

Maximizing Oracle Database Performance and Minimizing Licensing Costs in Virtualized Environments

Embarcadero Technologies, Inc.

October 2010

Americas Headquarters

100 California Street, 12th Floor
San Francisco, California 94111

EMEA Headquarters

York House
18 York Road
Maidenhead, Berkshire
SL6 1SF, United Kingdom

Asia-Pacific Headquarters

L7. 313 La Trobe Street
Melbourne VIC 3000
Australia

TABLE OF CONTENTS

Introduction	- 2 -
Q&A with Dave Welch and Scott Walz	- 3 -
VMware vSphere vs. Oracle VM	- 3 -
Reducing the Cost of Oracle Licensing.....	- 3 -
Virtualized Data Centers	- 5 -
Operating Systems and Enterprise Applications.....	- 6 -
Oracle Support for VMware.....	- 6 -
Oracle Performance on VMware.....	- 8 -
VMFS v. RDM	- 9 -
Virtual Server and RAC.....	- 9 -
Cultural Issues	- 10 -
Ensuring that peak performance is maintained post implementation.....	- 10 -
Tools to Maximize Oracle Database Performance and Availability in Virtualized Environments	- 12 -
About Us	- 13 -
About Embarcadero	- 13 -
About House of Brick	- 13 -
Additional Resources	- 14 -
Vendor Information	- 14 -
Independent research	- 14 -
Blogs and Articles	- 14 -

INTRODUCTION

On Sept. 30, 2010, Embarcadero Technologies and House of Brick Technologies hosted a Webcast on the topic of Oracle Database Performance and Licensing Dynamics in Virtualized Environments. As part of the program, the panel captured questions from registrants in advance and during the presentation, and shared valuable insights.

Among the panelists were Scott Walz, Senior Director of Product Management at Embarcadero as well as Dave Welch, founding partner and CTO of House of Brick. Welch is an Oracle Certified DBA, a former Oracle University RAC instructor, and a VMware Certified Professional. He started with Oracle technology as the lead production DBA at Intermountain Health Care in 1994, dealing with Oracle Parallel Server (RAC's predecessor). He specializes in Oracle enterprise license assessment, enterprise server assessment, business continuity options, performance, scalability, and system architecture.

What follows are highlights from the Webcast—including an extensive Q&A with Welch and Walz—as well as supporting references offered by Embarcadero to help DBAs maximize Oracle Database performance and minimize licensing costs in virtualized environments.

Q&A WITH DAVE WELCH AND SCOTT WALZ

In this Q&A from the Webcast, Welch and Walz discuss a variety of issues including virtualization platforms, licensing, performance, storage and support.

ON VMWARE VSPHERE VS. ORACLE VM

First, some background. Embarcadero's marketing of this webcast coincided with Oracle's hard launch of Oracle VM which included a series of events and announcements. As a result, we got many questions from registrants asking us about the differences between VMware vSphere and Oracle VM. For the reader not yet familiar with Oracle's marketing efforts on this subject, we have included those references here below. The inclusion of the Oracle links does not constitute an endorsement of their content by either Embarcadero or House of Brick.

Question: Why should we choose VMware's vSphere over Oracle VM?

Welch: We could take quite a while on this topic alone. I'll mention just three reasons. One is it's the best tooling in the industry. Another is standardization; it's running under pretty much every other workload in your data center other than Oracle.

Oracle VM is based on Xen. Since part of the Linux kernel is copied into the Xen "parent partition", those device drivers cannot be performance-optimized for virtual performance only. A third reason is reduced HA. There are Linux patches unrelated to virtualization that still force a down of the physical host.

For more information on Oracle VM and Xen, see:

[Oracle VM](http://www.oracle.com/technetwork/server-storage/vm/overview/index.html) <http://www.oracle.com/technetwork/server-storage/vm/overview/index.html>

[Xen](http://www.xen.org/) <http://www.xen.org/>

ON REDUCING THE COST OF ORACLE LICENSING

Question: How can we cut Oracle licensing costs with virtualization?


