# TUNING THE ORACLE GRID

*By Richard J. Niemiec, Rolta TUSC*

## ABSTRACT

With Oracle Enterprise Manager Grid Control, you now have a tremendous product at your side. This paper will look at a few of the screens that you can use to monitor the grid. BUT, I warn you, there are so many great screens and tools that I could never do the product justice with such a short paper. See the Oracle Database 10g Performance Tips and Techniques book for a detailed look at this product.

## CLUSTERING AND AN ACCELERATION TOWARD GRID COMPUTING

In June 1970, IBM's Ted Codd published the 11 page paper "A Relational Model of Data for Large Shared Data Banks." This article would lead to the relational databases that would hold the world's data. Advances in information storage/extraction/analysis as well as predicting customer needs has driven us deep into the information age. With 64-bit processing and using grid control, it is theoretically possible to store all of information currently in every database into a single Oracle10g database (Oracle10g allows an 8E database) and load all of it in memory (64 bit allows 16E). The advances of the last decade will be dwarfed by those of the next 10 years. Welcome to the 21st century DBA! This paper is a light hearted look at things NOT to do.

The information age is about to take a drastic step forward. The power of 64-bit computing can be imagined when you consider the theoretical limit for addressable memory. In unsigned 32-bit computing, we could directly address 4G of memory (the sign will cost you 2G). For a standard Oracle database, this allowed a tremendously increased System Global Area (SGA). The SGA is where my most often used data can be stored and kept in memory for fast access. The move to 64-bit starts to accelerate the information age exponentially. With 64-bit, the theoretical limit of addressable memory (2 to the power of 64) becomes 16E (Exabytes) or 18,446,744,073,709,551,616 bytes of memory. Since it is estimated that there is only 12 exabytes of information in the entire world, 16 exabytes is a pretty healthy amount memory (Larry now

can run the entire world in a single Oracle database - WOL). Imagine storing every single piece of information on earth in one database and IN MEMORY. It is theoretically possible, although the physical architecture has not been needed (and hence, not built). Author Note: If you did store all databases on a single system, in memory, I predict there would be a major CPU bottleneck.

Now that we know that the future of hardware theoretically solves any amount of data we will ever need to store in our system, let's move to the database. How we will we quickly and securely access that information? The answer, of course, is Oracle. While competing databases have *some* Oracle features, the information age requires a database that is incredibly fast and tune-able while the system is running, completely available 24x7x52, completely recoverable at a moments notice, allows maintenance on information that is being accessed, altered or even being recovered, and allows for test recoveries and resume-able recoveries of full or partial information. We also now require encrypted backups, encrypted data, and a way to manage it all with less resources. Welcome to the world of Oracle!

How do we scale the hardware that runs this database so that when we need more CPU power or want to service additional users. Welcome to the world of RAC (Real Application Clusters)! Real Applications allow us to share ONE database while having MULTIPLE System Global Areas (SGA) on multiple pieces of hardware. Using Oracle's "cache fusion," which is a process where we can move data from one SGA to another (saving costly disk I/O) when needed (via a high speed fiber interconnect), we get the most scalable Oracle to date. Imagine a single database running on a 8 machine cluster (8 machines hooked up by cache fusion to the same database) with 256G of memory on each instance. That's 2T of physical memory making 1+ terabyte of combined SGA not impossible to imagine along with 256 CPUs and 10T filesystems (required to get to an 8E database). I'm still anxiously awaiting the 16 Exabyte hardware that will run my 8E database. Consider the amount of directly addressable memory from 4-bit to 64-bit and you can see where we are headed now that most hardware is heading rapidly toward 64-bit computing.

| | Memory to Directly Address | Indirect/Extended |
|---|---|---|
| 4 Bit: | 16 | (640) |
| 8 Bit: | 256 | (65,536) |
| 16 Bit: | 65,536 | (1,048,576) |

| | |
|---|---|
| 32 Bit: | 4,294,967,296 |
| 64 Bit: | 18,446,744,073,709,551,616 |

The information age is about to take a drastic step forward. The power of 64-bit computing can be imagined when you consider the theoretical limit for addressable memory. In unsigned 32-bit computing, we could directly address 4G of memory (the sign will cost you 2G). For a standard Oracle database, this allowed a tremendously increased System Global Area (SGA). The SGA is where my most often used data can be stored and kept in memory for fast access. The move to 64-bit starts to accelerate the information age exponentially. With 64-bit, the theoretical limit of addressable memory (2 to the power of 64) becomes 16E (Exabytes) or 18,446,744,073,709,551,616 bytes of memory. Since it is estimated that there is only 12 exabytes of information in the entire world, 16 exabytes is a pretty healthy amount memory (Larry now can run the entire world in a single Oracle database - WOL). Imagine storing every single piece of information on earth in one database and IN MEMORY. It is theoretically possible, although the physical architecture has not been needed (and hence, not built). Author Note: If you did store all databases on a single system, in memory, I predict there would be a major CPU bottleneck.

