


ORACLE

Oracle Exadata Database Machine

Operational Best Practices

Exadata Product Management

January 2024



Exadata Vision

Extreme scalability, performance and availability for all data workloads, at the lowest cost



Ideal Database Hardware

Scale-out, database optimized compute, networking, and storage

Database Aware System Software










Unique algorithms vastly improve the performance of modern operational and analytic apps for all use cases at any scale

Automated Management

Fully automates and optimizes the end-to-end stack

Leading organizations in every industry rely on Exadata

10 of the Top 10 Telecoms | 10 of the Top 10 Banks | 9 of the Top 10 Retailers

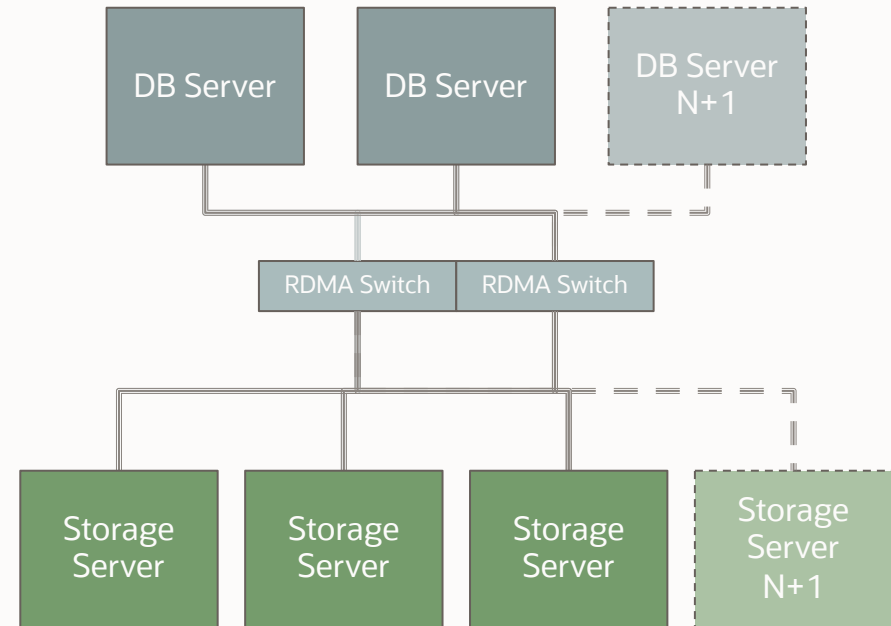
10 of top 10	Banking				
10 of top 10	Telecommunications				
10 of top 10	Food & Drug				
9 of top 10	Automotive				
9 of top 10	Healthcare				
9 of top 10	Retail				
8 of top 10	Technology				
7 of top 10	Energy				

* Ranking based on 2022 [Fortune Global 500](#)



Exadata Architecture

- Exadata is ***THE*** Best platform to run Oracle Database
- Redundancy built in across hardware
 - Redundant Database Servers
 - Redundant Storage Servers
 - Redundant RDMA Fabric Switches
 - Redundant PDUs
- Built specifically to run Oracle Database software
 - Automatic Storage Management
 - Oracle Real Application Clusters
 - Active Data Guard
 - Application Continuity



Exadata Architecture

- Exadata is ***THE*** Best platform to run Oracle Database
- Online* Everything!
 - RAC Rolling Database Updates
 - Standby First Database Upgrades**
 - Exadata Rolling SW Updates
 - Hardware expansion
 - Hardware component replacement

* Database remains available through almost all maintenance conditions

** Minimizes downtime to perform upgrades – use DBMS_ROLLING where applicable




Bare Metal vs Virtualization

Database Server Virtualization

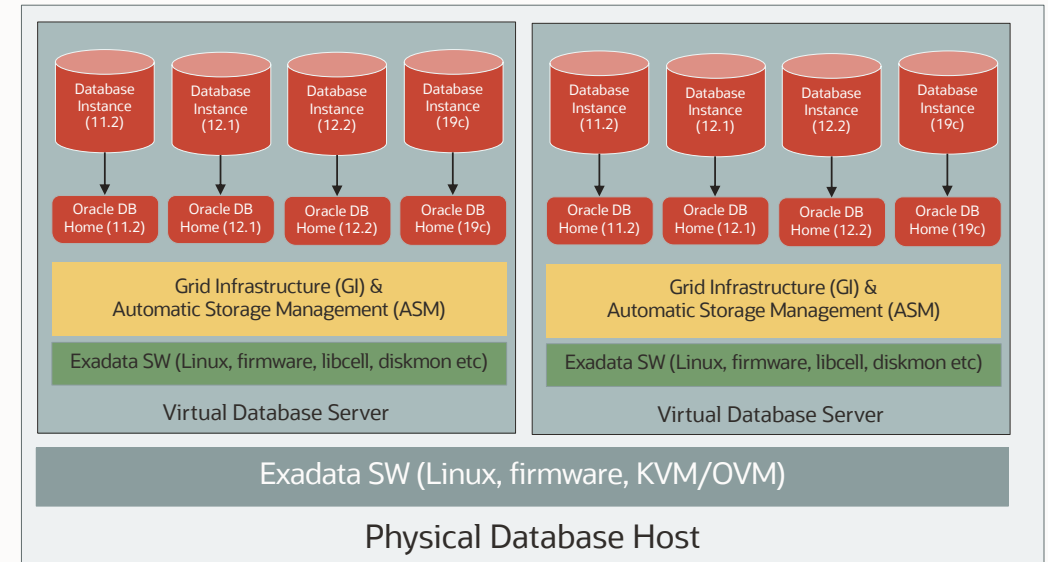
Bare Metal

- 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



Exadata Consolidation Options




- **Dedicated** Database Servers provide the **best isolation**
- Virtualization has good isolation but requires more management overhead and resource usage
 - VMs have separate OS, memory, CPUs, and patching
 - Isolation without need to trust DBA, System Admin
- Database consolidation in a single OS is **highly efficient** but less isolated
 - DB Resource manager isolation adds no overhead
 - Resources can be shared much more dynamically
 - But must trust admins to configure systems correctly
- **Best strategy is to combine VMs with database native consolidation**
 - Multiple trusted DBs or Pluggable DBs in a VM
 - Few VMs per server to limit overhead of fragmenting CPUs/memory/software updates/patching etc.

Exadata Scalability

Horizontal and Vertical Scalability Across Generations

- The physical network architecture of Exadata allows for multiple types of scalability
 - **Vertical scalability within the rack**
 - Add additional database or storage servers until the rack is physically full
 - **Horizontal scalability across multiple racks**
 - Connect multiple Exadata racks together using Clos 3-stage network
 - **Multigeneration scalability**
 - Add newer generations of database and storage servers to previous generations
 - eg Start with X8M Quarter Rack, add X9M or X10M generation database and storage capacity



Network and Data Center Resource Reservations

Reservations and Future Planning

- When deploying Exadata, ensure network and naming standards can scale
 - Review OEDA naming conventions and standardize hostnames
 - eg. <farm><physical rack><infrastructure><node type><network type><node id>{vm<vm number>}
 - If available, reserve whole network subnets for future use. Not just day-one requirements
 - Allocating subnets for client, DR, backup networks allow scaling quickly

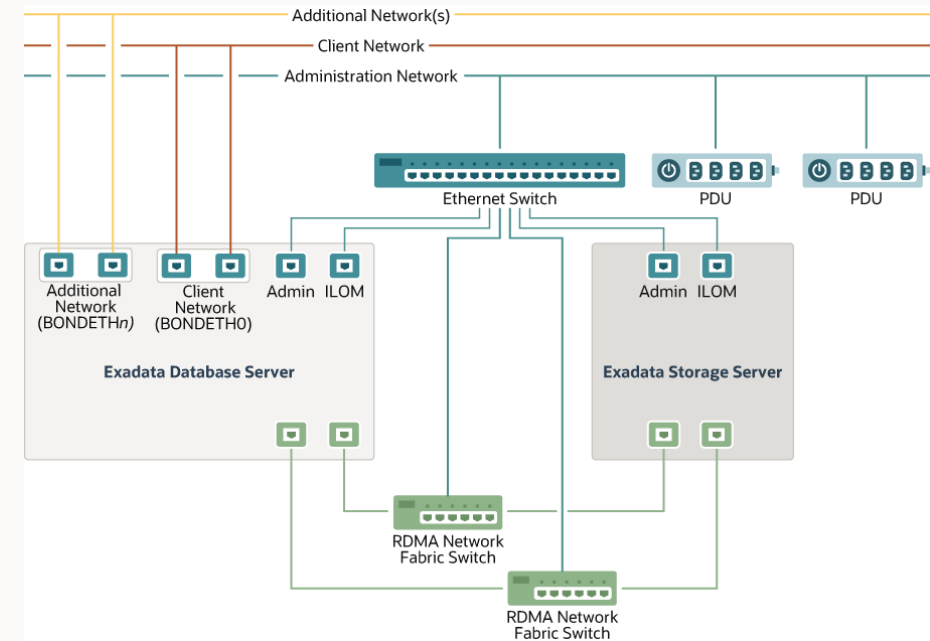
Exadata Farm: New York Production
Physical Rack: 03
Infrastructure Allocation: HR
Node Type: Database Server
Network Type: Admin
Exadata Node ID: 07
Server Type: VM
Server ID: 09
VM Server Name: nyp03hrdbadm07vm09



