
ORACLE®

ORACLE OPEN WORLD

experience
OPENWORLD

November 11–15, 2007

ORACLE

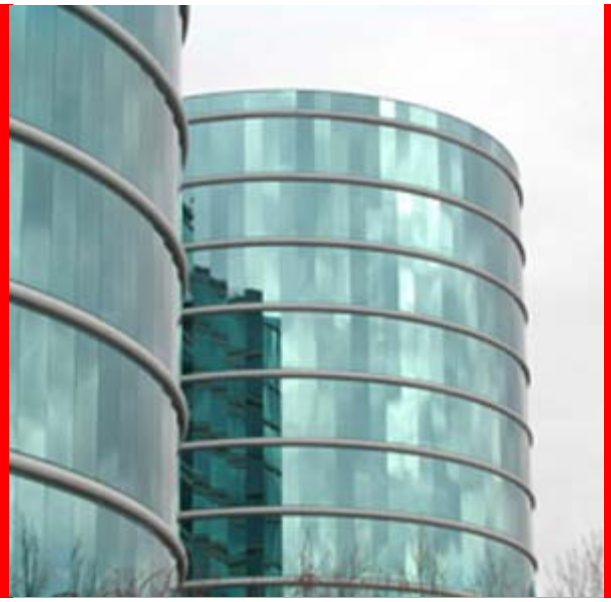

ORACLE

Backup and Recovery Best Practices for Very Large Databases (VLDBs)

Tim Chien
Senior Product Manager
Oracle

Greg Green
Senior Database Administrator
Starbucks

Agenda

- VLDB Trends and Requirements
- Backup and Recovery Design
  - Assess Recovery Requirements
- Backup and Recovery Deployment
  - Architect Backup Environment
  - Evaluate Backup and Recovery Tools
  - Plan Data Layout
  - Develop Backup Procedure
  - Develop Recovery Strategies
- Starbucks Case Study
- Summary/Q&A

VLDB Trends & Requirements

- Rapidly growing databases
  - 1 TB+ is ‘normal’ for initial deployments => 100 TB++
  - Growth of intranet and internet resources
- Data consolidation
  - Global view of the business
  - Integrated, targeted analysis critical to business decisions
- Concurrent, heavy access
  - Global user base – internal/partner/customer

Need to be cost-effective while scaling for data growth

VLDB Backup & Recovery

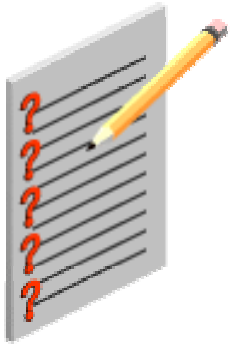

- Key database protection requirements:
  - Growing data – how to backup/restore in tolerable window?
  - Consolidated data – how to recover at an application level? How to protect data in 24x7 environment?
  - Cost efficiency with scale – how to manage software, infrastructure, operational costs for data protection/disaster recovery?

Agenda

- VLDB Trends and Requirements
- Backup and Recovery Design
  - Assess Recovery Requirements
- Backup and Recovery Deployment
  - Architect Backup Environment
  - Evaluate Backup and Recovery Tools
  - Plan Data Layout
  - Develop Backup Procedure
  - Develop Recovery Strategies
- Starbucks Case Study
- Summary/Q&A

Backup and Recovery Design

Assess Recovery Requirements

- Analyze and identify the cost of unavailable data
- Design recovery requirements around data criticality and logical relationships
  - Assess tolerance for data loss - Recovery Point Objective (RPO)
    - Prioritize data
    - Point-in-time recovery required?
  - Assess total recovery time - Recovery Time Objective (RTO)
    - Problem identification & recovery planning time + H/W & S/W recovery time + database recovery time (media + instance recovery)
  - Determine backup retention policy

Agenda

- VLDB Trends and Requirements
- Backup and Recovery Design
  - Assess Recovery Requirements
- Backup and Recovery Deployment
  - Architect Backup Environment
  - Evaluate Backup and Recovery Tools
  - Plan Data Layout
  - Develop Backup Procedure
  - Develop Recovery Strategies
- Starbucks Case Study
- Summary/Q&A

