



Official Publication of the Northern California Oracle Users Group

NoCOUG

J O U R N A L

Vol. 31, No. 3 • AUGUST 2017

Licensing Demystified

*House of Brick solves the puzzle.
See page 4.*

Oracle Cloud Revealed

*Brian Hitchcock takes notes.
See page 8.*

Parallel RMAN Backups with Standard Edition

*Easy-Peasy says Norbert Debes.
See page 20.*

Much more inside . . .

The New Phoenix by Axxana

For Oracle Databases and Applications

Zero Data Loss
at Unlimited Distances

Synchronous Protection
at Maximum Performance

- Low Cost Replication Lines
- Shortest Recovery Time
- Full Consistency Across Multiple Databases



AXXANA
B U I L T T O L A S T

www.axxana.com



ORACLE Gold Partner

ORACLE
NO DATA
APPLIANCE
READY

ORACLE
DATABASE
READY

ORACLE
EXADATA
READY

ORACLE
LINUX
READY

Professionals at Work

First there are the IT professionals who write for the *Journal*. A very special mention goes to Brian Hitchcock, who has written dozens of book reviews over a 12-year period. The professional pictures on the front cover are supplied by Photos.com.

Next, the *Journal* is professionally copyedited and proofread by veteran copy-editor Karen Mead of Creative Solutions. Karen polishes phrasing and calls out misused words (such as “reminiscences” instead of “reminisces”). She dots every i, crosses every t, checks every quote, and verifies every URL.

Then, the *Journal* is expertly designed by graphics duo Kenneth Lockerbie and Richard Repas of San Francisco-based Giraffex.

And, finally, David Gonzalez at Layton Printing Services deftly brings the *Journal* to life on an offset printer.

This is the 123rd issue of the *NoCOUG Journal*. Enjoy! ▲

—NoCOUG Journal Editor

Table of Contents

Interview	4	ADVERTISERS	
Book Notes	8	Axxana	2
Book Excerpt	12	Quest	22
Product Review	20	ORADBPRO	26
Session Descriptions	24	OraPub	28

Publication Notices and Submission Format

The *NoCOUG Journal* is published four times a year by the Northern California Oracle Users Group (NoCOUG) approximately two weeks prior to the quarterly educational conferences.

Please send your questions, feedback, and submissions to the *NoCOUG Journal* editor at journal@nocoug.org.

The submission deadline for each issue is eight weeks prior to the quarterly conference. Article submissions should be made in Microsoft Word format via email.

Copyright © by the Northern California Oracle Users Group except where otherwise indicated.

NoCOUG does not warrant the NoCOUG Journal to be error-free.

2017 NoCOUG Board

Craig Shallahamer
Director of Vendor Relations

Dan Grant
Member at Large

Eric Hutchinson
Webmaster

Hanan Hit
President Emeritus

Iggy Fernandez
President

Jeff Mahe
Vice-President

Kamran Rassouli
Social Director

Liqun Sun
NoCOUG Ambassador

Michael Cunningham
Training Director

Naren Nagtode
President Emeritus

Noelle Stimely
Membership Director

Roy Prowell
Publicity Director

Saibabu Devabhaktuni
Conference Chair

Tim Gorman
Board Advisor

Tu Le
Speaker Coordinator

Brian Hitchcock
Book Reviewer

ADVERTISING RATES

The *NoCOUG Journal* is published quarterly.

Size	Per Issue	Per Year
Quarter Page	\$125	\$400
Half Page	\$250	\$800
Full Page	\$500	\$1,600
Inside Cover	\$750	\$2,400

Personnel recruitment ads are not accepted.

journal@nocoug.org



Nathan Biggs

Licensing Oracle Software in Cloud Environments

by Nathan Biggs and Pamela Fulmer



Pamela Fulmer

Editor's Note: This article contains information on Oracle licensing that is provided as-is and without guarantee of applicability or accuracy. Given the complex nature of Oracle licensing and the ease with which license compliance risk factors can change significantly due to individual circumstances, readers are advised to obtain legal and/or expert licensing advice independently before performing any actions based on the information provided.

Nathan: I was pleased when invited by the Northern California Oracle Users Group to write an article addressing some of the complexities of licensing Oracle in virtualized and cloud environments. My first thought was that the content would be much more rich and beneficial if I were able to leverage the experience of one of our expert legal partners. Pam Fulmer, who is a partner at Fulmer Ware LLP based in San Francisco, kindly agreed to participate in this discussion. House of Brick has consulted with Pam and her firm many times when helping customers with the sometimes difficult task of defending themselves in audit and/or license negotiation situations with Oracle.

This article is intended to build on the material already provided by Mohammad Inamullah in the February 2017 and May 2017 issues of the *NoCOUG Journal*. We will try to explore the topic more fully by taking a deeper look into the intricacies of Oracle licensing in virtualized and cloud environments. Pam and I will be taking a roundtable discussion approach to our responses. I hope that this format proves beneficial to the reader.

There are so many things being said about Oracle licensing, especially if I am using virtualization or running Oracle in the cloud, that it is hard to keep track of all of the changes to the Oracle contracts. Is there an easy way to know what is in the contract and how that dictates what I should do?

Nathan: The short answer is that Oracle has made very few changes to the actual binding contracts for running in virtualized or cloud environments. This is contrary to the perception that things are changing all the time. Essentially, there are only a few key points that you need to understand about your binding contract (whether it is a newer OMA, an older OLSA, or an even older SLSA).

Somewhere in your agreement is the following definition of the term Processor: "Processor shall be defined as all processors where the Oracle programs are installed and/or running." This is the most fundamental and important definition in your agreement. "Installed" is a past-tense activity that presently applies.

"Oracle has made very few changes to the actual binding contracts for running in virtualized or cloud environments. This is contrary to the perception that things are changing all the time."

"Running" is a present-tense activity. There is nothing indicating a license requirement for prospective activities (things that might happen in the future), no matter how easy they may be to bring about.

Another important item in all Oracle agreements is the Entire Agreement clause. This states that only items cited in the agreement—or specifically referenced by it—are binding. Nothing else—verbal representations, policy documents, sales proposals, etc.—can be considered contractual.

Pam: As Nathan points out, the Oracle agreements are fully integrated, as defined by the Entire Agreement clause. Under California law, which is the law that applies to most Oracle licensing agreements in the U.S., when the parties to a written contract have agreed to it as an "integration"—a complete and final embodiment of the terms of their agreement—evidence of prior or contemporaneous negotiations and agreements that contradict, modify, or vary the contractual terms cannot be admitted. A court will look instead to the plain meaning of the words within the four corners of the contract. Under California law, the language of a contract is to govern its interpretation, if the language is clear and explicit, and does not involve an absurdity. As

"Somewhere in your agreement is the following definition of the term Processor: Processor shall be defined as all processors where the Oracle programs are installed and/or running. This is the most fundamental and important definition in your agreement."

“Installed is a past-tense activity that presently applies. Running is a present-tense activity. There is nothing indicating a license requirement for prospective activities (things that might happen in the future), no matter how easy they may be to bring about.”

Nathan says, “installed” is past tense and “running” is present tense. These words are generally clear and unambiguous, and do not involve future events that are purely speculative and may never happen. To argue otherwise would lead to an absurd result: forcing a customer to pay Oracle a royalty for software that was never used or installed.

One more potential absurdity concerning Oracle’s definition of “Processor” is that it is impossible to actually install Oracle software on a processor chip. Software is installed on computer systems, of which processors are a component. Processors only run the software, so the word “installed” does not make sense in this context. The thing that is clear in the definition is that the software is actually “running” (present tense activity) on processors in a computer system. It is well settled under California law that “language involving an absurdity is rejected . . .” by the courts. *Eucasia Schools Worldwide, Inc. v. DW August Co.* (2013) 218 Cal.App.4th 176, 182. Moreover, to the extent the inclusion of the word “installed” in the processor definition creates an ambiguity, under California law, that ambiguity should be construed against the drafter of the contract. In this case, the drafter of the contract is Oracle.

Nathan: Another area that causes confusion with many Oracle customers is the policy documents that Oracle publishes. Most of these documents (Partitioning Policy, Licensing Oracle Software in the Cloud Computing Environment, Licensing Data Recovery Guide, etc.) are not referenced by the agreement and are thus not binding in your contract with Oracle. The Partitioning Policy document is frequently cited by Oracle to customers running on VMware. Just remember that this document does not contain binding policy.

There are some non-contractual documents, however, such as the Licensing Oracle Software in the Cloud Computing Environment (Cloud Environment) policy from Oracle, that are fundamentally different. In this particular document, Oracle is granting additional privileges beyond the contract, rather than restricting them. In a blog post that I published on the House of Brick website last November, I stated the following:

A final consideration is that the Licensing Oracle Software in the Cloud Computing Environment policy document is a non-contractual reference. It is excluded by the Entire Agreement clause of your master agreement with Oracle. Even though it is not binding through the agreement, it has been the recommendation of House of Brick (to the legal teams of our customers) that this is a published policy, albeit non-contractual, that grants additional privileges. This is in contrast to a non-contractual document like the Partitioning Policy document that attempts to restrict privileges in a non-contractual way. (<http://houseofbrick.com/running-oracle-in-a-public-or-hybrid-cloud/>)

Pam: The Entire Agreement clause excludes Oracle’s ability to rely on documents that are not a part of the integrated license agreement to try to vary, restrict, or contradict the terms of the

license agreement. That clause also typically provides that the Agreement cannot be amended without a written amendment, submitted online through the Oracle store and signed by an authorized representative of the customer and Oracle. So an amendment to the license agreement requires that both parties indicate their consent to the amendment in writing. However, as Nathan points out above, nothing prevents Oracle from issuing policies outside of the contract that provide the customer with additional rights or privileges. Reliance on such documents may be risky, however, as Oracle expressly points out in the Licensing Oracle Software in the Cloud Computing Environment policy that it is non-binding and subject to change at any time. However, to the extent that Oracle is knowingly publishing extra-contractual documents on which its customers rely by making large investments, an argument can be made that Oracle should be estopped or prevented from changing course down the road, especially if such a change would cause injury to Oracle customers. Whether a court would accept this argument, or find that the customer proceeded at their own risk, is an open question.

Given the above definition of Processor, are virtual processors used in public cloud environments also considered Processors for Oracle licensing?

Nathan: Oracle’s binding contracts (not the non-binding policy documents) are completely devoid of any notion of virtualization or virtualized processors. We consider that the contracts refer

“The Licensing Oracle Software in the Cloud Computing Environment policy document is a non-contractual reference. It is excluded by the Entire Agreement clause of your master agreement with Oracle.”

specifically to hardware-based processors. This means that Oracle’s contracts do not allow for you to only count virtual processors in a cloud environment. The one consideration that varies from this is the Cloud Environment policy we discussed previously, which grants virtual CPU privileges for AWS – EC2 and RDS, and Azure public clouds. If a cloud provider can offer Oracle contract-compliant and audit-compatible mechanisms for tracking all discrete physical processors where Oracle software is installed and/or running, then the customer may be able to consider applying their hardware-based licenses in these clouds.

Pam: When attempting to interpret the terms of a contract, a court will seek to understand what the parties mutually intended and understood the term to mean at the time of contracting. Oracle has been using the Processor definition for a long time,

which has traditionally meant hardware-based processors and not virtual. This has been what Oracle customers have understood as well. Also, as Nathan points out, no mention is made in the license agreement of virtualization or virtualized processors. As the drafter of the license agreement, if Oracle wanted to include virtualization, it could have done so. In fact, Oracle could come out with new definitions and terms in subsequent iterations of the license agreement that specifically reference virtual-

“If a cloud provider can offer Oracle contract-compliant and audit-compatible mechanisms for tracking all discrete physical processors where Oracle software is installed and/or running, then the customer may be able to consider applying their hardware-based licenses in these clouds.”

ization. However, Oracle has not done so. The fact that Oracle specifically mentions virtual processors in the Cloud Environment policy also supports the fact that the master agreement (OMA, OLSA, or SLA) is referring to physical and not virtual processors.

Do I need to count hyper-threads in my count of processors?

Nathan: Again, the binding agreements do not mention hyper-threading (the Cloud Environment policy does introduce the concept for AWS – EC2 and RDS). Therefore, we consider the definition to refer to physical processor cores only. When Oracle deviates from physical processors, it specifically mentions it in the contract language. An example of this is for Standard Edition 2 (SE2) licenses. In this case, a physical socket in a server is used in place of a processor. After this definitional variance, the rest of the “installed and/or running” language still applies.

Pam: The fact that Oracle identifies in the license agreement the places where it is deviating from physical processors is strong evidence that the definition refers to physical processor cores only. When Oracle wants to deviate, it knows how to do so. A court, when interpreting a contract, looks at the entire contract, attempting to have it make sense as a whole. When interpreting a contract, the intention of the parties should control, and the best evidence of intent is the contract itself. Under California law, the whole of a contract is to be taken together, so as to give effect to every part of the agreement, with each clause helping to interpret the others.