Network Reservations

Administration Network

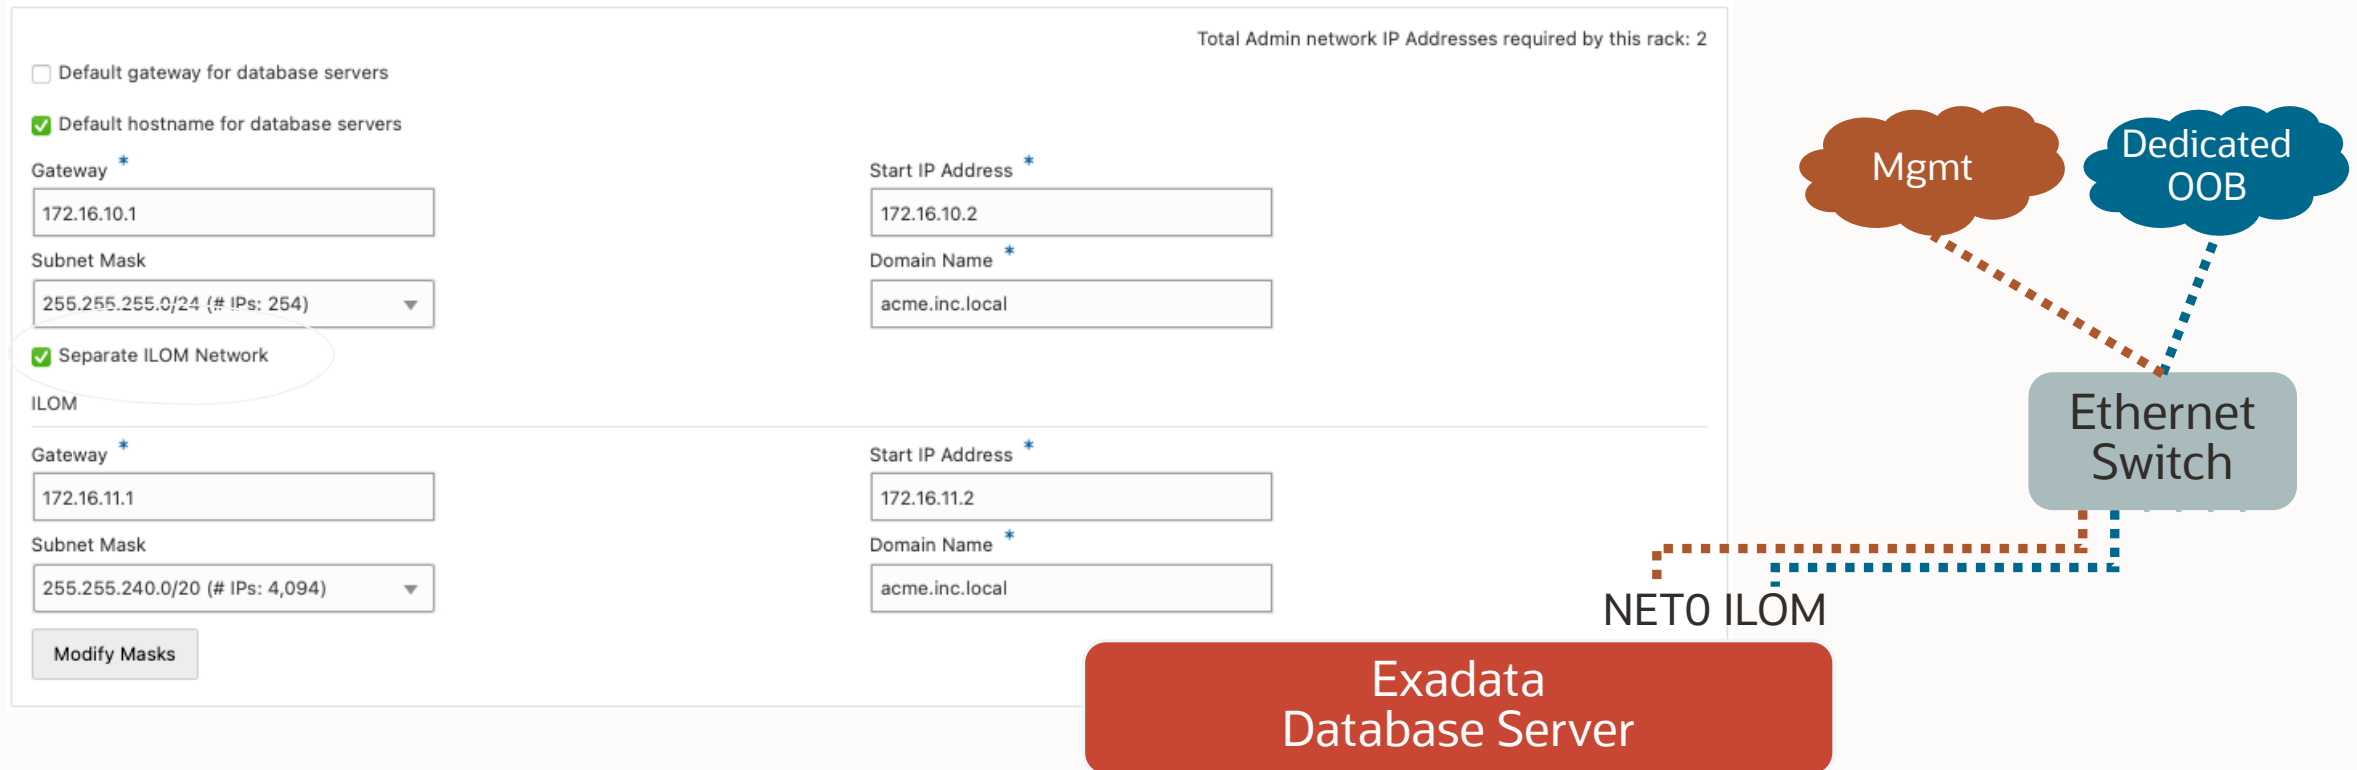
- Provides management and ILOM access to all hardware within the rack
 - CIDR recommendations
 - Ensure the network subnet is large enough to accommodate future expansion
 - An admin network subnet of /27 (30 IP addresses) only covers a quarter rack.
 - A CIDR of /24 (254 IP addresses) can provide easy-to-remember boundaries
 - If local requirements mandate, ILOM and MGMT can be segregated using separate subnets and VLANs through OEDA



Network Reservations

Dedicated Out of Band Network


- Some customers have TFTP / PXE dedicated networks for installation / provisioning
- Identified in OEDA




Network Reservations

Private Network

- RDMA Network Fabric connects the database servers and storage servers
 - Recommend leaving default CIDR for all rack sizes
 - /22 (1022 IP addresses)
 - All racks should use unique ranges within this subnet to reduce risk if future re-rack / multi-rack situations
 - Segregation can be achieved using Secure Fabric for virtual platform or Access VLAN configuration for physical platform



Secure Fabric



Access VLAN



Network Reservations

Client Network


- **Each Cluster Requires Network Configuration**
 - **Client Network** – Used for communicating from Oracle Database to application tier/user (via database listener)
 - **Backup Network** – Used for communicating from Oracle Database Server to backup infrastructure
 - **Other Additional Network(s)** – Used for any other customer requirements (replication, etc)
 - **All Client networks configured on Exadata MUST BE bonded**
- **Configured via OEDA using best practices**
 - Bonding, LACP (optional), VLAN (optional)
 - Upstream hardware configuration requirements
 - Redundant Switches (recommended), bonding, LACP (optional), VLAN (optional)
 - Requires data center services
 - NTP, DNS

Network Configuration

Client Network Interfaces – Bonding

- **Configuration**

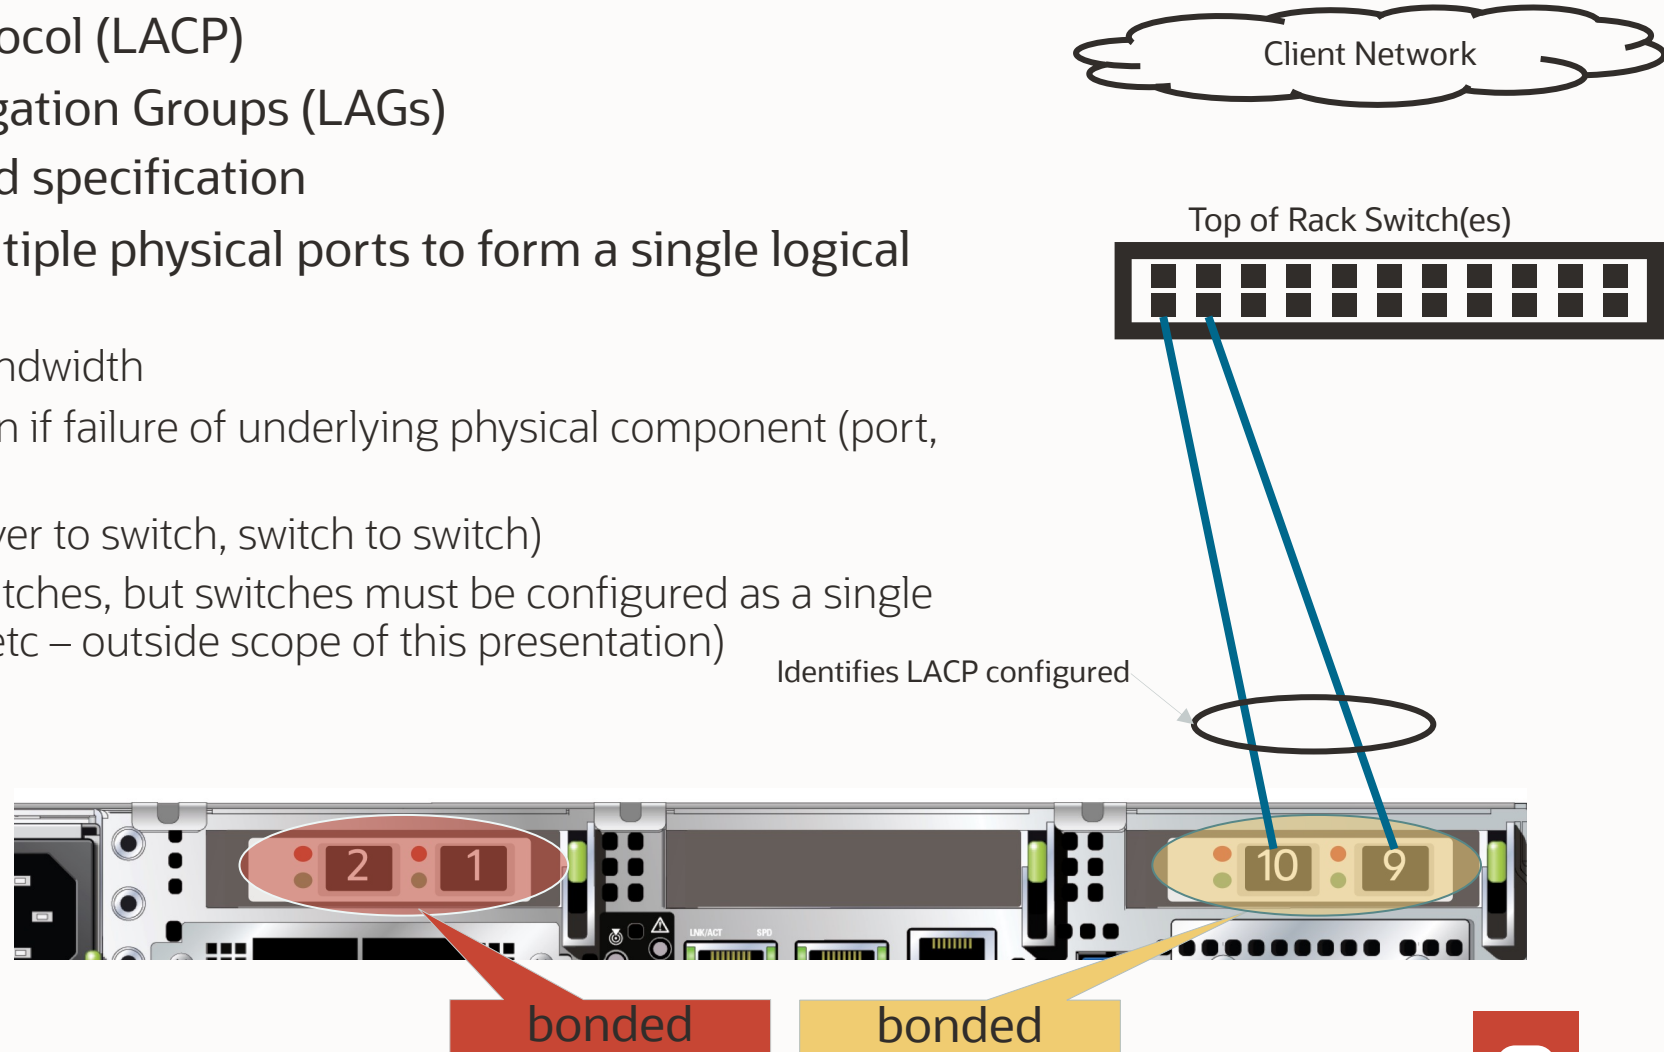
- Select the physical interfaces for network through Oracle Exadata Deployment Assistant (OEDA)
 - Bonded interfaces will be configured on the same NIC
 - The XML configuration file generated by OEDA includes detailed information that maps the bonded network interfaces to the underlying Ethernet ports
 - The bonded network interfaces are named “bondethN”
 - Client: bondeth0, Backup: bondeth1, Additional: bondeth2, AdditionalN: bondethM
- Linux bonding module is configured to use active-backup mode (mode=1).
 - If a different bonding policy is preferred, then reconfigure bonding module after initial configuration
 - For configuration details, refer to the "[Linux Ethernet Bonding Driver HOWTO](#)"
- Upstream network infrastructure (switches) must be capable of supporting the chosen bonding mode
 - Example: If mode 4 is configured (IEEE 802.3ad Dynamic link aggregation), then you must supply and configure the network switches capable of supporting this bonding mode



