



Official Publication of the Northern California Oracle Users Group

# NoCOUG

J O U R N A L

Vol. 31, No. 1 • FEBRUARY 2017

## Que Sera?

### A Bright and Cloudy Future

*The only constant is change.  
See page 4.*

### Oracle RAC on Amazon EC2

*To dream the impossible dream.  
See page 13.*

### Oracle Licensing in the Cloud

*Obi-Wan will teach you well.  
See page 19.*

**Much more inside . . .**

ARUP



# 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](http://www.axxana.com)



**ORACLE** Gold Partner

ORACLE  
NO DATA  
APPLIANCE  
READY

ORACLE  
DATABASE  
READY

ORACLE  
EXADATA  
READY

ORACLE  
LINUX  
READY

# Que Sera, Sera

The painting featured on the front cover is by Oracle Database expert Arup Nanda. It won first prize this year in a statewide competition in Connecticut, in the category of Expressive Realism: art that expresses the perceptions and emotions of the painter. This was the first time in my life that a painting really spoke to me. The boy in the picture is me. I am living in a surreal dream. I follow the path that lights up before me. I do not know where the path will lead me, but there is only one path, so I must follow it.

A song comes to mind:

Que sera, sera!

Whatever will be, will be.

The future's not ours to see.

Que sera, sera.

What will be, will be.

All of us face an uncertain future—especially if we don't keep up with the changing needs of our employers—but NoCOUG conferences and the *NoCOUG Journal* can help. At the winter conference on February 9 at the Oracle conference center at Oracle headquarters in Redwood City, there will be two tracks on the Cloud plus a workshop on Edition-Based Redefinition by distinguished product manager Bryn Llewellyn. In addition, there will be a hands-on RAC installation and configuration workshop ("RAC Attack") led by our friends at Database Specialists, now part of TriCore Solutions. *Can you name one good reason not to attend the conference?* We hope to see you there!

This is the 121<sup>st</sup> issue of the *NoCOUG Journal*. Enjoy! ▲

—NoCOUG Journal Editor

P.S. Beginning on April 24, Arup's painting will be displayed in the rotunda of the Connecticut State Capitol. About 100,000 paintings were submitted, but only 300 paintings were selected to be exhibited throughout Connecticut and then sold. Arup submitted 12 entries and all were selected! Arup has pledged to donate all his proceeds to Asha for Education ([new.ashanet.org](http://new.ashanet.org)). The Sanskrit word "asha" means hope. Asha for Education focuses on basic education for underprivileged children in India, in the belief that education is critical to effect socio-economic change.

## Table of Contents

Interview .....	4	<b>ADVERTISERS</b>
Book Notes .....	8	Axxana .....
Special Feature .....	13	TriCore Solutions .....
Special Feature .....	19	TriCore Solutions .....
Conference Report .....	24	OraPub .....
		2
		11
		27
		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](mailto: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*

Eric Hutchinson  
*Webmaster*

Hanan Hit  
*President Emeritus*

Jeff Mahe  
*Vice-President*

Kamran Rassouli  
*Social Director*

Michael Cunningham  
*Training Director*

Naren Nagtode  
*President Emeritus*

Noelle Stimely  
*Membership Director*

Omar Anwar  
*Track Leader*

Roy Prowell  
*Publicity Director*

Saibabu Devabhaktuni  
*Conference Chair*

Sri Rajan  
*Secretary, Treasurer*

Tim Gorman  
*Board Advisor*

Tu Le  
*Track Leader*

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](mailto:journal@nocoug.org)

# A Bright and Cloudy Future

by John Kanagaraj



John Kanagaraj

*It seems to me that “cloud” is the return of 1970s-style time sharing in which users paid a fee for seconds of CPU time and a fee for kilobyte-months of disk storage. Is there any difference between the two except that the CPUs and disks are faster?*

Only the concepts of chargeback by usage and offloading of skillsets are similar to the 1970s time-sharing systems. To answer this question, let’s take a brief look at what a typical cloud looks like and use it to serve as an introduction to this subject.

Basically, a cloud provides three types of computing services—Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). This is best described using a diagram.

