Oracle Maximum
Availability Architecture

Exadata Health and Resource Utilization Monitoring

System Baselining for Faster Problem Resolution

ORACLE WHITE PAPER | JANUARY 2017

ORACLE

Introduction

One of the most common causes of performance degradation is change - change in application logic and code, change in software and hardware/infrastructure configuration; even change in application usage patterns. Identifying that change, tracking it and correlating it with changes in application performance is a key technique in deterministically solving performance degradation.

To understand change, a solid baseline foundation is key and Oracle Enterprise Manager Cloud Control (Enterprise Manager) has a suite of tools for capturing current baseline information and using those captured metrics to quickly solve system performance problems. In general, the process records and validates system configuration data at all levels of the Exadata stack, validates databases schemas, and verifies application performance and resource usage is within defined ranges.

Configuration Management

An Exadata Machine encompasses many Enterprise Manager targets at multiple levels of the hardware and software stack, and when investigating performance problems within this stack, quickly identifying configuration changes can be a big time saver, especially when those changes may have happened without an administrator's knowledge. Enterprise Manager's Configuration Comparison and Drift Management tools can help to set up saved configurations and baselines for any target of the Exadata machine, from top level software down to the lowest hardware component.

With Configuration and Drift Management administrators can troubleshoot or investigate problems where a change is suspected, ensure target configurations remain the same as reference or saved targets, and ensure targets of a similar target type within a system remain the same.

Configuration Management allows administrators to run manual one-time comparisons, run automated comparisons with email notification of changes, and save configurations for later use. Saved configurations can be useful to capture a "Gold Copy" of a full system stack, which would represent the standard configuration for a company's architecture. When new systems come online, these new systems can be compared against the Gold Copy to ensure compliance with company standards. It is important to note that as systems change and evolve, such as with software or hardware upgrades or tech refreshes, new baselines should be captured to represent current proper configuration.

The basic Enterprise Manager setup captures a lot of configuration data to cover every target in the Exadata software and hardware stack, but it's also very extendable and customizable for customer environments. Configuration Extensions allow this customization through SQL queries, file parsing and custom host executed commands. If you need a tailored environment, Configuration Extensions will further extend Enterprise Manager's comprehensive functionality.

When running target comparisons, Templates define which configuration items to include or ignore. A template is associated with a specific target type, which determines the properties to be compared and also define constraints of acceptable values for specific properties. Included with Enterprise Manager are a set of comprehensive templates that can be copied and customized as needed to filter out any unwanted noise from known or expected differences in your architecture. A set of custom Exadata templates is also available through the following support document. These templates can also be used to compare Exadata targets across different Exadata machine versions (example: comparing X3-2 with X4-2)

- Enterprise Manager Configuration Comparison Templates for Exadata (Doc ID 2181282.1)

Every target in Enterprise Manager has a configuration, which is tracked by default with Drift Management. Administrators can use this Drift Management data to see what changes were made to a system over a defined window of time. When there is a need to be more focused on a particular aspect of your system, Saved Configurations allow administrators to be more specific with comparisons in conjunction with Comparison Templates. The following is an example of a saved configuration for a database instance, which among other items, will help keep track of changes to database initialization parameters.

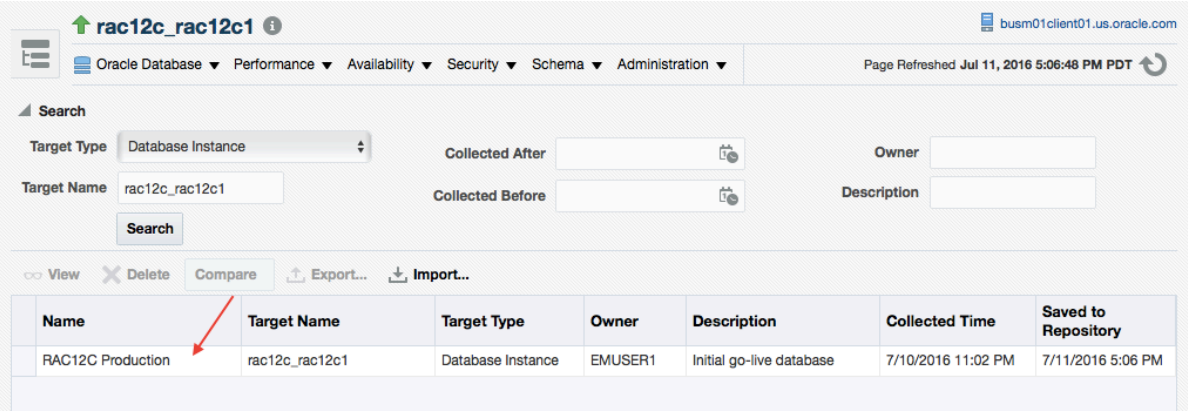

Figure 1

Viewing a saved configuration presents a tree list of configuration categories. Selecting a category will display the individual category properties. In the example below, System Global Area properties are selected. A saved configuration can also be exported to an Excel spreadsheet for additional analysis using the “Configuration Report” button for a full report, or the “Export” button for the properties of the selected category only.

