An Oracle White Paper
June 2012

Managing Oracle Exadata with Oracle Enterprise Manager 12*c*

ORACLE

Executive Overview

Oracle Enterprise Manager 12*c* is Oracle's integrated enterprise IT management product, providing the industry's first complete cloud lifecycle management solution. With Oracle Enterprise Manager 12*c*, you can create self-service IT, simplify and automate IT operations, and manage applications in a way that delivers maximum business value. Oracle Enterprise Manager 12*c*'s business-driven application management capabilities allow you to quickly set up, manage and support enterprise clouds as well as traditional Oracle environments—from applications-to-disk, helping you to achieve:

- *Best service levels for traditional and cloud applications* through management from a business perspective including Oracle Fusion Applications
- *Maximum return on IT management* investment through the best solutions for intelligent management of the Oracle stack and engineered systems such as Oracle Exadata Database Machine
- *Unmatched customer support experience* through real-time integration of Oracle's knowledgebase with each customer environment.

Introduction

Today's IT organizations want to reduce deployment costs and errors while ensuring maximum availability along with zero downtime application migration. They also want proactive monitoring for all components in order to ensure the highest service quality possible.

In the latest release of Oracle Enterprise Manager 12*c,* Oracle has designed a management solution that does just that. Oracle Enterprise Manager 12*c* uses a holistic approach to manage Oracle Exadata Database Machine, providing comprehensive lifecycle management, from testing and deployment, to proactive monitoring and ongoing maintenance across the entire engineered system. By replacing manual deployment tasks with powerful automation capabilities, Oracle Enterprise Manager 12*c* allows organizations to manage every stage and aspect of Oracle Exadata's lifecycle. The entire process is automated through Oracle Enterprise Manager 12*c*'s centralized management console, helping customers to become more proactive and more strategic at managing Oracle Exadata—which ultimately results in lower risks and higher availability for your applications. A comprehensive and consistent use of Oracle Enterprise Manager will not only accelerate problem resolution, but also reduce management efforts, complexity, costs and risks associated with implementing and maintaining business-critical applications, databases and related infrastructure.

Oracle Exadata Database Machine

Oracle Exadata Database Machine is a complete and fully integrated database system that includes all the components to quickly and easily deploy an Oracle Database Enterprise Edition delivering world class performance. Oracle Exadata Storage Server (Oracle Exadata storage or Oracle Exadata cells) is used as the storage for Oracle Database in Oracle Exadata and it runs the Oracle Exadata Storage Server software that provides the unique and powerful Oracle Exadata technology including Smart Scan, Smart Flash Cache, Smart Flash Logging, IO Resource Manager, Storage Indexes and Hybrid Columnar Compression. Oracle Exadata Storage Expansion Rack is a fast and simple means to grow the Oracle Exadata storage capacity and bandwidth of an existing Oracle Exadata or SPARC Super Cluster deployment.

Oracle Exadata is a pre-configured system ready to be turned on from day one, reducing integration work, cost and time in the database deployment process. Since Oracle engineers internally work on the same configuration that gets shipped to the customers Oracle Support is very familiar with how to service the system resulting in a superior support experience with the system. The benefit of a common infrastructure to deploy a database for any application, whether OLTP, DW, a mix of the two, or as a platform for consolidation several databases, creates tremendous opportunities for efficiencies in the datacenter. It truly is a "cloud in box".

There are two versions of Oracle Exadata, X2-2 and X2-8. Oracle Exadata X2-2 expands from 2 twelve-core database servers with 192 GB of memory and 3 Oracle Exadata Storage Servers to 8 twelve-core database servers with 1,152 GB of memory and 14 Oracle Exadata Storage Servers, all in a single rack. Oracle Exadata X2-8 is comprised of 2 eighty-core database servers with 4 TB of memory and 14 Oracle Exadata Storage Servers, in a single rack. The X2-2 provides a convenient entry point to the Oracle Exadata family with the largest degree of expandability in a single rack. The X2-8 is for large deployments with larger memory requirements or for consolidating multiple databases on to a single system. Both versions run Oracle Database 11*g* Release 2.

The hardware components of Oracle Exadata X2-2 are shown in Table 1.

	Oracle Exadata X2-2 Full Rack	Oracle Exadata X2-2 Half Rack	Oracle Exadata X2-2 Quarter Rack
Database Servers	8	4	2
Oracle Exadata Storage Servers	14	7	3
InfiniBand Switches	3	3	2

Table 1: Hardware Components of an Exadata X2-2

Managing Oracle Exadata

Oracle Exadata management and monitoring has been divided into four lifecycle categories which include:

- **Setup and Proactive Monitoring:** Once the machine has been delivered to your site, the initial task is to do the set up and establish proactive monitoring for Oracle Exadata targets and components.
- **Test:** Once monitoring has been completed, the next step is to identify and consolidate your application's performance on Oracle Exadata, making sure to properly test and resolve any performance regression within the application before deploying it to production.
- **Manage:** During this stage of the lifecycle, the focus is on application and database performance—making sure that optimal performance is consistently delivered within the specified service level agreements.
- **Maintain:** In the final lifecycle stage, we ensure that all the configurations are consistent and all the necessary patches have been applied in order to meet compliance and regulatory requirements.

In the following sections we will discuss how Oracle Enterprise Manager 12*c* addresses each stage of Oracle Exadata's lifecycle. [Figure 1]

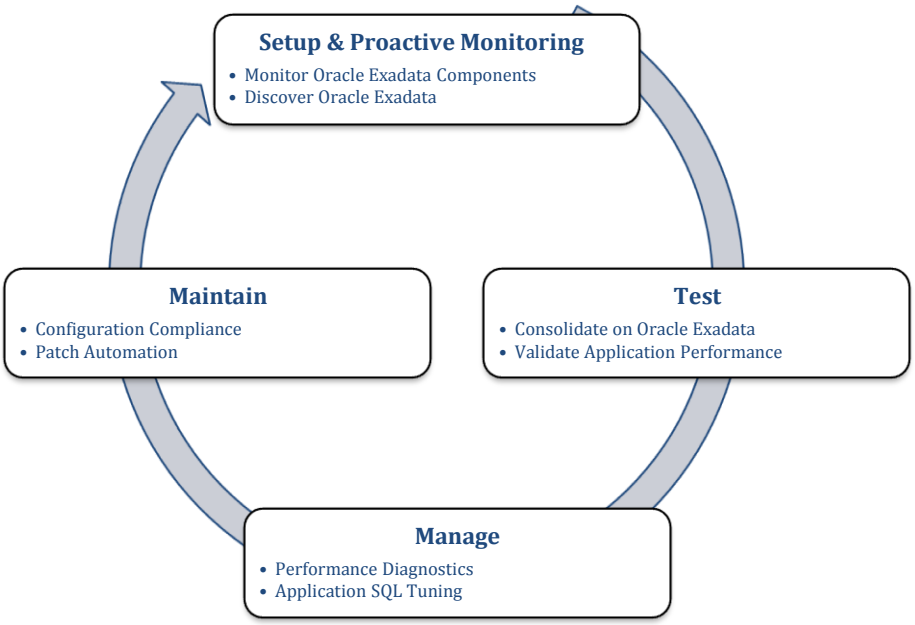

Figure 1: Four stages of the Exadata Management Lifecycle

Setup and Proactive Monitoring

Oracle Enterprise Manager 12*c* provides comprehensive monitoring and notifications to give administrators the ability to proactively detect and respond to potential problems quickly within Oracle Exadata, whether it's a software or hardware issue. Administrators can easily adjust these monitoring settings to suit the needs of their data center environment. If an alert is sent out, administrators can easily view its history and the associated performance metrics for the problem component, such as the network or an InfiniBand port or even the disk activity of an Oracle Exadata storage cell. Administrators can quickly identify the root cause of the problem. With direct connectivity into Oracle Exadata's hardware components, Oracle Enterprise Manager 12*c* can alert administrators to hardware-related faults and log service requests automatically through integration with Oracle Automatic Service Requests (ASR) for immediate review by Oracle Support.

Problems that would have required a combination of database, system and storage administrators to be detected in traditional systems, can now be diagnosed in minutes instead of hours or days because of Oracle Enterprise Manager 12*c*'s powerful and integrated systems monitoring capabilities for Oracle Exadata.

Setup

With Oracle Enterprise Manager 12*c* discovering Oracle Exadata for management and monitoring has become extremely easy. In Oracle Enterprise Manger 12*c* administrators can use wizard-driven guided discovery to perform the set up. Apart from the manual discovery from the Oracle Enterprise Manager console an Enterprise Manager Automation kit is also available to help speed up the discovery process. Using Oracle Enterprise Manager 12*c*, setup can be done in a matter of minutes taking minimum number of steps to discover Oracle Exadata components.

In Oracle Enterprise Manager 12*c* there has been a fundamental change in the agent architecture. The agent has been redesigned to be more modular and lightweight. Monitoring is done through a set of plug-ins that are deployed to the agent based on the applications and software products running on the server. Oracle Exadata is no exception to this method. All the code required for monitoring the various Oracle Exadata hardware components has been bundled into the Oracle Exadata plug-in that is pushed to the agents running on the Oracle Exadata Compute (or Database) nodes. Once the plug-in is deployed on the compute nodes, the database machine is ready for discovery with Oracle Enterprise Manager.

Along with the option to discover a brand new database machine, Oracle Enterprise Manager also provides an option to add components to an already discovered target. The latter is really useful when you are planning to upgrade from a quarter rack to a half rack or from a half rack to a full Oracle Exadata rack.

The following flow diagram illustrates the steps required to discover Oracle Exadata in Oracle Enterprise Manager 12c. The guided discovery process assumes that an installation of Oracle Enterprise Manager 12c already exists. The first step begins with deploying the agents on the Oracle Exadata Database Compute Nodes followed by pushing the Oracle Exadata Plug-in on to the agents. Once the plug-in has been deployed the next step is to start the automatic discovery of the target. In

each of the discovery steps, the user is requested to provide login credentials for the different Oracle Exadata components followed by agent assignment to all the components. Since agents only run on the compute nodes it is important to assign specific agents to different Oracle Exadata components. Agent assignment can be automatic using a pre-defined allocation or the user can force a manual assignment. Each component also has a backup agent assigned in case the primary monitoring agent is brought down or is not available. At the final step the user is asked to review the configuration details before submitting the process to register the target.

