Level 4 Exadata X11M & Oracle Database Appliance

Senior DBA Interview Questions and Answers

Architecture | Performance | Storage | Patching | Operations | Troubleshooting | Production Scenarios

Audience	Senior / Lead Oracle DBA, Exadata DBA, Engineered Systems Consultant
Coverage	59 detailed questions with operational and scenario-based answers
Product focus	Oracle Exadata Database Machine X11M and Oracle Database Appliance
Current reference point	X11M platform documentation and ODA 19.30/19.31-era documentation available in 2026

Interview guidance: Strong Level 4 answers should explain the architecture, show a repeatable diagnostic method, identify risk, name evidence to collect, and distinguish supported engineered-system operations from unsafe generic Linux or storage administration.

Contents

1. Exadata X11M Architecture and Hardware

Q1. What is Oracle Exadata X11M, and how is it different from a conventional Oracle RAC platform?

Answer: Exadata X11M is an engineered database platform that integrates Oracle database servers, intelligent Exadata Storage Servers, RDMA Network Fabric, optimized system software, and automated lifecycle management. A conventional RAC platform normally treats storage as generic block storage. Exadata pushes database-aware operations into the storage cells using features such as Smart Scan, storage indexes, Smart Flash Cache, Smart Flash Log, columnar flash cache, persistent memory-class caching through XRMEM, and I/O resource management. The result is reduced database-node CPU consumption, lower latency, higher throughput, and consistent consolidation behavior.

Q2. Describe the key X11M database-server hardware characteristics.

Answer: An X11M database server uses two 96-core AMD EPYC processors, DDR5 memory, local NVMe system storage, and dual-port 100 Gb/s RDMA Network Fabric connectivity operating active-active for up to 200 Gb/s aggregate fabric throughput. The platform uses PCIe 5.0 and supports multiple memory configurations. A Level 4 DBA should understand that usable database capacity depends on deployment model, virtualization, reserved resources, database licensing, and workload design rather than only the physical specifications.

Q3. What is new or important about X11M storage-server design?

Answer: X11M storage servers use newer AMD processors, DDR5 memory, high-performance flash, high-capacity disks on HC models, 200 Gb/s aggregate RDMA fabric connectivity, and XRMEM. XRMEM provides an RDMA-accessible memory tier in the storage server that can service eligible reads at very low latency without the database process performing traditional network and storage-stack work. The design preserves Exadata scale-out principles: capacity and bandwidth increase by adding storage servers.

Q4. What is XRMEM and how would you explain its value in an interview?

Answer: XRMEM is Exadata RDMA-accessible memory in the storage tier. Frequently accessed data can be served from storage-server memory through RDMA, reducing I/O latency and CPU overhead. It complements, rather than replaces, the database buffer cache and flash cache. The DBA should verify eligibility, cache effectiveness, database and system-software compatibility, and workload behavior instead of assuming every read is served from XRMEM.

Q5. Explain the RDMA Network Fabric in Exadata X11M.

Answer: The RDMA Network Fabric is the private high-speed interconnect used for database-to-storage traffic, RAC messaging, and other engineered-system communication. X11M uses RoCE-based 100 Gb/s ports in an active-active design. RDMA permits direct memory movement with low CPU overhead and low latency. Operationally, the DBA must monitor interface health, switch ports, bonding, MTU consistency, packet errors, congestion, and fabric topology; changes must follow Exadata-supported procedures.

Q6. What is the difference between a High Capacity and an Extreme Flash storage server?

Answer: High Capacity cells combine hard disks with flash and are designed for balanced capacity, performance, and cost. Extreme Flash cells use all-flash storage and target latency-sensitive or high-I/O workloads. Capacity planning must consider usable ASM capacity, redundancy, reserved free space, rebalance headroom, flash-cache effectiveness, write workload, recovery needs, and growth. The correct choice follows workload evidence, not a general belief that all-flash is always required.

Q7. How does Exadata scale?

