

Exadata Software 12.2.1.1.0

*The Best Database Cloud Platform
Just Got Better!*

Kothanda “Kodi” Umamageswaran
Vice President, Exadata Development

Gurmeet Goindi
Exadata Product Management




ORACLE®

Safe Harbor Statement

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

Exadata 12.2.1.1.0 Highlights

- Full storage offload functionality for Database 12.2
 - Database 11.2, 12.1, and 10.2 can coexist along side 12.2 on the same system
- Over 30 unique software features and enhancements
 - Better analytics, better transaction processing, better consolidation, more secure, faster and more robust upgrades, and easier to manage
- Complete investment protection
 - All new software features work on all supported Exadata hardware generations
- Updated Oracle Linux kernel and Oracle VM improve robustness and scalability
 - Oracle Linux 6.8 with UEK4, Oracle Virtual Machine 3.4.2






Better Analytics

Smart Analytics: In-Memory Formats in Columnar Flash Cache

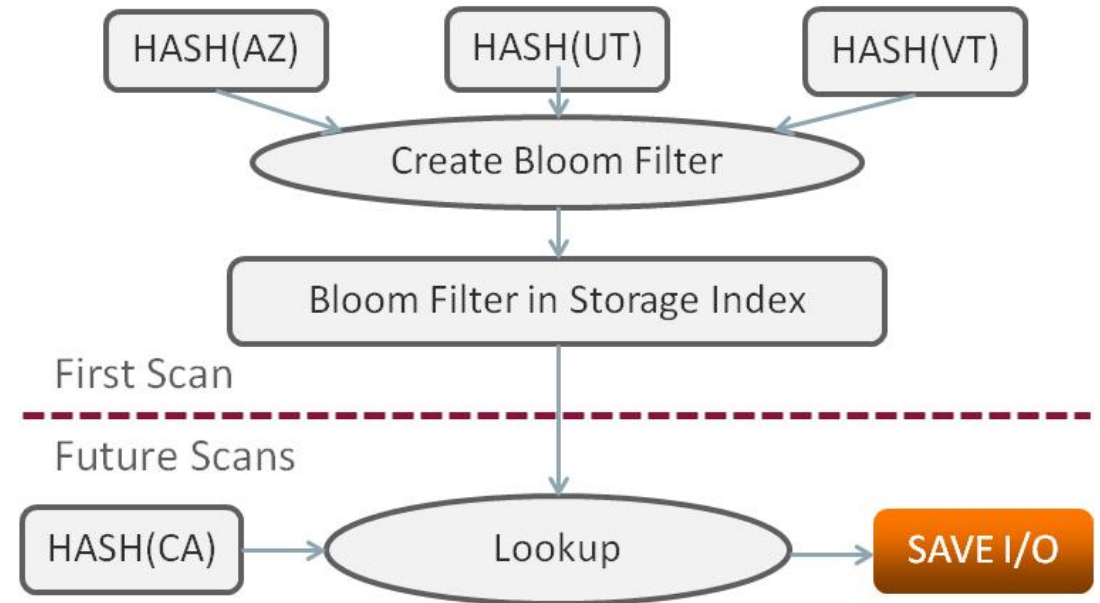
- In-Memory formats used in Smart Columnar Flash Cache
- Enables vector processing on storage server during smart scans
 - Multiple column values evaluated in single instruction
- Faster decompression speed than Hybrid Columnar Compression
- Enables dictionary lookup and avoids processing unnecessary rows
- Smart Scan results sent back to database in In-Memory Columnar format
 - Reduces Database node CPU utilization
- **In-memory performance seamlessly extended from DB node DRAM memory to 10x larger capacity flash in storage**
 - Even bigger differentiation against all-flash arrays and other in-memory databases
- **Supports Database 12.1.0.2 and Database 12.2.0.1**



