

Starbucks Enterprise Data Warehouse (EDW) VLDB Backup and Recovery Architecture



Greg Green
Senior Database Administrator
November 14, 2007

Starbucks EDW Backup and Recovery Architecture

- Current EDW Environment
 - Leveraging RMAN Features
- Oracle Database 11g Features Applicable to Our Environment
- Comments and Lessons Learned

Going Beyond Coffee



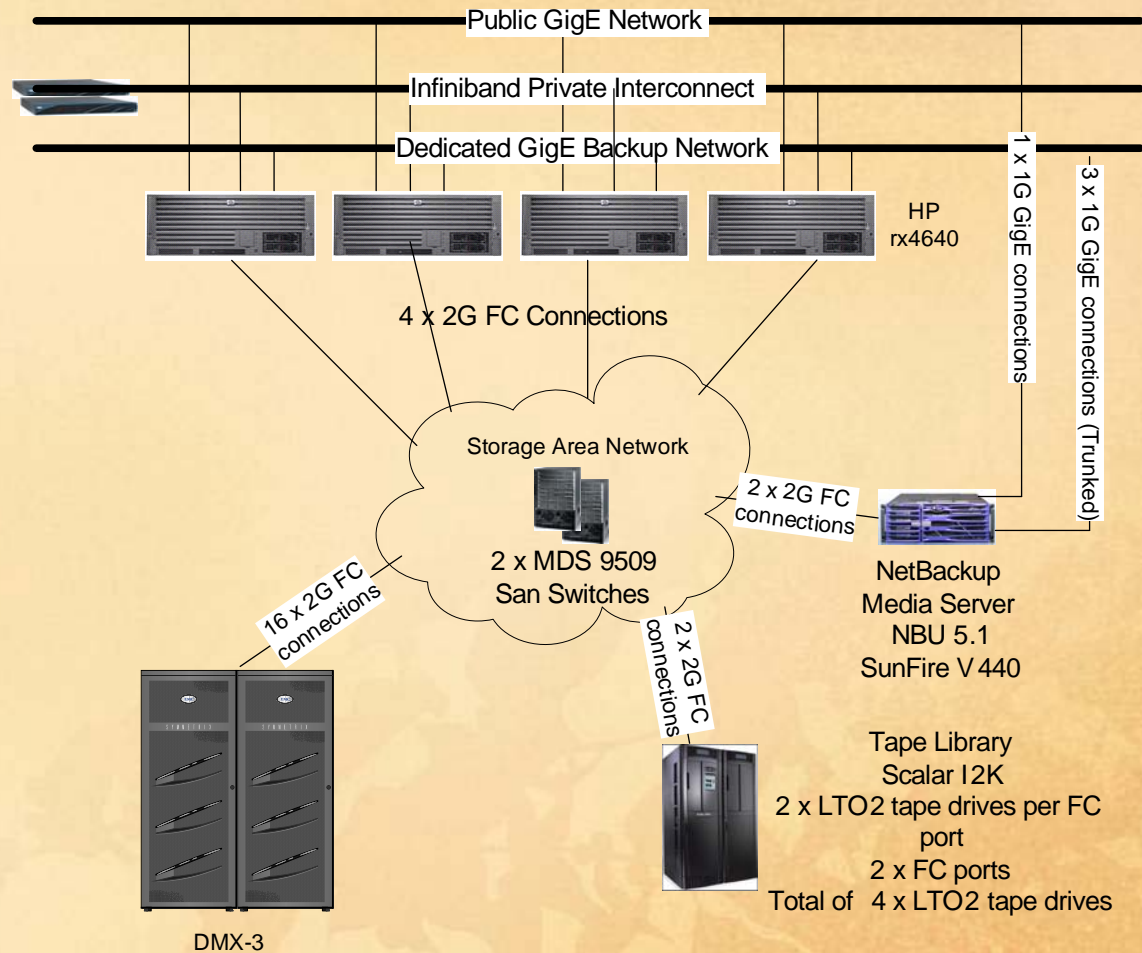
Who it Supports

- Production EDW supports Starbucks internal business users
 - 5 TB VLDB warehouse, growing 2-3 TB per year
 - Provides reports to the store level – sales, staffing, etc.
- Thousands of stores will directly access the EDW database by early 2008
 - Web-based dashboard reports via company intranet
- Front-end reporting with Microstrategy
- Leveraging Ascential Datastage ETL Tool
 - Toad, SQL Developer, and other ad-hoc tools used by developers and QA

