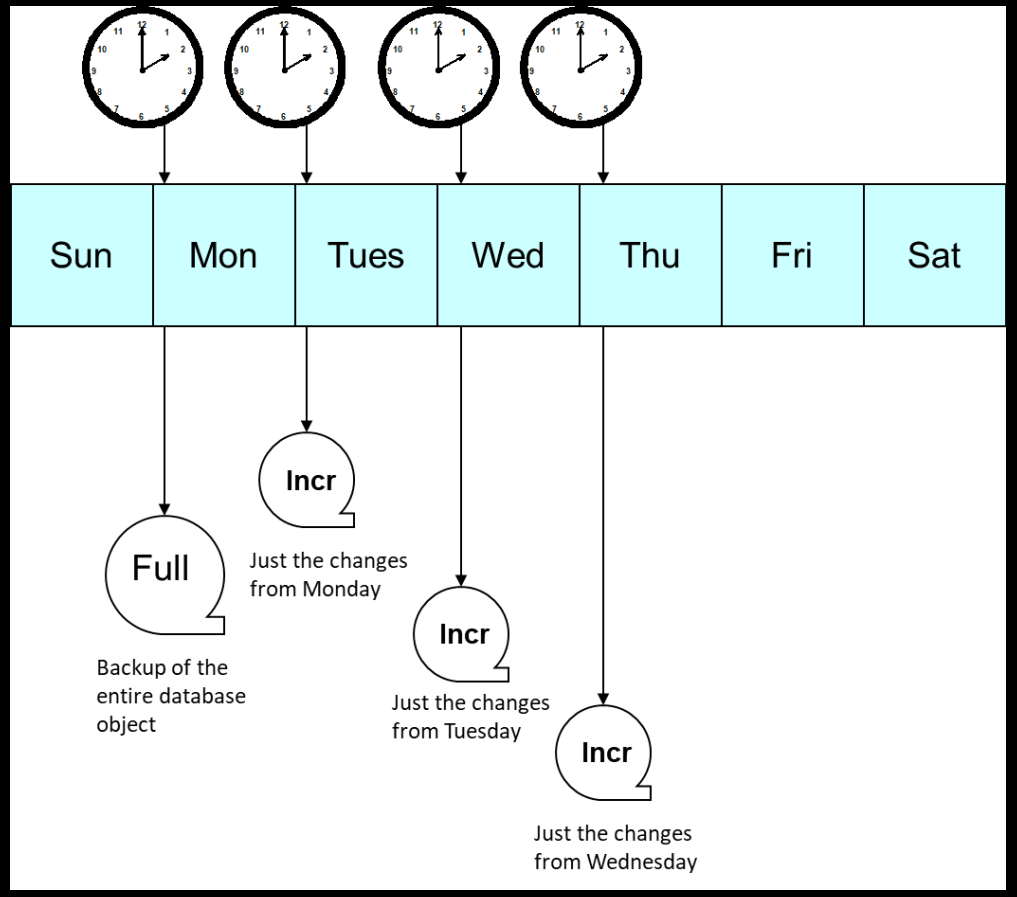
Sun
Mon
Tues
Wed
Thu
Fri
Sat
Full
Backup of the entire database object
Incr
Just the changes from Monday
Incr
Just the changes from Tuesday
Incr
Just the changes from Wednesday

MERGING INCREMENTAL IMAGE COPIES

- Some DBMS support merging incremental copies
- Combine multiple backups into one
  - Multiple incremental backups
  - One full with incremental backup
- Consider running after incremental backup if possible

BACKUP PLANNING CONSIDERATIONS

- Need for concurrent access
  - Allows to keep data online during backup but slows down recovery and the database
- Amount of time available for the backup process
- Speed of recovery utilities
- Need to access database logs
- Difference between ***hot*** and ***cold backup***

HOT VS. COLD BACKUP

- Cold backup
  - Database is shut down
  - Database is inaccessible during backup
  - The “easier” way
- Hot backup
  - Database remains online
  - Concurrent access is possible
  - Requires extensive testing
  - Can be complex to implement

BACKUP CONSISTENCY

- Create a consistent recovery point
- Relationships between database objects and other objects
  - Application relationships
  - Referential constraints
  - Triggers
- If you recover a database object to a previous point in time, you need to recover related objects as well