“AWS customers that use EC2 Dedicated Hosts may be able to choose the licensing model that works best for them. If those customers use dedicated hosts, then those hosts could be factored into the server inventory for Oracle licensing just as on-premises servers are. Since it is also on AWS EC2, then the Cloud Environment policy may also apply, and they could choose to license on a per-virtual CPU basis.”

The Cloud Computing Environment policy only applies to AWS and Azure. How should customers count processors in the Google Public Cloud and other cloud environments?

Nathan: Only AWS—EC2 and RDS, and Azure are granted per virtual CPU licensing privileges in the Cloud Environment policy document. Therefore, my answer here is based solely on that document and does not consider the possibility that any particular cloud provider has a special agreement with Oracle for additional privileges. Without the privilege for per virtual CPU licensing described in the Cloud Environment policy, every other cloud provider that wanted to allow customers to run Oracle software in their clouds would have to ensure that the customer could use standard hardware-based processor licensing instead. As an example, if a cloud provider could dedicate physical hosts to a customer and ensure that the Oracle workloads did not migrate off of those licensed hosts (or restrict live migration off of licensed hosts using means such as VMware’s DRS host affinity rules), then that provider could potentially qualify for hardware-based processor licensing. Overall, if a cloud provider can offer Oracle contract-compliant and audit-compatible mechanisms for tracking all discrete physical processors where Oracle software is installed and/or running, then the customer may be able to consider applying their hardware-based licenses in these clouds. Please carefully consider this with legal counsel and other experts before deploying Oracle in any cloud environment.

Are AWS customers who use EC2 Dedicated Hosts subject to the January 23 policy update, or can they use the on-premises licensing model?

Nathan: AWS customers that use EC2 Dedicated Hosts may be able to choose the licensing model that works best for them. If those customers use dedicated hosts, then those hosts could be factored into the server inventory for Oracle licensing just as on-premises servers are. Since it is also on AWS EC2, then the Cloud Environment policy may also apply, and they could choose to license on a per-virtual CPU basis.

One consideration in choosing a license model in this environment is whether an Unlimited License Agreement (ULA) from Oracle is also in play for certain Oracle products. In the Cloud Environment policy, Oracle states: “Licenses acquired under unlimited license agreements (ULAs) may be used in Authorized Cloud Environments, but customers may not include those licenses in the certification at the end of the ULA term.” For AWS customers who are using EC2 Dedicated Hosts and who also have a ULA from Oracle, it may be in their best interest to claim processor-based usage rather than cloud-based usage when certifying off of the ULA in order to maximize the number of licenses received in that process.

“The main point that I would like all NoCOUG readers to take from this discussion is that your contract with Oracle is what matters. You do not have to do something that is not in the best interest of your organization when it is outside of the scope of your binding agreement with Oracle.”

How does the January 23 Cloud Environment policy update affect VMware Cloud on AWS?

Nathan: That is an interesting question that may not be able to be fully answered at this time. Let's consider the possibility that VMware Cloud on AWS uses the dedicated host model described above (it appears from the public-facing information about this offering that this might be how it is architected). In this case, processor-based licensing and ULA entitlement could certainly be used. Then the question is whether VMware Cloud on AWS also falls under the Licensing Oracle Software in a Cloud Computing Environment policy. I discussed this in detail when the VMware Cloud on AWS announcement was made (<http://houseofbrick.com/running-oracle-in-a-public-or-hybrid-cloud/>), but the question to discuss with your legal advisors is the following: Does VMware Cloud on AWS run on “Amazon Web Services—Amazon Elastic Compute Cloud (EC2)”? This reference to EC2 is specifically cited by Oracle in their Cloud Environment policy. So, if Amazon makes a definitive statement that this is the case, then it appears that the Cloud Environment policy from Oracle, as stated in the January 23, 2017 version might apply. If not, and if the dedicated host model is not used, then you would need to evaluate options where Oracle workloads could be contained onto customer-licensed hosts in some fashion so that hardware-based licensing could be applied.

Can customers include cloud deployments in ULA certifications even though both the January 23 update and the previous version of the policy deny customers the ability to do this?

Nathan: Again, this is an area that you should discuss with your own legal advisors. The Cloud Environment policy document states that you cannot count virtual CPUs in your ULA certification. If you are claiming the additionally granted privilege from this non-contractual document as discussed previously, then this restriction would also seem to apply to you. Do not end your consideration there, however. You may have negotiated the specific privilege of counting cloud usage in the certification process into your ULA. You should have your attorneys look for additional ULA certification privileges for cloud usage in your own ULA agreement documents. If you are entering into a new ULA (admittedly not our recommendation for most customers), then it does not hurt to try to negotiate that privilege into the contract language before executing the agreement.

Closing Thoughts

Pam: As an attorney, I have retained House of Brick to assist me as my technical consultants on a number of Oracle-related matters. They are extremely knowledgeable, and I have been thoroughly impressed with their technical expertise and understanding of Oracle licensing.

Your contract with Oracle defines the legal rights of both parties. Although it can seem confounding at times, Oracle customers do have rights under the software license, and it is important to have a full understanding of these rights before going up

against Oracle. I recommend that any IT personnel receiving notice of an Oracle audit immediately seek help from your in-house legal team or outside counsel. And please get your lawyers involved before the audit begins, so that you are well advised of your legal rights under the terms of the contract.

Nathan: I appreciate Pam's input into this discussion. In our interactions with Pam, as we have worked together to help customers with their Oracle licenses and dealing with the sometimes uncomfortable task of interacting with Oracle sales, audit, or legal personnel, we have seen how her expertise and sensitivity bring clarity and resolution to the customers' worries.

The main point that I would like all NoCOUG readers to take from this discussion is that your contract with Oracle is what matters. You do not have to do something that is not in the best interest of your organization when it is outside of the scope of your binding agreement with Oracle. We strongly recommend, and can help customers with, a rigorous process of assessing Oracle license compliance, regular monitoring to ensure compliant operations, and preparation for the eventuality of an Oracle audit. As a best practice for our customers, we recommend a team approach to each of these things. This means that you should actively involve people from legal, purchasing, system and data architecture/administration, executive/business line management, and IT and operations management. As you work together, you will be united in eliminating risks and managing the complexity of Oracle licensing. ▲

Nathan Biggs is the CEO of House of Brick Technologies, a consulting firm focusing worldwide on virtualization and cloud computing for business/mission-critical systems. Nathan studied electrical engineering and computer science at Arizona State University, and Innovation and Entrepreneurship at Stanford University. In addition to being the CEO, Nathan frequently consults with customers on Oracle licensing issues, including defending against Oracle audits. He will be speaking at VMworld 2017 in the U.S., and Europe on the topics of Oracle licensing and audits.

Pam Fulmer is a partner at FulmerWare LLP, an IP and commercial litigation boutique located in San Francisco, California. Pam is admitted to practice law in California and has over 25 years of experience litigating all types of intellectual property and commercial disputes in California and across the United States. In addition to her litigation practice, Pam has a great deal of experience defending audits by software companies including Oracle Corporation, and has dealt with various issues involving Oracle software such as VMware virtualization, as well as hosting and related alleged areas of under-licensing. She has worked with her clients to develop strategies to mitigate these positions and to push back successfully on the audit findings.

NoCOUG does not warrant the NoCOUG Journal to be error-free. The statements and opinions expressed here are those of the authors and do not necessarily represent those of NoCOUG. Thanks to Franck Pachot, Karl Arao, Mark Farnham, Mohammad Inamullah, Mogens Nørgaard, Niall Litchfield, Sumit Sengputa, and Tim Hall for their reviews and comments.
Copyright © 2017, Nathan Biggs and Pamela Fulmer

Oracle Cloud Pocket Solutions Guide

Book Notes by Brian Hitchcock



Brian Hitchcock

Details

Author: Charles Kim, Nitin Vengurlekar, Jerry Ward, and Sudhir Balasubramanian

ISBN-13: 978-1537530277

Date of Publication: November 23, 2016

Publisher: Amazon Digital Services

Summary

This book is very short, only 110 pages and a lot of screenshots, but it does show you how to configure the various parts of Oracle Cloud. On page 2 we see the subtitle: *Real-Life Solutions for the Cloud*, and on page 3 we see a note that this book is the Oracle Open World 2016 Edition, so I assume this was published to hand out at OOW 2016. Three of the four authors are members of the consulting firm Viscosity North America; page 5 has an extensive list of the services they offer. The editor of the *NoCOUG Journal* asked me to review this book and sent me a physical copy. It has been some time since I have reviewed a physical book; it's a lot harder to highlight the text and make notes. It wasn't that long ago that I thought it was strange to be reading books online, but now it seems completely normal.

Preface

In this section we get some history of cloud computing. We are told that cloud computing helps IT organizations that are faced with growing IT costs due to legacy systems, software licensing, and power costs. Public and private cloud initiatives reduce these cost issues through better standardization and higher utilization. We also see comments about greater agility and improved responsiveness. There's not much new here, but of more interest is the observation that as most organizations go forward with their cloud projects, security is the high-priority concern. Many organizations also have to deal with lots of regulations and sensitive customer data. Public cloud environments are not directly controlled by the owners of the data, which makes securing information problematic. In 2016, the IOUG Survey on Database Cloud showed that security and privacy

were the biggest reasons that prevented more organizations from moving to the cloud.

And yet, more business moves to the cloud each year. The public cloud has become a disruptive technology. Any new development project must consider the cloud if only due to the significant savings in operational expenses. At the same time, the public cloud is getting faster. Cloud customers generally feel the cloud is slow, but progress is being made to address this using faster VMs and solid-state disks. It remains to be seen if any amount of cloud performance would actually meet customer expectations. So much is different, but nothing has changed? The cloud also promises the ability to expand and shrink an environment on demand to better meet performance needs while reducing budget issues.

Finally, Amazon Web Services (AWS) is identified as leading the way to take customers to the cloud, while Microsoft Azure is gaining ground. It will be interesting to see who controls the various segments of the cloud market as time goes by.

Chapter 1: Introduction to Oracle Cloud

Chapter 1 in this short book is less than one page long. We learn that Oracle is very much invested in the cloud and that Oracle wants to be a leader in the "magic quadrant." I had to look up the term. The statement means that Oracle wants to be highly rated in the Gartner market research reports of the same name. Oracle has cloud offerings in SaaS (Software as a Service), PaaS (Platform as a Service), and IaaS (Infrastructure as a Service) but hopes to capture more of the PaaS and IaaS areas. A list of Oracle products for each is provided.

For PaaS, Oracle offers Data Management Cloud (databases), Application Development Cloud (Java and other app-related software), Integration Cloud (SOA and GoldenGate), Business Analytics Cloud (BI and Big Data), Content and Process Cloud (Documents and Social Network), and Management Cloud (Analytics).

For IaaS, Oracle offers subscription-based services for Compute, Storage, Network, and Cloud Machine.

This short chapter ends with the observation that Database as a Service (DBaaS) is growing rapidly and will triple over the next

"It has been some time since I have reviewed a physical book; it's a lot harder to highlight the text and make notes. It wasn't that long ago that I thought it was strange to be reading books online, but now it seems completely normal."

two years, and that many enterprises are moving their data to the cloud.

Chapter 2: Oracle Cloud Fundamentals

Here we have three pages describing how to get started with Oracle Cloud. Note that Oracle Cloud is Oracle Public Cloud (OPC). Screenshots are shown covering the login process using your user ID, password, and Identity Domain. I find it interesting that you have to choose the location of “your” data center. At first I thought OPC should choose for me, but then I thought—depending on where my customers are geographically—perhaps I want to choose the data center location. Now I’m wondering how many cloud customers know where all their users are and how much difference the location of the data center makes. Overall, how does one make this decision and how important is this? The book does not comment on this.

After you choose your data center, you enter the Identity Domain that was assigned to you when your account was created. The domain that was assigned can be across all OPC Services and all data centers. Now I’m wondering if I want or need to have my OPC Services across multiple data centers and how I would decide this. Furthermore, as time moves along, do I revisit this periodically? Will OPC provide tools to help me monitor all this and recommend changes as needed? Again, this is not covered in the text.

At this point you will be at the Oracle Cloud dashboard page, which displays a summary of all the resources and services available to you. Note that the specifics of all of this are changing rapidly: the screens you see today may be different from what is shown in this book.

Chapter 3: Storage Cloud

This is one of Oracle’s IaaS services that provides scalable file and unstructured data storage that can also be used to back up files from on site to the cloud. After the data is uploaded to the Storage Cloud, the data can be shared with any part of your organization.