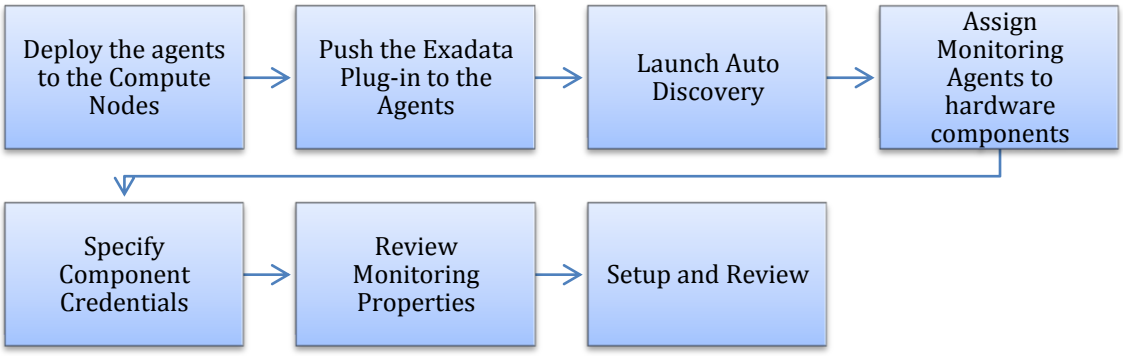

Figure 2: Oracle Exadata Discovery in Oracle Enterprise Manager 12*c*

Although the discovery of Oracle Exadata is extremely efficient the Enterprise Manager Automation kit for Oracle Exadata automates the first two steps of agent deployment and Oracle Exadata plug-in push. The kit comes in two patches available for download from support.oracle.com. Check support note 1110675.1 for the download links to the latest version of the kit.

The Readme of both the patches describe the steps required to install and execute the code. The second patch is only needed if there is no existing Oracle Enterprise Manager 12*c* installation available. Once Enterprise Manager is installed (or there is an existing installation) the first patch is used to deploy the agents and Oracle Exadata plug-in.

Proactive Monitoring

Oracle Enterprise Manager 12*c* provides a unified view of Oracle Exadata hardware and software with deep insight into the health and performance of all components such as the database servers, InfiniBand switches, Oracle Exadata storage cells, Oracle Databases, Automatic Storage Management (ASM), etc. Oracle Databases run transparently on Oracle Exadata without any changes. However, there are times when a database administrator (DBA) needs to drill down from the database tier into the storage system in order to identify and diagnose performance related bottlenecks or hardware faults. Oracle Enterprise Manager's integrated view of Exadata's hardware and software allows DBAs to navigate seamlessly from the familiar database performance pages to the associated Oracle Exadata

storage server to isolate the problem, whether they may be caused by a hardware component or other databases running on the same storage subsystem.

Oracle Exadata consists of the following components:

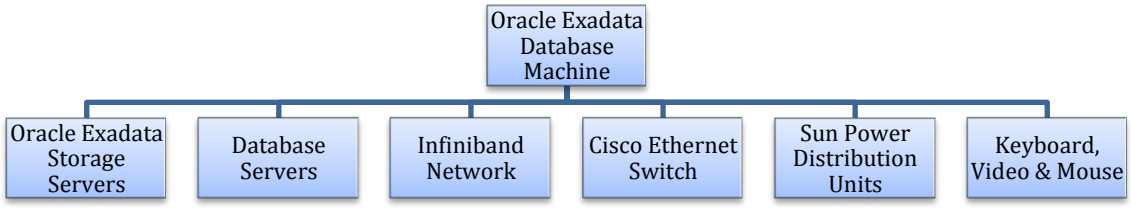

Figure 3: Oracle Exadata Components

Oracle Enterprise Manager 12*c* provides unique capabilities to manage each component individually and also collectively as one. The first step is to understand how Oracle Enterprise Manager is collecting all the monitoring information and how each component is individually monitored.

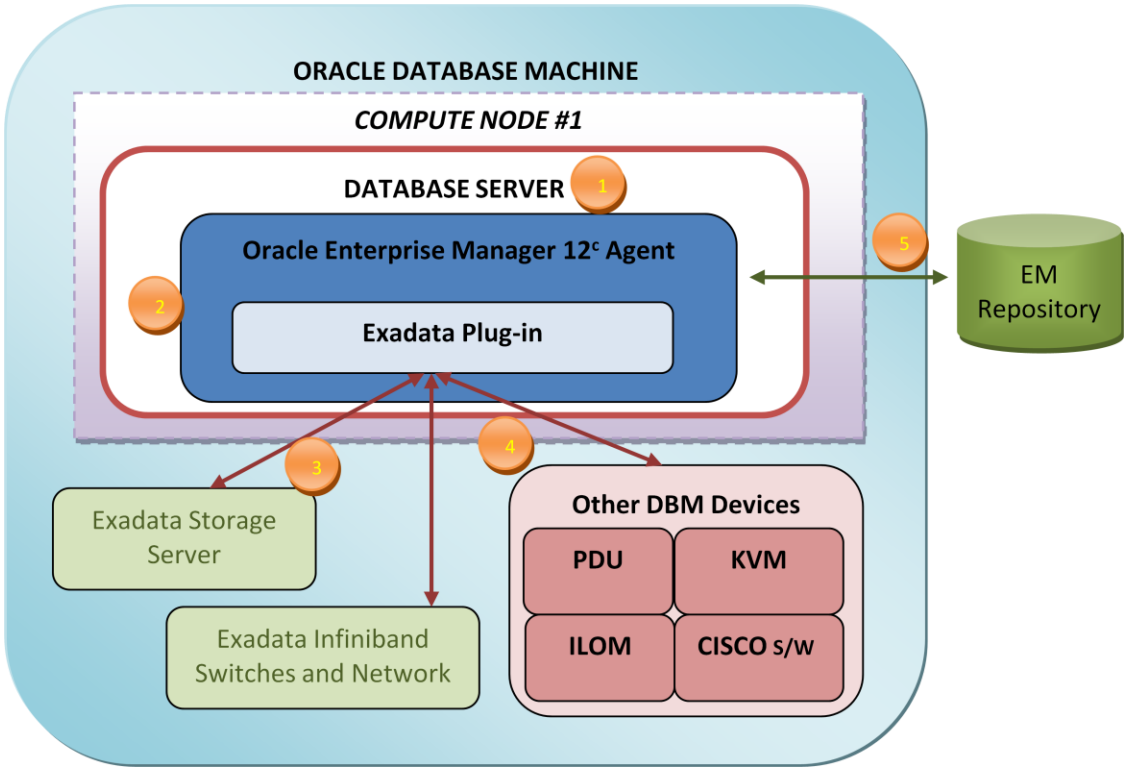

Figure 4: Oracle Enterprise Manager 12*c* agent communicating with Oracle Exadata components

How does the agent communicate with Oracle Exadata's different components?

1. Oracle Enterprise Manager 12*c* Agent deployed on the Compute Node.
2. Oracle Exadata Plug-in deployed with the Agent.
3. Agent communicates with Storage Server and Infiniband Switch targets directly.
4. Oracle Exadata Plug-in also monitors the other DBM components.
5. Oracle Enterprise Manager 12*c* agent collects data and communicates with the remote Enterprise Manager Repository.

Integrated Software and Hardware View

Once Oracle Exadata has been discovered in Oracle Enterprise Manager, the machine is physically represented in a schematic view (Figure 5). The schematic display also supports viewing multiple Exadata nodes that are connected to each other using the same InfiniBand Network. On the schematic display, the availability information for each component is indicated by a green (available) or red (not available) status. The schematic view also has the option to display the temperature for each component and the Edit option allows administrators to change the placement of any cell or compute node to match any physical movement in the server. Oracle Exadata's schematic page also comes with a target navigation bar on the left for directly drilling down to any particular component of your choice. The overview section at the top of the page also shows the summary of incidents and availability information for all targets.

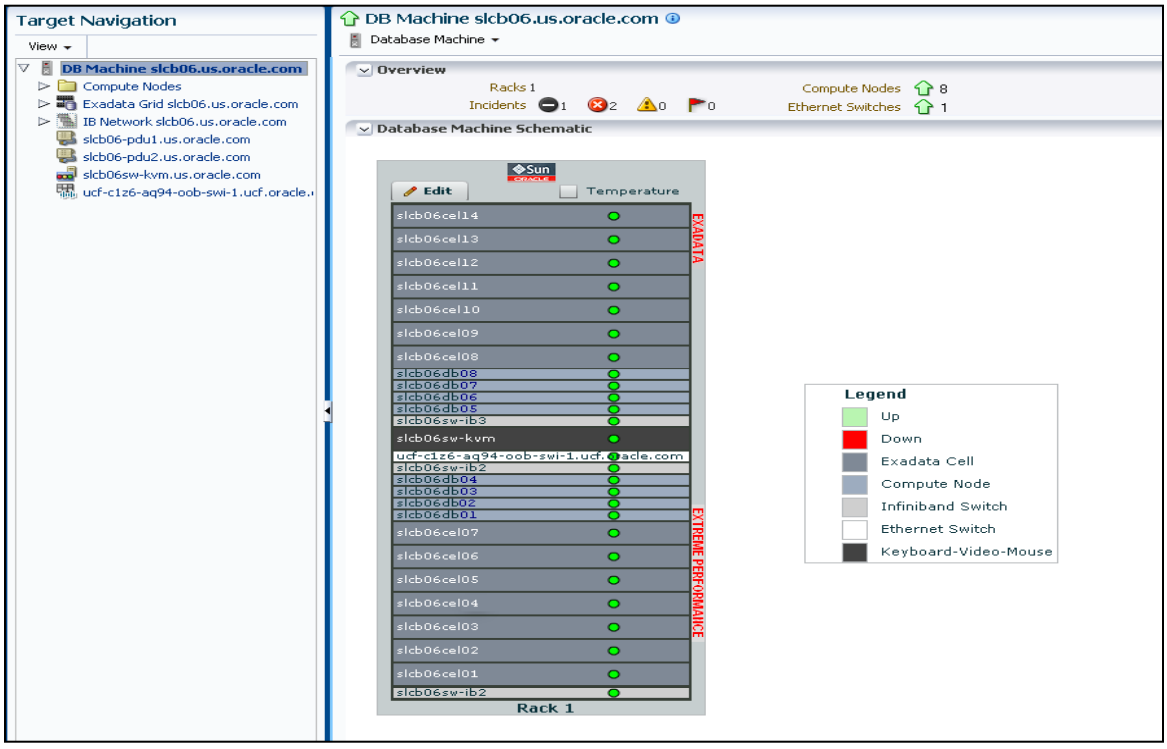

Figure 5: Oracle Exadata Schematic Page

The 'Incidents' region at the bottom of the page provides a list of all open incidents associated with Exadata.