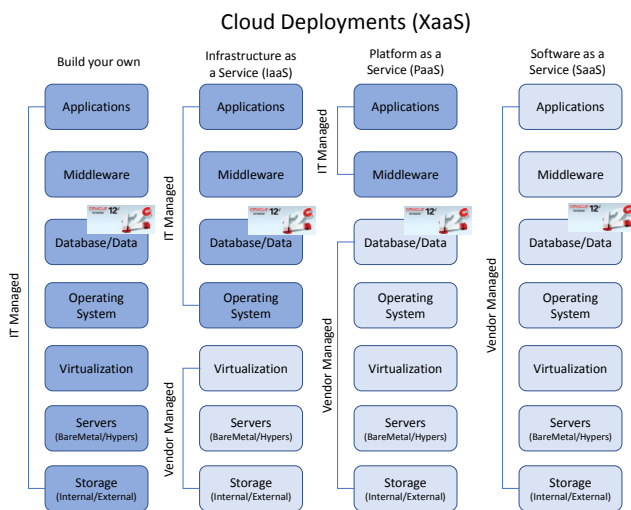


Figure 1.

Various layers in a typical database-based application are managed either by the internal IT department or by the external cloud vendor depending on the choice of “as a service.” Each type of service has varying implications for your IT staff, level of control, costs, and choice. If you use IaaS, you just want the hardware/storage/networking. You can select your database of choice, and your DBAs/app developers manage everything else. On the other end of the spectrum with SaaS, you are outsourcing every-

thing to the cloud vendor. With PaaS in the middle, you get all the underlying infrastructure and database support and are free to design/operate your own applications. With cloud-based services, there are also several third parties who provide what is known as “managed services.” These organizations are specialists who work with cloud vendors to support your organization’s journey to the cloud. They provide a spectrum of services and are a great way to start this journey.

Going to the cloud will help you get on the fast track to high availability and scalability with multi-data center support. The cloud vendor grows through acquisitions, and the maturity process is also pushed along by their customers. You benefit from all of this very quickly without investing your own resources. In other words, “a rising tide lifts all boats”!

This new paradigm, however, does come with its own set of challenges and problems.

*It seems to me that “cloud,” “loss of control,” “loss of visibility,” and “blind trust” are synonymous. Are organizations ready for it?*

Rather than ask if organizations are ready for this loss of control and visibility, let’s look at what has already happened in this area. If your organization uses Office (or Outlook) 365 or Salesforce, you are *already* in the cloud and are trusting Microsoft and Salesforce, respectively, with your critical email and sales/customer information. Before making the decision to use these services, your legal department and, indeed, your highest levels of management would (should!) have been involved in understanding the risks and weighing the costs, fallback/exit plans, and other factors before making that decision. When the choice to move IT to the cloud needs to be made, the same rigor and process should be followed but with some additional caveats.

The complexity and amount of effort involved in moving your internal IT application to the cloud are much higher. There’s a critical difference between using OL/Salesforce in the cloud and using the cloud to run your IT. No one in their right mind goes out to build their own email services (you just use OL or alternatives such as Gmail for Business), but converting/re-writing your specific IT application is an entirely different mat-

***Basically, a cloud provides three types of computing services—Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS).***



***There is an ongoing price war between the leading cloud vendors that is keeping margins either razor thin or even nonexistent, and only vendors with deep pockets and sustaining power remain in this business.***

ter. Here's where you need to use the right cloud model—i.e., choose between IaaS and PaaS to rebuild your application in the cloud or take the plunge to use standard Off-the-Shelf (OTS) software that happens to run in the cloud (SaaS). For example, if you are running an IT Service Management (ITSM) system such as BMC Remedy onsite, it may be a relatively easy task to move it to the cloud and let BMC manage this application end to end on its cloud.

On this topic, it is essential to keep in mind that your IT computing resources are intricately tied to a third party—namely your cloud vendor—and capabilities (or lack thereof) are dependent on this one vendor as well. For this reason, many organizations chose to build in multi-vendor capabilities in their journey to the cloud. Keep in mind that a few cloud vendors, including big ones such as Cisco and HP, have either announced their plans to exit this business or have already done so. This is primarily because there is an ongoing price war between the leading cloud vendors that is keeping margins either razor thin or even nonexistent, and only vendors with deep pockets and sustaining power remain in this business.