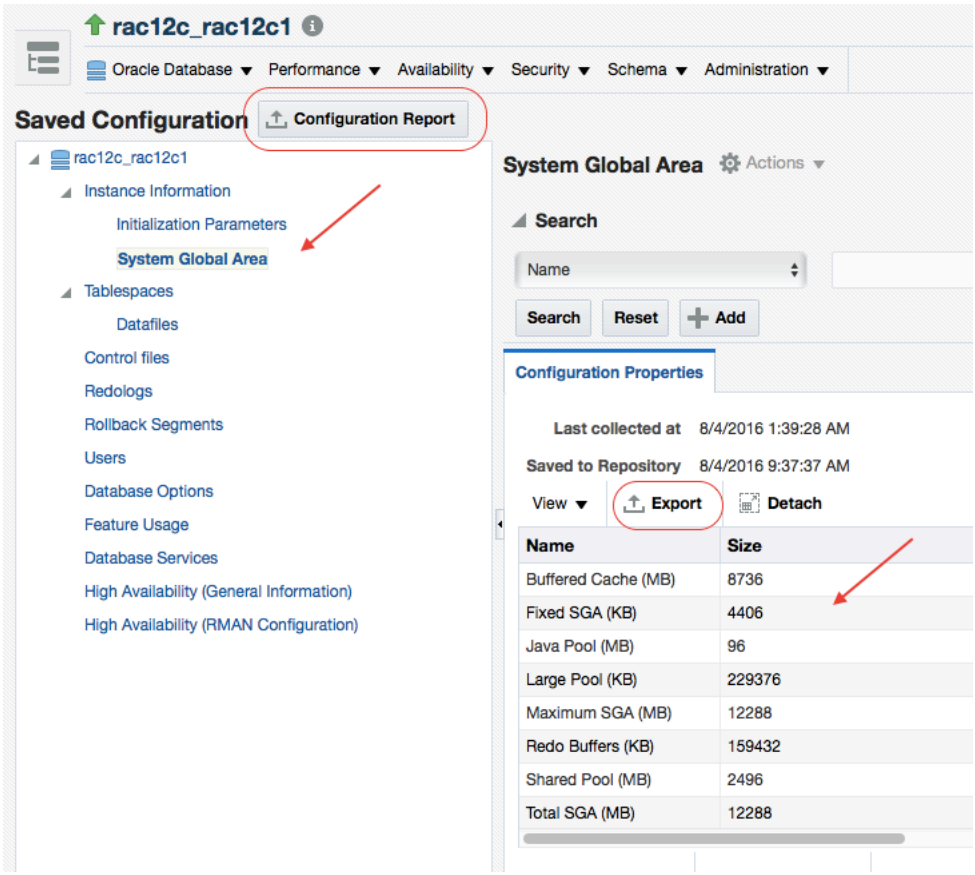

Figure 2

The Configuration History screens in Enterprise Manager can be used to actively query for any configuration changes that may have occurred in the past, track those changes over time, and configure email notification for any changes that have occurred. The example below tracks Database Instance changes which may have happened in the past 24 hours. Clicking on the “Schedule and Notify” button brings up a dialog box to schedule the running of this check and define email recipients.

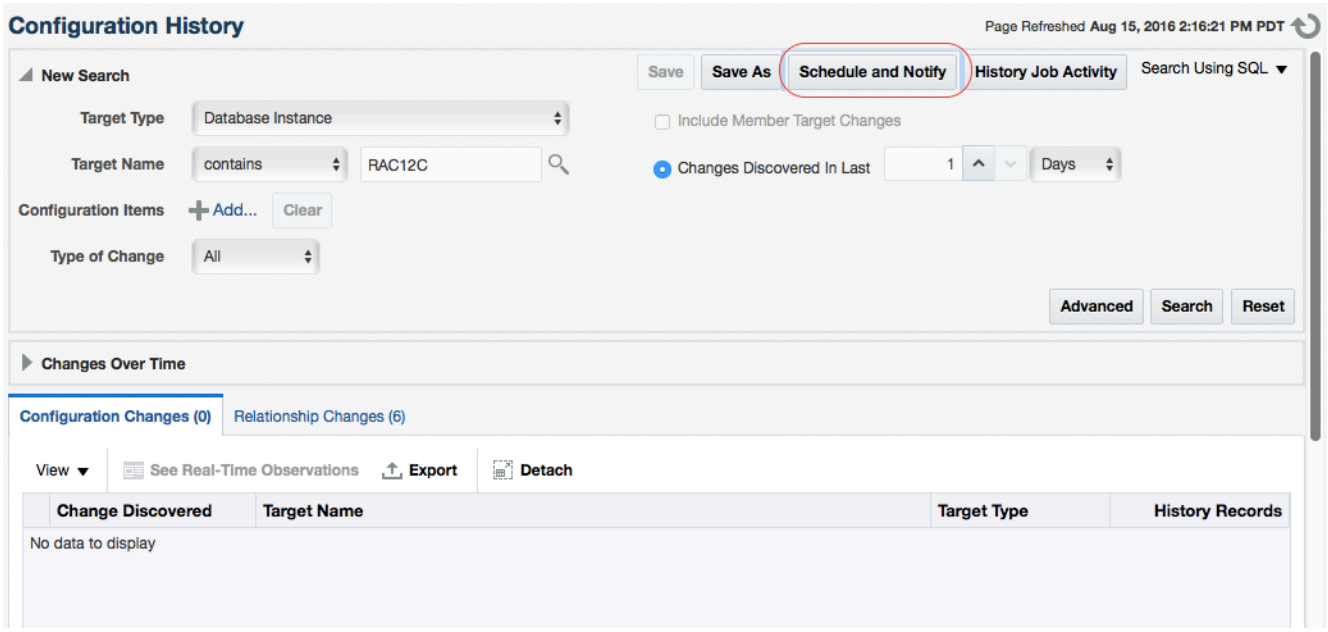

Figure 3

The email notification will give a summary of changes detected with a link to the Configuration History for the defined query.

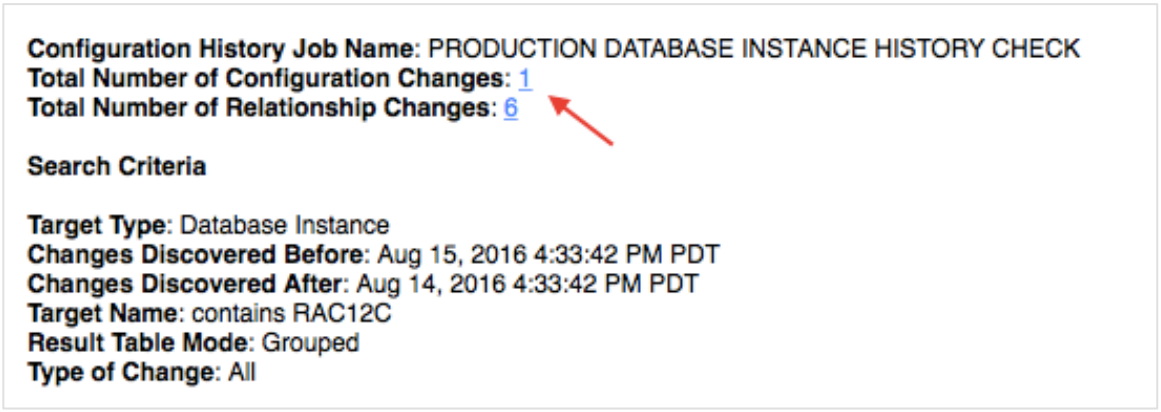

Figure 4

The History screen provides the ability to drill down to identify the change that was made. In this example a database initialization parameter was altered.

