

Exadata Software Maintenance

Doug Utzig
Exadata and MAA Best Practices

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Patching and Upgrading Oracle Exadata

- 1 Software Architecture Overview
- 2 Software Maintenance Planning
- 3 Updating Exadata Software


Patching and Upgrading Oracle Exadata

- 1 Software Architecture Overview
- 2 Software Maintenance Planning
- 3 Updating Exadata Software

Q: Why Exadata? A: Engineered Systems Value



Oracle Engineered Systems are the only fully tested full-stack configuration



Exadata Community Effect

Oracle Public Cloud
Oracle Development & Support
1000s of Customer and Partners

100%

Simplified Maintenance

Full Stack Patching
Full Stack Health Checks



Platinum Services

Oracle engineers perform remote patch installation at no additional cost

Exadata Database Machine

Software Architecture Review (Bare Metal / Physical)

Database Grid

- Oracle Database and Grid Infrastructure
- **Exadata** (firmware, Linux, Exadata)


Storage Grid

- **Exadata** (firmware, Linux, Exadata)

Networking

- **Exadata** (InfiniBand switch software)

Other: Ethernet switch, PDU



Software Architecture Comparison

Database Server: Bare Metal / Physical Deployment versus OVM Deployment

Bare Metal / Physical Database Server

Oracle GI/DB homes

Exadata (Linux, fw)

OVM Database Server

dom0

Exadata (Linux,
Xen, fw)

domU-1

Oracle GI/DB
homes

Exadata (Linux)

domU-2

/DB

nux)

domU-3

/DB

nux)

No change to **Storage Grid, Networking**, or **Other**

Patching and Upgrading Oracle Exadata

- 1 Software Architecture Overview
- 2 Software Maintenance Planning
- 3 Updating Exadata Software

Exachk for Planning Software Maintenance

Automated Exadata Health Check – MOS 1070954.1


Simplify software planning

1. Version recommendations
2. Critical Issue exposure report

Oracle Exadata Assessment Report

Table of Contents

- [Findings Needing Attention](#)
 - [On Database Server](#)
 - [On Storage Server](#)
 - [On InfiniBand Switch](#)
 - [Cluster Wide](#)
- [Maximum Availability Architecture \(MAA\) Scorecard](#)
- [Infrastructure Software and Configuration Summary](#)
- [Findings needing further review](#)
- [Platinum Certification](#)
- [Findings Passed](#)
 - [On Database Server](#)
 - [On Storage Server](#)
 - [On InfiniBand Switch](#)
 - [Cluster Wide](#)
- [Systemwide Automatic Service Request \(ASR\) healthcheck](#)
- [Skipped Checks](#)
- [Top 10 Time Consuming Checks](#)



Exachk Critical Issue Exposure Report (sample)

CRITICAL	Storage Server Check	System is exposed to Exadata Critical Issue EX51	All Storage Servers	View
PASS	Database Server Check	System is not exposed to Exadata Critical Issue EX50	All Database Servers	View

★ Exadata Critical Issues (Doc ID 1270094.1)

EX51	Storage servers running Exadata version 18.1.10, 18.1.11, or 18.1.12 using IORM to manage flash cache	Bug 29288067 - When I/O Resource Management (IORM) is configured to manage flash cache on storage servers, the cellsrv process may crash with error ORA-600 [FCGroupDesc::decLocalCnt_underflow].	Fixed in Exadata 18.1.13. See Document 2511918.1 for details.
------	---	---	---

Late-breaking issues - **MOS Alerts for Hot Topics** (See How To MOS 793436.2)

Exachk Version Recommendation (sample)

Component		Host/Location	Found version	Recommended versions	Status
DATABASE SERVER	Database Home	dm01db01,dm01db02: /u01/.../dbhome_1	11.2.0.3.28	11.2.0.4.190416	11.2.0.3 Error Correction Support ended Aug 2015.
		dm01db01,dm01db02: /u01/.../dbhome_2	11.2.0.4.10	11.2.0.4.190416	11.2.0.4 BP is older than recommended.
		dm01db01,dm01db02: /u01/.../dbhome_3	18.4.0.0.181016	18.6.0.0.190416	Version within recommended range.
	Grid Infrastructure	dm01db01,dm01db02: /u01/.../grid	18.6.0.0.190416	18.6.0.0.190416	Version within recommended range.
	Exadata	dm01db01,dm01db02	18.1.14.0.0	18.1.15.0.0	Version within recommended range.
STORAGE SERVER	Exadata	dm01cel01,dm01cel02	18.1.14.0.0	18.1.15.0.0	Version within recommended range.
		dm01cel03	11.2.3.3.1	18.1.15.0.0	Older than recommended version. Exception: Version is different from peers.
IB SWITCH	Firmware	dm01sw-iba0,dm01sw-ibb0	2.2.11-2	2.2.12-2 or higher	Version within recommended range.