Backup and Recovery Design

Architect Backup Environment

	Disk	Tape Library	Virtual Tape Library (VTL)
Compatibility, Provisioning	OS, drivers, storage must be compatible Disk should be provisioned for 'worst' case scenario	Well-known interface across heterogeneous systems Tapes easily added, as needed	Standard tape interface
Performance	Fast, random I/O access	Sequential-only access	Fast, random I/O access
Disaster Recovery	Optional I/O-based mirroring	Offsite tapes to DR site, long-term archival	Optional file-based replication
Cost	Price/capacity starts at few dollars/GB (ATA)	Best price/capacity, e.g. LTO-3 ($60 list) holds 800 GB compressed	VTL license + disk cost*

*File-level deduplication may reduce disk cost

Backup and Recovery Design

Architect Backup Environment

- Tape and disk considerations for VLDB
  - Allocate disk backup for most critical tablespaces
  - Utilize locally attached tape drives
    - If tapes must be shared, consider disk and/or tape backup SAN
  - Assess host resource utilization, production disk I/O, HBA/network & tape drive throughput
    - Minimum performant component of these will be bottleneck

Agenda

- VLDB Trends and Requirements
- Backup and Recovery Design
  - Assess Recovery Requirements
- Backup and Recovery Deployment
  - Architect Backup Environment
  - Evaluate Backup and Recovery Tools
  - Plan Data Layout
  - Develop Backup Procedure
  - Develop Recovery Strategies
- Starbucks Case Study
- Summary/Q&A

Backup and Recovery Deployment

Leverage Oracle Backup and Recovery Tools

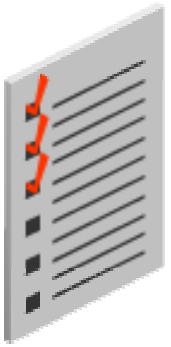

- Recovery Manager (RMAN)
  - Oracle-native tool for creating and managing physical backups, restoring and recovering databases
  - Multiplexed, parallel, block-validated full and incremental backups scale with VLDB needs
  - Supports data file, tablespace, and database backup and recovery, in addition to block media recovery
  - Flash Recovery Area manages all recovery-related files, e.g. backups, archived logs
- Data Pump
  - Logical backups of metadata and data (e.g. table, database)

Backup and Recovery Deployment

Leverage Oracle Backup and Recovery Tools

- Oracle Secure Backup (OSB)
  - Centralized tape backup management for database & file system
  - Fastest database backup to tape on the market via elimination of unused blocks and committed undo during backup
  - Low price of $3,000/tape drive
- Flashback Technologies
  - Broad suite of user error analysis and correction tools at the row, transaction, table, and database level
  - Optimized, continuous database snapshot with Flashback Database (built-in Continuous Data Protection (CDP) capability)
- Data Guard
  - Manages fast switchover and failover from production database to synchronized standby database on separate hardware
  - Serves as a remote, synchronized database ‘backup’