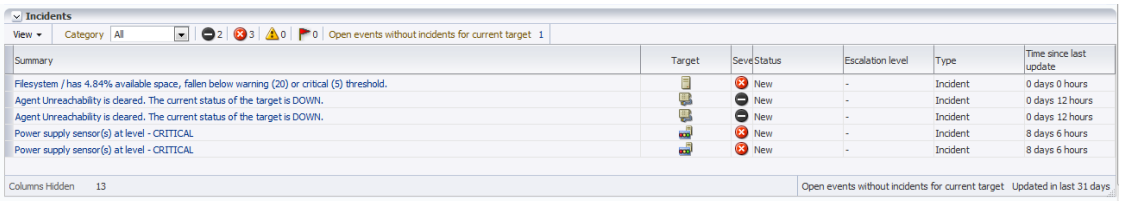

Figure 6: List of Incidents associated with Exadata Database Machine

Oracle Enterprise Manager 12*c* comes with the resource utilization views for the database compute nodes and Oracle Exadata Storage Cells. These views are unique to Exadata as they give an administrator the understanding of CPU consumption on each Compute Node along with which storage cell is experiencing the maximum I/O and the I/O utilization by databases.

Figure 7: CPU Utilization by Databases

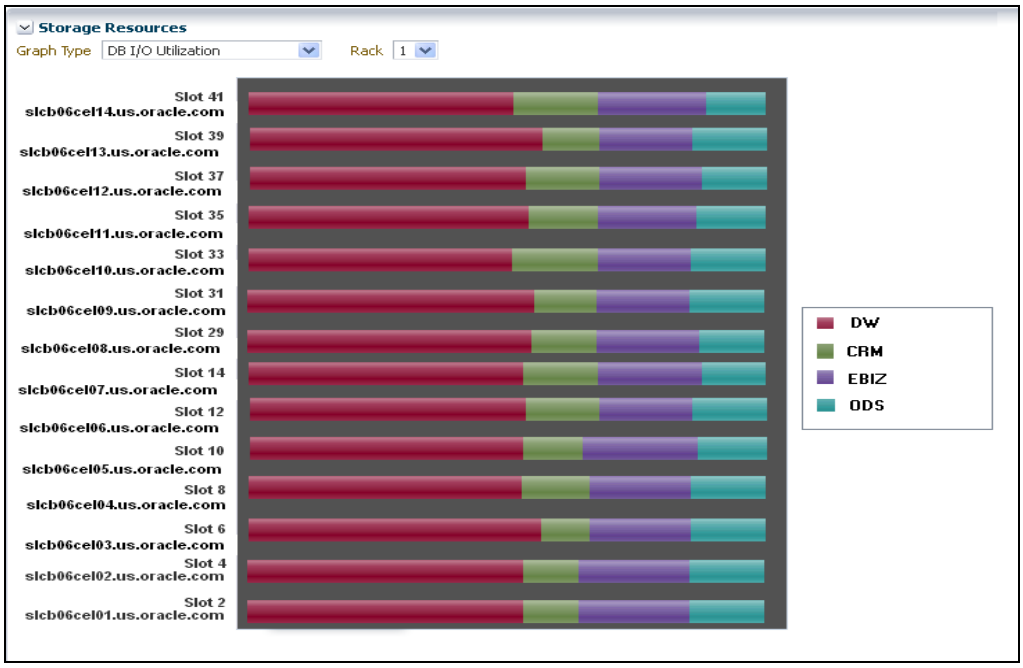

Figure 8: I/O Utilization by Databases

Using the Configuration view you can also explore the different configuration settings of Oracle Exadata.

Oracle Exadata Storage Server

An Oracle Exadata Storage Server is monitored as a single target in Oracle Enterprise Manager 12*c* and covers the following components:

- Hardware
- Operating system
- Oracle Exadata Storage Server software

Oracle Exadata Storage Server Management Server (MS) collects, computes, and manages storage server metrics. These storage server metrics are then gathered by the Oracle Exadata Plug-In and presented as Enterprise Manager metrics.

All Oracle Exadata Storage Server alerts are delivered by the storage server to Oracle Enterprise Manager 12*c* using Simple Network Management Protocol (SNMP). The communication between the Oracle Exadata Storage Server and Oracle Enterprise Manager 12*c* is done through the Oracle Enterprise Manager Agent (Oracle Exadata Plug-in).

There are two types of server alerts that come from Oracle Exadata Storage Server:

- For Integrated Lights Out Manager (ILOM)-monitored hardware components, ILOM reports a failure or threshold exceeded condition as an SNMP trap, which is received by MS. MS processes the trap, creates an alert for the storage server, and delivers the alert via SNMP to Oracle Enterprise Manager 12*c*.

- For MS-monitored hardware and software components, MS processes a failure or threshold exceeded condition for these components, creates an alert, and delivers the alert via SNMP to Oracle Enterprise Manager Cloud Control 12c.
-

From an end-user perspective there is no difference between these two kinds of alerts.

Oracle Enterprise Manager 12*c* has a rich new user experience that creates the Oracle Exadata Grid page. This is the consolidated view of I/O load, CPU utilization, Network utilization and average Response Time. The overview region displays the status along with the version of the storage software and status of the I/O Resource Manager (IORM). The target administration screens allow execution of Cell CLI commands directly on a single or a group of cells. The Manage IO resource screen allows the administrator to set directives on how the database utilize the disks and flash cache if they see excessive I/O resource consumption by one particular database affecting the performance of other databases on the same set of storage cells.

The Oracle Exadata Storage Grid performance page can be used to check the pattern of load on hard disks and flash disks over a selected (15 minutes, 2 hours, day etc) period of time. You can also check the flash cache usage and the workload distribution by databases. These performance charts help you to identify at what times of day a particular database becomes really busy and what times usage is below average. These utilization characteristics help the database administrators to decide how IORM settings should be applied.

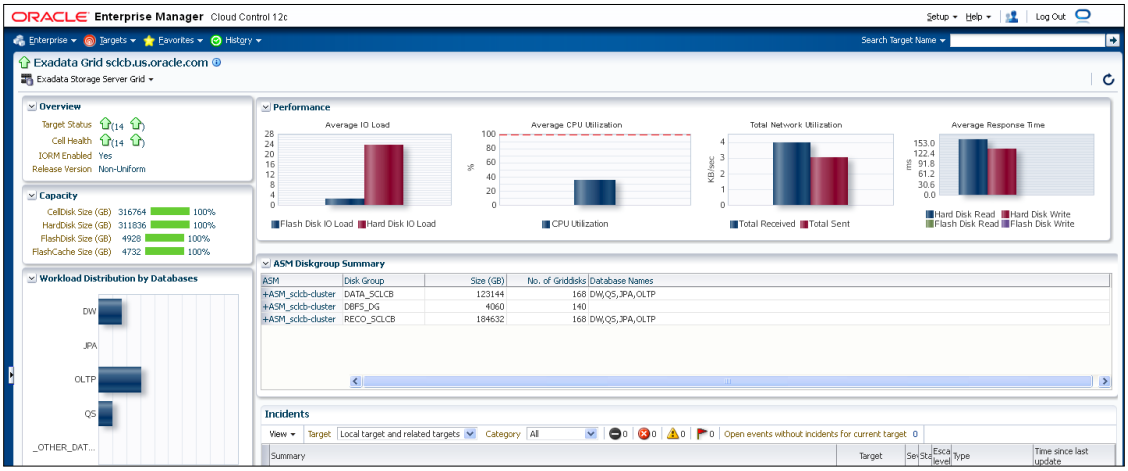

Figure 9: Oracle Exadata Storage Server Grid page

InfiniBand Network Management

InfiniBand switches provided with Oracle Exadata require checking for failed hardware components and sensors that have exceeded preset thresholds on the switch, and check for port errors that have occurred on switch ports. Oracle Enterprise Manager 12*c* uses a schematic display of the switches to display the status of ports. Each occupied normally operating ports are displayed as grey while empty or available ports are white. Ports that have encountered errors are yellow and the ones with degraded

performance are marked in red. The use of a color scheme to identify any problem greatly reduces time to investigate and fix any potential network problem. By moving the mouse on top of a port you can also identify the throughput of each port along with the peer port it is connected to. The InfiniBand home page also provides the throughput details of the ports of the Exadata components.

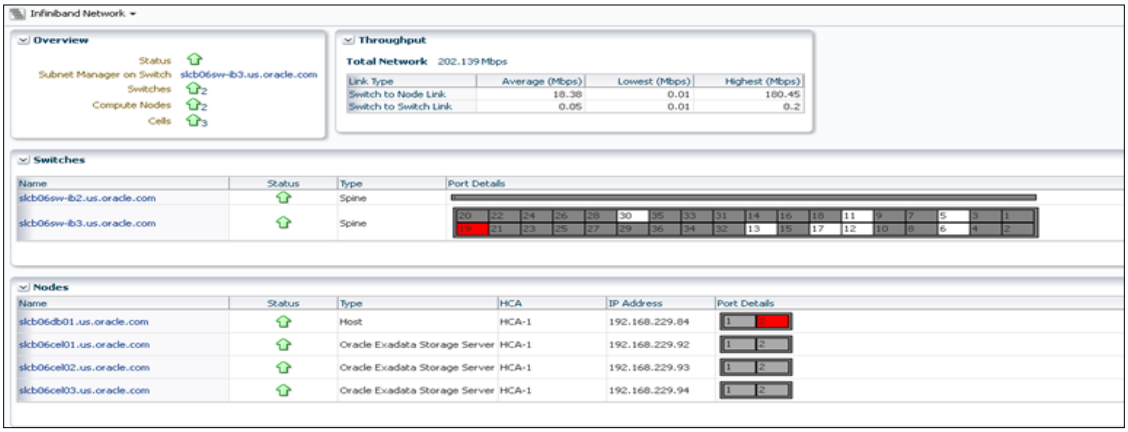

Figure 10: Infiniband Switch Homepage

Like the Oracle Exadata Storage Server page the Infiniband page also has a performance page for reviewing throughput across the Switch-to-Switch link, Switch-to-Node link and the Total Network. In addition to monitoring the performance of the InfiniBand network, administrators can also alter the port settings if Oracle Enterprise Manager detects port degradation.

The topology view of the Infiniband network also facilitates root-cause analysis of any network related issues. Whenever Oracle Enterprise Manager 12*c* identifies a network problem between the switch and one of the Nodes the network link turns red.

Cisco Ethernet Switch

The Cisco Ethernet does not participate in client connectivity to the database or database connectivity to the storage servers. Hence, it is not in the critical path. However monitoring and administrative traffic does depend on the availability of the Cisco Ethernet switch. The primary goal of monitoring the switch is to identify hardware component failure and environmental conditions that can lead to switch malfunction. The switch monitors availability and sensor thresholds for its hardware components. The metrics and default thresholds provided by the plug-in are sufficient for the level of monitoring necessary to ensure switch availability. Thresholds may be changed or set in the Metric and Policy Settings page in Enterprise Manager to handle site-specific requirements. Cisco Ethernet switch reports an alert condition as an SNMP trap, which is received by Oracle Enterprise Manager through the Oracle Exadata Plug-in.

