Hot ASM Migration

Downtime encountered: 45 seconds
Migration host: poweredge
OS: RedHat LINUX AS 5
DB source: SID=orcl, 11.1.0.6, data stored on disk in /u01/app/oracle/oradata
DB destination: same db home, ASM diskgroup +DATA, recovery destination /u01/app/oracle/oradata/reco (normally this would be a separate diskgroup called +RECO or +FRA on a separate array from data files).
RMAN repository: same host as database, SID=test.

1. Install ASM RDBMS into a new oracle home. The slide below shows /u01/app/oracle/product/11g/db_1, but normally it will be in a different home like /u01/app/oracle/product/11g/asm_1.
2. Make sure there is no ASM installed (ps –ef | grep asm). Start ./dbca from this ASM home. Choose Configure ASM, create ASM instance.

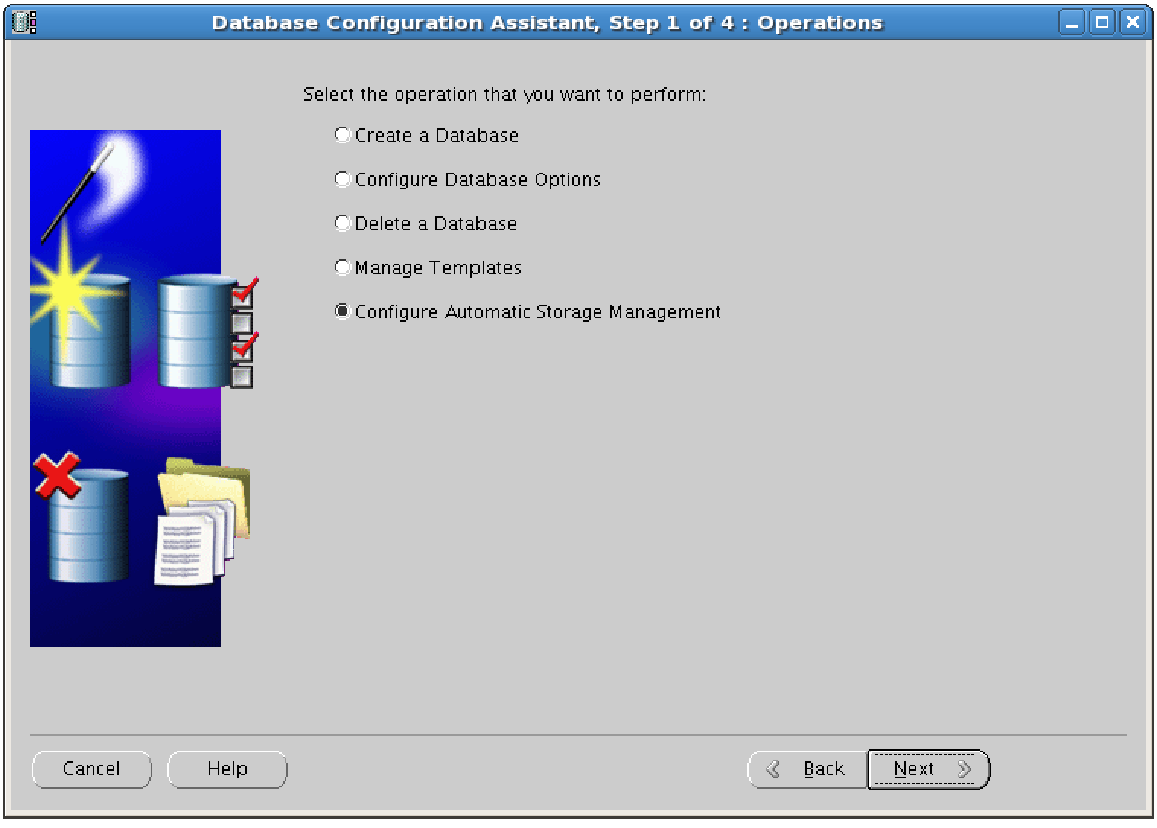

