

# Exadata Patching

## From Customer Experience to New Features

Daniele Massimi  
Senior Consultant




BASEL BERN BRUGG LAUSANNE ZUERICH DUESSELDORF FRANKFURT A.M. FREIBURG I.BR. HAMBURG MUNICH STUTTGART VIENNA

## ■ Unser Unternehmen

Trivadis ist **führend bei der IT-Beratung, der Systemintegration, dem Solution-Engineering** und der Erbringung von **IT-Services** mit Fokussierung auf **ORACLE®** und  **Microsoft** Technologien im D-A-CH-Raum.

Unsere Leistungen erbringen wir aus den strategischen Geschäftsfeldern:



Trivadis Services übernimmt den korrespondierenden Betrieb Ihrer IT Systeme.

## ■ Mit über 600 IT- und Fachexperten bei Ihnen vor Ort



12 Trivadis Niederlassungen mit über 600 Mitarbeitenden

200 Service Level Agreements

Mehr als 4'000 Trainingsteilnehmer

Forschungs- und Entwicklungsbudget: CHF 5.0 Mio. / EUR 4.0 Mio.

Finanziell unabhängig und nachhaltig profitabel


Erfahrung aus mehr als 1'900 Projekten pro Jahr bei über 800 Kunden

# ■ AGENDA

1. Overview
2. Preparation
3. Exadata Patching
4. New Features Storage Server SW 11.2.3.3 + 12.1.1.1

# Overview


## ■ Exadata Overview (Example X3-2)



2 Infiniband Switch	<ul style="list-style-type: none"> <li>● 36 port QDR (40 Gb/sec)</li> </ul>
1 Cisco Catalyst 4948	<ul style="list-style-type: none"> <li>● 48 1Gb Ports</li> </ul>
KVM Switch	<b>Not available on X3-2</b>
2 Database Server X3-2	<ul style="list-style-type: none"> <li>● 2 x Eight-Core Intel® Xeon® E5-2690 Processors (2.9 GHz)</li> <li>● 256GB Memory</li> <li>● Disk Controller HBA with 512MB Battery Backed Write Cache</li> <li>● 4 x 300 GB 10,000 RPM Disks</li> <li>● 2 x QDR (40Gb/s) Ports</li> <li>● 4 x 1/10 Gb Ethernet Ports (copper)</li> <li>● 2 x 10 Gb Ethernet Ports (optical)</li> <li>● 1 x ILOM Ethernet Port</li> <li>● 2 x Redundant Hot-Swappable Power Supplies</li> </ul>
3 Storage Server X3-2	<ul style="list-style-type: none"> <li>● 36 CPU cores for SQL processing (<b>18 cores enabled</b>)</li> <li>● 12 PCI flash cards (<b>6 cards enabled</b>) with 2.4 TB Exadata Smart Flash Cache</li> <li>● 18 x 3 TB 7,200 RPM High Capacity disks (<b>6 disks per storage server enabled</b>)</li> </ul>

# ■ Customer Exadata Environment

**Quarter Rack V2 (Test):**  
With MAA Architecture -  
RAC and Data Guard with  
several Database



**Half Rack X2 (Production):**  
With MAA Architecture -  
RAC and Data Guard with  
several Database

**Patching Approach:**  
1st Rack then 2nd Rack

# ■ Exadata Patching Overview

- Exadata has several layers that **requires** software maintenance
  - Storage Server
  - Database Server including GI and RDBMS Software
  - InfiniBand Switch
  - Ethernet Switch
  - PDU's
- It is **recommended** to maintain at least **Infiniband, Storage Server** and **Database Server** including GI and RDBMS.
- Patch in regular interval
- Evaluate patches in a test environment, preferably having the same HW and software components as the production environment.



## ■ Exadata Patching – Number of Components

Type	X2-2 Full X3-2 Full X4-2 Full	X2-2 Half X3-2 Half X4-2 Half	X2-2 Quarter/Eighth X3-2 Quarter/Eighth X4-2 Quarter/Eighth
Storage Server	14	7	3
Database Server	8	4	2
Ethernet Switch	1	1	1
Infiniband Leaf Switch	2	2	2
Infiniband Spine Switch	1 (X4-2, no more included)	1 (X4-2, no more included)	0
PDU	2	2	2
KVM	1 (X3, X4 no more included)	1 (X3, X4 no more included)	1 (X3, X4 no more included)

# ■ Exadata Patching Starting Point and Strategies

- Starting Point is MOS Note 888828.1
- Strategies:
  - QFSDP (Quarterly Full Stack Download Patch)
  - Individual Patch → (without important reason, use always recommended Version)
    - Infiniband
    - Storage Server
    - Database Server
    - Grid Infrastructure
    - RDBMS
- Check the compatibility and requirements for your Exadata DBM

## ■ Which component we will patch as first

- Basically all component can installed independently from each other, unless...