Smart Analytics: Storage Index Set Membership

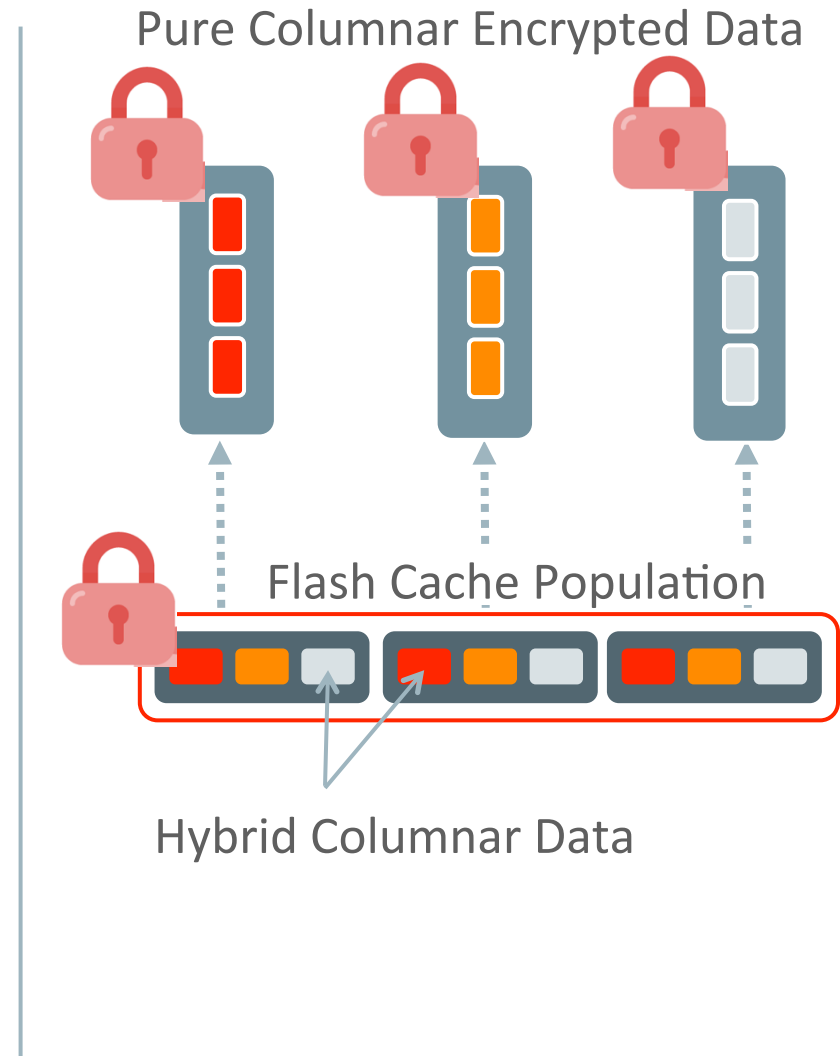
- Storage Index help skip unnecessary IOs
 - Currently contains up to 8 columns of min/max summary
 - Created automatically and kept in memory
- What about queries with low cardinality columns?
`select name, address from travels
where origin='Sierra Leone' and dest='CA'`
- Traditional min/max not good enough
- Database gathers stats and find that column has less than 400 distinct values
- Database requests storage to compute bloom filter
- Storage will compute distinct values and create a bloom filter
- Smart Scans check value 'CA' against bloom filter and saves performing I/O
- Supports Database 12.1.0.2 and Database 12.2.0.1

ORIGIN	DEST	NAME	ADDRESS
Sierra Leone	AZ	Alice	...
Sierra Leone	UT	Bob	...
Sierra Leone	VT	John	...