Now that we know that the future of hardware theoretically solves any amount of data we will ever need to store in our system, let's move to the database. How we will we quickly and securely access that information? The answer, of course, is Oracle. While competing databases have *some* Oracle features, the information age requires a database that is incredibly fast and tune-able while the system is running, completely available 24x7x52, completely recoverable at a moments notice, allows maintenance on information that is being accessed, altered or even being recovered, and allows for test recoveries and resume-able recoveries of full or partial information. We also now require encrypted backups, encrypted data, and a way to manage it all with less resources. Welcome to the world of Oracle!

How do we scale the hardware that runs this database so that when we need more CPU power or want to service additional users. Welcome to the world of RAC (Real Application Clusters)! Real Applications allow us to share ONE database while having MULTIPLE System Global

Areas (SGA) on multiple pieces of hardware. Using Oracle's "cache fusion," which is a process where we can move data from one SGA to another (saving costly disk I/O) when needed (via a high speed fiber interconnect), we get the most scalable Oracle to date. Imagine a single database running on a 8 machine cluster (8 machines hooked up by cache fusion to the same database) with 256G of memory on each instance. That's 2T of physical memory making 1+ terabyte of combined SGA not impossible to imagine along with 256 CPUs and 10T filesystems (required to get to an 8E database). I'm still anxiously awaiting the 16 Exabyte hardware that will run my 8E database. Consider the amount of directly addressable memory from 4-bit to 64-bit and you can see where we are headed now that most hardware is heading rapidly toward 64-bit computing.

| | Memory to Directly Address | Indirect/Extended |
|---|---|---|
| 4 Bit: | 16 | (640) |
| 8 Bit: | 256 | (65,536) |
| 16 Bit: | 65,536 | (1,048,576) |
| 32 Bit: | 4,294,967,296 | |
| 64 Bit: | 18,446,744,073,709,551,616 | |

## MANAGING THE GRID

One of the best screens to manage the grid is displayed below. It's the screen to click on a cluster and see whether the nodes are up or down as well as see the individual nodes. Here is the cluster "ioug" showing six nodes that are all up. To get to this screen, I just went to the Targets tab and clicked on the ioug cluster.

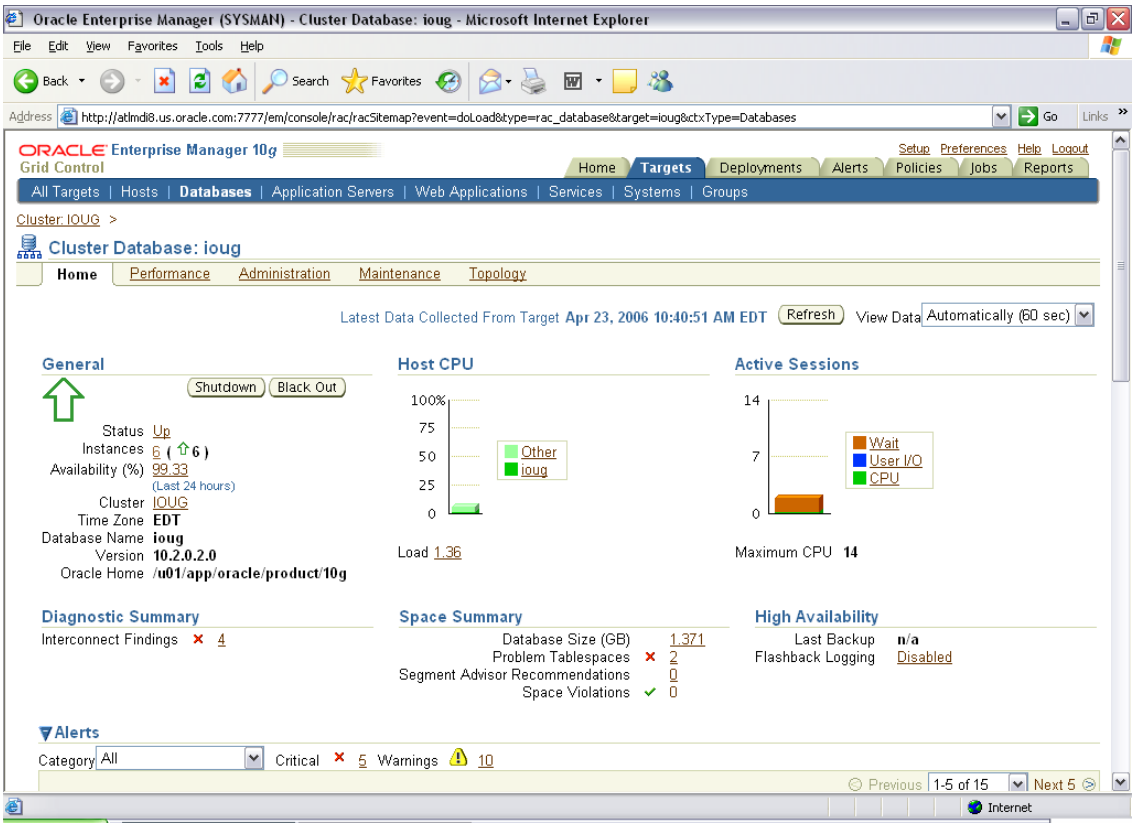


**Figure 1: Looking at the IOUG Cluster Database under Targets/Databases**

If you move down the page a bit, you can see the instances (all using ASM) that are associated with this cluster as shown in Figure 2.