LOG ARCHIVING AND BACKUP

- Database changes are stored (logged) into a log file
- Log currently in use is an *active log*
- When log is full, the log can be archived
  - Current information is moved offline to an archived file
  - Active log is reset
- Frequency of log archival can be typically controlled in the DBMS configurations

DBMS INSTANCE BACKUP

- In addition to database object failure, the entire DBMS instance can fail
- Backup crucial components
  - DBMS files
  - System catalog
  - Logs
  - Configuration and setup files
  - Libraries (system, program source, executable)
- Each DBMS have different key components

LOGICAL BACKUP

- Alternative to a physical backup
- Faster and simpler method
- Creates SQL commands from the existing database
- Useful for migrating between versions
- Cannot be used for point-in-time recovery

CT60A7650 – DATABASE SYSTEMS MANAGEMENT

DATABASE RECOVERY

Lecture

Jiri Musto, D.Sc.

RECOVERY

- Recovery can be a complex task
- Involves more than just restoring image of the data
- Bring back the data to its state at (or before) the time of the problem
  - Restore databases and then reapply changes that occurred
- In a successful recovery, data will be in the state you want it to be
  - With good planning, you should be able to recover from any type of failure

DETERMINING RECOVERY OPTIONS

- Type of failure?
- Cause of failure?
- Abort, crash, shutdown?
- Operating system errors?
- Server reboot?
- How critical is lost data?
- Existing backups?
- What needs to be recovered?
- Backup strategy?
- How much data needs to be recovered?

GENERAL STEPS FOR RECOVERY

1. Identify the failure
2. Analyze the situation
3. Determine what needs to be recovered
4. Identify dependencies between the database objects to be recovered
5. Locate the required image copy backup(s)
6. Restore the image copy backup(s)
7. Roll forward through the database log(s)

TYPES OF RECOVERY

- Recovery to current
  - Recover to at or just before the point of failure
  - Usually done after a disaster (media failure for example)
- Point-in-time recovery (or partial recovery)
  - Recover to a specific time
  - Deal with an application-level problem
- Transaction recovery
  - Recover a specific portion of the database
  - “Rollback” changes made to data
  - May require third-party software

