ORACLE

Oracle Exadata Database Machine

Software Maintenance

MAA Team & Exadata Product Management

April 2022

Patching and Upgrading Oracle Exadata
Agenda

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

Patching and Upgrading Oracle Exadata
Agenda

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

Q: Why Exadata?

A: Engineered Systems Value

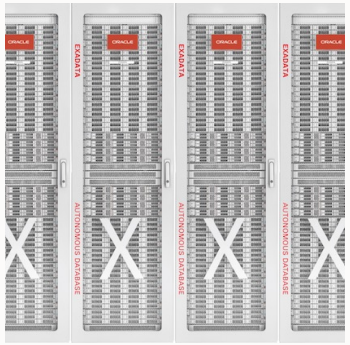

Oracle Engineered Systems are the only **fully tested full-stack configuration**

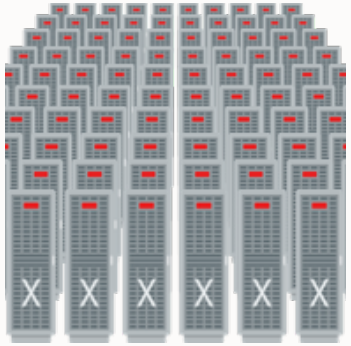

Exadata Community Effect

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

100%

Simplified Maintenance

Full Stack Patching
Full Stack Health Checks

ORACLE Support

Platinum Services

Oracle engineers perform remote patch installation at no additional cost

Exadata Database Machine
Architecture Overview

- Database Instances/Database Homes
- Grid Infrastructure & ASM
- Exadata Database Server
  - KVM Guest / OVM DomU / Bare Metal
  - Linux
  - Exadata specific libraries & binaries
  - Server firmware
  - ILOM
  - KVM Host / OVM Dom0 if virtualized
- RDMA Network Fabric Switch
- Exadata Storage Server
  - Linux
  - Exadata specific libraries & binaries
    - MS, RS, cellsrv, celloflsrv, diskmon
  - Server firmware
  - ILOM
- Other
  - Management Switch & PDUs

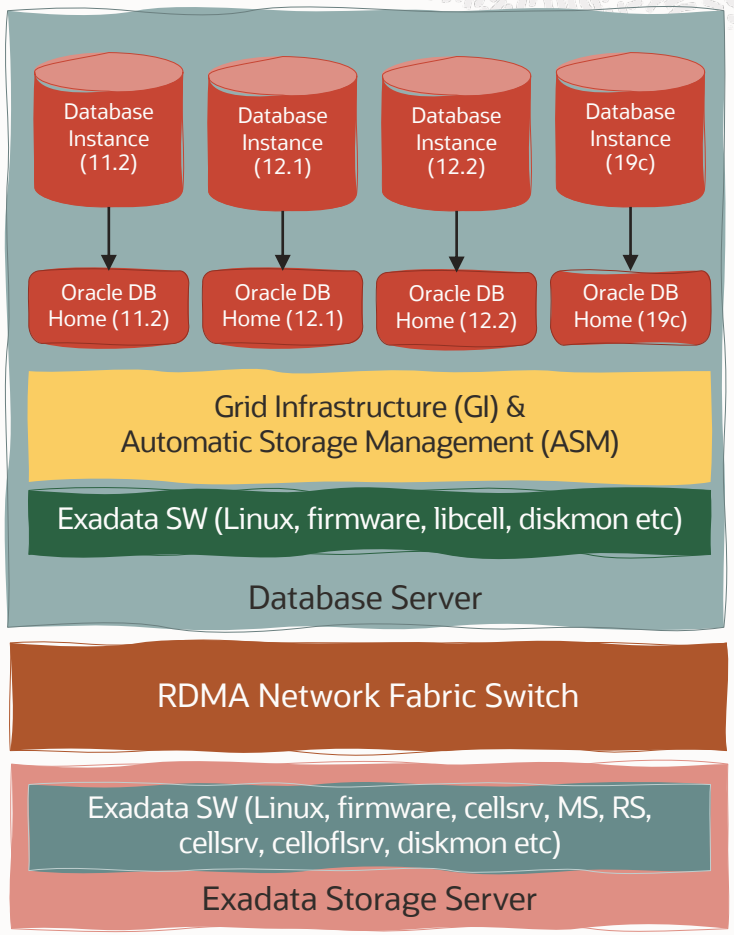

Bare Metal vs Virtualization

Database Server Virtualization

Bare Metal

- Strongest resource isolation
- Consolidate (if required) using multiple instances/homes and Multitenant (PDBs)

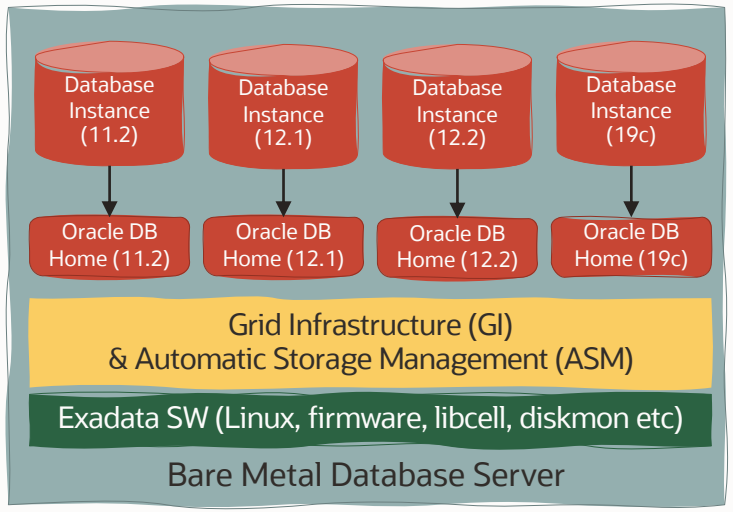

Virtualization

- Strong resource isolation – network, memory & OS boundaries while preserving CPU and sharing
- Consolidate (if required) using multiple instances/homes and Multitenant (PDBs)
- Increased management – more software to update

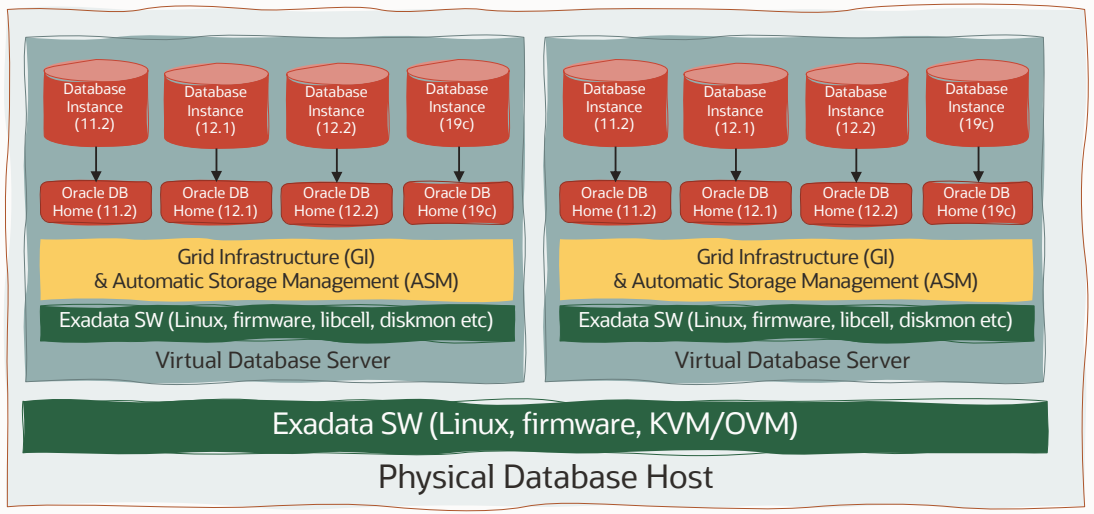

Patching and Upgrading Oracle Exadata
Agenda

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

Exachk for Planning Software Maintenance

Automated Exadata Health Check – MOS 1070954.1

- Exachk holistically evaluates Exadata Database Machine Engineered Systems
- Exachk content is continuously expanded and improved via an Maximum Availability Architecture (MAA) and Exadata Development collaborative process
- Simplify software planning with Exachk, which includes:
  - Version recommendations
  - Critical Issue exposure report