Network Configuration

Client Network Interfaces – Link Aggregation

- Link Aggregation Control Protocol (LACP)
 - Also known as Link Aggregation Groups (LAGs)
 - Subsection of IEEE 802.3ad specification
 - Allows the bundling of multiple physical ports to form a single logical channel
 - Typically used to increase bandwidth
 - Provides graceful degradation if failure of underlying physical component (port, cable) occurs
 - LAGs are point-to-point (server to switch, switch to switch)
 - Can be cabled to multiple switches, but switches must be configured as a single virtual switch (MC-LAG, VC, etc – outside scope of this presentation)



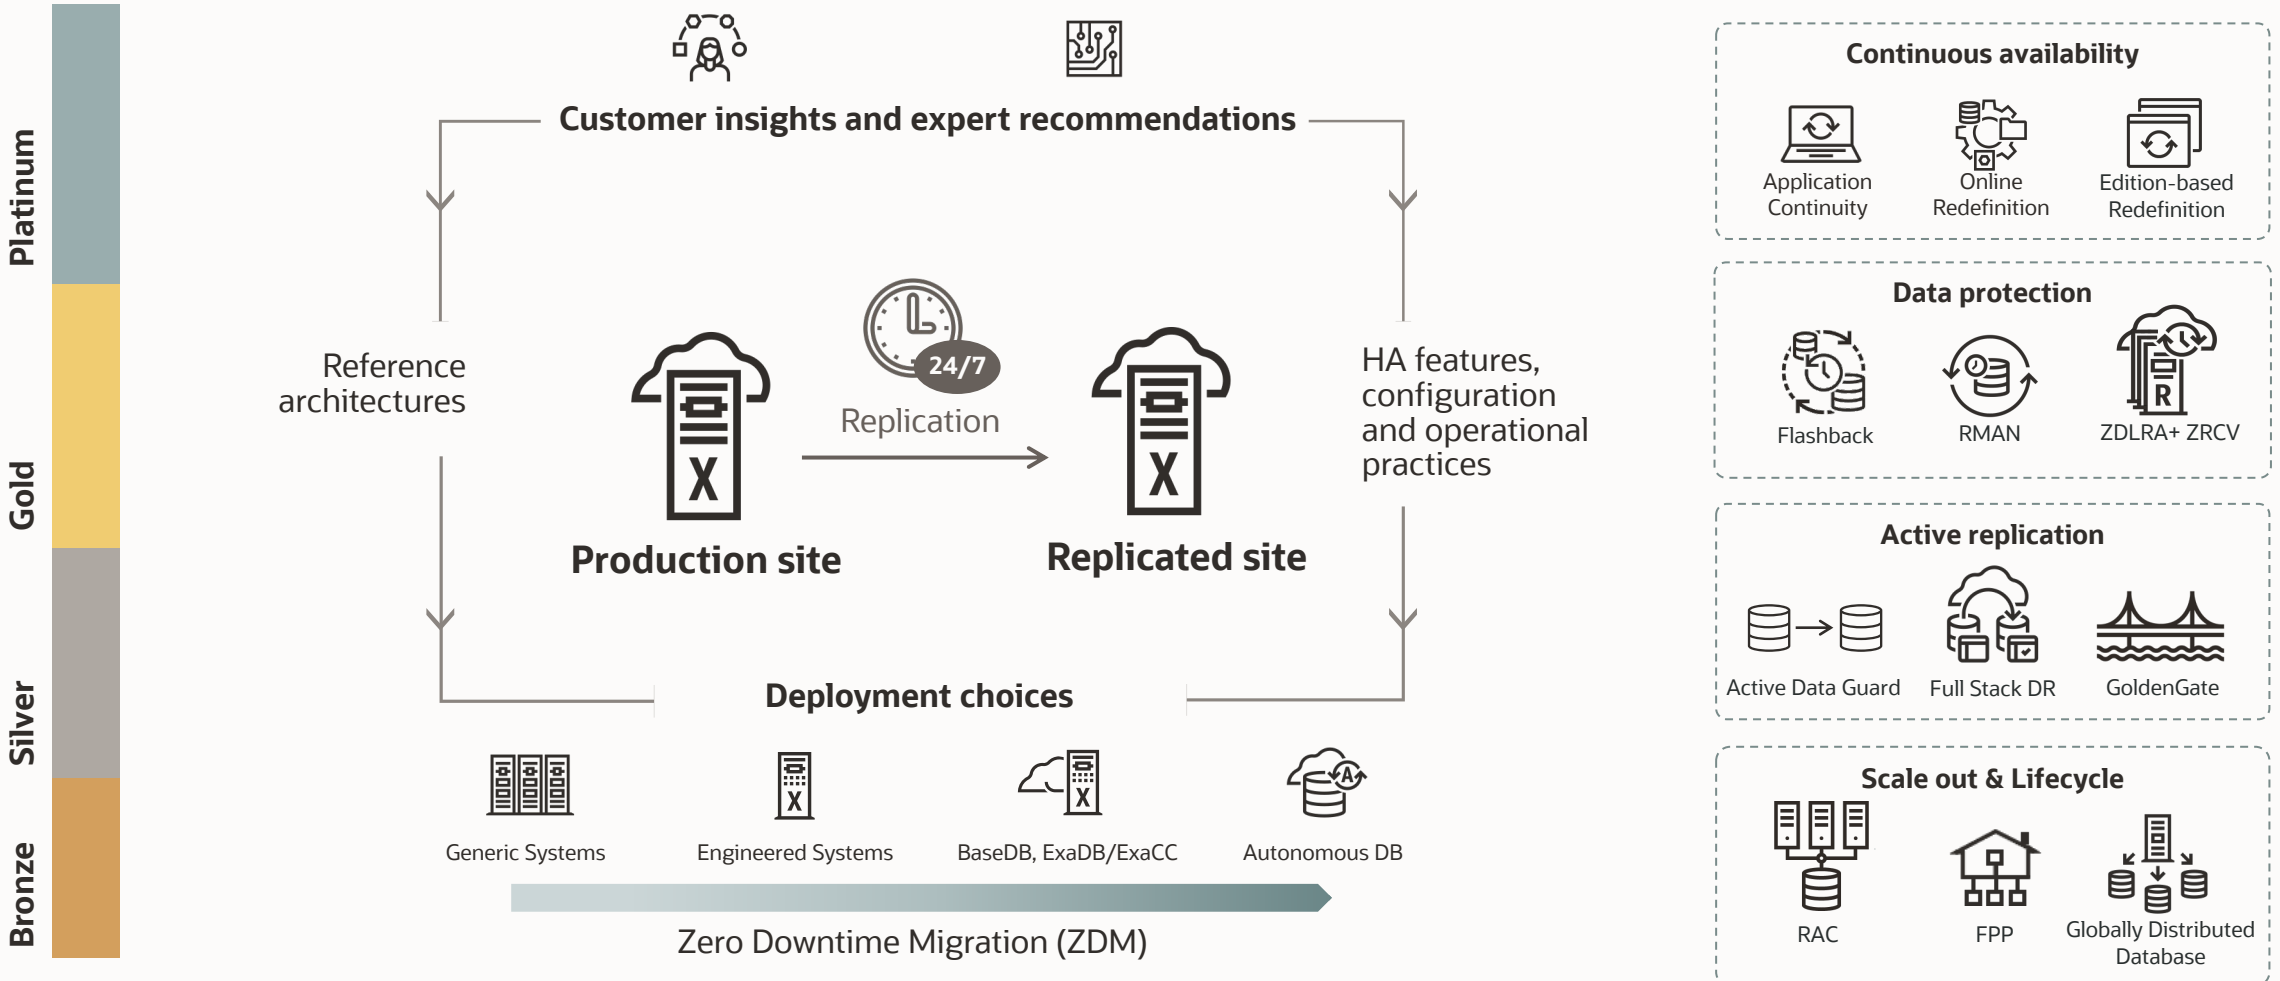
Always Verify Network Before Configuration Begins

- Use the checkip.sh script to ensure there are no IP address conflicts between the existing network and the new Oracle Exadata Rack
 - Available in the ZIP bundle from OEDA
 - The output from the script is a file that contains status messages such as **GOOD** or **ERROR**
- This should be run from a machine that has access to the same network where the Exadata will be deployed **before** the Exadata is connected to the network
 - Does not require Exadata to verify
 - Does not need to be run by Oracle Field Engineering
 - Should be validated before attempting to configure the Exadata

Maximum Availability

Oracle Maximum Availability Architecture (MAA)

Standardized Reference Architectures for Never-Down Deployments



Real Application Clusters

Let's talk about RAC

- RAC horizontally scales Oracle Database across multiple database servers in a shared everything architecture
- Each instance in a RAC Database can be independently controlled and patched while maintaining database and application availability
- RAC is **Highly|Strongly|Extremely** recommended but not mandatory
- What about databases that don't NEED or don't "LIKE" RAC?
 - See above point – but remember, you can use a SINGLETON service to force all connections to a specific instance
 - SINGLETON failover to an already running instance on another Database Server is faster than going through crash recovery
 - Useful for non-production databases



Storage Configuration

High Redundancy for the win!