Smart Analytics: More Enhancements


- Storage Index will now store information about 24 columns instead of 8
 - Metadata space shared by set membership summary and min/max summary
 - Workload determines which summary gets stored
- Columnar Flash Cache for encrypted tablespaces
 - Data stays encrypted in columnar flash cache
 - Works seamlessly with In-Memory Formats in Flash Cache
 - Benefits Exadata Cloud Service as all data is encrypted in Oracle Public Cloud
- **Supports Database 12.1.0.2 and Database 12.2.0.1**



Smart Analytics: Smart Write Bursts and Temp IO in Flash Cache

- Write throughput of four flash cards has become greater than the write throughput of 12-disks
- When database write throughput exceeds throughput of disks, Smart Flash Cache intelligently caches writes
- When queries write a lot of temp IO, Smart Flash Cache intelligently caches temp IO
 - Writes to flash for temp spill reduces elapsed time
 - Reads from flash for temp reduces elapsed time further
- Smart Flash Cache prioritizes OLTP data and does not remove hot OLTP lines from the cache
- Smart flash wear management for large writes
- Supports Database 11.2.0.4, 12.1.0.2 and 12.2.0.1

Write Bursts and
Temp IO in
Flash Cache



Accelerates Large Joins and Sorts
and Large Data Loads

Smart Analytics: Join and Aggregation Smart Scan

- Extend In-Memory Aggregation technique into storage (vector joins and vector aggregation)

- Find Sales per country

```
SELECT /*+ VECTOR_TRANSFORM */ country_id,  
      sum(amount_sold) amount_sold  
FROM customers, sales  
WHERE customers.cust_id = sales.cust_id  
GROUP BY customers.country_id  
ORDER BY customers.country_id;
```

- Storage cells scanning sales fact table return tuples


```
{country_id, sum_amount_sold }
```

- **Join and Aggregation offloaded to the storage server**

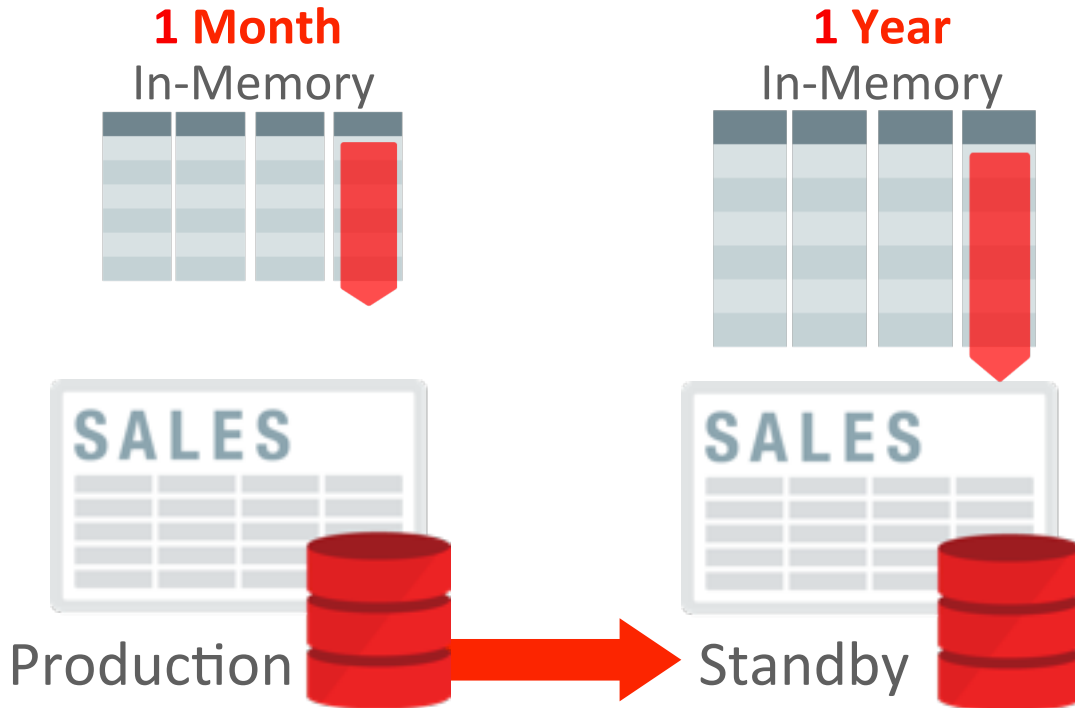


Smart Analytics: More Smart Scan Enhancements

- Smart Scan Offload for Compressed Index Scan
- Smart Scan enhancements for XML
 - Enhancements to XMLExists, XMLCast and XMLQuery
- Smart Scan offload enhancements for LOBs
 - Extended to “LENGTH, SUBSTR, INSTRM CONCAT, LPAD, RPAD, LTRIM, RTRIM, LOWER, UPPER, NLS_LOWER, NLS_UPPER, NVL, REPLACE, REGEXP_INSTR, TO_CHAR”



