How HP delivered a 3TB/hour Oracle™ backup & 1TB/hour restore

Andy Buckley
Technical Advocate
HP Network Storage Solutions

HP WORLD 2003
Solutions and Technology Conference & Expo

Why Oracle?

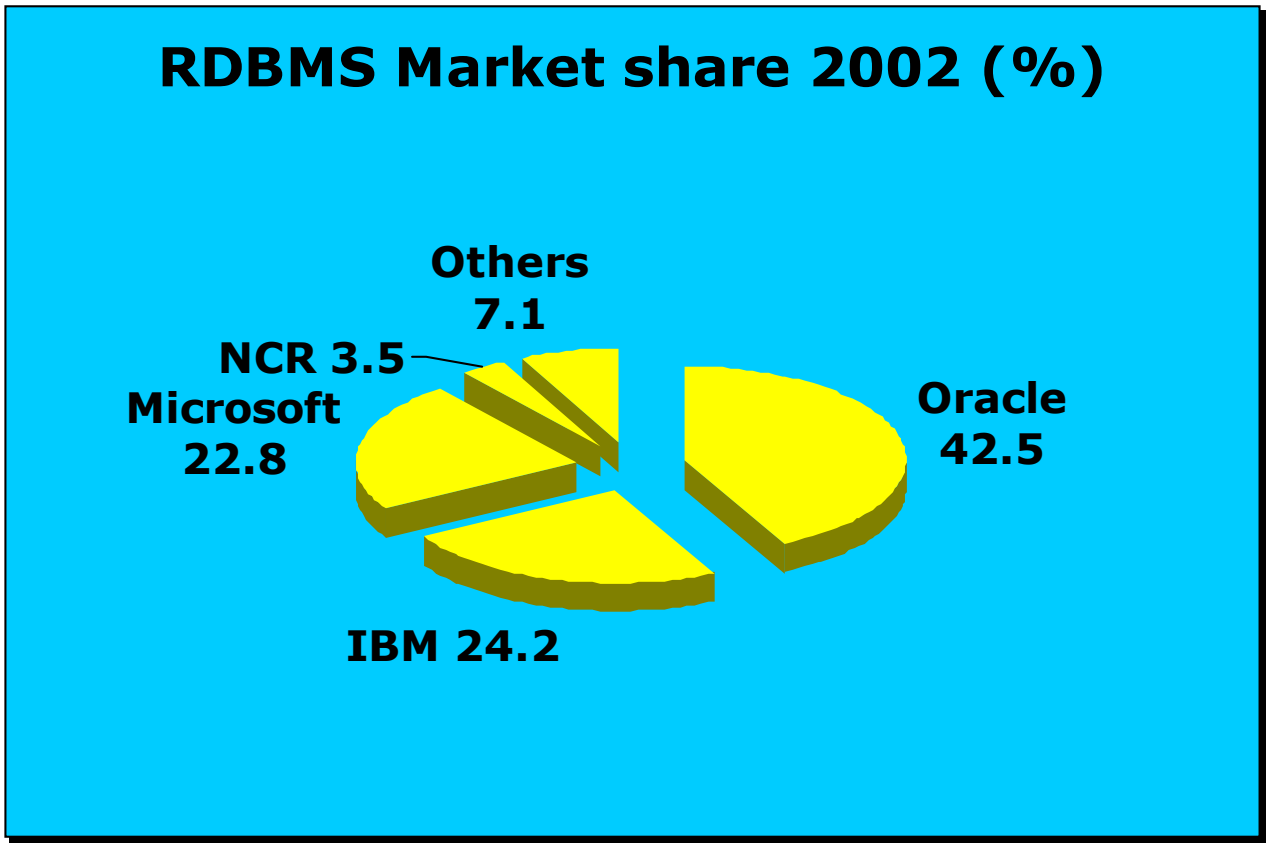

source: Gartner dataquest May 2003

Why Ultrium tape?

- Capacity – 200GB Native
- Performance – 30MB/sec Native
  - the only technology with adaptive tape speed
- Reliability
- Industry standard
- Solid 4 generation roadmap

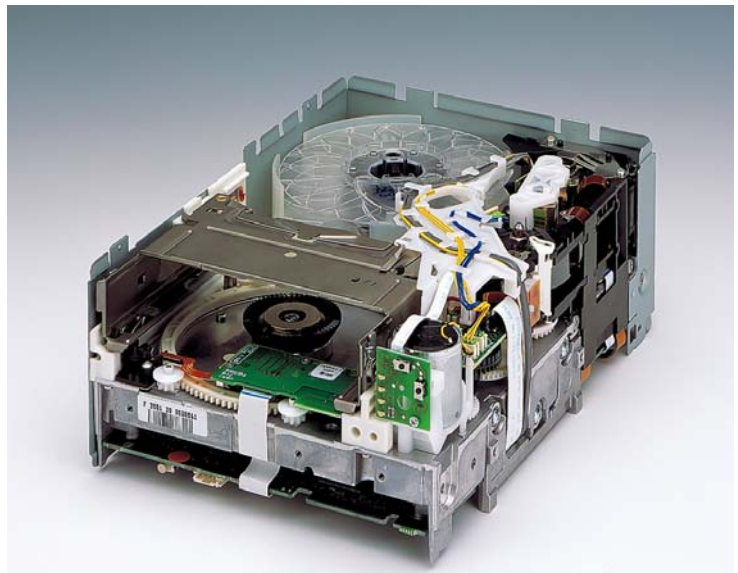

Ultrium 2 mechanism

Oracle Basics

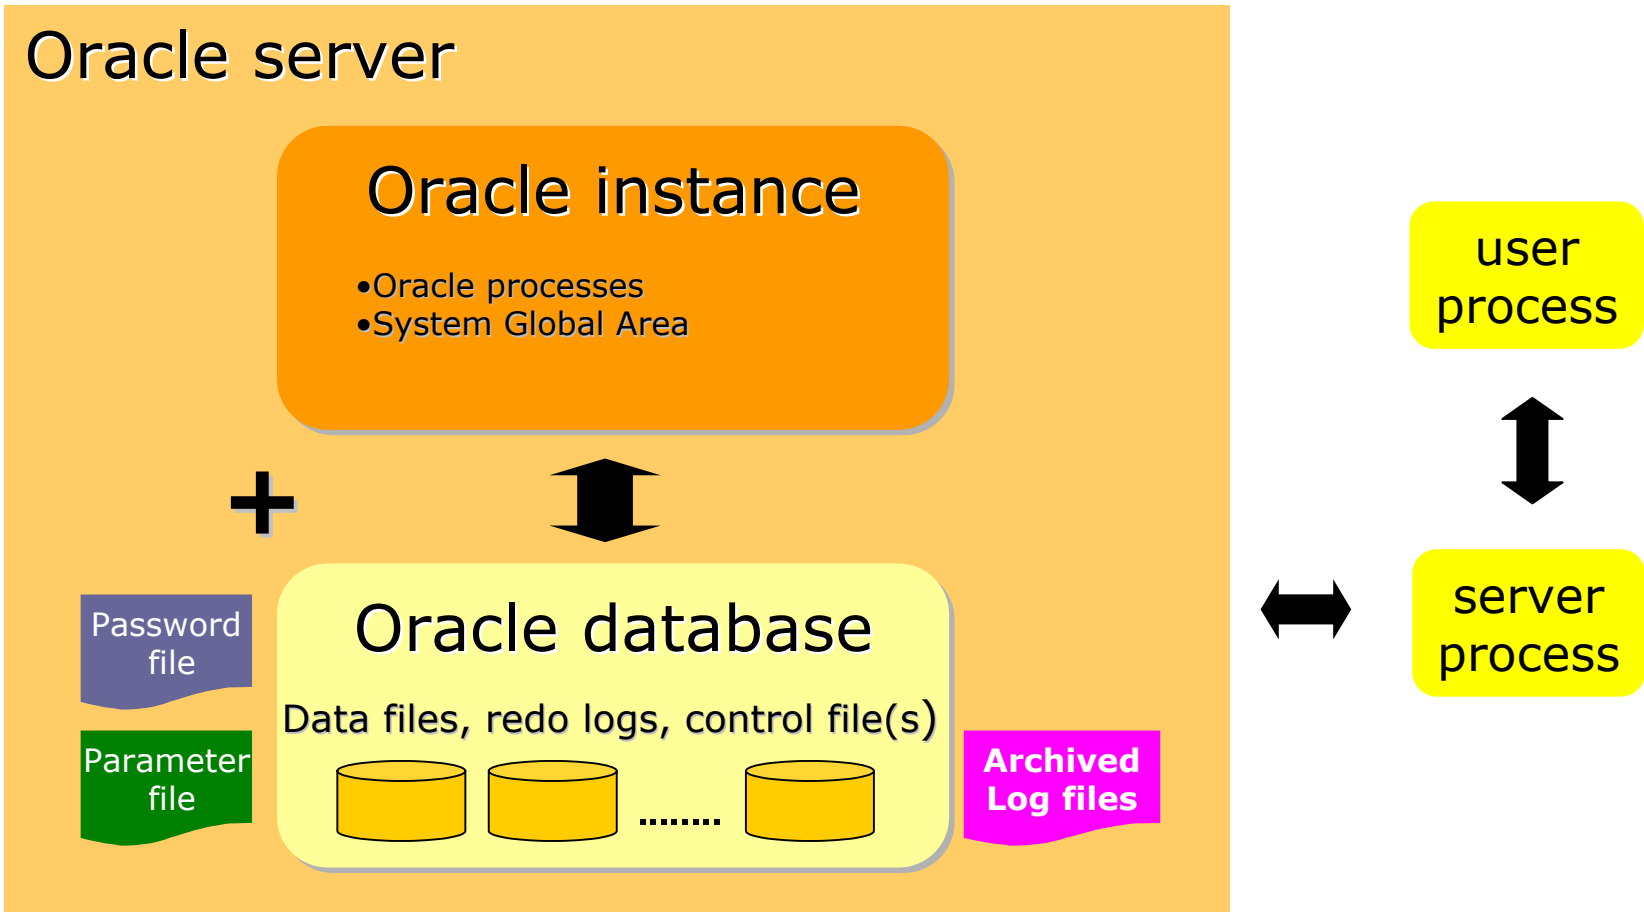

Database backup options

- Disk based
  - Snapshots & Mirroring
  - Oracle Recovery Manager backup to disk
- Tape based backup
  - offline (cold) backups
  - online (hot) backups ← **Area of focus**
- Most Backup & Recovery strategies use a combination of tape and disk

To find out more.....