- ASM used on Exadata to implement Stripe And Mirror Everything (SAME)
- Exadata supports multiple ASM redundancy modes
 - NORMAL (NR) – double mirroring
 - **HIGH (HR) – triple mirroring**
 - FLEX – mode that allows file/file group level redundancy attribute. i.e. mix of NR & HR defined at the datafile or group of datafiles
- Recommend **High Redundancy** for all deployments – treat everything as production, even development
 - OEDA default for all disk groups
- **To ensure better data protection, Oracle requires databases to be configured with ASM High Redundancy on Exadata systems that are older than 5 years since last ship date.**
 - Exadata System Software Certification (Doc ID [2075007.1](#))



ASM Disk Group Configuration

Best practice recommends keeping the number of Disk Groups to a minimum

- DATA Disk Group for data files, SPFILEs, and online redo logs
 - Separating files into DATA and RECO allows the database to survive the loss of either
- RECO Disk Group for backups, archive logs, and online redo logs
- SPARSE for Exadata Sparse Clones
- XTND for XT Storage

Number of Disk Groups

Reasons for More Than the Recommended Minimum Number of Disk Groups

- Separate disk groups for different storage tiers, e.g. HC/EF/XT
- Different Exadata Hardware generations – Disk Capacity in X8M - 14TB, X9M - 18TB, X10M – 22TB
 - Grid disks can be aligned to the same size/slice (eg 3TB) – but performance may vary
- A system has disks of varying capacity and performance. Because different drive types cannot be mixed in the same disk group and separate disk groups are required
 - You can't mix HC and EF in the same Disk group
- To isolate specific databases due to operational considerations such as E-Business Suite data and Siebel data

ASM Voting Files

- ASM requires 5 voting (quorum) files
 - Used for cluster membership management
 - Spread across 5 different storage servers
- Special case for High Redundancy Disk Groups with only 3 Storage Servers
 - Exadata Eighth and Quarter Rack Database Machines
 - Need to provide 5 quorum/voting disks
 - 2 located in the Database Servers in addition to 3 on Storage Servers
 - Automatically configured by OEDA
 - Make sure the iscsi quorum disks are online via `v$asm_disk_stat`



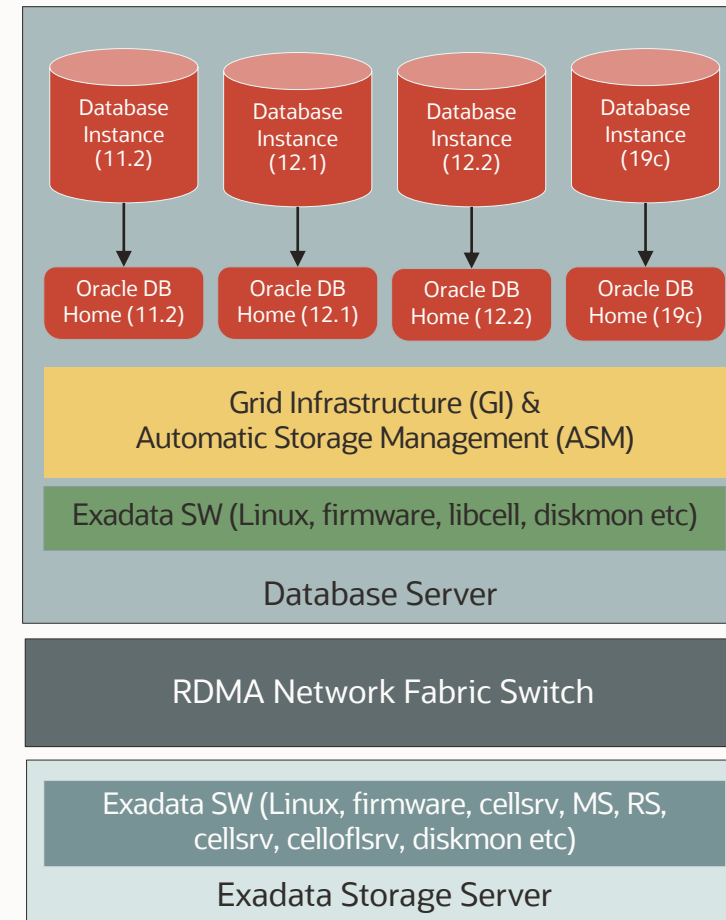
Software Currency

—

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



Stay Current

My Oracle Support Note [888828.1](#)

- All software should be updated regularly. Maintaining software at current or recent releases provides the following benefits:
 - Better software security
 - More stable sustaining releases
 - Continued compatibility with newer related software
 - Better support and faster resolution of issues
 - Ability to receive fixes for newly discovered issues.



Stay Current

Exadata System Software Recommendation

- Become familiar with the date for the Final Maintenance Release shown in this document in order to plan, test, and perform upgrades to a higher Exadata Software Release in a timely manner. Update from lower Exadata versions before the Date shown in the table above.
- The latest Exadata 23.1 release is recommended for all systems to get the latest functionality.
 - Updates to Exadata software from Exadata 21.2/22.1 to Exadata 23.1, perform Oracle Linux upgrade from OL7 to OL8. Verify compatibility of custom-installed software with OL8 before updating.
- The latest Exadata 22.1 release is recommended for all systems to get the latest bug fixes and latest security fixes.
 - Updates to Exadata software from Exadata 18, or lower, to Exadata 22.1, perform Oracle Linux upgrade from OL6 to OL7. Verify compatibility of custom-installed software with OL7 before updating.
- Monitor [Document 1270094.1](#) for Exadata Critical Issues and [Document 556.1](#) for Exadata and Linux Important Recommended Fixes. Plan to update Exadata System Software or apply interim fixes to proactively avoid issues described in those documents.
- Newer Exadata software may require updates to Grid Infrastructure and Database software to meet minimum version requirements. See the Exadata Supplemental README for the target Exadata release for details.
- Ksplice updates, delivered via the uptrack-updates rpm, are available for the kernel on Exadata database servers. Ksplice allows certain updates to be made to the kernel online without outage or reboot. See [Document 2207063.1](#) for details about updating database servers to a later uptrack-updates rpm than what is supplied with your Exadata release.

Grid Infrastructure / Database Software Recommendation

- Become familiar with Grid Infrastructure / Database Patching End Dates shown in [Document 742060.1](#) in order to plan, test, and perform upgrades to a new feature release in a timely manner. Upgrade from lower releases before support end dates to help you stay within the guidelines of Lifetime Support and Error Correction Policies.
- The latest Grid Infrastructure / Database 19c Release Update is recommended for all systems to get the latest bug fixes and latest security fixes.
- Monitor [Document 1270094.1](#) for Exadata Critical Issues and [Document 555.1](#) for Database and Grid Infrastructure 19c Important Recommended Fixes. Plan to update Database and/or Grid Infrastructure or apply interim fixes to proactively avoid issues described in those documents.
- Mission critical systems should update Grid Infrastructure and Database software regularly using Monthly Recommended Patches (MRPs, 19.17 and higher) or Release Update Revisions (RURs, 19.16 and lower). See [Document 2898740.1](#) for MRP details. See [Document 2898381.1](#) and [Document 2285040.1](#) for RUR details.
- Newer Grid Infrastructure / Database software may require updates to Exadata software to meet minimum version requirements.
- See [Document 2700151.1](#) for Transitional Support for Grid Infrastructure 11.2.0.4 on Exadata running Oracle Linux 7. This support on Exadata was added in order to enable Grid Infrastructure upgrade from 11.2.0.4 directly to Long Term Release 19c.

General Software Recommendation

- The recommended version depends on the state of deployment and may not be equal to the latest release specified in the Current Versions table below. The general guideline is the following:
- Run the latest exachk and review the MAA Scorecard section for version recommendations. Exachk evaluates current software versions for consistency, compatibility, and whether or not it is current.
- A software update should be installed on a production system only after it has been validated in a proper test environment.
- Systems in production or late testing stages before production should plan to periodically adopt more current software. It is not required or necessary to install every new update, but the installed software should not lag the latest update by more than 12 months.
- Systems that are in the early stages of testing before production or proofs-of-concept should adopt new updates as they are made available as indicated in the Current Versions table below.

Source: <https://support.oracle.com/epmos/faces/DocumentDisplay?id=888828.1>



Exachk

Exachk holistically evaluates Exadata Database Machine engineered systems.