Welch: Years ago I would have laughed if you had predicted that we'd ever encourage shops to run production and pre-production workloads on the same host. But VMware's resource separation for each of I/O, CPU, Memory, and Network is so good that there is no longer any technical incentive not to allow the workloads to co-mingle. Oracle processor-based licensing gives a very powerful incentive for the workloads to mingle on Oracle-licensed hosts. Historically, shops isolated production workloads on dedicated hardware sized for peak processing. On the other hand, today there are shops that happily run their Oracle-licensed hosts at around 70% CPU utilization sustained. That could represent over a three-fold improvement on a financial component that, for many shops, shadows the cost of the rest of the system stack.

Question: Can you address Oracle licensing dynamics and deltas within and across virtualization vendors?

Welch: Oracle continues to categorize VMware Infrastructure as “software virtualization” and therefore not eligible for licensing partitioning on the same host similar to RISC-based partitioning. Therefore, for shops using Oracle processor-based licensing, all of the host’s sockets must be licensed. Meanwhile about a year ago, Oracle announced that it was going to allow licensing of Oracle VM on only the subset of a host’s processors that the virtualization environment needs, through core pinning. But pinning has been technically possible with VMware Infrastructure for as long as we can remember. As we fail to see any relevant technical differentiator to cause the one platform to be eligible for licensing partitioning but not the other, this policy looks and smells like a pure revenue grab.

Question: What about Oracle licensing in a VMware ESX Cluster?

Welch: We have seen no contractual evidence to support the disturbingly-more-frequent assertion by the Oracle field that shops using processor-based licensing must license every node in the ESX cluster that Oracle *could* migrate to. It is only necessary to license hosts where the Oracle executables are “installed and/or running.”



For more information, see:

[Oracle Price List \(aka Software Investment Guide\) updated Nov 2009.](#) (Scan for “installed and/or running” and you’ll find the quote on page 15.)

[Oracle “Licensing Data Recovery Environments”, updated April 2010.](#) (Scan for “installed and/or run“.)

Question: Can you explain Oracle licensing with HA and/or Failover in vSphere?

Welch: The current published version of the Oracle “10-day rule” specifies that you can failover up to 10 calendar days a year to an unlicensed node as long as it meets all these conditions: It’s a named fail-over node; the node shares a cluster file system with the licensed original production node; and you fail-back.

ON VIRTUALIZED DATA CENTERS

Question: Is a virtualized data center the same as a “private cloud”?

Walz: When we speak of virtualized data centers, though it can encompass moving systems to a cloud, we focus on the combining of several physical servers—and the applications running on those servers—into a subset of physical servers running multiple versions of virtualized operating systems.

Question: For the DBA, how is a virtualized data center different than a “traditional” data center?

Walz: DBAs are used to having their database occupy a single physical machine. With that one-to-one relationship, a DBA can be assured that their system will have top billing on the machine. In a virtualized data center, their database may be one of many, thus introducing another X factor into the performance mix.

Virtualization of database environments is increasing. According to a Web-based survey of 381 members of the Independent Oracle Users Group (IOUG), published by Unisphere Research in January 2010, a majority of respondents said virtualization in database production environments is on the rise. Those that are embracing high levels of virtualization are more likely to be expanding their production database environments, according to the study.

Question: How do you properly determine resource allocation when designing a data center?

Welch: It would come as no surprise that VMware probably has the world’s best tools for capacity planning a virtual data center. As for the configuration of the hosts, the most important issue is to consider the ramifications of one of your more expensive data center components—your Oracle license if that is processor-based—when sizing your hardware. For some smaller enterprises, that could even mean introducing some single-socket nodes in your data center.

ON OPERATING SYSTEMS AND ENTERPRISE APPLICATIONS

Question: What operating systems do you recommend for Oracle database servers in a VMware ESX environment?