Power Distribution Units

The Power Distribution Unit (PDU) metering unit enables you to monitor the electrical current being used by equipment connected in the Oracle Exadata rack. You can monitor the current in person by viewing the LCD screen on the PDU itself or remotely from a system on the network. To remotely

monitor the PDUs, connect each PDU to the management network. PDUs should be connected with an Ethernet network cable directly to a data center switch, not to the Cisco Ethernet switch on the rack. Oracle Enterprise Manager 12*c* requires the PDU firmware to be 1.0.4 or higher. The PDU monitoring thresholds depend on several factors (the size of the Oracle Exadata configuration, the power and voltage type). Determine the values for those factors and then update the PDU monitoring information in the plug-in metrics using the values and procedures documented in Support Note 1299851.1.

Keyboard, Video and Mouse

The keyboard, video and mouse (KVM) provides remote console access to Oracle Exadata Storage Servers and database servers in the rack. Its primary purpose is to provide keyboard, video, and mouse control for servers when performing administrative functions when physically present at the Oracle Exadata Database Machine in the data center. Oracle Enterprise Manager 12*c* can monitor the status of the KVM and the event occurrences like Factory Defaults Set, Fan Failure, Aggregated Target Device Status, Power Supply Failure, Power Supply Restored, Reboot Started, and Temperature out of Range on the KVM target.

The following table provides the details of monitoring of each component:

COMPONENT	CATEGORY	MONITORED BY
Oracle Exadata Storage Servers	Hardware	Monitored by ILOM and MS
	Operating System	Monitored by MS
	InfiniBand Server Ports	Monitored by MS
	Exadata Software	Monitored by MS
Database Servers	Hardware	Monitored by EM Agent
	Operating System	Monitored by EM Agent
	InfiniBand Server Ports	Monitored by EM Agent
	Oracle Grid Infrastructure	Monitored by EM Agent
	Oracle Database	Monitored by EM Agent
InfiniBand Network	Sun InfiniBand switches	Monitored by EM Agent
	InfiniBand Fabric	Monitored directly
Cisco Catalyst Ethernet Switch	All	Monitored by EM Agent
Sun Power Distribution Units	All	Monitored by EM Agent
Avocent MergePoint Unity KVM	All	Monitored by EM Agent

For Oracle Exadata best practices please refer to support note: 1110675.1

Monitor Many as One

Using Oracle Enterprise Manager's system and services framework, you can create a dashboard that provides a single pane of glass view of usage and performance metrics across all Oracle Exadata components. IT executives can access this dashboard without even logging into Oracle Enterprise Manager.

In Oracle Enterprise Manager 12*c* automation kit for Oracle Exadata a dashboard creation script has been integrated.

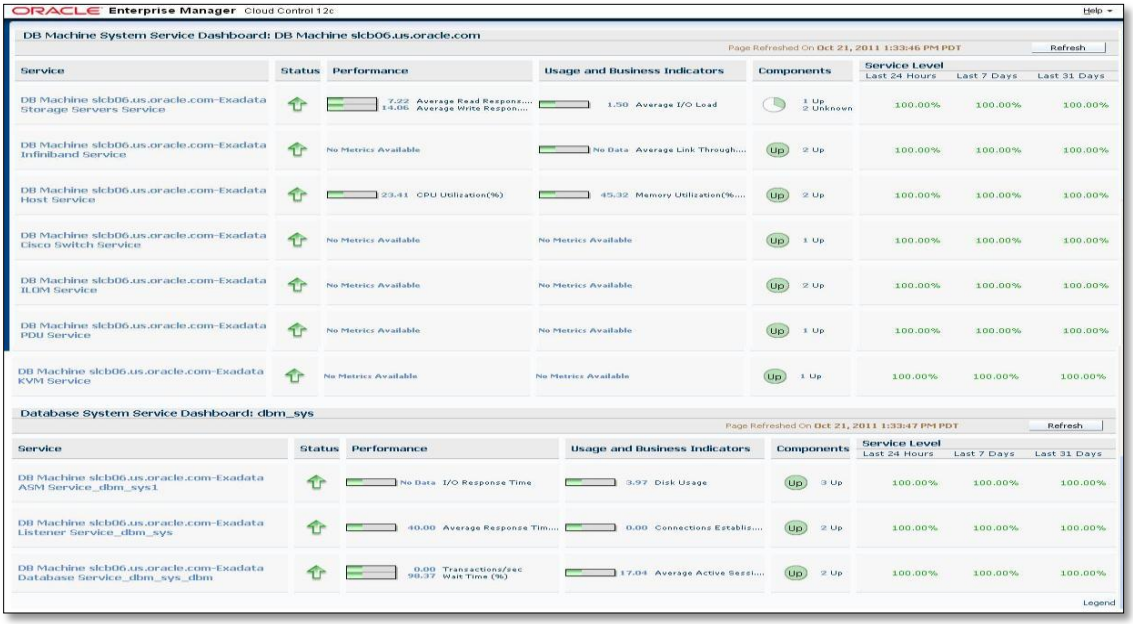

Figure 11: Oracle Exadata a reporting dashboard

Test

When you are migrating or upgrading your applications to Oracle Exadata there are several questions you want answered. These questions include whether Oracle Exadata will improve application response time and throughput—if yes, by how much? There are also changes associated with the migration as well. There may be an OS migration to Linux or Solaris, a possible database upgrade to Oracle Database11*g*R2 or a conversion from single instance to Oracle Real Application Clusters (RAC) database. It is very important to completely understand the risks and take appropriate action using proper testing to mitigate the risks before going live with your application on Oracle Exadata.

Using Oracle Enterprise Manager 12*c*'s Consolidation Planner capabilities and Oracle Real Application Testing, an Oracle Database 11*g* option, you can not only understand and mitigate your risks but also ensure a worry free application migration to Oracle Exadata.

Based on our past experience with several Oracle Exadata implementations we recommend using a 3-step testing process:

Identify

As enterprises increasingly look to consolidate their disparate databases onto Oracle Exadata infrastructure, administrators can use Consolidation Planner in Oracle Enterprise Manager 12*c* to identify optimal consolidation strategies for different Oracle Exadata configurations. Using the actual hardware configurations around CPU, Memory, I/O and the server workload history stored in Oracle Enterprise Manager, Consolidation Planner analyzes the workloads of the source systems and computes the expected utilization for the consolidation plan on the target Oracle Exadata system. Equipped with a rich library of hardware configurations, Consolidation Planner can guide administrators to define consolidation scenarios for even phantom Oracle Exadata servers, ranging from the different versions of X2-2 to X2-8. Now, businesses can make smarter and optimal decisions about the exact configurations of Oracle Exadata that is right for their database consolidation needs.

Each consolidation project contains one or more consolidation scenarios that are generated based on the inputs provided. Inputs provided to a scenario include:

- The source server resource requirements that a destination server must meet, including one or more of the following: CPU, memory, disk I/O, network I/O, and disk storage.
- Any business, compliance or technical constraints that must be considered.
- The destination servers to consider in the scenario.

A set of pre-configured consolidation scenarios are provided representing conservative, medium, and aggressive consolidation schemes. Each scenario is generated based on inputs you provide. Alternatively, you can create your own custom scenarios that best suit your situation. Once created, you can compare the various scenarios to determine which consolidation strategy best meets your requirements.

Each scenario also includes initial mappings between each source server and the destination server(s) it may be consolidated to. You can choose to create mappings manually, or allow Consolidation Planner to create them automatically. Once all inputs are specified, you can run the scenario and evaluate the results. You can change the values in a scenario, then re-run it to view the new results.

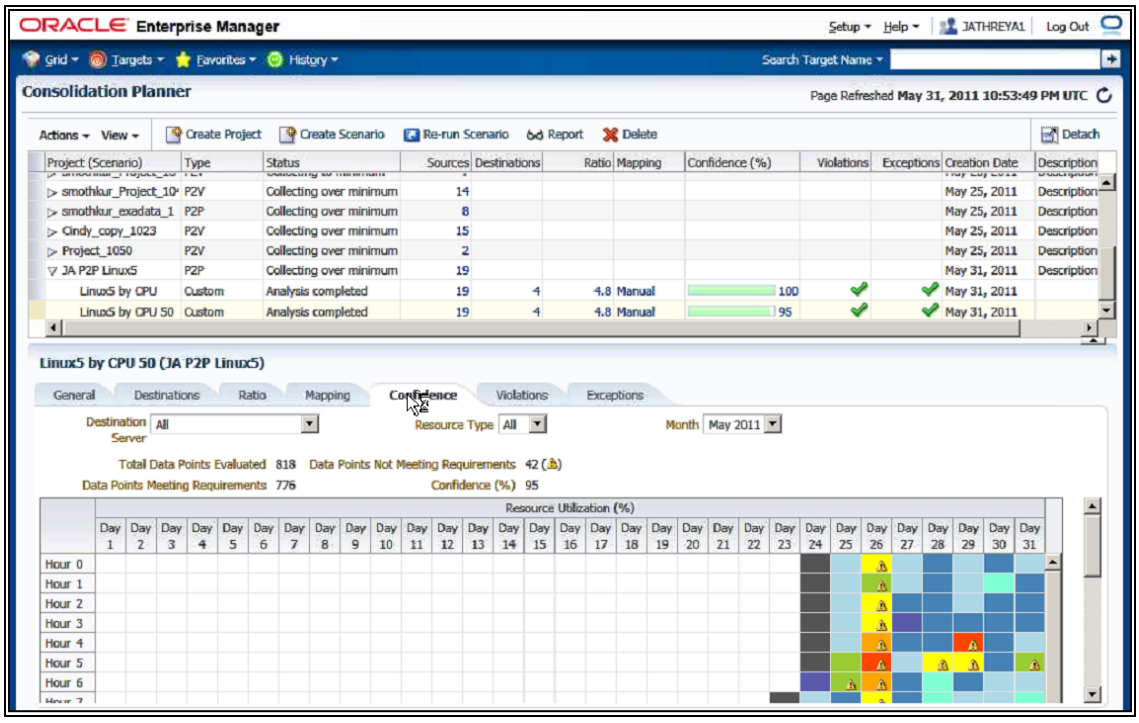

Figure 12: Oracle Enterprise Manager 12*c*'s Consolidation Planner Output

Creating Test Environments

Once you have identified the applications you would like to consolidate on to Oracle Exadata, the next step is to create a test environment. But all IT organizations need to protecting sensitive or personally identifiable aspects of the information stored in their production systems before replicating it in their test environments. As the number of applications increases, more and more data gets shared, thus further increases the risk of a data breach, where sensitive data gets exposed to unauthorized parties. Oracle Data Masking Pack, an Oracle Enterprise Manager 12*c* management pack, addresses this problem by replacing the original sensitive data with realistic-looking scrubbed data that has same type and characteristics as the original sensitive data thus enabling organizations to share fully masked data and still be in compliance with information security policies and government regulations.