- Configuration checks for Database Servers, Storage Servers, and Network Fabric Switches:
 - Firmware
 - Operating System (e.g. Oracle Linux)
 - Exadata software
 - Grid Infrastructure and ASM
 - Database
- MAA Scorecard:
 - MAA Configuration Review
 - Exadata Software Planner
 - Exadata Critical Issue alerts
- Automatic Correction (when applicable):
 - Configuration Correction
 - Critical Issue Avoidance
- Prerequisite checks for DB and GI software updates
- Prerequisite checks for DB and GI upgrades
- Prerequisite checks for application continuity readiness

Database Server

Status	Type	Message
CRITICAL	Database Check	Database parameter CLUSTER_INTERCONNECTS is not set to the recommended value
CRITICAL	Database Check	Database parameters log_archive_dest_n with Location attribute are not all set to recommended value
CRITICAL	OS Check	Hardware and firmware profile check is not successful. [Database Server]
CRITICAL	OS Check	The InfiniBand Address Resolution Protocol (ARP) Configuration on Database Servers should be as recommended
FAIL	SQL Check	Some data or temp files are not autoextensible
FAIL	OS Check	Memlock settings do not meet the Oracle best practice recommendations
FAIL	ASM Check	Fast recovery area allocation totals are greater than the total space of the DB_RECOVERY_FILE_DEST
FAIL	OS Check	Active kernel version should match expected version for installed Exadata Image
FAIL	OS Check	One or more database server has non-test stateless alerts with null "examinedby" fields
FAIL	OS Check	One or more database servers have stateful alerts that have not been cleared
FAIL	Database Check	Hidden database Initialization Parameter usage is not correct
WARNING	Database Check	Local listener init parameter is not set to local node VIP
WARNING	Database Check	Database parameter DB_BLOCK_CHECKING on PRIMARY is NOT set to the recommended value.
INFO	OS Check	Exadata Critical Issues (Doc ID 1270094.1):- DB1-DB4,DB6,DB9-DB41, EX1-EX54,EX56 and IB1-IB3
INFO	Database Check	One or more non-default AWR baselines should be created



Prevent Downtime By Relying on “Community Effect”

Health Checks using EXAchk Utility

- EXAchk provides configuration specific, up-to-date health check across the entire stack
 - Covers Exadata, Database, Grid Infrastructure, ASM critical issues
 - Provides MAA scorecard with MAA configuration gaps and guidance to mitigate
 - Automated periodic scheduled runs with email notifications
- **Continuous evolution of configuration checks**
 - EXAchk helps with saving a lot of time and money due to proactive health verification which dramatically reduces downtime
 - **Currently has over 1000 checks per target**
- Development recommends that the latest EXAchk be executed with the following frequency:
 - Monthly
 - Week before any planned maintenance activity
 - Day before any planned maintenance activity
 - Immediately after completion of planned maintenance activity or an outage or incident

Grid Infrastructure & Automatic Storage Management

- Grid Infrastructure (GI) & Automatic Storage Management (ASM) bind the Database Servers together
 - Underpinning of Real Application Clusters and Exadata MAA
 - Use the Latest Long Term Support Release – i.e. Version 19
 - Supports all Database versions 11.2.0.4 and higher
 - Starting with version 18 of Oracle Database and GI, GI RU/RUR versions can be lower than the DB RU/RUR
 - Eg GI 19.9.0 with DB 19.11.0 is supported
- Be as current as you can – low GI/ASM is an inhibitor to database currency
- OEDACLI supports major GI/ASM Upgrades eg 12.2 -> 19
- Database 21c requires GI/ASM 21

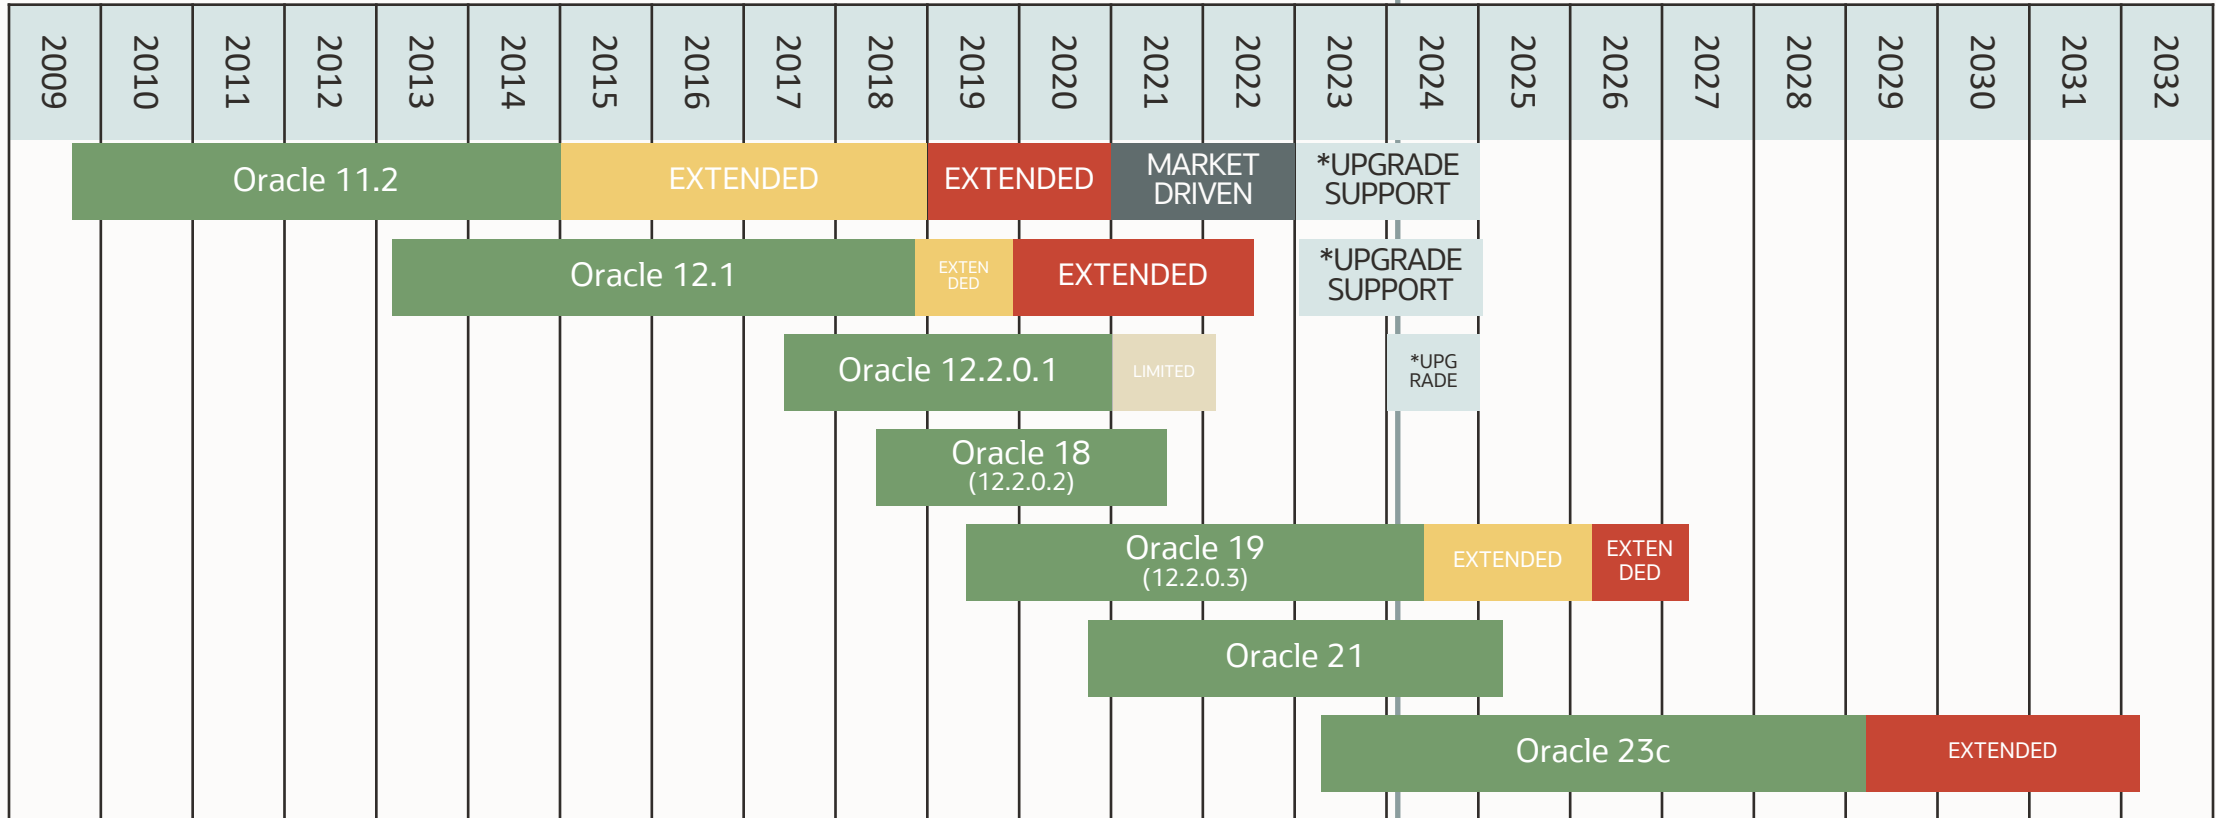
Clusterware	ASM	DB	Certified
19c	19c	19c	Y
19c	19c	18c	Y
19c	19c	12.2	Y
19c	19c	12.1	Y
19c	19c	11.2*	Y

* Pre-12.1 database instances require an ASM instance resident on the same node as the database instance. Pre-12.1 database instances cannot leverage the implicit HA of Flex ASM.



Lifetime Support Policy (Database)

Current as of January 2024



Exadata System Software versions

- Different Exadata System Software versions on different servers is supported
 - Examples
 - Some storage servers may run 22.1.19, while others are running 23.1.9
 - All storage servers running 23.1.9 while database servers running 22.1.19
 - Highly recommended that such configurations are for the purpose and duration of rolling upgrades

Patch Testing & Apply Guidelines

Exadata Patching Overview and Patch Testing Guidelines (Doc ID [1262380.1](#))

1. Review patch README for known issues, installation/deinstallation instructions and special notes
2. Run Exachk before and after patch installation – do not proceed until any issues are resolved
 - Always use the latest version
 - TEST and PRODUCTION
3. Apply and validate patch installation in TEST environment
 - Automate to avoid human error
 - Use technologies & practices to reduce downtime & risk
 - Logical transient standby, Standby First Patching, Rolling patch apply etc
 - Define, document & test fallback

Patch Testing & Apply Guidelines

Exadata Patching Overview and Patch Testing Guidelines (Doc ID [1262380.1](#))

4. Validate patch functionality in TEST environment

- Validate availability, performance & operational consistency (or improvement)
 - Use Real Application Testing, AWR, SQL Performance Management, SQL Performance Analyzer etc

5. Apply and validate patch in PRODUCTION environment

- Validate availability, performance & operational consistency (or improvement)

What should you use as a TEST environment

Recommend having an environment that is the same as production

- Equivalent to PRODUCTION environment
 - Primary Database
 - Standby Database
 - Middle/App Tier
- Production data set (prod database clone) with production schema stats
 - Masked appropriately for security
- Workload generation framework to mimic production
 - Real Application Testing
 - Application Testing Tools
- AWR Metrics from PRODUCTION for comparison with patched TEST environment



Tools to remain current

- Patchmgr is the provided and primary tool for updating Oracle Exadata components
 - Updates DB Servers – Bare Metal & Virtualized
 - RDMA Network Fabric Switches – RoCE & InfiniBand
 - Oracle supplied Management Switches
 - Storage Servers
 - Oracle Exadata Storage Server Cloud Scale Software Update
- Alternatives
 - Enterprise Manager – Exadata Patching
 - Fleet Patching & Provisioning


Software Maintenance

—

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



Oracle Linux Kernel on Exadata

- The Unbreakable Enterprise Kernel (UEK) is a Linux kernel built by Oracle and supported through Oracle Linux support.
- Its focus is performance, stability, and minimal backports by tracking the mainline source code as closely as is practical.
- UEK is well-tested and used to run Oracle's Engineered Systems, Oracle Cloud Infrastructure, and large enterprise deployments for Oracle customers.
- **Exadata uses and relies on Oracle's Unbreakable Linux Kernel (UEK)**
- Do not update or customize Oracle Linux Kernel or related software unless directed by Oracle Support
 - Other software may be installed, updated, or customized
 - Exadata updates may not carry newer version dependencies of customized components. You may be required to remove and subsequently reapply site-specific changes for successful Exadata update.