Software Maintenance

Recommended Update Schedule

Frequency	Database / Grid Infrastructure	Exadata
3-12 months	Quarterly Release Update	Monthly Sustaining Release
1-5 years	New Feature Release	New Feature Release

**Long Term
Support Release**

Patching End Date	11.2.0.4	12.1.0.2	12.2.0.1	18	19
	2020-Dec	2021-Jul	2020-Nov	TBD	2026-Mar

All software for Exadata
MOS 88828.1

Responses to security scan findings
MOS 1405320.1

* Upgrade now from versions (e.g. 11.2.0.3) where patching has ended.


Zero Downtime Software Maintenance

Rolling Software Update Support

Component to Update	How to Mitigate Impact and Risk
Database / Grid Infrastructure	Rolling GI / DB updates with Fleet Patching and Provisioning Application Continuous Availability Data Guard Standby First
Exadata Database Server	Rolling Database Server updates Application Continuous Availability and RHPHelper Data Guard Standby First
Exadata Storage Server	Rolling Storage Server updates ASM HIGH redundancy Data Guard Standby First
Exadata InfiniBand switch	Rolling InfiniBand switch updates Data Guard Standby First

Reduce Risk and Downtime with Data Guard


Data Guard Standby First Patching (MOS 1265700.1)



Standby First Patching Steps

1. Update software on Site B (Standby)
2. Test new software
3. Switchover (optional)
4. Update software on Site A
5. Run SQL portion of RU on Primary databases


Data Guard Standby First Patching



Standby First Patching Steps

- 1. Update software on Site B (Standby)**
- 2. Test new software**
3. Switchover (optional)
4. Update software on Site A
5. Run SQL portion of RU on Primary databases


Data Guard Standby First Patching



Standby First Patching Steps

1. Update software on Site B (Standby)
2. Test new software
- 3. Switchover** (optional)
4. Update software on Site A
5. Run SQL portion of RU on Primary databases


Data Guard Standby First Patching



Standby First Patching Steps

1. Update software on Site B (Standby)
2. Test new software
3. Switchover (optional)
- 4. Update software on Site A**
5. Run SQL portion of RU on Primary databases

Data Guard Standby First Patching




Standby First Patching Steps

1. Update software on Site B (Standby)
2. Test new software
3. Switchover (optional)
4. Update software on Standby
5. **Run SQL portion of RU on Primary databases**

Patching and Upgrading Oracle Exadata

- 1 Software Architecture Overview
- 2 Software Maintenance Planning
- 3 Updating Exadata Software**

High Level Software Maintenance Flow



Inputs

- Exachk
- Issue resolution
- Application compatibility

Applies to

- Oracle Database and Grid Infrastructure
- Exadata for Database Grid, Storage Grid, and Networking

Grid Infrastructure and Database Software

- Features Releases and Quarterly Updates
 - Use only certified releases specified in MOS 888828.1
- New Feature Release Upgrade
 - Confirm Exadata version requirements are met beforehand
 - (Example - GI/DB 18c requires Exadata 18.1.4)
 - Follow Exadata-specific step-by-step guides. See MOS 888828.1.

Grid Infrastructure and Database Software

Fleet Patching and Provisioning

- Fleet Patching and Provisioning (previously Rapid Home Provisioning RHP)
 - Simplifies and Automates out-of-place software update for GI/DB
 - Integrates with Application Continuous Availability

```
$ rhpctl move database -dbname ORCL \  
    -sourcehome /u01/.../dbhome1 -patchedwc DBHOME18_190416 \  
    -drain_timeout 180 ...
```

- Out-of-Place updates - Reduce maintenance window for GI/DB homes
 - Release Update install removed from maintenance window
 - Prepare new software home and apply updates while GI or DB remains up locally
 - Quick process to switch GI/DB to use new software home

Zero Downtime Software Maintenance

Application Continuous Availability

- Graceful Application Switchover in RAC with No Application Interruption (MOS 2440719.1)
 - No application interruption during dbnode planned maintenance
 - Grid Infrastructure software home patching and upgrades
 - Database software home patching
 - Exadata database server updates
 - Use modern Features and Configuration

Multi-instance RAC database	FAN-aware connection pools
Clusterware-managed services	TNS configuration