Welch: Only SuSE Enterprise Linux and Red Hat Enterprise Linux. We were enthusiastic about Oracle Enterprise Linux from its announcement as it was downstream from Red Hat and had a very nice support price point. A year ago, Oracle politicized the distribution and therefore I can no longer encourage its use.

For more information on Oracle and Linux, see Welch's blog at:

<http://www.houseofbrick.com/component/content/article/1/74>

But whatever you run, it has to be 64-bit. We imagine the only reason that VMware has allocated any fraction of its performance engineering budget to 32-bit operating systems in the last several years is to accommodate third-party application vendors that are dragging their feet on getting their software tested on 64-bit. There are massive performance deltas between 32-bit and 64-bit guests running in VMware.

ON ORACLE SUPPORT FOR VMWARE

Question: What is the real story on Oracle support for VMware?

Welch: We've never heard of a situation where Oracle Support didn't give all due diligence to a Support Request, including when the VMware customer was fully-disclosed with Oracle, and even on rare occasion when the issue isn't already in Oracle's knowledge base. Probably the best reference I can give is Indiana University (see below). I can't speak for Oracle, but all indicators are that support note 249212.1's reference to reproducing problems on native hardware is there as a legal hedge, and is not dissimilar to hedges by VMware's good partners HP & Dell. We encourage you to ask your Oracle Support rep as needs arise, rather than ask your Oracle field rep, who has a huge financial incentive to not see you running your Oracle processor-based license around 70% CPU utilization sustained.

Question: Can we virtualize Oracle on VMware when Oracle does not certify their products on it?

Welch: Historically, Oracle doesn't certify hardware platforms. They certify operating system distributions. Years ago, Indiana University's legal team looked at that, and told the Enterprise Infrastructure division, "proceed." They didn't ask Oracle in advance. VMware is hardware to Oracle.

“With the help of House of Brick and VMware Professional Services, we have been able to virtualize our most demanding Oracle Databases on x86 servers. Where initially we believed that these databases would be too demanding for a virtual machine, we now have the confidence that vSphere can handle our largest transaction-processing databases with ease.”—Rob Lowden, Director of IT at Indiana University.

Learn how Indiana University’s IT department, with the help of VMware and House of Brick, increased flexibility and reduced costs for large, business-critical Oracle databases, at:

<http://www.vmware.com/solutions/partners/alliances/oracle-database-customers.html>

Question: Are you aware of any problems with Oracle Support on ESX?

Welch: Again, we’ve never seen or even heard of a shop that’s had problems with Oracle Support, including when they are fully-disclosed with Oracle, and including on occasion when the issue is nowhere to be found in Oracle’s knowledge base. House of Brick has been a lightning rod for issues surrounding Tier-1 Oracle workloads on VMware. If there were technical issues, we would have heard about them.

Question: Any issue of running Oracle Enterprise Linux, E-business suite and database on VMware VM?

Welch: I’ll limit my comments to E-Business Suite on VMware. In Q1 of 2006, a food giant with thousands of retail outlets went live with all tiers virtual of their EBS 11.5.10.2 stack that we had installed just months before. They had great performance out of the box on an ESX Server release whose performance characteristics VMware now likes to run down. The shop has been fully-disclosed with Oracle Support for years. They have as many Oracle support requests as other comparable organizations. And of course none of them has ever turned out to be remotely VMware-related.

For more on Oracle Support, see House of Brick’s Oracle-on-VMware support position paper at: www.houseofbrick.com/oracle-on-vmware.

ON ORACLE PERFORMANCE ON VMWARE

Question: Can the virtualization platform handle the I/O demands of OLTP/Data Warehouse Systems?