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

Customization	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. OL7 RPM on OL8)

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 login should not include interactive menus or prompts

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

What is OEDACLI?

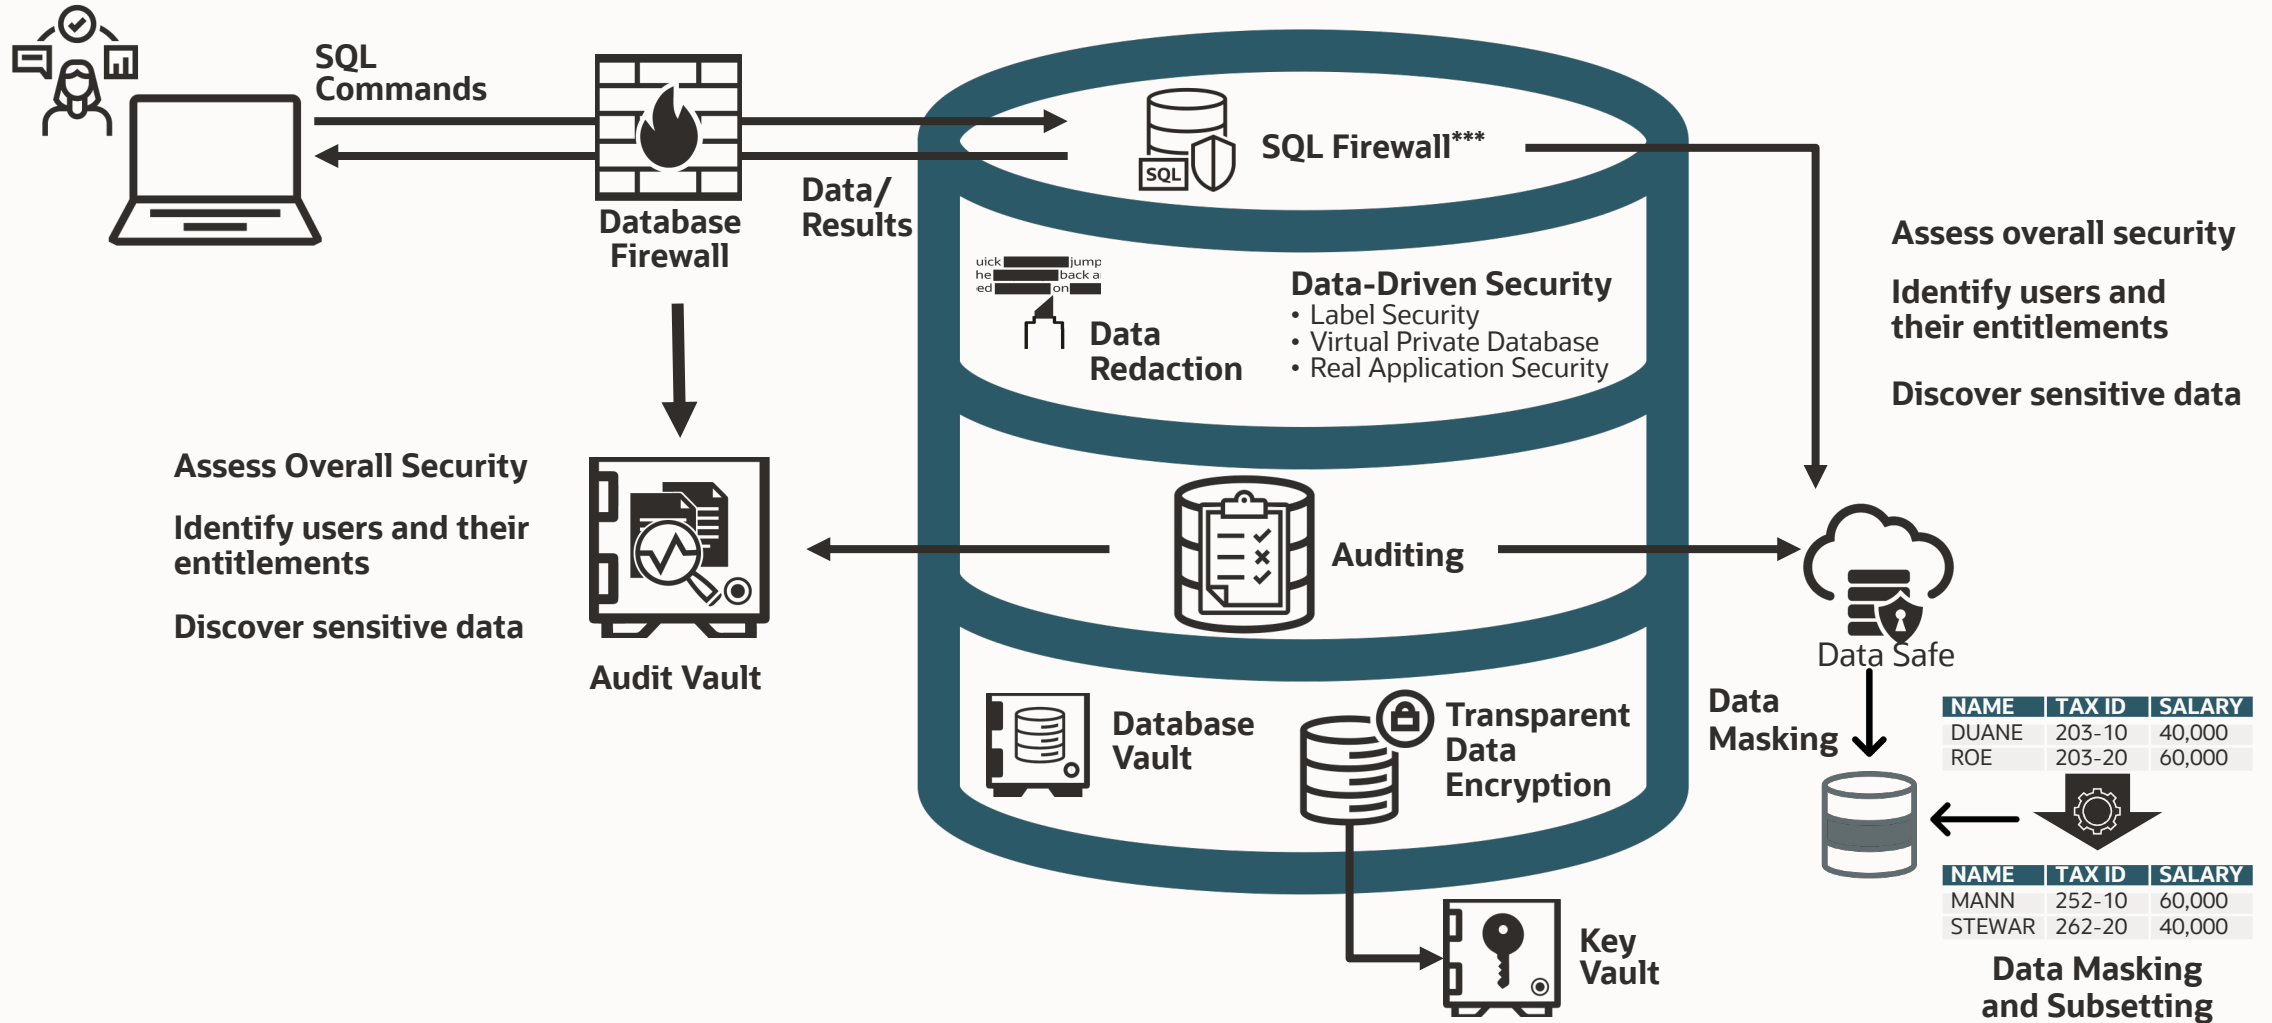
- One of the best kept secrets in Exadata
- Command line utility of Oracle Exadata Deployment Assistant
- Used for Initial Deployment and On-going lifecycle operations
 - Add database home
 - Upgrade VM Cluster GI
 - Alter ASM disk groups (including grid disks)
 - Create|Remove Databases
 - Add|remove VM Guests from existing cluster
 - And many more operations

oedacli is recommended for lifecycle operations

Security




Oracle Database Maximum Security Architecture



*** Only available in Oracle Database 23c



Exadata Development Focus



Security Optimized

- ✓ Optimize the installation by removing access to features that do not accomplish Exadata's primary purpose
- ✓ Restrict access to features when not required



Security Focused

- ✓ Disable all unnecessary and insecure services unless required for system functionality
- ✓ Modify resulting services to run with lowest privileges



Security Hardened

- ✓ Conduct regular, extensive security scans across the entire stack using industry leading security scanners
- ✓ Integrate security fixes into each release and provide emergency fixes to address zero-day vulnerabilities



Exadata Maximum Security Architecture

Nano Linux Kernel Installation

Security: Optimized

Exadata uses a custom, nano (micro) kernel with removed dependencies that reduce size and features that are not needed in an enterprise data center.