Production Hardware

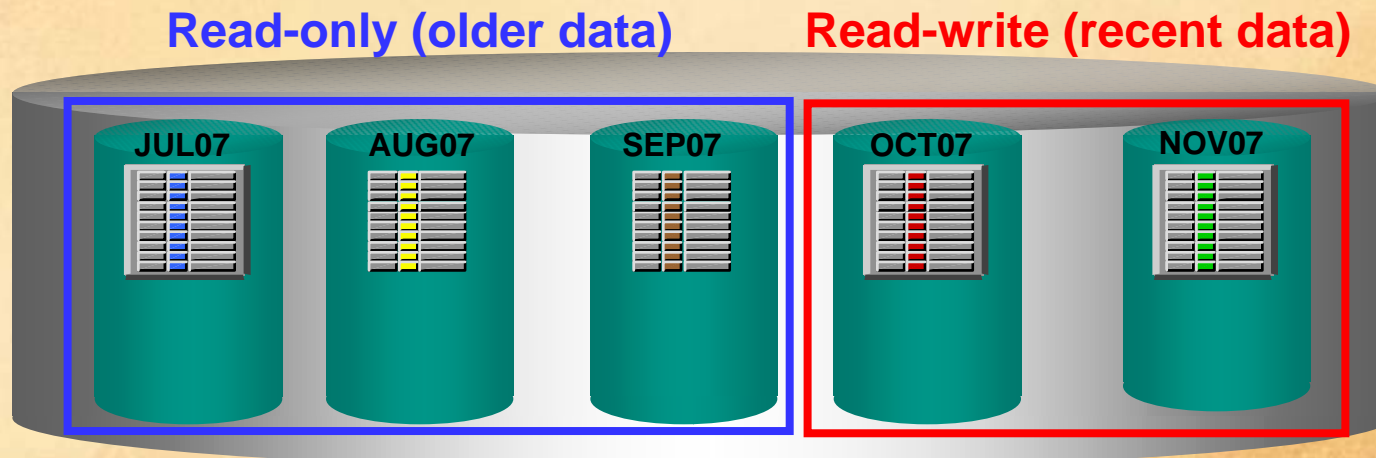
4 Node RAC Database

- Servers – 4 CPU
HP ia64
1.5 GHz CPU 16 GB memory
- Network – Private Interconnect
Infiniband Ib
- Public Network – Gigabit Ethernet
- Storage – SAN
- 16 (8 Usable) Terabytes Raid 1+0 Storage 15k 146 GB Drives
- 8 Terabytes Raid 5 (3+1) 300 GB Drives
- DB Software – Oracle 10.2.0.3 EE
- Media Mgmt – NetBackup 5.1
- RMAN Backup



Data Layout

- Data Layout for Backups is Key – View the EDW database as having an ACTIVE and PASSIVE portion.
- Large Partitioned Fact Tables are spread across month level tablespaces (Critical for my RMAN Duplication Strategy)
- Older tablespaces can be regularly moved to read-only



Backup Criteria

- In addition to RTO and RPO we also considered other factors
 - Minimize the load on the servers and array to perform the backup tasks
 - Minimize overall backup time
 - Keep disk space consumed by backups to a minimum
 - Minimize backup scripting & infrastructure setup time
 - Reduce hardware and storage costs
 - Ensure backup process can scale as the warehouse grow
- RMAN+Flash Recovery Area (FRA) or Split Mirror Backups could meet these requirements
 - However, RMAN was chosen for it's ease of implementation, greater flexibility, and lower cost.