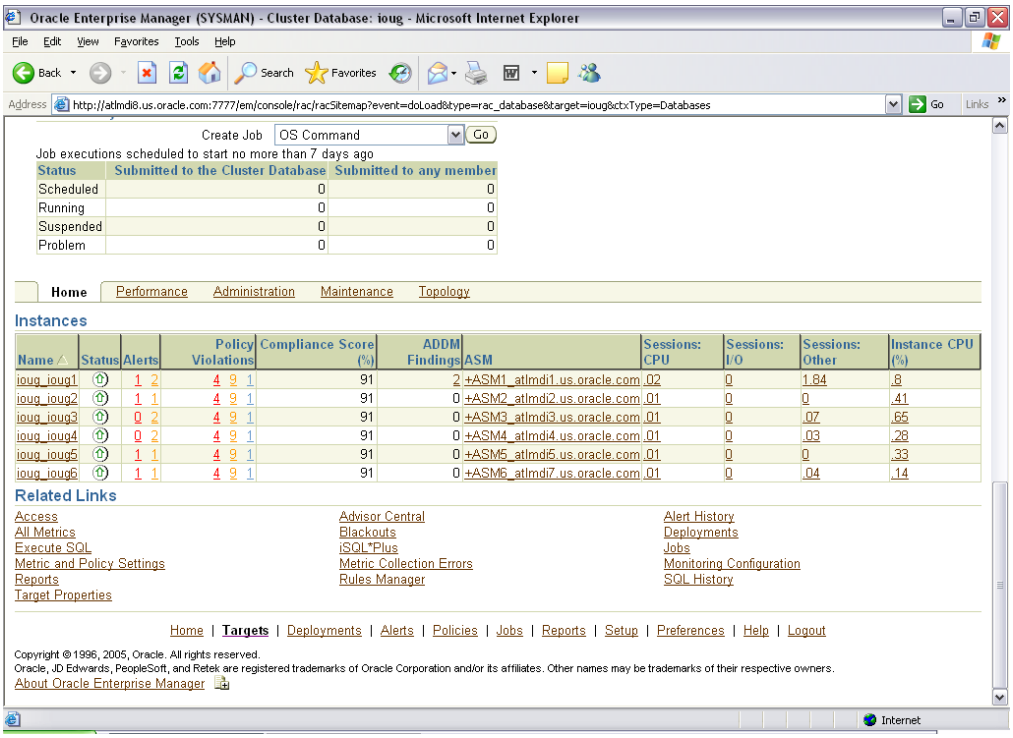


**Figure 2: Further down the page of Figure 1 we see the individual Nodes 1-6**

If I click on the topology tab (see Figure 3), we can see the topology for all six instances (each instance is on a separate node, so there are also six separate nodes. Notice that my mouse if over one of the instances and additional information about the instance is provided.