RECOVERY TO CURRENT

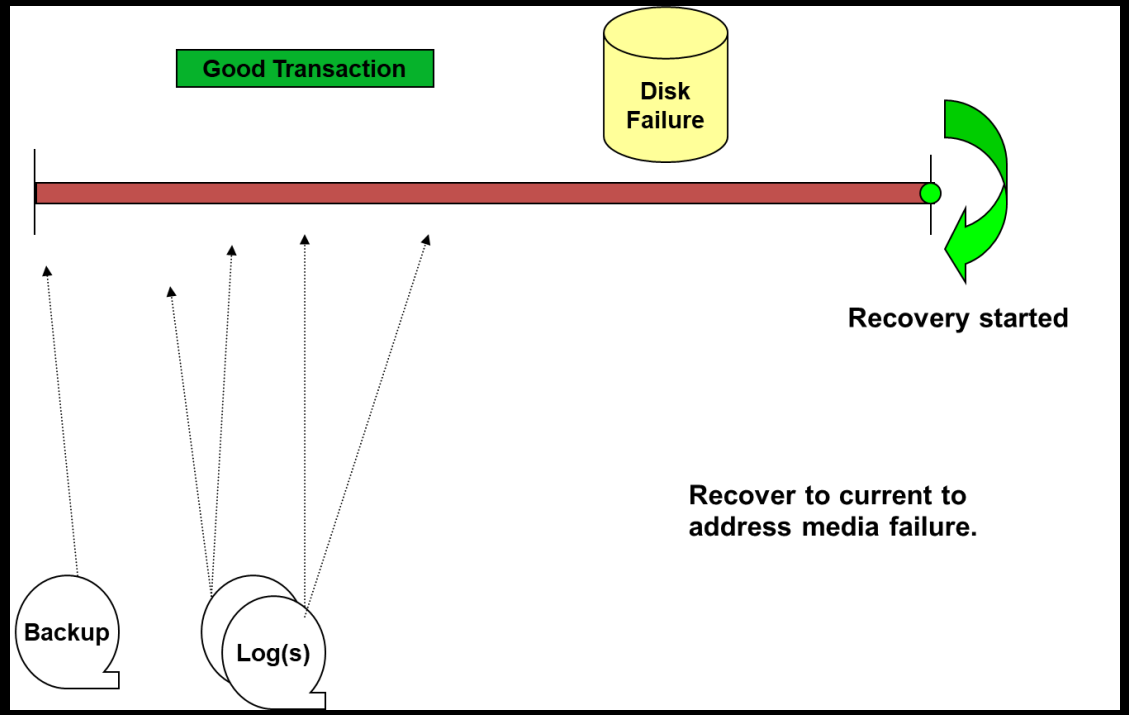

POINT-IN-TIME RECOVERY

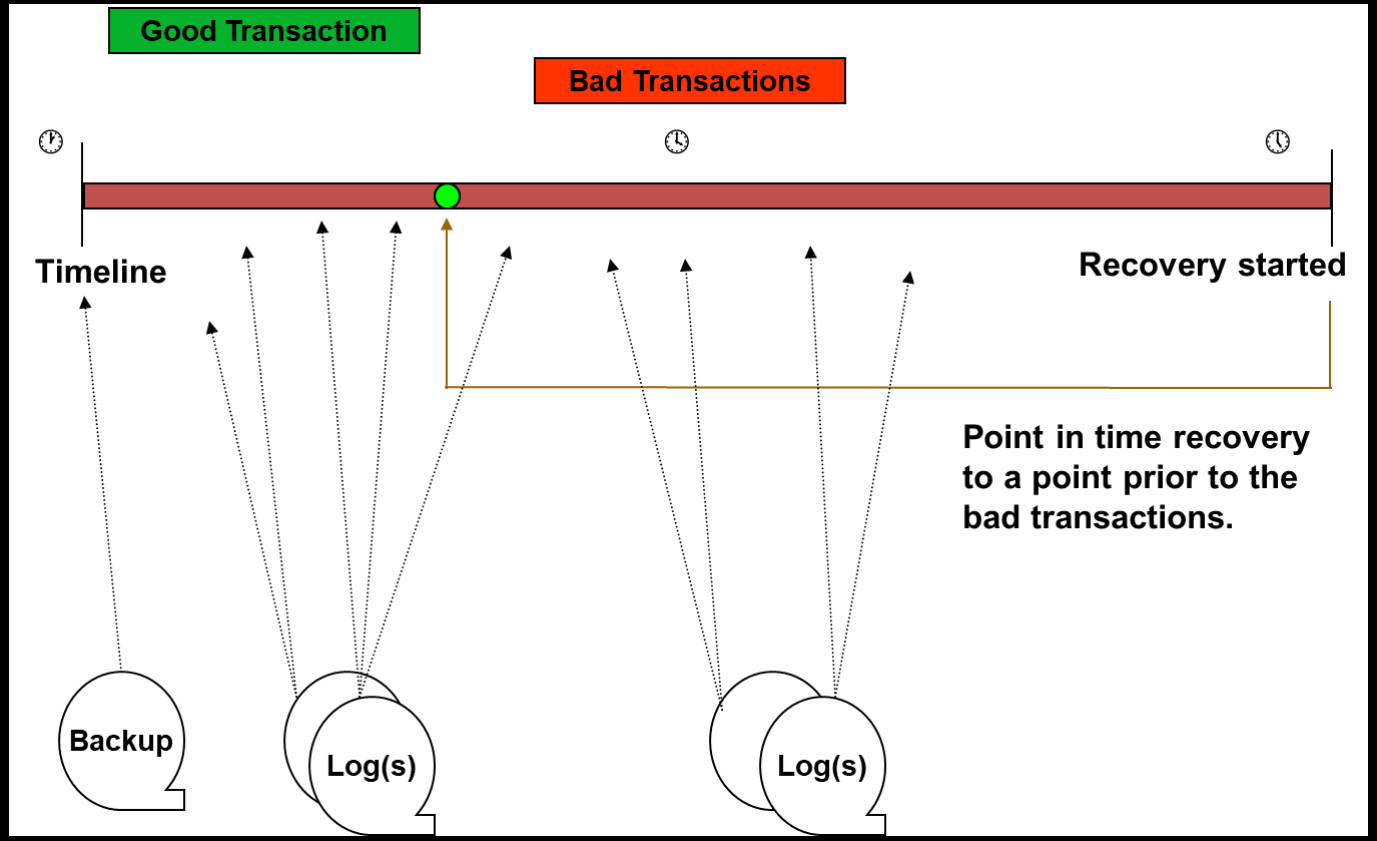

TRANSACTION RECOVERY

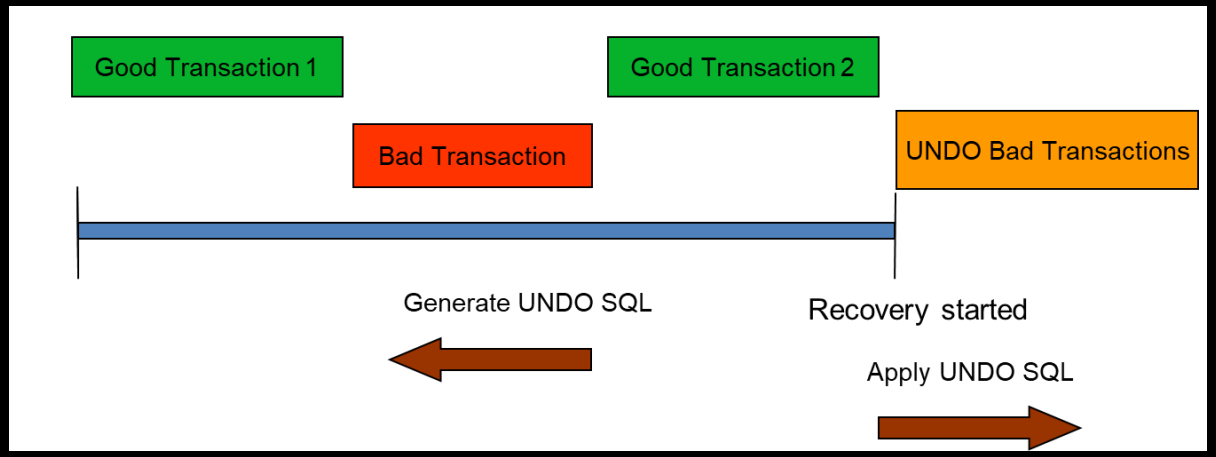

UNDO Transaction recovery

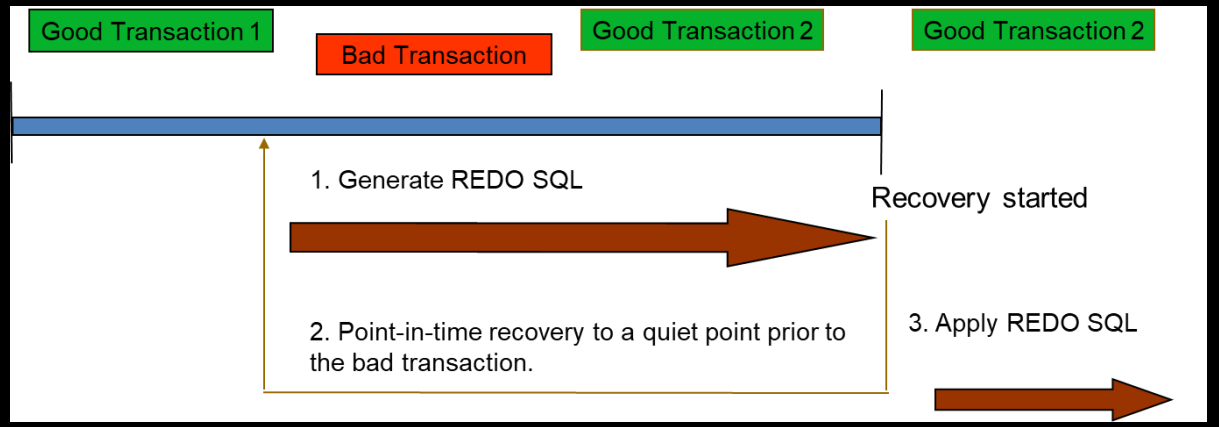

REDO Transaction recovery

CHOOSING THE RECOVERY STRATEGY

- Transaction identification
  - Can all problem transactions be identified?
- Data integrity
  - Has updates been done since the problem occurred? Is all required data available? Will data be lost?
- Speed
  - If multiple options are viable, which is fastest?
- Availability
  - How soon the applications becomes available? Can you go offline?
- Invasiveness
  - How invasive the failure was? Were decisions made based on bad data? Can subsequent work be trusted?