- Fewer device drivers
- Smaller footprint
- Improved upgrade time

Standard kernel for OL7:
kernel-3.10.0-1127.13.1.el7.x86_64

- DomU kernel size **167MB**



Small Installation Footprint

Security: Optimized

Exadata **reduces the attack surface** by only including the software components required specifically to run the Oracle database (e.g. minimum Linux distribution)

Channel Detail

Name	Oracle Linux 7 Latest (x86_64)
Description	All packages released for Oracle Linux 7 (x86_64) including the latest errata packages. (x86_64)
Label	ol7_x86_64_latest
Last Modified	2021-05-06
Architecture	x86_64
Packages	5439

Channel Detail

Name	Exadata release 21.2.0.0.0 db server installation packages (x86_64)
Description	All packages released on the Exadata release 21.2.0.0.0 (x86_64) OL7 installation media.
Label	exadata_dbserver_21.2.0.0.0_x86_64_base
Last Modified	2021-05-06



Access Control For RESTful Service

Security: Focused

Oracle Exadata System Software release 19.1.0 introduces a new capability for users to configure access control lists on the HTTPs access to the RESTful service

- Specify a list of IP addresses or subnet masks to control access to the RESTful service via HTTPs
- If not used, RESTful service can be disabled altogether
- Applies to both Oracle Exadata Database and Storage Server

```
# lsof -i -P -n | grep LISTEN | grep java
java      <pid> dbmsvc  55u  IPv4  40193      0t0  TCP *:7879 (LISTEN)
```

```
# dbmcli -e alter dbserver httpsAccess=none
This command requires restarting MS. Continue? (y/n): y
Stopping MS services...
The SHUTDOWN of MS services was successful.
Updating HTTPs access control list.
Starting MS services...
The STARTUP of MS services was successful.
DBServer successfully altered
```

```
# lsof -i -P -n | grep LISTEN | grep java
```



Pre-scanned full stack

Security-Hardened

Every Exadata release includes security and emergency fixes to address zero-day vulnerabilities discovered by our internal scanning tools.

- Static/Dynamic code analyzing
- Malware scans
- Third-party software checks
- Vulnerability scans
- How to research Common Vulnerabilities and Exposures (CVE) for Exadata packages (Doc ID 22568871)
- System hardening reviews (STIG)
- Exadata OL7 System Hardening for STIG Security Compliance (Doc ID 2614471.1)

Customers take advantage of these fixes out of the box by just upgrading to the latest release

- Number of issues reported should be much less compared to a custom configuration



Monitoring

—

Exadata Manageability: Best Practices

Enterprise Manager

- MAA Best Practice is to use EM for Exadata and ZDLRA Manageability
- We certify Exadata and ZDLRA Manageability with EM
 - For each Exadata hardware generation and software feature release
 - For Customers, Platinum, and Oracle Cloud Operations
 - To ensure EM raises incidents for hardware and software faults and provides necessary monitoring and management capabilities
 - We certify Exadata with ASR (separate from EM) to ensure hardware faults automatically generate SRs
- Ensure your EM environment is as available as the Exadata and ZDLRA environments it manages
 - EM HA levels provide High Availability and Disaster Recovery
- Keep your EM environment patched and upgraded as necessary to support latest releases
 - [Exadata System Software and Hardware Versions Supported by Oracle Enterprise Manager Plug-ins \(Doc ID 1626579.1\)](#)
 - [Support Note 2542836.1 - ZDLRA Software Versions Supported by the Oracle Enterprise Manager ZDLRA Plug-in](#)



Exadata and Exadata Cloud Manageability with EM 13.5

- Core existing EM Exadata manageability updated, works with X10M and earlier models
 - Guided discovery, ILOM and HALRT incidents, metrics, performance monitoring, navigation tree
 - Exadata 23.1 software release and below
 - XRMEM metrics, Event Compression to reduce noise, Command line Exadata discovery via EMCLI
 - Improved Exadata PDU monitoring stability with PDU firmware version v2.12
 - MAA KPIs integrated as first class metrics, updated technical brief
 - <https://www.oracle.com/docs/tech/database/exadata-database-machine-kpis.pdf>
 - EM 13.5 Exadata Getting Started Guide
 - <https://docs.oracle.com/en/enterprise-manager/cloud-control/enterprise-manager-cloud-control/13.5/emxig/index.html>
- EM Exadata Cloud manageability
 - OCI API integration, support for Exadata Infrastructure, Exadata VM Cluster, Autonomous Exadata VM Cluster
 - EM 13.5 Exadata Cloud document
 - <https://docs.oracle.com/en/enterprise-manager/cloud-control/enterprise-manager-cloud-control/13.5/emxcs/index.html>
- See the following MOS Note for specific plug-in version and patching requirements:
 - [Exadata System Software and Hardware Versions Supported by Oracle Enterprise Manager Plug-ins \(Doc ID 1626579.1\)](#)



Comprehensive Support from One Integrated Team

Premier & Platinum Support

Complete. Proactive. Premier.

- Specialized Engineered Systems support team
- 24/7 support
- 2-hour onsite response to hardware issues¹
- Updates and upgrades for database, server, storage, and OS software via quarterly patch bundle deployment
- Proactive support portal
- "Phone home" automated service requests

Integrated. No Additional Cost. Platinum.

Only available for Oracle Engineered Systems

- Oracle engineers perform quarterly patching and 24/7 fault monitoring
- Faster response and restore times:
 - ✓ 5-minute fault notification
 - ✓ 15-minute restoration or escalation to development
 - ✓ 30-minute joint debugging with development

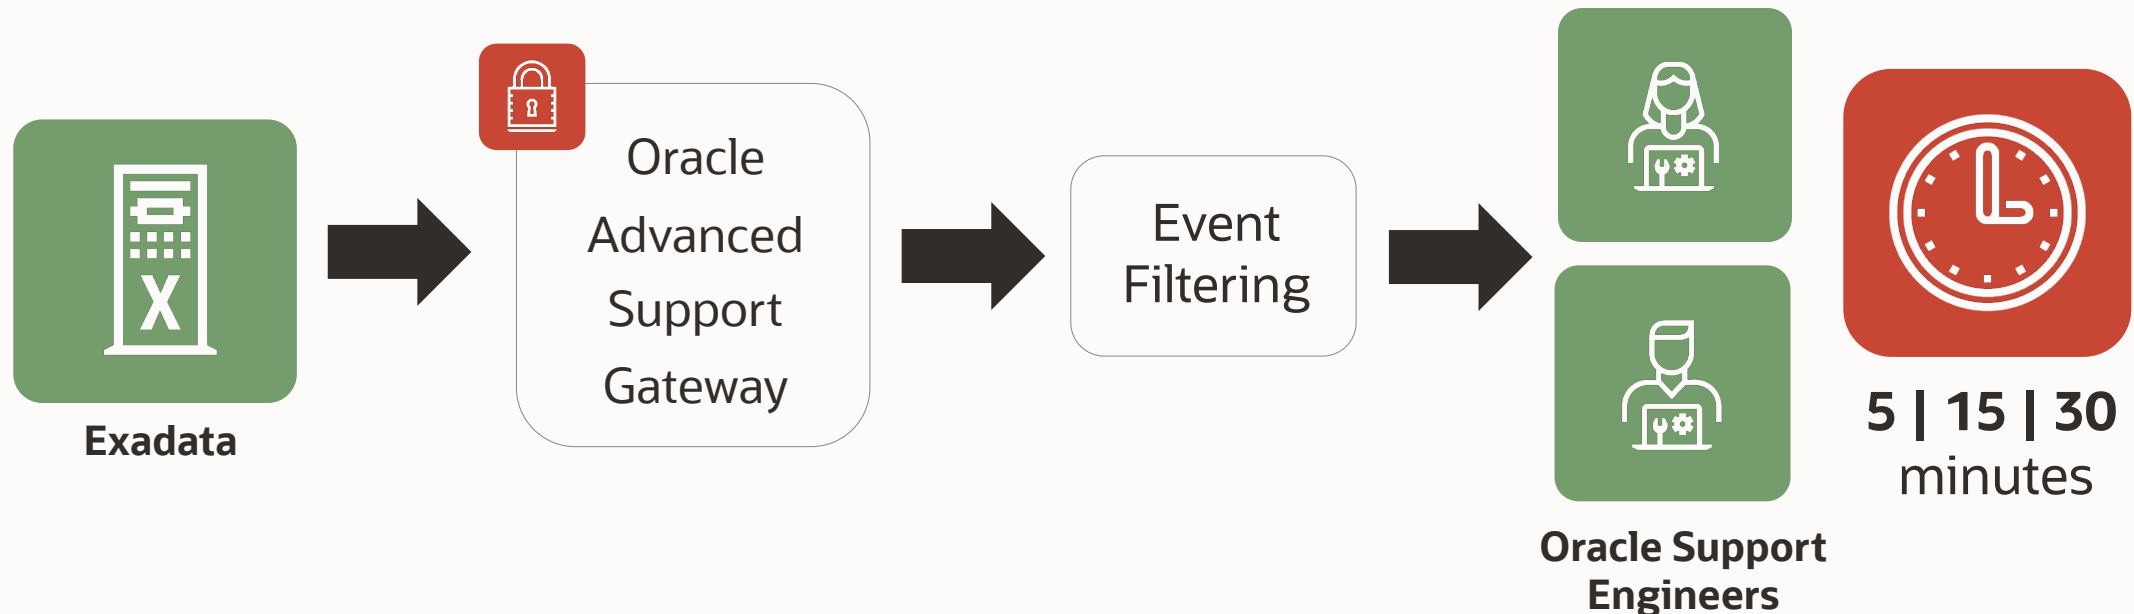


¹ Covered system must be within an Oracle two-hour service area to receive two-hour response as a standard service.

Automatic Service Request

24/7 remote fault monitoring

- Monitoring for **faults in the hardware, database, operating system and networking components** of covered systems
- Focused on **identifying issues** with the ability of core system components to function properly in order to **maintain system availability**



ExaWatcher

Automatically & continually watches and collects system diagnostics information for

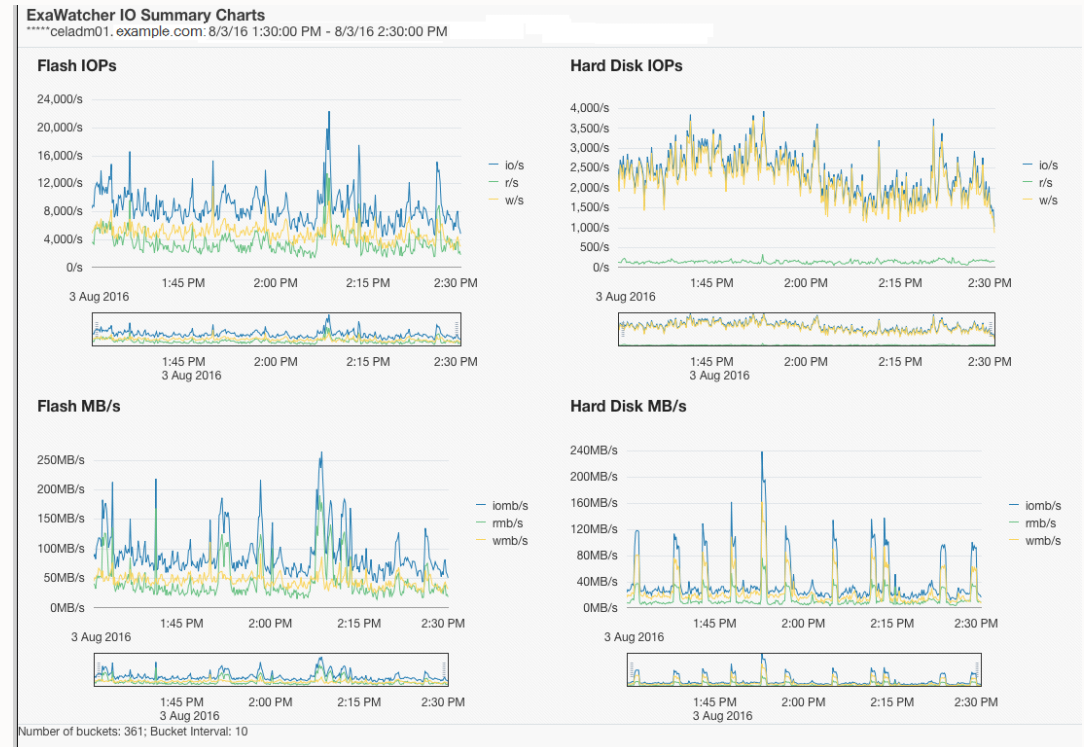
- Database servers
 - Bare Metal
 - KVM Hosts & OVM dom0
 - KVM Guests & OVM domU
- Exadata Storage Servers