There are two versions of this service: Standard and Archive. Standard provides 1Tb of storage for \$30/month. As always, these numbers can and will change over time. Archive is \$1/month for 1Tb but has “numerous caveats and rules.” We aren’t told what these rules are, but we are told that Amazon Web Services (AWS) has similar restrictions. Given how much lower the cost is for Archive, I assume the restrictions are pretty severe. The text tells us that all these rules will be covered in an update of this book, and that Archive is much less costly than AWS Glacier Service. We aren’t told what AWS Glacier is, but my first guess was that it’s glacially slow. I was joking of course, but when I looked it up on the AWS website, I found that AWS Glacier is a low-cost archive storage service where access to data can take from a few minutes to a few hours. When you access Storage Cloud you are accessing the Oracle Storage Cloud Storage Appliance, which makes the storage look like NFS-mounted storage.

There are sections covering many aspects of this storage service. To store data you need to create a container, which is like a directory but can’t be nested. You define the access control for each container as well as the replication policy. We are told that when we create or change a replication policy we must wait 10 to

15 minutes. I’m curious about why this is so, but we aren’t given the answer.

The section on ReST describes it as architecture for stateless, client-server protocol for HTTP. ReST is an alternative to Remote Procedure Calls and Web Services like SOAP/WSDL. Curl is a command-line utility for moving files using URL syntax. A shell script that uses curl—curl.ksh—is described. Examples are shown using curl for creating and archiving a container as well as uploading a file, listing the contents, deleting a file, and determining the space consumed by a container.

Chapter 4: Database Cloud

This is the longest chapter in the book. Oracle Database Cloud services sets Oracle apart from all other cloud services providers, and we are told that Oracle provides Database as a Service better than anyone else. You can choose from having a

“Public cloud environments are not directly controlled by the owners of the data, which makes securing information problematic. In 2016, the IOUG Survey on Database Cloud showed that security and privacy were the biggest reasons that prevented more organizations from moving to the cloud.”

service that supports a single schema to having Exadata Service for multiple instances. As a customer of Oracle Database as a Service you get a dedicated virtual machine for your database instance(s), running database 11gR2 or 12c. You can also run RAC, and data is encrypted using TDE. You will have full root access and the SYSDBA role, so you will have full access to all aspects of managing your database. On a personal note, you may not want this full access; I’ve seen many cases where customers want full access, get full access, and then make a real mess of things.

The next section covers the process of provisioning a database, and many screenshots are included. The process includes choosing the instance type (for example, Oracle database or MySQL) and creating the service to start the database wizard. All of the Oracle software is pre-installed. If you choose the Database Cloud Service you use the Database Creation Assistant (DBCA) to create your database, or you can choose the Virtual Image option where you create the database using SQL*Plus commands. We are told that by default billing is set to non-metered monthly. We aren’t told more about these billing details. The database will have the latest PSU and required one-off patches installed. You also have to choose the software edition, which means choosing between Standard Edition, Enterprise Edition, High Performance, or Extreme Performance. Enterprise is the default. High Performance is Enterprise with Active Data Guard (ADG), In-Memory Database, and RAC. Extreme Performance is Enterprise with all the enterprise management packs. Next is the Service Details screen, where you have more database configuration options as

“Amazon Web Services (AWS) is identified as leading the way to take customers to the cloud, while Microsoft Azure is gaining ground. . . . We are told that Oracle provides Database as a Service better than anyone else.”

well as choices for backups. You need to upload or generate SSH keys for your new database. For backups you choose from cloud and local, cloud only, or none. For the cloud and local option, 30 days of backups are made to the cloud with the seven most recent backups placed in local storage as well. The cloud-only option does 30 days of backups but nothing locally. You need to create a container in the cloud to store the cloud backups. There are notes in this section describing how to create a database up to 12 TB, but to do so you have to provision and allocate 2 TB storage volumes. I don't understand why this is a manual process; why can't this be automated? You know how much space you want; the rest should be handled behind the scenes. At this point you can create your database, a process that can take up to 30 minutes.

This chapter of the book brings up an important issue. Screenshots can be very useful when describing how to use GUI-based software, but those screenshots have to be clear enough that they can be read. I could not read most of the text of most of the screenshots. If the intent was that this book would be read on a platform where the screenshots could be enlarged, then this issue might be an oversight for the printed copies. However, this

“The Dedicated Compute option is recommended This implies that performance in the cloud isn't as predictable as customers might like. You won't know how many other customers are competing for the resources of the site.”

book is only available in printed form. I wonder if those who looked at this book before it was made available to the public can read the text of these screenshots.

Pricing out the database cloud is covered next. A chart is shown where the cost per month for the various flavors of Enterprise Edition range from \$3,000 to \$5,000. For the first time we see that Standard Edition is also available for \$600/month. I don't know if the current prices are different. We also see a new option called High-Memory Compute that brings more memory, although it isn't clear how much more. There are also columns in the chart for the cost of these options per hour. I don't know how this would be measured. For instance, what does “per hour” really mean? Is that connect time, CPU time, or something else?

The steps needed to connect to a database using ssh as oracle and root are shown. By default, port 1521 is not open, so you have to deal with this. It isn't clear to me if you must use port 1521 or if you can specify a different port.

You probably have existing databases that are not in the cloud. In order to move data between them and your new cloud database you need to create a database link between them. The commands to do this are shown. To be more secure, you may not want to open port 1521 at all, and in this case we have the steps needed to create a secure tunnel with ssh to your cloud database.

The next section of the chapter covers how to add database storage to DBaaS. There are screenshots, and a lot of steps are required. I don't understand why I have to be concerned with partitions and offsets, and creating labels. I think I should be able to specify how much storage I want to add, and that should be the end of it. Doesn't the Oracle Cloud have massive amounts of processing power? Can't the robots set up the disk space on their own? Having added space, we now need to cover the basics of Linux logical volume manager (LVM). The commands that are shown cover creating a physical volume, a logical volume, and a volume group.

This chapter ends with a conclusion section that warns us that if we don't adapt to the cloud our jobs may be outsourced to India. Other advice includes the statement that Oracle is a dominant force in cloud computing. I'm not sure how AWS would react to that.

Chapter 5: Oracle Database Backup Cloud

For your databases that are not in the cloud, you can still use the Oracle Cloud to store backups of those databases using the Oracle Database Backup Cloud. These backups are required to be encrypted and can also be compressed. Personally, I think all backups should be encrypted. You never know where those backup tapes are going to wind up. All of this can be yours for the low, low price of \$33 per month per TB for the non-metered option. There is a metered option for \$24/month/TB, but the authors advise us not to use this as the metered transfer charges add up quickly.

To use this service, you first need to download the Database Cloud Backup Module. This module supports backup tapes and is fully integrated with RMAN. Creating a backup to the Oracle Cloud is just the same as making a backup to any other tape device. The process of installing this module is explained, followed by an example of configuring RMAN to make backups to the Oracle Cloud.

Just when it was sounding easy, things got complicated. If you want to back up Oracle Standard Edition databases to the Oracle Cloud, you have to apply a specific patch. I'm curious: does this special patch survive future database upgrades?

Next is a section covering best practices for making backups to OPC, followed by an extensive example of how to restore and recover a database from OPC.

Chapter 6: Oracle Storage Cloud Software Appliance

This is a software product, installed in their environment by a customer, that acts as a gateway to Oracle Cloud. The hardware and software requirements for installation are reviewed. The OSCSA encrypts the data before it is replicated to the cloud. This means all your data is encrypted all of the time it is outside your environment. Screenshots are shown for the login process and selecting configuration options. There are sections on best practices for OSCSA and how to find additional information. Once installed, you can copy any files from your organization to the OSCSA, and those files are encrypted and backed up into Oracle Cloud.

This chapter focuses on Infrastructure as a Service (IaaS), which—unlike the other two service models, PaaS and SaaS—really doesn't provide any application framework. IaaS provides compute power, network services, and storage, while the customer provides the applications. A customer may have multiple virtualized servers across multiple data centers. A site is defined to be a set of physical servers. Customers can choose to be one of many tenants using the site, known as Compute Capacity, or they can be the only tenant on the site, known as Dedicated Compute. For customers that want reliably predictable performance, the Dedicated Compute option is recommended. While the text doesn't raise the following concern, I will: this implies that performance in the cloud isn't as predictable as customers

“I believe the goal of the Cloud has been that anyone with a credit card and an internet connection can set up compute resources. If you have to debate data centers and storage volumes, I don't think we have realized all the promise of the cloud. I remain unclear as to why all of this has not been automated.”

might like. You won't know how many other customers are competing for the resources of the site. I wonder how much of an issue this is in practice.

Once connected, customers see a dashboard where they can provision and manage their Oracle Compute Cloud Service instances. Screenshots show us how to create an instance, which is an Oracle Compute Cloud virtual machine created using a machine image with assigned memory and CPU resources. The Instance Creation Wizard offers a range of images for different applications and CPU/memory profiles. Customers can choose from Oracle-branded images, private images, and marketplace images. All Oracle-branded images are free, while those in the marketplace may or may not be free. This sounds similar to the App Store for a smart phone. Customers can create their own images and upload them to their IaaS service.

A key feature of IaaS is called Cloud Burst, which allows you to temporarily use capacity beyond your subscription rate. You will be billed for each hour of excess capacity you use. I don't see any discussion of specific costs of all of this, but I guess it would be difficult to cover pricing with so many options.

There are multiple layers to the Oracle Compute Cloud, each of which needs to be configured. First we see screenshots for the steps necessary to configure the orchestration model, which defines the objects and attributes of our Oracle Compute Cloud Service instance, including instances and storage volumes. Next we see screenshots for configuring an instance, which is a virtual machine running a specific OS. Each instance has CPU and memory resources and is part of an orchestration. Next is the setup of security applications, rules, and lists, all of which control access to your instance. Storage volumes are configured as are

storage snapshots, which can be used for creating other instances or as backups.

Appendix A: Sign up for a free trial account

Here we have screenshots showing the process of creating a trial account. This account last for 30 days, but what happens then? Do you have the option of saving your data or moving to a permanent account?

Appendix B: Generating SSH Keys on Windows

If you are running on Windows, the specific steps to generate SSH keys are shown. This process was covered for Linux systems as part of setting up an Oracle Cloud Database Service.

Conclusion

This short book gives a good description of what it is like to actually set up a service in Oracle Public Cloud. I am puzzled by the steps of the process that are not automated. I believe the goal of the Cloud has been that anyone with a credit card and an internet connection can set up compute resources. If you have to debate data centers and storage volumes, I don't think we have realized all the promise of the cloud. I remain unclear as to why all of this has not been automated. ▲

Brian Hitchcock works for Oracle Corporation where he has been supporting Fusion Middleware since 2013. Before that, he supported Fusion Applications and the Federal OnDemand group. He was with Sun Microsystems for 15 years (before it was acquired by Oracle Corporation) where he supported Oracle databases and Oracle Applications. His contact information and all his book reviews and presentations are available at www.brianhitchcock.net/oracle-dbafmw/. The statements and opinions expressed here are the author's and do not necessarily represent those of Oracle Corporation.

Copyright © 2017, Brian Hitchcock

Be the first to correctly identify the flowers and insects on the front cover of the NoCOUG Journal and win a one-year individual membership to NoCOUG worth \$95. A high-resolution image can be downloaded from http://www.nocoug.org/download/2017-08/NoCOUG_Journal_201708-page-001.jpg. Tweet your entry to @NoCOUG.

Oracle SQL Tuning with Oracle SQLTXPLAIN

by Stelios Charalambides

This is the introductory chapter from the book *Oracle SQL Tuning with Oracle SQLTXPLAIN: Oracle Database 12c Edition* published by Apress, April 2017, ISBN 978-1484224359; Copyright 2017. For a complete table of contents, please visit the publisher site: <http://www.apress.com/us/book/9781484224359>.

Welcome to the world of fast Oracle SQL tuning with SQLT. Never heard of SQLT? You're not alone. I'd never heard of it before I joined ORACLE, and I had been a DBA for more years than I care to mention. That's why I'm writing this book. SQLT is a fantastic tool because it helps you diagnose tuning problems quickly. What do I mean by that? I mean that in half a day, maximum, you can go from a 'slow' SQL to having an understanding of why the SQL is malfunctioning, and finally, to knowing how to fix the SQL. This of course assumes that your SQL can run faster. Some SQLs are just doing their best with the data and what you are asking them to do. It's not miraculous, but it's pretty close.

Will SQLT fix your SQL? No. Fixing the SQL takes longer. Some tables are so large that it can take days to gather statistics. It may take a long time to set up the test environment and roll the fix to production. You may find that you need to make design changes that affect other application SQLs. The important point is that in half a day working with SQLT, you will have an explanation. You'll know what needs to be done (if anything) to improve the SQL, and if not you'll be able to explain why it can't go any faster. That may sound like a losing scenario, but when you know something can't be improved, you can move on to other tasks.

You need to know about SQLT because it will make your life easier. But let me back up a little and tell you more about what SQLT is, how it came into existence, why you probably haven't heard of it, and why you should use it for your Oracle SQL tuning.