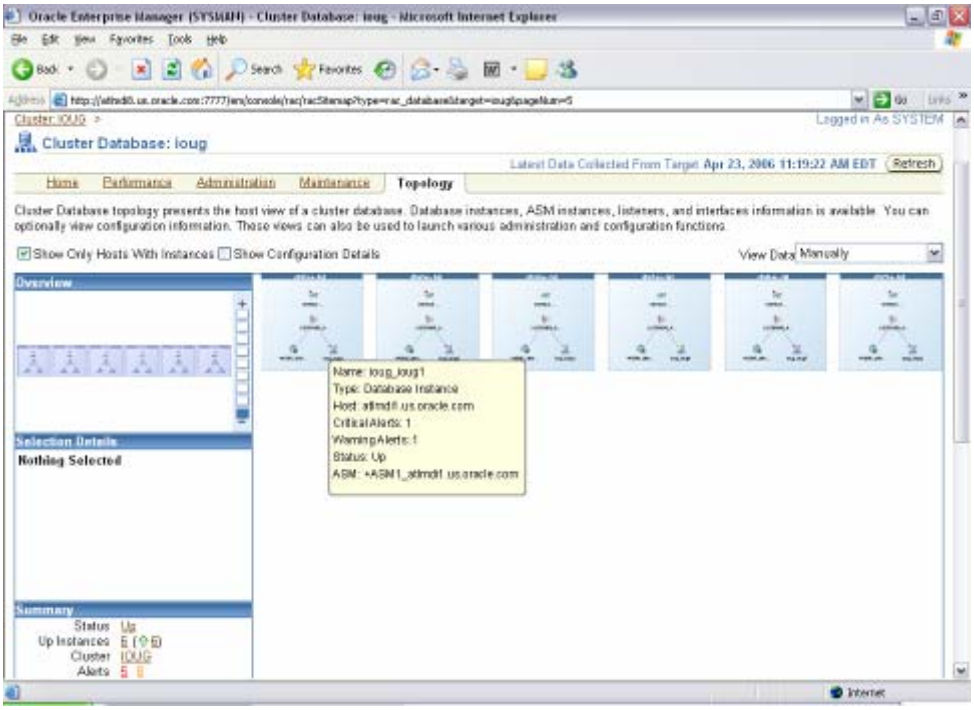


**Figure 3: Looking at the Topology of the 6 Nodes in the IOUG Cluster**

If I click on the Performance tab and then click onto the CPU Used chart (see Figure 4), I can see performance all nodes in the "ioug" cluster, each in a different color.

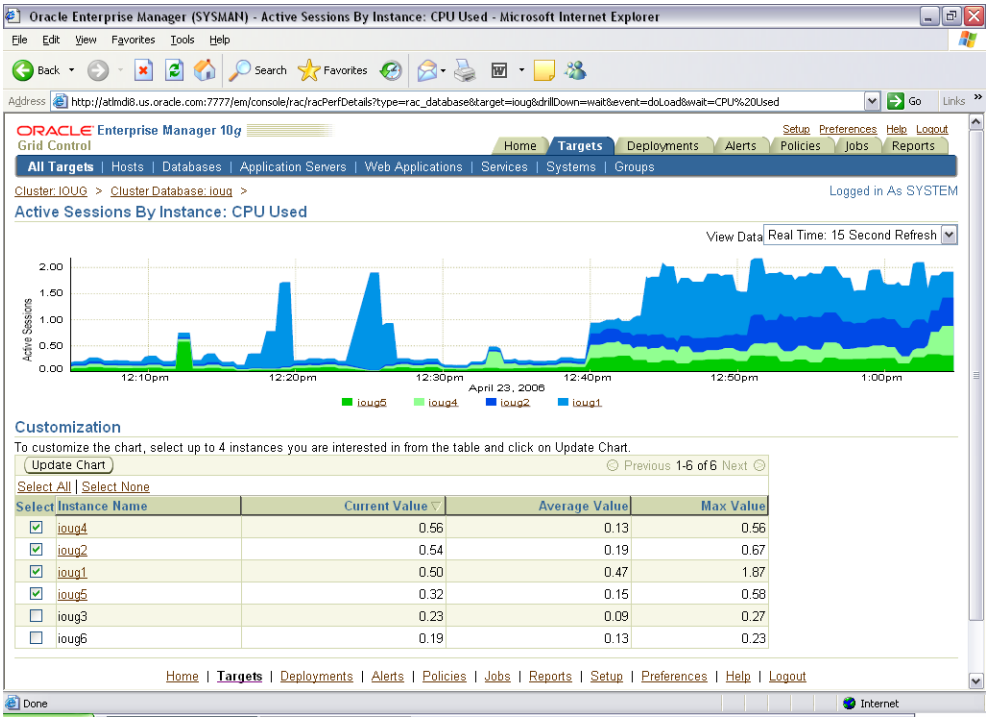


**Figure 4: Looking at CPU for 4 of the selected nodes in the IOUG Cluster**

## RUNNING THE AWR REPORT FROM ENTERPRISE MANAGER

The Database Administration tab of Enterprise Manager can also be used at the Instance level to run the Automatic Workload Repository (AWR) Report. An option from Administration Screen only at the instance level is the link to the Automatic Workload Repository (AWR). Once the AWR option from the Administration screen is clicked, the AWR General information is displayed. This screen includes information on all Snapshots and Collection Levels.

