

Performance Tuning Corporation News

March 1, 2009
Volume 1 Issue 1

www.perftuning.com info@perftuning.com (800) 887-4513

Welcome

By Edward Whalen

INSIDE THIS ISSUE

- 1 Welcome
- 2 Oracle RAC and Parallel Query
- 4 Oracle on Linux: Oracle VM
- 5 Oracle on Windows: Hyper-V on Windows Server 2008
- 7 Oracle on Windows: 64-bit is required
- 8 About Us; Virtual Class Schedule

For a while we have been kicking around the idea of starting up a newsletter. We have participated in several newsletters provided by some of our partners, but the focus of these newsletters didn't entirely align with ours. In upcoming newsletters we will be providing useful information on three specific areas; Oracle RAC, Oracle on Linux (including Oracle VM) and Oracle on Windows.

We are going to be producing this newsletter each quarter and are interested in your feedback in its format, content and topics. For example, if you have a specific topic that you would like to see covered in future newsletters, please email us and let us know. In addition, we're interested in hearing what type of articles you would like to see, such as question and answer, tips, etc.

We will also be including information on the latest Oracle releases and any specific issues that we have encountered, such as configuration problems or problems with Critical Patch Updates or hot-fixes. If you have a specific problem that you have encountered, please let me know and I will include it in our next newsletter.

We hope you enjoy our newsletter and find its contents to be interesting and useful in both work and towards your personal development.

Oracle RAC and Parallel Query

By Edward Whalen

Oracle RAC and the Parallel Query Option can be a powerful combination, but if you don't entirely understand how it works; it might not be doing exactly what you think it is. The parallel query option allows some long running operations such as table scans to be divided into smaller chunks and run by multiple server processes. Since operations such as table scans are very I/O intensive, this essentially allows other processes to be doing work on your behalf while some of them are simply waiting on the I/O to return. Please see the following example:

Let's say that you have a table scan that is going to read 1,000,000 blocks. Ignoring all disk caching, driver optimizations, etc., let's assume that each physical I/O takes 2ms. In order to read this entire table into the SGA, it will take 1,000,000 I/Os x 0.002 sec/I/O = 2,000 sec which is 33.3 minutes. It will take a half an hour (plus overhead) just to read the data into the SGA, most of the time which is spent idle waiting on the I/O to return.

Now let's turn on parallel query with a parallelism of 10. Oracle will internally divide this 1,000,000 block table into multiple pieces and pass those pieces off to the 10 parallel query processes. If you are using anything other than a single SATA drive for your database, the I/O subsystem can easily handle the parallelism and will now return those same 1,000,000 blocks in 1/10th of the time, or 3.33 minutes.

Note: I prefer to have control over my parallelism. Therefore, I will use the parallel hint to invoke parallel query as such; `SELECT /*+ PARALLEL(table_alias, degree) */ <columns> FROM table;`

Taking it up a notch, let's consider that the operations are consuming all of the CPU in the system. I can now add additional nodes to use in a RAC cluster for parallel query, thus taking advantage of parallelism and RAC. This can be done by extending the hint to include the number of nodes to use in a parallel query operation as such; `/*+ PARALLEL(table_alias, degree, instances) */ <columns> FROM table;` This will allow me to specify not only the number of query slaves to use, but the number of instances to use as well.

This is only part of the story. Now it is necessary to specify whether we are talking about 10g or 11g. Let's start with Oracle 10g.

“Since operations such as table scans are very I/O intensive, parallel query essentially allows other processes to be doing work on your behalf while some of them are simply waiting on the I/O to return. .”

Parallel query and RAC work differently depending on whether you are running on Oracle 10g or Oracle 11g.

Parallel Query and Oracle 10g RAC

With Oracle 10g, the number of RAC nodes that will participate in a parallel query operation depends on a couple of parameters; `instance_groups` and `parallel_instance_group`. The `instance_groups` parameter is used to set up `instance_groups` and the `parallel_instance_group` parameter is used to specify which group you belong to.