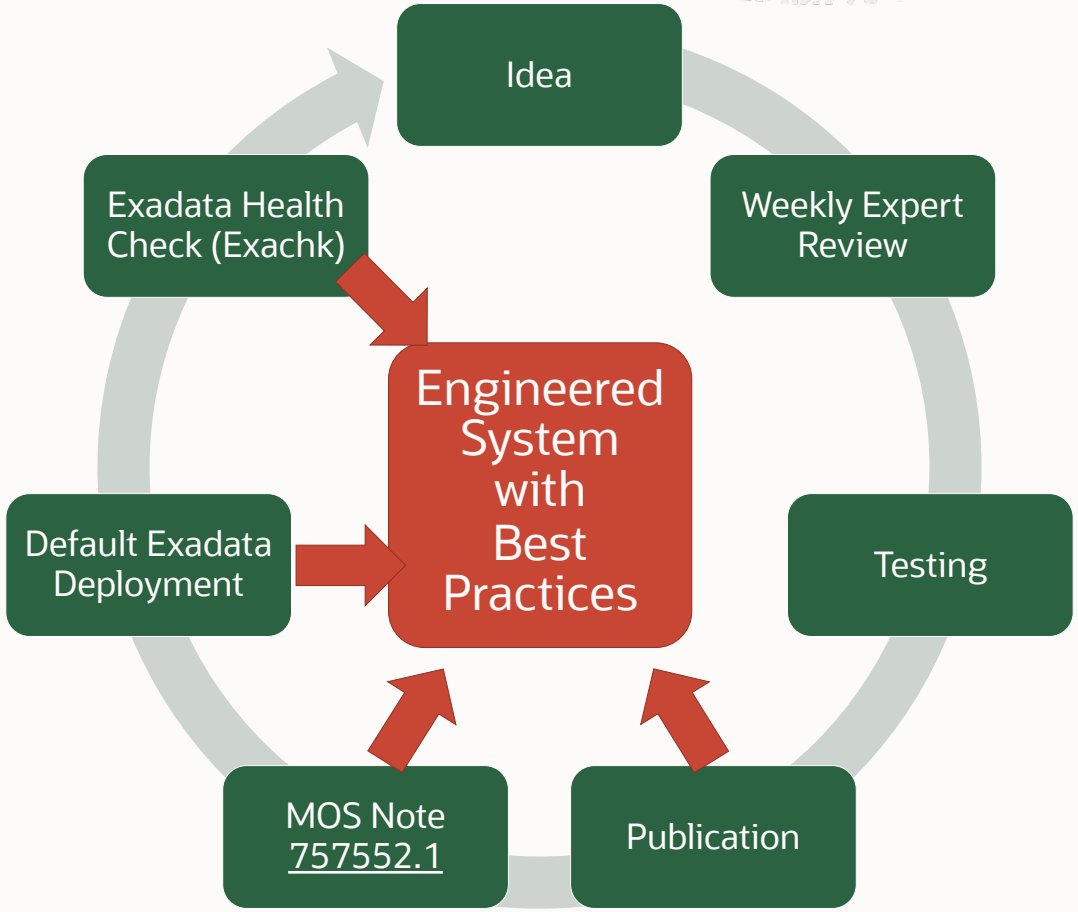

Exachk Critical Issue Exposure Report (sample)

Exadata Critical Issues (Doc ID 1270094.1)

CRITICAL	Storage Server Check	System is exposed to Exadata Critical Issue EX72	All Storage Servers	View
PASS	Database Server Check	System is not exposed to Exadata Critical Issue EX50	All Database Servers	View

EX72	Database servers on RoCE-based systems running Exadata 21.2.2 through 21.2.7.	Bug 33703438 - A database node (a physical server or a virtual KVM guest) on a RoCE-based (i.e. X8M or X9M) Exadata system does not automatically restart if it is evicted by Oracle Clusterware Cluster Synchronization Services (CSS). The evicted node will remain unavailable until manual action is taken to reset the node.	Fixed in Exadata 21.2.8. See Document 2833252.1 for details.

Late-breaking issues - MOS Alerts for Hot Topics (See HowTo MOS 793436.2)

Exachk Version Recommendation (sample)

Component		Host/Location	Found version	Recommended versions	Status
DATABASE SERVER	Database Home	dm01db01,dm01db02: /u01/…/dbhome_1	11.2.0.4.10	11.2.0.4.210119	No longer under Error Correction Support. See My Oracle Support Document 742060.1.
		dm01db01,dm01db02: /u01/…/dbhome_2	18.4.0.0.181016	18.14.0.0.210420	No longer under Error Correction Support. See My Oracle Support Document 742060.1.
		dm01db01,dm01db02: /u01/…/dbhome_3	19.10.0.0.210119	19.14.0.0.220118	19 RU is older than recommended.
		dm01db01,dm01db02: /u01/…/dbhome_4	19.13.0.0.211019	19.14.0.0.220118	Version within recommended range.
	Grid Infrastructure	dm01db01,dm01db02: /u01/…/grid	19.14.0.0.220118	19.14.0.0.220118	Version within recommended range.
	Exadata	dm01db01,dm01db02	21.2.9.0.0	21.2.10.0.0	Version within recommended range.
STORAGE SERVER	Exadata	dm01cel01,dm01cel02	21.2.9.0.0	21.2.10.0.0	Version within recommended range.
		dm01cel03	21.2.3.0.0	21.2.10.0.0	Older than recommended version. Exception: Version is different from peers.
RoCE SWITCH	Firmware	dm01sw-rocea0,dm01sw-roceb0	7.0(3)I7(9)	7.0(3)I7(9)	Version within recommended range.

Lifetime Support Policy (Database)