Error Correction with Flashback

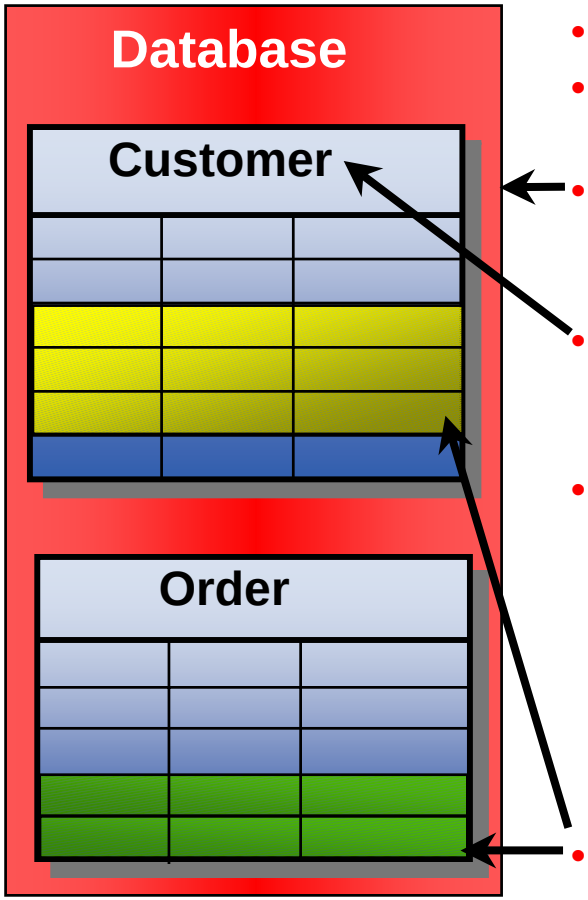

- Correct errors at any level
- Recovery time is function of changes, not database size
- **Flashback Database** – restore database to point-in-time (flashback log-based)
- **Flashback Table** – recover contents of tables to point-in-time (undo-based)
- **Flashback Drop** – restore accidentally dropped tables (based on free space in tablespace)

ORACLE 11g DATABASE

- **Flashback Transaction** – back out transaction and all subsequent conflicting transactions (redo-based)

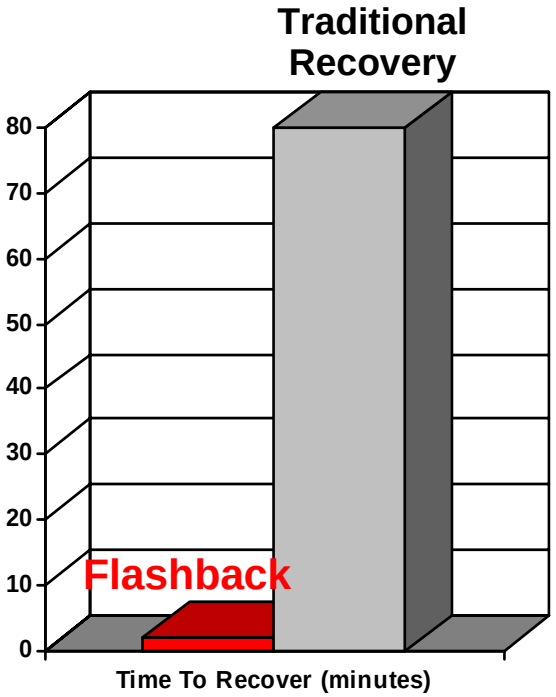

Backup and Recovery Deployment

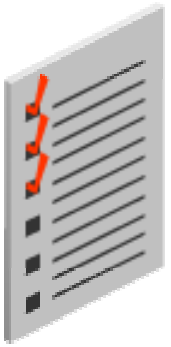

Third Party Disk Backup Solutions

- Snapshots **Additional Cost**
  - Fast point-in-time copy, but exposed to production array failures
- Split Mirror Backup **Additional Cost**
  - Tape backup can be offloaded, but additional setup required
  - Starbucks will present their criteria evaluating RMAN+Flash Recovery Area with split mirror backup
- CDP Appliances **Additional Cost**
  - Continuous snapshot via copy-on-write or allocate-on-write

	Third Party CDP	**Flashback Database**
I/O consumption	• All file changes must be logged (control file, redo, data file) • Each and every block change must be logged	• Only data file block changes are tracked • Only one before-image is logged per 30 min interval, regardless of number of changes to the block

Agenda

- VLDB Trends and Requirements
- Backup and Recovery Design
  - Assess Recovery Requirements
  - Architect Backup Environment
- Backup and Recovery Deployment
  - Evaluate Backup and Recovery Tools
  - Plan Data Layout
  - Develop Backup Procedure
  - Develop Recovery Strategies
- Starbucks Case Study
- Summary/Q&A

Backup and Recovery Deployment

Plan Data Layout

- Exploit partitioning and read-only tablespaces
  - Older partitions can be moved to read-only tablespaces
  - Backup read-only tablespaces once, then periodically, depending on tape retention policy
- Table Compression – reduce space by 2-3X, no performance degradation accessing compressed data
  - OLTP Table Compression – minimize write overhead

