

Official Publication of the Northern California Oracle Users Group

NoCOUG

J O U R N A L

Vol. 23, No. 1 · FEBRUARY 2009

\$15

Weather Tough Times with NoCOUG

Spotlight on Oracle

Candid questions get candid answers from Dr. Bert Scalzo. See page 4.

Oracle on VMware

Brian Hitchcock reviews Dr. Bert Scalzo's latest book. See page 6.

Bolt-On Power

Chris Lawson discusses the perils and pitfalls of RAC. See page 18.

Much more inside . . .

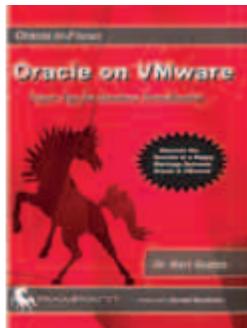
Oracle on VMware

Expert Tips for Database Virtualization

A Book Review by Brian Hitchcock

Details

Author: Dr. Bert Scalzo
ISBN: 978-0979795145
Pages: 232
Year of Publication: 2008
Edition: 1
Price: \$34.95
Publisher: Rampant Techpress



Summary

Overall review: Worthwhile as a brief introduction to Oracle databases running on VMware.

Target audience: Almost anyone who supports or manages Oracle databases.

Would you recommend to others: Yes.

Who will get the most from this book: Anyone new to virtualization for Oracle databases.

Is this book platform specific: Yes, only Linux and Windows are discussed.

Why did I obtain this book: I was asked to review this book for NoCOUG.

Overall Review

First, I have no experience with VMware. This book discusses VMware using Windows and Linux. Again, I don't have experience with Windows or Linux. My experience is on Solaris. This means that I can't comment on the details of the content that refer to aspects of Linux and Windows, but it also means that I am able to review this book as someone completely new to virtualization. I think many potential readers of this book will be in the same position.

This book is a quick and easy read. For anyone who is wondering what virtualization for Oracle might be like, I recommend reading this book. If you are already in the midst of implementing VMware, this book probably won't tell you much that you haven't already experienced. In addition, I

would recommend this book as a general overview of the issues that not only apply to VMware for Oracle but also will come up for anyone wondering about cloud computing, the grid, or any other setup where you are trying to gain the benefits of multiple databases on multiple shared servers.

On the minus side, the index is minimal and when I tried to look something up, it wasn't in the index. I saw many typos in the text, and there were several places where the same exact string of text was used over again. More thorough editing would have taken care of these issues. At multiple points in the text we are told about the "DVD that is included with the book." The problem is, there is no DVD. Using the URL for the Online Code Depot from page 1 of the book, I found that the contents of the missing DVD are available for download.

Prologue

Here Don Burleson (the series editor) takes us through the history of virtualization from the early days on IBM mainframes to VMware today. He tells us that the "2nd age of mainframe computing" is arriving soon in the form of large servers split into virtual machines using VMware. Being able to run multiple operating systems on the same physical server is a step toward OS independence. An overview of Oracle's development of its own version of VMware is described, followed by a description of the more common concerns about VMware—which we are told are misconceptions. These include single point of failure and a single application hogging the whole server. Specific shortcomings of Oracle VM are offered, namely resource sharing, overhead, and that DBA jobs will suffer. While I'm concerned about DBA jobs being eliminated, I'm not clear who, if not more DBAs, will determine the overhead imposed by Oracle VM and which applications can coexist in the same physical server, will move applications between servers to attempt to improve performance, and so on. As always, new technologies promise to simplify the IT infrastructure and reduce headcount. From what I've seen of clustered environments and RAC specifically, the promised benefits of simplification and reduced headcount continue to be slightly out of reach.

“As always, new technologies promise to simplify the IT infrastructure and reduce headcount. From what I've seen of clustered environments and RAC specifically, the promised benefits of simplification and reduced headcount continue to be slightly out of reach.”