Some prerequisites must be fulfilled ! → RTFM (read the «full» manual) 😊

Best Practice:

1. Infiniband
2. Storage Server
3. Database Server
4. GI and RDBMS
5. PDU
6. Ethernet Switch and KVM (Firmware are not delivered from Oracle)

## ■ Patching Tools (1/2)

- dcli
  - Executes command in distributed mode
  - Commands runs in parallel

```
$ dcli -g cell_group -l root 'imageinfo'  
$ ssh-keygen -t rsa  
$ dcli -g cell_group -l root -k
```

- patchmgr
  - Utility for patching the Storage Servers
  - Group File like in dcli are needed (e.g. cell\_group)

## ■ Patching Tools (2/2)

- `dbnodeupdate.sh`
  - Utility for Patching Database Node (OS Update)
  - It is possible to upgrade, rollback, backup, relink Oracle Component on Computing Node
- `oplan`
  - Utility that facilitates patching of GI and RDBMS
  - Generate patching installation and rollback instruction
- `opatch`
  - Utility for installing GI and RDBMS patches
  - Normally using with «auto» option as root user

# ■ Preparation of Fallback Strategies

- Infiniband
  - Make sure that the old Firmware are available and accessible
- Storage Server
  - Make sure the old Image is on the inactive Partition and rollback is possible
- Database Server
  - Create a backup of root and boot partition
- Grid Infrastructure
  - Create a filesystem backup or TAR-Balls of GI Home
  - Backup OCR and OLR Dumps
- RDBMS
  - Create a filesystem backup

# ■ Backing Up InfiniBand Switch Settings

- ILOM console has backup/recover function from 1.1.3-2 upwards
- For earlier releases:
  1. Make copies of the following files:
    - /etc/sysconfig/network-scripts/ifcfg-eth0, /etc/resolv.conf, /etc/ntp.conf, /etc/localtime, /etc/opensm/opensm.conf, /etc/sysconfig/network
  2. Run the nm2version command and save the output

```
NM2-72p version: 0.1.0-1
Build time: Aug 24 2009 16:41:03
FPGA version: 0x94
ComExpress info:
Board Name: "NOW1"
Manufacturer Name: "JUMP"
Manufacturing Date: 2009.02.19
Last Repair Date: 1980.01.01
Serial Number: "NCD2S0240"
Hardware Revision: 0x0100
Firmware Revision: 0x0102
Jida Revision: 0x0103
Feature Number: 0x0001
```

## ■ Restoring InfiniBand Switch

1. Apply the firmware
2. Restore the settings
  - Recommended: Download current version of InfiniBand Switch software before patching



# ■ Backing Up Storage Server

- No need to backup storage servers OS before patching
- Three structures
  - Active partition
  - Inactive partition
  - CELLBOOT USB flash drive contains current image and last good configuration
- Patches
  - In-Partition Patches
    - Patches are applied to the active partition
    - Changed files are saved for rollback
  - Out-of-Partition Upgrades
    - Patch is applied to inactive partition
    - After success this partition becomes active
    - To rollback, make the old partition active
- During patching if something goes wrong Oracle restores the previous image
- Contact Oracle support for any doubts during restore

## ■ Fallback Check - Storage Server

- Check if Storage Server has valid Storage Server Image on inactive Partition

```
[root@dm01cel01 ~]# imageinfo

Kernel version: 2.6.39-400.128.1.el5uek #1 SMP Wed Oct 23 15:32:53 PDT 2013 x86_64
Cell version: OSS_12.1.1.1.0_LINUX.X64_131219
Cell rpm version: cell-12.1.1.1.0_LINUX.X64_131219-1

Active image version: 12.1.1.1.0.131219
Active image activated: 2014-02-06 10:33:40 +0100
Active image status: success
Active system partition on device: /dev/md6
Active software partition on device: /dev/md8

In partition rollback: Impossible

Cell boot usb partition: /dev/sdml
Cell boot usb version: 12.1.1.1.0.131219

Inactive image version: 11.2.3.2.1.130109
Inactive image activated: 2014-02-05 16:00:49 +0100
Inactive image status: success
Inactive system partition on device: /dev/md5
Inactive software partition on device: /dev/md7

Boot area has rollback archive for the version: 11.2.3.2.1.130109
Rollback to the inactive partitions: Possible
```