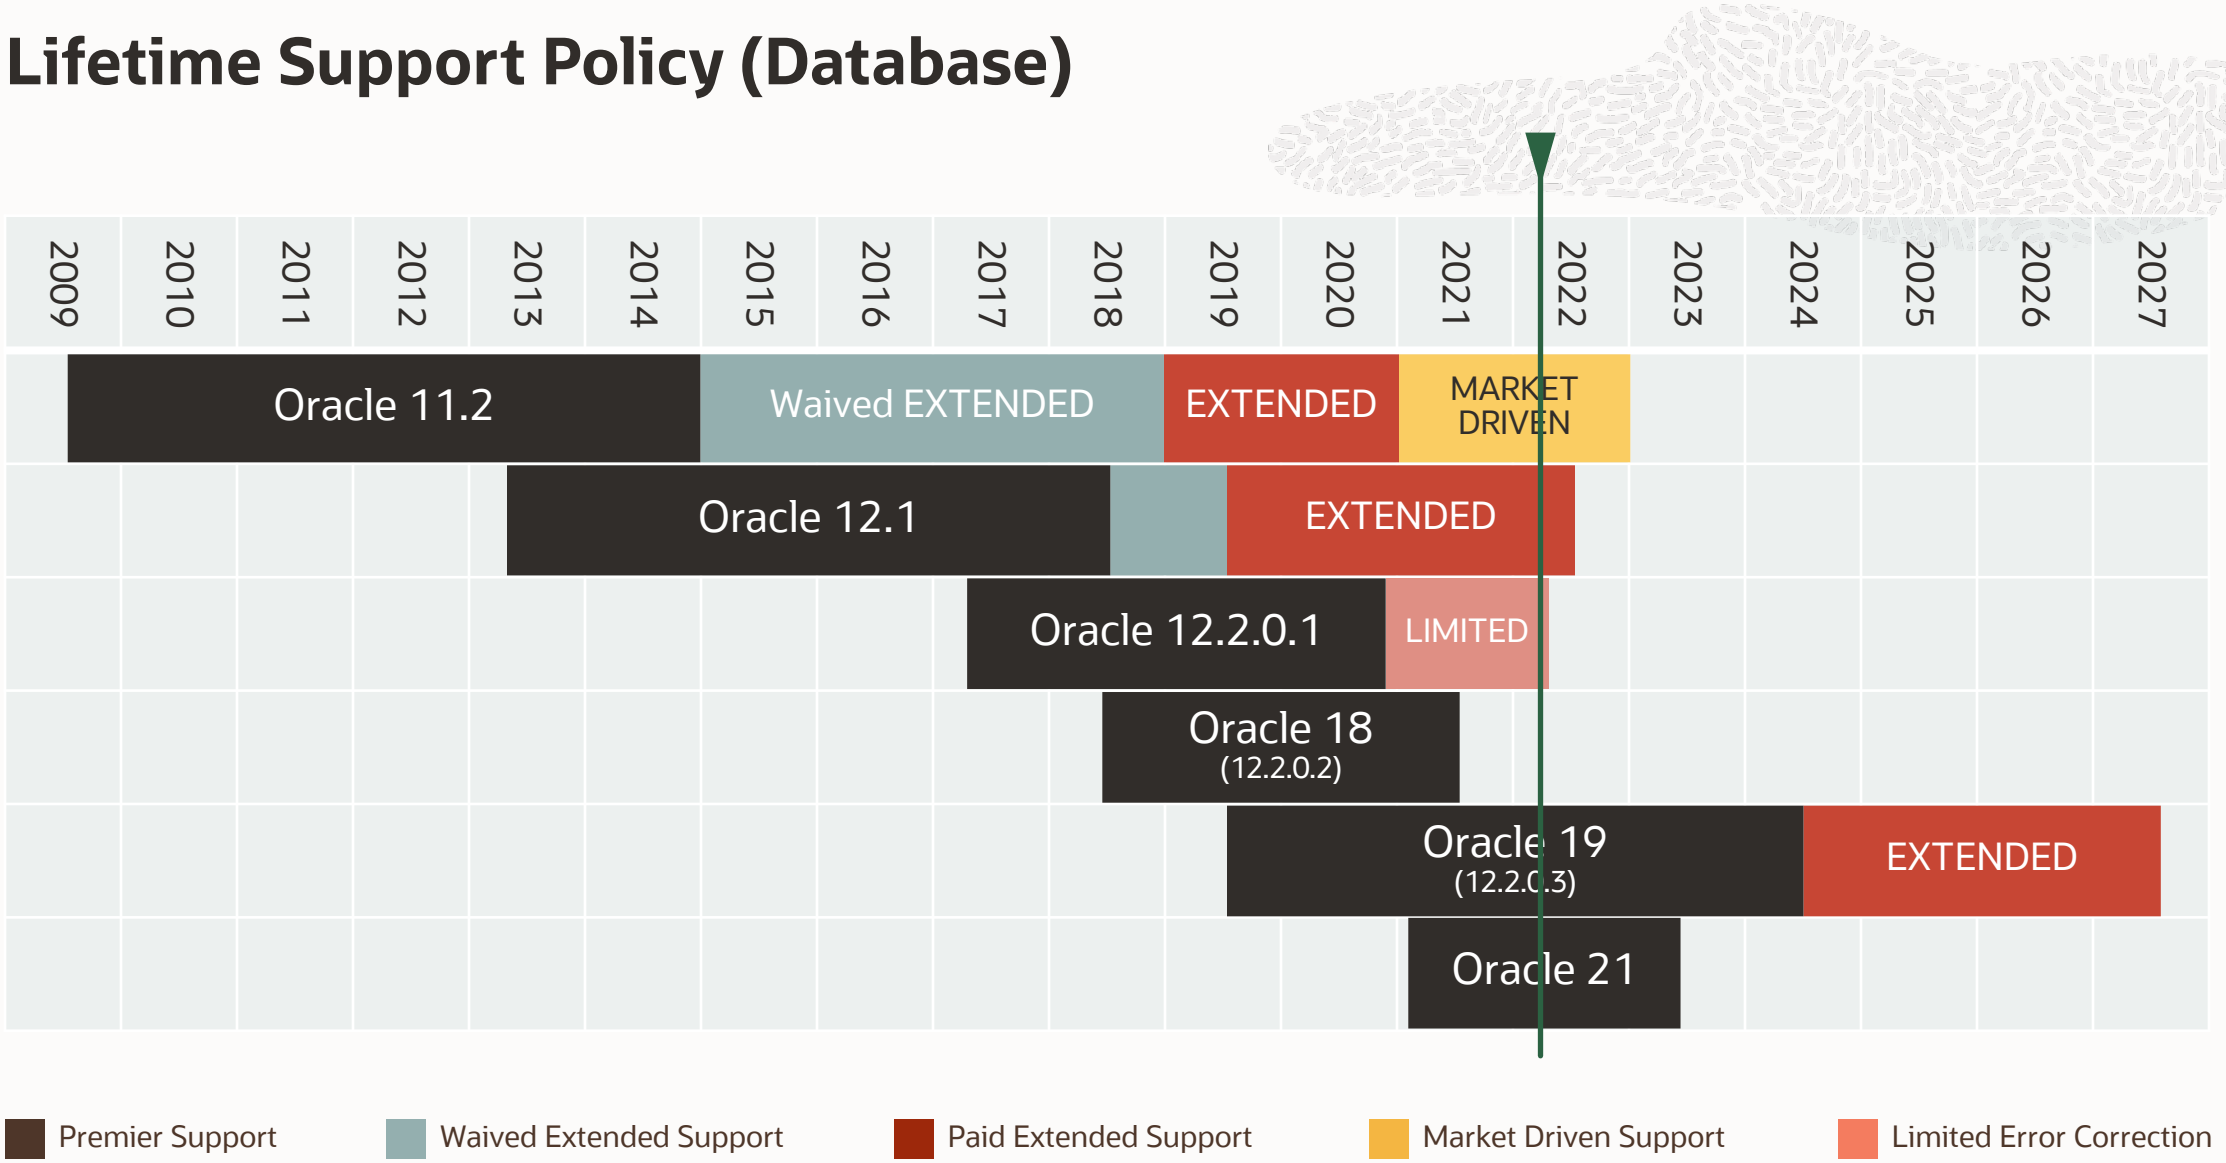

[MOS Note: 742060.1 - Release Schedule of Current Database Releases]

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 RDMA over Converged Ethernet (RoCE) switch	Rolling RoCE switch updates Data Guard Standby First

Reduce Risk and Downtime with Data Guard

Data Guard Standby First Patching (MOS 1265700.1)

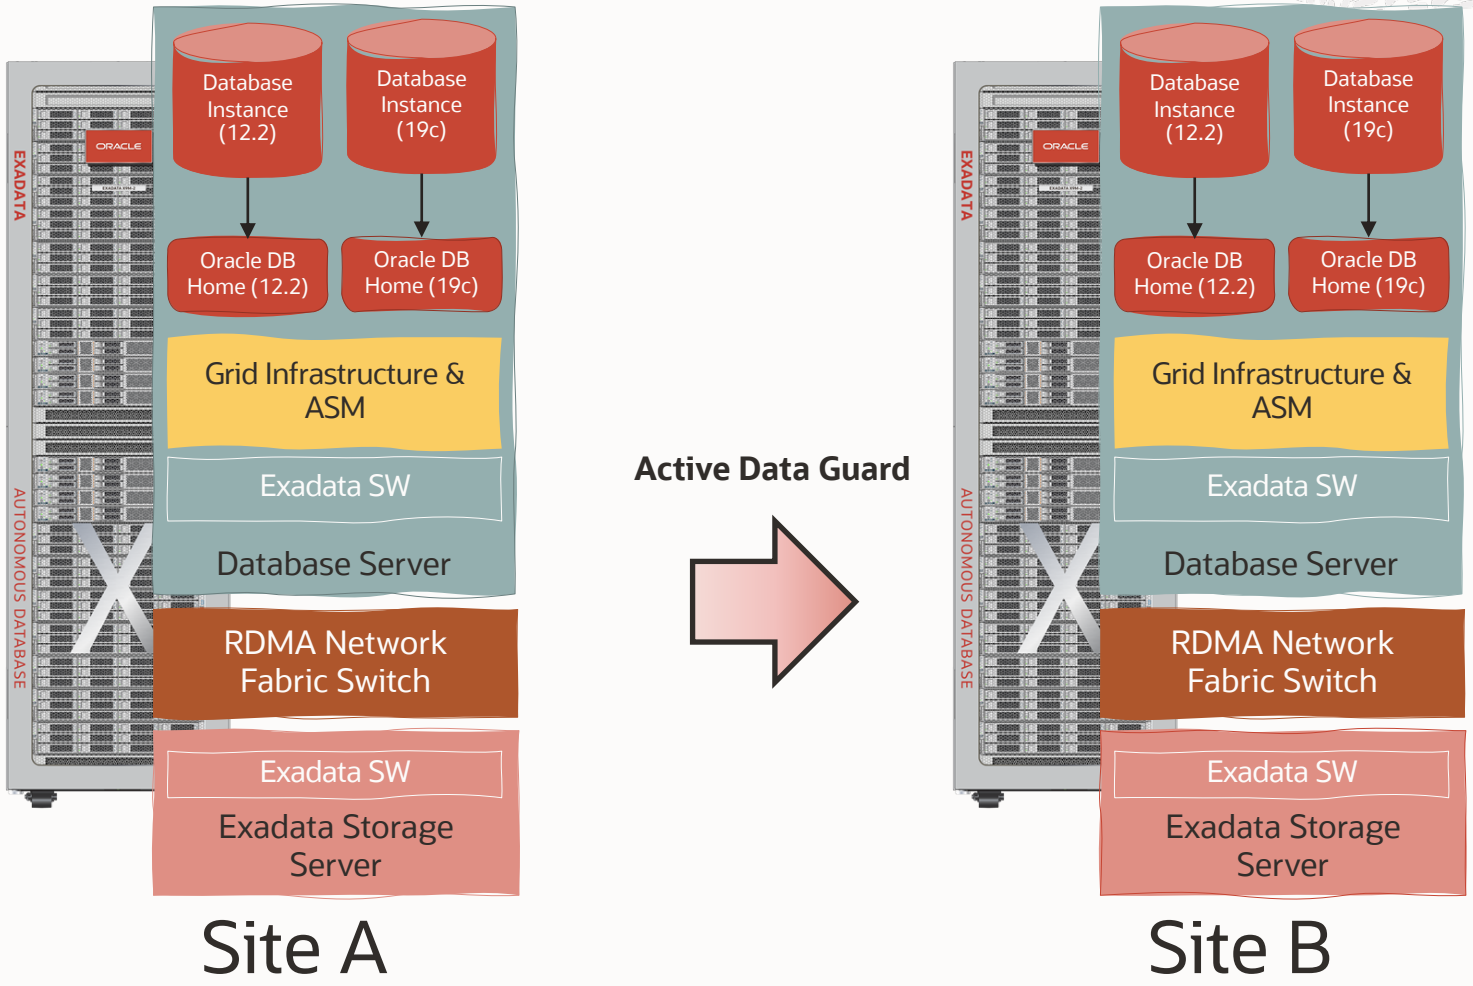

Standby First Patching Steps
1. Update software on Site B (Standby)
2. Test new software
3. Switchover (optional)
4. Update software on Site A
5. Run SQL portion of RU on Primary databases
6. Switchback (optional)