Finally, you also have the option to build a private cloud inside your own data centers. Microsoft, VMWare, and Dell provide such options if you want a commercial solution, or you can use open-sourced OpenStack if you want to roll your own. In this case, you get the advantages of a standardized, scalable computing environment that your organization fully owns and operates, and you need not be overly concerned about loss of control. Oracle also provides an option called “Oracle Applications Managed Cloud Service” that runs Oracle E-Business Suite–based applications in a private cloud onsite.

***Is the cloud ready for prime time? Can my 100-terabyte business-critical database run in the cloud? Or is the cloud suitable for small- and medium-sized businesses only? Dev databases without much data?***

Is the cloud ready for prime time? The answer is, “It depends.” As we noted before, SaaS applications such as Office 365 and Salesforce are *already* in prime time and have been so for a while. However, high-scale and large Oracle database–based applications haven't had as much success or as wide an adoption yet. Specifically, there is very limited support for Oracle RAC outside of the Oracle Cloud for reasons such as lack of support for Shared Storage and Multicast IP as well as Private Interconnect or dedicated networking. Amazon AWS and Microsoft Azure have stated that Oracle RAC can be supported, and VMWare has been supporting Oracle RAC for a while. (As far as I know, Google has not made any statements on this topic thus far about its Google Cloud Platform.)

In this context, the following statement from Oracle deserves consideration when you want to use Oracle RAC on other cloud vendors:

*“Oracle RAC is supported on all cloud environments supported by the Oracle Database, as long as the environment is able to provide the hardware, storage, and networking requirements as specified in the Oracle RAC and Grid Infrastructure documentation.*

*With the exception of the Oracle Cloud, Oracle has not tested nor certified Oracle RAC in these environments.”*

One obvious recommendation, then, would be to seriously consider using the Oracle Cloud platform when you require the scale or storage that only an Oracle RAC—or Oracle Exadata—based environment can provide. Oracle does provide IaaS and PaaS as well as SaaS for Oracle-specific applications, so you can choose OCP as a one-stop shop. Keep in mind that, in this case, you are then potentially locked in to Oracle Cloud.

So what is the scaling pattern for non-Oracle Cloud Platform–based solutions? One answer lies in the new Oracle Sharding feature that has been made available in Oracle Database 12c Release 2. Sharding splits the workload by the key so that the shard is spread out among many individual, loosely coupled databases. This new feature will allow you to scale out using standard single nodes with limited storage that are available in the non-Oracle world. In this case, rather than scaling *up* using

***There is very limited support for Oracle RAC outside of the Oracle Cloud for reasons such as lack of support for Shared Storage and Multicast IP as well as Private Interconnect or dedicated networking.***

Oracle RAC on a single cluster, you scale *out* by spreading the workload among many smaller single nodes. If the workload increases, you just add more shards. However, both the product version and the concept of sharding are relatively new and complex. As well, the sharding pattern is suited only for a limited set of applications.

Using the cloud for developing and testing databases is ideal. Such environments are not required to be highly scalable or to perform as well as production environments, and typically require scaled-down copies of data. As well, for Oracle/MySQL databases, you can use products such as Delphix in the cloud to save both DBA grunt work for creating database copies for dev/test and reduce/eliminate dev/test infrastructure on-site (this can be significant, as production involves a single copy, while dev/test can require up to three or four copies).

***How safe is it? From hackers? From foreign states? From dear old Uncle Sam?***

Safety on the big, bad internet is relative. The larger cloud providers have two significant advantages over the smaller IT departments when it comes to security—they can afford to employ extremely skilled and knowledgeable IT security professionals who can build and maintain secure systems. Since the environment is standardized, such cloud providers can respond to security threats and breaches quickly through mass patching. Cloud providers are also held to much higher security standards

***Using the cloud for developing and testing databases is ideal. Such environments are not required to be highly scalable or to perform as well as production environments, and typically require scaled-down copies of data.***

and are periodically tested/audited to ensure compliance. On the other hand, it can be argued that since there is concentration of computing and data from multiple organizations all in the same place, such cloud providers present a big target and that such visibility attracts the wrong kind of constant attention from hackers. This is true, and hence there is no such thing as a 100% secure cloud.