First, set up the instance groups. In this example I have a four node RAC cluster:




Illustration of Instance Groups

- `INSTANCE_GROUPS`
- Assign the node a set of instance groups that it participates in
 - `node1.instance_groups = node1,first_two,all_nodes`
 - `node2.instance_groups = node2,first_two,all_nodes`
 - `node3.instance_groups = node3,last_two,all_nodes`
 - `node4.instance_groups = node4,last_two,all_nodes`




Illustration of using Parallel Instance Group "node1"

The graphic on the left illustrates the configuration.

Setting the `parallel_instance_group` parameter will define which group you belong to, thus defining where your parallel queries can run.

For example, if I set my `parallel_instance_group` parameter to Node1, my parallel queries will only run on node 1 as shown to the left:




Illustration of using Parallel Instance Group "all_nodes"

Setting the parameter `parallel_instance_group` in the `init.ora` file will specify the default value for each user who logs into that instance. An `alter session` command can be used to modify that parameter as shown here: `SQL> alter session set parallel_instance_group=all_nodes;`

This will allow the parallel query to span all nodes in the RAC cluster as shown to the left.

Note: You don't have to be on the node that parallel query is configured to run on. If you happen to log onto node 3 and your `parallel_instance_group` is Node1, non-parallel queries will run on node 3, but any parallel queries will run on node 1.

Parallel Query and Oracle 11g RAC

What about Oracle 11g? If you are running on Oracle 11g, forget everything I just told you. One of the best features about Oracle 11g is that it now uses services to specify which nodes your parallel query runs on. Depending on which service you connect to, the preferred nodes for that service are the nodes that parallel queries will use.

Conclusions

Depending on the operation that you are doing, you might be able to achieve huge performance gains by parallelizing across multiple nodes in a RAC cluster. Of course your mileage might vary, and if you aren't careful you might increase the load on nodes to a point where you are affecting other users. I have a saying, "If you aren't careful with parallelism, you might paralyze your system."

Oracle on Linux: Oracle VM

By Edward Whalen

Virtualization has recently become a standard part of almost every data center. Recently Oracle has entered the virtualization market with Oracle VM. Based on Xen technology, the Oracle VM product is specifically enhanced for Oracle products and is the only virtualization product supported by Oracle. Oracle VM utilizes core Oracle Virtual Machine Servers dedicated to supporting virtual guest machines, as well as utility servers used for installation and maintenance of virtual machines.

Oracle's product offering includes the VM Server that is made up of a small linux distribution and the Xen server products and the VM Manager. There is one VM Manager that supports a web application for managing all of the VM Servers. You can have multiple VM Servers to support virtual machines. If you have the ability to support shared disks, HA features are available, including a live migration.

In upcoming newsletters I will be covering topics related to Oracle on Linux, including Oracle VM. Please feel free to email suggestions for future topics to info@perftuning.com.

"Based on Xen technology, the Oracle VM product is specifically enhanced for Oracle products and is the only virtualization product supported by Oracle."

Oracle on Windows: Hyper-V on Windows Server 2008

By Gary Parker

Microsoft's current server operating system offering, Windows Server 2008, includes its first enterprise-worthy virtualization product bundled with the 64-bit version of the OS. Hyper-V, which is included as a role in the operating system for a small additional cost, is the next generation hypervisor-based virtualization technology. It allows you to make the best use of your hardware investments by consolidating multiple server roles as several separate virtual machines running on a single physical machine. Hyper-V will efficiently support multiple operating systems - Windows, Linux and others - in parallel on a single physical server to fully leverage the power of 64-bit computing.

"Recently, PTC conducted a benchmark comparison of Oracle 10g performance on a Windows Server 2008 native platform and a Hyper-V virtual platform..."

Hyper-V Server 2008 is especially suited to meet testing and development virtualizations needs. Using virtual machines, development staffs can create and test a wide variety of scenarios in a safe, self-contained environment that accurately approximates the operation of physical servers and clients. This is especially beneficial in an Oracle environment, where patches and / or new releases of the Oracle software can sometimes break an existing application.

Recently, PTC conducted a benchmark comparison of Oracle 10g performance on a Windows Server 2008 native platform and a Hyper-V virtual platform also running Windows Server 2008 using SwingBench v2.3. Analysis of the benchmark results led to the following observations:

- The Hyper-V virtual machine Oracle 10g database was able to provide an almost identical performance to the Oracle 10g database running on the physical database server.
- Hyper-V guest virtual machines are limited to a maximum of four CPU cores; therefore, you should run Oracle 10g within Hyper-V guest virtual machines only if your workload performance can be satisfied by no more than four CPUs.
- When compared against native configurations with comparable hardware resources, the same throughput can be achieved within a guest virtual machine at a cost of slightly increased CPU utilization.