Data Guard Standby First Patching

Update Software on Site B and perform testing

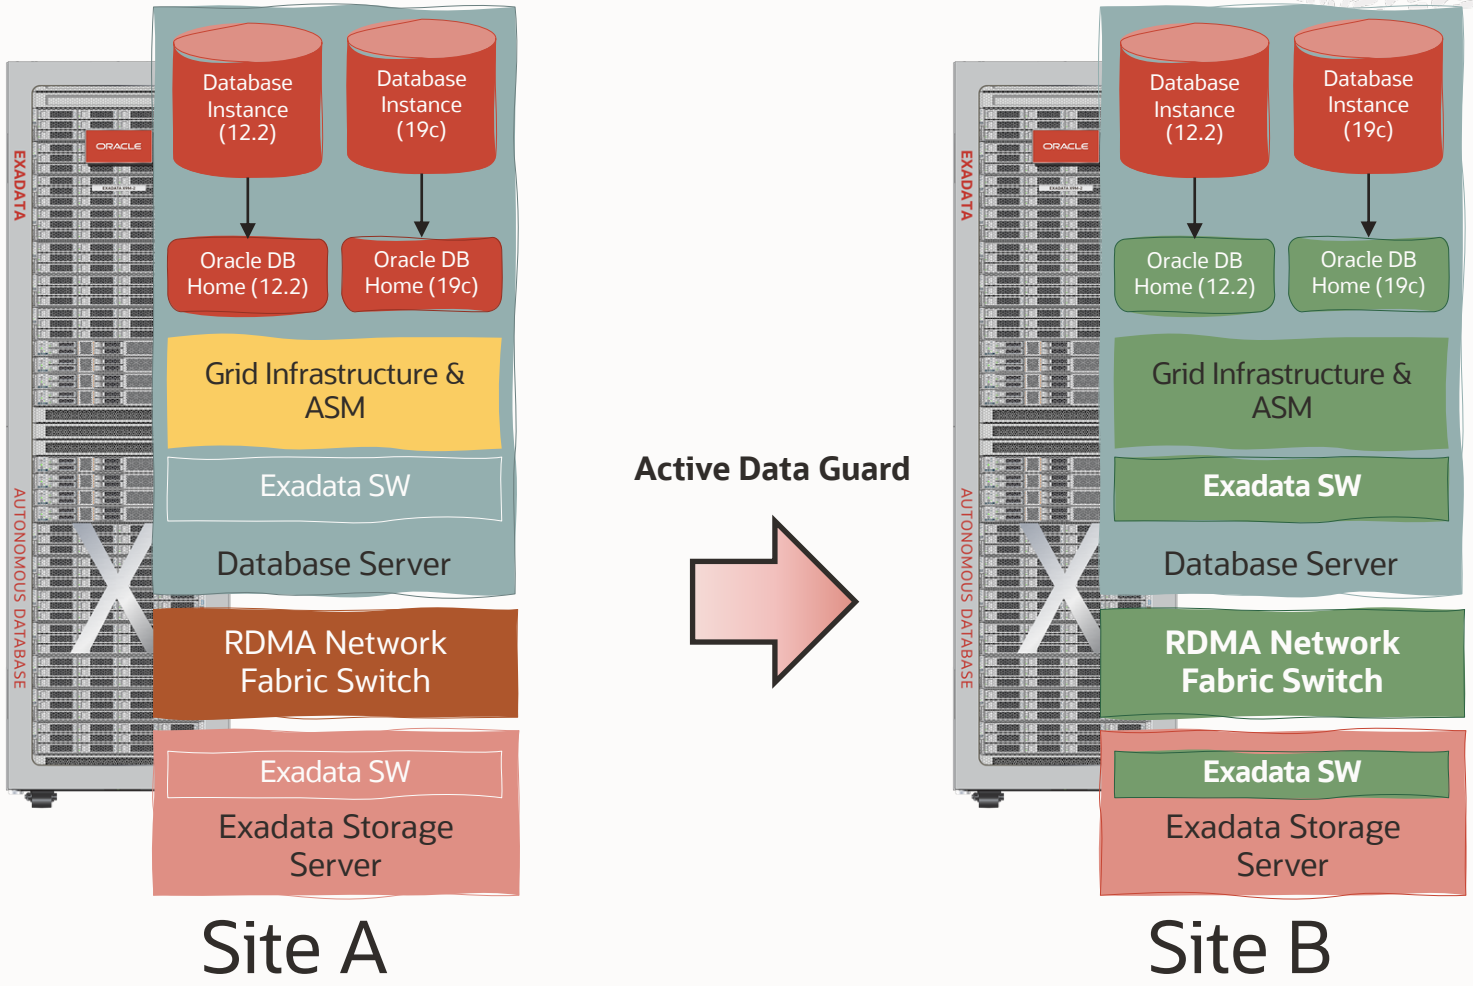

Standby First Patching Steps
1. **Update software on Site B (Standby)**
2. **Test new software**
3. Switchover (optional)
4. Update software on Site A
5. Run SQL portion of RU on Primary databases
6. Switchback (optional)

Data Guard Standby First Patching

Switchover

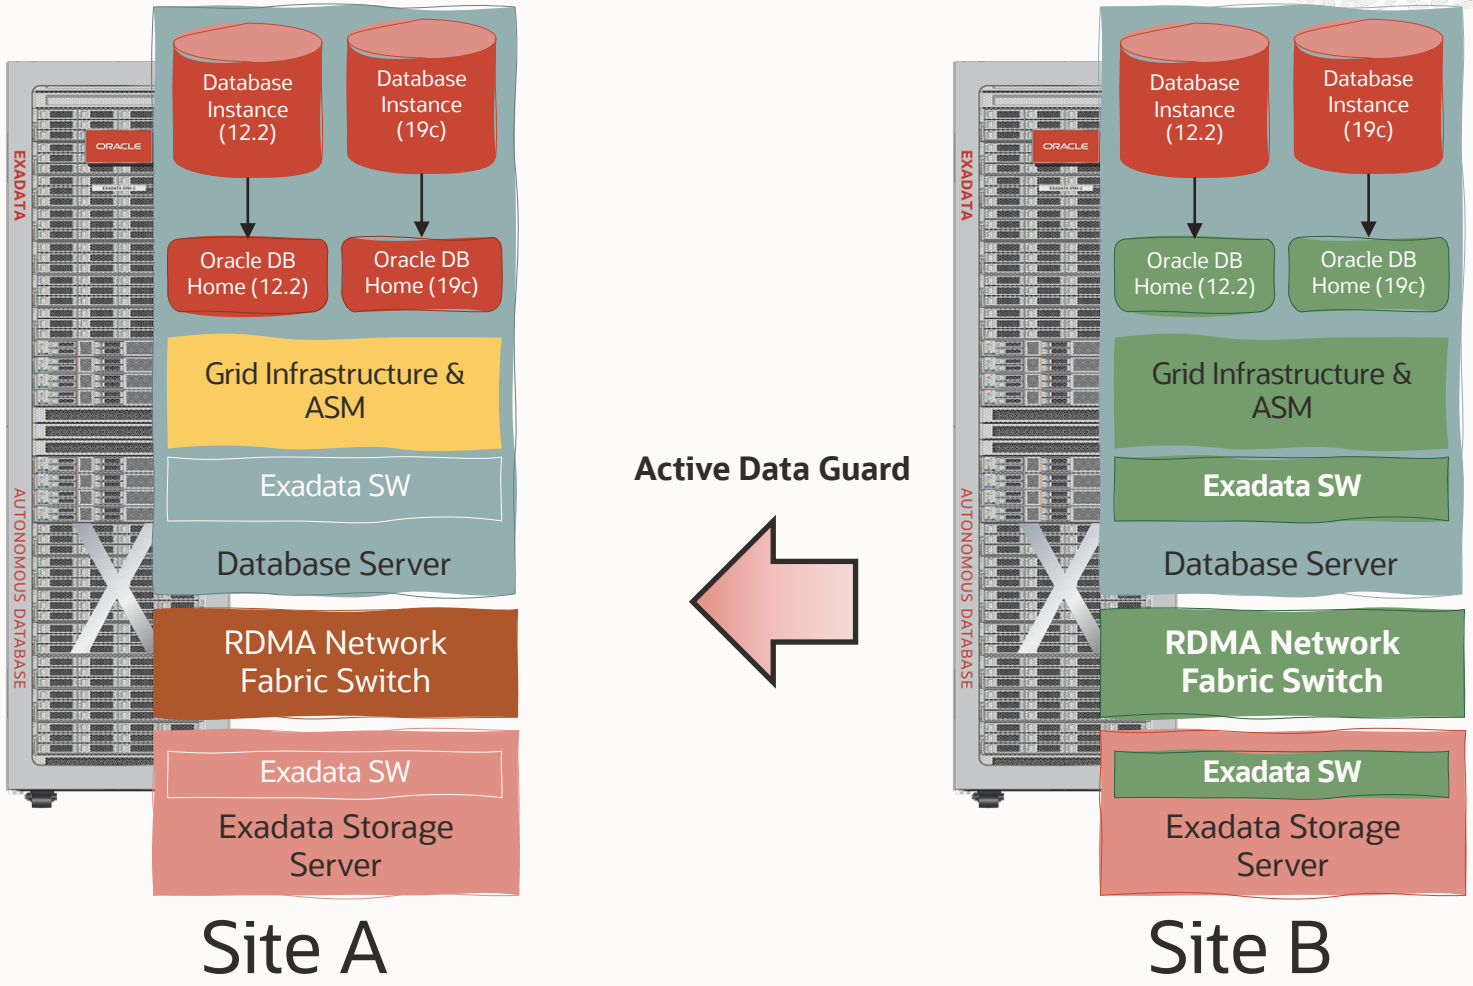

Standby First Patching Steps
1. **Update software on Site B (Standby)**
2. **Test new software**
3. **Switchover (optional)**
4. Update software on Site A
5. Run SQL portion of RU on Primary databases
6. Switchback (optional)