Welch: We hit performance nirvana three years ago when ESX Server 3.5 went GA. I/O is the ultimate frontier of performance. The ESX 3.5 I/O scalability trial that VMware ran in May of 2008 was no surprise to us. A single 16-core ESX host was loaded with 500 spindles in three Clarion cabinets configured with some intentionally-ridiculous I/O burdens. It scaled linearly to over 100K IOPS with marvelous average I/O times. The I/O virtualization latency wedge was a whopping 1.4%. And that was three million development hours ago in ESX GA releases. Today, a 10ms I/O spends a mere 100 *micro*-seconds in the vSphere kernel. If you don't get great Tier-1 performance out of the box, you need to assume that your issues are configurable and not inherent the VMware Infrastructure. That's been the case in 100% of the shops we've triaged, and we've been doing this for years.

Question: Are there any specific configuration parameters that should be considered when running on VM?

Welch: The two big things that differ from native hardware Oracle implementations are to give the virtual machines no more memory or CPU than they truly need. It's not a big deal if you do, but either can induce unnecessary modest latencies.

Question: Do you have to put the database in backup mode, in order to rely on VM snapshots for recovery?

Welch: Yes, should you use VMware snapshots as a backup method in non-production environments, and then you can either put the database in backup mode, or Alter System Quiesce before the VMware snapshot. Warning: don't leave snapshots hanging around after production environments' maintenance windows as there can be a significant performance price to be paid for doing so.

Question: What about Oracle licensing and how to prevent SCSI queue saturation?

Welch: The bane of VMware proof-of-concept Tier-1 storage configurations would appear to be the fantastic server consolidation ratios that tier-2 and tier-3 workloads have enjoyed on VMware for years. That success came primarily from optimizing and consolidating CPU and memory. Many organizations fall victim to the temptation to assume to consolidate their storage in a similar or worse fashion when they experiment with Tier-1 virtualization. We frequently encounter a bunch of workloads including some tier-1 I/O workloads fanned into LUNs (or Data Stores) in a fashion that no one would dream of doing in a native hardware environment. That problem is frequently exasperated with apples-to-oranges comparisons; in other words, when the native storage for example is RAID 1+0 with a bunch of cache and spindles, yet the VMware proof-of-concept environment is RAID-5 with far less cache and far fewer spindles. Each hardware device in the storage path has a SCSI queue. And when workloads are fanned in too heavily through single storage paths, those SCSI queues can saturate. Often at this point, folks say, "we tried Oracle on VMware but the performance wasn't good, so we backed out." But they haven't done the analysis to determine that the problem had nothing at all to do with anything inherent in VMware Infrastructure. For Tier-1 workloads, we recommend a 1:1:1 ratio of LUN/Datastore to VMFS to VMDK to workload, until such time that you either get benchmarking done or have a couple months of production metrics under your belt. At that point, you may discover that your I/O spreadsheet shows that you can safely fan in more workloads to a storage path without compromising performance.

ON VMFS v. RDM

Question: Which should we choose, VMFS or RDM?

Welch: Absolutely VMFS, per the strong recommendation of VMware, EMC and its competitors, and House of Brick. Don't orphan yourself from the mature storage tooling capabilities by having the virtual machine directly mounted storage. The *only* time we would want you using Raw Device Mapping is with RAC-on-VMware for both technical and Oracle Support reasons, which are beyond the scope of today's discussion. We strongly encourage Oracle shops to run Automatic Storage Management (ASM). In a VMware environment, that means running all your storage through ESX and the VMFS file system, then carving it up within ESX as raw slices that ASM then mounts.

ON VIRTUAL SERVER AND RAC

Question: What is the benefit of using virtual server and RAC combined?

Welch: If RAC is appropriate for the workload at all, then yes, it absolutely makes sense to layer RAC on VMware. That enables RAC cloning at the click of a button back to the earliest stages of the product development lifecycle. That radically optimizes the product development lifecycle and reduces product investment risk. That holds true whether or not RAC is virtualized in production. We did our first production-grade implementation of RAC-on-VMware in May of 2007.