Zero Downtime Software Maintenance

Application Continuous Availability

- High-Level Glance - Move Services and Drain Sessions

1. Stop/relocate and disable services on target node
 - Services move to other instance(s) and immediately accept new connections
 - Existing connections finish their work (within drain_timeout)
2. Disconnect remaining long-running sessions and stop database instance(s)
3. Perform desired maintenance
 - Examples: Patch database home software, Update grid infrastructure home software, Exadata dbnode update
4. Start database instance(s) on target node
5. Enable and start services on target node
 - Services move back and begin accepting new connections
6. Repeat steps on next node

**Fully automatic when using
Fleet Patching and Provisioning**

Database Software

Oracle Java VM is RAC Rolling Installable

- MOS 2217053.1 - RAC Rolling Install Process for the "Oracle Java VM Component Database PSU/RU" (OJVM PSU/RU) Patches
- Requires
 - Out-of-place software update
 - Clusterware-managed database services

GI/DB Upgrade Keys to Success


GI/DB Software Maintenance Rules

- Use Application Continuous Availability to eliminate application impact
- Leverage Data Guard to evaluate new software before using on primary system
- Use Fleet Patching and Provisioning to simplify and reduce patching time
- Check patch conflicts weeks before maintenance
- Qualify maintenance readiness - Upgrade only when exachk and prereq check are clean

Exadata Database Server

Simple dbnode update steps

- Exachk Health Check
- Perform a first prerequisite check with `–nomodify_at_prereq`
 - `patchmgr -dbnode -precheck -nomodify_at_prereq`
- Perform a “backup only” run using the `-backup` flag
 - `patchmgr -dbnodes dbs_group –backup`
- Perform second prerequisite check
 - `patchmgr -dbnodes dbs_group –precheck`
- Update database servers
 - `patchmgr -dbnode -dbnode_upgrade -nobackup [-rolling]`
 - * **Steps are same for dom0 or domU, but are performed independently**



Zero Downtime Software Maintenance

Application Continuous Availability

- Using RHPhelper to Minimize Downtime During Planned Maintenance on Exadata (Doc ID 2385790.1)
 - Patchmgr manages Oracle Clusterware shutdown and startup during Exadata updates
 - No application interruption during dbnode planned maintenance
 - Grid Infrastructure software home patching and upgrades
 - Database software home patching
 - Exadata database server updates



Database Server Ksplice Updates

Install Important Linux Fixes without Reboot

- Install important Oracle Linux kernel security, stability, and performance fixes on database servers without reboot
- Intended for installing fixes in between Exadata Sustaining Releases

Example:

1. Quarterly - Install Exadata Sustaining Release
2. Monthly - Install Ksplice updates for Oracle Linux kernel security compliance

Jan	Exadata update
Feb	Ksplice update
Mar	Ksplice update
Apr	Exadata update

- Ksplice Offline Client Procedure

1. Obtain uptrack-updates RPM (mirrored ULN repo or download from ULN)
2. Install using yum / rpm command (no reboot)

- HOWTO: Install ksplice kernel updates for Exadata Database Nodes (Doc ID 2207063.1)

Database Server - Standard vs. Custom Configuration

- **Standard Configuration**

- Best practice configuration to run Oracle Database and Grid Infrastructure
- Minimal by design
- Full pre-release testing coverage
 - **Predictable, low risk update**

- **Custom Configuration**

- Customer-specific changes made to database servers after deployment
 - Allowed (sometimes required for given environment), but resist, test, track, and automate
- Limited / minimal pre-release testing coverage
- Increases admin cost and risk

Exadata Standard Configuration for Database Server

Software	Exact list of Oracle Linux packages and their versions, and firmware versions
Configuration	Best practice configuration (e.g. sysctl, network, ssh, pam, modules, drivers, etc.)
Disk	RAID, Logical volume (LVM), and file system configuration

Database Server Custom Configuration

Customization Examples and Impact to Update

	Impact to Update
Exadata Standard Configuration (i.e. not customized)	None
Using all free space in VGExaDb	Low
Customized file system – different mount points	Low
Updating packages shipped with the current Exadata Image	Low
Installation of additional (non-Exadata) rpm packages	Low
Customizing configuration files, removing / changing basic O.S. functionality	Medium
Installation of additional (non-Exadata) non-rpm packages	Medium
Setting up interactive shell profile / menus	High
Changing LVM layout	High

Database Server Customization

Best Practices

- Resist customization, keep it minimal if you must
- Firmware
 - Maintained automatically during Exadata update process – do not customize
- Linux packages (RPMs)
 - Acceptable to update supplied packages to later versions (ULN or public-yum)
 - **Except** kernel and boot related packages
 - Acceptable to add new packages
 - Automate install / removal - some Exadata updates require custom package remove / reinstall
 - New package dependencies introduced must be customer-managed
 - Do not install 32bit packages, or packages for wrong OL release (e.g. OL6 RPM on OL7)