Oracle has developed a comprehensive 4-step approach to implementing data masking called Find, Assess, Secure, and Test (FAST). You can use these steps to create a test environment on Oracle Exadata that has all the sensitive information masked before proceeding with the next step of validating the performance of the system.

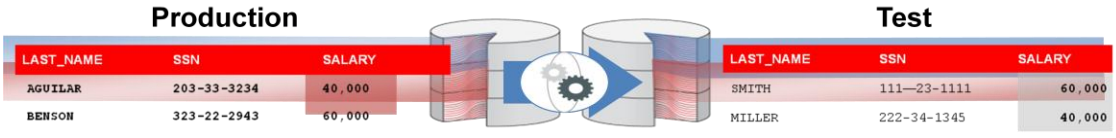

Figure 13: Oracle Data Masking Workflow

Validate Performance

Oracle Database 11*g*'s Real Application Testing option, introduces two new solutions, Database Replay and SQL Performance Analyzer. These two capabilities combine to make an ideally solution for validating production performance in test environments.

Database Replay

Database Replay provides DBAs and other system administrators with the ability to accurately capture real production workloads, including online user and batch workloads, and rerun them in a test environment. Database Replay enables administrators to realistically test system changes, including all concurrency, dependencies and timing, automatically in the test system without having to hand code or use customized testing scripts. This saves both time and effort. DBAs don't have to spend months getting a functional knowledge of the application and developing test scripts. With a few point and clicks, DBAs have a full production workload available at their fingertips to test and rollout any change necessary. This cuts down on testing cycles from many months to days or weeks and brings significant cost savings benefits to the businesses.

With Database Replay, DBAs and system administrators can test the following scenarios for a variety of Oracle Exadata implementations and use cases.

- Database upgrades, patches, parameter, and schema changes, etc.
- Configuration changes such as conversion from a single instance to RAC, ASM.
- Storage, network, and interconnect changes.
- Operating system patches, hardware migrations, upgrades, and parameter changes.

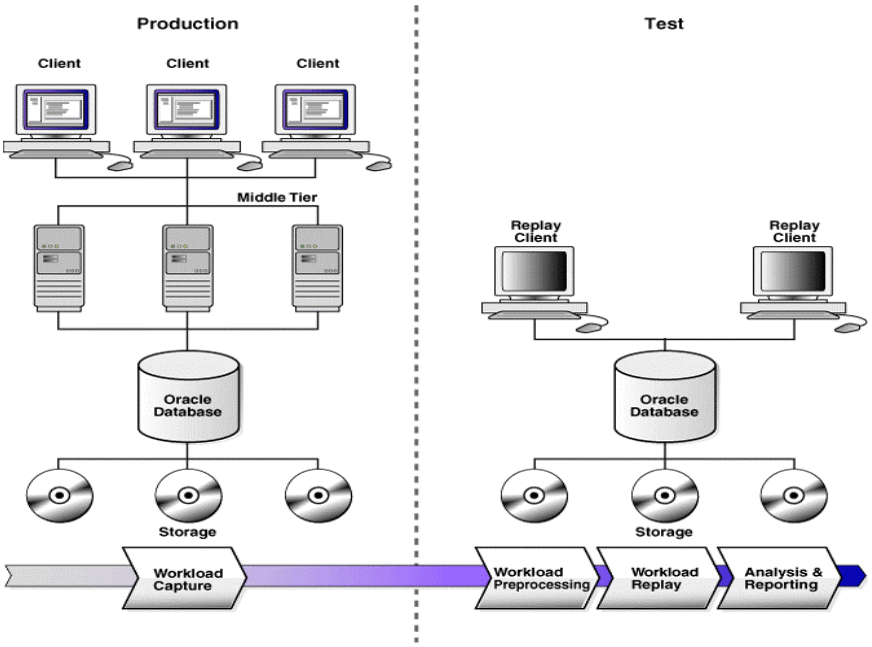

Figure 14: Database Replay workflow

Both the workload capture and replay process support a filtering capability that is useful for targeting workloads of interest, such as by service, action, or module. Oracle Enterprise Manager 12*c* significantly enhances the value of Oracle Real Application Testing by supporting end-to-end Database Replay automation on Oracle Exadata. This simplifies the process of saving and transferring the captured workload and performance data to the test system, setting up the test system and replay clients correctly. In addition, Oracle Enterprise Manager 12*c* orchestrates the entire replay workflow through its centralize Cloud Control console, making life easier for administrators.

SQL Performance Analyzer

Changes that affect SQL execution plans can severely impact application performance and availability. As a result, DBAs spend enormous amounts of time identifying and fixing SQL statements that have regressed due to the system changes. SQL Performance Analyzer (SPA) can predict and prevent SQL execution performance problems that may occur during a migration.

SPA provides a drill-down and granular view of the impact of environment changes on SQL execution plans, and statistics by running the SQL statements serially before and after the changes. SQL Performance Analyzer generates a detailed report outlining the net benefit on the workload due to the system change as well as the set of regressed SQL statements. For regressed SQL statements, appropriate executions plan are highlighted along with recommendations from Oracle Enterprise Manager 12*c* Tuning Advisor to optimize and tune the SQL statements for the best possible performance.

SQL Performance Analyzer is well integrated with existing SQL Tuning Set (STS), SQL Tuning Advisor and SQL Plan Management functionality. SQL Performance Analyzer completely automates and simplifies the manual and time-consuming process of assessing the impact of change on extremely large SQL workloads that may contain thousands of SQL statements. DBAs can use SQL Tuning Advisor to fix regressed SQL statements in their test environments and generate new plans. These plans are then seeded in SQL Plan Management baselines and exported back into production. Thus, using SQL Performance Analyzer, businesses can validate with a high degree of confidence that a system change to a production environment will in fact results in net positive improvements at a significantly lower cost.

SQL Performance Analyzer can test the following system changes associated with an Oracle Exadata migration:

- Database upgrade, patches, and initialization parameter changes
- Configuration changes to the operating system, hardware, or database
- Schema changes, such as; adding new indexes, partitioning or materialized views
- Gathering optimizer statistics
- SQL tuning actions, for example, creating SQL profiles

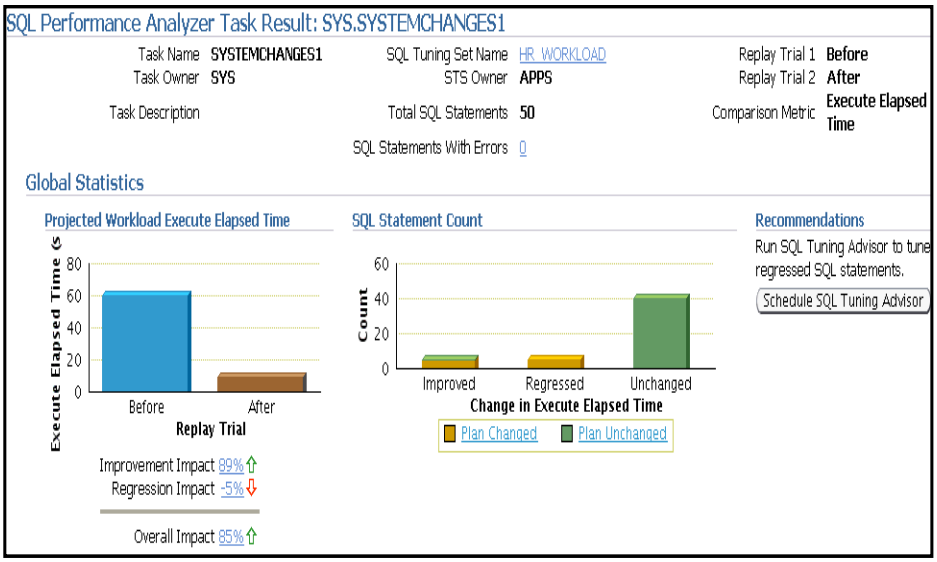

Figure 15: Oracle Real Application Testing's SQL Performance Analyzer Execution Report

The SPA comparison report in Figure 15 shows significant performance improvement of overall SQL workload after a proposed system change but with a few execution plan regressions. SQL Performance Analyzer takes into account the number of executions of a SQL statement when measuring its impact. A SQL statement that completes in seconds but is frequently executed may have a higher impact on the system than a long running statement executed only once. SPA takes these factors into account when analyzing and predicting overall performance improvements and regressions. If any regressions are encountered, SPA allows the user to fix them using SQL Tuning Advisor or with SQL Plan Baselines, a new plan stability feature first introduced in Oracle Database 11*g*.

Oracle Exadata Simulation

In Oracle Database 11*g* you can also use the SQL Performance Analyzer's special Oracle Exadata Simulation tool to simulate the behavior of SQL statements under Oracle Exadata. This is a powerful feature for those looking to take advantage of Oracle Exadata's extreme performance capabilities. While executing SQL queries Oracle Exadata directly filters out unnecessary data at the storage cell level and thus the amount of information handled by the Oracle RAC nodes is significantly reduced which in turn reduces interconnect traffic. In addition, since the Oracle RAC nodes have fewer buffers to process, they will consume less CPU as well, further reducing the elapsed time.

All this, of course, depends on how much filtering is done at the Oracle Exadata storage cell tier. This is what the simulation tool will help you determine. You will see which SQL statements that have improved in execution time and which ones that have not, if any at all. The effect on each SQL statement individually and cumulatively is shown clearly in the SPA report (Figure 15), which makes it very useful to see the potential impact when implementing Oracle Exadata.

The latest release of Oracle Enterprise Manager 12*c* seamlessly integrates both Oracle Real Application Testing and Oracle Data Masking to help maintain and preserve sensitive data for regulatory

compliance purposes. This enhancement not only helps mask sensitive information in the test databases but also masks sensitive information captured in the production workloads as well.