FAILURE AND RECOVERY

- Media failure?
  - Recover to current
- Transaction failure?
  - Point-in-time or transaction recovery
- Database or subsystem failure?
  - Recover to current

TESTING RECOVERY PLAN

- Develop a recovery plan
- Test the plan at least twice per year
- Recovery plan:
  - Write aspects of the recovery plan in detail, document each step
  - Include all scripts required to back up and recover
  - Review the plan with people who may need to implement it
  - Include contact list
  - Keep the plan up-to-date

ALTERNATIVES TO BACKUP & RECOVERY

- Standby databases
  - Identical copy of the active database, close to being up-to-date
  - In the event of a failure, control is transferred to the standby database
- Replication
  - Maintain data in a separate copy of the database
  - Can be a subset of data
- Disk mirroring
  - Allocate a secondary device that contains a duplicate copy
  - If primary device fails, secondary device can be used

CT60A7650 – DATABASE SYSTEMS MANAGEMENT

DISASTER PLANNING

Lecture

Jiri Musto, D.Sc.

WHAT IS A DISASTER

- “Unplanned, extended loss of critical business applications due to lack of computer processing capabilities for more than 48-hour period”

- “Any event that has a small chance of transpiring, a high level of uncertainty, and a potentially devastating outcome”

NEED FOR PLANNING

- Disaster does not need to have global consequences to be a disaster for you
- How a disaster might impact you is the sole purpose of disaster recovery planning
- Recognize likely situations
  - Likelihood of floods, hurricanes and tornadoes next to a coast
  - Blizzards and severe cold weather
  - Earthquakes
  - Storms
- Even if they are unlikely, you should still plan for them

DICTATE PRIORITIES

- Very critical applications
  - Require current data upon recovery
  - Try to limit the number to a handful
- Business-critical applications
  - Applications important to the organization
- Critical applications
  - May not need to be available immediately compared to the previous
- Required applications
  - Noncritical but need to be backed up so they can be recovered if needed
- Noncritical applications
  - Do not need support in the event of a disaster
  - Very few applications fall into this category

GENERAL DISASTER RECOVERY GUIDELINES

- Minimize downtime and loss of data
- Planning for a disaster is an enterprise-wide task
- DBMS and database recovery is just one component
- Organization needs to look at all of the business functions and activities
  - Customer interfaces
  - Phone centers
  - Networks
  - Applications

THE REMOTE SITE

- Off-site location to setup operations
- Far enough from the primary site to not be impacted by the disaster
  - May need multiple remote sites
- For example
  - Dual data centers
  - Backup data center
  - Recovery service provider

WRITTEN PLAN

- Disaster recovery plan needs to be in writing
  - Explicit actions taken in the event of a disaster
  - The complete step-by-step procedures for the recovery of each piece of system software, every application, and every database object, and the order in which they should be restored
  - Provides blueprint for others to follow
- Distribute to all key personnel
- Needs to be updated when business and IT environment changes
- Destroy all outdated copies of the plan and replace them with the new plan

TESTING THE DISASTER PLAN

- Disaster recovery test can discover weaknesses and errors in the plan
- Test does not need to end in a successful recover (though desired)
- Should test the plan if:
  - System hardware changes
  - DBMS is upgraded or changed
  - Backup procedures are changed
  - Primary data center is relocated
  - Major new applications or significant upgrades of existing critical applications

BACKING UP FOR A DISASTER

- External media can be sent to remote sites
  - Relatively “safe” method
  - Includes a lot of work (need to store database logs between backups)
- Storage management system backups
  - System-wide point-in-time backup
  - Simplifies recovery preparation and execution
  - May require a longer time for the database to be offline
- Remote standby databases and mirroring
  - Requires different locations physically
  - Simple method as long as there is a constant connection to the remote sites

GUIDELINES

- Establish procedures and policies to prevent problems
  - Natural disasters cannot be prevented but man-made disasters can
  - Surge protectors, backup generators, etc.
- Adhere to the written plan
- Pay attention to the order of recovery
- Remember vital data
- Post-recovery image copies


LUT
University