Mixed Workload: In-Memory on Active Data Guard



- Real-time analytics with no impact on production database
- Make productive use of standby database resources
- Can populate different data from production database
 - Use new `DISTRIBUTE BY SERVICE` to determine where to populate a table
 - Increase total columnar capacity
- **Exclusively available on Exadata and DB PaaS**




Better Consolidation



Smart Consolidation: Up To 4,000 Pluggable Databases

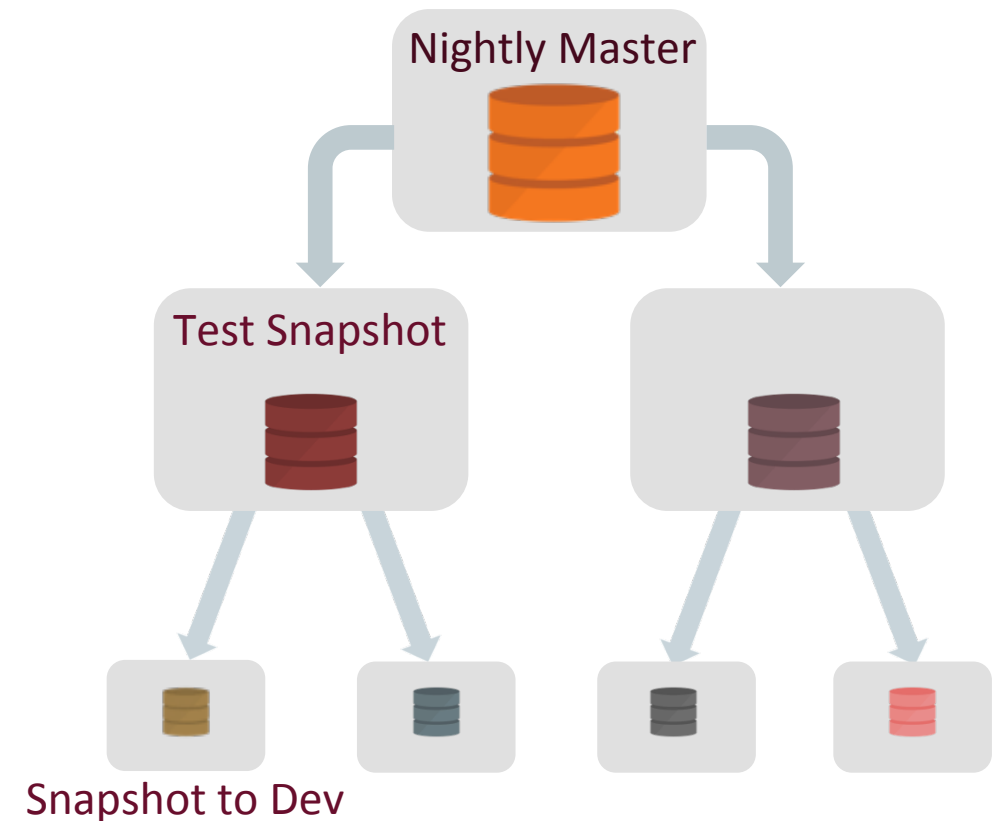
NEW IN
DB 12.2

- Exadata offers unique end-to-end resource management and consolidation capabilities
- Multitenant Option now allows greater than 252 Pluggable Databases within a single Oracle Multitenant Container Database
 - Up to thousands
- **Exclusively available on Exadata and DB PaaS**