Chapter 1—Introduction

The meaning of virtualization is covered, including platform versus resource virtualization. Five specific types of platform virtualization are illustrated, and the book focuses on the first three. The benefits of virtualization are described, including the usual suspects of lower cost and improved resource utilization. It is pointed out that virtualization is similar to Oracle's RAC and grid products. The overhead needed for virtualization is stated to be 10–20%, but we are told that this isn't an issue because the hardware is ten times faster. From comments made later in the book, I assume the author means that the newest servers are faster and cheaper, and that this will easily cover the overhead due to virtualization. The text doesn't offer any insights into how this figure of 10–20% was generated. There is mention that Oracle hasn't (yet?) certified the database for virtual machines. There is no further discussion of this issue, and while I will assume that this certification has, or will soon, be done, you must verify this if you are about to implement any form of virtualization. The reasons to use VMware, why virtualization is the future, and some details about Oracle VM are covered as well.

Chapter 2—Architecture

The focus here is on two of the architectures described in the previous chapter. The two are Para-Virtualization using the full-cost VMware ESX hypervisor and Full Virtualization using the freeware VMware Server (GSX). The former is for full installations and the latter is for test and development use. The hypervisor is described as a stripped-down operating system specifically optimized to run virtual machines and nothing else. The 10–20% overhead for virtualization is brought up again, but there are still no details on whether this is a real metric or a number simply to illustrate the author's contention that new hardware is so much faster and cheaper that even a 20% overhead doesn't matter. From my inexperienced viewpoint, virtualization should help reduce underutilized servers, but how much better utilization will I have if I first pay a 20% utilization penalty? I would like to have seen a more thorough discussion of the overall utilization improvement when virtualization, its overhead penalty, and new hardware are all taken into account. We are also told that hardware is now so cheap that it is better to "throw hardware" (this is the author's term) at performance problems before spending person time. This will be controversial for some Oracle tuning experts. Note that I am not taking sides on this issue other than to agree that person time spent on tuning exercises is very expensive.

Next there is a worthwhile discussion about the continuing trend to more and more abstraction. The conclusion from this discussion is well stated: "nothing maps one-to-one anymore." Going back to the subject of tuning, the author offers that it is no longer possible to look for a "hot" disk, and perhaps one should no longer try. More generally, the point is that with virtualization, resources are being allocated dynamically. It becomes very hard to look at tuning one SQL statement in one Oracle database. Most Oracle tuning books still focus on specific ways to tune SQL statements. Is this relevant in the virtualized world? Again, this will be controversial but it is a topic that must be addressed.

Finally, virtualization means that the DBA will have to deal

"We are told that hardware is now so cheap that it is better to "throw hardware" at performance problems before spending person time. This will be controversial for some Oracle tuning experts."

with more and more people to get things done. To implement a database in such an environment, the DBA may have to work with administrators for storage, network, OS (one or more?), virtualization, and perhaps an infrastructure architect. This also means that the environment is more complex. Remember how virtualization was going to reduce DBA headcount and simplify the environment?

Chapter 3—Host Setup

Here we see that the performance of the host, i.e., the physical server that is hosting all the virtualized Oracle servers, is the limiting factor on the performance of the Oracle databases. The author tells us that any database server that is a guest on the host server can't be more optimized (tuned) than the host server itself. This means that you must set up the host correctly, and then, when there are database performance problems, you need to look at the host server first. This assumes that you know how to correctly set up the host, and it will cause issues for DBAs that want to tune any database issue by looking for the SQL statement that is consuming the most resources within the database. We are also told that if you can improve host server performance by 10%, this will improve all the hosted databases by 10%. I'm not sure this is true. I think it's an assumption that all databases will be using the same host resources in the same way at the same time. No details or examples are given to support this. I'd like to see some detailed studies of this issue to see exactly what the relationship is between changes in host server performance and changes in database performance for a selection of differently configured databases (OLTP, data warehouse, IO intensive, memory intensive, etc.).

More discussion of the need for lots of hardware follows—specifically the recommendation for as much CPU and memory as budget will allow, since the author believes that adding hardware to solve performance issues is cheaper than the cost of the person hours to attempt to tune the existing hardware. The author brings up a good point that many people will start to evaluate virtualization using old hardware that isn't up to the task and then conclude that virtualization isn't a good solution. It's one of many chicken-and-egg situations in the real world. Until you know enough about virtualization to competently choose and configure the needed hardware, you can't do a competent evaluation of virtualization, which is something to keep in mind when you think you will just add "evaluate virtualization" to your group's project task list. Despite the marketing claims (simple! faster! fewer DBAs!), it really is more complex, and it will not be simple to evaluate for your business. The rest of this chapter has many specific details for

“Virtualization means that the DBA will have to deal with more and more people to get things done. To implement a database in such an environment, the DBA may have to work with administrators for storage, network, OS (one or more?), virtualization, and perhaps an infrastructure architect. This also means that the environment is more complex. Remember how virtualization was going to reduce DBA headcount and simplify the environment?”

installing, configuring, and optimizing the Linux (or Windows) host operating system to support virtual machines.