Backup Strategy

- RMAN Rolling Image Copy Backups
 - Disk - Flash Recovery Area (FRA)
 - Monthly Level 0 Image Copy backups to FRA
 - Daily Level 1 Differential Incremental Backups
 - Daily Roll Image Copy forward for the rest of the month with `'SYSDATE - 1'`
 - Daily Disk Script:

```
{  
    RECOVER COPY OF DATABASE WITH TAG  
      'WEEKLY_FULL_BKUP'  
    UNTIL TIME 'SYSDATE - 1';  
    BACKUP INCREMENTAL LEVEL 1 FOR RECOVER OF  
      COPY WITH TAG WEEKLY_FULL_BKUP DATABASE;  
    BACKUP AS BACKUPSET ARCHIVELOG ALL NOT  
      BACKED UP DELETE ALL INPUT;  
    DELETE NOPROMPT OBSOLETE RECOVERY WINDOW OF  
      2 DAYS DEVICE TYPE DISK; }
```
 - Tape
 - Weekly “Backup Recovery Area”
 - Rest of the week “Backup Backupset ALL”

Backup Performance to FRA

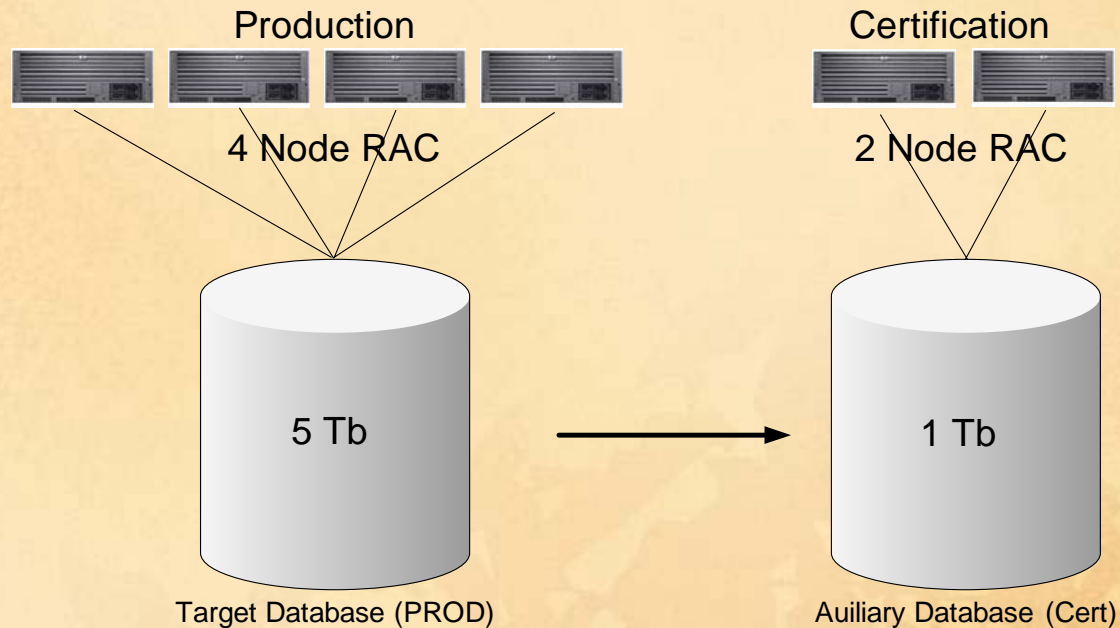
- **Nightly Incremental Backups**
 - 45-60 minutes for recovery of image copy to 'sysdate - 1' and new incremental backup + archive log backup. The uncompressed backup set is typically 200-250 GB on a nightly basis.
 - 4 RMAN disk channels running on one cluster node
- **Full Backups**
 - Approximately 45 minutes to delete existing Image Copy (5 TB)
 - Approximately 5.5 hours to create a new level 0 backup (1 TB every 70 minutes)