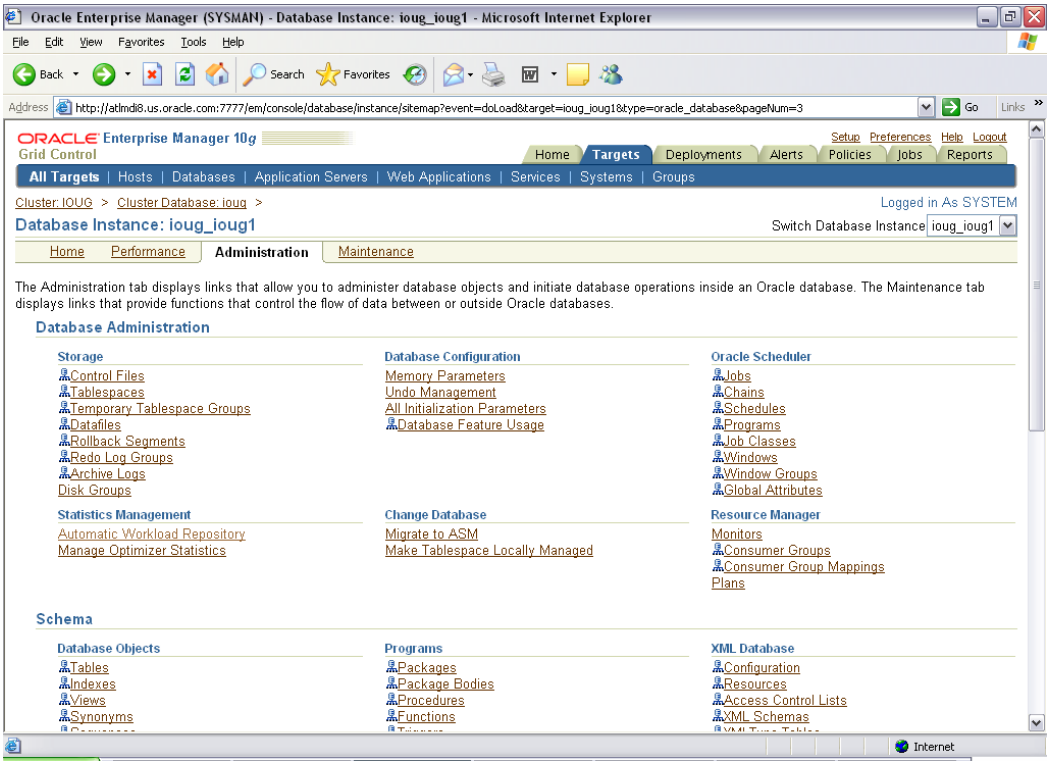


**Figure 5: Database Administration links Instance level**

In the example in Figure 6, there are 40 snapshots with a Retention of 25 days and an interval of 10 minutes (way too often - an hour may be a better interval).

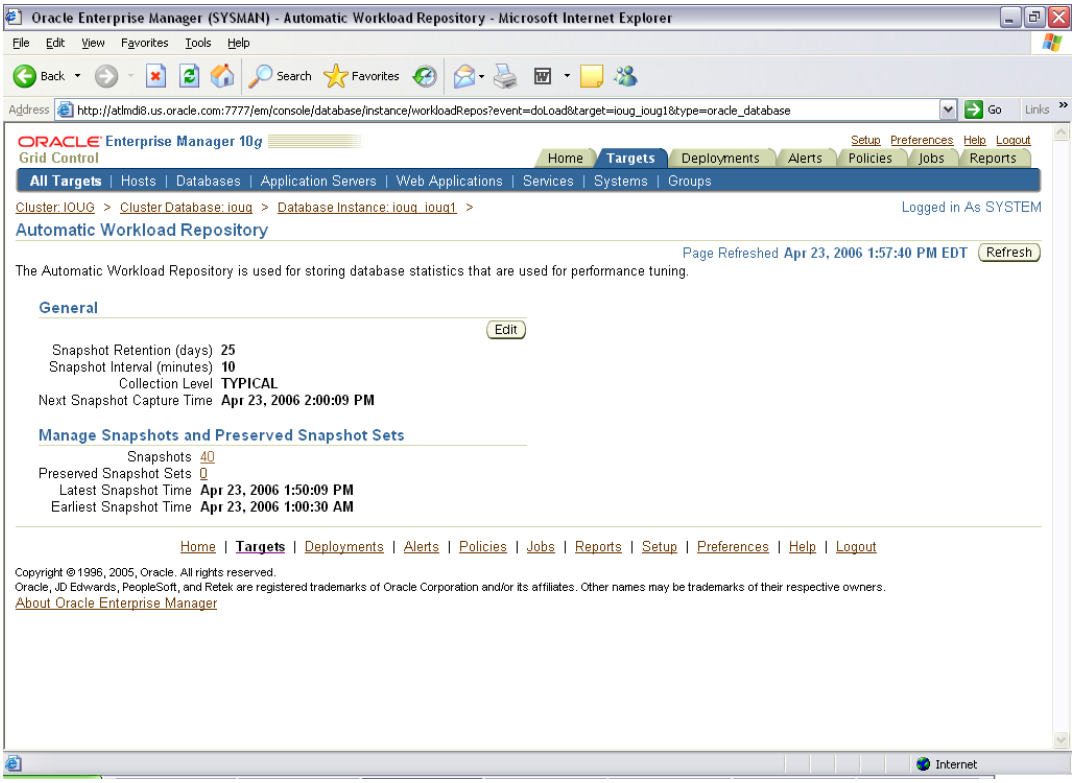


**Figure 6: Automatic Workload Repository (AWR)**

By clicking on the "Edit" button (see Figure 7), the interval or retention of the information may be changed. The collection level can also be edited here.