Data Guard Standby First Patching

Update Software on Site A

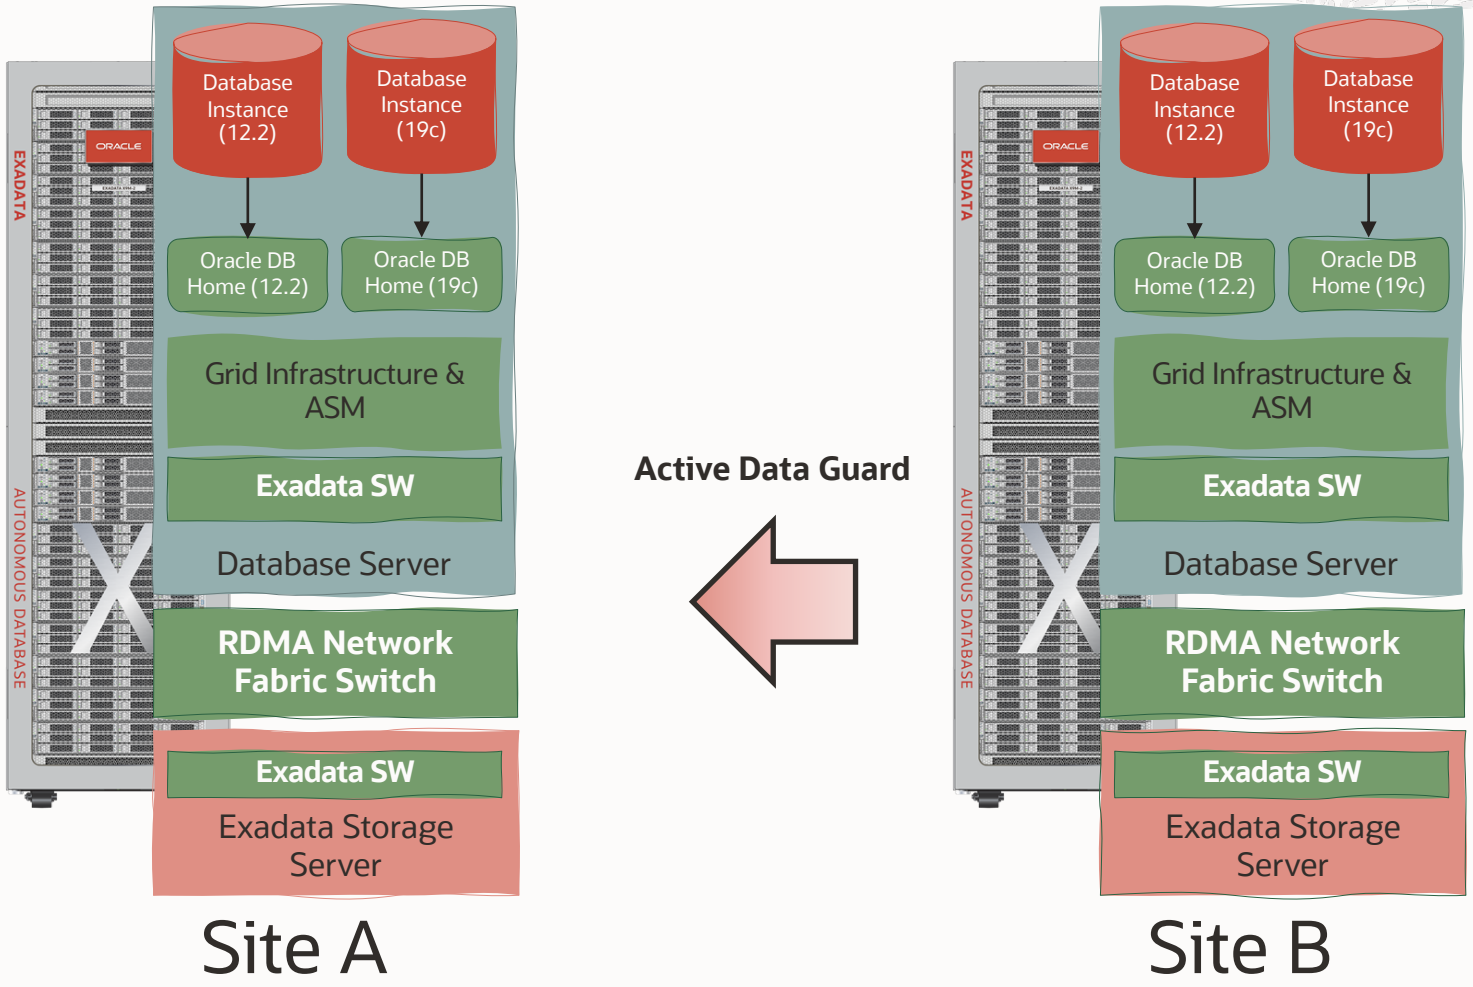

Standby First Patching Steps
1. **Update software on Site B (Standby)**
2. **Test new software**
3. **Switchover (optional)**
4. **Update software on Site A**
5. Run SQL portion of RU on Primary databases
6. Switchback (optional)

Data Guard Standby First Patching

Apply SQL portion of RU on primary databases

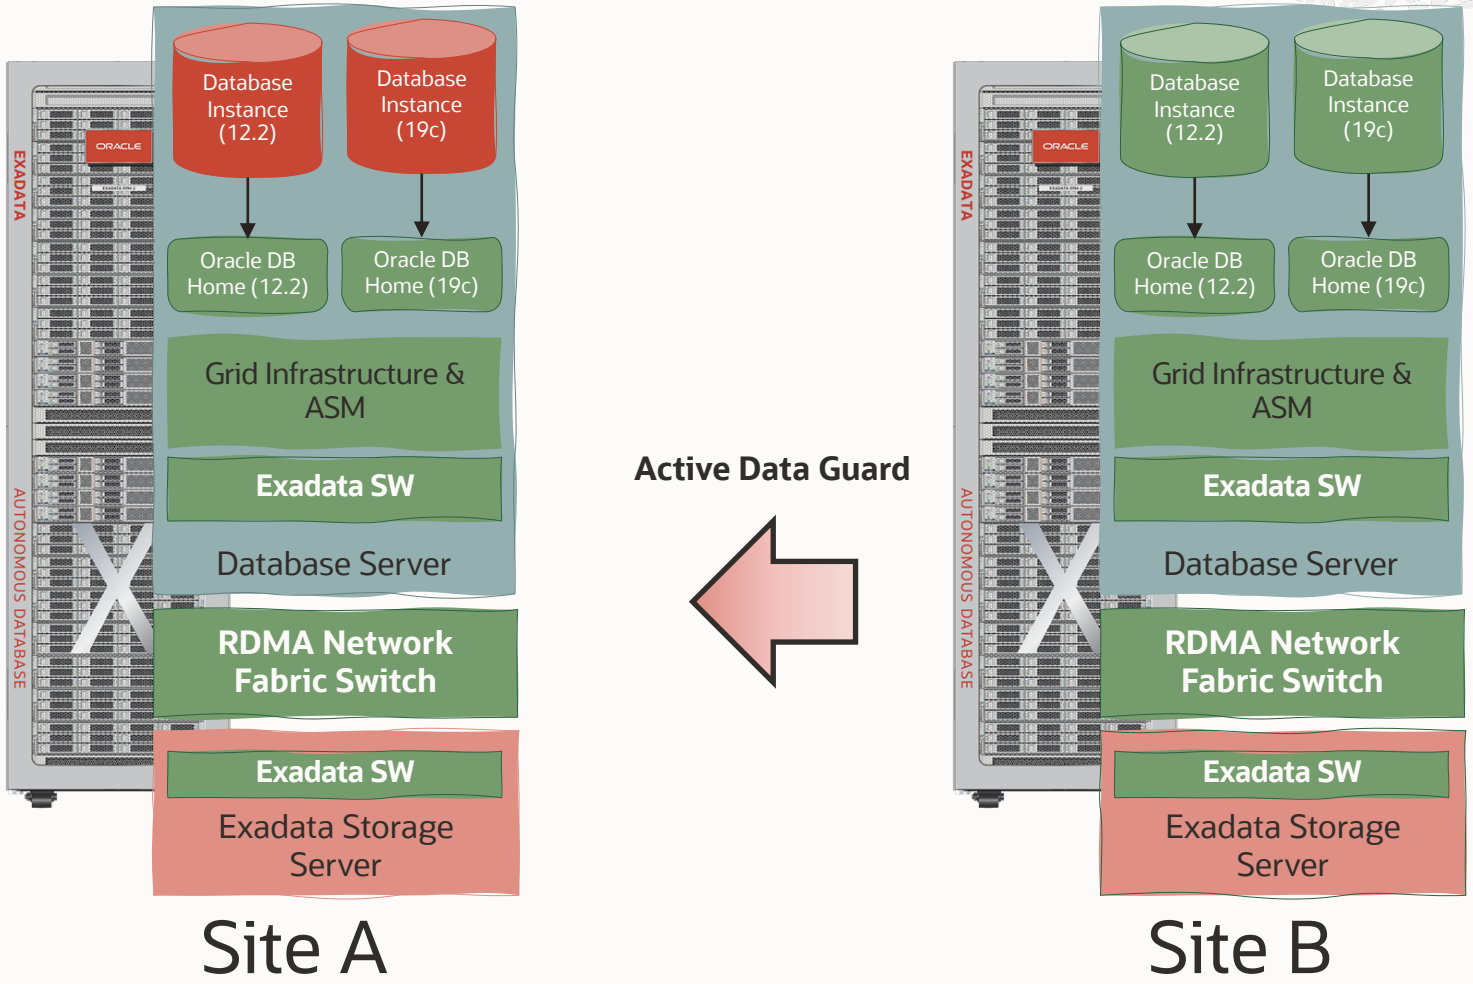

Standby First Patching Steps
1. Update software on Site B (Standby)
2. Test new software
3. Switchover (optional)
4. Update software on Site A
5. Run SQL portion of RU on Primary databases
6. Switchback (optional)