Backup-to-Disk, and Multi-Level Data Protection

Session 1619
Harald Burose
Senior Architect
HP Nearline Storage Division

Session 1597
Glenn Wuenstel
Solutions Systems Engineer
HP Nearline Storage Division

Test lab - system diagram

HP Superdome server

32 x 2Gbps fibre channels

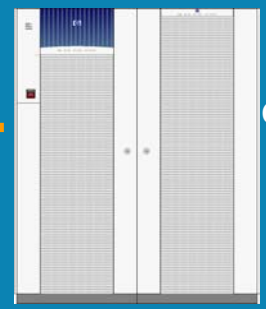

XP1024 disk array

HP fc16b Fibre Switch X 2

storage routers

m2402 m2402
m2402 m2402
m2402 m2402
m2402

FC FC

HP m2402 storage router

SCSI SCSI SCSI SCSI

Ultra-3 SCSI

Ultra-3 SCSI

1 2

robotics

3 16

tape drive tape drive tape drive

16 Ultrium 460 tape drives in ESL9595 library

Selective Storage Presentation – optimises performance by maximising FC bandwidth

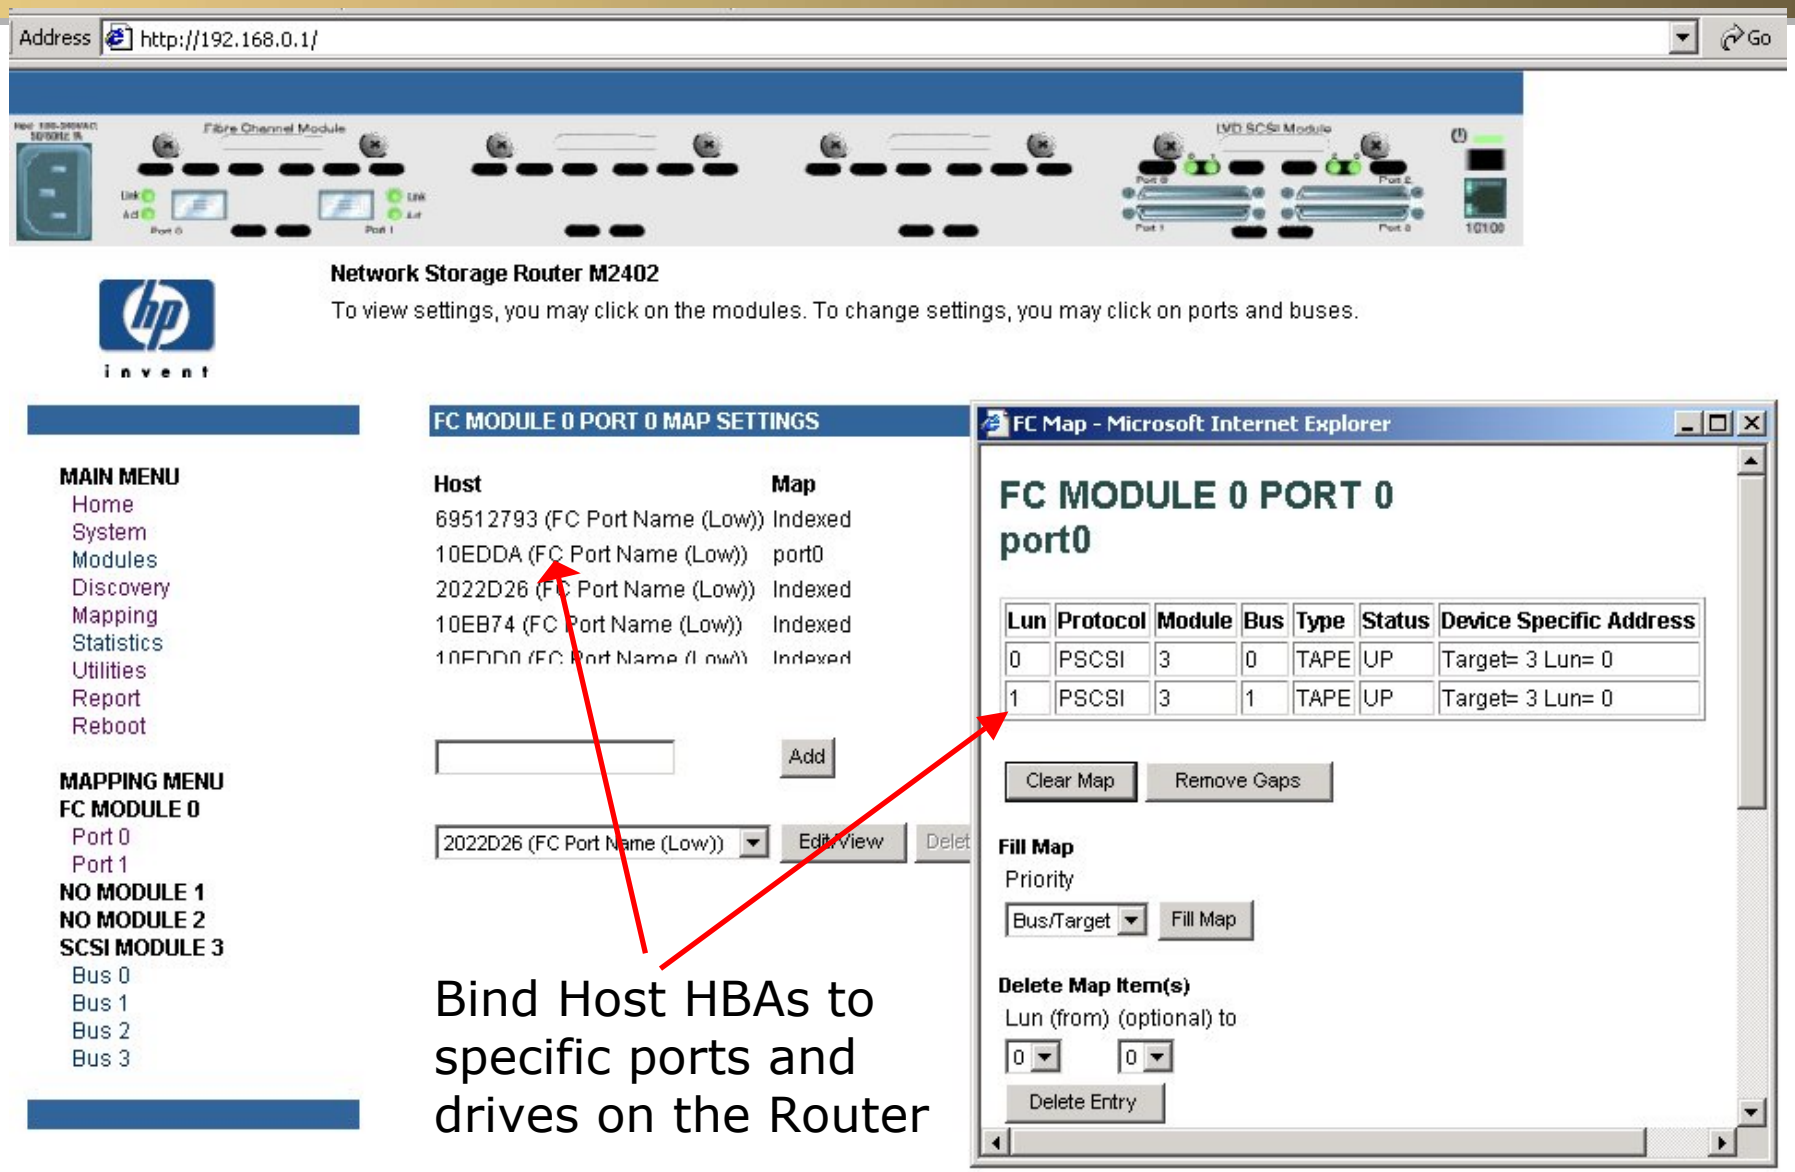

HP-UX/Oracle9i

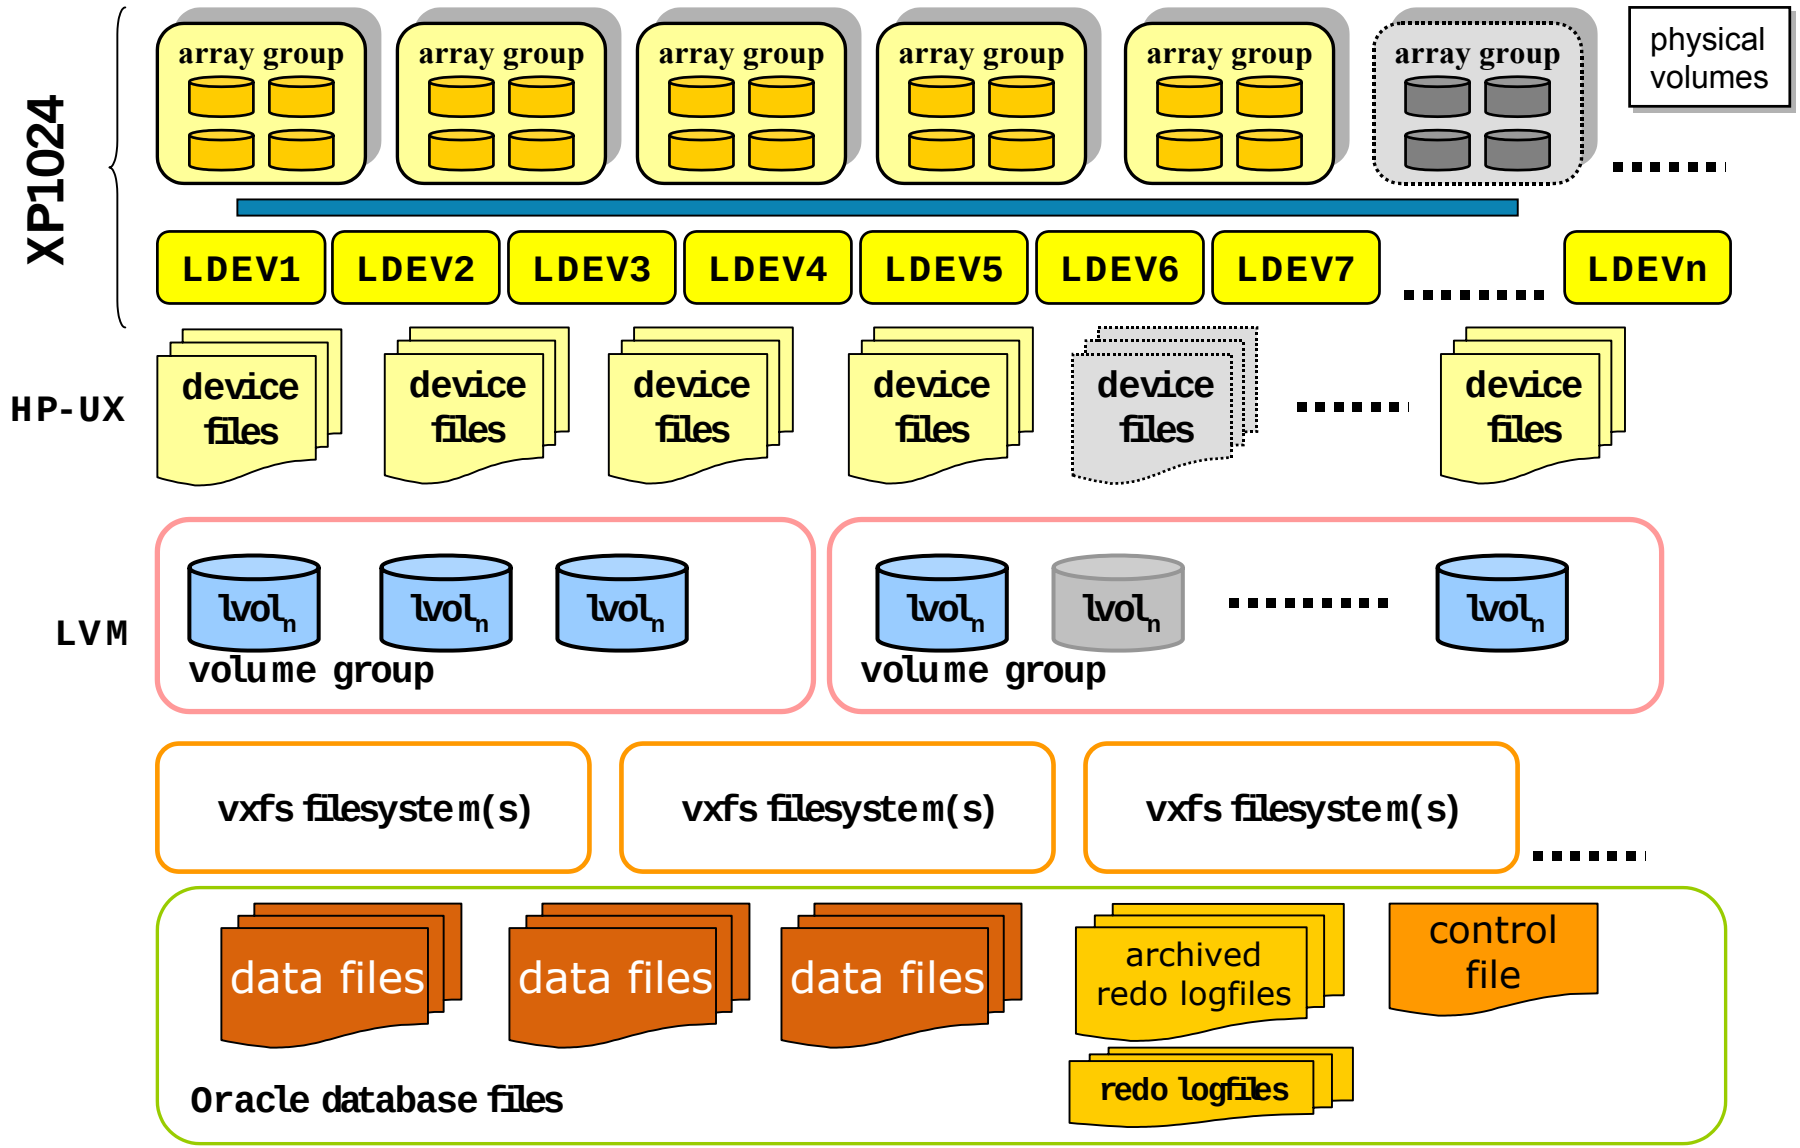

Generating the test database

- Test databases are available from the Transaction Processing Council TPC

  http://www.tpc.com

- Used TPC-H dbgen tool

- Limit of 535GB per mount point as a result of the 8MB extent size. This is a limitation set by the HP-UX logical volume manager.

Test database structure

- 9TB of space allocated with 4.3TB of real data loaded
- 8 tables within the TPC schema
  - Line item 15billion records partitioned over 45 tablespaces
  - 3 datafiles per tablespace
  - Order : 4.5 billion records over 15 tablespaces
  - Partsupp : 2.4 billion records over 20 tablespaces
  - Part: 600 million records 0ver 20 tablespaces
  - Customer : 400 million records over 20 tablespaces
  - Supplier: 30 million records over 20 tablespaces
  - Nation: 25 records in a single tablespace
  - Region: 5 records in single tablespace
- Total: 259 datafiles and 146 tablespaces