An additional concern that most organizations would have is this: since cloud providers offer shared computing resources, who else is sharing the network, storage, and computing resources that you are using? The possibility that a competitor who is sharing the same computing resource with you might deliberately hack into your network is very real—and concerning. Even more prevalent is data leakage due to bugs in the cloud software layers. The larger and more mature cloud vendors provide for isolation at various layers as a means of im-

***The cloud is a positive development for the Oracle Database professional. It takes away the low-value-added grunt work required to build, install, and maintain infrastructure and databases, and forces DBAs to move up the value chain.***

proving security in a shared environment via Virtual Private Cloud (VPC) capabilities.

It's interesting that you should ask about Uncle Sam. I read in a recent article that Bentonville, Ark., police issued a warrant to Amazon asking for data stored in its cloud from an Alexa Echo device to help prosecute a suspected murderer. While Amazon has declined to hand over said information, you can never be sure whether new laws related to terrorism could force cloud providers to hand over your information without your consent. Laws change as political changes occur, and since most cloud providers span the globe and run data centers in multiple countries, we will continue to see news items related to such issues from all over.

The question that follows, then, is "Who owns my data?" Usually the answer is "Your organization!" Cloud vendors should build in clauses that specifically state this and provide guarantees about who has what type of access to your data, and for what purposes. While you do run the risk of a rogue internal employee of the cloud vendor hacking away at your data, the more mature cloud vendors provide adequate audit and access logging that keeps such access transparent.

Before we close out this item, we do need to talk about disaster recovery and high availability. For small/medium businesses that run their IT from a single data center for various reasons

such as cost and resources, going to the cloud enables a heightened level of safety: single-data center high availability using multiple Availability Zones (AZs) and remote disaster recovery through geographically separated datacenters. This is a huge advantage.

***So who is running their workloads in the cloud?***

A well-known example is Netflix—they started onsite and migrated to the cloud, recently completing the migration of their billing system to the cloud. Today, the big cloud vendors host workloads from pretty much all segments, company sizes, and industries. Customer case studies can be found on their websites.

***Is the cloud cheaper? How much cheaper?***

The answer to this question lies in the way that using the cloud shifts the cost model from CAPEX to OPEX. Your finance department usually loves OPEX, as payments are somewhat predictable. The price you pay depends on what you want to do in the cloud. The cost model for SaaS applications will generally be much lower for an equivalent in-house-developed and -hosted application. You, however, lose control of the features and depend on the roadmap of the vendor. Computing and storage costs are generally much lower in the cloud, and as I mentioned before, cloud vendors are currently in a price war to take over as much territory as possible, so customers benefit. AWS has gone a step further with the introduction of the Lambda architecture in which computing/storage resources are triggered "just in time" and accumulate costs only when used.

***Do you recommend PaaS or IaaS?***

As you can see in the diagram at the beginning of this article, the models provide various levels of control. In general, the more customized your applications are and the more control you require, the more IaaS would meet your needs, and vice versa. I recommend that you choose PaaS over IaaS, if possible, and outsource as much work as possible to the cloud. And don't forget to review your applications in view of the plethora of SaaS applications available—and be prepared to make some hard decisions to let go.

***Who are the cloud service providers out there? How do I choose between them?***

There are many cloud service providers out there, but Amazon Web Services, Microsoft Azure, and Google Cloud Platform are the biggest and most well-known. For Oracle-centric workloads, you will need to add Oracle Cloud Services to this list. There are a lot of considerations, such as cost, ease of migration and use, company stability and commitment (see my earlier comment about big names exiting the business), availability of specific features such as Oracle RAC, and so on. I recommend that you consult reliable research and advisory firms such as Gartner, Forrester, and IDC to get some insights, and browse the case studies to look for patterns/industries and organizations that resemble yours.