Data Guard Standby First Patching

Switchback

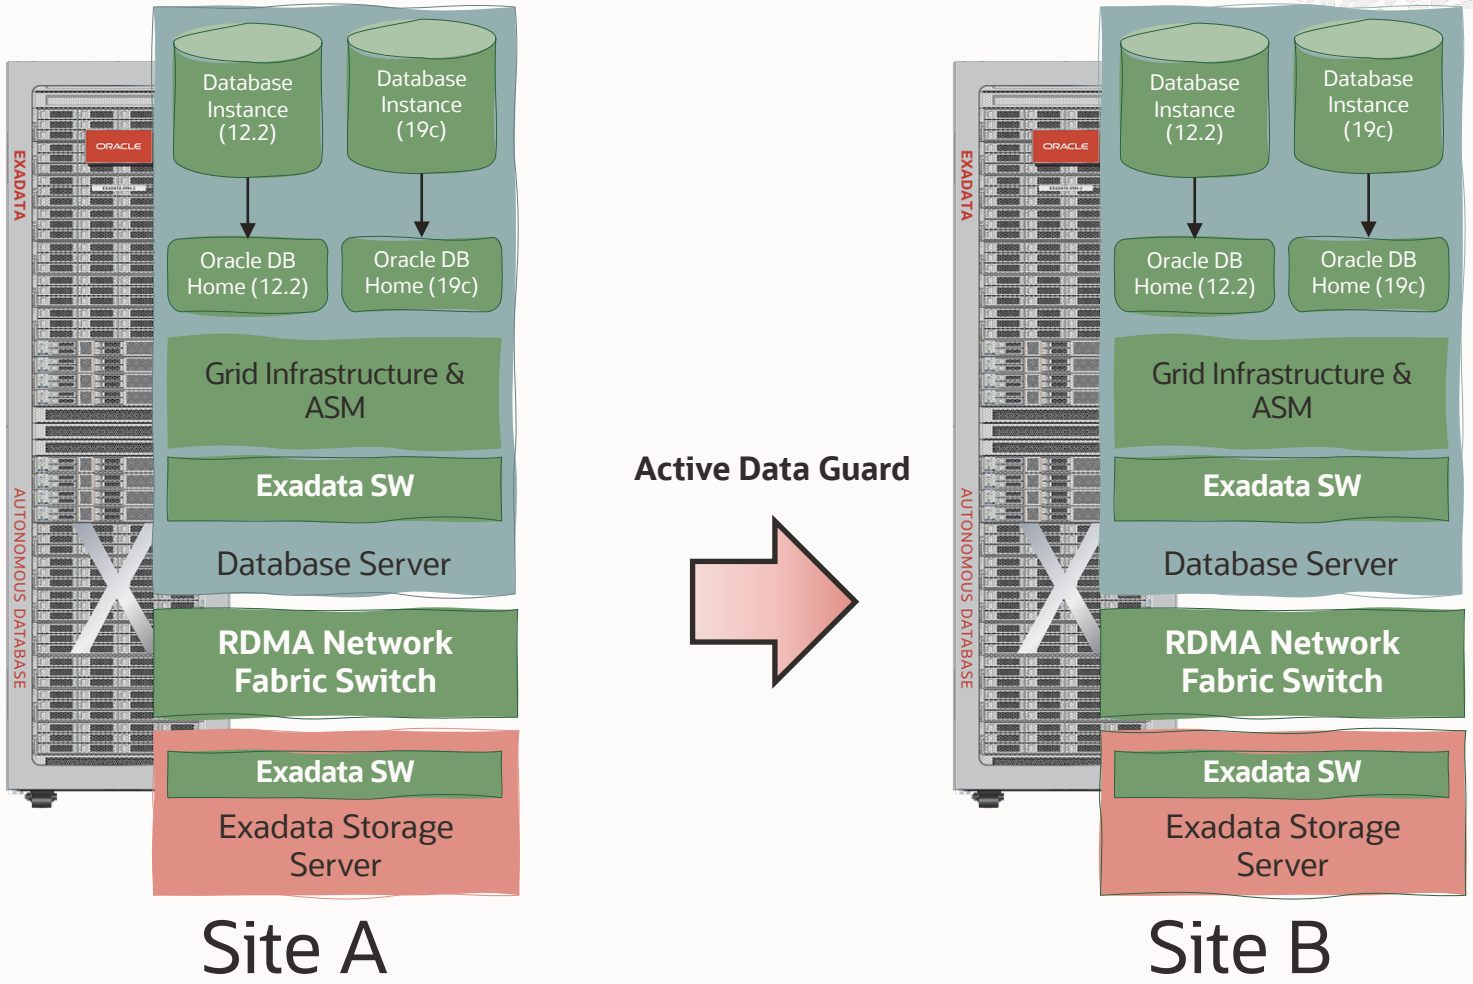

Standby First Patching Steps
1. Update software on Site B (Standby)
2. Test new software
3. Switchover (optional)
4. Update software on Site A
5. Run SQL portion of RU on Primary databases
6. Switchback (optional)

Patching and Upgrading Oracle Exadata
Agenda

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

Updating Exadata Software

High Level Software Maintenance Flow

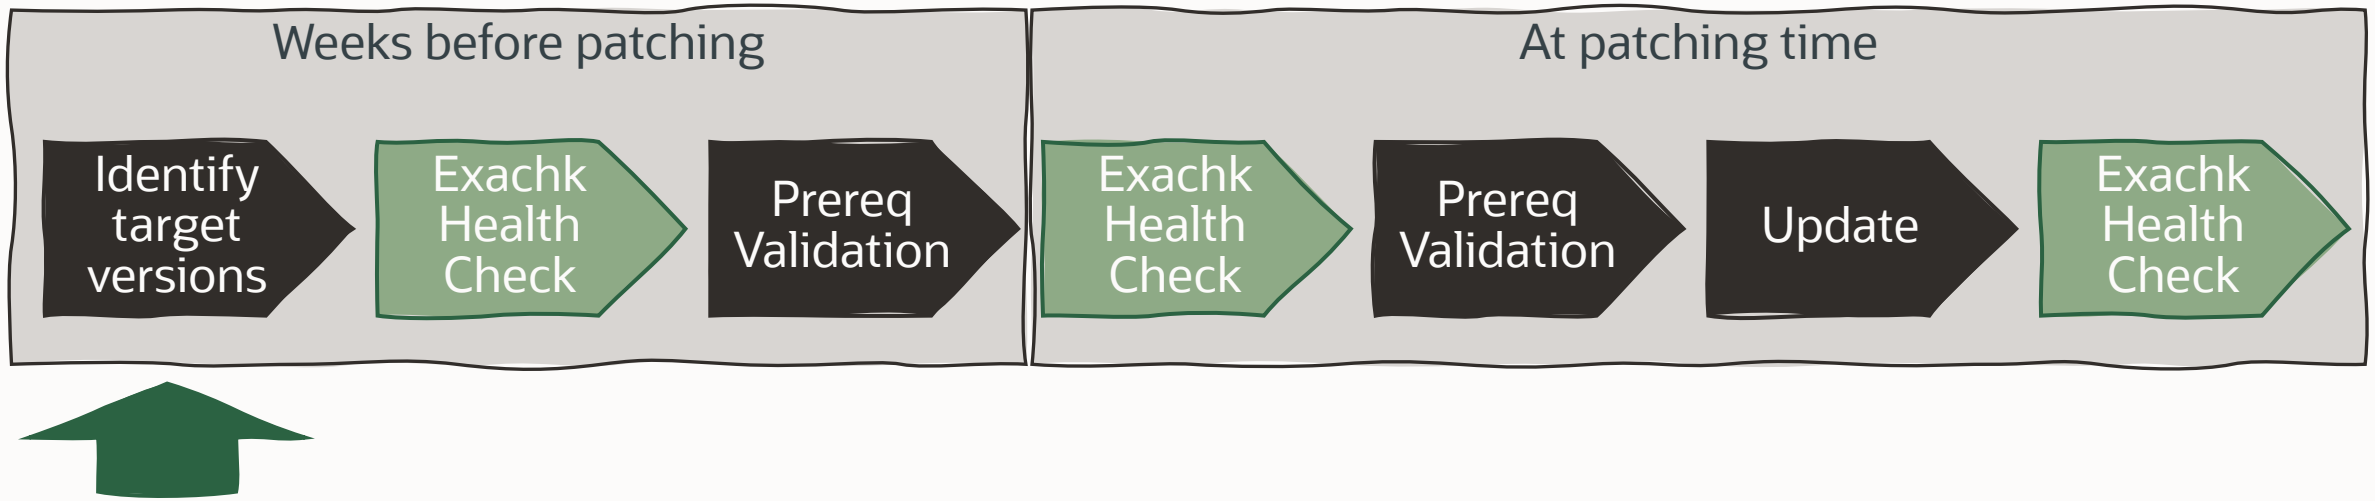

Inputs

- 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

Feature Releases

- **Feature 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 19c requires Exadata 19.3.0)
  - 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 DBHOME19_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 Best Practices

- Graceful Application Switchover in RAC with No Application Interruption (MOS 1593712.1)
  - 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
    - Clusterware-managed database Services
    - FAN-aware connection Pools
    - TNS Configuration
    - Transparent Application Continuity

Zero Downtime Software Maintenance

Application Continuous Availability Best Practices

- High-Level Glance - Move Services and Drain Sessions
  - 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)
  - Disconnect remaining long-running sessions and stop database instance(s)
  - Perform desired maintenance
    - Examples: Patch database home software, Update grid infrastructure home software, Exadata dbnode update
  - Start database instance(s) on target node
  - Enable and start services on target node
    - Services move back and begin accepting new connections
  - Repeat steps on next node

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

Database Software

Oracle JavaVM is RAC Rolling Installable

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

Grid Infrastructure and Database Upgrade Keys to Success

Software Maintenance Rules

- Use Application Continuous Availability Best Practices 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 of GI/ASM and Database Homes
- Use AutoUpgrade for Database upgrades
- 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 KVM Host or KVM Guest, but are performed independently

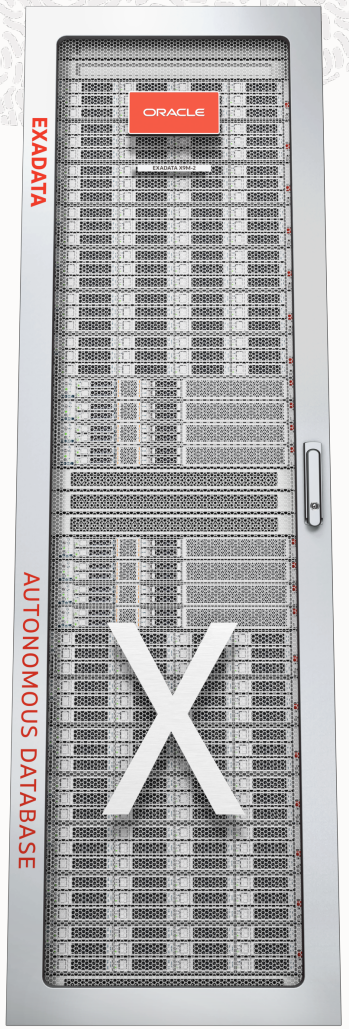

Zero Downtime Software Maintenance

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

Exadata 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:
  - Quarterly - Install Exadata Sustaining Release
  - Monthly - Install Ksplice updates for Oracle Linux kernel security compliance
- **Not Applicable to Exadata Storage Servers**

- Ksplice Offline Client Procedure
  - Obtain uptrack-updates RPM (mirrored ULN repo or download from ULN)
  - Install using yum / rpm command (no reboot)
- HOWTO: Install ksplice kernel updates for Exadata Database Nodes (Doc ID 2207063.1)

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

Exadata 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

Exadata Database Server Custom Configuration

Customization Examples and Impact to Update

	Impact
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

Exadata 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)

Exadata 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

Exadata 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

Exadata Database Server Upgrade

Keys to Success

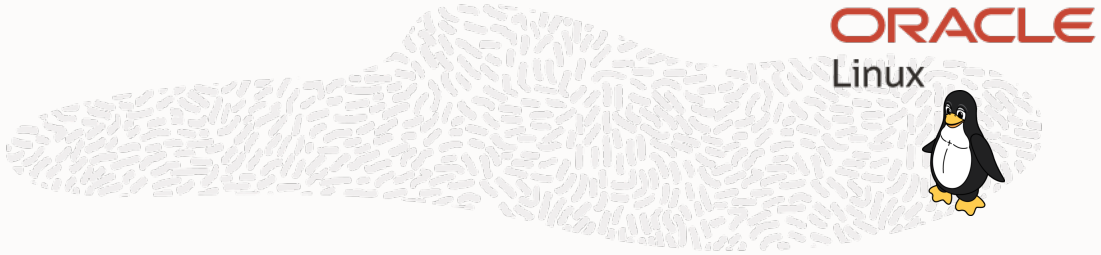

- Database Server Software Maintenance Rules
  - Use Application Continuous Availability to eliminate application impact
  - Customization allowed, but resist - test changes to avoid latent patching failures
  - Closely track customizations, automate build-up and teardown
  - Qualify maintenance readiness
  - Always use latest patchmgr (Doc ID 1553103.1)