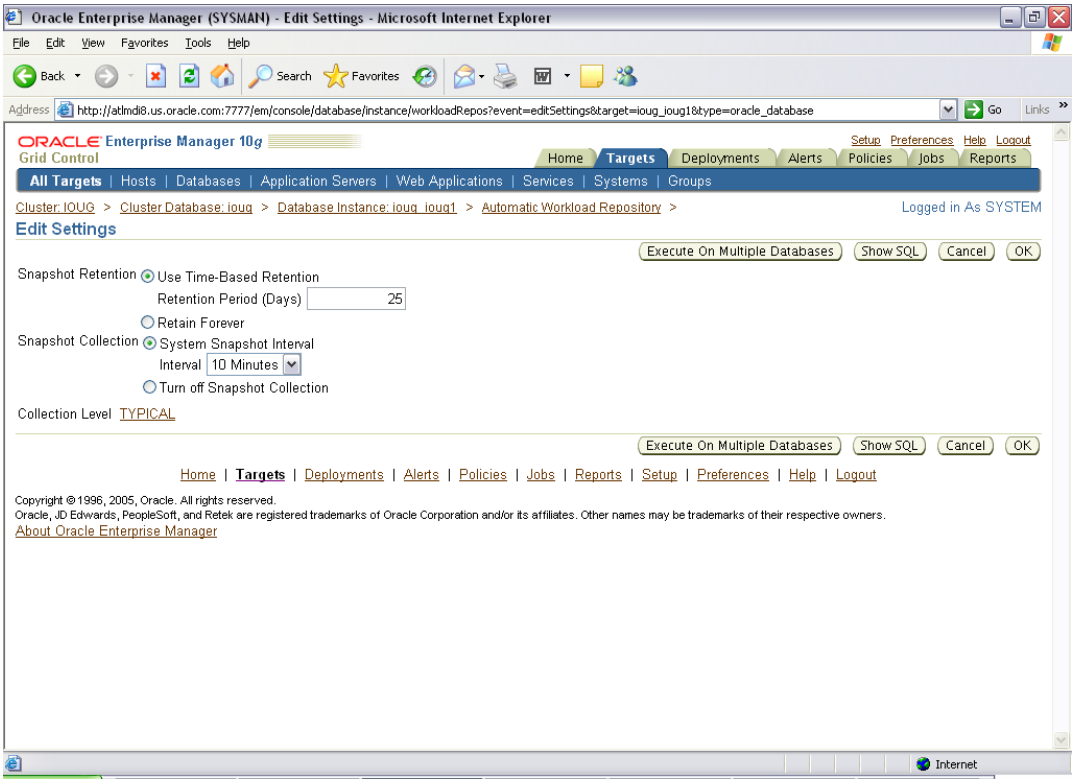


**Figure 7: Automatic Workload Repository (AWR) Edit Settings**

By clicking on the number of snapshots displayed in the AWR General information screen (the number 40 as shown in Figure 6), the 40 snapshots will then be displayed one at a time as shown in Figure 8. The time that the snapshot was generated is listed along with the collection level.