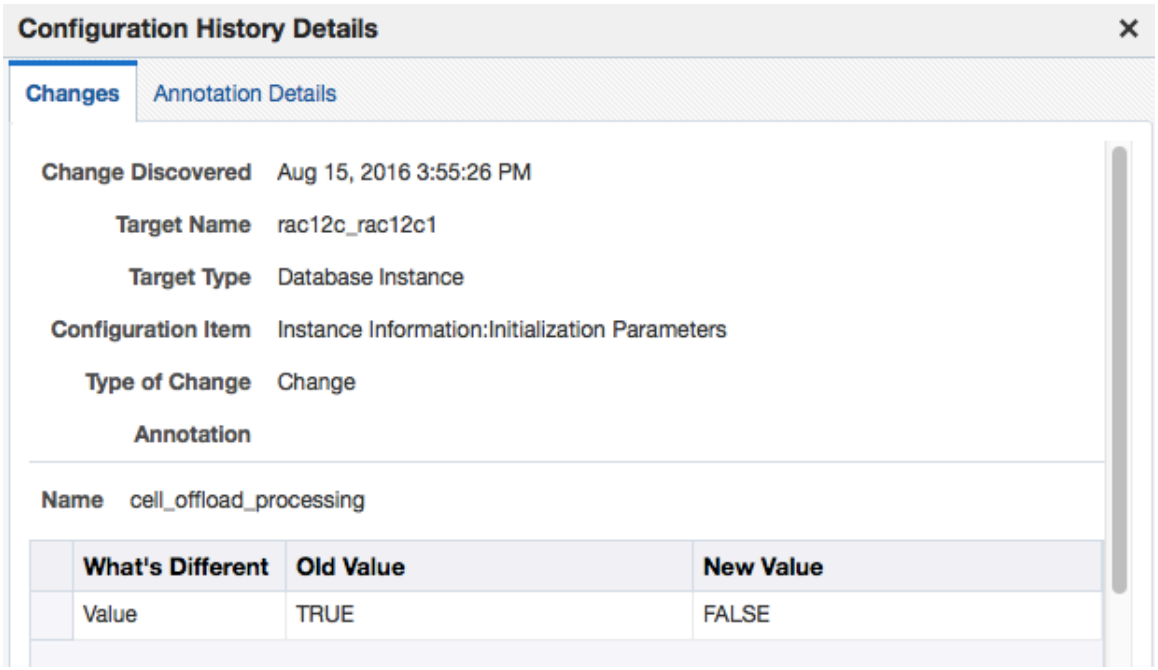

Figure 5

In summary, the Configuration Management tools within Oracle Enterprise Manager provide the means to track, record, notify and analyze changes for any target across the Enterprise.

Schema Baselines

Saving database schema baselines as an application evolves, from initial go-live to the current release version, is a great tool to track and report on schema changes over time. When comparing schemas, there are no Comparison Templates to use. Instead, the schema objects to compare are selected when creating the baseline. These schema objects will be used when running the compare. Items like indexes, packages and triggers, and also some supporting objects like profiles and roles can be included. As new versions of an application are released and it's associated schema evolves, additional baselines should be captured to represent and keep up with the latest version of the application. Eventually you will have a documented history of schema changes, which can be used for other reporting purposes as well, like security and regulatory audits. To accommodate applications, which span schemas, multiple schemas can be included in a single baseline capture.

The following example creates a baseline for the schema BENCH1.

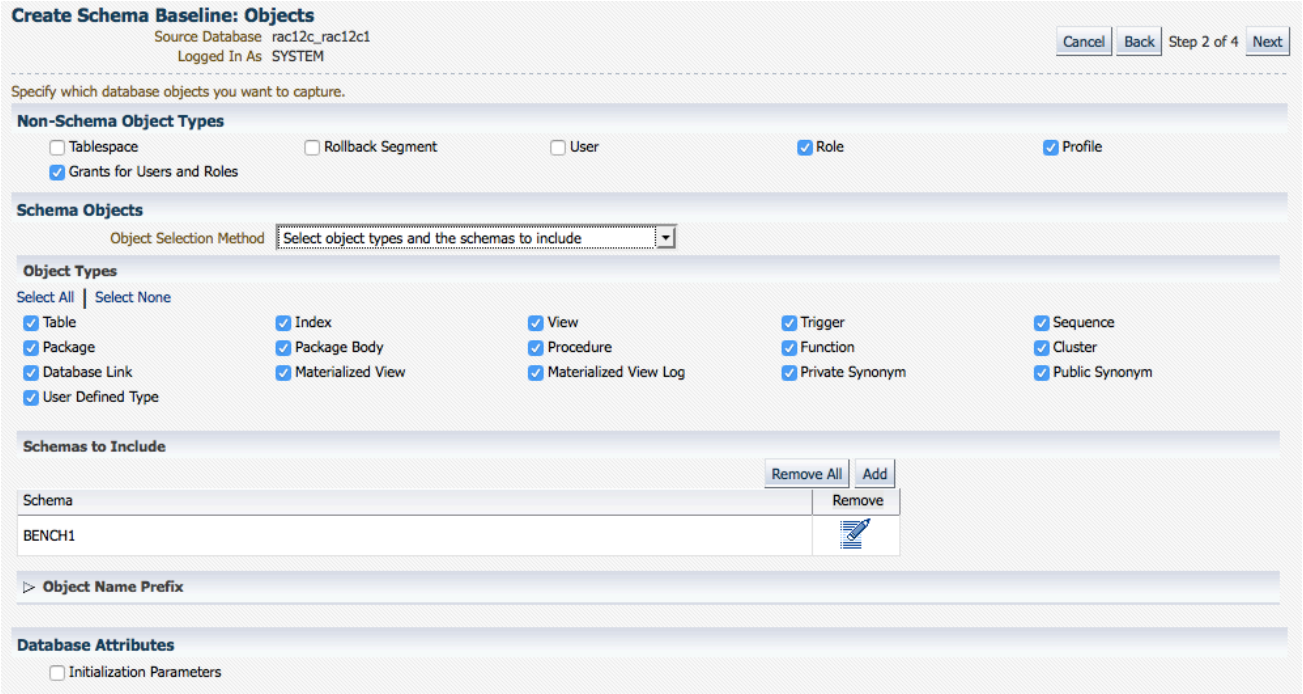

Figure 6

Workload Baselines

In Enterprise Manager, AWR Baselines can capture database resource usage for a defined time window, using static one-time baselines, a regularly scheduled repeating baseline or a rolling window. Using all three methods can provide additional options when comparing AWR reports.

Understanding application behavior will help in deciding which baseline methods will be most effective. The start and end times of the AWR Baselines will depend on this behavior. It may be fine to have a small baseline window of 10 minutes for consistent steady applications, but application behavior that might last for hours, like an extract, transform, and load (ETL) job, would require a larger window. When there are longer running regularly scheduled jobs with defined start times, capturing a static baseline of normal behavior and then setting up repeating scheduled baselines is a great way to look at past performance in comparison to recent activity. This can be useful for database loads which start late in the evening or early morning when administrators may not be regularly available.

As an example, for a steady workload that rises and falls depending on the number of users running queries, a baseline lasting 30 minutes would be sufficient to capture a typical load.