Test Database structure – partitioned data

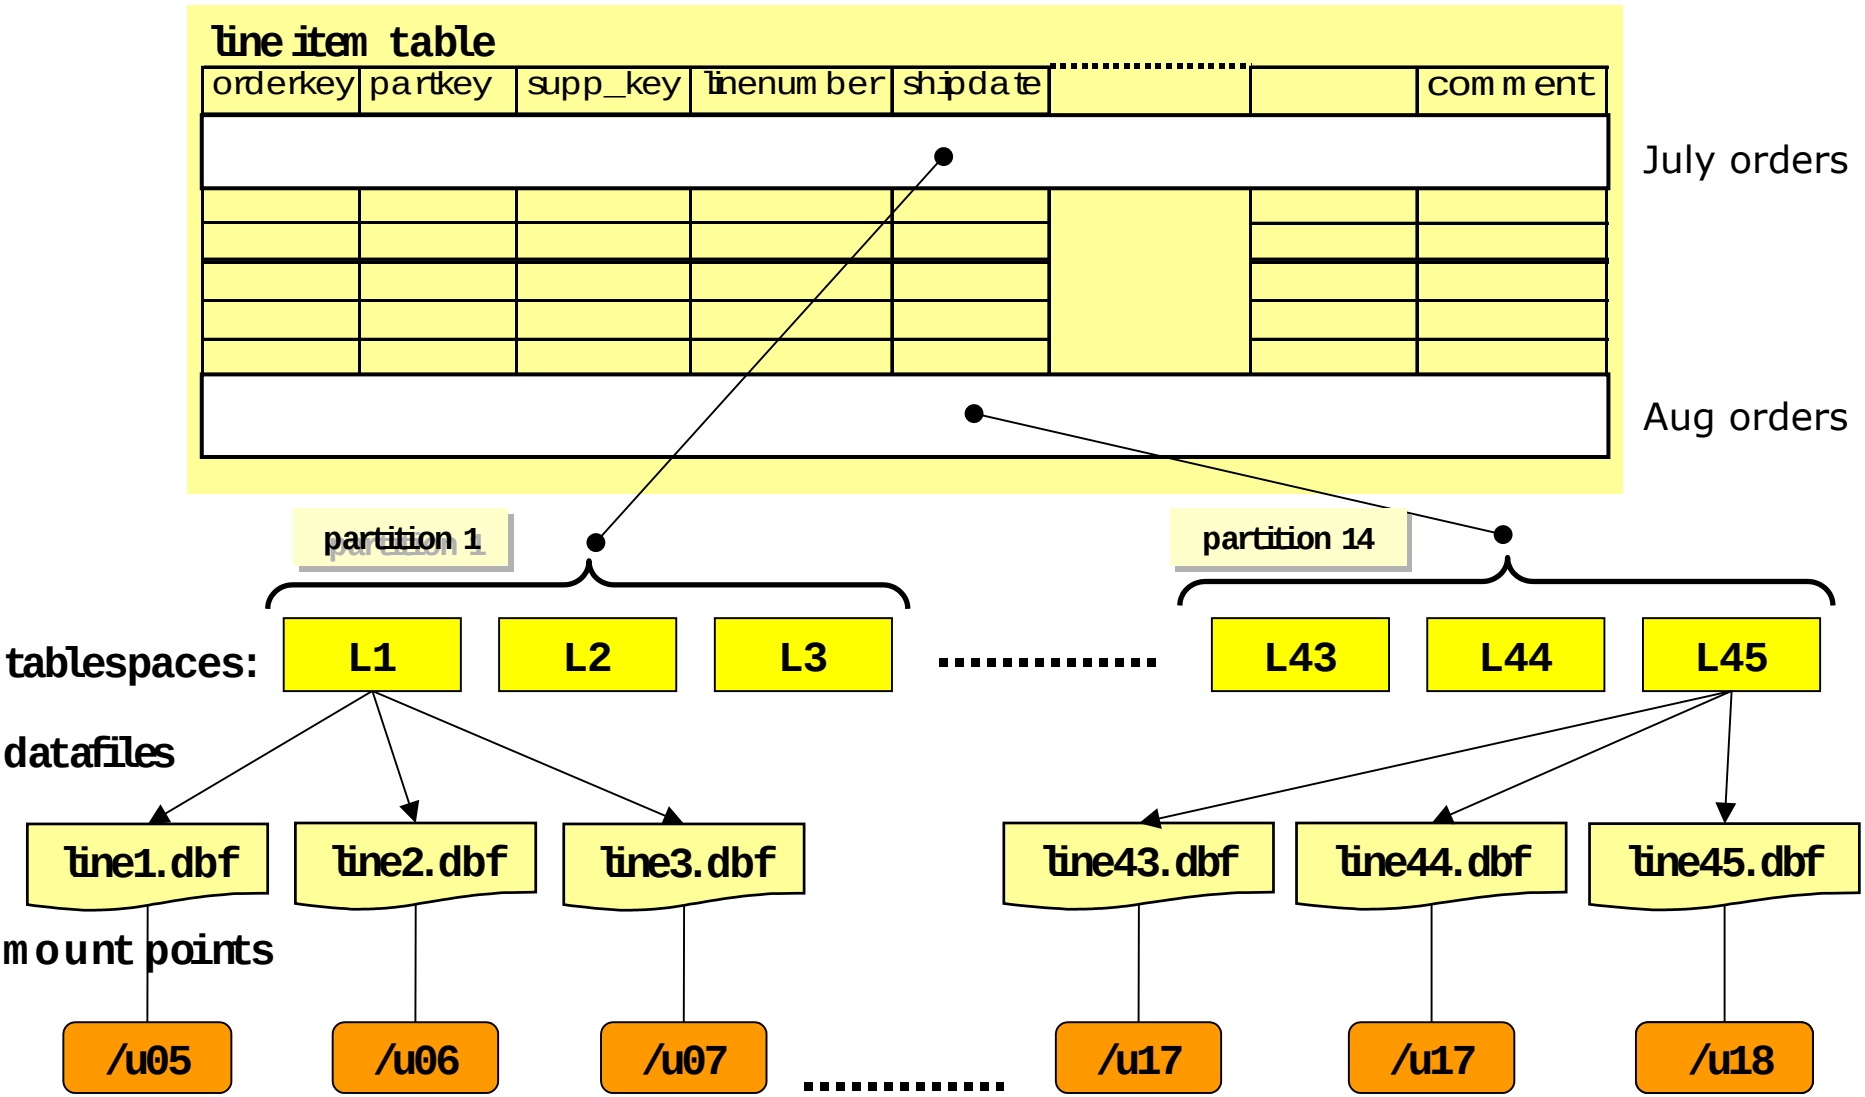

Oracle9i Recovery Manager

- commonly called RMAN
- available from Oracle 8 onwards
- can perform online & offline backups
- ability to back up ‘raw disk’ as well as filesystem
- provides block level data integrity checks
- eliminates backup of ‘whitespace’
- provides incremental backup functionality
- provides ‘point in time’ backup
- platform independent
- database driven process
- scriptable

What does a media manager do?

- The Oracle recovery manager requires a media manager to backup to a tape device.
- Controls the tape library.
- Provides RMAN with access to tape devices.
- Tracks the location of media in the library
- HP Data Protector was used as the media manager for the 3 TB/hr project.