***What does all this mean for Oracle Database professionals? Are their jobs safe? Do they have to reinvent themselves?***

In general, I see the cloud as a positive development for the Oracle Database professional. It takes away the low-value-added grunt work required to build, install, and maintain infrastructure and databases, and forces DBAs to move up the value chain, so they definitely have to learn new skills and adapt themselves. However, there are obviously many shades of gray here. In the end, fewer DBAs will be needed overall as this lower-level work is automated and outsourced to the cloud vendor, but the overall value of a skilled DBA who can adapt to the new order and serve as a bridge between the cloud vendor and the IT/application user base will increase.

***Thanks for taking the time to answer my questions. Honestly, my company is not ready for the cloud, but I want to stay ahead of the game by learning about it. Where do you recommend that I start?***


Thank you for the opportunity to share my thoughts here! The World Wide Web is full of news items, opinions, and blog entries as well as “marchitecture” (marketing/architecture) about

this new technology trend, so you need to be careful to apply appropriate filters to what you read. In general, use respected advisory sites such as Gartner and IDC to understand what’s happening in this segment of the industry, and then look at the top three or four cloud vendors to read up on their offerings. Most of these vendors provide a limited trial period, so you can get your feet wet before committing. Don’t forget to review the SaaS offerings from various vendors, and be prepared to make some hard decisions. If possible, use the managed service vendors and consultants to help you get started. Good luck to you and your organization in this journey! ▲

*John Kanagaraj is a longtime Oracle DBA and UNIX professional with more than three decades of experience in the IT industry. He is an Oracle ACE, an author and technical editor for many Oracle books, as well as a long-standing editor of SELECT Journal. Kanagaraj currently works at PayPal, providing data architecture and NoSQL database consultation, and can be reached via email at [john.kanagaraj@gmail.com](mailto:john.kanagaraj@gmail.com) or his LinkedIn account at <https://www.linkedin.com/in/johnkanagaraj>.*



## Dr. DR



## By Rich Parsons



Dr. DR is brought to you by Axxana.

# Study Guide for 1Z0-060: Upgrade to Oracle Database 12c

Book Notes by Brian Hitchcock

## Details

**Author:** Matthew Morris

**ISBN-13:** 978-1941404027

**Pages:** 262

**Date of Publication:** March 21, 2014

**Edition:** 1

**List Price:** \$14.99

**Publisher:** ODB Press

## Summary

This book is a good review of all the material that you need to know for the 12c DBA Upgrade exam. I have not (yet) taken the exam, so I can't (yet) comment on how much help I received from this book. As I described in my previous book review, this OCP exam is different from DBA Upgrade exams of the past. Previously, the exam was almost all about the new features of the latest version of the database. There were some questions covering general DBA topics not specific to the new database version, but you could miss all of them and still pass the exam. The 12c Upgrade exam is divided into two parts. The first part covers 12c new features as one would expect, but the second part is dedicated to what are called Core DBA Skills. You have to get a passing score for both sections to pass the exam.

Similar to the OCP exam guide I reviewed previously, this book does not try to introduce or explain the new features of the 12c database or the general DBA topics. This book is even more focused on reviewing material, and it is assumed that the reader already knows and understands what is being discussed.

I found many fewer errors than the previous book I read. It is interesting to me that this book, written and published independently, without any connection to Oracle Press, has so many fewer errors. Perhaps the author was more invested?

I want to stress that this OCP Study Guide, unlike the previous one I read, covers all the topics that will be on the exam. The official Exam Guide from Oracle Press makes no effort to cover the topics relevant to the second part of the exam, namely the general DBA topics.

I purchased the Kindle version of this book and, therefore, instead of chapters we simply have a list of topics or sections for each of the two parts.

Before we get to the first part, we have a Contents list, which are links to the sections instead of the traditional page number you would have for the beginning of each chapter in a printed book.