Smart Consolidation: Hierarchical Snapshots


- Hierarchical Snapshots
 - Create snapshots of databases on previously created snapshots
 - Syntax and technology remain unchanged
 - Works with pluggable and non-pluggable databases
 - Use case example
 - Development releases nightly build of the database
 - Tester creates a snapshot for himself and finds a bug
 - Tester creates a snapshot of his snapshot
 - Tester provides the new copy back to development for analysis
- Sparse backup of snapshots
 - RMAN backs up the modified blocks and not the unchanged blocks from parent



Smart Consolidation: Additional Ethernet Connectivity



- Additional Ethernet card offered for customers that require additional connectivity
 - Dual port 10GE copper connectivity (part number 7100488)...OR
 - Dual port 10GE optical connectivity (part number X1109A-Z)



- Installed in PCIe slot 1 on the Oracle Exadata X6-2 database server
- Exadata SW automatically recognizes the new card and configures the two ports as eth6 and eth7 interfaces on the database server
- Additional ports can be used for any purpose



Better OLTP and Availability

Smart OLTP: End-to-End IO Latency Capping

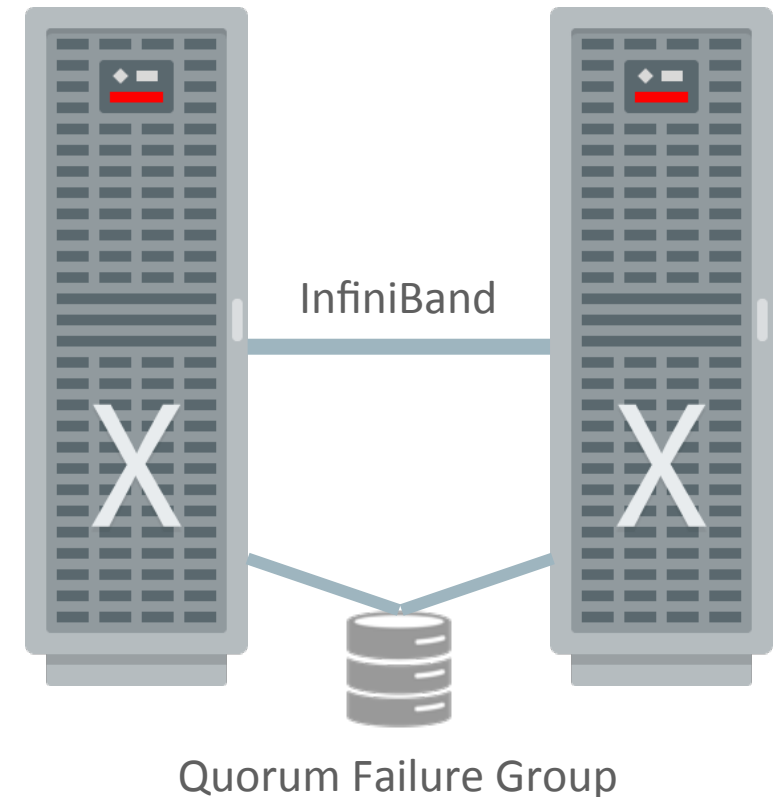
- Exadata Storage Server software detects and automatically eliminates IO latency outliers on disk and flash media
- On very rare occasions network outliers can deteriorate latency between database and storage servers
- Database 12.2 automatically redirects slow read I/O operations to another Exadata storage server
- Ensures end-to-end low latency for OLTP read I/Os



Smart OLTP: Extended Distance Clusters

NEW IN
DB 12.2

- Data Mirroring between two nearby sites
- InfiniBand connected for high performance
 - Limited to 100m optical cables in 2016 (best for fire cells)
- Implemented using 12.2 ASM Extended Diskgroups
 - Nested failure groups
- Compute nodes at each site read data local to that site
- Data is written to all sites
- Smart Scans scan across cells on both sites increasing throughput
 - Row filtering, column projection, storage index, and flash cache provide extreme performance
- Data Guard continues to be the recommended DR solution