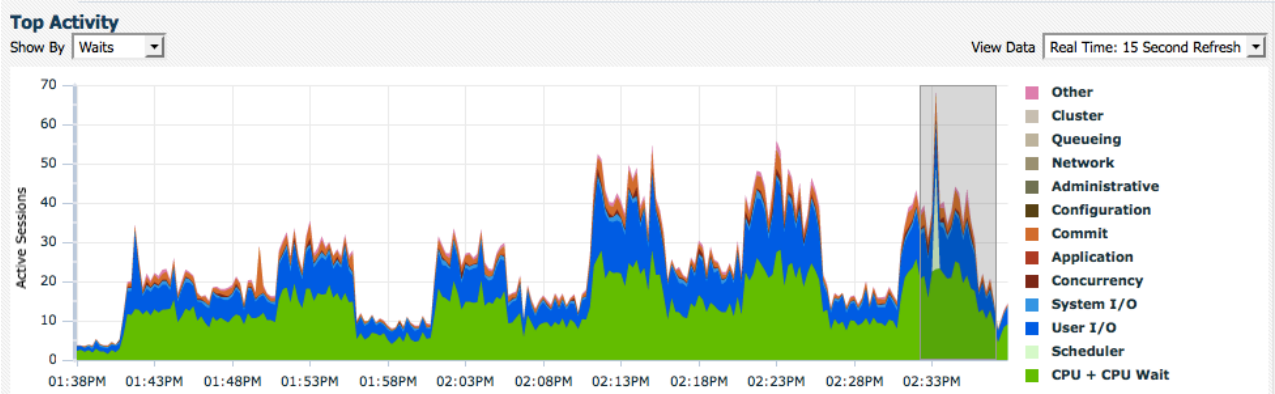

Figure 7

To create AWR baselines, navigate to the AWR Administration screen and click on the baseline link.

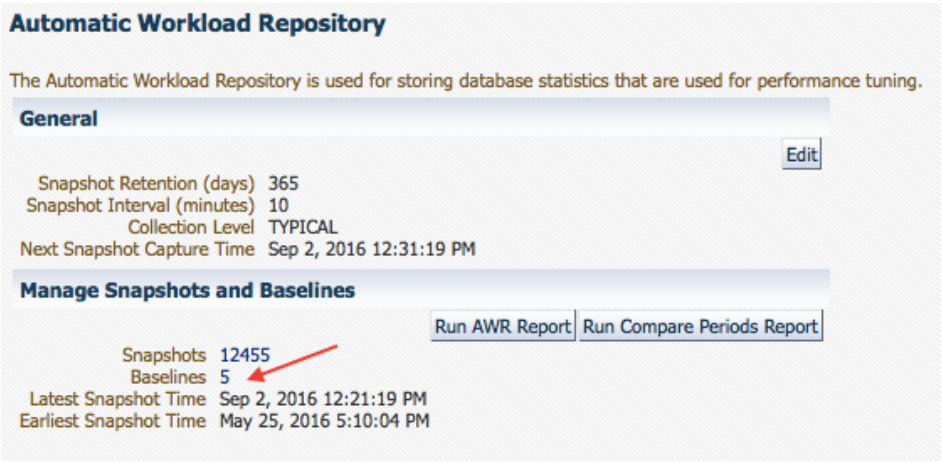

Figure 8

The example in the following figures consists of three rolling baselines which track nightly ETL jobs, a one-time baseline that captures 30 minutes of a typical day, and a daily rolling window, which is a default baseline provided by Enterprise Manager. These various baselines are defined by clicking on the Related Link at the bottom of the screen.

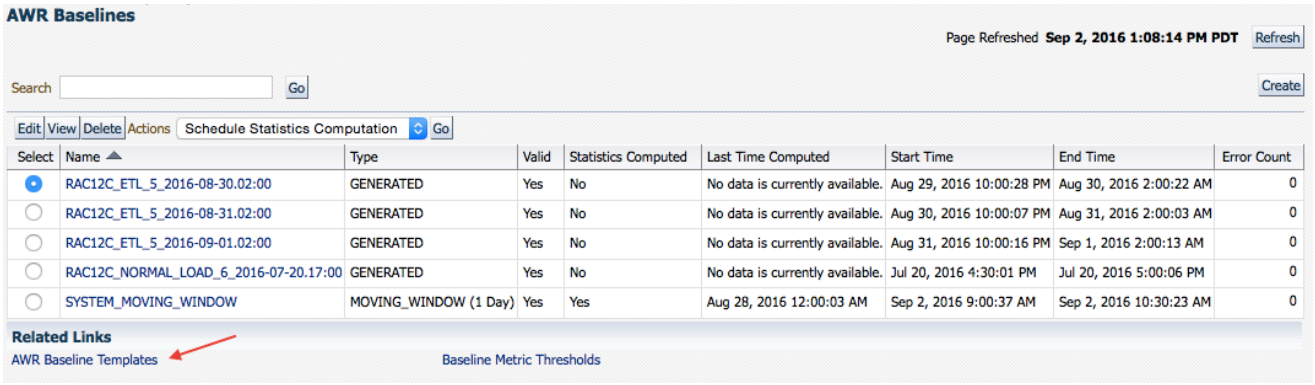

Figure 9

The following example shows a repeating baseline template that captures an ETL job and a one-off static baseline template that is scheduled to capture a typical day's load over a 30-minute window.

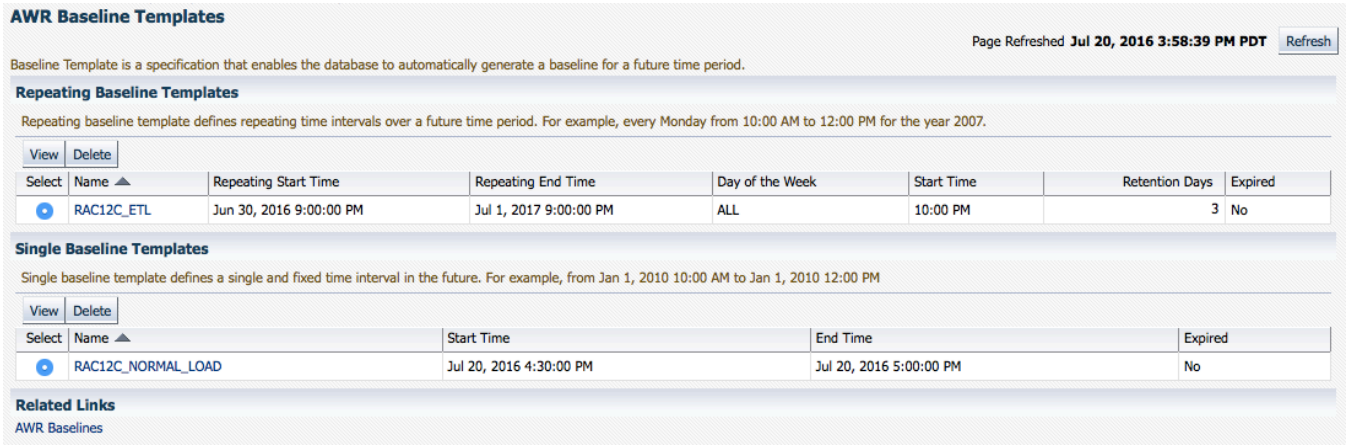

Figure 10

Using multiple static and rolling window baselines, it is possible to capture typical application behavior at various times in the application life cycle and to use that information for comparison against encountered anomalies. These baselines should also be used to help configure monitoring and alerts for various Exadata targets. For more information on Exadata monitoring, see References.