HP Data Protector & Oracle Integration

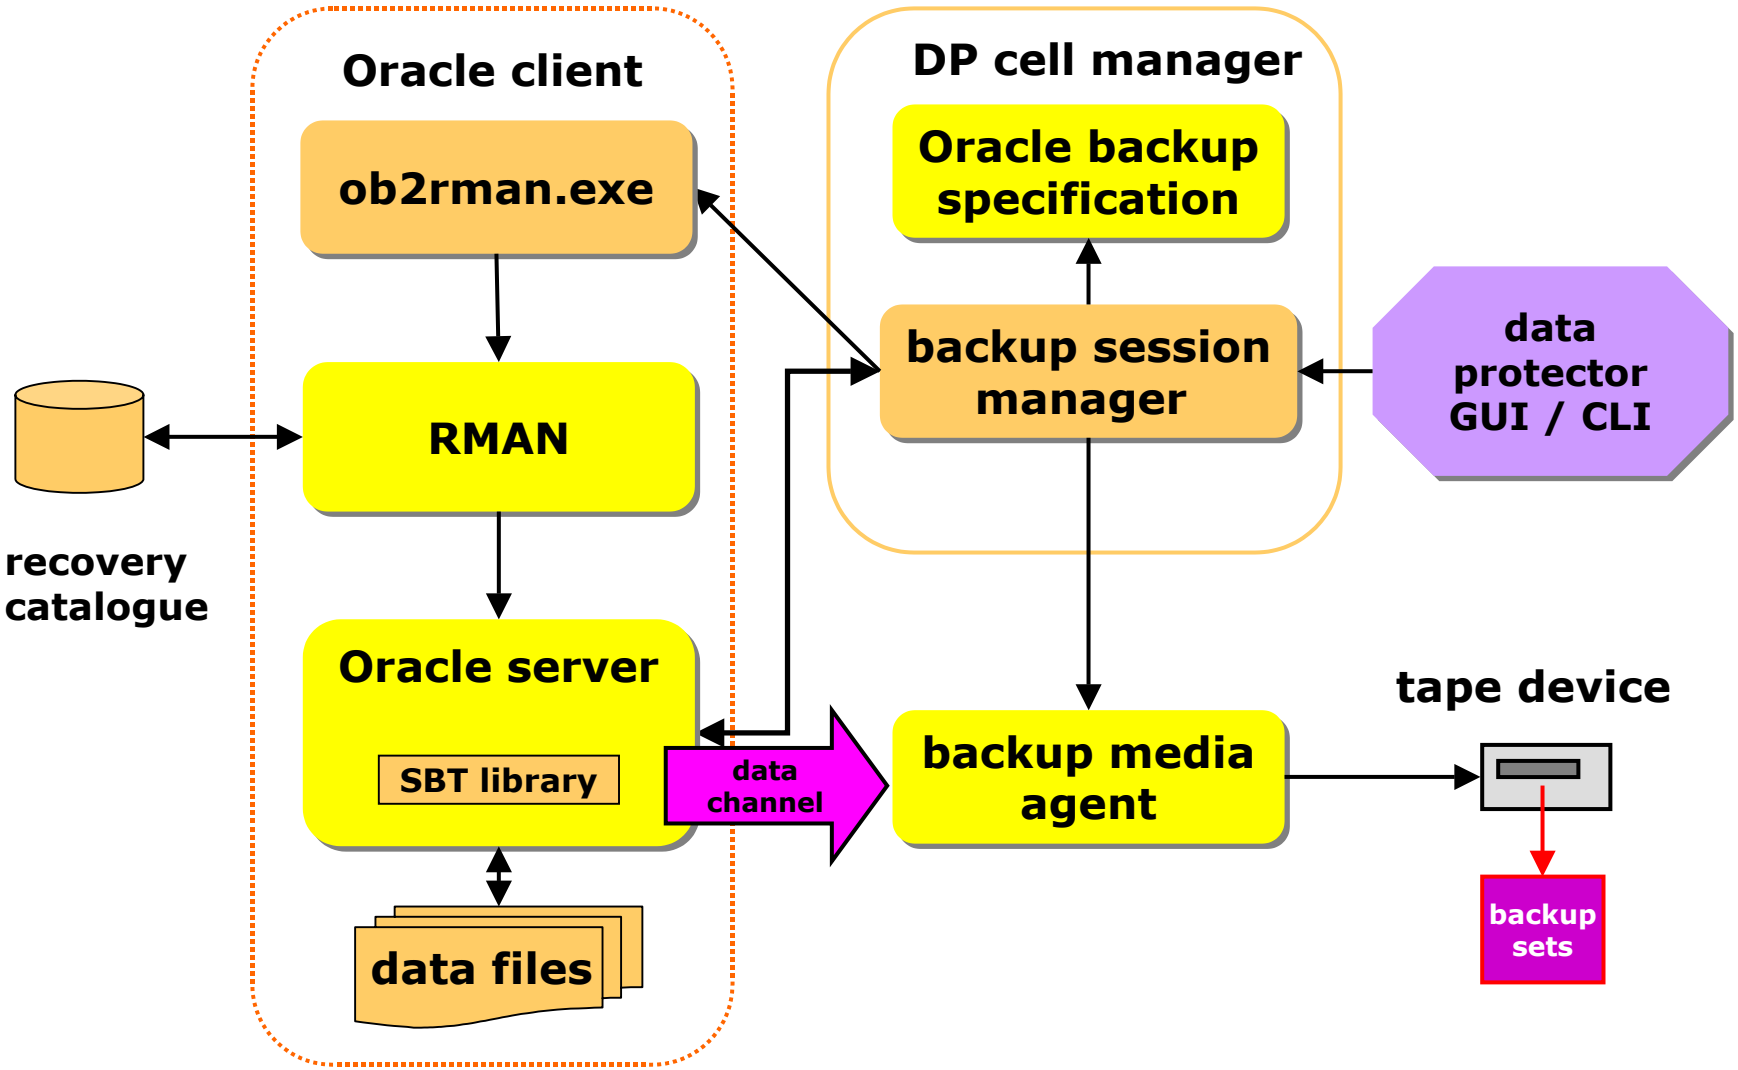

RMAN – Data Protector integration

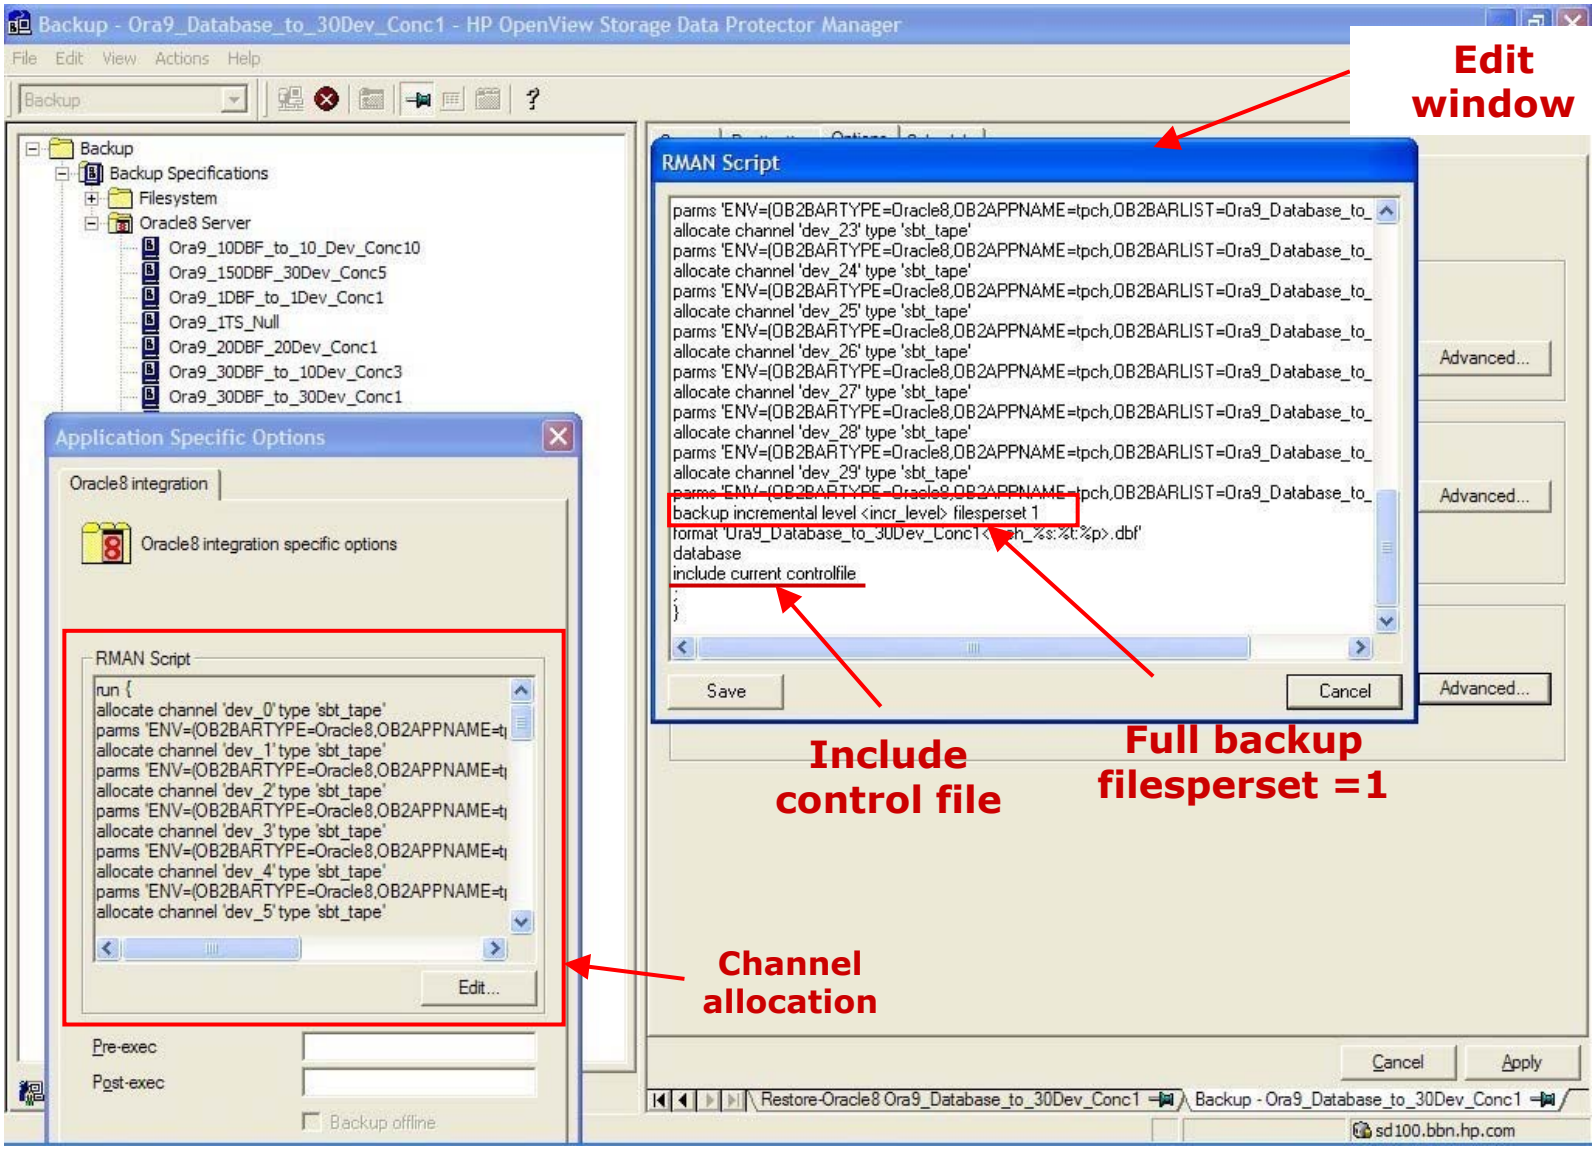

RMAN & Data Protector tuning

- concurrency
- filesperset
- blocksize
- maxopenfiles
- tape I/O slaves
- disk I/O slaves
- contolfile autobackup

Measuring the performance

- Glance/XL utility monitors CPU and disk utilization
- RMAN statistics (SQL script queries the recovery catalog)
- HP Data Protector statistics
- FC Switch throughput