# ■ Backing up Linux Database Server – Method 1

- Using `dbnodeupdate.sh`
  - Backup recovery process on MOS Doc Id 1553103.1, Patch 16486998

```
[root@dm01db01-admin oracle.SupportTools]# ./dbnodeupdate.sh -u -l p17997668_121110_Linux-x86-64.zip
.
.
.
Active Image version      : 11.2.3.2.1.130109
Active Kernel version     : 2.6.32-400.11.1.el5uek
Active LVM Name           : /dev/mapper/VGExaDb-LVDbSys1
Inactive Image version    : 11.2.3.2.1.130109
Inactive LVM Name        : /dev/mapper/VGExaDb-LVDbSys2
Current user id           : root
Action                    : upgrade
Upgrading to              : 12.1.1.1.0.131219
Baseurl                   : file:///var/www/html/yum/unknown/EXADATA/dbserver/140214181859/x86_64/ (iso)
Iso file                  : /u01/app/oracle/stage.140214181859/121110_base_repo_131219.iso
Create a backup           : Yes
Shutdown stack           : No (Currently stack is down)
RPM exclusion list       : Not in use (add rpms to /etc/exadata/yum/exclusion.lst and restart dbnodeupdate.sh)
RPM obsolete list        : /etc/exadata/yum/obsolete.lst (lists rpms to be removed by the update)
                          : RPM obsolete list is extracted from exadata-sun-computenode-12.1.1.1.0.131219-1.x86_64.rpm
Logfile                  : /var/log/cellos/dbnodeupdate.log (runid: 140214181859)
Diagfile                 : /var/log/cellos/dbnodeupdate.140214181859.diag
Server model             : SUN FIRE X4170 M3
dbnodeupdate.sh rel.     : 2.44 (always check MOS 1553103.1 for the latest release of dbnodeupdate)
Automatic checks incl.   : Known Issue - Database Server upgrades may hit network routing issues after the upgrade
                          : Issue - Yum rolling update requires fix for 11768055 when Grid Infrastructure is below 11.2.0.2
BP12
Note                     : After upgrading and rebooting run './dbnodeupdate.sh -c' to finish post steps
```

# ■ Backing up Linux Database Server 1/2 - Method 2

- Backup using LVM snapshot-based backup on NFS Share
- Using `dbserver_backup.sh` – MOS 1556257.1, Patch 13741363
- Backup recovery process can be found at Oracle Exadata Database Machine Owner's Guide

## Create Backup Directory:

```
[root@bdmlldb01 ~]# mkdir -p /root/tar  
[root@bdmlldb01 ~]# mount -t nfs -o rw,intr,soft,proto=tcp,nolock ip_address:/dbutil2/ /root/tar
```

## Create LVM Snapshot for root partition and give them a Name:

```
[root@bdmlldb01 ~]# lvcreate -L1G -s -n root_snap /dev/VGExaDb/LVDbSys1  
[root@bdmlldb01 ~]# e2label /dev/VGExaDb/root_snap DBSYS_SNAP
```

## Mount the Snapshot:

```
[root@bdmlldb01 ~]# mkdir /root/mnt  
[root@bdmlldb01 ~]# mount -t ext3 /dev/VGExaDb/root_snap /root/mnt
```

## Optional - Create LVM Snapshot for Oracle partition and give them a Name:

```
[root@bdmlldb01 ~]# lvcreate -L5G -s -n u01_snap /dev/VGExaDb/LVDbOral  
[root@bdmlldb01 ~]# e2label /dev/VGExaDb/u01_snap DBORA_SNAP
```

## Optional - Mount the Snapshot:

```
[root@bdmlldb01 ~]# mkdir -p /root/mnt/u01  
[root@bdmlldb01 ~]# mount /dev/VGExaDb/u01_snap /root/mnt/u01 -t ext3
```