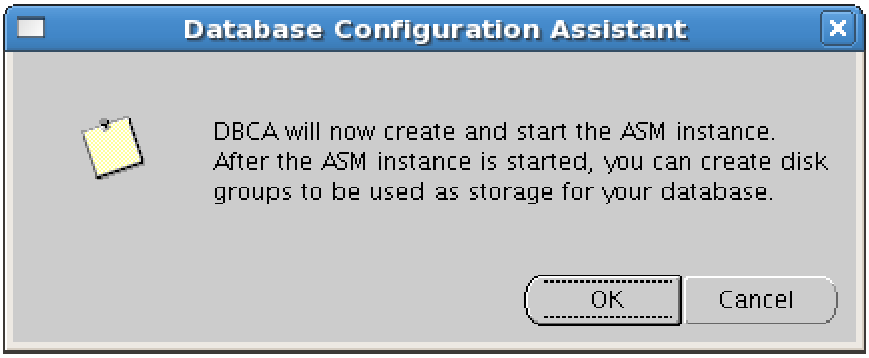

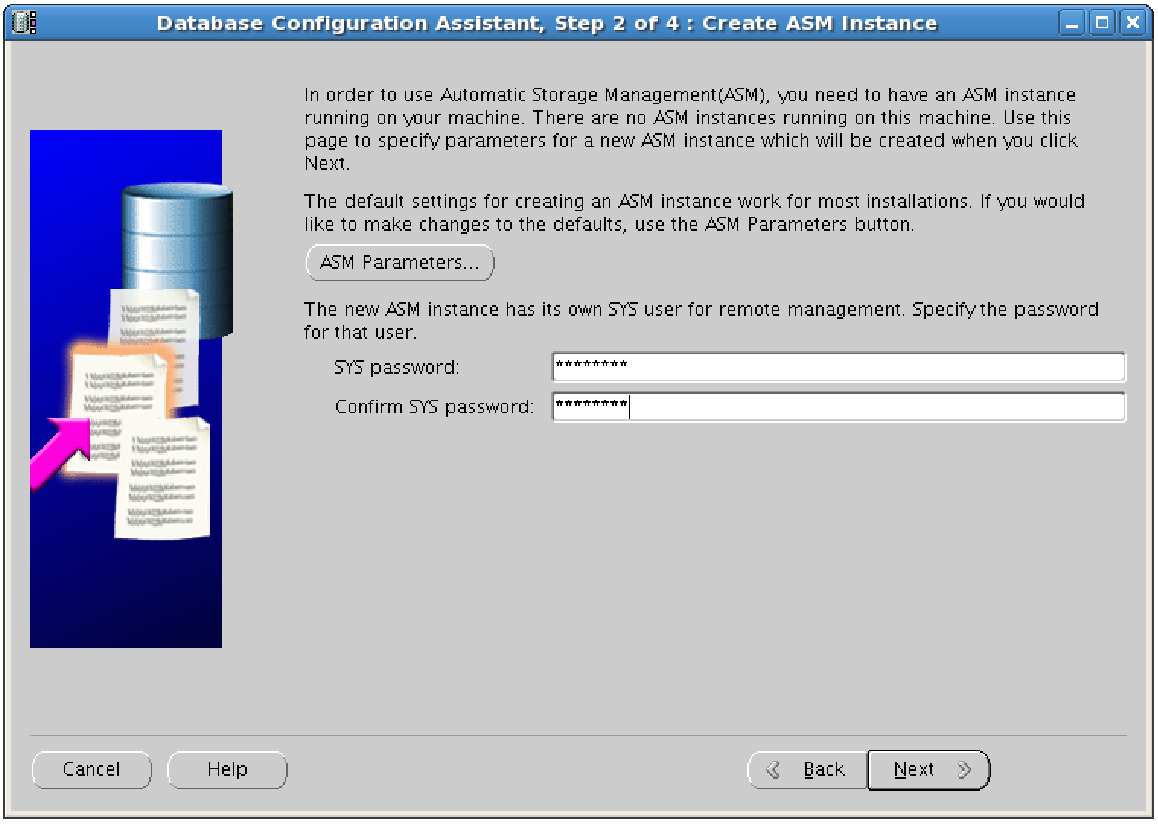

3. Log on as root and run required localconfig add script

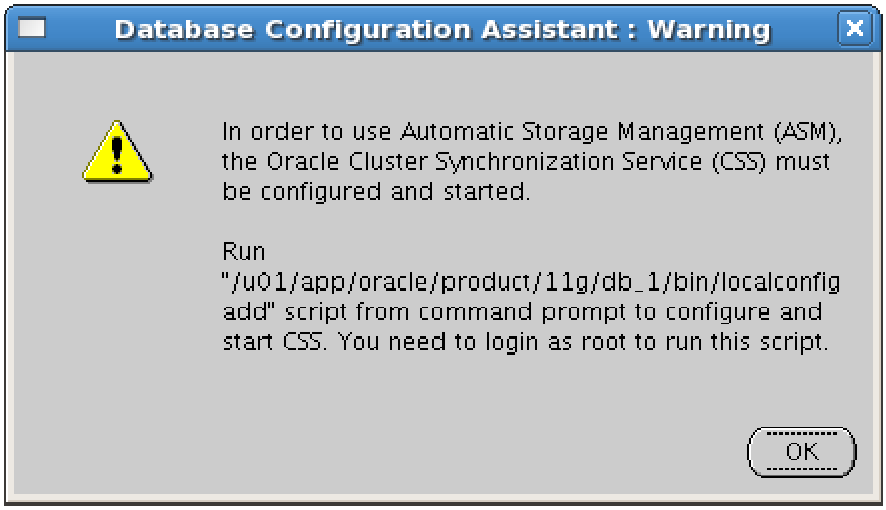

4. Localcofig add creates or adds CSS for ASM

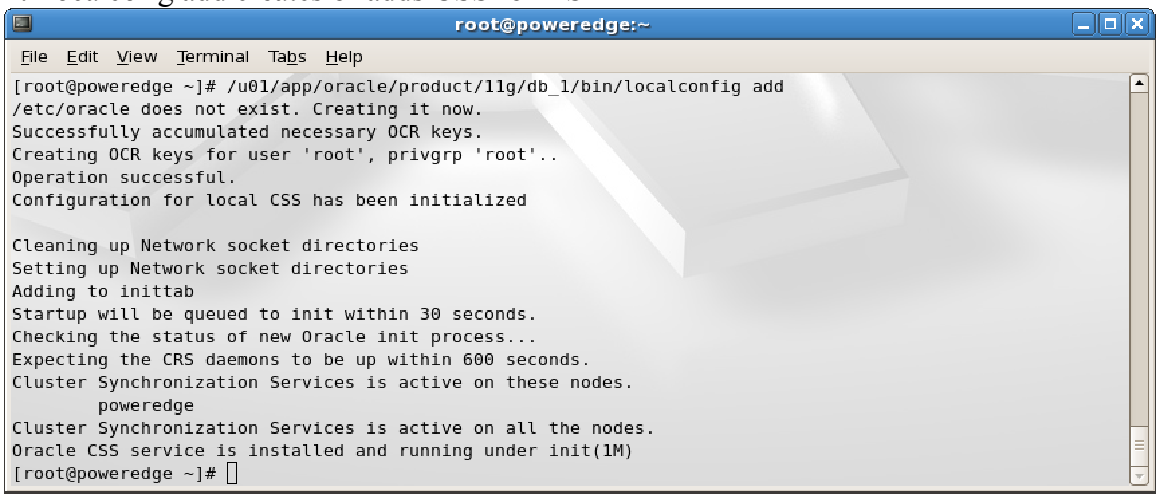

5. Next dbca will show available disks (on RedHat we configure ASMlib first, which is not shown here because this for HP/UX platform). Normally, we would create 2 disk groups: +DATA diskgroup (for datafiles, temp/undo, redo logs, controlfiles, etc.) and +RECO diskgroup (for backups, image copies, archivelog copies/backups, etc.). Optionally, the Flash Recovery Area diskroup can be stored on disk, as in this case. This is not recommended it was done to save space on this home test system.