ORACLE 11g DATABASE

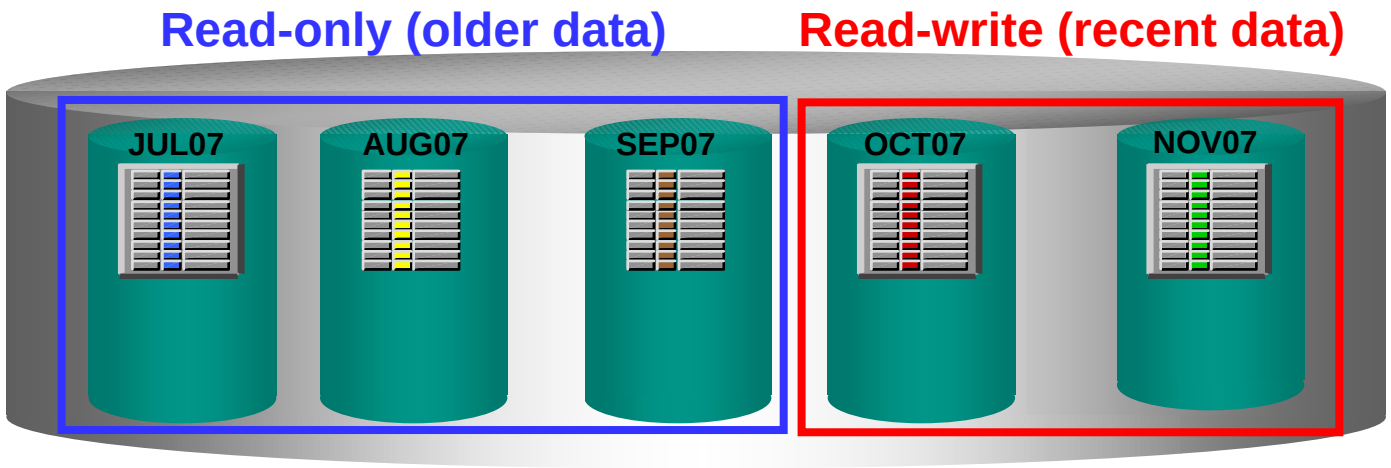

Agenda

- VLDB Trends and Requirements
- Backup and Recovery Design
  - Assess Recovery Requirements
  - Architect Backup Environment
- Backup and Recovery Deployment
  - Evaluate Backup and Recovery Tools
  - Plan Data Layout
  - Develop Backup Procedure
  - Develop Recovery Strategies
- Starbucks Case Study
- Summary/Q&A

Backup and Recovery Deployment

Develop Backup Procedure

- **Option 1: Use Level 0 (Full) and Fast Incremental Backups (Oracle Database 10g+)**
  - Weekly level 0 and nightly 'differential' incremental backups to tape, with optional compression
  - Only changed blocks are read and written during incremental
  - Archived logs are backed up and retained on-disk, as needed
  - Best for:
    - VLDBs that can tolerate hours/days for restore & recovery
    - Tape backup-only environments, where disk is premium
    - Low-medium change frequency between backups

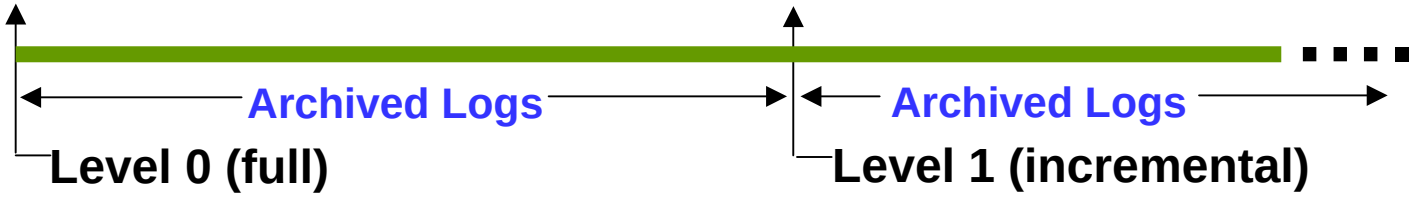

Backup and Recovery Deployment

Develop Backup Procedure

- **Option 2: Use Level 0, Fast Incremental Backups, Incrementally Updated Backup**
  - Initial level 0 image copy to disk, followed by nightly incremental backups
  - Roll forward image copy with incremental, to produce new on-disk full backup, on regular basis (e.g. daily/weekly/monthly)
  - Full backup archived to tape, as needed
  - Archived logs are backed up and retained on-disk, as needed
  - Best for:
    - VLDBs that can tolerate no more than several hours for restore & recovery
      - Fast recovery from disk or directly use image copies
    - Environments where disk can be allocated for full database image copy or image copy of most critical tablespaces