Question: Is RAC support from Oracle under VMware coming soon?

Welch: Oracle has been perennially hesitant to make any public statements as to what's behind this statement in the published RAC FAQ: "There are technical restrictions that prohibit the support of RAC in a VMware environment." Anecdotally, we're certain the primary issue has been Linux Guest clock-drift. Read up on that by searching "clock drift" in the VMware Knowledge Base. It can happen on older Linux kernels when the host CPU heats up to 100%. Needless to say, RAC can be uncomfortable when various instances' clocks get out of sync. So I wouldn't blame Oracle Support for having been about this. But those days are gone. I designed and we executed a ground-breaking RAC-on-VMware Break/Fix lab earlier this year, which I presented at IOUG in April. The summary is that our testing concludes that clock drift is a thing of the past on vSphere, even with the old virtualization-illiterate 2.6.9 kernel. We were also unable to elicit any RAC node evictions or even eviction warnings under the most ridiculously-hostile, CPU-starved conditions.

Read more about vSphere's apparent elimination of the clock drift issue at:

<http://www.houseofbrick.com/oracle-on-vmware>

According to our proofing, any technical concerns about RAC-on-VMware have been paper tigers for almost two years now. So we would envision a situation where Oracle is pretty much cornered on the RAC support restriction, especially since they have announced support for RAC on what we consider to be their very less capable Oracle Virtual Machine. I'll mention here that House of Brick has multiple methods for complying with the lingering RAC support prohibition in production environments.

Question: How can I guarantee a smooth migration of distributed databases to a centralized database on Exadata?

Welch: Let me preface by saying there are some things about Oracle's Exadata II that I really like, as well as other things that I really don't. If House of Brick were to answer this question, we'd tell you "Don't." I'll mention one issue only: Exadata hardware appears to be 100% compatible with what we hold is the world's premier platform, VMware Infrastructure. But Exadata requires more Oracle Support involvement as a percentage of overall operations compared to non- Exadata stacks. And Oracle currently appears to be politically hostile to VMware Infrastructure.

ON CULTURAL ISSUES

Question: How can I get my operations people to stop predicting doom if we virtualize Oracle workloads?

Welch: Start by virtualizing pre-production Oracle workloads and get them hooked on the operational agility and benefits.

Question: How do you prevent people from feeling a loss of control or worrying about performance of their essential systems?