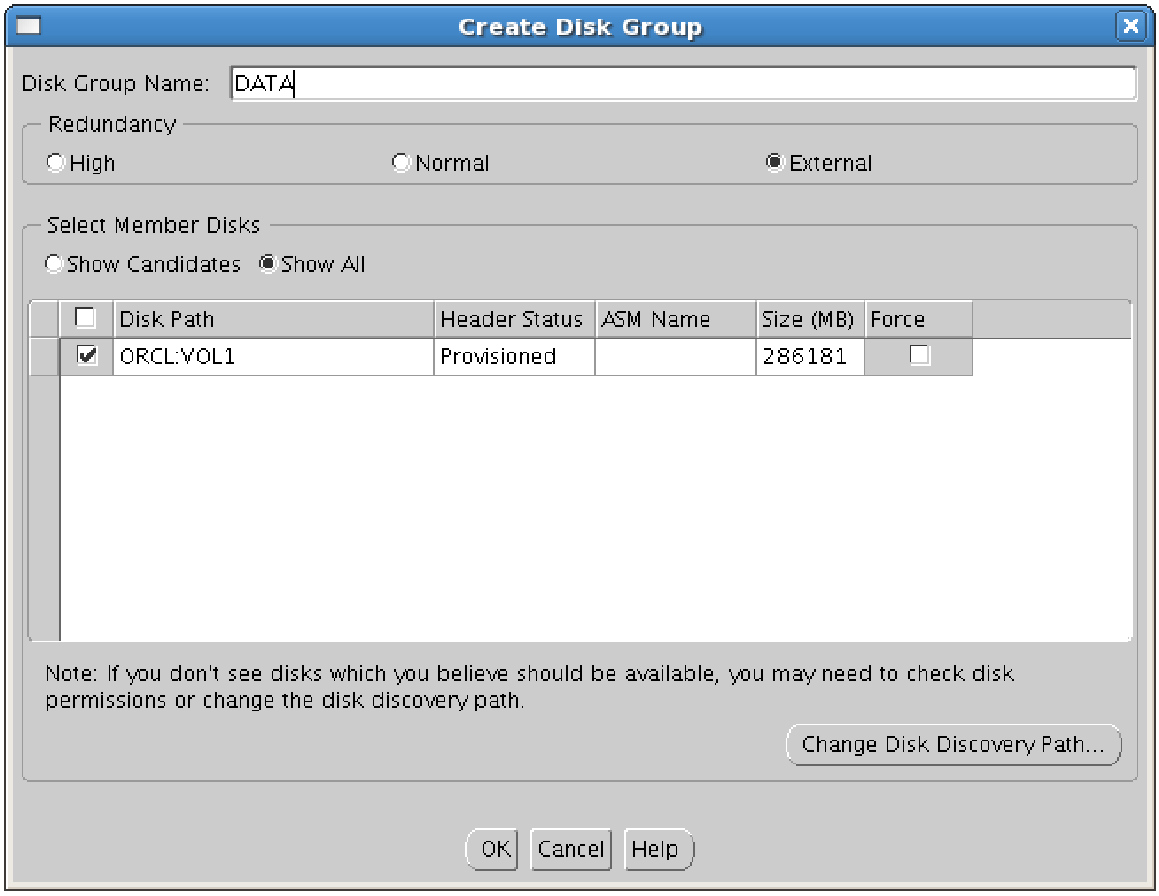
Create Disk Group
Disk Group Name: DATA
Redundancy
High
Normal
External
Select Member Disks
Show Candidates
Show All
Disk Path
Header Status
ASM Name
Size (MB)
Force
ORCL:VOL1
Provisioned
286181
Note: If you don't see disks which you believe should be available, you may need to check disk permissions or change the disk discovery path.
Change Disk Discovery Path...
OK
Cancel
Help

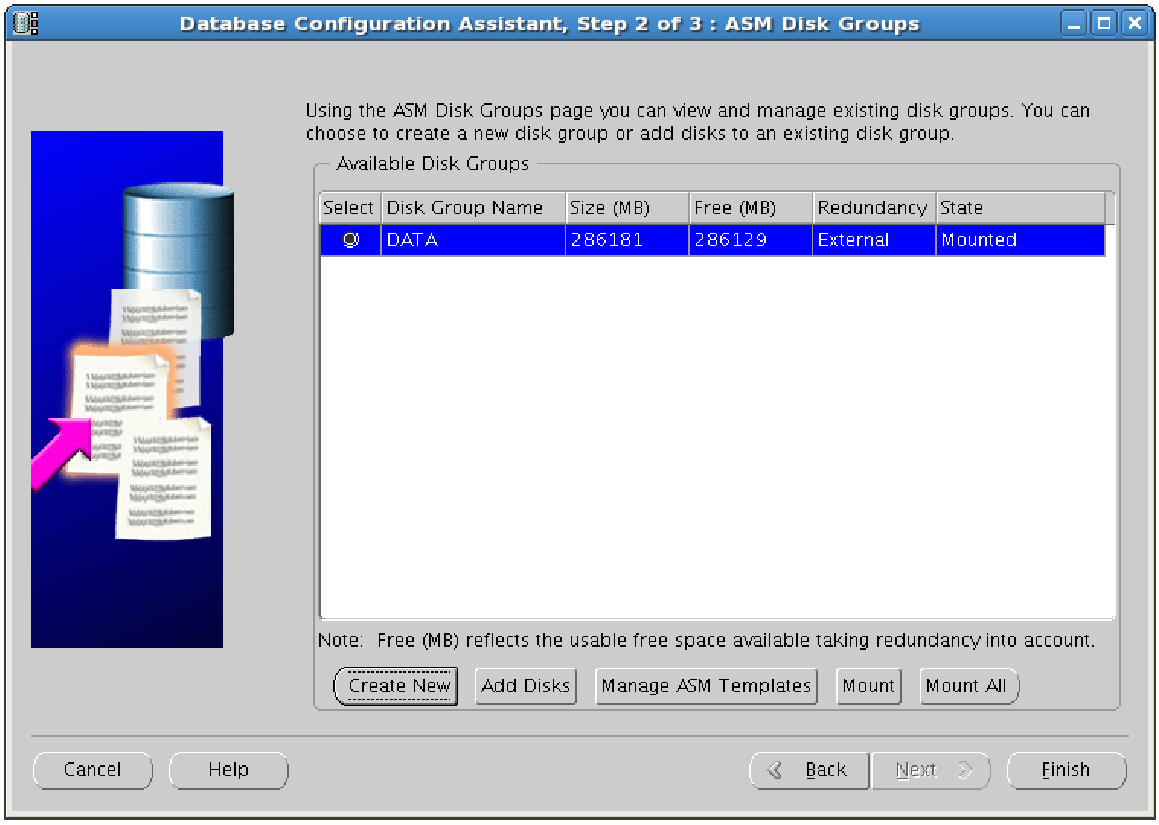

7. Now, after we have created the +DATA diskgroup, we have to instruct database to store newly created/restored datafiles in ASM instead of file system. We also set recovery file destination to /u01/app/oracle/oradata/reco. This is not recommended it should be another diskgroup in a separate storage array for redundancy reasons sized 2-3 times the size of +DATA.