Upgrade only when exachk and prereq check are clean

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]
```

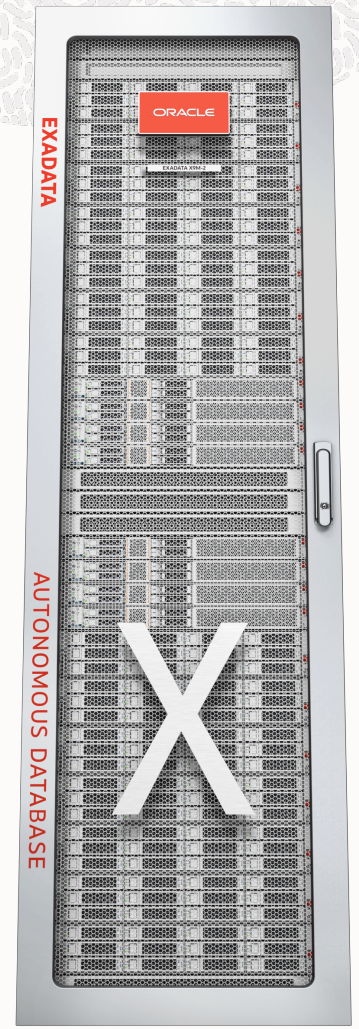

Storage Server Update Flow

Non-Rolling and Rolling

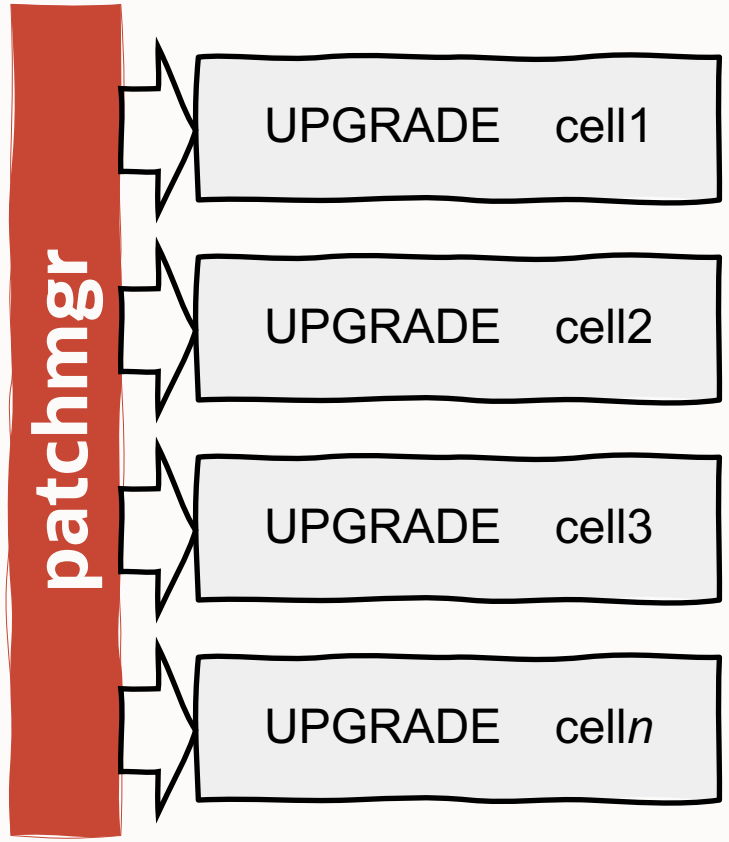

Non-Rolling

- Databases/CRS down
- All cells upgraded in parallel

Rolling

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

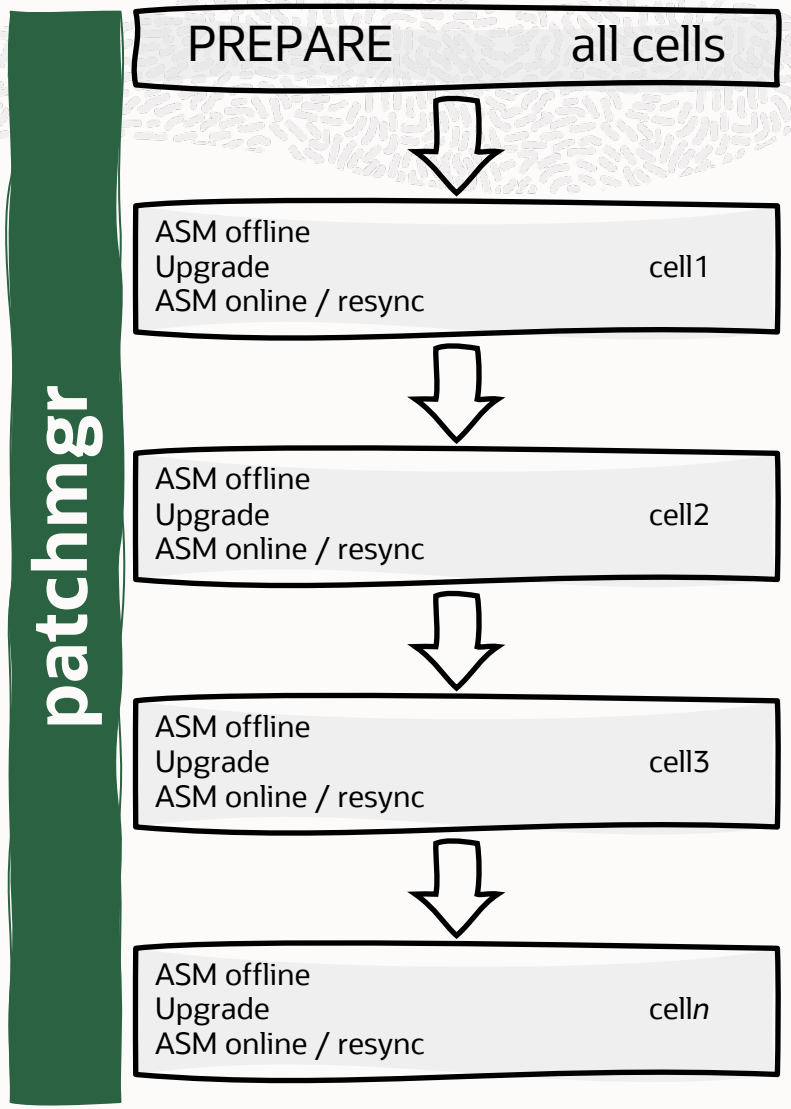

Rolling Storage Server Updates

High-Level Flow

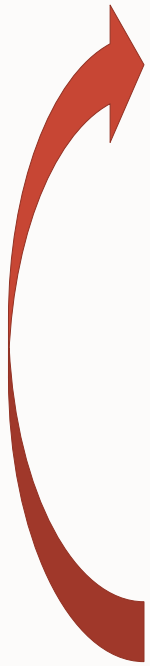

Step	Scope	All Cells in Service?
Prereq check	All cells in parallel	Yes
Copy and stage new software	All cells in parallel	Yes
Inactivate disks	One cell at a time	No
Reboot into new partitions	One cell at a time	No
Complete software update	One cell at a time	No
Update firmware	One cell at a time	No
Validate	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 (sample)

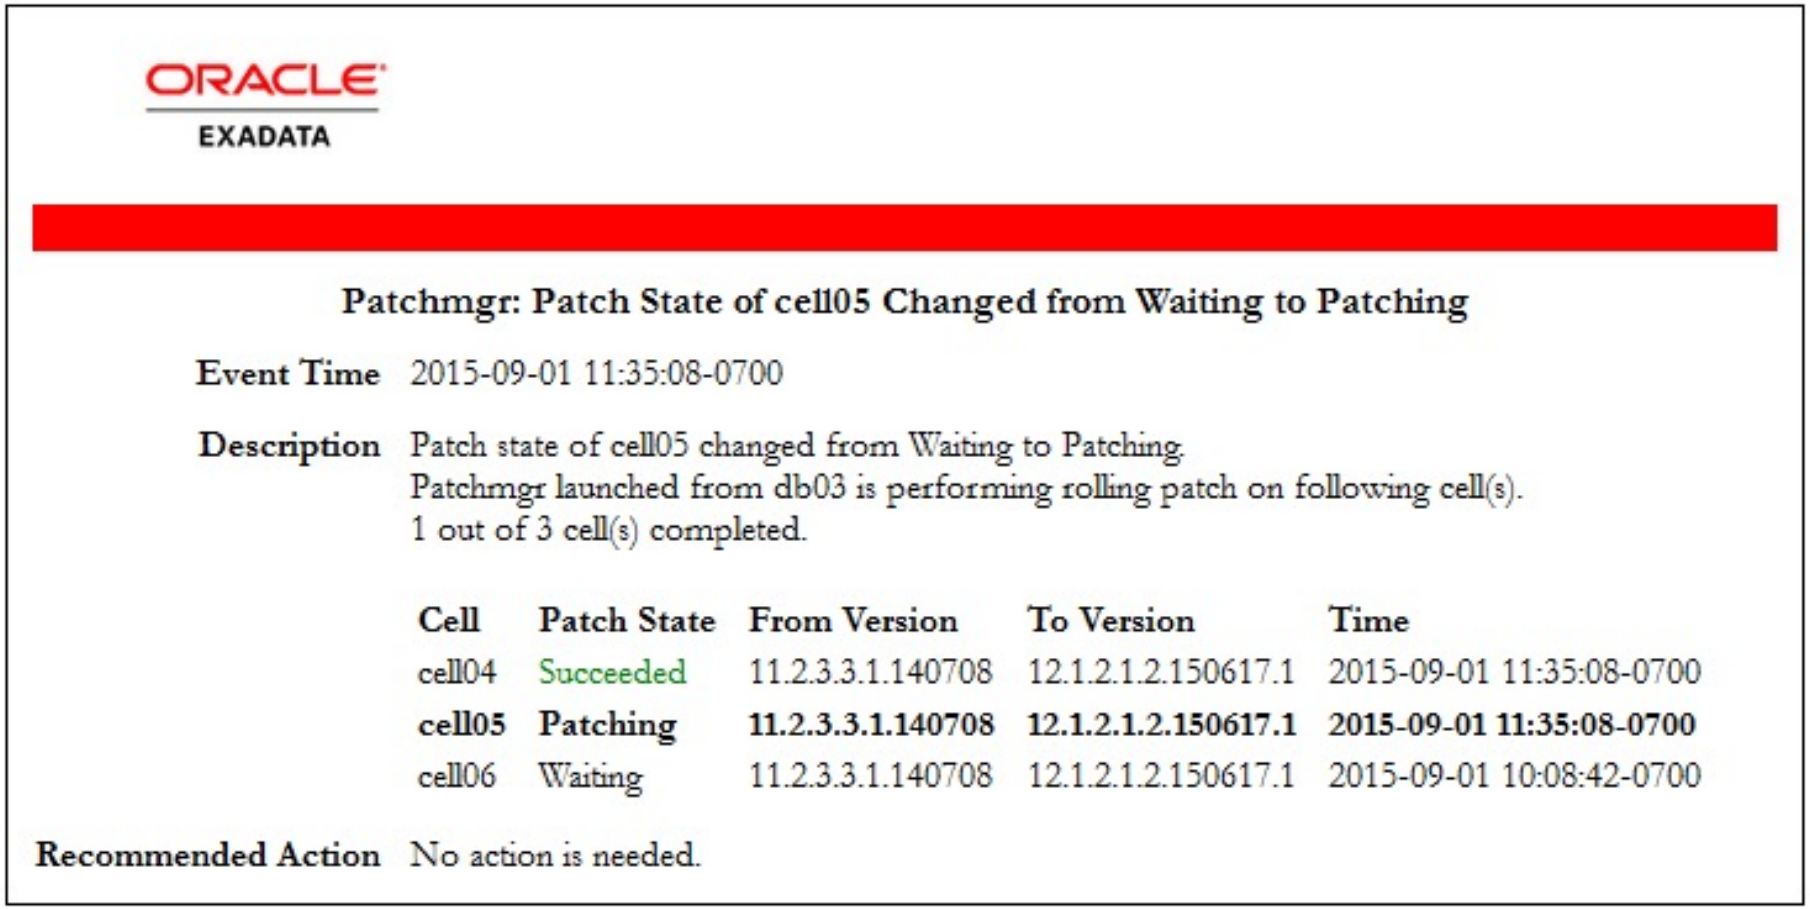

ORACLE EXADATA

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.

Exadata Storage Server Software Update at Scale

Scheduling Automated Updates

- Cell Software Update (see Exadata documentation)
  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
     a) 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

Exadata Storage Server Update

Keys to Success

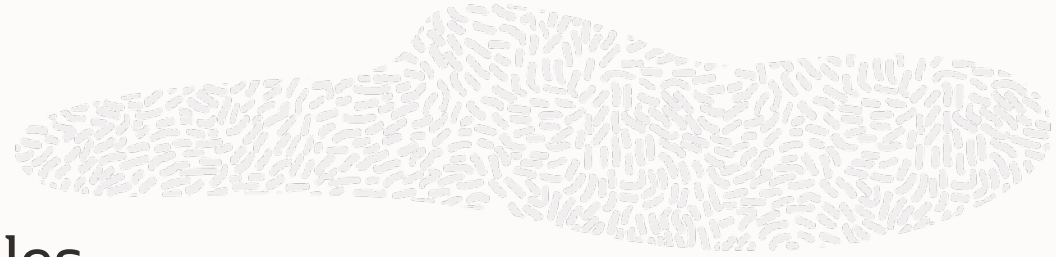

- Storage Server Software Maintenance Rules
  - Do NOT make unsupported configuration changes
  - Qualify maintenance readiness
  - Reduce rolling update disk failure risk w/ ASM HIGH redundancy (or DG)

Upgrade only when exachk and prereq check are clean

Exadata RDMA Network Fabric

Switch Update Steps

- Simple switch update steps
  - Exachk Health Check
  - Prereq validation