Backup Performance to Tape

- Nightly Backup of Backupsets
 - Nightly Backup time to tape for 2 tape channels on one cluster node takes 50 minutes (roughly 200 GB at 240 GB/hr)
- Weekly Backup of Recovery Area
 - Weekly Backup of Recovery area with 4 channels (2 channels on 2 nodes) takes 15-16 hours (currently ~5.5 TB)
 - LTO2 tape drives => 4 * (30-35) MB/s
 - Should see at least 500 GB/hr going to tape
 - Observing 330-340 GB/hr
 - Tuning exercise showed the bottleneck is on the Bus of V440

RMAN Duplication

- Clone 4-node RAC production to 2-node RAC non-production (certification)
- Lessons learned from our specific environments



Starbucks EDW Backup and Recovery Architecture

- Current EDW Environment
 - Leveraging RMAN Features
- Oracle Database 11g Features Applicable to Our Environment
- Comments and Lessons Learned

Oracle Database 11g RMAN New Features for Our Environment

- Multisection Backups
- RMAN Command Script Substitution Variables

Multi-Section Backups

- “Divide and Conquer” bigfiles

Example

```
rman> backup tablespace test_data01 section size 2G;
starting backup at 02-OCT-07
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=153 device type=DISK
allocated channel: ORA_DISK_2
channel ORA_DISK_2: SID=154 device type=DISK
channel ORA_DISK_1: starting full datafile backup set
channel ORA_DISK_1: specifying datafile(s) in backup set
input datafile file number 00006 name=/opt/oracle/oradata/orcl/test_data01.dbf
Backing up blocks 1 through 262144
channel ORA_DISK_1: starting piece 1 at 02-OCT-07
channel ORA_DISK_2: starting full datafile backup set
channel ORA_DISK_2: specifying datafile(s) in backup set
input datafile file number 00006 name=/opt/oracle/oradata/orcl/test_data01.dbf
Backing up blocks 262145 through 524288
```

RMAN Substitution Variables

- Per the RMAN 11g Backup and Recovery Guide
 - “You can create RMAN command files and stored scripts that accept user input at runtime. Thus, backup scripts can use RMAN substitution variables for tags, filenames, restore point names, and so on.”
- How could this apply to our environment?
 - Read-Only Data Driven Backups
 - Spread Backup of Read-Only Tablespaces over a month
 - Evens out usage of tape drives over the month with no “spikes”.
 - Ensures all data has been backed up with the last 30 days and I’m not relying on a backup tape from 2 1/2 years ago.

Read-Only Tablespace Backups

- Script Example

```
Create script read_only_tbs_backup
{
  backup device type sbt copy of tablespace &1
  section size 50G;
  backup device type sbt copy of tablespace &2
  section size 50G;
}
```

Shell Script:

```
Runbackup.sh read_only_tbs_backup HST_200703
HST_200704
```


Oracle Database 11g RMAN Features of Interest

- Backup Compression
- Undo Optimization
- Network Duplicate from Active Database
- Improved Long-Term Backups

Starbucks EDW Backup and Recovery Architecture

- Current EDW Environment
 - Leveraging RMAN Features
- Oracle Database 11g Features Applicable to Our Environment
- Comments and Lessons Learned

Comments and Lessons Learned

- Read OTN articles to find latest backup features
- Read 11g Release Notes on RMAN
- Size Flash Recovery Area for the worst case scenario
 - Filling up the FRA is not fun, especially when it is in an ASM diskgroup
- Be prepared for a debate – Ask 10 dbas how to design a large backup and they'll give you 10 different answers. Everyone has his/her own opinion on backups.
- You may have to break your existing standards to get the job done.

Summary

- Current EDW Environment
- Oracle Database 11g Features Applicable to Our Environment
- Comments and Lessons Learned