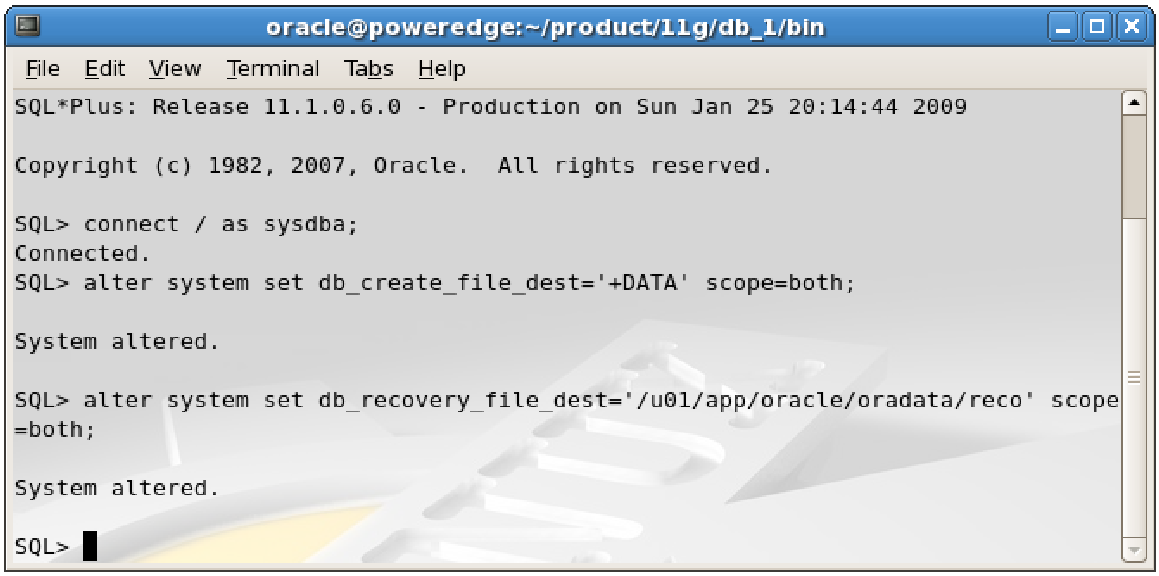

8. Next we connect to RMAN recovery catalog instance TEST (optional), target ORCL in RMAN command line interface and backup database.