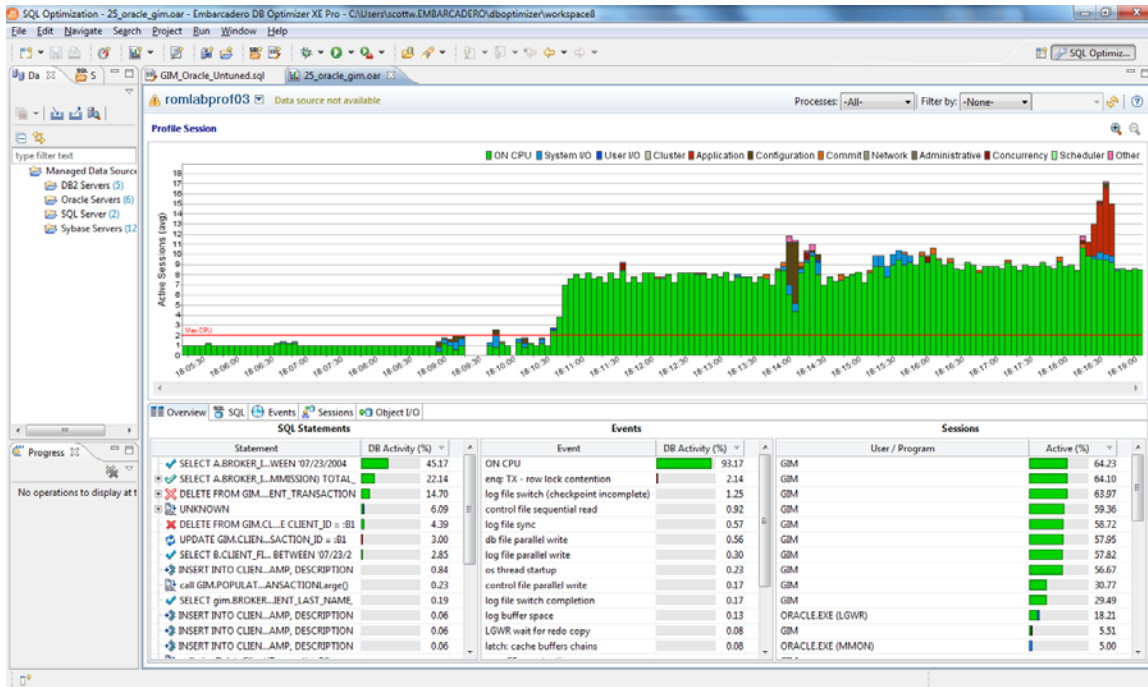
Walz: Provided the applications are performing at pre-virtualization levels, there usually are no issues. At the end of the day, the owners of those systems expect top performance, and provided the performance is meeting the agreed upon SLAs, the owners are usually fine. If performance issues do arise, the normal course of action will be taken, and the troubleshooting will begin. The introduction of virtualization will require the DBA to identify if the trouble lies within the database—think missing index or poor performing SQL code—or if it's an issue of the database hitting the upper limits of the allocated CPU.

ON ENSURING THAT PEAK PERFORMANCE IS MAINTAINED POST IMPLEMENTATION

Question: How do you set up clients so that they can be successful in monitoring their databases in virtualized environments?


Welch: The first thing we do is make sure they have a firm understanding at a concept level of IO, CPU, memory, and network resources. Then we show them how these performance metrics are represented at the command line level through OS-bundled utilities like vmstat, iostat, SAN-provided utilities, or whatever they actually use. That way, eventually when this data gets pulled into a GUI tool like Embarcadero DB Optimizer or DBArtisan, our client has a clear understanding of what they are looking at and can really leverage the ROI of the tool.

Walz: Let me show you what these concepts might look like in tools like DB Optimizer and DBArtisan. An engineer can read all the text reports. But when he wants to show it to his manager or the people who make the big decisions with hardware, it's super powerful to have something they can immediately understand.



Customers are able to dig in there for the command prompt, and take advantage of firms like House of Brick Technologies that can provide the excellent instruction to get you where you need to be with the command prompt. They can get to that information, but to be able to graphically and visually show this, to be able to take it to somebody with authority and to be able to plead your case is really powerful.

It's important, and probably just as important if not more important in a virtual environment to be able to pinpoint those performance bottlenecks. And from a real-time perspective, DBArtisan can help you do that. When you combine a tool like DBArtisan with DB Optimizer, you've got an unbelievably powerful combination.



TOOLS TO HELP MAXIMIZE ORACLE DATABASE PERFORMANCE AND AVAILABILITY IN VIRTUALIZED ENVIRONMENTS

As organizations move their Oracle databases into virtualized environments, they need better tools to help ensure high performance. Embarcadero offers solutions that can help companies maintain the performance they need.

When it comes to Oracle performance issues, Embarcadero tools know where to look, including harnessing Oracle's wait interface, and high speed sampling of session state and activity, to render various views of Instance and Session based performance metrics.