Answer: Exadata scales database compute and storage independently within supported configurations. Adding database servers increases CPU, memory, and RAC processing capacity. Adding storage servers increases usable storage, flash, cell CPU, scan bandwidth, and I/O concurrency. Scale-out requires reviewing ASM rebalance impact, network topology, cluster configuration, database services, IORM plans, licensing, and operational windows.

Q8. What should be checked before expanding an Exadata rack?

Answer: Confirm model and software compatibility, rack power and cooling, physical slots, switch-port capacity, supported cable topology, IP addresses, firmware and image levels, ASM free space, disk-group redundancy, rebalance power, expected data redistribution time, database workload windows, backup status, and rollback procedures. Run current health checks and preserve baseline metrics before the change.

2. Exadata Storage and Smart Features

Q9. Explain Smart Scan and the conditions required for it.

Answer: Smart Scan offloads eligible table and index scans to storage cells. Cells perform column projection, predicate filtering, and in some cases other processing, returning a reduced result set to database nodes. It generally requires direct-path reads and eligible objects and operations. It may not occur for small buffer-cache reads, unsupported datatypes or functions, certain chained-row conditions, or execution paths that do not use direct path. Verify using execution-plan storage predicates and cell-related session statistics.

Q10. What are storage indexes?

Answer: Storage indexes are in-memory metadata maintained automatically by cells, containing value-range information for storage regions. They allow cells to skip regions that cannot satisfy predicates. They are not database indexes, are not manually created, are not persisted like schema objects, and consume no application-maintained storage. Effectiveness depends on data clustering, predicate selectivity, and workload access patterns.

Q11. Explain Smart Flash Cache and how it differs from database buffer cache.

Answer: The database buffer cache resides in database-node memory and caches Oracle blocks managed by the instance. Smart Flash Cache resides in storage-server flash and caches frequently or strategically accessed data below ASM. It is shared by databases using the cells and is controlled by Exadata caching algorithms and object-level hints such as KEEP or NONE where appropriate. A read can miss the buffer cache but still be served from flash without disk access.

Q12. What is Smart Flash Log?

Answer: Smart Flash Log accelerates redo-write completion by writing redo in parallel to flash and disk, acknowledging when one protected path completes according to the implementation. It reduces the impact of occasional disk latency spikes on log file sync. It is not a replacement for redo logs and does not change database recovery semantics.

Q13. What is Exadata Smart Write Back Flash Cache?

Answer: Write-back mode allows changed blocks to be acknowledged after being safely mirrored in flash, with disk writes completed later. It can significantly improve write-heavy workloads. The DBA must confirm normal cache mode, flash health, cell redundancy, battery or power protection state where relevant, and absence of degraded conditions. A failed or disabled write-back cache can materially change latency.

Q14. What is Hybrid Columnar Compression and where should it be used?

Answer: HCC stores data in compression units optimized either for query performance or archival compression. It is best suited to bulk-loaded, relatively static data such as warehouse partitions, historical data, and archives. Frequent singleton DML can cause row migration to less-compressed formats and reduce benefits. Choose QUERY LOW/HIGH or ARCHIVE LOW/HIGH based on scan performance, space, load time, and update pattern.

Q15. Explain IORM and why it is important in consolidated Exadata environments.

Answer: I/O Resource Management controls how cell resources are allocated among databases, consumer groups, categories, and workloads. It prevents a noisy workload from monopolizing disks, flash, cell CPU, or bandwidth. A mature design aligns IORM with Database Resource Manager and business service priorities. The DBA validates plans under load, monitors objective compliance, and avoids setting rigid limits without understanding peak and failover behavior.

Q16. What is the purpose of cellcli, dcli, and exadcli?

Answer: cellcli manages an individual storage cell. dcli executes commands across a defined set of servers using SSH and is often used for operating-system level checks. exadcli provides authenticated, structured management across Exadata components and is preferred for many modern operations. A Level 4 DBA uses these tools carefully, limits command scope, captures output, and avoids unsupported manual modifications.