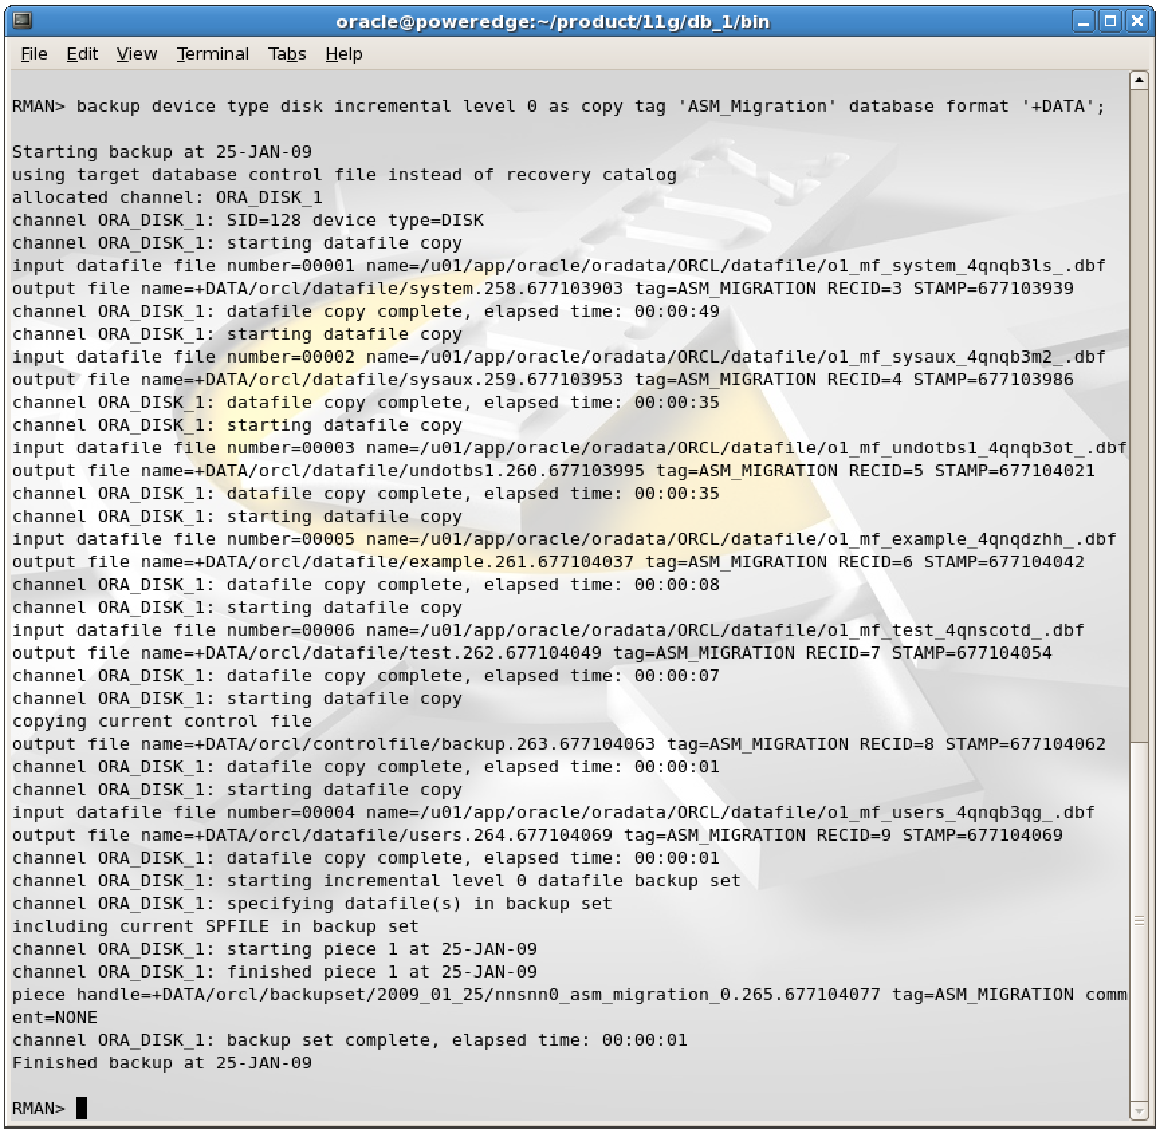

9. Next we query current online redo log file locations, add new redo log files stored in +DATA instead of file system, drop old ones until all redo logs are in the ASM.