HP Glance/XL results

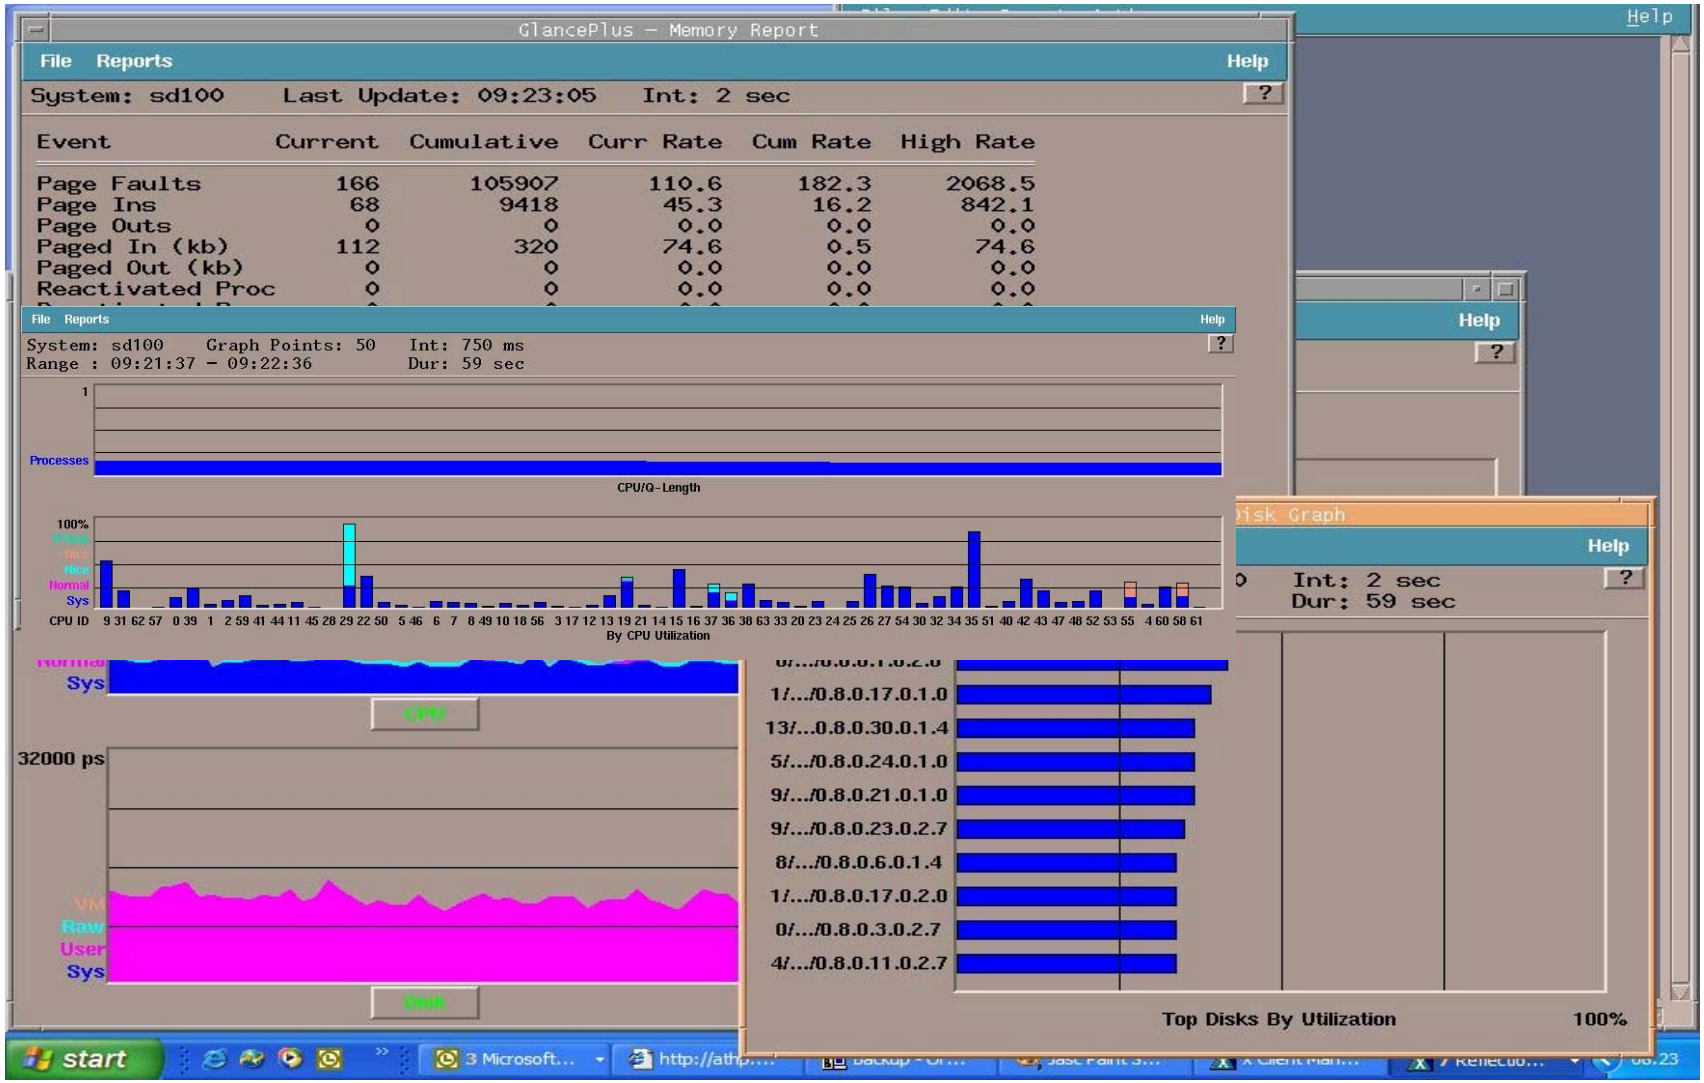

Results: 3TB/hr backup

# tape drives	# processors	CPU idle	Data protector parameters	RMAN settings	backup performance
16 in HP ESL9595 tape library	32	34%	concurrency=1 (1 RMAN channel/tape drive)	disk I/O slaves=16 filesperset=no set maxopenfiles=1 tapeblocksize=1MB backup_tape_io_ slaves=disabled controlfile autobackup=ON	**3.16 TB/hr** (including tape load/unload time) **3.62 TB/hr** (excluding tape load/unload time)

Results: backup (16 way CPU)

# tape drives	# processors	CPU idle	Data protector parameters	RMAN settings	backup performance
16 in HP ESL9595 tape library	16	15%	concurrency=1 (1 RMAN channel/tape drive)	disk I/O slaves=16 filesperset=no set maxopenfiles=1 tapeblocksize=256 KB backup_tape_io_ slaves=disabled controlfile autobackup=ON	**2.87 TB/hr** (includes tape load/unload time)

Results: restore

# tape drives	# processors	CPU idle	Data protector parameters	RMAN settings	backup performance
16 in HP ESL9595 tape library	32	45%	concurrency=1 (1 RMAN channel/tape drive)	disk I/O slaves=16 filesperset=no set maxopenfiles=1 tapeblocksize=1MB backup_tape_io_ slaves=enabled controlfile autobackup=ON	**1.23 TB/hr** (including tape load/unload time) **1.29 TB/hr** (excluding tape load/unload time)

- Restore times subsequently improved to 2.46TB/hr !

Database Recovery

- Use RMAN script to recover database
- Recovery Catalogue must be available (normally resides on another server)
- Media Manager must be available
- GUI based recovery is available from Data Protector 5.1 onwards

Future directions

- Oracle Real Application Clusters
- Serverless Backup
- Advanced Filesystems
- Ultrium 3 (late 2004) +60Mb/s

Lessons Learnt

- Backup performance determined in the main by disk subsystem performance. The 3:1 rule held true.
- Use SAME for best Oracle file system performance
- Storage Router configuration could be simplified.
- Memory usage was low 19GB Max (of 256GB)
- Could have been achieved with 16 Processors (15% idle)

Lessons Learnt

- In a high performance environments set 1 RMAN channel per tape drive, don't use external "multiplexing"
- Restore performance subsequently increased to 2.4 TB/HR – restore 40-60% of backup rate.
- Data Protector 5.1 integration with Oracle is now even better – restore GUI.
- Only Ultrium technology has adaptive tape speed to optimize performance and increase media life.

Conclusions

- 3TB/hour online database backup of a real life system is possible and cost effective with standard tape library configurations
- Ultrium tape has the growth path in future generations for even greater capacity and performance
- The Oracle9i Recovery Manager provides performance, flexible and reliable backup and recovery
- Tape based backup is still essential and viable

Reference material

- Oracle9i RMAN Backup & Recovery
  by RG Freeman & M Hart (ISBN 0-07-222662-5)
- HP Dataprotector Oracle 9i Integration Guide
- HP Dataprotector Unix installation Guide
  (www.hp.com)
- Oracle9i Recovery Manager User Guide and Reference
  (http://otn.oracle.com)
- Copies of HP 3TB/Hr Whitepaper available after this presentation

HP WORLD 2003
Solutions and Technology Conference & Expo

Interex, Encompass and HP bring you a powerful new HP World.