Q17. How do you check cell health and key storage objects?

Answer: Use cellcli or exadcli to review cells, physical disks, LUNs, cell disks, grid disks, flash devices, flash cache, and alerts. Typical checks include LIST CELL DETAIL, LIST PHYSICALDISK, LIST CELLDISK, LIST GRIDDISK, LIST FLASHCACHE DETAIL, and LIST ALERTHISTORY. Correlate cell status with ASM disk state, database wait events, I/O latency, and hardware fault telemetry.

3. Exadata Performance Diagnosis

Q18. A query is not using Smart Scan. How do you troubleshoot it?

Answer: First confirm the actual execution plan and whether the operation is a full or fast full scan. Check storage predicates, direct-path read statistics, cell offload eligibility, bytes eligible for predicate offload, bytes returned by cells, and storage-index savings. Review object size, buffer-cache behavior, parallelism, function usage, datatype limitations, encryption or compression implications, and optimizer transformations. Do not force full scans merely to obtain offload; compare total response time and resource cost.

Q19. Which Exadata statistics are useful for proving offload benefit?

Answer: Useful session or system statistics include physical read total bytes, physical read bytes optimized, cell physical IO interconnect bytes, cell physical IO bytes eligible for predicate offload, cell physical IO bytes saved by storage index, and cell physical IO bytes saved by columnar cache. Interpret deltas over a controlled interval and correlate them with SQL_ID, plan, AWR/ASH, and cell metrics.

Q20. How do you investigate high cell single block physical read latency?

Answer: Identify affected SQL and objects, compare latency across cells and time, and distinguish flash hits, XRMEM hits, and disk reads. Check flash-cache mode and population, cell alerts, physical disk status, flash wear or errors, IORM throttling, RDMA fabric health, ASM imbalance, and database queueing. Also verify whether the real issue is excessive random I/O caused by a poor access path rather than storage hardware.

Q21. How do you diagnose high cell smart table scan time?

Answer: Check scan volume, offload ratio, returned bytes, cell CPU, IORM allocation, flash or disk throughput, storage-index effectiveness, predicate quality, parallel degree, skew, and concurrent workloads. Review whether the query reads much more data than expected because of stale statistics, missing partition pruning, datatype conversion, or a plan regression. A fast storage platform cannot compensate for scanning unnecessary terabytes.

Q22. What causes high interconnect traffic on Exadata?

Answer: Potential causes include poor offload, returning too many rows or columns, inefficient parallel distribution, RAC block transfers, backup or rebalance activity, large result sets, and workload imbalance. Separate database-to-cell traffic from RAC global-cache traffic. Review cell interconnect bytes, gc waits, PX distribution, network interface statistics, and SQL-level offload metrics.

Q23. How would you tune a mixed OLTP and data warehouse workload on Exadata?

Answer: Separate services and consumer groups, define Database Resource Manager and IORM priorities, right-size parallelism, protect redo and latency-sensitive OLTP, schedule or govern large scans, use partitioning and compression appropriately, and validate flash-cache behavior. Capacity tests must include node or cell failure because the design must meet service levels in degraded mode, not only at full capacity.

Q24. How does ASM allocation affect Exadata performance?

Answer: ASM distributes extents across grid disks and cells. Balanced allocation is essential for parallel I/O and predictable latency. Check disk-group imbalance, grid-disk sizing, rebalance status, failgroup design, and free-space distribution. Large differences can occur after capacity expansion, disk replacement, or insufficient rebalance power. Avoid manual layouts that defeat Exadata recommendations.

Q25. What is your Exadata performance baseline?

Answer: A baseline should include database DB time, wait classes, top SQL, CPU, commit latency, RAC waits, cell single-block and scan latency, flash and XRMEM effectiveness, offload ratios, IORM utilization, RDMA errors, ASM capacity and balance, backup throughput, and batch completion times. Capture normal, peak, month-end, and failure-mode behavior so incident analysis has a meaningful comparison.