Manage

Once a single application or a number of consolidated applications go live on Oracle Exadata, the next important step is to ensure that optimal performance is achieved in production which meets the service level agreements. Oracle Enterprise Manager 12*c*'s integrated view of Oracle Exadata's hardware and software components allow DBAs to navigate from the database performance page to the Oracle Exadata system health page. For all databases running on Oracle Exadata, Oracle Enterprise Manager 12*c* adds two separate buttons, an Oracle Exadata System Health and Oracle Exadata System Performance under the I/O tab in the database performance page. When Oracle Exadata is running correctly without issues, the system health button is green but when Oracle Enterprise Manager detects an error, the system health button turns red indicating a problem may be impacting the database performance. Certain hardware problems may impact performance including:

- Load imbalance
- ASM related problems
- Cell software or hardware failures
- Cell configuration issues
- Network related failures

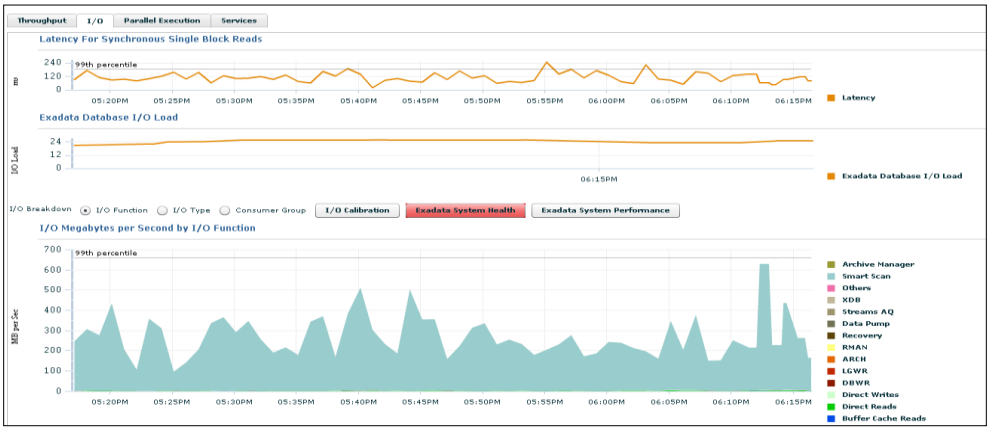

Figure 16: Oracle Exadata System Health Button in Oracle Enterprise Manager 12*c*

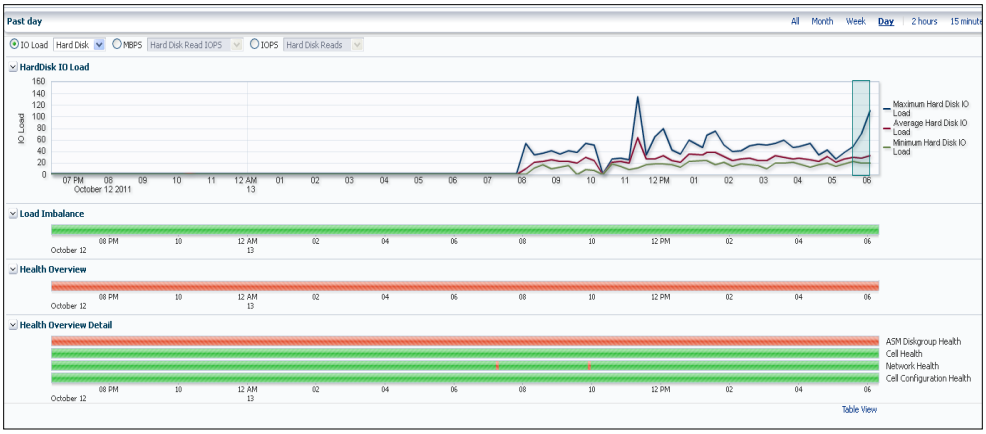

Figure 17: When a problem occurs, DBAs can drilldown in to the Oracle Exadata System Health page.

Database Performance Analysis

Oracle Database 10*g* first introduced Automatic Database Diagnostic Monitor (ADDM), which was a revolutionary feature that helped create the first self-managing database. ADDM uses an integrated approach to provide database-wide performance analysis, which covered storage, system resource, space, application & SQL and backup & recovery management. It provides proactive analysis to DBAs and is available on demand to troubleshoot problems related to performance.

Oracle Database 11*g*, running on Oracle Exadata, extends ADDM by offering cluster-wide performance analysis for Oracle RAC databases. For Oracle RAC environments ADDM analyses the RAC cluster and reports on issues that are affecting the entire database as well as its individual instances. DBAs can now use ADDM to perform database-wide analysis of global resources, such as high-load SQL, global cache interconnect traffic, network latency issues, skew in instance response times, I/O capacity, etc. DBAs also have the ability to restrict ADDM analysis on a few specified instances of an Oracle RAC cluster as well. With ADDM for Oracle RAC, performance analysis becomes as simple as a single instance database.

In Oracle Database 11*g*, ADDM findings can be suppressed by DBAs using directives to filter and display only findings of interest. To better understand the impact of the findings over time, each finding has a descriptive name that facilitates search, a link to the number of previous occurrences of the finding in last 24 hours, and affected instances.

Automatic Workload Repository (AWR) is a built-in repository within every Oracle Database running either on Oracle Exadata or on a non-Oracle Exadata platform that contains operational statistics about that particular database and other such information. AWR has built in support for Oracle Exadata specific wait events and the data captured allows both system level and user level performance analysis. AWR forms the foundation for all self-management functionality of Oracle Database. It is the source of information that gives the Oracle Database an historical perspective on how it is being used and enables it to make decisions, which are accurate and specifically tailored for the environment that system is operating in.

Oracle Enterprise Manager 12*c* includes a new tool to explore the Active Sessions History (ASH) data. ASH Analytic allows administrators to rollup, drilldown, and slice or dice performance data across various performance dimensions. With the ability to create filters on various dimensions, identifying performance issues has never been easier. The built-in Treemap view allows administrators to explore performance data using predefined performance dimension hierarchies. The ASH Analytics feature can be used on any database running on Oracle Exadata specifically using the Wait event dimension to zero into Oracle Exadata specific wait events. In Figure 19, you will find a drill down using the Service → Module → Action dimension hierarchy to the Wait event dimension revealing waits on cell smart table scan.

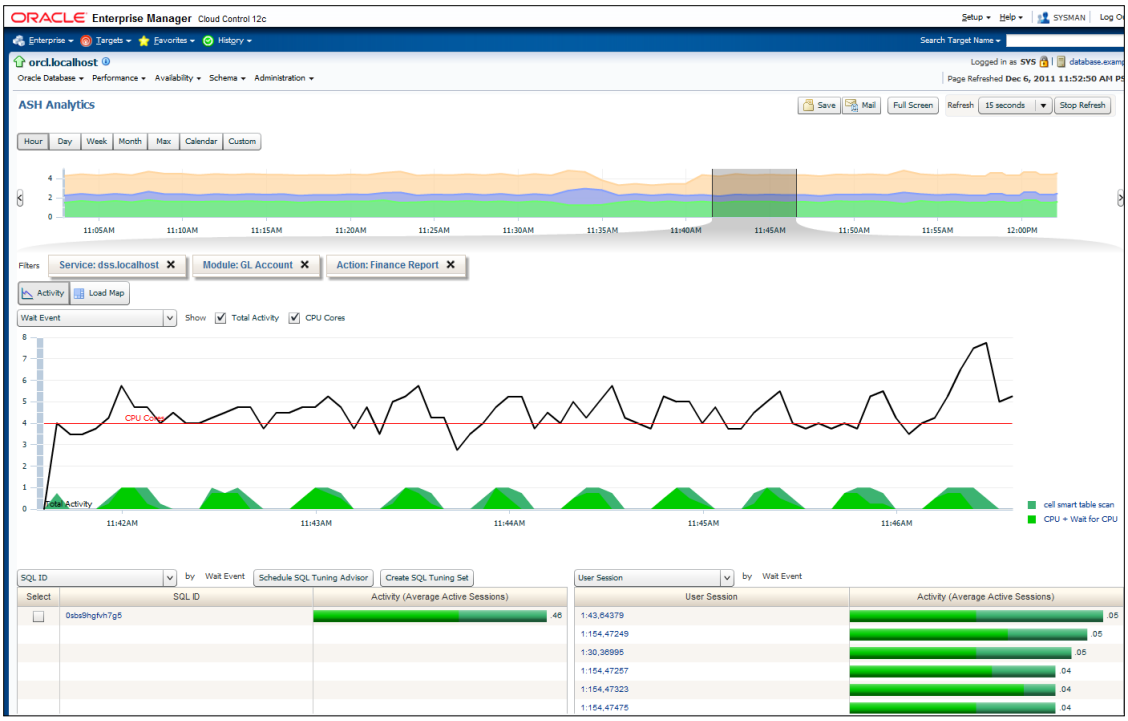

Figure 18: Active Session History (ASH) Analytics

Automatic SQL Tuning

In Oracle Database 11*g*, the SQL tuning process has been further enhanced and automated to keep databases running at their peak performance. The SQL Tuning Advisor now runs automatically during the system maintenance windows as a maintenance task. In each run, it auto-selects high-load SQL queries in the system, and generates recommendations on how to tune them. To validate the recommendation, SQL Tuning Advisor in Oracle Database 11*g* performs a test execute of the SQL statements with the new execution plan for which a SQL Profile is recommended. This dramatically increases the accuracy and reliability of SQL Profile recommendations.

Starting in Oracle Database 11*g* R2 the SQL Tuning Advisor also searches real-time and historical performance data for alternative execution plans for the statement. If plans other than the original exist, then SQL Tuning Advisor reports an alternative plan finding. The SQL Tuning Advisor may

recommend accepting a profile that uses the Automatic Degree of Parallelism (Auto DOP) feature. A parallel query profile is only recommended when the original plan is serial and when parallel execution can significantly reduce the elapsed time for a long-running query. The Automatic DOP recommendations are especially important for databases running on Oracle Exadata as running queries in parallel significantly enhances cell offloading via smart scans for full table reads.

Real-Time SQL Monitoring

Real-Time SQL Monitoring capabilities built into Oracle Database 11*g* enables monitoring of the SQL statement's performance while they are executing. Live execution plans of long running SQL are automatically displayed on the SQL Monitor page in Oracle Enterprise Manager 12*c* using new, fine-grained SQL statistics that are tracked out-of-the-box.

By default, SQL Monitoring is automatically started when a SQL statement runs in parallel which is quite often the case in a database running on Oracle Exadata, or when it has consumed at least 5 seconds of CPU or I/O time in a single execution. DBAs can observe the SQL statement step through the execution plan, displaying statistics for each step as it executes. Row source information at each step of the execution plan is tracked by means of key performance metrics, including elapsed time, CPU time, number of reads and writes, I/O waits and other wait times. SQL Monitoring gives the DBA information on what steps long running SQL are executing and allows DBAs to decide if additional tuning action is needed.