Smart Availability: Efficient Rebalance and Restore

- Intelligent and flexible rebalance power setting
 - Dynamically change `ASM_POWER_LIMIT`
- ASM rebalance restores redundancy first
 - Drastically reduces secondary failure exposure window
 - Exposed via new `REBUILD` phase in `v$asm_operation`
- Exadata leverages flash cache for rebalance reads
 - Improves performance of redundancy restoration by up to 30%



Smart OLTP: New Redo Log Write Metrics

- Redo Log Write response times are very critical to large scale OLTP systems
- New metrics to measure
 - Overall IO latency
 - Networking and other overhead
 - IOs serviced by Flash Log
 - Overall latency per storage server
- Improves visibility of redo log write performance



Transformational Full System Availability with Oracle



“Exadata and SuperCluster both achieve AL4 fault tolerance in a **Maximum Availability Architecture** configuration”



Only other AL4 Systems

- IBM - z Systems
- HPE - Integrity NonStop & Superdome
- Fujitsu – GS & BS2000
- NEC – FT Server/320 Series
- Stratus ftServer & V Series
- Unisys – Dorado

A New Gold Standard




Better Security and Manageability



Smart Manageability: Super Fast and Robust Software Updates

- Oracle Public Cloud is the largest Exadata deployment with hundreds of Exadata Database Machine deployed
 - Each software release goes through thousands of upgrade cycles
 - All upgrade utilities get exercised thousands of times as well
 - Contributes to the robustness of the software release and utilities
- **5x speed up** in Storage Server Software Update
 - Parallel firmware upgrades across components such as hard disks, flash, ILOM BIOS, InfiniBand card
 - Reduced reboots for Software updates, use kexec where possible
- Manage a Cloud instead of managing a single rack
 - Use single patchmgr utility to upgrade hundreds of racks
- Enable patchmgr to run from a non-Exadata system and run as low privileged user



Better Security: Full Stack Secure Erase

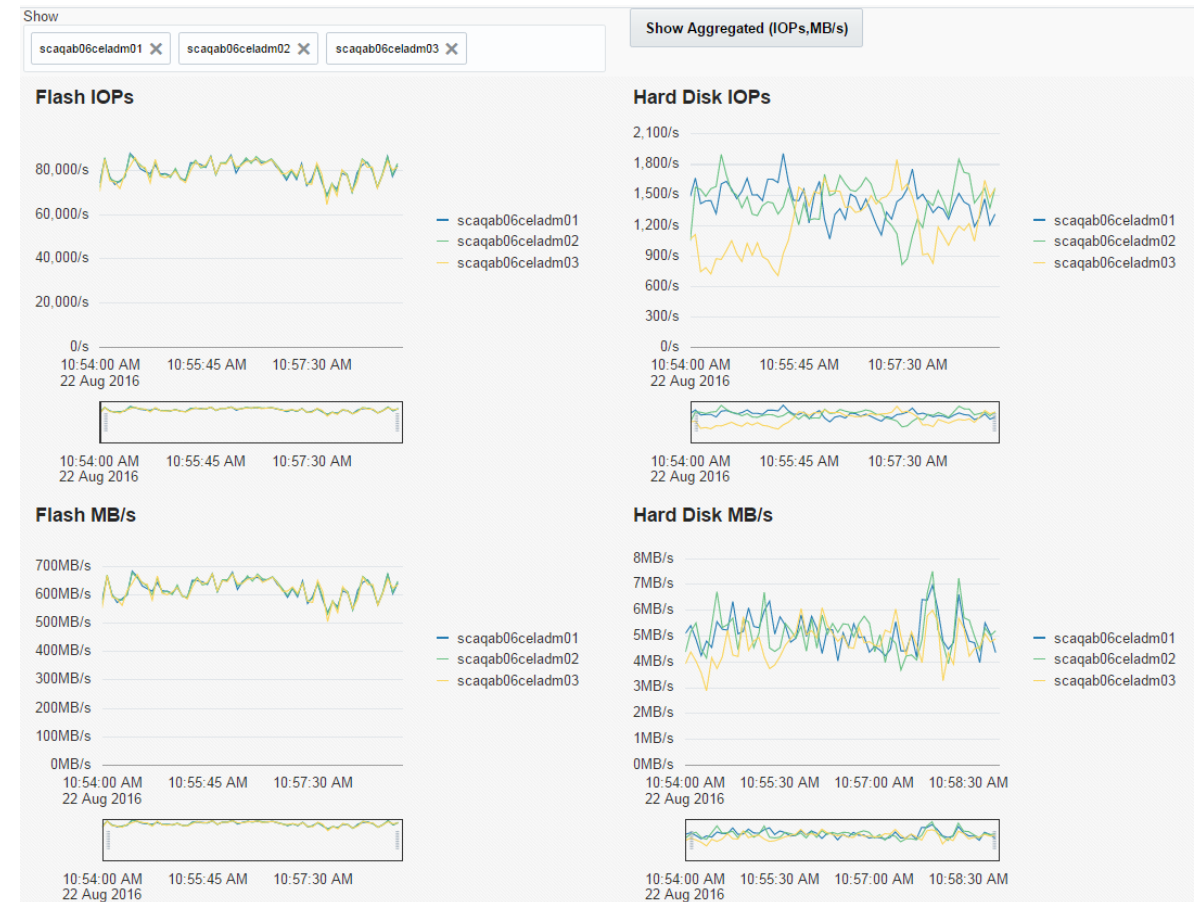
- Efficient, secure, verifiable mechanism to erase all user information and data across the Exadata stack during:
 - Relocation or End of life operations of systems
 - Replacing failed media devices
 - Onboarding new private or public cloud users
- Components covered: Media, storage and compute servers, switches, and PDUs
- Only comprehensive solution across the stack, others offer a piecemeal approach
- Leverages crypto erase and other modern techniques where possible
- Complies with NIST and DoD standards