4. Exadata Operations, Patching, and Troubleshooting

Q26. What is the difference between database RU patching and Exadata system software patching?

Answer: Database RUs patch Oracle Database and Grid Infrastructure homes and are managed with Oracle patching tools and supported automation. Exadata system software patching updates storage cells, database-server operating-system components, firmware, switches, and related engineered-system layers using Exadata-supported tooling and sequencing. They have different prerequisites, rollback methods, and compatibility matrices. Never treat the appliance as a collection of unrelated Linux servers.

Q27. Explain a safe Exadata patching approach.

Answer: Validate target versions and compatibility, read release notes and known issues, confirm backups and recovery, run exachk, clear critical alerts, preserve configuration, verify space and network prerequisites, and establish a rollback plan. Use supported patchmgr or current lifecycle tooling, patch in the recommended order, monitor rolling operations, then run post-checks, exachk, database validation, and performance comparison.

Q28. What is exachk and when do you run it?

Answer: exachk is Oracle Autonomous Health Framework tooling that validates Exadata, RAC, database, storage, network, configuration, and best-practice conditions. Run it before and after patching, expansions, migrations, major configuration changes, and during periodic health reviews. Findings must be assessed by severity, relevance, and approved remediation; not every informational recommendation should be applied blindly.

Q29. What diagnostic bundles would you collect for an Exadata incident?

Answer: Collect AHF/TFA diagnostics for cluster and database events, exachk, cell alert history, cell metrics, relevant OS logs, ILOM or hardware fault data, switch diagnostics where appropriate, ASM and CRS logs, AWR/ASH, and SQL monitor reports. For hardware or complex engineered-system incidents, use Oracle-supported tools such as sundiag or component-specific bundles as directed by documentation or Support.

Q30. A storage cell fails. What happens and what do you do?

Answer: ASM mirroring and Exadata redundancy should keep databases available if the failure is within design tolerance. Confirm database impact, ASM disk and failgroup state, cell reachability, alerts, and whether a second failure risk exists. Avoid unnecessary ASM drops that trigger large rebalances. Engage hardware support, protect redundancy, monitor I/O redistribution, and replace or recover the cell using supported procedures.

Q31. How do you handle a predictive disk failure?

Answer: Confirm the physical disk and associated cell disk/grid disks, review alert history, and check ASM redundancy. Follow Exadata disk replacement procedures; in many cases the cell software automatically handles poor performance and disk status transitions. Ensure no concurrent fault compromises redundancy, avoid ad hoc OS disk operations, and verify post-replacement recreation and ASM rebalance completion.

Q32. What is your approach to an RDMA fabric issue?

Answer: Identify whether the problem is host, adapter, cable, transceiver, switch port, bonding, MTU, congestion, or configuration. Compare both fabric paths and affected nodes, inspect link and error counters, correlate with database and cell timeouts, and preserve diagnostics before resets. Maintain redundancy and use supported switch and fabric commands; a seemingly small manual network change can destabilize the entire rack.

Q33. How do you plan Exadata capacity?

Answer: Model CPU, memory, database licenses, usable ASM capacity, redundancy, flash working set, XRMEM and flash effectiveness, disk and flash throughput, cell CPU, fabric throughput, backup and recovery windows, growth, rebalance headroom, and failure-mode capacity. Include consolidation peaks, Data Guard role transitions, maintenance, and expected compression ratios based on representative data.

5. Oracle Database Appliance Architecture

Q34. What is Oracle Database Appliance?

Answer: Oracle Database Appliance is an engineered, integrated platform combining compute, storage, networking, Oracle Linux, Grid Infrastructure, ASM/ACFS, database software, and appliance lifecycle tooling. It targets simplified deployment and operation for Oracle databases. ODA models include high-availability two-node systems and single-node systems, with support depending on model and release for bare metal, DB systems, KVM application VMs, and integrated backup or management functions.

Q35. How does ODA differ from Exadata?