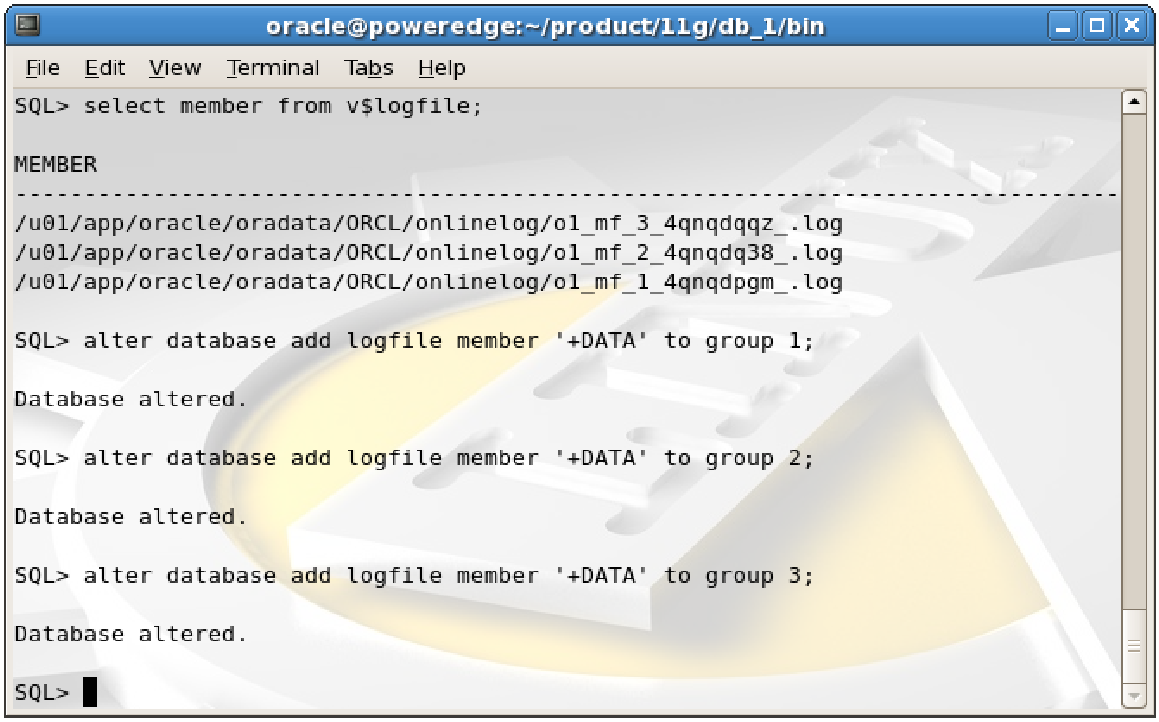

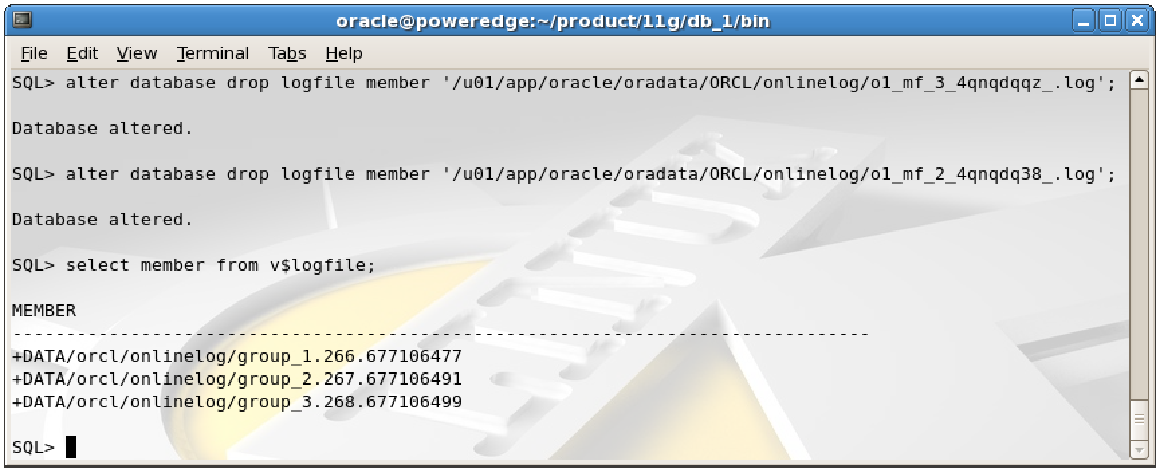

10. We do likewise with temp file

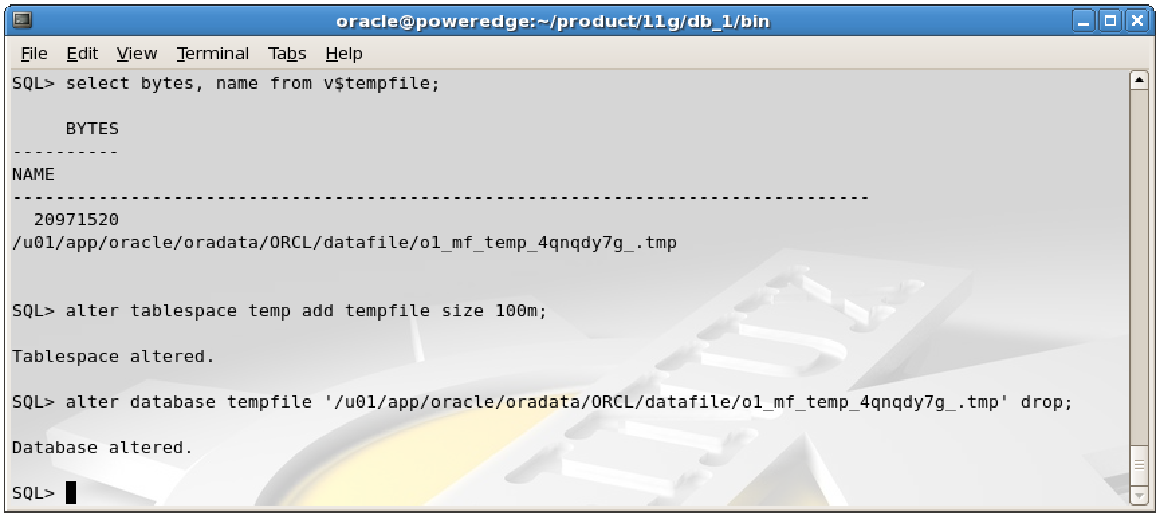

11. We then recover copy of database we created (backed up) above, since we made some changes.

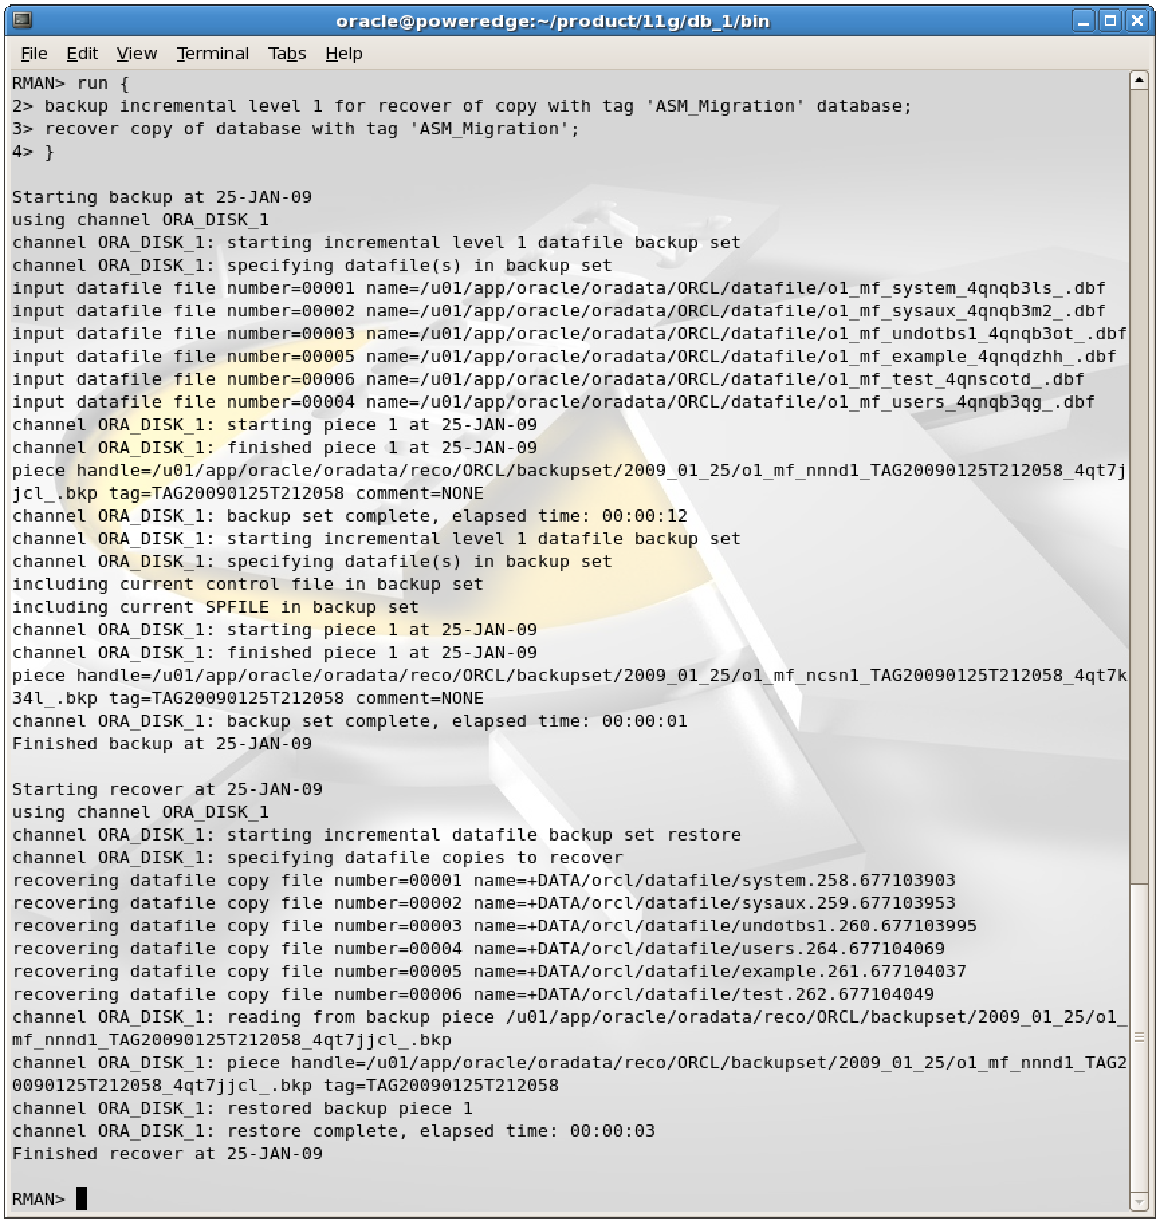

12. Next, we backup controlfile to +DATA ASM diskgroup and set new location of conntrolfile to +DATA.