What Is SQLT?

SQLT is a set of packages and scripts that produces HTML-formatted reports, some SQL scripts, and some text files. The entire collection of information is packaged in a zip file and often sent to Oracle Support, but you can look at these files yourself. There are just over a dozen packages and procedures (called "methods") in SQLT. These packages and procedures collect different information based on your circumstances. We'll talk about the packages suitable for a number of situations later.

What's the Story of SQLT?

They say that necessity is the mother of invention, and that was certainly the case with SQLT. Oracle support engineers handle a huge number of tuning problems on a daily basis; problem is, the old methods of linear analysis are just too slow. You need to see the big picture fast so you can zoom in on the detail and tell the customer what's wrong. As a result, Carlos Sierra, a support engineer at the time created SQLT. The routines evolved over many visits to customer sites to a point where they can gather *all* the information required quickly and effectively. He then provided easy-to-use procedures for reporting on those problems.

The Oracle SQLTXPLAIN tool (that this book is about) was always an unsupported product until recently. Any fixes to the free tool were done by Oracle on a best efforts basis. Now you can log a bug against this useful tool in case there are problems.

Why Haven't You Heard of SQLT?

If it's so useful, why haven't you heard about SQLT? Oracle has tried to publicize SQLT to the DBA community, but still I get support calls and talk to DBAs who have never heard of SQLT—or if they have, they've never used it. This amazing tool is free to supported customers, so there's no cost involved. DBAs need to look at problematic SQL often, and SQLT is hands down the fastest way to fix a problem. The learning curve is nowhere near as high as the alternatives: interpreting raw 10046 trace files or 10053 trace files. Looking through tables of statistics to find the needle in the haystack, guessing about what might fix the problem and trying it out? No thanks. SQLT is like a cruise missile that travels across the world right to its target. Over the past few years SQLT has certainly gained some prominence (I hope partially due to the first edition of this book), but still there are too many sites that do not use it on a regular basis, and that, in my opinion is a lost opportunity.

Perhaps DBAs are too busy to learn a tool, which is not even mentioned in the release notes for Oracle. It's not in the documentation set, and even though it is officially part of the product now, it's just a tool that happens to be better than any other tool out there. Let me repeat. It's free.

It's also possible that some DBAs are so busy focusing on the obscure minutiae of tuning that they forget the real world of fixing SQL. Why talk about a package that's easy to use when you could be talking about esoteric hidden parameters for situations

you'll never come across? SQLT is a very practical tool. It even collects other tuning information, such as AWRs and SQLMonitor reports if it can. So SQLT saves you even more time.

Whatever the reason, if you haven't used SQLT before, my mission in this book is to get you up and running as fast and with as little effort from you as possible. I promise you that installing and using SQLT is easy. Just a few simple concepts, and you'll be ready to go in 30 minutes.

How Did I Learn about SQLT?

Like the rest of the DBA world (I've been a DBA for many years), I hadn't heard of SQLT until I joined Oracle. It was a revelation to me. Here was this tool that's existed for years, which was exactly what I needed many times in the past, although I'd never used it. Of course I had read many books on tuning in years past: for example, Cary Millsaps's classic *Optimizing Oracle Performance*, and of course *Cost-Based Oracle Fundamentals* by Jonathan Lewis.

The training course (which was two weeks in total) was so intense that it was described by at least two engineers as trying to drink water from a fire hydrant. Fear not! This book will make the job of learning to use SQLT much easier.

Now that I've used SQLT extensively in day-to-day tuning problems, I can't imagine managing without it. I want you to have the same ability. It won't take long. Stick with me until the end of the book, understand the examples, and then try and relate them to your own situation. You'll need a few basic concepts (which I'll cover later), and then you'll be ready to tackle your own tuning problems. Remember to use SQLT regularly even when you don't have a problem; this way you can learn to move around the main HTML file quickly to find what you need. Locate the useful extra files that SQLT collects for you, including AWRs, special Exadata Cell configuration information, and many other goodies. Run a SQLT report against SQL that isn't a problem. You'll learn a lot. Stick with me on this amazing journey.

Getting Started with SQLT

Getting started with SQLT couldn't be easier. I've broken the process down into three easy steps.

1. Downloading SQLT
2. Installing SQLT
3. Running your first SQLT report

SQLT will work on many different platforms. Many of my examples will be based on Linux, but Windows or Unix is just as easy to use, and there are almost no differences in the use of SQLT between the platforms. If there are, I'll make a note in the text.

How Do You Get a Copy of SQLT?

How do you download SQLT? It's simple and easy. I just did it to time myself. It took two minutes. Here are the steps to get the SQLT packages ready to go on your target machine:

1. Find a web browser and log in to My Oracle Support (<http://support.oracle.com>)
2. Go to the knowledge section and type "SQLT" in the search box. Note 215187.1 entitled "SQLT (SQLTXPLAIN)" – Tool that helps to diagnose a SQL statement performing

poorly [ID 215187.1]" should be near the top of the list.

3. Scroll to the bottom of the note and choose the version of SQLT suitable for your environment. There are currently versions suitable from 9.2 to 12c.
4. Download the zip file.
5. Unzip the zip file. The current version is 12.1.160429 (from April 29, 2016).

You now have the SQLT programs available to you for installation onto any suitable database. You can download the zip file to a PC and then copy it to a server if needed.

How Do You Install SQLT?

So without further ado, let's install SQLT so we can do some tuning:

1. Download the SQLT zip file appropriate for your environment (see steps above).
2. Unzip the zip file to a suitable location.
3. Navigate to your "install" directory under the unzipped area (in my case it is /home/oracle/sqlt/install, but your locations will be different).
4. Connect as sys, for example, sqlplus / as sysdba.
5. Make sure your database is running.
6. Run the sqcreate.sql script.
7. Select the default for the first option. (We'll cover more details of the installation in Appendix A.)
8. Enter and confirm the password for SQLTXPLAIN.
9. Assuming you know which Tablespaces you want to use to keep the SQLT objects and procedures, select the default "NO" again to avoid listing all the Tablespaces.
10. Select the tablespace where the SQLTXPLAIN will keep its packages and data (in my case, USERS).
11. Select the temporary tablespace for the SQLTXPLAIN user (in my case, TEMP).
12. Then enter the username of the user in the database who will use SQLT packages to fix tuning problems. Typically this is the schema that runs the problematic SQL (in my case this is STELIOS).
13. Then enter "T", "D", or "N." This reflects your license level for the tuning and diagnostics packs. Most sites have both so you would enter "T" (this is also the default). My test system is on my private server (an evaluation platform with no production capability) so I would also enter "T." If you have the diagnostics pack, only enter "D"; and if you do not have these licenses, enter "N."

The last message you see is "SQCREATE completed. Installation completed successfully."

Make sure you have granted sys the appropriate privilege on SQLTXADMIN (This is for 12c databases only).

SQL> grant inherit privileges on user sys to sqltxadmin;

Running Your First SQLT Report

Now that SQLT is installed, it is ready to be used. Remember that installing the package is done as sys and that running the reports is done as the target user. Please also bear in mind that although I have used many examples from standard schemas

available from the Oracle installation files, your platform and exact version of Oracle may well be different, so please don't expect your results to be exactly the same as mine. However, your results will be similar to mine, and the results you see in your environment should still make sense.

1. Now exit SQL and change your directory to /home/oracle/sqlt/run. From here log in to SQLPLUS as the target user.
2. Then enter the following SQL (this is going to be the statement we will tune):
SQL> select count(*) from dba_objects;
3. Then get the SQL_ID value from the following SQL:
SQL> select sql_id from v\$sqlarea where sql_text like 'select count(*) from dba_objects%';
In my case the SQL_ID was g4pkmrqrgxg3b.
4. Now we execute our first SQLT tool sqltextract from the target schema (in this case STELIOS) with the following command:
SQL> @sqltextract g4pkmrqrgxg3b
5. Enter the password for SQLTXPLAIN (which you entered during the installation). The last message you will see if all goes well is "SQLTXTRACT completed".
6. Now create a zip directory under the run directory and copy the zip file created into the zip directory. Unzip it.
7. Finally from your favorite browser navigate to and open the file named sqlt_s<nnnnn>_main.html. The symbols "nnnnn" represent numbers created to make all SQLT reports unique on your machine. In my case the file is called sqlt_s89906_main.html

Congratulations! You have your first SQLT XTRACT report to look at.

When to Use SQLTXTRACT and When to Use SQLTXECUTE

SQLT XTRACT is the easiest report to create because it does not require the execution of the SQL at the time of the report generation. The report can be collected after the statement has been executed. SQLTXECUTE, on the other hand, executes the SQL statement and thus has better runtime information and access to the actual rows returned. This happens when statistics_level=all or "_rowsource_execution_statistics=true". This means it can assess the accuracy of the estimated cardinality of the steps in the execution plan (see "Cardinality and Selectivity" later in this chapter). SQLTXECUTE will get you more information, but it is not always possible to use this method, perhaps because you are in a production environment or perhaps the SQL statement is currently taking three days to run, which is why you are investigating this in the first place. Another reason for not running SQLTXECUTE for some SQL statements is if they are DML (insert, update, delete, or merge), they will change data. We will look at both SQLTXECUTE and SQLTXTRACT report (and other SQLT options also). For now we will concentrate on one simple SQLTXTRACT report on a very simple SQL statement. So let's dive in.

Your First SQLT Report

Before we get too carried away with all the details of using the SQLT main report, just look at Figure 1. It's the beginning of a

whole new SQLT tuning world. Are you excited? You should be. This header page is just the beginning. From here we will look at some basic navigation, just so you get an idea of what is available and how SQLT works, in terms of its navigation. Then we'll look at what SQLT is actually reporting about the SQL.

215187.1 SQLT XTRACT 12.1.160429 Report: sqlt_s46414_main.html

Review log and fix following errors:

sqltSa: *** icompute_estim_size_if_rebuilt: ORA-01031: insufficient privileges

Global	Plans	Tables
<ul style="list-style-type: none"> Observations SQL Text SQL Identification Environment CBO Environment Fix Control CBO System Statistics DBMS_STATS Setup Initialization Parameters NLS Parameters I/O Calibration Tool Configuration Parameters 	<ul style="list-style-type: none"> Summary Performance Statistics Performance History (delta) Performance History (total) Execution Plans <p>Plan Control</p> <ul style="list-style-type: none"> Stored Outlines SQL Patches SQL Profiles SQL Plan Baselines SQL Plan Directives 	<ul style="list-style-type: none"> Tables Statistics Statistics Extensions Statistics Versions Modifications Properties Physical Properties Constraints Columns Indexed Columns Histograms Partitions Indexes
Cursor Sharing and Binds	SQL Execution	Objects
<ul style="list-style-type: none"> Cursor Sharing Adaptive Cursor Sharing Peeked Binds Captured Binds 	<ul style="list-style-type: none"> Active Session History AWR Active Session History SQL Statistics SQL Detail ACTIVE Report Monitor Statistics Monitor ACTIVE Report Monitor HTML Report Monitor TEXT Report Segment Statistics Session Statistics Session Events Parallel Processing 	<ul style="list-style-type: none"> Objects Dependencies Fixed Objects Fixed Object Columns Nested Tables Policies Audit Policies Tablespaces Metadata
SQL Tuning Advisor		
<ul style="list-style-type: none"> STA Report STA Script 		

Figure 1. The top part of the SQLT report shows the links to many areas

Some Simple Navigation

Let's start with the basics. Each hyperlinked section has a Go to Top hyperlink to get you back to the top. There's a lot of information in the various sections, and you can get lost. Other related hyperlinks will be grouped together above the Go to Top hyperlink. For example, if I clicked on Indexes (the last link under the Tables heading), I would see the page shown in Figure 2.

Indexes

#	Table Name	Owner	Count ¹	Num Rows ²	Sample Size ²	Blocks ²	Last Analyzed ²	Indexes
1	LINKS	SYS		0	0	0	07-JUL-14	1
2	OBJ\$	SYS	92190	19416	19416	300	07-JUL-14	5
3	SUM\$	SYS	2	0	0	0	07-JUL-14	2
4	USERS	SYS		63	63	3	07-JUL-14	2
5	USER_EDITIONS	SYS	2	2	2	1	07-JUL-14	1

(1) SELECT COUNT(*) performed in Table as per tool parameter "count_star_threshold" with current value of 10000.

(2) CBO Statistics.

[Go to Indexed Columns](#)

[Go to Tables](#)

[Go to Top](#)

SYS.LINKS - Indexes

#	In Plan	Index Name	Owner	Index Type	Uniqueness	Col ID	Column Name	Column Name ¹	Num Rows ²	Sample Size ²	Last Analyzed ²
1	TRUE	LLINK1	SYS	NORMAL	NONUNIQUE	1	OWNER#	OWNER#	0	0	07-JUL-14
						2	NAME	NAME			

(1) Column names including system generated names.