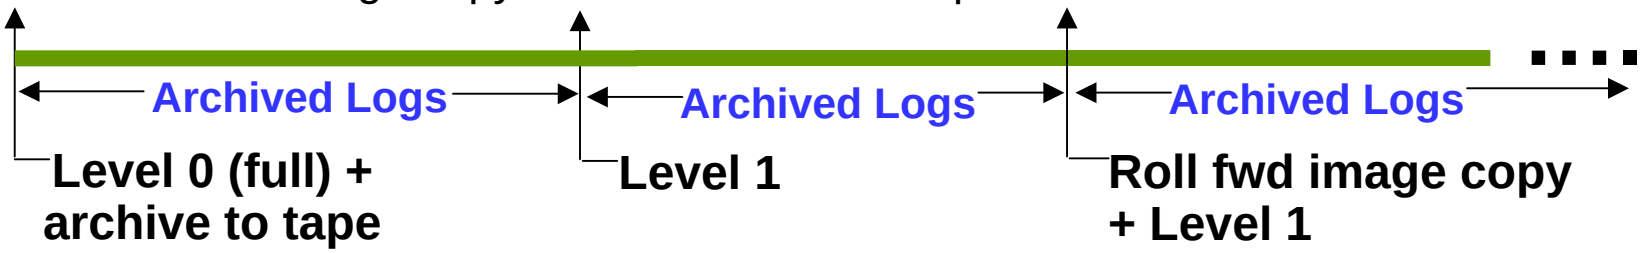

Backup and Recovery Deployment

Develop Backup Procedure

- **Option 3: Data Guard + Full/Incremental Backups**
  - Full and incremental backups can be offloaded to physical standby database and used for restore at primary or standby database
    - Backups can be taken at each database for optimal local protection
  - Best for:
    - VLDBs that require no more than several minutes of recovery time, in event of any failure
    - Environments that can allocate symmetric hardware and storage for physical standby database
    - Environments whose tape infrastructure can be shared between primary and standby database sites

Backup and Recovery Deployment

Develop Backup Procedure

- **Option 4: Maintain ETL loads**
  - Take full backup and maintain *n* ETL loads, where *n* is number of days between full backups
  - Restore last full backup and run needed ETL loads
  - Manual backup and recovery process versus Options 1-3
  - Best for:
    - VLDBs that do not require point-in-time recovery or do not have space for archived logs
    - VLDBs whose changes are confined to controlled batch jobs
    - Environments that cannot keep both incremental backups and load files
    - Legacy systems that will not be upgraded

Backup and Recovery Deployment

Develop Backup Procedure

- NOLOGGING considerations
  - Data warehouse environments commonly use NOLOGGING to increase performance and save space
  - Recovery to point-in-time within NOLOGGING results in some corrupt blocks
  - Best practice is to take incremental backup when NOLOGGING finishes, so ideal for controlled batch load environments
  - E.g. Option 1: Level 0 + Fast Incremental Backups

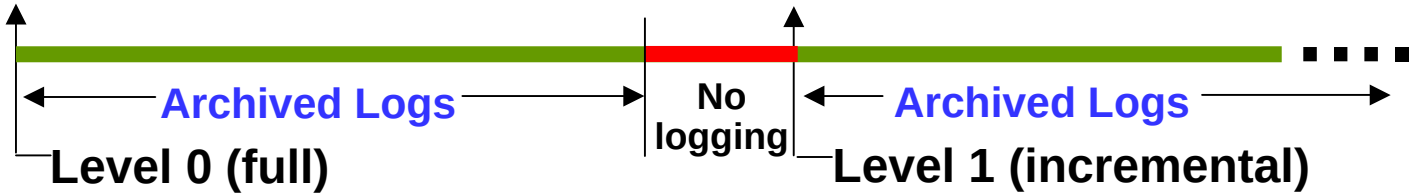

Backup and Recovery Deployment

Develop Backup Procedure

- Other considerations
  - Divide full backup workload across multiple days
    - `BACKUP DATABASE NOT BACKED UP SINCE 'SYSDATE-3' DURATION 06:00 PARTIAL MINIMIZE TIME;`
    - Manage backup window and user impact
    - Full restore requires data file backups potentially taken on several tapes
  - Save time with tablespace level backups
    - Backup index tablespaces less frequently than data tablespaces
    - Backup little used/less critical tablespaces less frequently
    - Reduce restore time for most critical tablespaces, by grouping them together in separate backups