## ■ Backing up Linux Database Server 2/2 – Method 2

Change to root partition mountpoint:

```
[root@bdm1db01 ~]# cd /root/mnt
```

Create backup of root and boot partition:

```
[root@bdm1db01 ~]# tar -pjcvf /root/tar/mybackup.tar.bz2 * /boot --exclude  
tar/mybackup.tar.bz2 > /tmp/backup_tar.stdout 2> /tmp/backup_tar.stderr
```

Remove all mountpoints and LVM Snapshots:

```
[root@bdm1db01 ~]# cd /  
[root@bdm1db01 ~]# umount /root/mnt  
[root@bdm1db01 ~]# /bin/rm -rf /root/mnt  
[root@bdm1db01 ~]# lvremove /dev/VGExaDb/root_snap  
  
[root@bdm1db01 ~]# umount /root/tar
```

# ■ Recovering Linux Database Server – Method 1

- Recovery with dbnodeupdate.sh
  - Very simple 😊
  - Using Parameter `-r` for Rollback

```
[root@dm01db02-admin oracle.SupportTools]# ./dbnodeupdate.sh -r
```

# ■ Recovering Linux Database Server – Method 2 (1/2)

- Recovery with non-modified partitions
  1. Boot the database server from diagnostic.iso (Available at /opt/oracle.SupportTools/diagnostics.iso) from ILOM console
  2. Follow the instructions

```
Choose from following by typing letter in '()':
(e)nter interactive diagnostics shell. Must use credentials from Oracle
support to login (reboot or power cycle to exit the shell),
(r)estore system from NFS backup archive,
Select:r
Are you sure (y/n) [n]:y
The backup file could be created either from LVM or non-LVM based compute node
versions below 11.2.1.3.1 and 11.2.2.1.0 or higher do not support LVM based partitioning
use LVM based scheme(y/n):y
Enter path to the backup file on the NFS server in format:
<ip_address_of_the_NFS_share>:/<path>/<archive_file>
For example, 10.10.10.10:/export/operating_system.tar.bz2
NFS line:nfs_ip:/export/mybackup.tar.bz2
IP Address of this host:<IP address of the DB host>
Netmask of this host:<netmask for the above IP address>
Default gateway:<Gateway for the above IP address>
```

## ■ Recovering Linux Database Server – Method 2 (2/2)

- Recovery with modified partitions
  1. Request password from oracle support
  2. Boot the database server from diagostic.iso
  3. Create File Systems
  4. Restore each volume
  5. Setup the boot loader
- Recommended: Request password before patching
- Step by step backup recovery guide can be found at [Oracle Exadata Database Machine Owner's Guide](#)



## ■ GI and Database Backup

- Before RDBMS Upgrade
  - Is additional backups needed for fallback ?
  - Flashback Database ?
  - Disk Backups ?
- Before GI Upgrade
  - File System Backup or create TAR-Ball of GI Home
  - Service, Instance status, Resources → OCR and OLR backups
  - Spfile Backup
  - SQL\*Net configuration
  - How to Proceed from Failed Upgrade to 11gR2 Grid Infrastructure on Linux/Unix [ID 969254.1]

# Exadata Patching

## ■ Exadata Patching – Let's do it

- Customer starting position is to upgrade from 11.2.3.1.1 to 11.2.3.2.1

### Our Patching Steps:

- Executing exachk → we focussed more on HW Failure
- Infiniband → they are just on Version 1.3.3.2, so nothing was to do...
- Storage Server → was the first component we've patched
- Database Server
- Grid Infrastructure
- RDBMS

## Table of Contents

# ■ Exachk

- Download: MOS 1070954.1
- Generates HTML report over 150 pages
- Especially hardware errors has to be examined and corrected before patching
- Report should be generated before and after patching for documentation.
- Exachk commands
  - Normally in interactive mode with `-a` (All Checks)
  - `-a` Parameter list Best Practices and recommended Patches

- Findings Needing Attention
  - On Database Server
  - On Storage Server
  - Cluster Wide
- MAA Scorecard
- Findings Passed
  - On Database Server
  - On Storage Server
  - Cluster Wide
- Systemwide firmware and software versions
- Systemwide Automatic Service Request (ASR) healthcheck

```
[root@bdm2db01 exachk_224_bundle]# ./exachk -a
```