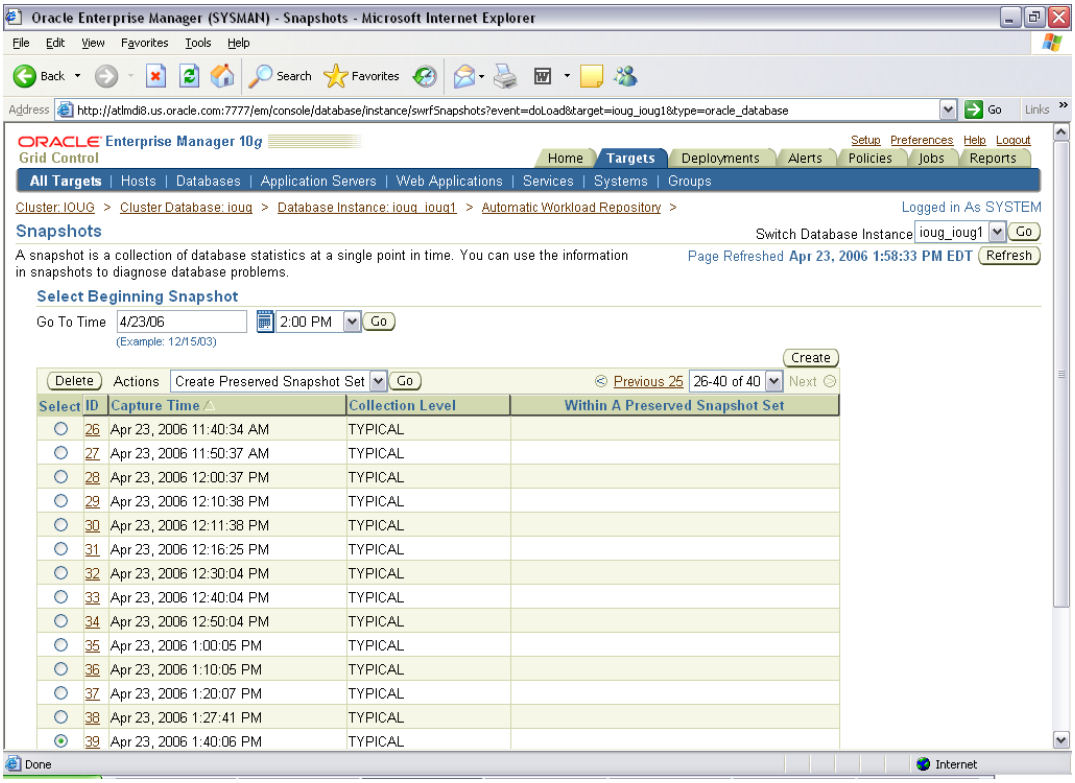


**Figure 8: Automatic Workload Repository (AWR) Snapshot Listing**

Clicking on any specific snapshot to begin and end with will generate some basic snapshot details listed in Figure 9 (like a very mini-statspack), or we can run a report by clicking on Report. This will run and display the AWR Report (Figure 10).