Database Server Customization

Best Practices

- System configuration
 - Avoid changing kernel parameters
 - Standard Linux server settings often do not apply
 - Do not change driver / module configuration
 - Do not shutdown running services
 - Track customizations closely
 - Exadata updates may apply new best practices (overwriting customizations)
 - Ensure system boots properly after customizing (broken boot → upgrade failure)
 - Shell profile should not be interactive

Database Server Customization

Best Practices

- Local disk RAID, LVM, filesystem configuration
 - Do not change RAID configuration
 - Do not change supplied LVM configuration
 - Acceptable to add volumes
 - Leave free space for dbnodeupdate.sh backup snapshot
 - Do not change supplied filesystem configuration
 - Acceptable to add filesystems
 - Do not place Oracle Database software in /opt/oracle



Database Server Upgrade Keys to Success

Database Server Software Maintenance Rules

- Use Application Continuous Availability to eliminate application impact
- Customization allowed, but resist - test changes to avoid latent patching failures
- Closely track customizations, automate build-up and teardown
- Qualify maintenance readiness - Upgrade only when exachk and prereq check are clean
- Always use latest patchmgr (Doc ID 1553103.1)

Exadata Storage Server


Simple cell update steps

- Exachk Health Check
- Prereq validation

```
# patchmgr -cells -patch_check_prereq [-rolling]
```


- Update all storage servers

```
# patchmgr -cells -patch [-rolling]
```



Storage Server Update Flow

Non-Rolling and Rolling

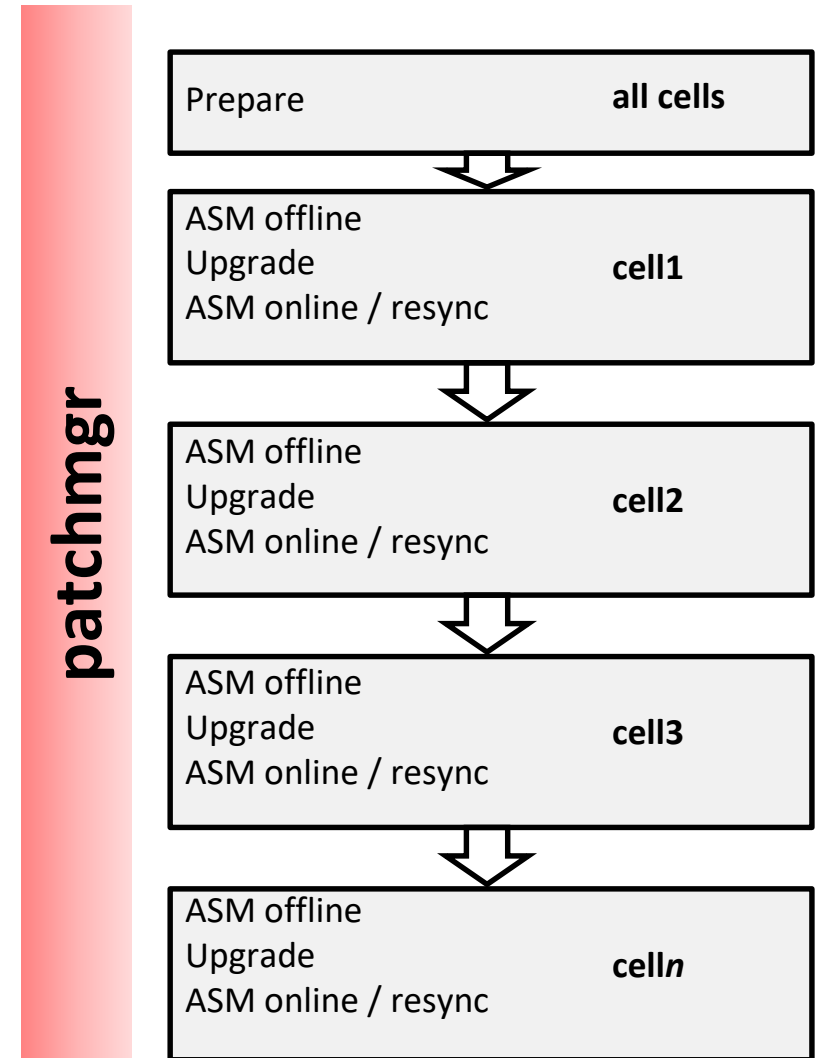


Non-Rolling

- Databases/CRS down
- All cells upgraded in parallel

Rolling

- Databases/CRS up
- One cell upgraded at a time
- patchmgr manages ASM offline / online