Real-Time SQL Monitoring has been enhanced in Oracle Database 11*g* Release 2 to support execution plans that are being executed in part by the Oracle Exadata. In Figure 19, you can see the cell offload efficiency (70%) along with other Oracle Exadata specific wait events indicating time spent in the storage layer. The cell offload efficiency is the ratio of data accessed by the storage cell to the data that is processed by the database. Larger values mean more offloading to the storage cell and results in better operational efficiency.All Exadata specific wait events are usually prefixed with the word cell.

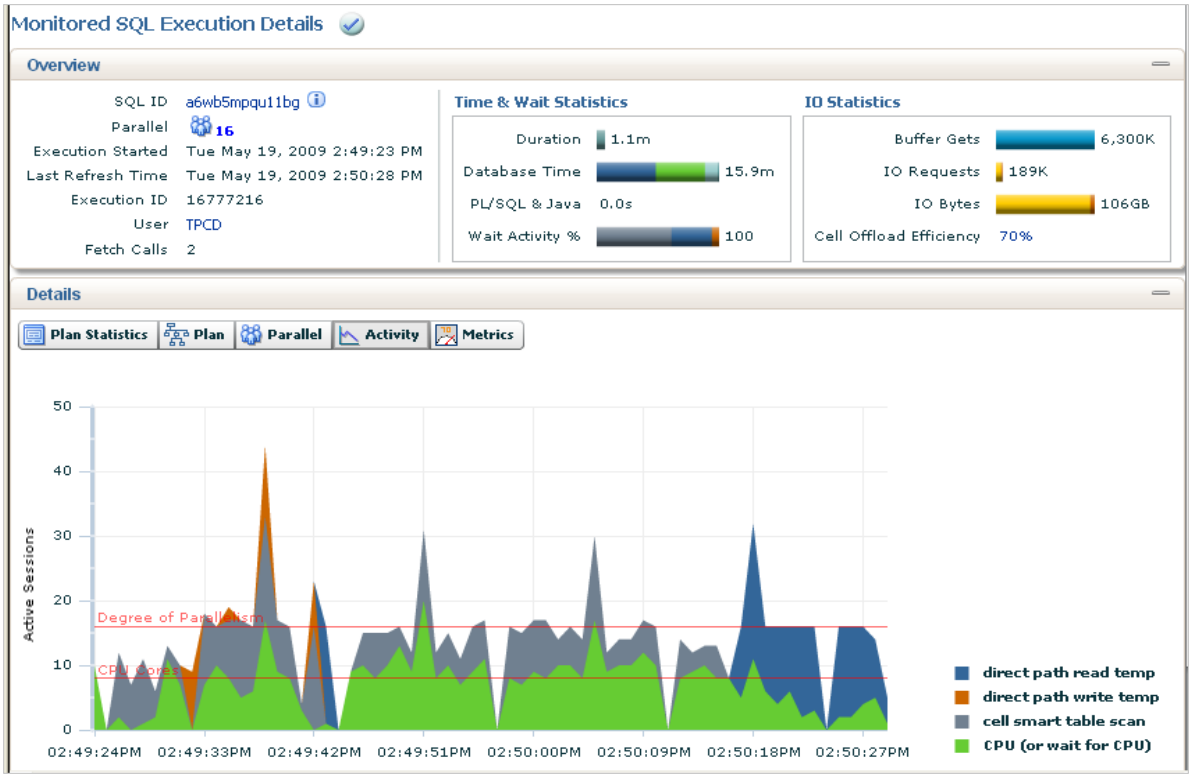

Figure 19: SQL Monitoring Report

Maintain

Once your application is live on Oracle Exadata along with ensuring optimal performance you also need to maintain configuration compliance across all components, apply all necessary patches and proactively resolve any issues that may stem either from the hardware or software.

Fault Diagnostics

Support Workbench

Oracle Database 11*g* includes an advanced fault diagnostic infrastructure for preventing, detecting, diagnosing, and resolving problems. The problems that are targeted in particular are critical errors that can affect the health of the database. When a critical error occurs, it is assigned an incident number, and diagnostic data for the error (i.e. traces, dumps, etc) are immediately captured and tagged with this number. The data is then stored in the Automatic Diagnostic Repository (ADR)—a file based repository outside the database—where it can later be retrieved by incident number and analyzed. Oracle Database 11*g* extensive fault diagnostics infrastructure provides the following benefits:

- Respond proactively to small problems and prevent catastrophic system failure by alerting DBAs using Health Checks.
- Limiting damage and interruptions after a problem is detected using Data Recovery and SQL Repair Advisor.

- Reducing problem diagnostic time through ADR and Test Case Builder.
- Simplifying customer interaction with Oracle Support using IPS and Oracle Configuration Support Manager.

Support Workbench is a facility in Oracle Enterprise Manager 12*c* that enables you to interact with the new fault diagnostic infrastructure of Oracle Database 11*g*. With it you can investigate, report, and where appropriate, repair problems, all with an easy-to-use graphical interface. The Support Workbench provides a self-service means for you to package diagnostic data using IPS, obtain a support request number, and upload the IPS package to Oracle Support with minimum effort, thereby reducing time-to-resolution for problems. With Oracle Enterprise Manager 12*c* all automated interactions with Oracle Support, such as support number creation or IPS package upload are done via the OCM API that is bundled with the agent.

Since Oracle Exadata runs Oracle Database 11*g*, correlated packaging and reporting of incidents from the database through Automatic Storage Management are all supported via the Support Workbench.

Incident Management

Oracle Enterprise Manager 12*c* greatly expands target monitoring and management capabilities beyond previous releases by letting DBAs focus on what is important from a broader perspective rather than focusing on discrete events that may or may not be relevant to a particular situation. By definition, an incident is an event or a set of closely correlated events that represent an observed issue requiring resolution through (manual or automated) immediate action or root-cause problem resolution. Incident Manager provides the ability to search, view, manage, and resolve incidents and problems impacting Oracle Exadata. Incident Manager can be used in the following situations:

- Respond and work on an incident
- Filter incidents, problems, and events
- Manage and automate incident workflow
- Suppress incidents and problems

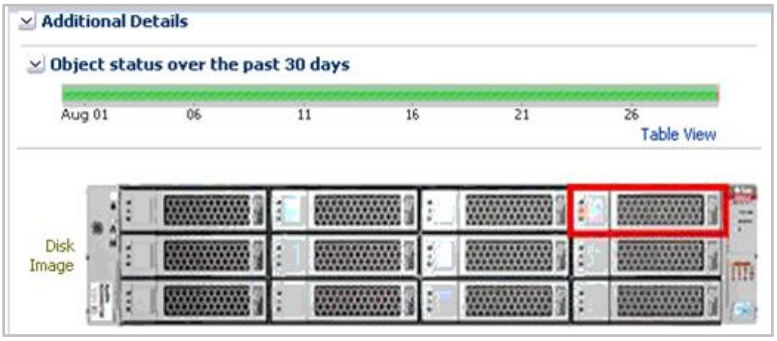

Figure 20: Hardware Failure Location

Oracle Exadata hardware failures come with a picture indicating the failed part location thus helping in immediate identification and remediation of problem.

Auto Service Requests

Auto Service Request (ASR) is a secure, scalable, customer-installable software feature and Oracle Support service that provides auto-case generation when common hardware component faults occur. ASR is designed to enable faster problem resolution by eliminating the need to initiate contact with Oracle Support for common hardware component failures, reducing the number of phone calls needed. ASR is applicable only for component faults. Not all component failures are covered, though the most common components (such as disk, fan, and power supplies) are covered, including:

- CPU
- Disk controllers
- Disks
- Flash Cards
- Flash modules
- InfiniBand
- Cards
- Memory
- System Board
- Power supplies
- Fans

ASR also simplifies support operations by using electronic diagnostic data. Easily installed and deployed, ASR is completely controlled by customers to ensure security. The following diagram illustrates the ASR workflow:

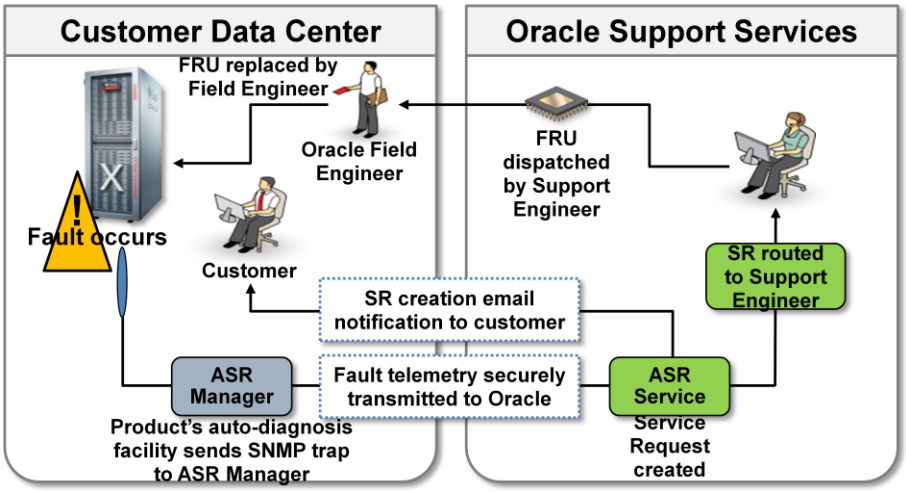

Figure 21: Auto Service Request Workflow

Configuration Compliance

Using Oracle Enterprise Manager 12*c*'s configuration management capabilities to help automate the discovery and monitoring of Oracle Exadata is a very powerful feature that helps DBAs detect configuration drift—which occurs when a target or component's standard configuration differs or "drifts" from a pre-defined gold standard or baseline between environments such as production and development and test which may affect the availability and performance of the environment. A very common example is the mismatch in IORM settings across cells shared by the same group of databases or even database parameter mismatch across instances in the same database.

Using Oracle Enterprise Manager 12*c*, DBAs can enforces compliance to many industry and regulatory standards using automated out-of-box policies and features such as, search, compare and change history tracking. The configuration management inventory capability is another way Oracle helps DBAs streamline managing Oracle Exadata and includes version summary of Exadata's underlying components.

In addition to these capabilities, Oracle Enterprise Manager 12*c* configuration management also provides:

- Continuous and real time event capture for changes made to Oracle Exadata
- The ability to perform reconciliation with Change Management Systems
- Comprehensive reporting of changes: who made the change, what change was made, when it was made, and where the change occurred.