```
# patchmgr ---roceswitches --roceswitch-precheck
```

  - Update all RoCE switches

```
# patchmgr --roceswitches --upgrade
```

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

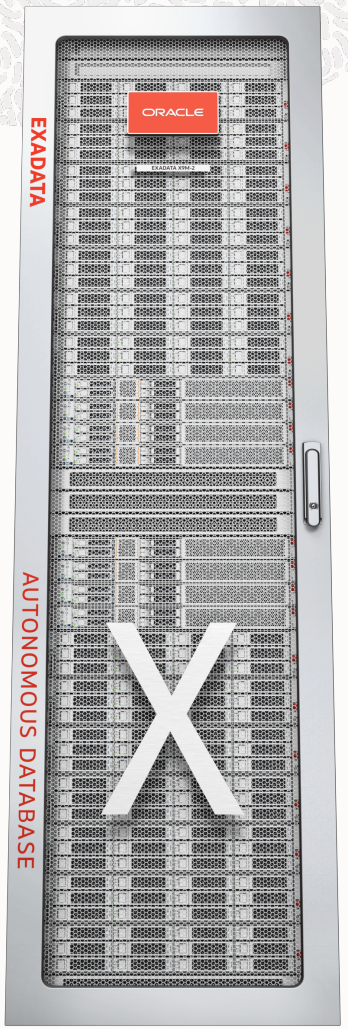

Exadata RDMA Network Fabric Update

Keys to Success

- RDMA Network Fabric Switch Software Maintenance Rules
  - Do NOT make unsupported configuration changes
  - Qualify maintenance readiness

Upgrade only when exachk and prereq check are 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)

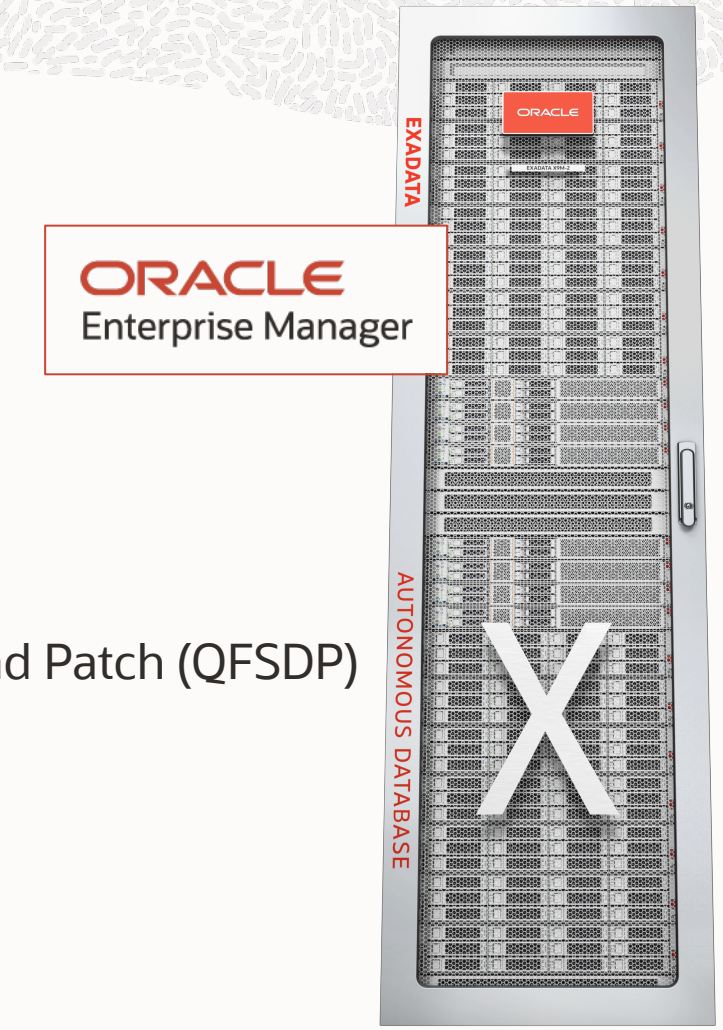

ORACLE Enterprise Manager

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

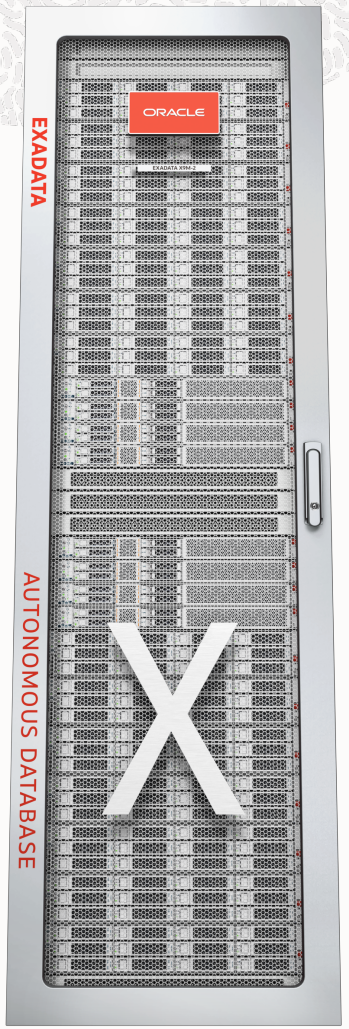

Exadata Software Maintenance

References

- Exadata Product Management
  - Blog: Monitoring Oracle Exadata MOS Notes
    - https://blogs.oracle.com/exadata/post/exadata-mos-notes

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


ORACLE