Smart Manageability: Miscellaneous Enhancements

- Storage Server Rescue Plan
 - Create a per server repository of customizations such as IORM plans, thresholds, etc.
 - Invoke the rescue plan upon restoration
- Make DNS, NTP, and other IP address changes online
- Storage server to storage server offload support for ASM Scoped-Security
 - Enable keys, called cell keys, on storage servers to authenticate offload operation
 - Applies to ASM resync, rebalance, resilver, rebuild and database high throughput write operations
- Quarantine Manager support for high throughput write
 - Automatically detects and quarantines any anomalous operation and uses alternate IO path

Smart Manageability: Miscellaneous Enhancements

- IPv6 + Virtual machine + VLAN deployments
- Get graphs from Exawatcher
- Seamless customer service with Automatic Service Requests sending diagnostic attachments
- Oracle installer support for Quorum Disks in Database 12.2.0.1
- Manage Compute nodes using a RESTful service
 - ExaCli enabled for compute nodes in addition to storage cells



Exadata Advantages Increase Every Year

Transformational OLTP, Analytics, Consolidation

Cloud Without Compromise

Smart Software

- Smart Scan
- InfiniBand Scale-Out

Smart Hardware

- Scale-Out Servers

- Database Aware Flash Cache
- Storage Indexes
- Columnar Compression

- IO Priorities
- Data Mining Offload
- Offload Decrypt on Scans

- DB Processors in Storage

- Scale-Out Storage

- Network Resource Management
- Multitenant Aware Resource Mgmt
- Prioritized File Recovery

- PCIe NVMe Flash

- Unified InfiniBand

- Tiered Disk/ Flash

- In-Memory Fault Tolerance
- Direct-to-wire Protocol
- JSON and XML offload
- Instant failure detection

- **Exadata Cloud Service**
- **In-Memory Columnar in Flash**
- **Smart Fusion Block Transfer**

- **3D V-NAND Flash**
- **Software-in-Silicon**



Integrated Cloud

Applications & Platform Services

ORACLE®