## What to expect from the test

This section covers the structure of the exam, the number of questions in each section, and the passing score for each as well as the time allowed for the exam. The author provides some very specific advice on how to manage your time during the exam based on his own experience. The section ends with the author's observation that most people who fail this exam will likely do so on the Core DBA Skills section. I haven't been a full-time core DBA in many years, so it will be interesting to see how I do on this section.

## What to expect from this study guide

This section covers some details about what you will get out of reading this study guide. We are told that while Oracle provides specific topics for the 12c new features, the topics provided for the Core DBA Skills sections are entirely vague. One example is discussed: Perform Daily Administration Tasks. What does this mean? What is included? How do you prepare to answer questions related to a topic this broad? I will do my best, and when I take the exam, I will be better able to comment.

We are reminded that this guide really does assume that the reader has a "reasonable level of familiarity with Oracle." I assume that when the author refers to "Oracle," he means the Oracle database. This guide and the exam are set up in a way that assumes those taking the exam have previously taken and passed one or more OCP DBA exams and therefore have experienced the OCP exam format. Perhaps the best summary of what to expect from reading this book is this: the guide will help the reader focus on what is most likely to be found on the exam.

## Additional Study Resources

In addition to this study guide, the author has set up a companion website with other resources that you can use to prepare for this exam. Note that the author has prepared study guides for other exams, and the website has additional material for those exams as well. We are told that this website does not, and will not, link to any unauthorized content such as brain dumps or purported copies of actual exam questions. I agree with the author's advice to read and study a variety of study materials while preparing for any OCP exam.

## Practice Questions

While we have a section that discusses practice exam questions, it turns out that this study guide does not provide any practice questions! The explanation is reasonable in that adding enough practice questions to really help would double the size of



the study guide. It turns out that the author has prepared practice tests to go along with this guide, but those practice tests need to be purchased separately. I have no issue with this, but I can see how some readers might, because the Oracle Press OCP study guides do have some practice questions. On the other hand, I am very happy to have multiple practice tests available to me. The practice exams cost \$14.50, which I think is fine. I realize others may think this is too much, but compared to the value of the time I am putting into preparing to take this OCP exam, spending \$15 once or twice is, in my opinion, a good value. If you decide to purchase and take one or more of these practice exams, please read all of the sections on the website to make sure you understand how they work. For example, there is no study mode; once started, you can't pause the exam. The goal of the author's practice tests is to simulate the actual OCP exam experience—not to be part of your study or review process.

---

### First Part: New Features of Oracle Database 12c

---

Each of the sections in this first part focuses on one aspect of the new features in 12c, reviewing what you need to know for the exam. Again, it is assumed that the reader already knows what the new features are and how they work. The information in each section is very brief. This isn't a criticism at all—the whole point of this book is to get the reader focused on what will be asked on the exam.

---

### Enterprise Manager and Other Tools

---

This section covers using EM Express, which is built into the database and doesn't rely on any background processes. A screenshot shows us what the main screen looks like, where we find four main menus for configuration, storage, security, and performance. Each of these menus is discussed. How to configure the HTTPS port for EM Express is described, and examples of the needed entry in the `tnsnames.ora` as well as the `init.ora` file are shown.

Specific aspects of Oracle Universal Installer (OUI) and the Database Creation Assistant (DBCA) are covered next. The minimum required disk space and memory for installing 12c are shown, followed by the pre-installation tasks needed. If you are running Oracle Linux 6, Oracle provides a pre-install RPM that makes the installation easier. Screenshots are presented covering the various steps of using the OUI. Two scripts need to be executed to complete the installation.

The next section covers the DBCA, and the author reminds us of the bad old days when we had to use the `CREATE DATABASE` statement to set up a database. The DBCA can be used in the interactive or silent mode, and the five steps it goes through to create a new database are shown with screenshots.

---

### Basics of the Multitenant Container Database (CDB)

---

We are told that the CDB is good for consolidation, reducing costs, faster implementation, simplified management, and better security. Tuning and patching are also improved. The details of what the terms “root,” “container,” “container database (CDB),” and “pluggable database (PDB)” mean in Oracle 12c are discussed. Each CDB has one root container, one seed PDB, and zero or more user-created PDBs. These terms can be confusing, but the material in this section is clear.