Compliance With Exachk

Compliance Management in conjunction with Exadata's Exachk tool adds significant value by verifying your full Exadata stack is configured correctly based on Exadata Best Practices complied directly from Oracle's development team. Although not a true baseline tool, it does guarantee your systems will remain within Oracle's recommended Best Practice framework as new patches and software upgrades become available. Critical issues and related patches are also tracked for your system so you know right away if your systems are vulnerable and allow you to be proactive and plan upcoming maintenance accordingly.

Compliance results can be access through the Enterprise menu as shown in figure 11

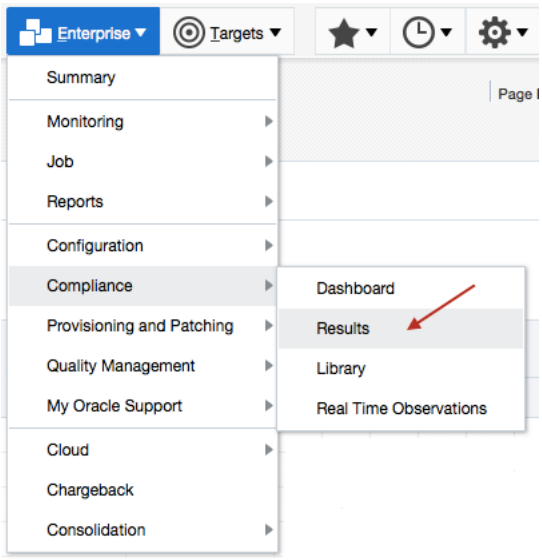

Figure 11

Compliance Standards are grouped by target type. You can further drill down into any problem areas by selecting a listed Standard, Target Evaluation or Violation.

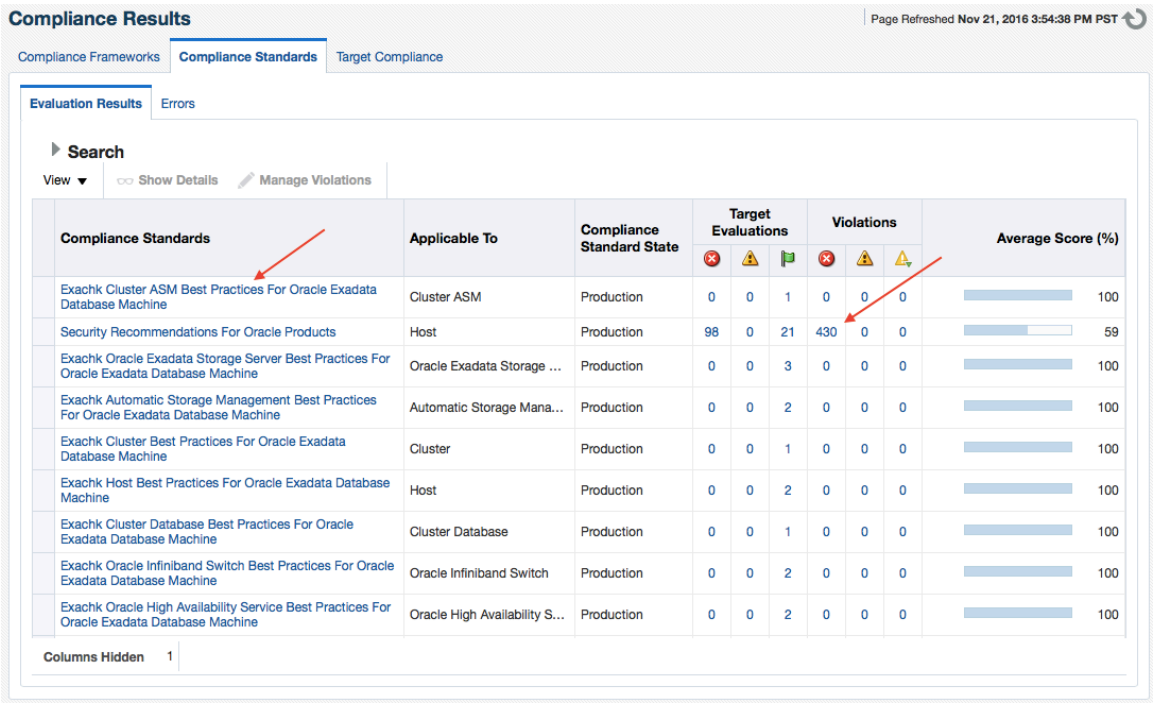

Figure 12

When Performance Degrades

In the lifetime of an application, it is inevitable that there will be system problems at some level, so it is important to have a plan in place to quickly identify where the problem exists and work towards resolution. The Exadata Health and Resource Usage Monitoring whitepaper covers many aspects of this performance troubleshooting process, which includes:

1. **Hardware Alerts** - With Oracle's Auto Service Request technology, hardware problems will be quickly identified and scheduled for diagnostics or replacement if needed.
2. **Configuration Changes** - Enterprise Manager's Configuration Management tools expose any unknown differences in a system configuration for quick identification and resolution.
3. **Application Behavior** - AWR Report Baselining and comparisons help to identify changes in application behavior which result in changes to system resource usage and potential changes in performance.

Additionally, system problem details need to be tracked and reported, and Enterprise Manager's "Incident Manger" provides a central location from which to view, manage, diagnose, and resolve incidents. Details on the Incident Manager can be found in the Enterprise Manager Cloud Control Administrator's Guide.

Capturing a system baseline allows for a repeatable documented approach to diagnosing performance problems when they occur. In addition to ruling out hardware faults, identified via ASR or through the steps defined in the Exadata Health and Resource Usage Monitoring whitepaper, comparing the current system to saved baselines allows fast identification of changes in configuration and workload.

When comparing baselines, the recommended approach is to:

1. Validate the current system against saved configuration baselines
2. Validate current schemas against schema baselines
3. Validate application behavior against AWR baselines

Validating configurations can be automated through Enterprise Manager's Drift and Consistency Management tools with a recommended proactive frequency of once a week.