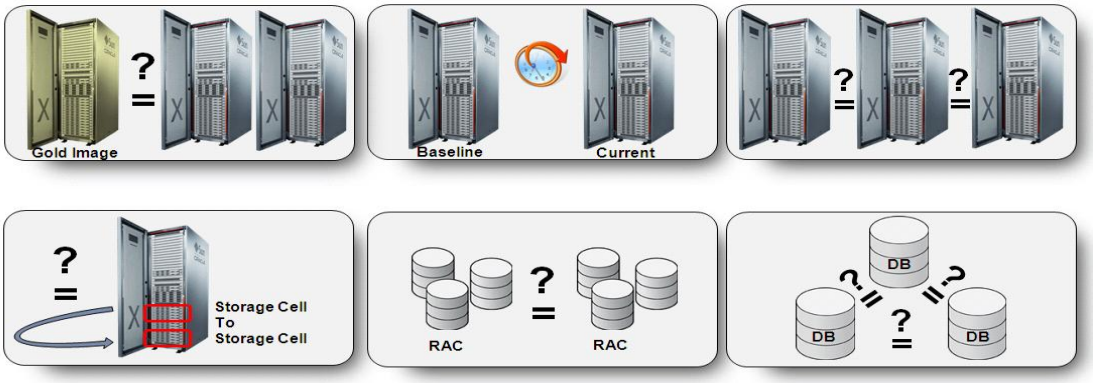

Figure 22: Different Types of Configuration Comparison

Using the built-in configuration comparison wizard you can compare the following Oracle Exadata components:

- Gold Image vs. Current
- Baseline Vs. Current
- Across multiple database machines
- Storage Cell to another Storage Cell or between two InfiniBand Switches
- Cluster Configurations
- Database Configurations

The configuration comparison wizard is used to submit the comparison job that uses pre-defined templates to compare the entire database machines or an individual component of Oracle Exadata. You can also setup rules to ignore (e.g. Storage Cell Name) certain attributes that you know will always be different. Oracle Enterprise Manager 12*c* also provides out-of-the-box compliance standards for Oracle Exadata to check whether any component's configuration has a violation.

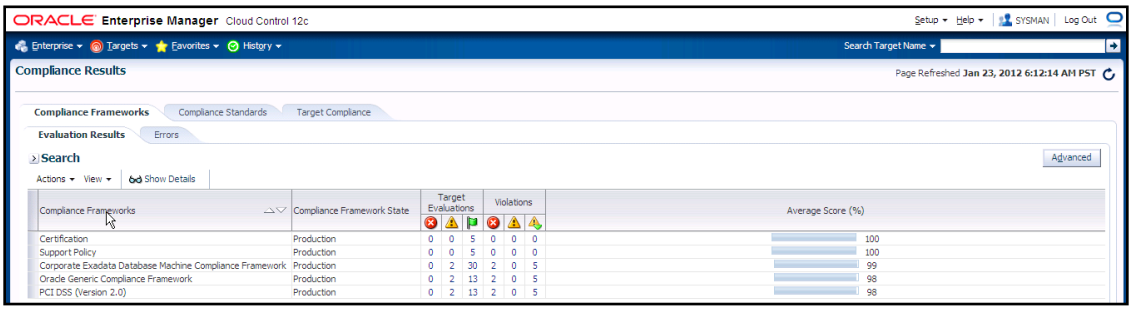

Figure 23: Oracle Exadata Configuration Search and Compliance Feature

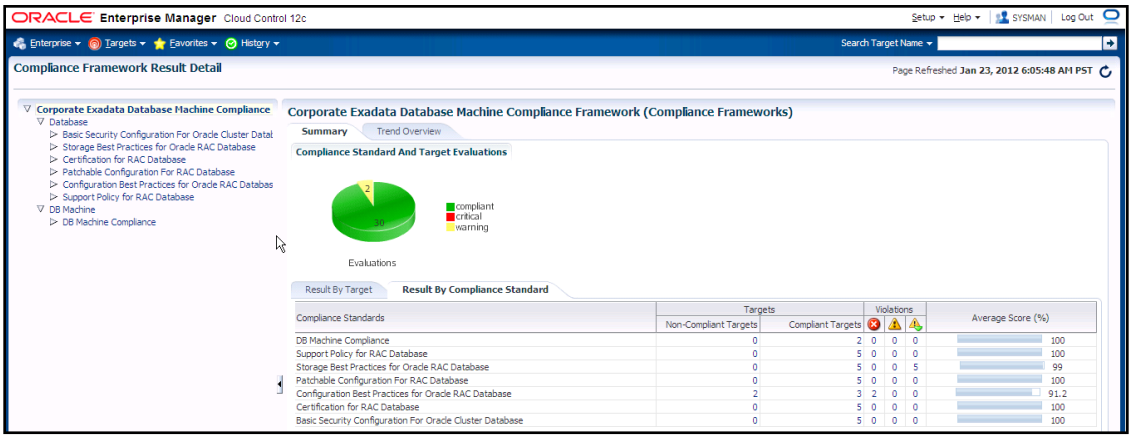

Figure 24: Output from Oracle Exadata's Compliance Framework

Patching Automation

Patching can be a very time consuming and error prone activity that often takes a lot of the database administrator's time and causes disruption to usual business activity. It is very important to make sure all recommended patches are applied on a timely manner otherwise applications may be missing critical application and security updates. Some of the typical software patching challenges include:

- Identifying which patches are available
- Identifying which patches are applicable to individual applications
- Identifying which targets need to have patches applied
- Scheduling the best time to apply patches

It is not uncommon for organizations to choose to delay their patching cycles and opt to apply multiple patches at half yearly intervals or apply patches during a single maintenance window. This can be a major undertaking when you consider there are hundreds of servers spread across testing and production to patch. This is an even a greater challenge when manual scripts and/or command line tools are used as part of the patching process.

Oracle Enterprise Manager 12*c* helps solve this problem by automating the deployment of database infrastructure patches to your Oracle Exadata. The patching application automates the deployment of Oracle patches for the database infrastructure including, Clusterware, Grid Infrastructure, Oracle Real Application Clusters, and Automatic Storage Management.

Oracle Enterprise Manager 12*c* Database Lifecycle Management Pack supports the entire patch management process and is integrated with My Oracle Support to provide a synchronized view of available and recommended patches. The Pack includes key capabilities such as, patch advisories, pre-deployment analysis, rollout and reporting. These recommended patches can then be analyzed for conflicts before deployment. DBAs can then apply multiple patches to multiple databases in a single downtime window, saving both time and effort.

Deployment Procedures

Deployment Procedures for patching are designed to enable maximum ease and minimum downtime. Oracle Enterprise Manager 12*c* introduces patching for Oracle Real Application Clusters on Oracle Exadata through 'Out of Place' method. In this method, clones of the GI and Oracle homes are created, and the patches are applied to the cloned home instead of the original home. Once the cloned home is patched, user can selectively switch the GI and instances to run from the cloned home. Using this method provides the following benefits:

- Minimal downtime is incurred only during the switch over to other instances.
- Flexibility, each instance can be selectively migrated based on maintenance schedules.
- Recoverability, in case of any issues the databases can be switched back to the old Oracle Home.

Oracle Enterprise Manager 12*c* provides complete automation to patch the cluster in zero-downtime using both in In-Place and Out of Place methods. Oracle recommends that customers apply these patches Out of Place.

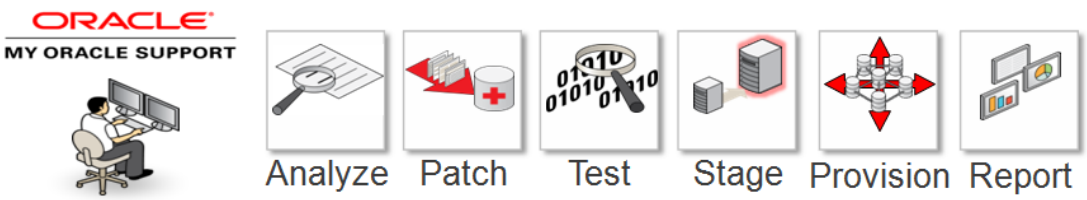

Figure 25: Patch workflow

Apart from the database infrastructure DBAs will also need to patch other Oracle Exadata components such as, the Oracle Exadata Storage Server. An Oracle Exadata Storage Server patch typically contains updates to the firmware, operating system, and/or Oracle Exadata Storage Server software. Updates to Oracle Exadata Storage Server must occur only with an Oracle Exadata Storage Server patch provided as a single downloadable patch from My Oracle Support (MOS). It is not recommended to manually update firmware or software on storage servers. An Oracle Exadata Storage Server patch may contain updates to the firmware, operating system and database servers. Oracle Exadata Storage Server patches are typically installed using a script supplied with the patch called 'patchmgr'. They may be installed in one of two ways: rolling and non-rolling fashion. Similarly an InfiniBand switch patch contains updates to the software and/or firmware for InfiniBand switches. Updates to InfiniBand switch software and/or firmware must occur only with the patch downloaded from MOS. Prerequisites and instructions for installing a patch are provided in a README supplied with the patch. MOS article 888828.1 and Note 1262380.1 provide an excellent overview of patching and testing guidelines for all Oracle Exadata patches.

Ops Center

Oracle Enterprise Manager Ops Center is the most comprehensive management solution for Sun hardware infrastructure, and supports heterogeneous environments as well. Offering a single console to manage multiple server architectures and multiple operating systems, Oracle Enterprise Manager Ops Center capabilities include, asset discovery, provisioning of firmware and operating systems, automated patch management, patch and configuration management, virtualization management, and comprehensive compliance reporting. An open, extensible system that can be integrated with existing data center management tools, Oracle Enterprise Manager Ops Center manages across the entire infrastructure stack—from firmware, operating systems and virtual machines, to servers, storage, and network fabrics. If you are already using Ops Center, you can choose to use Ops Center for all your OS Patch management on Oracle Exadata compute nodes. With Ops Center you can choose to create a central repository of ISO images, track and audit change history, patch and configuration compliance. Ops Center comes with a unique knowledge base that contains metadata on consistent, supported configurations for Oracle Solaris and Oracle Linux operating systems.

Conclusion

Oracle Exadata Database Machine not only offers the fastest time to value with the lowest possible risks—but ensures extreme and predictable performance. Oracle Enterprise Manager 12*c* maximizes Oracle Exadata performance with automated database diagnostics, tuning and self-management. The automated change and configuration management capabilities built into Oracle Enterprise Manager 12*c* help customers drive down IT operation costs while enforcing compliance with industry and regulatory standards. Oracle Enterprise Manager 12*c* automates testing of patches, changes and upgrades while keeping the data secure—helping organizations unlock the full value of their Oracle Exadata investment.

Oracle Enterprise Manager 12*c*
Oracle Exadata
June 2012
Author: Deba Chatterjee
Contributing Authors:
Scott McNeil , Pete Sherman, Andrew Bulloch

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com/enterprisemanager

Oracle is committed to developing practices and products that help protect the environment

Hardware and Software, Engineered to Work Together