Oracle DBAs and DBA Managers can use Embarcadero's tools to deliver the highest possible levels of performance and reliability, allowing you to:

- See detailed metrics that identify performance, capacity planning, and storage management issues before they become major problems
- Automate 24/7 monitoring of the health index of all your databases, alerting you to any performance issues in real-time
- Eliminate performance bottlenecks by visually profiling key metrics inside the database (CPU, I/O, wait times), relating resource utilization to specific queries, and helping you visually tune problematic SQL

Embarcadero database performance management and SQL tuning products for the Oracle data management professional:

- [DBArtisan for Oracle and DBArtisan XE](#) – the leading database administration tool among Oracle DBAs includes powerful performance management capabilities
- [DB Optimizer for Oracle and DB Optimizer XE](#) – prevent, find and fix performance issues via visual SQL profiling and tuning. Enterprise edition includes [DB Performance Center](#) for 24x7 monitoring, customized alert thresholds and notifications, and a unified view of the health of all your databases

While Embarcadero's products are known for providing deep, native support for Oracle, they also provide in-depth support for Sybase ASE, Microsoft SQL Server, DB2, MySQL (and more). With Embarcadero's new XE line of tools, support for all major database platforms is included in a single, affordable license, so you can extend your Oracle skills into less familiar databases while enabling your entire organization to standardize on a common, powerful solution.

ABOUT Us



ABOUT EMBARCADERO

Since 1993, Embarcadero has been the trusted tooling provider to thousands of Oracle customers. Explore our [Oracle solutions page](#) or [download a free trial](#). You will quickly discover why experienced data management professionals place their confidence in the depth of Embarcadero's Oracle-specific capabilities to get their jobs done quickly.

Embarcadero Technologies, Inc. is the leading provider of software tools that empower application developers and data management professionals to design, build, and run applications and databases more efficiently in heterogeneous IT environments. Over 90 of the Fortune 100 and an active community of more than three million users worldwide rely on Embarcadero's award-winning products to optimize costs, streamline compliance, and accelerate development and innovation. Founded in 1993, Embarcadero is headquartered in San Francisco with offices located around the world. Embarcadero is online at www.embarcadero.com.



ABOUT HOUSE OF BRICK

House of Brick Technologies offers complete IT solutions and consulting to a variety of large and small businesses. Since 1998, the Omaha, Nebraska firm has delivered the innovative tools, expert advice and comprehensive mentoring necessary to maximize clients' capabilities in areas ranging from scalable architecture, project planning and component integration to backup, disaster recovery and managed services. The company, which derives its name from the indestructible house in the "Three Little Pigs" children's story, has served more than 100 clients throughout the United States, as well as Asia, Australia, Europe, North America and South America. Contact us at (402) 445-0764 or visit our Web site, www.houseofbrick.com.

ADDITIONAL RESOURCES

VENDOR INFORMATION

Learn more about tools for maximizing database performance at Embarcadero's site:

<http://www.embarcadero.com/>

Read blogs and other content on database and virtualization technology at House of Brick Technologies' site: <http://www.houseofbrick.com/>

Information about Oracle on VMware, including white papers, customer success stories and product information:

<http://www.vmware.com/oracle>

<http://vmware.com/solutions/partners/alliances/oracle-database.html>

Information about Oracle's support for VMware environments:

<http://www.vmware.com/go/oracle-support>

VMware vSphere Essential Database Deployment Tips:

<http://vmware.com/resources/techresources/10101>

INDEPENDENT RESEARCH

Here's what Gartner has to say about Oracle and virtualization, and a link to its report: Many organizations are using virtualization software, such as VMware, with their Oracle database management system and infrastructure licenses. Oracle has created policies for how it counts usage under virtualization, but these policies are not a formal part of the Oracle license agreement.

http://www.gartner.com/DisplayDocument?doc_cd=165003

BLOGS AND ARTICLES

Data Access and Virtualization:

<http://blogs.datadirect.com/2010/01/data-access-and-virtualization.html>

Chris Wolfe's Virtualization Tips and Ramblings:

<http://www.chriswolf.com/>

Computerworld article: Oracle Virtualization 'Tax' Squeezes Customers:

http://www.computerworld.com/s/article/9133623/Oracle_Virtualization_Tax_Squeezes_Customers