The following is an example comparison using the Databases Instance Target Saved Configuration. There was a change made to the initialization parameter sga_target, but the change was not propagated to both database instances.

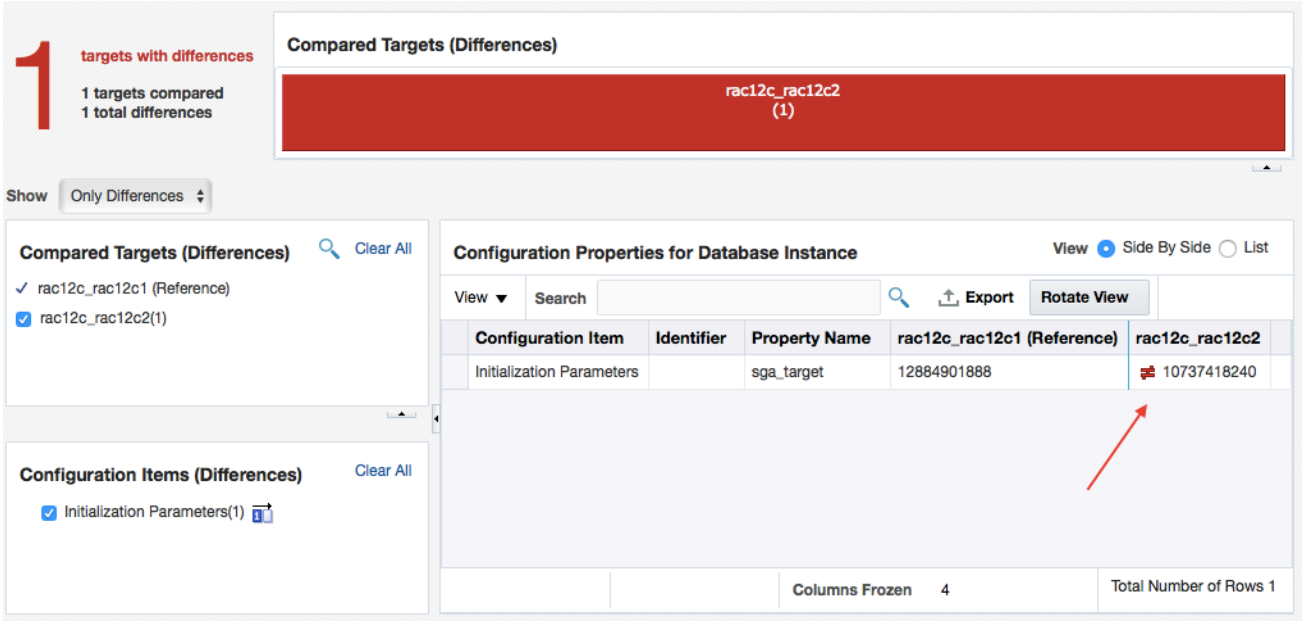

Figure 13

It's important to note that target configuration data is refreshed on a periodic basis, so any system configuration changes will not be immediately visible within Enterprise Manager. A manual refresh would be required to make those changes visible and to use them for any comparison. The Latest Configuration screen for a target includes the Last Collected date/time and also provides a means to refresh the configuration data.

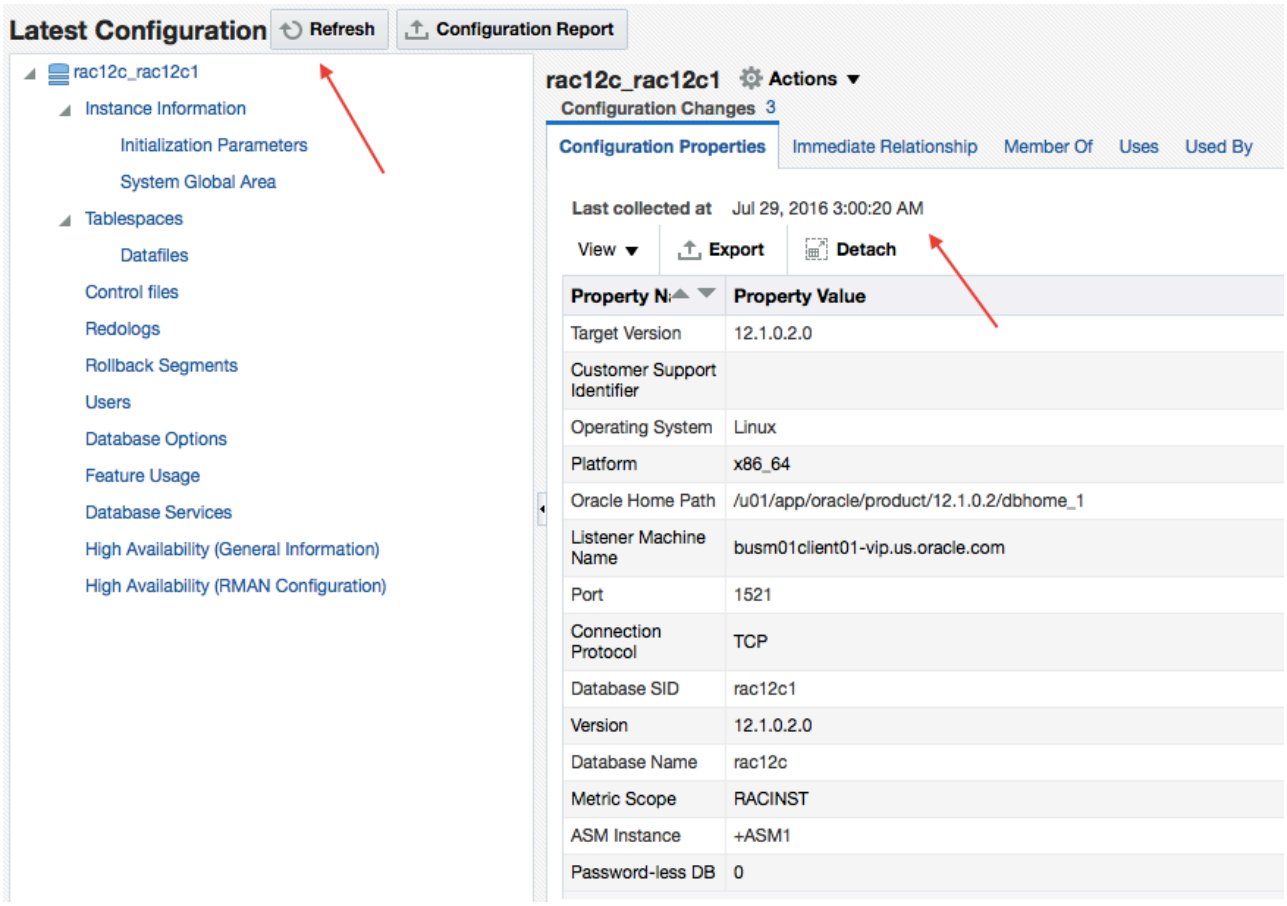

Figure 14

Using the previously saved Schema Baseline, administrators can easily check at the schema level for any changes that may have occurred to cause performance problems. Each schema comparison is saved for historical tracking and later review. The example below depicts a schema comparison using the saved baseline against the current live schema.