Answer: Exadata is designed for extreme scale and uses intelligent storage cells with Smart Scan, storage indexes, flash-log technology, IORM, RDMA storage access, and scale-out database and storage tiers. ODA focuses on simplicity, integrated HA, predictable sizing, and lifecycle automation using conventional Oracle database storage architecture with ASM and ACFS. ODA does not provide Exadata Smart Scan or Exadata storage-cell offload.

Q36. Explain ODA bare metal, DB systems, and application VMs.

Answer: Bare metal runs databases directly on the ODA host operating system. DB systems are KVM-based database environments managed by ODA tooling, providing stronger isolation and flexible database-home or OS separation. Application VMs run non-database application tiers on supported ODA configurations. The DBA must understand resource allocation, CPU pools, storage repositories, networking, patch boundaries, high availability, and backup responsibilities for each model.

Q37. What are odacli and odaadmcli used for?

Answer: odacli is the primary appliance lifecycle command-line interface for creating and managing databases, homes, networks, DB systems, VMs, backups, and patching jobs. odaadmcli is used mainly for lower-level appliance hardware and storage administration and diagnostics. Always confirm the command syntax for the installed ODA release because capabilities and JSON payloads evolve.

Q38. What is the role of DCS components on ODA?

Answer: The Database Configuration Assistant service stack, often referred to through DCS services, maintains appliance metadata and orchestrates lifecycle operations. Components manage agent communication, jobs, repositories, provisioning, and state. When odacli jobs fail, review job details and appropriate DCS logs before manually modifying resources. Unsupported manual changes can create a mismatch between real infrastructure and appliance metadata.

Q39. What storage technologies are used on ODA?

Answer: ODA uses ASM disk groups for database storage and ACFS for relevant appliance repositories, database homes, VM storage, and other managed files depending on deployment model and release. Typical database storage separates DATA and RECO. The DBA must size redundancy, recovery area, backups, snapshots, VM repositories, and free space while respecting ODA-managed layouts.

Q40. How is high availability achieved on an ODA HA system?

Answer: An HA ODA uses two server nodes, shared storage, Oracle Clusterware, RAC or RAC One Node as licensed and configured, redundant networking, and appliance-managed infrastructure. High availability still depends on correct services, database configuration, quorum, backups, client connect descriptors, and operational procedures. ODA simplifies deployment but does not eliminate architecture responsibility.

Q41. What is capacity-on-demand on ODA?

Answer: Capacity-on-demand allows licensed CPU capacity to be controlled in supported increments by limiting enabled cores. It helps align physical deployment with Oracle licensing requirements. Changes must follow ODA-supported procedures and organizational license governance. A DBA should never interpret technical core settings as a substitute for contractual license review.

6. ODA Deployment, Patching, and Management

Q42. Describe the high-level ODA deployment process.

Answer: Rack and cable the appliance, configure ILOM and network prerequisites, verify hardware, install or reimage the required ODA software if necessary, create the appliance configuration, deploy infrastructure, validate Grid Infrastructure and storage, create database homes and databases, configure backups and monitoring, and run acceptance tests. Preserve the deployment JSON and document all addresses, VIPs, SCAN, DNS, NTP, VLANs, and credentials.

Q43. What prechecks are critical before an ODA deployment?

Answer: Validate supported model and software release, power and cooling, cabling, public and private networks, DNS forward and reverse resolution, NTP, IP allocation, VLANs, gateway, hostnames, SCAN where applicable, disk health, firmware, installation media, and database version compatibility. Network mistakes are among the most common causes of failed or unstable deployments.

Q44. How do you create a database on ODA?

Answer: Create or select a supported database home, prepare a database creation payload or use Browser UI, specify database name, version, shape, storage, character set, workload, redundancy, RAC option, and backup configuration, then submit and monitor the odacli job. Validate database, listener, services, storage, backup policy, and appliance metadata after completion.

Q45. What are ODA database shapes?