(2) CBO Statistics.

[Go to Indexes](#)

[Go to Tables](#)

[Go to Top](#)

SYS.LINKS - Index Statistics

Figure 2. The Indexes section of the report

Before we get lost in the SQLT report, let's again look at the header page (Figure 1). The main sections cover all sorts of aspects of the system.

- CBO environment
- Cursor sharing

- Take a minute and browse through the report. Did you notice the hyperlinks on some of the data within the tables? SQLT collected all the information it could find and cross-referenced it all.
- So, for example, continuing as before from the main report at the top (Figure 1):

USERS - indexes

Index Name	Owner	Index Type	Uniqueness	Col #1	Column Name1	Column Name2	Num Rows	Sample Size%	Last Analyzed	District Keys ¹	Buckets ¹	Segment Extents ²	Segment Blocks ²	Total Segment Blocks ²	DBSEG_SPACE Allocated ³	Leaf Blocks ³	Leaf Extent Target Size ⁴
TRUE_LUSER#	SYS	NORMAL	UNIQUE	1	USER#	USER#	63	63	07-JUL-14	Stats	Prox	Distx	Cols	Mats			
				2	TITLE#	TITLE#			20-SPR-21								
				21	SPACE#	SPACE#											
FALSE_LUSER#	SYS	NORMAL	UNIQUE	2	NAME#	NAME#	63	63	07-OCT-14	Stats	Prox	Distx	Cols	Mats			

Index names include system-generated names.

No Distincts.

INDEXES

USERS - Index Statistics

Plan	Index	Owner	Index Type	Parti	Parti Temp	Sum Rows	Sample Size%	Pctns	Last Analyzed	Distinct Keys ¹	Buckets ¹	Segment Extents ²	Segment Blocks ²	Total Segment Blocks ²	DBSEG_SPACE Allocated ³	Leaf Blocks ³	Leaf Extent Target Size ⁴
TRUE_LUSER#	SYS	NORMAL	N	I	63	63	100.0	2014-07-0705:56:06	63	0	1	8	0	1			
FALSE_LUSER#	SYS	NORMAL	N	I	63	63	100.0	2014-07-0705:56:06	63	0	1	8	0	1			

No Distincts.

Considers the blocks from all partitions (if the index is partitioned).

¹The estimate size of the index if it were rebuilt, as computed by DBSEG_SPACE.CREATE_INDEX_COST.

²Number of leaf blocks with a 90% index efficiency. Only evaluated for nonpartitioned normal indexes with more than 10000 leaf blocks.

³Estimate for nonpartitioned normal indexes with less than 10000 leaf blocks.

⁴Estimate for nonpartitioned normal indexes with less than 10000 leaf blocks.

⁵Assumes default CBO environment, including optimizer_index_causd_4sg-10 and optimizer_index_causd_4sg-10 among others.

⁶For example, if the index has 10000 leaf blocks, the estimate would be 10000 * 10 = 1000000 bytes.

⁷Size of Plan T was used to report; the estimation comes from EXPLAIN PLAN FOR the CREATE INDEX command so it's dependent on good statistics.

INDEXES STATISTICS INFORMATION

INDEXES

TABLES

TOP

- Now, in the Index Meta column (far right in Figure 3), click on the Meta hyperlink for the I_USER2 index to display the index metadata shown in Figure 4.

Here we see the statement we would need to create this index. Do you have a script to do that? Well SQLT can get it better and faster. So now that you've seen a SQLT report, how do you approach a problem? You've opened the report, and you have one second to decide. Where do you go?

Well, that all depends.

As with any methodology, different approaches are considered for different circumstances. In the next chapter we look at

Synonym - Metadata

- [DBA_OBJECTS](#)

[Go to Metadata](#)
[Go to Top](#)

PUBLIC.DBA_OBJECTS - Synonym Metadata

```
CREATE OR REPLACE NONEDITIONABLE PUBLIC SYNONYM "DBA_OBJECTS" FOR "SYS"."DBA_OBJECTS"
```

[Go to Metadata](#)
[Go to Top](#)

AWR how that helps us decide if we should be tuning SQL or the system. After all there's no point in trying to tune a SQL if your system is not able to run it properly. Once you've decided there is something wrong with your SQL, you could use a SQLT report. Once you have the SQLT report, you are presented with a header page, which can take you to many different places (no one reads a SQLT report from start to finish in order). So where do you go from the main page?

If you're absolutely convinced that the execution plan is wrong, you might go straight to "Execution Plans" and look at the history of the execution plans. We'll deal with looking at those in detail later.

Suppose you think there is a general slowdown on the system. Then you might want to look at the "Observations" section of the report.

Maybe something happened to your statistics, so you'll certainly need to look at the "Statistics" section of the report under "Tables."

All of the sections I've mentioned above are sections you will probably refer to for every problem. The idea is to build up a picture of your SQL statement, understand the statistics related to the query, understand the cost-based optimizer (CBO) environment, and try and get into its "head." Why did it do what it did? Why does it not relate to what you think it ought to do? The SQLT report is the explanation from the optimizer telling you why it decided to do what it did. Barring the odd bug, the CBO usually has a good reason for doing what it did. Your job is to set up the environment so that the CBO agrees with your worldview and run the SQL faster!

Cardinality and Selectivity

My objective throughout this book, apart from making you a super SQL tuner, is to avoid as much jargon as possible and explain tuning concepts as simply as possible. After all we're DBAs, not astrophysicists or rocket scientists.

So before explaining some of these terms, it is important to understand why these concepts are key to the CBO operation and to your understanding of the SQL running on your system. Let's first look at cardinality. It is defined as the number of rows expected to be returned for a particular column if a predicate selects it. If there are no statistics for the table, then the number is pretty much based on heuristics about the number of rows, the minimum and maximum values, and the number of nulls. If you collect statistics then these statistics help to inform the guess, but it's still a guess. If you look at every single row of a table (collecting 100 percent statistics), it might still be a guess because the data might have changed, or the data may be skewed (we'll cover skewness later). That dry definition doesn't really relate to real life, so let's look at an example. Click on the "Execution Plans" hyperlink at the top of the SQLT report to display an execution plan like the one shown in Figure 5.

Execution Plan phv:1203538133 [B]
[W] sqlt_phv:47043 sqlt_phv2:24547 source:GV\$SQL_PLAN inst:1 child:1(00000000D21DBF88) *

SQL Text: [L]
select count(*) from dba_objects

ID	Exec Ord	Operation	Go To	More	Cost	Estim Card	Work Area
0	10	SELECT STATEMENT			1199	1	
1	10	SORT AGGREGATE	[L]		1199	1	
2	17	VIEW DBA_OBJECTS			1199	409740	
3	10	UNION-ALL			401		
4	11	FILTER	[L]		401		
5	5	NASH JOIN	[L]		385	411335 [L]	
6	1	INDEX FULL SCAN I_USERS	[L]		1	132	
7	4	NASH JOIN	[L]		303	106319 [L]	
8	2	INDEX FULL SCAN I_USERS	[L]		1	132	
9	3	TABLE ACCESS FULL OBJ\$	[L]		382	93698	
10	6	TABLE ACCESS FULL USER_EDITIONING\$	[L]		2	1	
11	9	NOSTED LOOP JOIN	[L]		2	1	
12	7	INDEX SKIP SCAN I_USERS	[L]		1	1	
13	8	INDEX RANGE SCAN I_OBJ\$	[L]		1	1	
14	10	TABLE ACCESS FULL USER_EDITIONING\$	[L]		2	1	
15	15	NESTED LOOPS			0	1	
16	12	INDEX FULL SCAN I_LINK\$	[L]		0	1	
17	14	TABLE ACCESS CLUSTER UNES	[L]		0	1	
18	13	INDEX UNIQUE SCAN I_USERS	[L]		0	1	

Performance statistics are only available when parameter 'statistics_level' was set to 'ALL' at next parse time, or SQL contains 'gather_plan_statistics' *

(1) If estim_card * starts < output_rows then under-estimate. If estim_card * starts > output_rows then over-estimate. Color highlights when exceeding * 10x

(2) Largest contributors for cumulative-statistics columns are shown in red.

Other XML (id=1): [L]
Outline Data (id=1): [L]

Figure 5. An execution plan in the "Execution Plan" section

In the "Execution Plan" section, you'll see the "Estim Card" column. In my example, look at the TABLE ACCESS FULL OBJ\$ step. Under the "Estim Card" column the value is 93,698. Remember cardinality is the number of rows returned from a step in an execution plan. The CBO (based on the table's statistics) will have an estimate for the cardinality. The "Estim Card" column then shows what the CBO *expected* to get from the step in the query. The 93,698 shows that the CBO expected to get 93,698 records from this step, but in fact got 92,681. So how good was the CBO's estimate for the cardinality (the number of rows returned for a step in an execution plan)? In our simple example we can do a very simple direct comparison by executing the query shown below.

```
SQL> select count(*) from dba_objects;  
COUNT(*)  
-----  
92,681  
SQL>
```

So cardinality is the actual number of rows that will be returned, but of course the optimizer can't know the answers in advance. It has to guess. This guess can be good or bad, based on statistics and skewness. Of course, histograms can help here.

For an example of selectivity, let's look at the page (see Figure 6) we get by selecting Columns from the Tables options on the main page (refer to Figure 1).

Table Columns

#	Table Name	Owner	Count1	Num Rows2	Sample Size2	Blocks2	Last Analyzed2	Column Stats	Column Usage	Column Prop	Column Histogram	Single Table SQL Plan Directives
1	LINKS	SYS		0	0	0	07-JUL-14	11	2	Prop		
2	OBJ\$	SYS	92190	10416	10416	300	07-JUL-14	25	14	Prop	8	5
3	SUM\$	SYS	2	0	0	0	07-JUL-14	40	4	Prop		
4	USERS	SYS	63	63	63	3	07-JUL-14	30	11	Prop	6	2
5	USER_EDITIONING\$	SYS	2	2	2	1	07-JUL-14	2	2	Prop	1	

(1) SELECT COUNT(*) performed in Table as per tool parameter "count_star_threshold" with current value of 10000.

(2) CBO Statistics

Go to Tables

Go to Top

SYS.LINK\$ - Table Column

- Column Statistics
- Column Usage
- Column Properties
- Histograms

Go to Table Columns

Go to Tables

Go to Top

Figure 6. The "Table Column" section of the SQLT report

Look at the "SYS.OBJ\$ - Table Column" section. From the "Table Columns" page, if we click on the "25" under the "Column Stats" column, we will see the column statistics for the SYS.OBJ\$. Figure 7 shows a subset of the page from the "High Value" column to the "Equality Predicate Cardinality" column. Look at the "Equality Predicate Selectivity" and "Equality Predicate Cardinality" columns (the last two columns). Look at the values in the first row for OBJ\$.

High Value?	Last Analyzed	Avg Col Len	Density	Num Buckets	Histogram	Fluctuating Endpoint Count?	Popular Values	Global Stats	User Stats	Equality Predicate Selectivity	Equality Predicate Cardinality
"50"	2014-07-07 05:55:43	3	2.575196e-05	9	FREQUENCY	FALSE	0	YES	NO	0.111111	2158
"11"	2014-07-07 05:55:43	4	2.575196e-05	44	FREQUENCY	FALSE	0	YES	NO	0.022727	442
"50"	2014-07-07 05:55:43	3	2.575196e-05	9	FREQUENCY	FALSE	0	YES	NO	0.111111	2158
"19940"	2014-07-07 05:55:43	2	2.440212e-04	1	NONE	FALSE	0	YES	NO	0.000002	2
"obj_values49_1"	2014-07-07 05:55:43	20	6.900000e-05	254	HYBRID	FALSE	0	YES	NO	0.000009	2
"7405500"	2014-07-07 05:55:43	6	2.575196e-05	17	FREQUENCY	FALSE	0	YES	NO	0.000024	1143
"6"	2014-07-07 05:55:43	3	5.000000e-01	1	NONE	FALSE	YES	YES	NO	0.000000	0708
"55550"	2014-07-07 05:55:43	5	9.000000e-02	1	NONE	FALSE	YES	YES	NO	0.000009	1706

Figure 7. Selectivity is found in the "Equality Predicate Selectivity" column

Selectivity is 0.111111, and cardinality is 2158.

This translates to "I expect to get 2158 row back for this equality predicate, which is equivalent to a 0.111111 chance (1 is certainty 0 is impossible) or in percentage terms I'll get 11.11 percent of the entire table if I get the matching rows back."