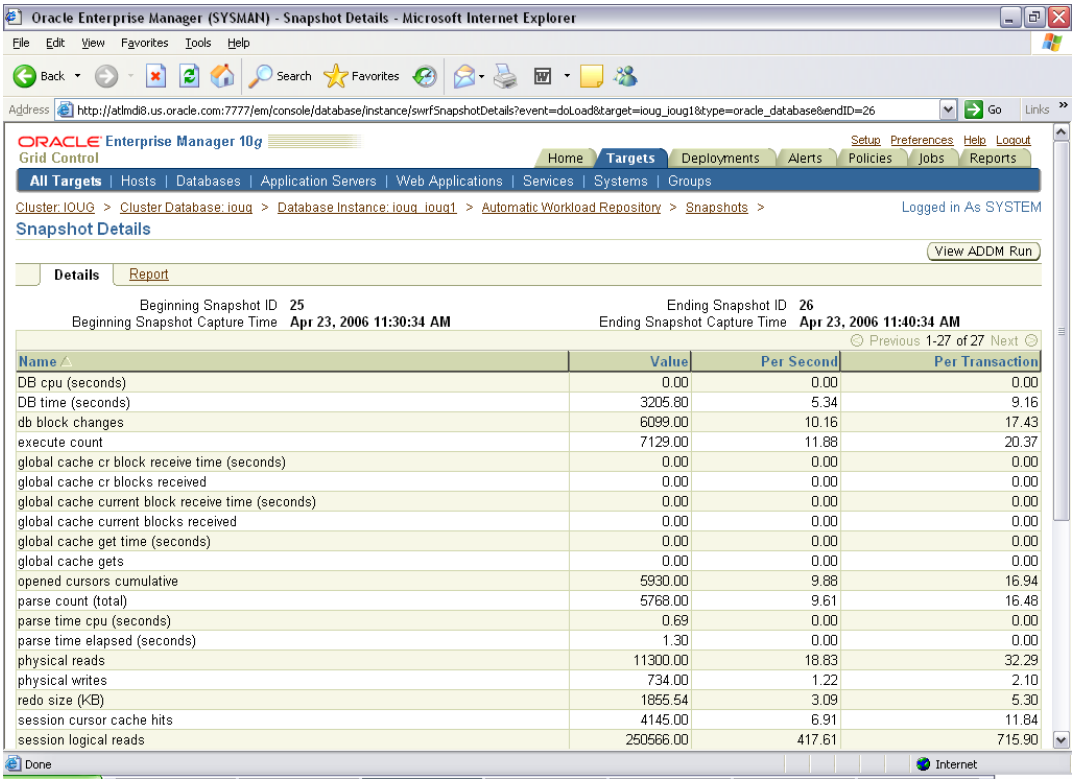


**Figure 9: Automatic Workload Repository (AWR) Snapshot Listing**

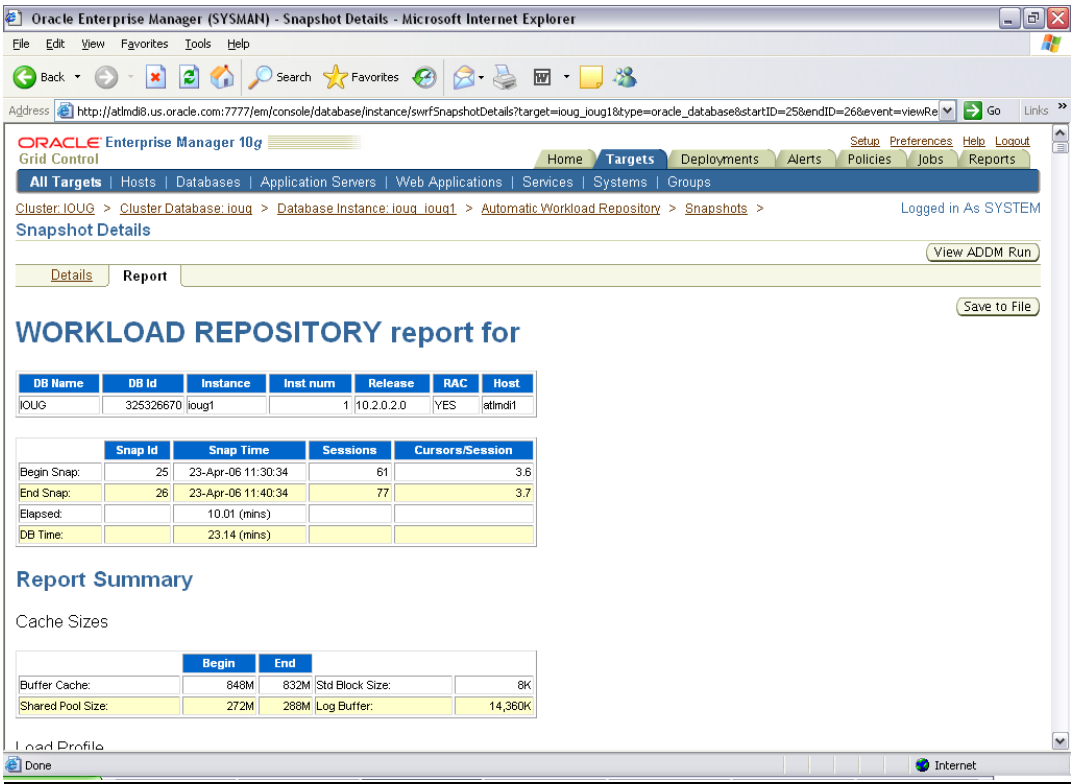


**Figure 10: AWR Report Output**

And there's much more to this tool including interconnect information and global block transfer information. It is a great tool for monitoring the grid. See the Collaborate '08 talk by Rich Niemiec on Grid Control for a lot more information on this great tool.

## REFERENCES:


Oracle11g New Features; Niemiec, IOUG conference 2010

Oracle10g Performance Tuning Tips and Techniques; Niemiec, 2007

Oracle9i Performance Tuning Tips and Techniques, Niemiec, 2003

Oracle Performance Tuning Tips and Techniques, Niemiec, 1999

RAC SIG Event Update, Rich Niemiec & Murali Vallath

Oracle Clusterware, Oracle

Using Oracle Clusterware to Protect 3rd Party Applications, Oracle

IOUG Masters Tuning Class, Rich Niemiec, 2002

www.oracleracsig.org


## AUTHOR BIOGRAPHY

*Rich Niemiec, 48, Rolta's President of International EICT (Enterprise Information & Communications Technology) and President of TUSC – A Rolta Company. TUSC is a Chicago-based systems integrator of Oracle-based business solutions since 1988; TUSC was the Oracle Partner of the Year in 2002, 2004, 2007 & 2008. Rolta is an international market leader in IT-based geospatial solutions, and caters to industries as diverse as infrastructure, telecom, electric, airports, defense, homeland security, urban development, town planning and environmental protection. Rich is the past President of the International Oracle Users Group (IOUG) and the current President of the Midwest Oracle Users Group (MOUG). Rich is one of six originally honored worldwide Oracle Certified Masters. In 2007, he authored the Oracle bestseller "Oracle10g Performance Tuning Tips & Techniques," an update of his previous 2 Oracle best sellers on Oracle8i*

*and Oracle9i Performance Tuning. Rich was inducted into the Entrepreneurship Hall of Fame in 1998.*

*Rolta TUSC is an expert level consultancy that helps companies optimize their investment in Oracle technology. We provide integrated functional and technical solutions since 1988 in the areas of Oracle's E-Business Suite, Business Intelligence/Data Warehousing, Custom Development, Managed Services/Remote DBA, Database Services, Training & Mentoring and Oracle Licensing. Rolta also has software that works with the Oracle database in the areas of GIS, SOA Services and Business Intelligence. Please report errors in this article to rich@tusc.com. Neither TUSC nor the author warrants that this document is error-free. TUSC © 2010. This document cannot be reproduced without expressed written consent from an officer of TUSC except Collaborate 2008 may make copies and make this paper available as needed for the conference and proceedings.*

*Thanks Oracle: Ken Jacobs, Debbie Migliore, Penny Avril, Maria Colgan. and Linda Smith.*



---

Special Thanks to:

Angelo Pruscino, Murali Vallath, Sohan DeMel, Erik Peterson, Debbie Migliore, Sudip Datta, Philip Newlan, Jagan Athreya, Anil Khilani, Bharat Paliwal,Ron Weiss, Ara Shakian, Prabhaker Gongloor (GP), Barb Lundhild, Kirk McGowan, Carol Colrain, Stefan Pommerenk, Troy Anthony, Sue Tang, Joshua Solomin, Eunhei Jang, Supratim Chowdhury, Sourav Mukherjee and Khethavath Singh (KP).