Backup and Recovery Deployment

Develop Backup Procedure

- Oracle Database 11g Enhancements
  - Multi-section backups (intrafile parallel backup)
    - Best for databases composed of a few data files that are significantly larger than the rest (e.g. bigfile tablespace)
      - Or, if there are fewer data files than # of tape drives
    - Scales linearly with addition of channels
    - Start with section size = average data file size / # of channels
  - Fast backup compression (ZLIB) -- 40% faster than Oracle Database 10g
    - Suitable for daily, incremental backups
    - Use BZIP2 for full backup -- better compression ratio
  - Data Pump export compression, including metadata & table data

Agenda

- VLDB Trends and Requirements
- Backup and Recovery Design
  - Assess Recovery Requirements
  - Architect Backup Environment
- Backup and Recovery Deployment
  - Evaluate Backup and Recovery Tools
  - Plan Data Layout
  - Develop Backup Procedure
  - Develop Recovery Strategies
- Starbucks Case Study
- Summary/Q&A

Backup and Recovery Deployment

Develop Recovery Strategies

Recovery Scenario	Oracle Technologies
Media Failure	RMAN – restore all files to new storage location
Block Corruption	RMAN Block Media Recovery, Trial Recovery, LogMiner
User/Logical Error	Flashback Technologies, RMAN TSPITR, LogMiner
Disaster	Data Guard; RMAN -- restore all files to new host/storage

- **Data Recovery Advisor – built-in database failure diagnosis, analysis, & repair tool**

ORACLE DATABASE 11g

Starbucks Case Study

Agenda

- VLDB Trends and Requirements
- Backup and Recovery Design
  - Assess Recovery Requirements
  - Architect Backup Environment
- Backup and Recovery Deployment
  - Evaluate Backup and Recovery Tools
  - Plan Data Layout
  - Develop Backup Procedure
  - Develop Recovery Strategies
- Starbucks Case Study
- Summary/Q&A

Resources

- **Maximum Availability Architecture white papers:**
  http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm
- **Oracle HA Portal on OTN:**
  http://www.oracle.com/technology/deploy/availability/
- **Oracle HA Customer Success Stories on OTN:**
  http://www.oracle.com/technology/deploy/availability/htdocs/HA_Case Studies.html

Summary

- VLDBs growing rapidly
  - More multi-PB systems in few years?
- Methodically work through backup and recovery design and deployment
  - Understand Data Layout and Change Characteristics
  - Assess Recovery Requirements
  - Architect Backup Environment
- Put design to the test with deployment
  - Leverage Oracle Backup Tools & Strategy
  - Develop and Test Recovery Strategies

QUESTIONS
ANSWERS

Database HA Sessions From Oracle Development

Wednesday, Nov 14

- S291487 - Backup and Recovery Best Practices for Very Large Databases (VLDB), 11:15 am - 12:15 pm, Moscone South 104
- **S291920 - Oracle Active Data Guard: How to Utilize Your Standby Databases for Production Workload - What They Didn't Print in the Manuals, 3:00 - 4:00 pm, Moscone South 304**
- **S291917 - Oracle Data Guard Tips and Tricks: Direct From Oracle Development, 4:30 - 5:30 pm, Moscone South 102**

Thursday, Nov 15

- S291495 - Oracle Streams Replication and Advanced Queuing (AQ): What's New in Oracle Database 11*g*, 8:30 - 9:30 am, Moscone South 304
- S291499 - Best Practices for Implementing Replication with Oracle Streams in Oracle Database 10*g* and 11*g*, 10:00 - 11:00 am, Moscone South 304
- S291525 - Maximum Availability Architecture (MAA) Best Practices: Online Patching, Rolling Upgrades and Planned Maintenance with Minimal Downtime with Oracle Database, 11:30 am - 12:30 pm, Moscone South 104
- S290542 - Maximum Availability Architecture (MAA) Best Practices for Siebel 8.0, 2:30 pm - 3:30 pm, Marriott Salon 10 & 11

ORACLE IS THE INFORMATION COMPANY