Includes:

- top, process, memory, CellSrv Status, network, disk, flash and other system related data

Automatically started at boot and space consumption automatically managed

Charts of subset of metrics automatically generated



Autonomous Health Framework (AHF)

Superset tool including Oracle EXAchk and Trace file Analyzer (TFA)

EXAchk

- Lightweight and non-intrusive compliance check framework for the Oracle stack of software and hardware components
- Recommended to run regularly, before and after system changes
- ALWAYS USE THE LATEST VERSION
 - Ensures latest critical issues are included in checks

TFA

- Monitors your logs for significant problems, automatically trims logs and collects relevant diagnostics across cluster nodes and consolidates everything in one place
- Collections are analyzed for known problems
- Includes DBA tools for problem triage

Streamlines diagnostics collection when issues arise



Automated Performance Monitoring

Uses Machine Intelligence to Detect Performance Issues and Determine Root Cause

- Covers a wide-range of sub-systems, including CPU, memory, file system, IO, and network
- Combines artificial intelligence, years of real-world performance triaging experience, and best practices

Process is running slowly

- CPU on server is maxed out?
- Process is spinning / stuck?

Automated Performance Monitoring Checks
Server: High CPU Usage
Process: Spinning
Process: Stuck in state 'D'
kswapd: High CPU Usage

Memory Usage

- Incorrect config of HugePages
- Application/Agent memory leak

Automated Performance Monitoring Checks
Process: High Memory Usage
System: High Memory
Swapping
HugePages

Network Usage

- Inconsistent client response times
- Flaky client-side network

Automated Performance Monitoring Checks
Port up/down
TCP retransmits



Exadata Real-Time Insight

New in Exadata System Software 22.1

- Automatically stream up-to-the-second metric observations from all servers in your Exadata fleet
- Feed customizable monitoring dashboards for real-time analysis and problem-solving

• Comprehensive

- 200+ Exadata Software & Hardware Metrics
- Fine-grained metrics can be collected as often as every 1 second

• Integrated

- Integrated with popular time-series and observability platforms
- Stream fine-grained metrics to user-defined endpoints in real time

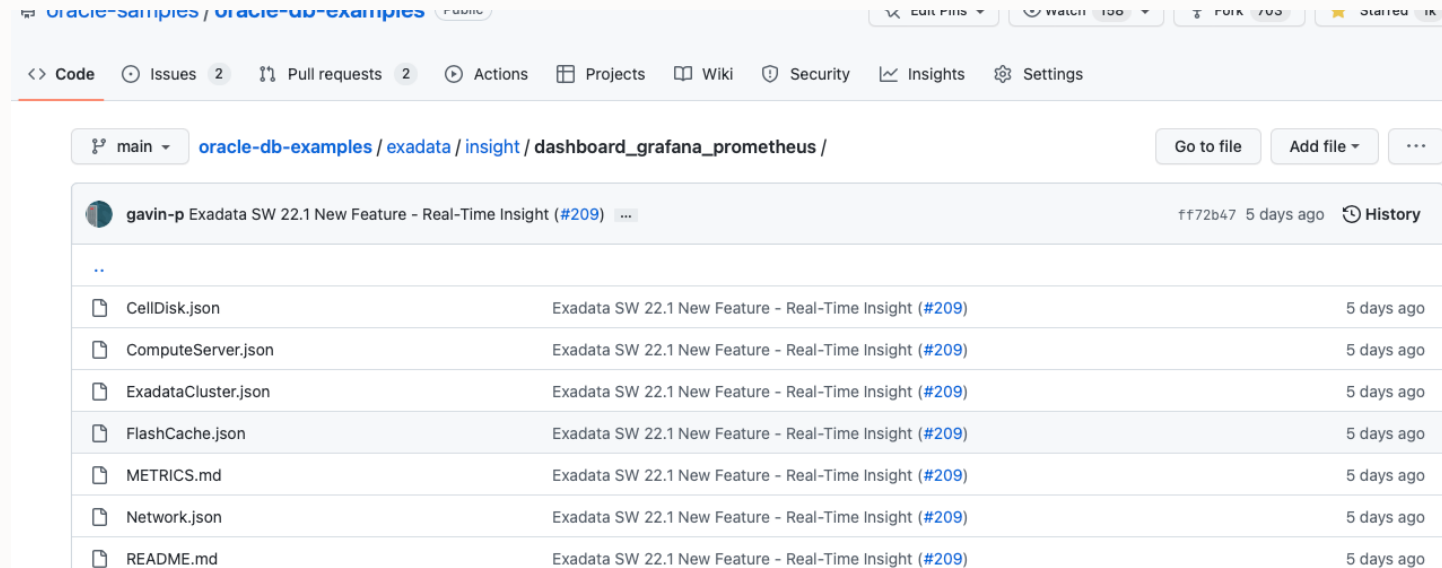
• Insightful

- Enables proactive issue detection and real-time decision making



Exadata Real-Time Insight – Sample Dashboards Code

- Oracle Samples repository on GitHub.com contains example Real-Time Insight dashboards.
 - The following dashboard code is included (Grafana/Prometheus):
 - Exadata Cluster
 - Compute
 - Storage Server
 - Cell Disk
 - Flash Cache
 - Smart Scan
 - Network



- <https://github.com/oracle-samples/oracle-db-examples/tree/main/exadata/insight>



Exadata Real-Time Insight – Sample Dashboards

- **Exadata Cluster:** Provides a cluster-wide view that shows metrics for compute nodes and storage servers.
- **Compute:** Provides a compute-node view that shows CPU and network utilization for the compute nodes.
- **Storage Server:** Provides a storage-server-centric view that focuses on storage server CPU and I/O metrics, as well as Exadata metrics for Smart Flash Cache, Smart Flash Log, and Smart I/O.
 - Cell Disk: Shows cell disk I/O metrics on the storage server.
 - Flash Cache: Shows flash cache metrics on the storage server.
 - Smart Scan: Shows smart scan metrics on the storage server.



Resources



Useful Docs & links

Check these documents regularly

Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)

<https://support.oracle.com/epmos/faces/DocumentDisplay?id=888828.1>

Exadata Critical Issues (Doc ID 1270094.1)

<https://support.oracle.com/epmos/faces/DocumentDisplay?id=1270094.1>

Master Note for Database Proactive Patch Program (Doc ID 888.1)

<https://support.oracle.com/epmos/faces/DocumentDisplay?id=888.1>

Oracle Database 19c Important Recommended One-off Patches (Doc ID 555.1)

<https://support.oracle.com/epmos/faces/DocumentDisplay?id=555.1>

Exadata Maximum Availabilty Architecture - Whitepaper

<https://www.oracle.com/a/tech/docs/exadata-maa-wp.pdf>

Exadata My Oracle Support Notes – Core Infrastructure

Note ID	My Oracle Support Note Title
888828.1	Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)
1270094.1	Exadata Critical Issues (Doc ID 1270094.1)
556.1	Oracle Exadata: Exadata and Linux Important Recommended Fixes (Doc ID 556.1)
2724126.1	Exadata X9M/X8M (RoCE-based systems) Software Requirements and Recommendations (Doc ID 2724126.1)
2075007.1	Exadata System Software Certification (Doc ID 2075007.1)
1570460.1	Exadata Software and Hardware Support Lifecycle (Doc ID 1570460.1)
1070954.1	Oracle Exadata Database Machine EXAchk (Doc ID 1070954.1)
1306791.2	Information Center: Oracle Exadata Database Machine (Doc ID 1306791.2)

Review the latest My Oracle Support Notes related to Exadata [here](#).



Exadata My Oracle Support Notes – Updates/Patching/Upgrades

Note ID	My Oracle Support Note Title
888828.1	Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)
1262380.1	Exadata Patching Overview and Patch Testing Guidelines (Doc ID 1262380.1)
2207063.1	HOWTO: Install ksplice kernel updates for Exadata Database Nodes (Doc ID 2207063.1)

Review the latest My Oracle Support Notes related to Exadata [here](#).



Exadata My Oracle Support Notes – Security

Note ID	My Oracle Support Note Title
2751741.1	Oracle Exadata Database Machine Security FAQ (Doc ID 2751741.1)
2256887.1	How to research Common Vulnerabilities and Exposures (CVE) for Exadata packages (Doc ID 2256887.1)
1405320.1	Responses to common Exadata security scan findings (Doc ID 1405320.1)

Review the latest My Oracle Support Notes related to Exadata [here](#).



Exadata My Oracle Support Notes – Database and Grid Infrastructure

Note ID	My Oracle Support Note Title
742060.1	Release Schedule of Current Database Releases (Doc ID 742060.1)
2700151.1	Grid Infrastructure 11.2.0.4 on Exadata / Oracle Linux 7 Transitional Support (Doc ID 2700151.1)
2285040.1	Release Update Introduction and FAQ (Doc ID 2285040.1)
2542082.1	19c Grid Infrastructure and Database Upgrade steps for Exadata Database Machine running on Oracle Linux (Doc ID 2542082.1)
1919.2	19c Database Self-Guided Upgrade with Best Practices (Doc ID 1919.2)

Review the latest My Oracle Support Notes related to Exadata [here](#).

ORACLE

Our mission is to help people see
data in new ways, discover insights,
unlock endless possibilities.