Notice that as the cardinality increases the selectivity also increases. The selectivity only varies between 0 and 1 (or if you prefer 0 percent and 100 percent), and cardinality *should* only vary between 0 and the total number of rows in the table (excluding nulls). I say *should* because these values are based on statistics. What would happen if you gathered statistics on a partition (say it had 10 million rows) and then you truncate that partition, but don't tell the optimizer (i.e., you don't gather new statistics, or clear the old ones). If you ask the CBO to develop an execution plan in this case it might expect to get 10 million rows from a predicate against that partition. It might "think" that a full table scan would be a good plan. It might try to do the wrong thing because it had poor information.

To summarize, cardinality is the count of expected rows, and selectivity is the same thing but on a 0-1 scale. So why is all this important to the CBO and to the development of good execution plans? The short answer is that the CBO is working hard for you to develop the quickest and simplest way to get your results. If the CBO has some idea about how many rows will be returned for steps in the execution plan, then it can try variations in the execution plan and choose the plan with the least work and the fastest results. This leads into the concept of “cost,” which we will cover in the next section.

What Is Cost?

Now that we have cardinality for an object we can work with other information derived from the system to calculate a cost for any operation. Other information from the system includes the following:

- Speed of the disks
- Speed of the CPU
- Number of CPUs
- Database block size

These metrics can be easily extracted from the system and are shown in the SQLT report also (under the “Environment” section). The amount of I/O and CPU resource used on the system for any particular step can now be calculated and thus used to derive a definite cost. This is the key concept for all tuning. The optimizer is always trying to reduce the cost for an operation (even when the lowest cost, which after all is only an estimate, is not a guarantee of the best plan). I won’t go into details about how these costs are calculated because the exact values are not important. All you need to know is this: higher is worse, and worse can be based on higher cardinality (possibly based on out-of-date statistics), and if your disk I/O speeds are wrong (perhaps optimistically low) then full table scans might be favored when indexes are available. Cost can also be directly translated into elapsed time (on a quiet system), but that probably isn’t what you need most of the time because you’re almost always trying to get an execution time to be reduced, that is, lower cost. As we’ll see in the next section, you can get that information from SQLT. SQLT will also produce a 10053 trace file in some cases, so you can look at the details of how the cost calculations are made.

Reading the Execution Plan Section

We saw the execution plan section previously. It looks interesting, and it has a wobbly left edge and lots of hyperlinks. What does it all mean? This is a fairly simple execution plan, as it doesn’t go on for pages and pages (like SIEBEL or PeopleSoft execution plans). There are of course many different ways of getting an execution plan, which don’t involve SQLT, but I prefer SQLT’s presentation (see Figure 5) because it easily identifies the execution order (and much else besides).

There are a number of simple steps to reading an execution plan. I’m sure there’s more than one way of reading an execution plan, but this is the way I approach the problem. Bear in mind in these examples that if you are familiar with the pieces of SQL being examined, you may go directly to the section you think is wrong; but in general if you are seeing the execution plan for the first time, you will start by looking at a few key costs.

Philosophically speaking ‘cost’ is always an estimate. The optimizer derives a cost, but it is always an estimate for a ‘true’ cost, which can never be determined. So for the sake of brevity always assume that when I refer to cost, I am talking about the optimizer’s estimated cost.

The first and most important cost is the overall cost of the entire query. This is always shown as “ID 0” and is always the first row in the execution plan. In our example shown in Figure 5, this is a cost of 1199. So to get the cost for the entire query, just look at the first row. This is also the last step to be executed (“Exec Ord” is 19). The execution order is not top to bottom.

The Oracle engine will carry out the steps in the order shown by the value in the “Exec Ord” column.

So if we followed the execution through, the Oracle engine would do the execution in this order:

1. INDEX FULL SCAN I_USER2
2. INDEX FULL SCAN I_USER2
3. TABLE ACCESS FULL OBJ\$
4. HASH JOIN
5. HASH JOIN
6. TABLE ACCESS FULL USER_EDITIONING\$
7. INDEX SKIP SCAN I_USER2
8. INDEX RANGE SCAN I_OBJ4
9. NESTED LOOP SEMI
10. TABLE ACCESS FULL USER_EDITIONING\$
11. FILTER
12. INDEX FULL SCAN I_LINK1
13. INDEX UNIQUE SCAN I_USERS#
14. TABLE ACCESS CLUSTER USER\$
15. NESTED LOOPS
16. UNION-ALL
17. VIEW DBA_OBJECTS
18. SORT AGGREGATE
19. SELECT STATEMENT

However, nobody ever represents the plan of a SQL statement like this. What is important to realize is that the wobbly left edge gives information about how the steps are carried out. The less-indented operations indicate parent (also called outer) operations that are being carried out in child (also called inner) (more indented) operations. So for example steps 2, 3, and 4 would be read as “An index full scan is carried out using I_USERS2, then a full table scan of OBJ\$ and the results of these are HASH JOINED to produce a result set.” Each operation produces results for a less-indented section until the final result is presented to the SELECT (ID=0).

The “Operation” column is also marked with “+” and “-” to indicate sections of equal indentation. This is helpful in lining up operations to see which result sets an operation is working on.

So, for example, it is important to realize that the HASH JOIN at step 5 is using results from steps 1, 4, 2, and 3. We’ll see more complex examples of these later. It is also important to realize

Just above the first execution plan is a section called “Execution Plans.” This lists all the different execution plans the Oracle engine has seen for this SQL. Because execution plans can be stored in multiple places in the system, you could well have multiple entries in the “Execution Plans” section of the report. Its source will be noted (under the “Source” column). Here is a list of sources I’ve come across:

- GV\$SQL_PLAN
- GV\$SQLAREA_PLAN_HASH
- PLAN_TABLE
- DBA_SQLTUNE_PLANS
- DBA_HIST_SQL_PLAN

SQLT will look for plans in as many places as possible so that it can give you a full range of options. When SQLT gathers this information, it will look at the actual elapsed time associated with each of these plans and label them with “W” in red (worst) (worst Elapsed Time) and “B” in green (best) (best Elapsed Time). In my simple test case, the “Best” and “Worst” are the same, as there is only one execution plan in play. However you’ll notice there are two records: one came from mining the memory GV\$SQL_PLAN, and one came from the PLAN_TABLE (i.e., an EXPLAIN PLAN). You could also have one from DBA_SQLTUNE_PLANS, (SQL Tuning Analyzer).

When you have many records here, perhaps a long history, you can go back and see which plans were best and try to see why they changed. Noting the timing of a change can sometimes be crucial, as it can help you zoom in on the change that made things worse.

Before we launch into even more detailed use of the “Execution Plans” section, we’ll need more complex examples.

Join Methods

This book is focused on very practical tuning with SQLT. I try to avoid unnecessary concepts and tuning minutiae. For this reason I will not cover every join method available or every DBA table that might have some interesting information about performance or every hint. These are well documented in multiple sources, not least of which is the Oracle Performance guide (which I recommend you read). However, we need to cover some basic concepts to ensure we get the maximum benefit from using SQLT. So, for example, here are some simple joins. As its name implies, a join is a way of “joining” two data sets together: one might contain a person’s name and age and another table might contain the person’s name and income level. In which case you could “join” these tables to get the names of people of a particular age and income level. As the name of the operation implies, there must be something to join the two data sets together: in our case, it’s the person’s name. So what are some simple joins? (i.e., ones we’ll see in our SQLT reports).

HASH JOINS (HJ) – The *smaller* table is hashed and placed into memory. The *larger* table is then scanned for rows that match the hash value in memory. If the larger and smaller tables are the wrong way around this is inefficient. If the tables are not large, this is inefficient. If the smaller table does not fit in memory, then this is more than inefficient: it’s really bad!

NESTED LOOP (NL) – Nested Loop joins are better if the tables are smaller. Notice how in the execution plan examples above there is a HASH JOIN and a NESTED LOOP. Why was each chosen for the task? The details of each join method and its associated cost can be determined from the 10053 trace file. It is a common practice to promote the indexes and NL by adjusting the optimizer parameters `Optimizer_index_cost_adj` and `optimizer_index_caching` parameters. This is not generally a winning strategy. These parameters should be set to the defaults of 100 and 0. Work on getting the object and system statistics right first.

CARTESIAN JOINS – Usually bad. Every row of the first table is used as a key to access every row of the second table. If you have a very few number of rows in the joining tables this join is OK. In most production environments, if you see this occurring then something is wrong, usually statistics.

SORT MERGE JOINS (SMJ) – Generally joined in memory if memory allows. If the cardinality is high then you would expect to see SMJs and HJs.

Summary

In this chapter we covered the basics of using SQLTXTRACT. This is a simple method of SQLT that does not execute the SQL statement in question. It extracts the information required from all possible sources and presents this in a report.

In this chapter we looked at a simple download and install of SQLT. You’ve seen that installing SQLT on a local database can take very little time, and its use is very simple. The report produced was easy to unzip and can be used to investigate the SQL performance. In this first example we briefly mentioned cardinality and selectivity and how these affect the cost-based optimizer’s plans. Now let’s look at the bigger picture and what SQLT can do for you. ▲

The Northern California Oracle Users Group is a volunteer-run 501(c)(3) organization that has been serving the Oracle Database community of Northern California for more than thirty years. We organize four conferences a year and publish a quarterly Journal.
www.nocoug.org

Controlling and Monitoring Oracle Recovery Manager—Part II

by Norbert Debes



Norbert Debes

Introduction

This article is the second of a two-part series. In this installment I will address parallelizing RMAN operations in an Oracle Standard Edition (SE) environment. Once again RmanJ is used as a blueprint for implementing the features discussed. For those who would like a more thorough and complete understanding of RmanJ features, I recommend reading the *RmanJ User Guide* that is available at www.oradbpro.com/rmanj.html.

Serial Backup Limitation with Oracle Database Standard Edition

Any RMAN-based copy, backup, or restore operation requires a so-called channel to transfer data between a database file and an RMAN backup piece (i.e., part of a larger backup) or data file copy, and vice versa. While Oracle Enterprise Edition places no practical restriction on the number of parallel channels used by a single invocation of RMAN, Oracle Standard Edition does not allow parallel use of channels per RMAN invocation at all. Hence RMAN operations in an Oracle SE environment cannot fully leverage the speed of today's CPUs as well as the data transfer bandwidths of today's disk storage subsystems and (virtual) tape drives.

Oracle Standard Edition has always limited RMAN operations to a single channel. The quick test below demonstrates that Oracle12cR2 (version 12.2) is no exception. Note that since release Oracle12cR1, RMAN has been capable of executing SELECT statements.

```
$ rman target /
Recovery Manager: Release 12.2.0.1.0 - Production on Wed Jul 5 14:03:28 2017
Copyright (c) 1982, 2017, Oracle and/or its affiliates. All rights reserved.
connected to target database: SE122 (DBID=679484639)
RMAN> SELECT banner FROM v$version WHERE banner LIKE 'Oracle%';
using target database control file instead of recovery catalog
BANNER
-----
Oracle Database 12c Standard Edition Release 12.2.0.1.0 - 64bit Production
RMAN> SHOW DEVICE TYPE;
RMAN configuration parameters for database with db_unique_name SE122 are:
CONFIGURE DEVICE TYPE DISK PARALLELISM 1 BACKUP TYPE TO BACKUPSET; # default
RMAN> CONFIGURE DEVICE TYPE DISK PARALLELISM 2;
new RMAN configuration parameters:
CONFIGURE DEVICE TYPE DISK PARALLELISM 2 BACKUP TYPE TO BACKUPSET;
new RMAN configuration parameters are successfully stored
```

```
RMAN> BACKUP TABLESPACE system, sysaux;
Starting backup at 05-Jul-2017 14:07:27
using target database control file instead of recovery catalog
RMAN-06908: warning: operation will not run in parallel on the allocated channels
RMAN-06909: warning: parallelism require Enterprise Edition
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=59 device type=DISK
...
Finished backup at 05-Jul-2017 14:07:31
```

Note the grammatical error in RMAN message 6909 that Oracle Corp. hasn't fixed in more than a decade (since at least 10gR2).