(cont.)

- The information captured is specific to performance considerations; for your deployment, take functional considerations (i.e., supported configurations, options to achieve high availability, and so on) also into account.
- For consolidation scenarios, the amount of storage resources available as well as the scenario will drive your decision. You should size your storage with your workload and response time requirements in mind. Always follow best practices with respect to the underlying storage in Hyper-V environments, just as you would with any Oracle 10g deployment.

The results achieved in this test demonstrated that the Oracle 10g database can deliver an almost identical performance regardless of whether the underlying platform is a physical server or a Hyper-V virtual machine.

For more information regarding Oracle on Hyper-V and the details of the benchmark comparison, read our full whitepaper at:

<http://www.perftuning.com/pdf/Oracle%2010g%20on%20Hyper-V%202008.pdf>

Oracle on Windows: 64-bit is required

By Edward Whalen

I've mentioned several times that I strongly recommend running Oracle for Windows using 64-bit. I've had a number of questions about this, so I wanted to clarify the importance of 64-bit Oracle on Windows.

Oracle on Microsoft Windows is unlike any other platform that Oracle supports. When Oracle was originally ported to windows, the decision was made to take advantage of Windows threads. A thread, sometimes known as a lightweight process, shares the same context as the calling process. Thus, the context switch is replaced by a thread switch (which uses much less CPU resources). With Oracle for windows, all traditional processes (background and server) are implemented as threads. This provides for more efficient processing, but causes another problem.

“Within the thread model, all threads within a process share the same virtual memory address space.”

Within the thread model, all threads within a process share the same virtual memory address space. With Windows 32-bit, this virtual address space is 3 GB in size (with /3GB flag, 2 GB otherwise). Thus, all of the Oracle processes share the same 3 GB virtual memory limit. In a process environment such as Unix, each process has its own 3 GB virtual memory address space. This limitation can cause the virtual memory space to be completely used, thus causing failures. The symptom of this in 32-bit Oracle is a failure to allocate memory for new dedicated server processes. This problem has been solved with the 64-bit version of Oracle for Windows. The 64-bit version of Oracle for Windows supports a 16 Terabyte virtual memory limit. In addition, memory above 4 GB can be accessed directly and indirect data buffers need not be used. This makes the use of memory above 4 GB much more efficient and faster.

When implementing Oracle with an SGA that is more than 4 GB in size I always recommend the use of large memory pages. Large memory pages are much more efficient for large SGAs and make a huge performance difference, which will be the subject of my next blog. For information on how to use large memory pages in Oracle on Windows see Metalink note 422844.1.

The limitations in this article lead me to strongly recommend the use of 64-bit Oracle on 64-bit Windows. So, “Just say no, to Oracle on 32-bit Windows”.

About Us

Performance Tuning Corporation (PTC) is a technology services company that designs, implements and optimizes IT database and application solutions for companies worldwide. In business since 1997, we service small and large companies, U.S state and local government and various federal agencies. In addition, we staff classified projects with consultants cleared to the Top Secret level.

Our Consulting Services Include:

- Consulting Projects
- Staffing
- Technical Training
- Remote DBA Support

Performance Tuning Corporation

13492 Research Blvd.,
Suite 120-242
Austin, TX 78750-2254

Phone:
(800) 887-4513

E-mail:
info@perftuning.com

Please refer to our web site for more details on the services we provide at: <http://www.perftuning.com>

For a list of our training courses, please go to:
<http://www.perftuning.com/pages/training.php>

Virtual Training Class Schedule

Below is our Spring 2009 schedule of virtual training classes:

March 17-19: Oracle 10g RAC for Administrators

March 24-26: Oracle 10g RAC for Administrators

April 14-16: NEW CLASS!!! - Oracle VM for Administrators

For details on our virtual classes pricing and how to register, please go to: http://www.perftuning.com/pages/t_virtual_class1.php