Component	Host/Location	Found version	Recommended versions	Status
DATABASE SERVER	Database Home	<u>bdm1db01.bdm1db02</u> /u01/app/oracle/product/11.2.0.3/dbhome_1	11.2.0.3.5	11.2.0.3.20 Write-back Flash Cache required minimum 11.2.0.3.9 not met. Older than recommended version.
	Grid Infrastructure	<u>bdm1db01.bdm1db02</u> /u01/app/11.2.0.3/grid	11.2.0.3.5	11.2.0.3.20 Older than recommended version.
	Exadata	<u>bdm1db01.bdm1db02</u>	11.2.3.1.1	11.2.3.2.1 Older than recommended version.
STORAGE SERVER	Exadata	<u>bdm1cel01.bdm1cel02.bdm1cel03</u>	11.2.3.1.1	11.2.3.2.1 Older than recommended version.
IB SWITCH	Firmware	<u>bdm1sw-ib1.bdm1sw-ib2.bdm1sw-ib3</u>	1.3.3-2	1.3.3-2 Version within recommended range.

# ■ InfiniBand Upgrade to the Latest Version

- Current version is at least 1.1.3-2
- Follow Patch readme

```
# On a database server
$ unzip p12373676_112100_Linux-x86-64.zip
$ cd 12373676
$ cp /tmp/downloads/p11891229_133_Generic.zip .
$ unzip p11891229_133_Generic.zip
$ tar -zxvf SUN_DCS_36p_1.3.3.tar.gz
$ scp sundcs_36p_repository_1.3.3_2.pkg root@dm01sw-ib1:/tmp
$ scp ibswitchcheck.sh root@dm01sw-ib1:/tmp

# On the IB-switch, as root user
$ disablesm                #disable subnet manager
$ spsh                      #enter the ILOM CLI
$ load -source /tmp/sundcs_36p_repository_1.3.3_2.pkg
...
# check
$ /tmp/ibswitchcheck.sh post
```

## ■ Storage Server Patching

- Storage Server Patching is possible in online or offline Mode
  - Customer has MAA Infrastructure, so offline Mode was chosen
  - Customer has also a Test Environment, so we started with this Environment
  - We've begun with Standby Site
  
- As first we start with one Storage Server → then all other in parallel
  - This procedure to assure that delivered Patch is installable as expected

# ■ Storage Server Patching – Rolling vs. Non Rolling

## ■ Rolling Mode

- Online
- Patching time depends on number of Storage Servers and disk activity during patching
- Approximately 2 hour per Storage Servers
- A problem can occur only at one cell at a time
- In sequential mode over all Storage Server
- No or less redundancy during patching

## ■ Non Rolling Mode

- Offline
- Downtime is independent of number of storage servers
- In parallel over all Storage Servers
- Approximately 2 hours for all Storage Servers
- A problem may affect more than one Storage Cells at the same time

## ■ Preparation

- Patch manager is executed from database server as root
- ssh equivalence for the root user must be configured

```
# dcli -g /opt/oracle.SupportTools/onecommand/cell_group -l root 'imageinfo'
```

- Do not use /opt/oracle or any directory below for patch storage
- Do not monitor the log files in a writable mode. Use view, tail, ...
- Upgrade progress can be monitored by ILOM console



## ■ Non-Rolling patching (Offline)

- Make sure that every software, including Grid Infrastructure are stopped
- Make sure that every Storage Server Cells are stopped
- All Storage Cells are patched in parallel
- Recommendation: Patch one cell first and the rest in parallel

```
dcli -g dbs_group -l root "/u01/app/11.2.0/grid/bin/crsctl stop crs -f"
dcli -g cell_group -l root "cellcli -e alter cell shutdown services all"

cd patch_11.2.3.2.1.130109

[root@bdm2db01 patch_11.2.3.2.1.130109]# ./patchmgr -cells /tmp/cell_group -
cleanup → clean up all temporary file and rollback files
[root@bdm2db01 patch_11.2.3.2.1.130109]# ./patchmgr -cells /tmp/cell_group -
patch_check_prereq → check Prerequisites
[root@bdm2db01 patch_11.2.3.2.1.130109]# ./patchmgr -cells /tmp/cell_group -
patch → will patch the system
```

## ■ Post Patching Checks

```
[root@bdm2db01]# dcli -g /tmp/cell_group -l root "imageinfo"
[root@bdm2db01]# dcli -g /tmp/cell_group -l root "imagehistory"
[root@bdm2db01]# ./patchmgr -cells /tmp/cell_group -cleanup
[root@bdm2db01 ~]# ssh bdm2cel101 imageinfo
```

```
Kernel version: 2.6.32-400.11.1.el5uek #1 SMP Thu Nov 22 03:29:09 PST 2012 x86_64
Cell version: OSS_11.2.3.2.1_LINUX.X64_130109
Cell rpm version: cell-11.2.3.2.1_LINUX.X64_130109-1
```