interex
shared knowledge • shared power

encompass
AN HP USER GROUP

hp
invent

Supplementary slides

HP WORLD 2003
Solutions and Technology Conference & Expo

RMAN – Data Protector integration

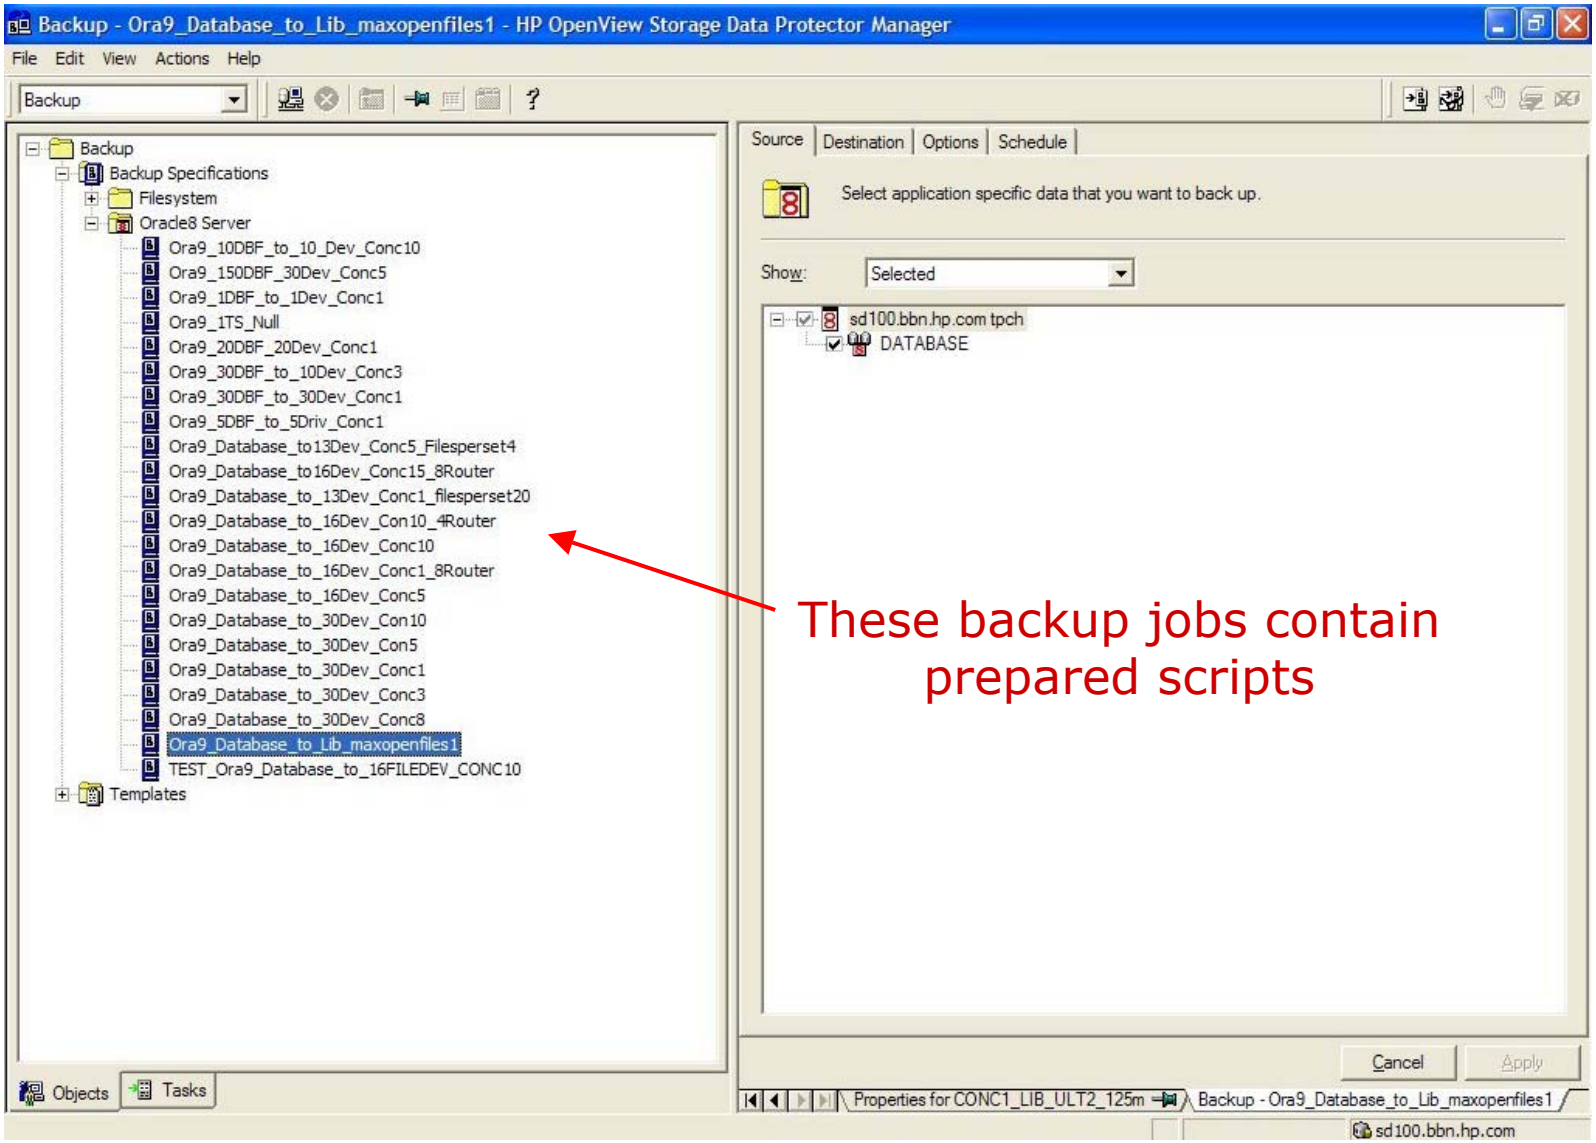

Filesystem vs Raw disk

- Oracle9i can use the filesystem supplied by the OS or what is known as raw disk. Raw disk is generally considered to be faster but is less manageable.
- The 3TB/hr project used filesystem