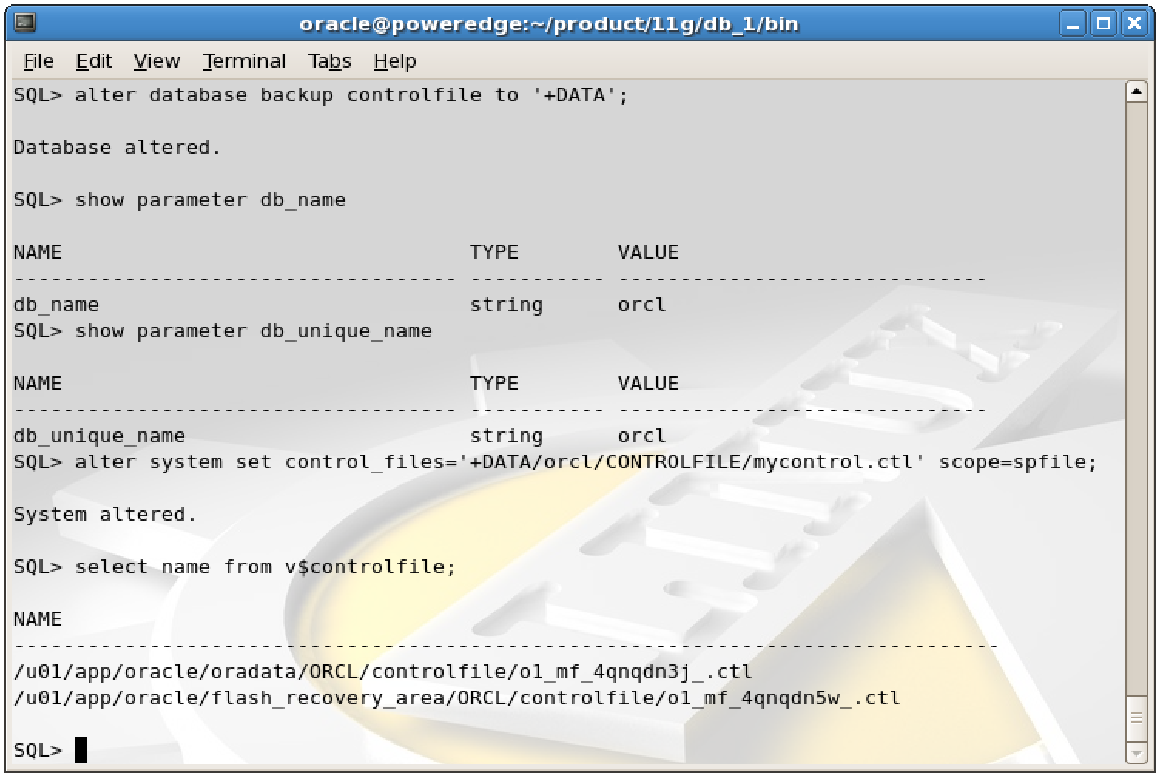

13. This change is not immediate as database has to be down for this to occur. Up till this point all of activities were executed with the database up and running. Its time to change that. At this point we are shutting down database to point to its copy in ASM and restore controlfile to the same location.