Bash Code for Parallel RMAN Backup

Figure 1 shows 28 lines of Bash, SQL*Plus, and RMAN code that suffice to implement a parallel backup for Standard Edition. Certainly the code is no more than a proof of concept. In a nutshell, the code:

- Accepts the parallel degree as a command-line argument (line 2).
- Initializes the remainder for a subsequent modulus calculation (line 3).
- Loops over remainder values from 0 up to, but not including, the parallel degree (line 2).
- Selects a numbered list of data files sorted by decreasing size, and filters only those files where the remainder calculation on the file's position matches the current value of the remainder inside the loop. This ensures that the workload is spread more or less evenly (lines 5 to 24).
- Invokes RMAN in the background handing over a list of data files (line 25).
- Increments the remainder such that the next iteration will retrieve another subset of data files to be handled in parallel (line 26).
- Waits for all background processes (RMAN invocations) to finish (line 28).

```

1 #!/bin/bash
2 PARALLEL_DEGREE=$1
3 REMAINDER=0
4 until [ $REMAINDER -eq $PARALLEL_DEGREE ]; do
5     sqlplus -s /nolog <<EOF
6         CONNECT / AS SYSDBA
7         set heading off
8         set feedback off
9         set pagesize 0
10        set verify off
11        define parallel_degree=$PARALLEL_DEGREE
12        define remainder=$REMAINDER
13        spool rman_process_&&remainder..rcv
14        SELECT 'BACKUP DATAFILE ' || file# || '; # size=' || size_gb || ' GiB'
15        FROM
16            (SELECT rownum AS row_nr, f.*
17             FROM (
18                 SELECT file#, round(bytes/1048576/1024, 1) AS size_gb
19                 FROM v$datafile
20                 ORDER BY bytes DESC) f)
21        WHERE mod(row_nr, &&parallel_degree)=&&remainder;
22        spool off
23        exit
24    EOF
25    rman target / cmdfile rman_process_${REMAINDER}.rcv log rman_process_${REMAINDER}.log &
26    REMAINDER=$((REMAINDER + 1))
27 done
28 wait
"se_par_bkp.sh" 28 lines --100%--
28,1 All

```

Figure 1: Bash proof of concept code for parallel backup with Oracle Database Standard Edition

The database in the test environment has 65 data files. The following command uses four parallel invocations of RMAN to speed up the backup:

```
$ ./se_par_bkp.sh 4
```

The workload is split across four RMAN scripts.

```

$ head -1 rman_process_?.rcv
==> rman_process_0.rcv <==
BACKUP DATAFILE 7; # size=32 GiB
==> rman_process_1.rcv <==
BACKUP DATAFILE 32; # size=32 GiB
==> rman_process_2.rcv <==
BACKUP DATAFILE 5; # size=32 GiB
==> rman_process_3.rcv <==
BACKUP DATAFILE 46; # size=32 GiB

```

All scripts combined contain 65 BACKUP commands for a single data file. Thus the backup covers the entire database:

```

$ wc -l rman_process_?.rcv
16 rman_process_0.rcv
17 rman_process_1.rcv
16 rman_process_2.rcv
16 rman_process_3.rcv
65 total

```

Of course there are no duplicate commands inside the RMAN scripts:

```

$ cat rman_process_?.rcv|sort -u|wc -l
65

```

So here's a complimentary parallel backup for Standard Edition environments. Admittedly it's merely a prototype, but it could easily be integrated into an existing RMAN control script.

RmanJ and Parallel Operations in an Oracle Standard Edition Environment

RmanJ excels when used to enable parallel copy, backup, and restore operations with Oracle Database Standard Edition, thus allowing RMAN-based operations to complete orders of magnitude faster than serial operations offered by Oracle RMAN in a Standard Edition environment.

RmanJ overcomes the single-channel limitation of Oracle SE by controlling many invocations of RMAN in parallel using Java's support for multithreading.

Load balancing across many invocations of RMAN occurs at the data file level. All the data files involved in an RMAN operation are added into a thread-safe queue. Multiple threads then dispatch individual data files to available RMAN invocations. Whenever an RMAN invocation finishes handling one of the data files, it signals completion using the message code RMAN-03091. The message is received by one of RmanJ's threads, which then hands off the next data file to RMAN until all data files have been processed. Note that RMAN invocations are reused to optimize performance and to avoid the overhead of starting and stopping RMAN processes. There is a limitation of about 64 concurrent RMAN processes. See errors RMAN-06900 and ORA-19921: maximum number of 64 rows exceeded. The RMAN command DUPLICATE cannot be parallelized by RmanJ. However a manual parallel database duplication can be implemented using RMAN commands and RmanJ.

RmanJ accepts two different ways of setting the parallel degree: by property value or by command-line switch argument. The property `parallelism` sets the parallel degree for copy, backup, or restore operations, i.e., the number of RMAN invocations with one channel per RMAN process. The default parallel degree is one (i.e., no parallelism).

Alternatively, the command-line switch `-p` may be used to set the degree of parallelism. If used, the argument of the command-line switch `-p` overrides the property value.

RMAN supports automatic control file backups. If enabled, the control file is backed up at the conclusion of each BACKUP command. Thus a recent control file backup is always available, and synchronization of control file records with an RMAN catalog schema becomes optional. When no catalog is available and backups are written to tape, automatic control file backups are essentially mandatory. As a last resort the package `DBMS_BACKUP_RESTORE` may be called directly to retrieve a control file backup from tape, but this approach is undocumented and cumbersome. RMAN locates the most recent automatic control file backup simply by searching for files that conform to the naming pattern `c-%I-<current-date>-<hexadecimal sequence nr.>` where `%I` is the RMAN substitution variable for the database ID and the hexadecimal sequence number is in the range `0x00` to `0xFF`, which allows for up to 256 unique file names per day. Should the need arise to create more than 256 automatic control file backups within a single day, the file with sequence number `0xFF` is overwritten.

Automatic control file backups aren't crucial for backups to a file system since backup pieces in a file system can always be registered into a control file by RMAN. However, automatic control file backups are recommended since they ensure that each

backup operation concludes with a backup of the control file.

When parallelizing backups with RmanJ, automatic control file backups should be disabled. Otherwise, a control file (and probably SPFILE) backup will occur after each BACKUP command completes. That would cause unnecessary delays and wasted disk or tape storage. Since automatic control file backups provide an extra level of safety, they should not be disabled permanently but only during backups that are parallelized using RmanJ. The RmanJ properties `backup_prologue` and `backup_epilogue` may be used to temporarily disable automatic control file backups before the actual parallel backup and to re-enable them when the backup is complete.

Maximum Degree of Parallelism

The Oracle RDBMS has a built-in limitation regarding the bookkeeping performed by each invocation of RMAN. The status of RMAN operations is recorded in the control file, and data access is provided via the view `V$RMAN_STATUS`. Access to the underlying storage area in the control file is limited to 64 concurrent RMAN processes. Data loss occurs when this limit is exceeded. As a result reports provided by RmanJ will not reflect the correct amount of data handled by all concurrent RMAN sessions. To avoid data loss a parallel degree greater than 60 should not be used. In practice it is quite unlikely to observe a further speedup beyond a parallel degree of around 16 processes; hence, this limitation has no practical relevance. The following error messages will be reported when the parallel degree is too high:

```
RMAN-06900: WARNING: unable to generate V$RMAN_STATUS or V$RMAN_OUTPUT row
RMAN-06901: WARNING: disabling update of the V$RMAN_STATUS and V$RMAN_OUTPUT rows
RMAN-06003: ORACLE error from target database:
ORA-19921: maximum number of 64 rows exceeded
```

Monitoring and Multiple Block Sizes

An Oracle database may have tablespaces with multiple block sizes. Note that the block size of archived logs is always 512 bytes. The ensuing discussion applies solely to data files and not to archived logs. Oracle RMAN does not provide information on the block size of the data files it is processing via the view `GV$SESSION_LONGOPS`. In order to report the throughput of an RMAN operation in MB/s (megabytes per second), RmanJ needs the block size of the data files that RMAN is processing. There are two scenarios where the block size can be determined:

- Single block size database (i.e., the database has only one block size that is specified using the initialization parameter `db_block_size`)
- Multiple block size database with parallelization performed by RmanJ.

When parallelization is performed by RMAN itself, as in Oracle Enterprise Edition, RmanJ will not be able to calculate throughput metrics for operations against databases with multiple block sizes. As a workaround, use RmanJ parallelization instead of multiple channels in RMAN itself. Thus throughput metrics will be available for parallel backups of multiple block size databases using Oracle Enterprise Edition.

Figure 2 shows sample monitoring output from an incremental Standard Edition database backup that ran with a parallel



DATABASE MANAGEMENT SOLUTIONS

Develop | Manage | Optimize | Monitor | Replicate

Maximize your
Database Investments.



Quest™

Sid	Serial#	Spid	Start Time	Operation	Compl.	Wait Status	Days Remaining	Elas.(s)	MB/s
16	33719	28447	2017-07-12 21:32:52.0	incremental datafile backup	25.8%	resmgr:cpu quantum>0s	0 1:40:44.0	2100	4.02
1540	39559	28454	2017-07-12 21:35:57.0	incremental datafile backup	24.1%	resmgr:cpu quantum=.300145s (0s ago)	0 1:40:33.0	1915	4.03
29	34389	28451	2017-07-12 21:28:09.0	incremental datafile backup	52.6%	resmgr:cpu quantum=.293767s (0s ago)	0 0:35:52.0	2383	3.79
1548	23351	28450	2017-07-12 21:27:02.0	incremental datafile backup	59.0%	resmgr:cpu quantum>0s	0 0:28:24.0	2450	
46	47051	28453	2017-07-12 21:30:24.0	incremental datafile backup	59.7%	resmgr:cpu quantum>0s	0 0:25:15.0	2248	4.01
1546	51209	28449	2017-07-12 21:29:25.0	incremental datafile backup	80.3%	resmgr:cpu quantum>0s	0 0:9:25.0	2307	3.73
1533	48829	28438	2017-07-12 21:26:59.0	incremental datafile backup	98.6%	resmgr:cpu quantum>0s	0 0:0:36.0	2453	
25	14753	28437	2017-07-12 21:26:59.0	incremental datafile backup	98.7%	resmgr:cpu quantum>0s	0 0:0:33.0	2453	
<hr/>									
Status	Operation			Object Type	Input GB	Output GB	Input/Output GB		
COMPLETED	BACKUP			DATAFILE INCR	13.414063	0.935104	14.349167		

Figure 2: RmanJ Monitoring Output

degree of eight based on RmanJ. Note that Standard Edition also lacks block change tracking, rendering incremental backups orders of magnitude slower than with Enterprise Edition.

Throughput is under the column heading “MB/s.” RmanJ lacked sufficient input data for calculating the current throughput for some of the data files. Hence only five of eight data files that are backed up in parallel have an associated throughput value. The throughput itself of around 4 MB/s per data file is quite low since in this case backups were written to an NFS volume on an aging NFS server in a test environment.

Summary

The second and final installment on the subject of controlling Oracle Recovery Manager shows that it is relatively easy to overcome the serial backup limitation of Oracle Database Standard

Edition. Shell scripting is sufficient to attain the goal. However, more elaborate features can be built more easily using a more sophisticated programming language such as Java. The latter approach is employed by RmanJ in order to calculate backup throughput and remaining time of RMAN operations. An advanced feature of RmanJ that was not covered is the ability to trigger archived log backups when a certain fill level in an archive destination is reached. This feature supports file systems as well as ASM disk groups as archive destinations. Maybe a future release of RMAN will support fill level-dependent archived log backups natively, but for now additional software such as RmanJ is required for this and other features presented in the article. ▲



Norbert Debes is the principal at ORADBPPO GmbH, Germany

© 2017, Norbert Debes



Dr. DR

by Rich Parsons



Dr. DR is brought to you by Axxana.

NoCOUG Summer Conference

Session Descriptions

For the most up-to-date information, please visit <http://www.nocoug.org>.

—Keynote—

Why You Need to Upgrade to Oracle 12c to Easily and Seamlessly Migrate to Oracle Cloud—Charles Kim, Viscosity. . . 9:30

In 2017, Oracle adopted its Cloud First policy. All new innovations in their flagship database technology stack will be deployed first in Oracle Public Cloud (Oracle Cloud) and eventually released for general availability to the public. Oracle deployed Oracle Database 12c Release 2 to Oracle Cloud in September 2016 and released it to the public in March 2017. Having your database on the cloud (Platform as a Service or PaaS) in Oracle 12c R2 gives you access to new features much earlier than others. Companies can see immediate and significant ROI by upgrading to Oracle 12c while preparing to migrate to Oracle Cloud:

- With Oracle's multitenant option, you can leverage all the latest features, such as cloning and replication features for PDBs (Pluggable Databases), that enable the journey to a private, hybrid, or public cloud environment. With 12c R2, you can have hot-cloning (the source no longer needs to be read-only), refresh clones, and flashback PDB only, and build PDB-level standby plans. You can tightly control SGA, PGA, and I/O resources for individual PDBs, including how to construct standard PDB resource profiles that can be leveraged by multiple PDBs with similar resource demands.
- Utilize Information Lifecycle Management (ILM) features like Automatic Data Optimization (ADO) and heat maps, which work within individual PDBs for effective management of storage resources based on actual data usage patterns.
- Achieve another level of data scalability and horizontal partitioning with sharding.
- Eliminate downtime with enhanced online features for table and partitioning and index maintenance.
- Scale performance exponentially with In-Memory (IM) column store; in Oracle 12c R2 there are additional In-Memory enhancements (including IM virtual columns)
- Tackle partitioning and index advancements in 12c and 12c R2.
- Secure your entire database with online encryption capabilities.
- In Oracle 12c R2, object names for users, roles, tables, columns, indexes, constraints, and so on increase from 30 bytes to 128 bytes with few limitations, to easily migrate from other databases.