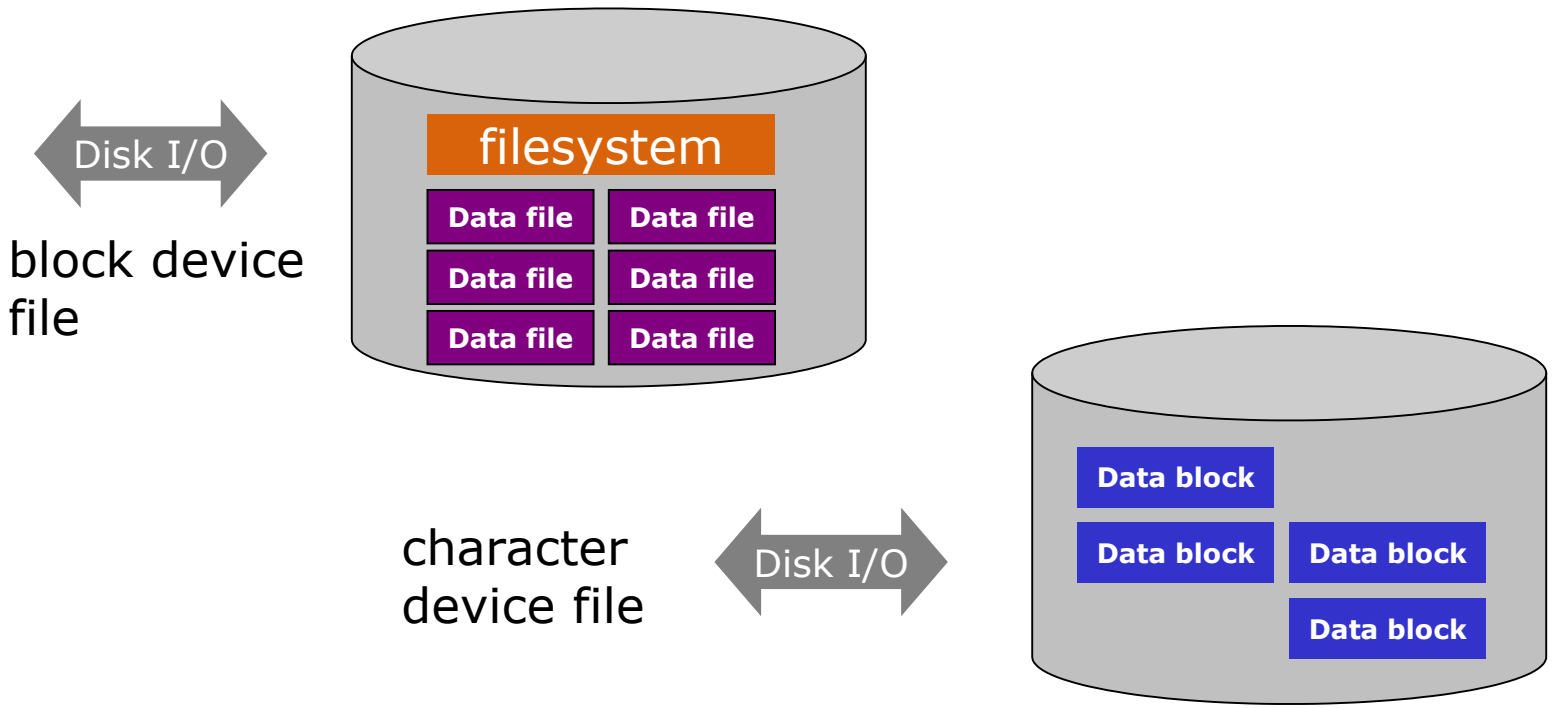

Multiplexing backup sets

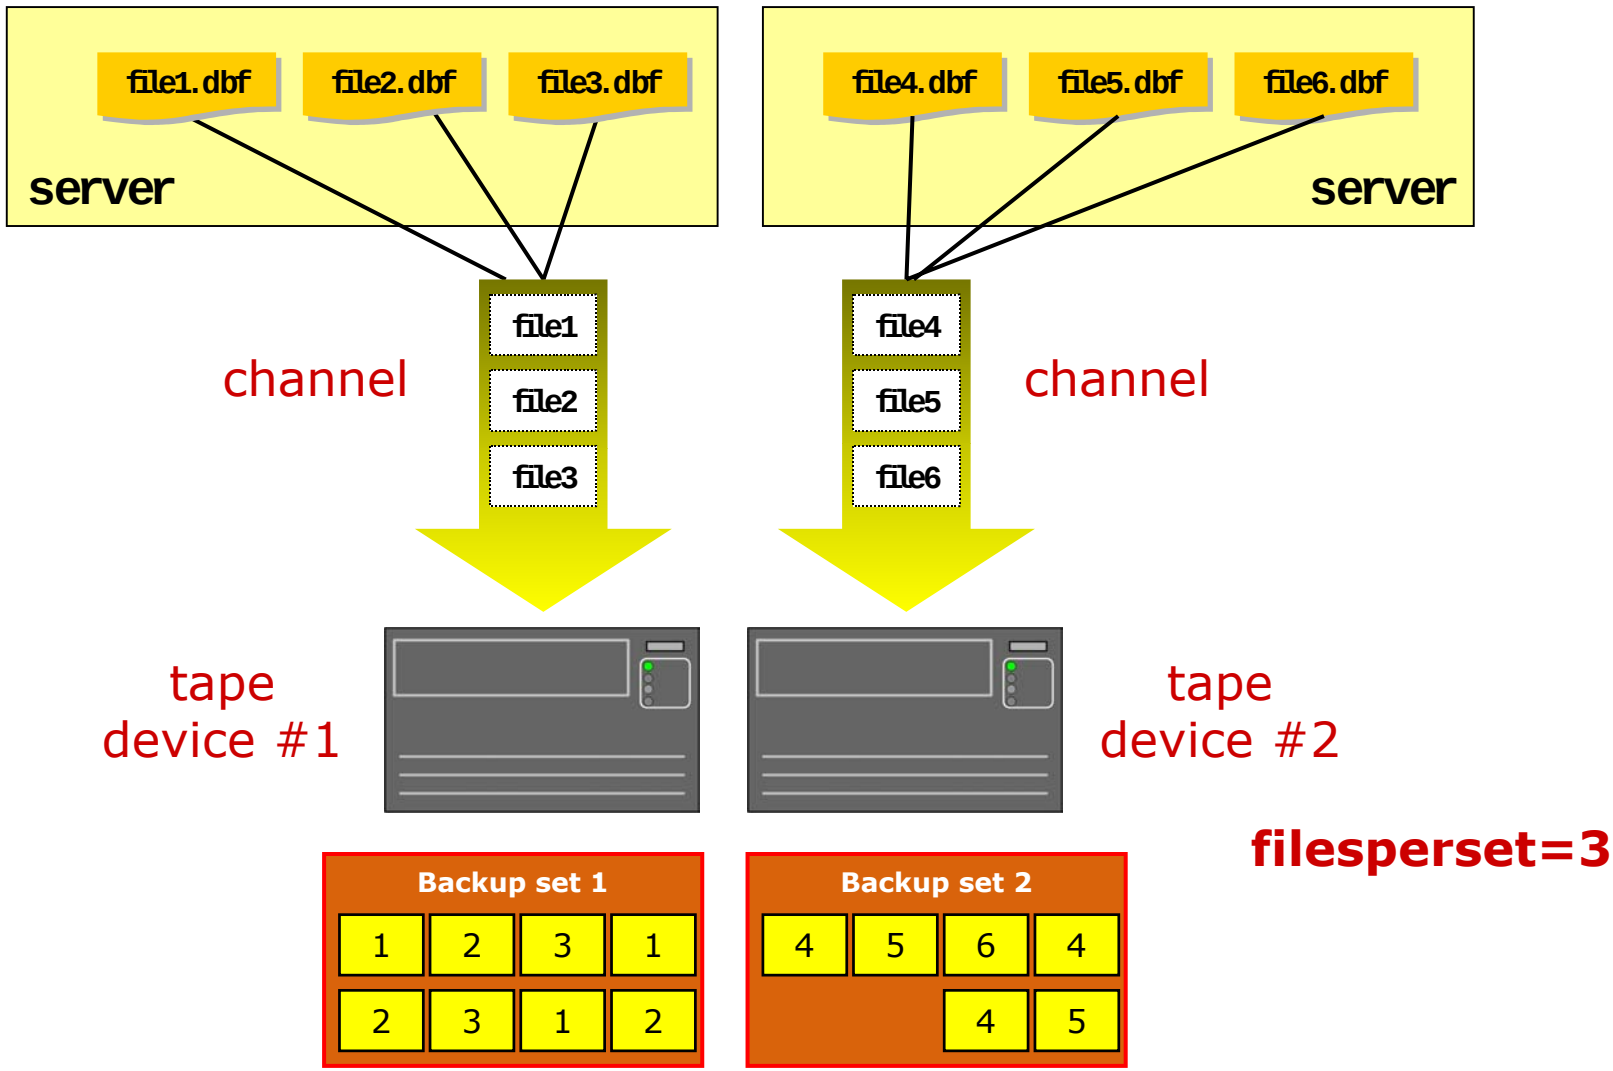

3TB/hr directory structures

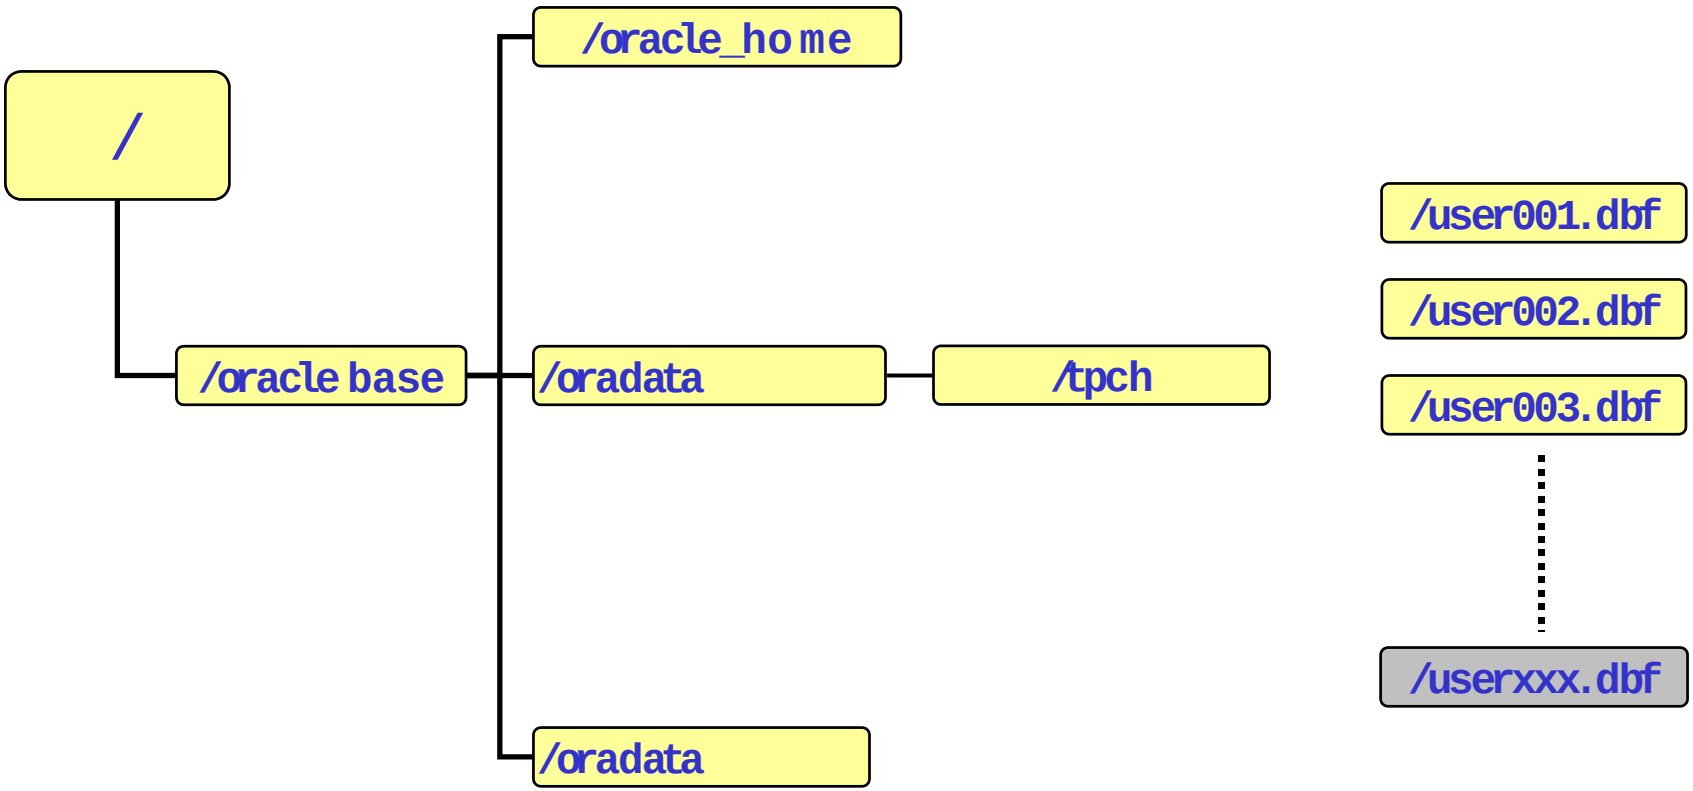

as per Oracle9i OFA model (Optimal Flexible Architecture)