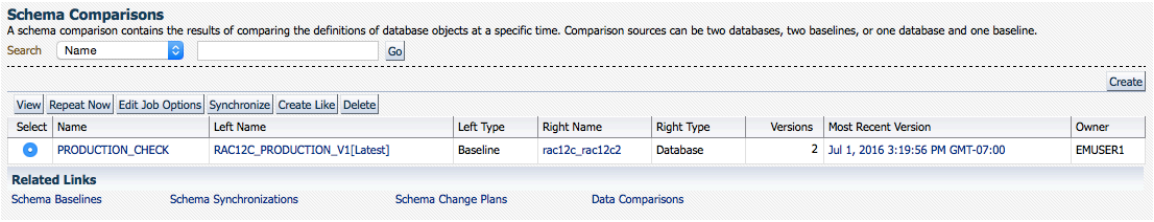

Figure 15

Drilling down further exposes an index with changes made to the leading column, potentially changing an optimizer execution plan. This level of detail also allows quick development of a rollback or roll-forward strategy if necessary, based on SQL which is easily accessible.

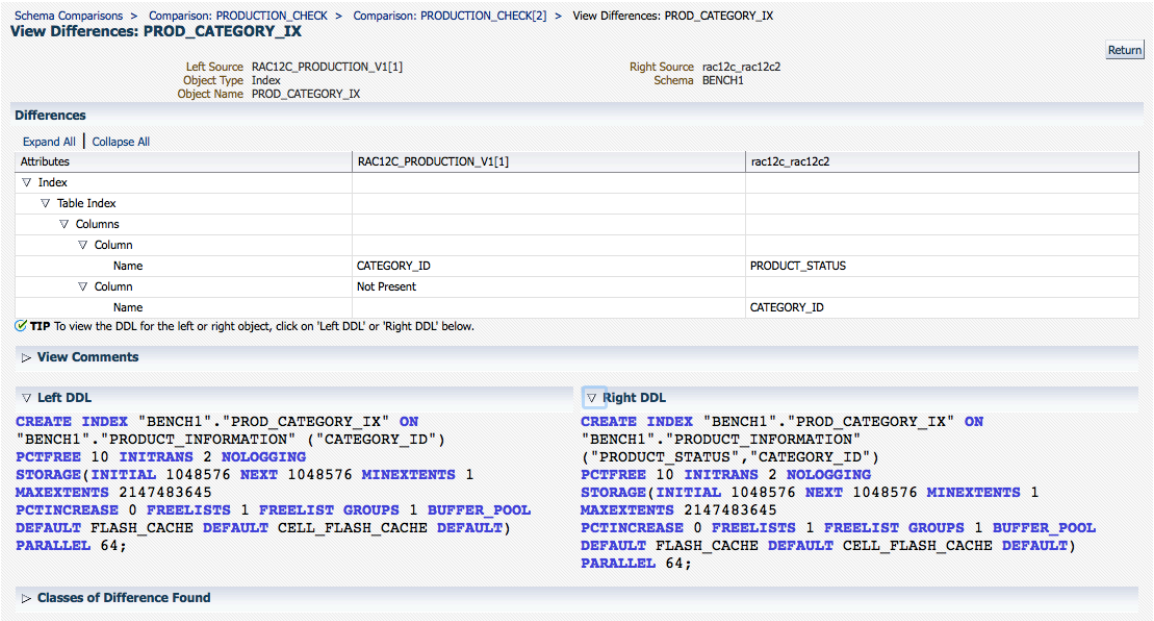

Figure 16

After ruling out any changes to the system's configuration, compare the load characteristics. Using saved AWR Baselines can help identify key areas of the system in which load characteristics have changed significantly and the resultant Workload Repository Compare Period Report highlights key areas in system resource usage to help easily identify those differences.

As an example, in this Top Activity graph something changed over a 10-minute window.

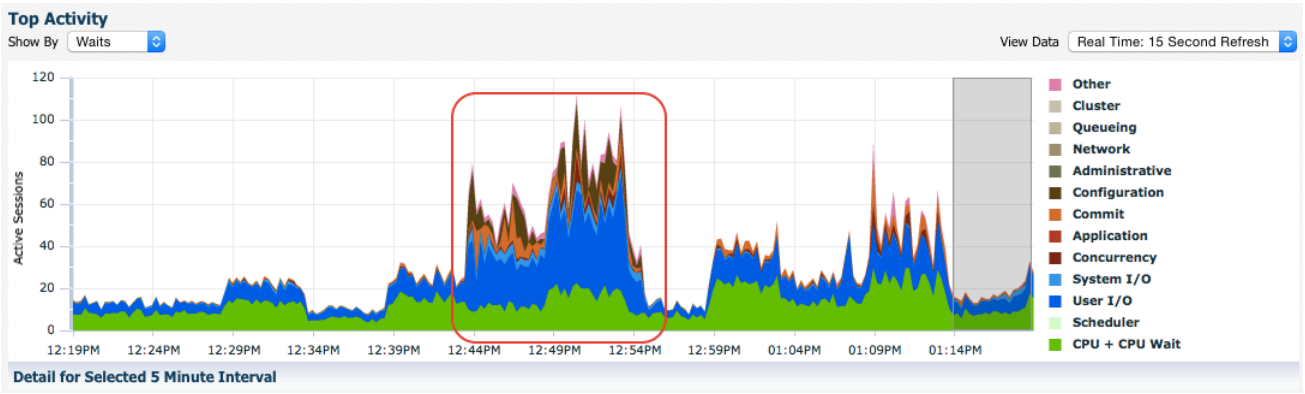

Figure 17

To investigate further, use the saved AWR Baseline of a normal load to make a comparison.