Leverage approximate query functions for performance.

By adopting Oracle Database 12c Release 2 on-premises at your data center, you will have already mitigated the majority of the security concerns for cloud adoption. Not only will you reduce your operating expenses (OPEX) and achieve new levels of availability, but you can also easily migrate your databases to Oracle Cloud and start reducing capital expenses (CAPEX).

Charles Kim is an Oracle ACE Director, a VMware vExpert, an Oracle Certified DBA, a Certified Exadata Implementation Specialist, and a Certified RAC Expert. Charles specializes in RAC, Exadata, Cloud, and Virtualization and is the co-author of nine books on Oracle Databases. Charles holds certifications in Oracle, VMware, Red Hat Linux, and Microsoft, and has over 25 years of Oracle experience on mission- and business-critical databases. Charles presents regularly at local, regional, national, and international Oracle conferences, including IOUG Collaborate, VMware World, and Oracle OpenWorld, on topics of RAC, ASM, Linux best practices, Data Guard best practices, VMware best practices, Oracle on VMware virtualization, and 7x24 high-availability considerations. Charles is the founder and president of Viscosity North America, a niche Oracle consulting company specializing in Exadata, RAC, virtualization, and Oracle Cloud. Charles is also the president of the Cloud Computing SIG for the Independent Oracle User Group.

—Breakout Track I—

Containerizing Oracle: Why You Should Be Thinking About It If You Aren't—Debaditya Chatterjee, Robin Systems . . . 11:30

Currently, DBAs mostly use traditional methods like hypervisor-based virtualization to consolidate databases in their data centers. With the introduction of the Oracle Multitenant option, DBAs have another choice for consolidation, but there are still concerns about performance loss and manageability challenges. In this session we will explore these prevalent approaches and elaborate on how to use container technologies like Docker and LXC to overcome some of the current roadblocks and consolidate databases without compromising performance while guaranteeing isolation and no manageability changes. We will demonstrate how to run an Oracle database using Docker and how we can simplify some of the most common lifecycle management tasks using the robin cloud platform.

Deba Chatterjee is the director of product management at Robin Systems. Prior to his current role he was the product manager for the Oracle multitenant option and Oracle diagnostics and tuning packs. Before product management, Deba worked for the performance services team in Oracle Product Development IT, where he was responsible for the performance of large data warehouses. He has previously worked at Oracle Consulting; Oracle India; Michelin Tires in Clermont-Ferrand, France; and Tata

Consultancy Services. Deba has a master's degree in technology management, a joint program by Penn Engineering and Wharton Business School.

Are Your End Users Happy with Your Test Data Refresh

Solution?—Alberto Sigismondi, Veritas Technologies13:00

Our end users are demanding faster access to Oracle data for TestDev. They want it when they need it, and they don't want to ask for it: they want to access the data themselves or via their development application toolchain. Also, many of our application development teams are looking to move their infrastructure to the cloud, so data flow on-premises becomes more and more difficult, if not impossible. And yes, how do we deal with native DBs deployed in the cloud? This session will give you a wider perspective on how to approach these situations in your organizations and how Veritas might be able to help.

Alberto Sigismondi is a senior manager of product management at Veritas Technologies.

Goodbye Passwords: Securing Oracle with SSL

—Harris Baskaran, Google14:30

At Google we take security seriously. Excessive privileges, shared accounts, lack of meaningful audit trails, and plain text passwords over the wire are common nightmares for DBAs everywhere who are trying to secure data. Security is hard. Given the spectrum of constant data breaches, securing database access is a big challenge. We will explore how database engineers can address modern Oracle database authentication and authorization challenges. We will show how two-factor authentication, which has become the standard for your corporate account, can be used for SQL developer and other SQL tools to connect to databases securely and access data in a secure way. We will also explore how to automate user and privilege management. Finally, we will show how to put all components together and sunset passwords. Stop worrying and start loving database security.

Harris Baskaran is a database administrator with extensive experience in Oracle databases and supporting tools.

—Breakout Track II—

Oracle Backup and Restore Taken to the Next Level

—Dan Fischer, Veritas Technologies11:30

Oracle's Recovery Manager (RMAN) feature is recognized for its rock-solid backup and restore functionality. But Veritas has taken RMAN to a higher level. In this technical session we will talk about new and innovative ways that Veritas engineers have created technologies based on RMAN to improve your Oracle backups. Are you using the cloud as part of your Oracle data-protection strategy? We've got you covered. We make the cloud easy. We will show you an easy solution to integrate cloud backups as part of your backup processes. We will also offer guidance for the suggested methods of utilizing the cloud.

Dan Fischer is an information protection architect at Veritas.

Advanced SQL Techniques: Advanced Grouping

—Jared Still, Pythian13:00

The CUBE, ROLLUP, and GROUPING SETS extensions to SQL make querying and reporting easier and faster. CUBE, ROLLUP, and grouping sets produce a single result set that is

equivalent to a UNION ALL of differently grouped rows. ROLLUP calculates aggregations such as SUM, COUNT, MAX, MIN, and AVG at increasing levels of aggregation, from the most detailed up to a grand total. CUBE is an extension similar to ROLLUP, enabling a single statement to calculate all possible combinations of aggregations. The GROUPING SETS extension lets you specify just the groupings needed in the GROUP BY clause.

Recognized for his numerous presentations at Oracle conferences like RMOUG, Hotsos, IOUG, and more, Jared Still is well-known within the Oracle community. His colleagues and peers describe him as meticulous, as he values the importance of detailed, organized documentation in order to provide full transparency and clarity to his clients. Jared's dedication to identifying the root causes of issues strengthens his ability to troubleshoot intractable problems. When he isn't relishing in low-level, in-depth technical projects, Jared can be found auto racing, which has been a passion of his for many years.

Advanced SQL Techniques: Recursive SQL

—Jared Still, Pythian14:30

Version 11g Release 2 introduced recursive subquery factoring or the recursive with clause. This is an extension to the SQL syntax with which you can do recursive/hierarchical queries. However, since version 2, Oracle has had the connect-by clause for hierarchical queries. At first glance, the connect-by and the recursive-with seem very similar in what they can do. But on a second look, there are some really interesting differences.

Recognized for his numerous presentations at Oracle conferences like RMOUG, Hotsos, IOUG, and more, Jared Still is well-known within the Oracle community. His colleagues and peers describe him as meticulous, as he values the importance of detailed, organized documentation in order to provide full transparency and clarity to his clients. Jared's dedication to identifying the root causes of issues strengthens his ability to troubleshoot intractable problems. When he isn't relishing in low-level, in-depth technical projects, Jared can be found auto racing, which has been a passion of his for many years.

—Breakout Track III—

Migrating PeopleSoft to Oracle Public Cloud

—Marcus Bode, SpearMC Consulting11:30

Organizations face comprehensive technical and business challenges in assessing, selecting, and deploying various cloud strategies for their PeopleSoft applications. This session covers our experience with a current PeopleSoft customer. We will explain how we planned for the migration to Oracle Public Cloud. We will discuss best practices from a planning, architecture, and execution standpoint—best practices co-developed by SpearMC and Oracle to “lift and shift” on-premises customers to the cloud.

Marcus Bode is a founder and managing director at SpearMC.

AWS on Steroids: Cloud Formation Templates

—Paul Marcelin, Independent Consultant13:00

Learn to write Cloud Formation templates to create, secure, and reconfigure an Amazon Relational Database Service (RDS) Oracle database instance. After reviewing different infrastructure as code options, I will outline the structure of a Cloud For-

mation template, propose naming conventions, and show you how to deploy and reconfigure basic AWS resources with your templates. You will receive samples to use as a starting point for your own work, and pointers to sources of additional information. My goal is to help you reduce manual work in the AWS Console.

Paul Marcelin is an expert in SQL, PostGIS, Python, C, UNIX/Linux shell scripting, SaltStack, and Amazon cloud infrastructure (AWS). He has 20 years' experience working with databases, from Oracle to Postgres to MongoDB, for employers and clients ranging from Lawrence Livermore Lab, Pacific Bell, Levi's, and Weotta.com (a natural language search startup) to Sindeo (an online mortgage marketplace).

Hands-on with Oracle MySQL Cloud Service

—Kathy Forte, Oracle14:30

This hands-on lab explores setting up a sample application using Oracle MySQL Cloud Service. Attendees step through each of the options available and get recommendations on how to size instances and monitor them with MySQL Enterprise Monitor.

Kathy Forte is a MySQL Solutions Architect at Oracle. ▲

If your Member ID is 6445130, 6517753, 6517753, 6544564, 6561843, 6562416, 7251931, 8626876, 11266275, 29395063, or 29431017, you have won a 64 GB flash drive emblazoned with the NoCOUG logo. You must attend the summer conference on August 3 in order to collect your prize.

RmanJ by ORADBPRO


Blazing fast parallel RMAN backups for Oracle® Database Standard Edition. Professional-grade RMAN control-utility for all editions (EE, SE, XE). No programming required to control RMAN.

- Fast multi-channel RMAN operations for Oracle SE (backup, restore, duplication)
- Standby database solution for Oracle SE for high availability (optional delayed apply)

RmanJ Features:

- Password encryption (target, catalog, and auxiliary connections)
- Monitors throughput (MB/s), data volume, and remaining time of RMAN operations
- Archived log backup depending on the destination's fill level (ASM, file system)
- E-mail notification
- Timing of all RMAN operations
- Platform-independence through Java
- Licenses available from 490 USD

Please visit <http://www.oradbpro.com> for additional information on RmanJ. Trial licenses available. Please send requests to info@oradbpro.com.



TriCore Solutions: DBA Pro Service



DBA PRO BENEFITS

- *Cost-effective and flexible extension of your IT team*
- *Proactive database maintenance and quick resolution of problems by Oracle experts*
- *Increased database uptime*
- *Improved database performance*
- *Constant database monitoring with Database Rx*
- *Onsite and offsite flexibility*
- *Reliable support from a stable team of DBAs familiar with your databases*

CUSTOMIZABLE SERVICE PLANS FOR ORACLE SYSTEMS

Keeping your Oracle database systems highly available takes knowledge, skill, and experience. It also takes knowing that each environment is different. From large companies that need additional DBA support and specialized expertise to small companies that don't require a full-time onsite DBA, flexibility is the key. That's why TriCore Solutions offers a flexible service called DBA Pro. With DBA Pro, we work with you to configure a program that best suits your needs and helps you deal with any Oracle issues that arise. You receive cost-effective basic services for development systems and more comprehensive plans for production and mission-critical Oracle systems.

DBA Pro's mix and match service components

Access to experienced senior Oracle expertise when you need it

We work as an extension of your team to set up and manage your Oracle databases to maintain reliability, scalability, and peak performance. When you become a DBA Pro client, you are assigned a primary and secondary TriCore Solutions DBA. They'll become intimately familiar with your systems. When you need us, just call our toll-free number or send email for assistance from an experienced DBA during regular business hours. If you need a fuller range of coverage with guaranteed response times, you may choose our 24 x 7 option.

24 x 7 availability with guaranteed response time

For managing mission-critical systems, no service is more valuable than being able to call on a team of experts to solve a database problem quickly and efficiently. You may call in an emergency request for help at any time, knowing your call will be answered by a TriCore Solutions DBA within a guaranteed response time.

Daily review and recommendations for database care

A TriCore Solutions DBA will perform a daily review of activity and alerts on your Oracle database. This aids in a proactive approach to managing your database systems. After each review, you receive personalized recommendations, comments, and action items via email. This information is stored in the Database Rx Performance Portal for future reference.

Monthly review and report

Looking at trends and focusing on performance, availability and stability are critical over time. Each month, a TriCore Solutions DBA will review activity and alerts on your Oracle database and prepare a comprehensive report for you.

Proactive maintenance

When you want TriCore Solutions to handle ongoing proactive maintenance, we can automatically access your database remotely and address issues directly — if the maintenance procedure is one you have pre-authorized us to perform. You can rest assured knowing your Oracle systems are in good hands.

Onsite and offsite flexibility

You may choose to have TriCore Solutions consultants work onsite so they can work closely with your own DBA staff, or you may bring us onsite only for specific projects. Or you may choose to save money on travel time and infrastructure setup by having work done remotely. With DBA Pro we provide the most appropriate service program for you.



CALL (888) 239-7775 OR VISIT www.TriCoreSolutions.com

NoCOUG

P.O. Box 3282
Danville, CA 94526

RETURN SERVICE REQUESTED

SOLVE ORACLE DB NIGHTMARES

TRAINING THAT GIVES YOU SOLUTIONS

- Online training by Craig Shallahamer
- Skills assessment & certification
- How-to webinars
- 24/7 unlimited access
- Priority response
- Learning paths



Use coupon code **NOCOUG10** for 10% off!
Questions? Contact support@orapub.com.



Find out more at orapub.com.