Answer: Database shapes are predefined resource profiles that assign CPU and memory characteristics to databases or DB systems. They simplify standardization but must be selected using workload evidence. Review SGA/PGA needs, CPU concurrency, consolidation, licensing, failover capacity, and growth. A shape is a starting allocation, not a replacement for performance engineering.

Q46. Explain the ODA patching sequence.

Answer: Patch planning normally covers appliance infrastructure, server components, storage, Grid Infrastructure, database homes, databases, and DB systems or VMs as applicable. Run repository update and prechecks, review component versions, clear failed jobs, confirm backups, patch in Oracle-documented order, monitor odacli jobs, and perform post-validation. Do not mix arbitrary RPM, firmware, GI, or database patching outside the supported ODA workflow.

Q47. An odacli patch job fails. What is your troubleshooting workflow?

Answer: Review odacli describe-job output and failed task, check component versions and precheck reports, inspect DCS and patch logs, verify repository integrity, free space, active databases, cluster health, and network access. Correct the supported root cause and resume or rerun according to documentation. Avoid deleting metadata or manually marking jobs successful unless directed by Oracle Support.

Q48. How do you manage backups on ODA?

Answer: Configure appliance-supported backup destinations and policies, which may include local, NFS, or object storage depending on release and architecture. Define retention, encryption, credentials, archivelog handling, and recovery objectives. Regularly test restore and recovery, monitor backup jobs and destination capacity, and maintain an independent disaster-recovery strategy; appliance-integrated backup is not automatically a complete DR solution.

Q49. How would you migrate a database to ODA?

Answer: Choose RMAN backup/restore, active duplicate, Data Guard, transportable methods, or logical migration according to size, endian format, version, downtime, and rollback needs. Pre-create supported homes, validate character set and options, size DATA/RECO, configure network and services, test application connectivity, and integrate the migrated database into ODA-supported lifecycle management without creating metadata inconsistencies.

Q50. How do you troubleshoot ODA storage-space pressure?

Answer: Determine whether pressure is in ASM DATA, RECO, ACFS repositories, database homes, logs, backups, VM storage, or staging. Check ASM usable capacity and imbalance, FRA usage, guaranteed restore points, archivelogs, obsolete backups, diagnostic data, and abandoned clones. Reclaim only through supported database and appliance procedures; deleting files directly can break recovery or appliance metadata.

Q51. What ODA health checks should be part of routine operations?

Answer: Review odacli jobs and component versions, appliance and hardware status, cluster resources, ASM and ACFS capacity, disk and network health, database and listener status, backups, DCS services, alert logs, OS capacity, ILOM faults, time synchronization, and security findings. Run Oracle health-check tooling appropriate to the release and retain trend reports.

7. Advanced Scenario-Based Questions

Q52. After an Exadata patch, OLTP commit latency increases. How do you investigate?

Answer: Compare pre- and post-patch AWR for log file sync and log file parallel write, redo size, commit rate, CPU, and storage latency. Check Smart Flash Log status, cell alerts, flash health, IORM, redo-log placement, ASM balance, fabric errors, and any database parameter or plan changes. Confirm whether latency is system-wide or limited to particular services. Preserve evidence and validate known issues for the exact database and Exadata image levels before rollback.

Q53. A full-rack Exadata reports high CPU on database nodes but low cell CPU. What does that suggest?

Answer: It often suggests database-layer work rather than storage processing: inefficient joins, parsing, PL/SQL, encryption, decompression, excessive row return, RAC block processing, or SQL that is not offloadable. Review DB CPU, top SQL, execution plans, rows processed, offload eligibility, PX distribution, and gc waits. Increasing storage capacity would not address a database-node CPU bottleneck.

Q54. One Exadata cell shows much higher I/O than others. What do you check?

Answer: Check ASM allocation and grid-disk balance, hot objects, disk-group layout, rebalance state, cell or disk degradation causing retries elsewhere, flash-cache distribution, IORM, and whether a database is pinned to a subset of disks through configuration. Compare object extent distribution and cell metrics. Correct the underlying imbalance using supported ASM and Exadata procedures.