Chapter 4—Guest Setup

This chapter focuses on creating and configuring the virtual machines, the operating system installations that are guests on the host server. Specifics are covered for both Linux and Windows. The author describes how virtualization provides more options when database performance is lacking. Instead of moving the database to a different server, you can assign more host server resources to the existing guest virtual machine. And while this new level of flexibility is more complicated, the author suggests that this is a good thing because it means the DBA has more job security. I’m confused, but that isn’t unusual. I guess this means that the DBAs that don’t get laid off have more security.

The process of creating the virtual machines is covered in great detail, including screen shots of the windows seen in VMware during the process. The importance of the DBA understanding and having input in the selection of the VM settings is emphasized. The DBA will have to learn some new tricks to have the aforementioned increase in job security. ’Twas ever thus! One of the windows tells us that the maximum memory that can be assigned to a virtual machine is 3580 MB. I hope I’m misunderstanding this, because many databases will require much more than 3.6 GB of memory, and I assume the virtual machine will require some of this memory itself.

Once installed, the steps for tuning the virtual machine for virtualization and the Oracle database are covered. This covers optimization of CPU, memory, IO, and network resources. Specifics are given for both Linux and Windows. There is an

¹ Editor’s note: Brian meant “skillsets” of course. NoCOUG editor Karen Mead asked if database administrators use frying pans in their work. Perhaps Brian wrote this review right after eating breakfast!

interesting section that explains that neither Linux nor Windows is superior in terms of TPC performance, and therefore the choice of which to use for the virtual machine is more a matter of available skillsets¹ in your organization. Again, I won’t take sides.

Chapter 5—Oracle 11g Setup

This chapter offers a process to create a demo installation for 11g, including VMware on minimal hardware such as a laptop. The reason this is needed is that the DBA needs to tune the virtualized database installation and then use this “golden” installation as a source of files for further installations. I can see why you would need demo systems—perhaps for training, for instance—but I’m not clear that time spent tuning Oracle on a laptop would directly apply to the full-up installations of Oracle that will be needed for the real servers. Because the installation is being done on a laptop, details are given for choosing BIOS settings along with all the other aspects of the operating system and the database. I wonder how many DBAs are qualified to tinker with BIOS settings. I’m told that I should only install the minimum Windows services needed. More things for the DBA to know about. The installation of 11g is also covered in great detail.

Chapter 6—Pseudo RAC setup

Similar to the previous chapter, here we see how to set up a RAC installation on a single node (a laptop). We are told that this is for demo and learning purposes. The fact that Oracle doesn’t support this is explained, which leaves me unclear on how this will be useful for the actual RAC setup of 11g on VMware. As an aside, the author offers statistics that 70% of users currently using 9i or 10g don’t plan to move to 11g, and therefore the database version installed in this chapter will be 10g. But we installed 11g in the previous chapter, didn’t we?

Chapter 7—True RAC Setup

We start with the issue of RAC and VMware compatibility. We are told that most database servers are overpowered in terms of CPU because database performance is limited by I/O. Therefore, these database servers have underutilized CPU capacity, and virtualization will allow better use of these servers’ CPUs. I’m curious about whether most database servers are underutilizing their CPUs by more than 20%. If virtualization takes 10–20% overhead (as described earlier), then we need to be sure that our RAC database servers have more than 20% CPU available before we can be sure that RAC on VMware will indeed provide more CPU resources.

This results in a figure that shows nodes of multiple RAC installations installed on a single physical server. Each virtual machine on the server is one node of a RAC installation. I would like to see a discussion of how this would all play out as the various RAC nodes failover. What would happen to performance if all the RAC clusters failed over to the same physical server?

The process of setting up virtualized servers for a true RAC setup is described, and this is characterized as a “complex technical undertaking.” One of the more complex aspects is the host’s network setup, in which you need to establish public, private, and storage networks. Clearly I don’t have the necessary background, as I couldn’t understand the sections

describing virtual switches and NIC Teaming. Again, virtualization will require more technical resources.