### **Active image version: 11.2.3.2.1.130109**

```
Active image activated: 2013-11-06 12:01:44 +0100
Active image status: success
Active system partition on device: /dev/md6
Active software partition on device: /dev/md8
```

```
In partition rollback: Impossible
```

```
Cell boot usb partition: /dev/sdk1
Cell boot usb version: 11.2.3.2.1.130109
```

### **Inactive image version: 11.2.3.1.1.120607**

```
Inactive image activated: 2012-10-03 12:09:37 +0200
Inactive image status: success
Inactive system partition on device: /dev/md5
Inactive software partition on device: /dev/md7
```

```
Boot area has rollback archive for the version: 11.2.3.1.1.120607
Rollback to the inactive partitions: Possible
```

# ■ Patching Database Server Overview

- Prepare YUM Repository
- Stop CRS Resources
- Disable CRS
- Patch OS
- Relink GI and RDBMS Home
- Enable CRS
- Restart CRS Resources

# ■ Patching Database Server 1/4

- We follow the instruction of „Exadata ISO Cheat Sheet“
  - MOS Doc ID 1545789.1 → cover 11.2.3.1.1 and 11.2.3.2.1 Version
- „Manual “ procedure like as `dbnodeupdate.sh`
- As of 11.2.3.1.0 no more „local“ YUM Repository is needed
- YUM Repository can be „mounted“ from ISO Image
- Looking out →
  - if you adding supplemental RPMs, this are not included in ISO Image, you've to upgrade manually
  - if you update with `dbnodeupdate.sh`, you will receive an error message

```
ERROR: Yum update failed due to dependency issues. Resolve conflicting packages issue(s) and rerun
dbnodeupdate.sh.
```

```
The following next steps are suggested :
```

```
RPM required           : bind-libs (release 30)
RPM installed          : bind-libs-9.3.6-20.P1.el5.x86_64
Requirement enforced by rpm : 30:bind-9.3.6-20.P1.el5.x86_64
Possible next step to solve dependency issue : Downgrade or remove bind-libs
```

## ■ Patching Database Server 2/4

- Copy the ISO file to each compute node
  - e.g. /u01/patches/iso/
- Create directory for ISO Mount Point

```
mkdir -p /mnt/iso/yum/unknown/EXADATA/dbserver/11.2/latest
```

- Mount the ISO Image

```
mount -o loop /u01/patches/iso/112_latest_repo_130302.iso /  
/mnt/iso/yum/unknown/EXADATA/dbserver/11.2/latest
```

- Adjust YUM configuration

```
[exadata_dbserver_11.2_x86_64_latest]  
name=Oracle Exadata DB server 11.2 Linux $releasever - $basearch - latest  
baseurl=file:///mnt/iso/yum/unknown/EXADATA/dbserver/11.2/latest/x86_64  
gpgcheck=1  
enabled=0
```

## ■ Patching Database Server 3/4

- Disable CRS and stop CRS Resources

```
[root@bdm2db01]# $GI_HOME/bin/crsctl disable crs  
[root@bdm2db01]# $GI_HOME/bin/crsctl stop crs -f
```

- Disabling all YUM Repositories

```
sed -i 's/^[ \t ]*enabled[ \t ]*=[ \t ]*1/enabled=0/g' /etc/yum.repos.d/*
```

- If you need to update UEK Kernel, remove following from /etc/yum.conf

```
exclude=up2date kernel-uek-headers  
must be replaced with  
exclude=up2date
```

- Executing Update

```
[root@bdm2db01]# yum clean all  
[root@bdm2db01]# yum --enablerepo= exadata_dbserver_11.2_x86_64_latest reposit  
[root@bdm2db01]# yum --enablerepo= exadata_dbserver_11.2_x86_64_latest update
```

# ■ Patching Database Server 4/4

## ■ Relink GI and RDBMS Home

```
[root@bdm2db01]# $GI_HOME/crs/install/rootcrs.pl -unlock  
  
[root@bdm2db01]# export ORACLE_HOME=/u01/app/11.2.0.4/grid  
[root@bdm2db01]# export PATH=$ORACLE_HOME/bin:$PATH  
[root@bdm2db01]# relink all  
[root@bdm2db01]# make -C $ORACLE_HOME/rdbms/lib -f ins_rdbms.mk ipc_rds ioracle
```