Rolling Storage Server Updates

High-Level Flow

Step	Scope	All Cells in Service?
Prereq check	All cells in parallel	Yes
Copy and stage new software	All cells in parallel	Yes
Inactivate disks	One cell at a time	No
Reboot into new partitions	One cell at a time	No
Complete software update	One cell at a time	No
Update firmware	One cell at a time	No
Validate	One cell at a time	No
Activate disks and resync	One cell at a time	Yes
Update complete - go to next cell	One cell at a time	Yes



Zero Downtime Software Maintenance

ASM High Redundancy

- Define DATA and RECO ASM disk groups with HIGH redundancy
 - Defined during OEDA configuration
 - To change after deployment requires data reload
- Use HIGH for best protection during rolling storage server updates
 - HIGH – data mirrored on 3 cells; NORMAL – data mirrored on 2 cells
 - HIGH can tolerate disk failure while a storage server is offline for update
 - NORMAL may result in ASM disk group offline and data loss

Storage Server Update Monitoring

Patchmgr Progress Email Notification

ORACLE
EXADATA

Patchmgr: Patch State of cell05 Changed from Waiting to Patching

Event Time 2015-09-01 11:35:08-0700

Description Patch state of cell05 changed from Waiting to Patching.
Patchmgr launched from db03 is performing rolling patch on following cell(s).
1 out of 3 cell(s) completed.

Cell	Patch State	From Version	To Version	Time
cell04	Succeeded	11.2.3.3.1.140708	12.1.2.1.2.150617.1	2015-09-01 11:35:08-0700
cell05	Patching	11.2.3.3.1.140708	12.1.2.1.2.150617.1	2015-09-01 11:35:08-0700
cell06	Waiting	11.2.3.3.1.140708	12.1.2.1.2.150617.1	2015-09-01 10:08:42-0700

Recommended Action No action is needed.

Storage Server Software Update at Scale (Exadata 18)

Scheduling Automated Updates

- Cell SoftwareUpdate (see Exadata doc)
 1. Point cells to software store URL (HTTP(S) location)
 2. Stage new cell software update in store
 3. Set time when cells will automatically apply update to themselves
 - Prereq validation automatically runs before scheduled time
 4. To perform subsequent cell update – repeat steps 2 and 3
- Cells use ASM redundancy checks to determine when it is safe to deactivate griddisks and update themselves
 - With larger systems (e.g. full rack) multiple cells will update simultaneously



Storage Server Update Keys to Success

Storage Server Software Maintenance Rules

- Do NOT make unsupported configuration changes
- Qualify maintenance readiness - Update only when exachk is clean
- Reduce rolling update disk failure risk w/ ASM HIGH redundancy (or DG)

Exadata InfiniBand Switch

Simple switch update steps


1. Exachk Health Check

2. Prereq validation

```
# patchmgr -ibswitch -ibswitch_precheck
```

3. Update all InfiniBand switches

```
# patchmgr -ibswitch -upgrade
```



Method	Patchmgr Orchestration	Database Downtime
Rolling	One switch updated at a time	None

InfiniBand Switch Update Keys to Success


InfiniBand Switch Software Maintenance Rules

- Do NOT make unsupported configuration changes
- Qualify maintenance readiness - Update only when exachk is clean

Updates Using Enterprise Manager Cloud Control 13c

Database Lifecycle Management Pack


- Patch All Exadata Software
 - Grid Infrastructure and Database homes
 - Rolling or non-rolling
 - In-place or out-of-place
 - Exadata Storage Servers, database servers, and InfiniBand switches
 - Rolling or non-rolling
- Patch Recommendations based on the Quarterly Full Stack Download Patch (QFSDP)



Summary

Best Practices for Exadata Planned Maintenance

- Leverage Exachk for simple software planning
- Use available features to Mitigate Impact and Risk
 - Application Continuous Availability
 - Data Guard Standby First
 - Out-of-place software update
 - ASM HIGH Redundancy
- Take advantage of Exadata Engineered defaults - Resist customization



Exadata Software Maintenance

MOS References

- **MOS 888828.1 - Supported and Recommended Versions**
- MOS 1270094.1 - Critical Issues
- MOS 1405320.1 - Responses to Common Security Scan Findings
- MOS 1553103.1 - Database Server Update Tool
- MOS 1070954.1 - Exachk
- MOS 1262380.1 - Software Maintenance Overview and Guidelines
- MOS 2207063.1 - HOWTO: Install ksplice kernel updates for Exadata Database Nodes
- MOS 2440719.1 - Graceful Application Switchover in RAC with No Application Interruption

ORACLE®