Q55. An ODA HA node is down and the surviving node is overloaded. What actions do you take?

Answer: Prioritize service stability: confirm cluster and database state, stop nonessential workloads, adjust services or Resource Manager, validate storage and network health, and assess whether CPU or memory oversubscription is causing risk. Restore the failed node using supported procedures. Capacity planning should already have defined which critical services must meet SLA on one node and which can be degraded or stopped.

Q56. An ODA database was manually patched with OPatch and now odacli shows an inconsistent version. What do you do?

Answer: Do not conceal the change or edit metadata manually. Capture OPatch inventory, odacli database-home and component information, job history, and logs. Determine whether the home can be brought to a supported state through an ODA-managed clone or patch operation. For metadata inconsistency, open an Oracle Support service request and follow supported remediation, because future appliance patching may otherwise fail.

Q57. How would you choose between Exadata X11M and ODA for a client?

Answer: Use workload and operational requirements. Choose Exadata for very large consolidation, extreme throughput, low-latency scale-out, data warehouse offload, high flash bandwidth, advanced I/O governance, or strategic mission-critical growth. Choose ODA for smaller to mid-sized Oracle estates needing integrated HA, simpler lifecycle management, predictable capacity, and lower operational complexity. Compare five-year capacity, licensing, DR, support, skills, maintenance windows, and failure-mode performance rather than acquisition price alone.

Q58. How do you demonstrate Level 4 ownership during a major engineered-systems incident?

Answer: Establish incident command, protect data and redundancy, define impact and timeline, freeze risky changes, assign database, storage, network, and hardware workstreams, and collect synchronized evidence. Communicate facts and uncertainty clearly, engage Oracle Support with complete diagnostics, test each corrective action, maintain rollback, and produce a root-cause report with measurable preventive actions, monitoring improvements, and runbook changes.

Q59. What questions would you ask before accepting handover of an Exadata or ODA estate?

Answer: Ask for architecture diagrams, versions, patch history, support identifiers, rack and network topology, capacity and growth, database inventory, licensing boundaries, backup and recovery tests, Data Guard, monitoring, alert routes, known defects, security baselines, custom scripts, maintenance procedures, vendor contacts, recent incidents, performance baselines, passwords and key custody, and evidence of single-component failure tests. Validate documentation against the live system.

Reference Commands for Interview Discussion

Purpose	Example
Exadata cell status	`cellcli -e "list cell detail"`
Physical disks	`cellcli -e "list physicaldisk attributes name,status,disktype"`
Grid disks	`cellcli -e "list griddisk attributes name,status,asmdeactivationoutcome"`
Cell alerts	`cellcli -e "list alerthistory where severity != clear"`
Flash cache	`cellcli -e "list flashcache detail"`
ASM disk groups	`asmcmd lsdg`
Cluster resources	`crsctl stat res -t`
ODA appliance details	`odacli describe-appliance`
ODA component versions	`odacli describe-component`
ODA jobs	`odacli list-jobs`
ODA job detail	`odacli describe-job -i`
ODA databases	`odacli list-databases`
ODA database homes	`odacli list-dbhomes`
ODA hardware checks	`odaadmcli show server; odaadmcli show storage`

Command syntax and availability vary by image and ODA release. Confirm against the documentation for the installed version before production use.

Official References

- Oracle Exadata Database Machine X11M Data Sheet.
- Oracle Exadata System Overview and X11M Hardware Component documentation.
- Oracle Exadata Database Machine Documentation Library and Exadata System Software release documentation.
- Oracle Database Appliance Documentation Library.
- Oracle Database Appliance Release Notes, including 19.30 and 19.31 documentation available in 2026.
- Oracle Database Appliance X11 Deployment and User Guide.
- Oracle Autonomous Health Framework and Exadata health-check documentation.

Prepared as an interview and study guide. Product capabilities, support matrices, commands, and patching workflows can change by release; always validate production actions against the exact Oracle documentation and My Oracle Support notes for the installed system.