## ■ Lock GI Home and enable CRS

```
[root@bdm2db01]# $GI_HOME/crs/install/rootcrs.pl -patch  
[root@bdm2db01]# $GI_HOME/bin/crsctl enable crs
```

## ■ Reboot the Computing Node

```
[root@bdm2db01]# shutdown -r now
```

# ■ GI and RDBMS Infrastructure Patching - Overview

- Stop Database with Cluster commands (auto\_start=restore)
- Create OCM Response File if not exist
- Check Patch Conflicts

```
$ORACLE_HOME/OPatch/opatch prereq CheckConflictAgainstOHWithDetail -phBaseDir ./16869210/16869205  
$ORACLE_HOME/OPatch/opatch prereq CheckConflictAgainstOHWithDetail -phBaseDir ./16869210/16619898  
$ORACLE_HOME/OPatch/opatch prereq CheckConflictAgainstOHWithDetail -phBaseDir ./16869210/16937759  
$ORACLE_HOME/OPatch/opatch prereq CheckConflictAgainstOHWithDetail -phBaseDir ./16869210/16869205
```

- Apply Patch as root (opatch auto)
  - Implicitly all CRS resources will be stopped
  - Patches will be applied

```
$ORACLE_HOME/OPatch/opatch auto ./11.2.0.3.20_QDPE_Jul2013/16869210 -ocmrf /tmp/ocm.rsp
```



## ■ Data Guard Standby-First Patch Apply

- As of 11.2.0.1 Oracle supports different software releases between a primary database and its physical standby database under following conditions:
  - If patch is certified. Check readme
  - Same compatible parameter values primary and standby
  - Not longer than a month
  - Patch Release date difference not longer than a year
  - See MOS Oracle Patch Assurance - Data Guard Standby-First Patch Apply [ID 1265700.1] for more information

## ■ Conclusion

- Even Exadata as an engineered solution needs software maintenance
- Patching should be done regularly to proactive avoid problems
- Integrate it into the Lifecycle-Management with a well defined patch process
- Patching involves near countless steps, define and use checklists
- Regularly check your Systems state with the provided Healthcheck Scripts (exachk)
- Evaluate the latest patch feature, as there is a fast improvement cycle
- Having redundancy in terms of a MAA system or Test Environment reduces patch risks and time

# Storage Server New Features

## 11.2.3.3 and 12.1.1.1

## ■ Exadata New Features 12.1.1.1.0 (Jan. 2014)

- Support for RDBMS 12.1.x and 11.2.x (also as mixed combination)
- IORM Support for Container Databases and Pluggable Databases
- Cell to Cell Data Transfer
  - Offload Data movement for ASM will be transferred without going through the Database Server (rebalance, resynchronisation, etc.)

## ■ Exadata New Features 11.2.3.3.0 (Dec. 2013)

- Flash Cache Compression (up to 4x for uncompressed Table)
- Automatic Flash Caching for Table Scan Workloads
  - Special algorithm ensure that only part of table or partition is cached
- Fast Data File Creation
- Network Resource Management (Database- , Storage-Server, Infiniband)
  - i.e. RAC Cache Fusion Messages are higher prioritized
- Active Bonding Network (active/active) → from X4 upwards
- Oracle ASM Disk Group in Appliance Mode
  - New Diskgroup attribute `appliance.mode` (not available with 12.1.0.1 !!)
  - Improves Disk rebalance completion time, so redundancy is restored faster
- Automatic Hard Disk Scrub and Repair
  - Detects and repairs bad sectors automatically
  - Storage Server SW send requests to ASM to repair from mirror copy
- Oracle Exadata Database Machine Eighth Rack Configuration
  - Enable and disable a eighth resp. quarter rack with `alter cell` command
  - Before we need the `resourcecontrol` utility

# Questions and answers ...

Daniele Massimi

Senior Consultant

Tel. +41 58 459 50 92

[daniele.massimi@trivadis.com](mailto:daniele.massimi@trivadis.com)



BASEL BERN BRUGG LAUSANNE ZUERICH DUESSELDORF FRANKFURT A.M. FREIBURG I.BR. HAMBURG MUNICH STUTTGART VIENNA