It's clear that the process of setting up RAC on virtualized servers is complicated. We are told, again, that if (when?) there are performance issues, the best approach is simply to move the entire virtualized database to another physical server to see if performance improves. I don't see moving a complete RAC setup from one set of virtualized servers to another as an easy task. Do you have the resources to be ready at any time to perform such a move? Further, who will know exactly which node of which RAC database is running on which physical server? Who will know which other physical servers have enough spare capacity to support such a move? How many "extra" physical servers do you keep in reserve to support moves like this? Aren't spare servers, by definition, underutilized? How many such servers would represent enough underutilized capacity to offset the utilization gains promised by virtualization? I assume some of my concerns would be addressed by cloud computing, where you assume a third party will have spare resources at all times. But then you are relying on this third party to always have the spare servers you may need, along with the ability to quickly move your database between servers.

Chapter 8—Performance Comparisons

This chapter tells us that the performance of a database on a virtualized host is similar to that of a non-virtualized host. Most of the chapter is spent describing the benchmarking process; it covers some benchmarking myths, best practices, and how to select a benchmark. The results of all this show that while virtualization does add some overhead, the database performance is only slightly reduced due to virtualization.

Chapter 9—Additional Tuning Thoughts

While virtualization adds another layer of abstraction, the material here tells us that most of what you already know about database performance tuning still applies. Tuning methodologies are broken down into two basic approaches, measuring and reducing ratios, and measuring and reducing response time. While the author prefers the response time method, a third way is also described wherein a consultant can ask simple questions and make simple changes to improve performance in a single day. Several additional tuning issues that are specific to virtualization are then discussed.

One of the specific issues addressed is very interesting. We are told to examine the workload characteristics of each of the virtualized databases running on the physical server to see if they are compatible. First, how will you determine this? A single physical server could be hosting many virtual databases,

and those databases may have changing workloads. Second, if you have to know the specifics of each database on each physical server, doesn't this contradict the flexibility that virtualization is offering? How can you simply apply additional (or reassign existing) resources when you must first determine exactly what is using the resources? What happens when you use cloud computing where, in theory, you have no idea what is running, which server it is running on, etc?

Chapter 10—Oracle Virtual Appliances

Here we are told that since virtualization is quickly becoming mainstream, it will soon lead to other innovations. One of these is software delivery. Instead of installing new software, you will simply copy the files for the entire virtual machine and start the application. Similarly, vendors that deliver large data sets to customers could simply supply a complete database ready to run in your virtualized environment. Other possibilities would affect training and education. The contents of the DVD, which include a virtual machine with Linux and Oracle 10g database already installed, are described in detail at the end of the chapter.

Conclusion

This book is worth reading to gain an understanding of what virtualization will require of your organization. There is a great deal of detail that isn't covered, but for someone new to the subject, the amount of information in this book is about right. I would like to have seen some discussion of support and maintenance of a real-world virtualized environment. I wonder how it will really work when you need to apply patches to the hypervisor, which supports multiple virtualized databases. When the time comes to troubleshoot problems, I wonder who will know all the details of what runs where, how each physical server and each virtual machine are configured, and what to change. I am also curious to see a detailed study of the overall reduction of unused server capacity that results after the overhead of virtualization and the need for extra resources to handle future performance needs are taken into account. ▲

Brian Hitchcock has worked at Sun Microsystems in Newark, California, for the past 11 years. He is a member of a DBA team that supports 2400+ databases for many different applications at Sun. He frequently handles issues involving tuning, character sets, and Oracle applications. Other interests include Formula One racing, finishing his second Tiffany Wisteria lamp, Springbok puzzles, Märklin model trains, Corel Painter 8, and watching TV (TiVo rules!). Previous book reviews by Brian and his contact information are available at www.brianhitchcock.net.

Copyright © 2009, Brian Hitchcock

“We are told that since virtualization is quickly becoming mainstream, it will soon lead to other innovations. One of these is software delivery. Instead of installing new software, you will simply copy the files for the entire virtual machine and start the application. Similarly, vendors that deliver large data sets to customers could simply supply a complete database ready to run in your virtualized environment.”