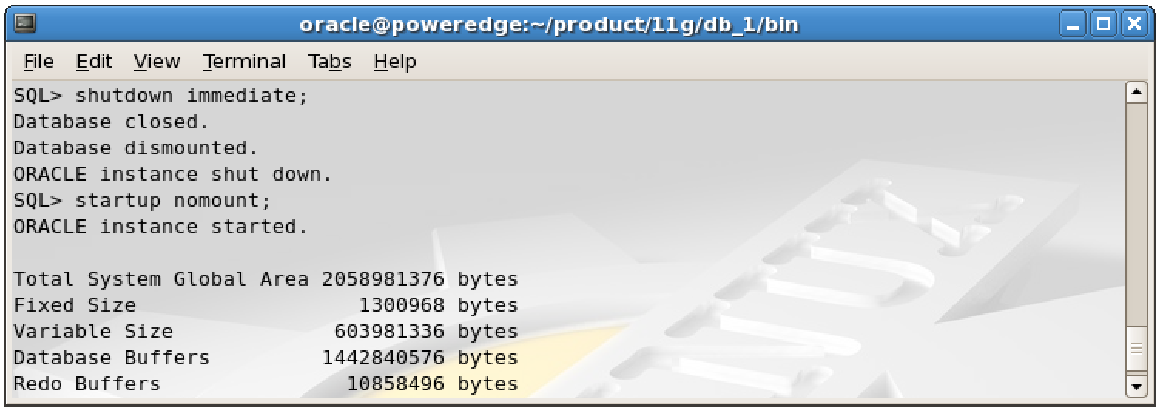

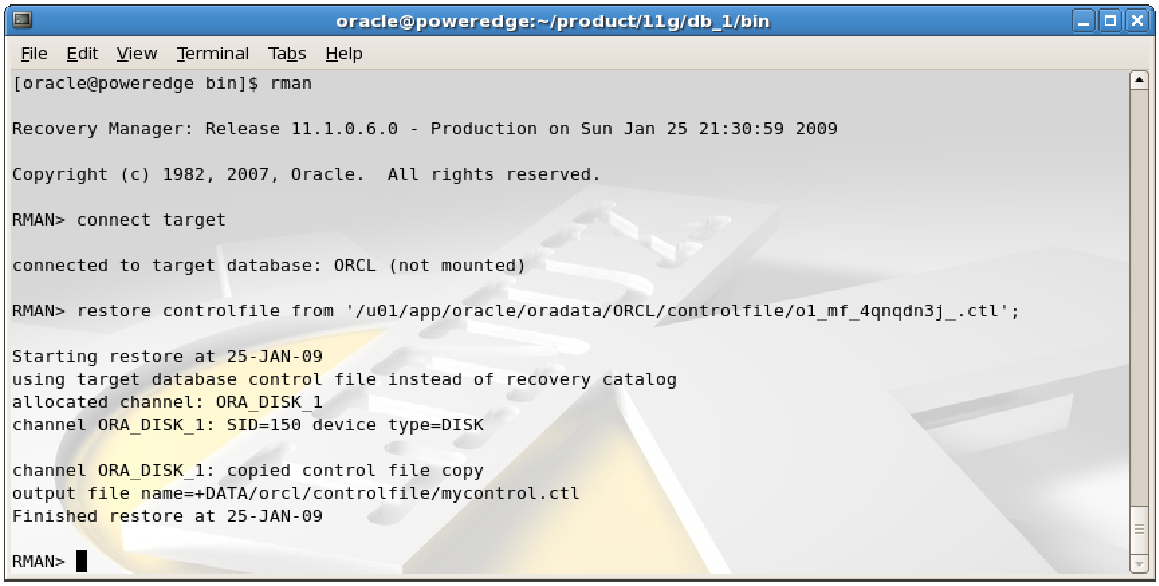

14. Now its time to switch to copy and recover database. As you can see we are switching from /u01/app/oracle/oradata after only a half a minute this database spent in downtime. This time will be longer for larger databases (this one is only 2-3GB), but should not exceed 15 mins even for a TB database.

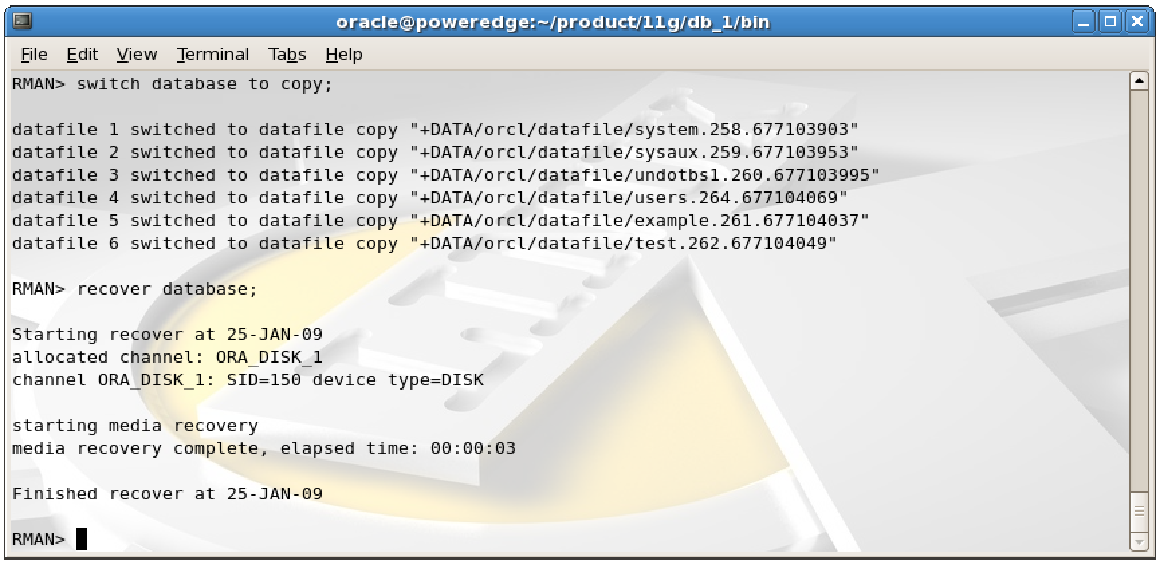

15. Now we are completed database hot migration to ASM. To make sure everything is in ASM lets run the following query:

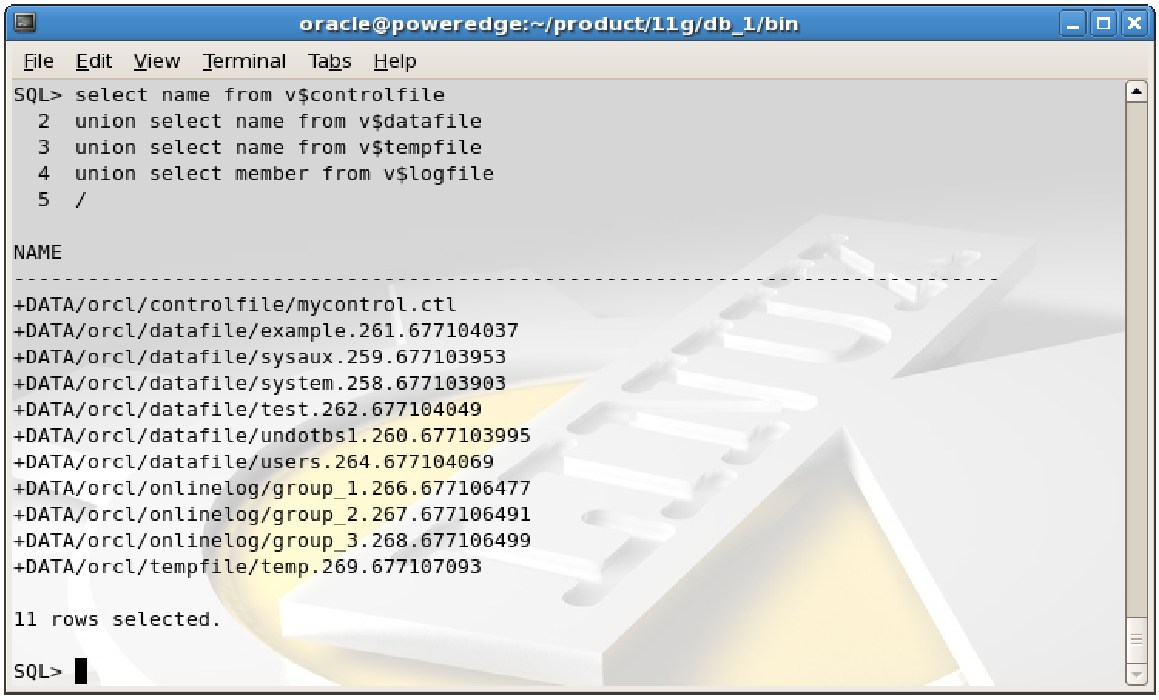

16. As you can see, all data structures are in ASM. In production, however, you would also make sure that controlfiles are multiplexed (there is only one above) and existing backup sets are moved into the newly created +RECO, which we did not create in this pre-test.