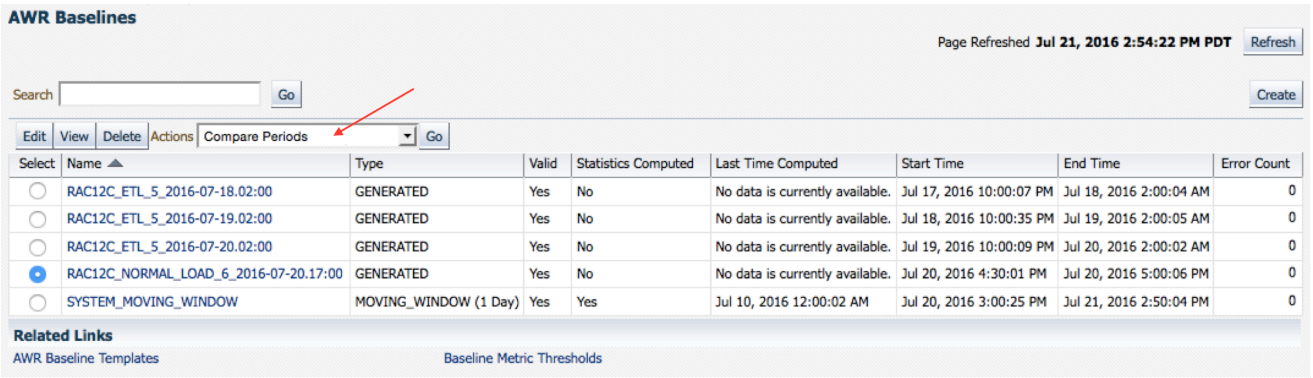

Figure 18

Although the Workload Repository Compare Period Report contains a lot of information, the Top Timed Events and Load Profile sections are the best places to start identifying differences at a high level. In this example using the Load Profile, differences in IO usage are clearly visible.

Load Profile

	1st per sec	2nd per sec	%Diff	1st per txn	2nd per txn	%Diff
DB time:	12.5	18.8	51.2	0.0	0.0	100.0
CPU time:	8.5	5.8	-31.8	0.0	0.0	0.0
Background CPU time:	0.9	0.7	-22.8	0.0	0.0	0.0
Redo size (bytes):	7,723,839.9	101,711,917.6	1,216.9	3,874.6	81,146.8	1,994.3
Logical read (blocks):	244,127.3	180,432.8	-26.1	122.5	144.0	17.5
Block changes:	43,571.9	28,715.9	-34.1	21.9	22.9	4.8
Physical read (blocks):	5,116.7	15,566.4	204.2	2.6	12.4	383.3
Physical write (blocks):	5,143.4	15,770.3	206.6	2.6	12.6	387.6
Read IO requests:	5,115.5	3,968.0	-22.4	2.6	3.2	23.3
Write IO requests:	2,830.4	2,584.2	-8.7	1.4	2.1	45.1
Read IO (MB):	40.0	121.6	204.3	0.0	0.1	400.0
Write IO (MB):	40.2	123.2	206.6	0.0	0.1	400.0
IM scan rows:	0.0	0.0	0.0	0.0	0.0	0.0

Figure 19

Drilling down to the SQL layer in relation to IO, we see two queries are identified that are not part of the normal AWR baseline and thus not part of a normal daily workload. If these SQL statements were discovered to be a new addition to the daily load profile, a new AWR baseline should be captured to include this new load.

Top SQL Comparison by I/O Time

	I/O Time % of DB time					I/O Time (ms) per Exec		Elapsed Time (ms) per Exec		#Exec/sec (DB time)		CPU Time (ms) per Exec		Physical Reads per Exec		#Rows Processed per Exec		#Executions		#Plans	
SQL Id	1st	1st Total	2nd	2nd Total	Diff	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st	2nd	1st/2nd/Both	SQL Text
0w2qpuc6u2zsp	14.78	14.78	24.92	24.92	10.14	2	10	8	18	62.38	25.95	6	5	4.19	4.95	1.00	1.00	1,402,457	881,445		BEGIN :1 := orderentry.neworde...
3fw75k1snsddx	8.41	23.19	14.41	39.33	6.01	1	6	2	7	62.38	25.95	1	1	2.17	2.54	1.00	1.00	1,402,539	881,536	1/ 1/ 1	INSERT INTO ORDERS (ORDER_ID,...
f7rxuxzt64k87	4.96	28.14	7.40	46.73	2.45	0	1	1	2	160.18	66.64	0	0	0.62	0.70	1.00	1.00	3,601,585	2,263,637		INSERT INTO ORDER_ITEMS (ORDE...
g81cbrq5yamf5	0.53	28.67	2.08	48.81	1.55	0	1	0	1	70.18	29.17	0	0	0.14	0.31	2.36	2.36	1,577,819	991,034	1/ 1/ 1	SELECT ADDRESS_ID, CUSTOMER_ID...
1jmf3ggv3xbvq			0.96	49.77	0.96		325,383		417,344		0.00		62,562		1,287,753.00		0.00		1		begin loop insert /*+ append *...
da53t3dxrp2ss			0.94	50.70	0.94		318,310		415,699		0.00		62,607		1,287,752.00		0.00		1		begin loop insert /*+ append *...
apgb2g9q2zjh1	4.18	32.85	4.80	55.50	0.62	5	15	9	21	7.79	3.22	4	5	16.92	17.56	1.00	1.00	175,222	109,530		BEGIN :1 := orderentry.browsea...

Figure 20

In this example the newly added SQL is easily identified, but there will be times when the differences are subtle and Enterprise Manager has additional tools to help with this analysis. This toolset includes the SQL Tuning Advisor, the SQL Access Advisor, SQL Tuning Sets and SQL Plan Control. More information on these tools can be found in the Database SQL Tuning Guide as part of the larger Database Administration documentation for your databases.

Summary

Managing consistency and compatibility across large hardware and software deployments across the enterprise at Cloud scale can be a difficult task. Enterprise Manager helps to simplify this task and enables administrators to view, save, track, and compare configuration information for all managed targets, and provide auditing as an automated process which can be further extended for incident ticketing systems. Using this information and taking time to capture the physical architecture, work load characteristics, and resource usage metrics of the Exadata Database Machine will greatly enhance the speed at which administrators can solve performance discrepancies. Proactive monitoring of configuration changes will also further reduce unforeseen downtime and potential outages, providing the benefit of long term stability in a controlled stable environment.

References

This paper is part of the larger Maximum Availability Architecture document series for Exadata available at: http://oracle.com/goto/maa.

For more information, see the following related publications:

- Enterprise Manager Exadata Management - Getting Started Guide
- Exadata Health and Resource Usage Monitoring
- Exadata Health and Resource Utilization Monitoring - Exadata Storage Server KPIs
- Exadata Health and Resource Utilization Monitoring - Adaptive Thresholds

Additional Exadata Target Comparison Templates

- Enterprise Manager Configuration Comparison Templates for Exadata (Doc ID 2181282.1)

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

ORACLE

CONNECT WITH US


blogs.oracle.com/oracle


facebook.com/oracle


twitter.com/oracle


oracle.com

Integrated Cloud Applications & Platform Services

Exadata Health and Resource Utilization Monitoring: System Baselining for Faster Problem Resolution January 2017
Author: Kurt Schmutz