---

### Configuring and Creating CDBs and PDBs

---

How to create and configure CDBs and PDBs is shown next, starting with the planning process necessary to decide exactly what your application needs. There are some specific initialization parameters you must set as well as memory and disk space considerations. Using the DBCA simplifies the process. If you decide to use the `CREATE DATABASE` command, there are extra manual steps you must perform. The files that are created on disk are specified using new parameters `SEED_FILE_NAME_CONVERT` and `PDB_FILE_NAME_CONVERT`. There are prerequisites that must be met before creating a PDB and there are four different ways to create the PDB, each of which is discussed in detail. The optional clauses of the `CREATE PLUGGABLE DATABASE` command are covered in detail. Dropping a PDB and how to migrate a non-CDB into a PDB are also shown.

---

### Managing CDBs and PDBs

---

The section on managing CDBs and PDBs starts with explaining how to connect to each using a database service. There are a number of specific aspects of these services and how to set them up, and examples are shown using SQL\*Plus. Startup and shutdown of CDBs and PDBs is different from the pre-12c days. You can't start a PDB up without starting the CDB first, for example. The various open modes are covered and more SQL\*Plus examples are provided. We use instance parameters to configure the CDB, but there are some parameters that can be changed for each PDB within the CDB. Again, this material moves quickly because it is assumed that you already know how CDBs and PDBs work.

---

### Managing Tablespaces, Common and Local Users, Privileges and Roles

---

Once we have created a CDB and the PDBs that live inside it, we need to manage the associated tablespaces. There are many specifics: a CDB can only have one active UNDO tablespace, for example. The many options of the `ALTER PLUGGABLE DATABASE` command are explained. Managing users also becomes more complex. There are common users that exist in all the PDBs within the CDB and local users that exist only in one PDB. How to create both kinds of users and assign privileges to these users is described with examples. Common users have usernames that must begin with `C##`—except for Oracle-supplied common users, such as `SYS`, that don't.

---

### Backup, Recovery and Flashback for a CDB/PDB

---

Now that we have a CDB and multiple PDBs, how do we handle backup and recovery of the database changes? `RMAN` can make backups of the entire CDB and all of the PDBs inside of it in one operation. You can also back up one or more PDBs separately from the CDB. There are many ways to handle all of this, and various combinations of backup and recovery using `RMAN` are shown. It seems that only `RMAN` is discussed, so I assume this means that any exam questions regarding backup and recovery will be limited to using `RMAN`. Since I haven't used `RMAN`, I paid close attention to the examples shown.

---

### Information Lifecycle Management and Storage Enhancements

---

Oracle database 12c has two new features for Information Lifecycle Management (ILM). The first is the Heat Map, which

tracks which data is being used most often; this is needed to decide which data to move to slower, less expensive storage. The second is Automatic Data Optimization (ADO), which supports data compression and data movement policies. Once created, the database can use these policies and the Heat Map to work in the background to move and compress data. How to set this up is discussed and examples are shown. Several storage enhancements are included in 12c, such as moving a data file while the database is open.

### **In-Database Archiving and Valid-Time Temporal**

Another new feature of 12c is called Valid Time Temporal. Using two hidden date-time columns, specific rows of a table can be given a start and stop time that defines the time range when the rows are temporally valid. Applications can use this information to decide which rows are displayed. An example offered is insurance policies that have dates when coverage starts and expires. Somewhat related is In-Database Archiving, which uses a single hidden column to record which rows are valid. This is useful for data that isn't relevant anymore but needs to be kept for compliance purposes. The data that has been marked as archived can be compressed to reduce storage space. For both of these features, a detailed example is provided, including all of the needed SQL.

### **Auditing**

Auditing within the Oracle database is not a new feature, but it has been improved in that all the auditing is in one unified place. In previous versions the audit trails were in multiple locations. The new unified audit trail records audit information from many sources within the database, each of which is shown in a detailed list. Various security improvements have been created to make it difficult for anyone to alter the records stored in the audit trail. Many details of how to implement this are covered with SQL examples. This new auditing scheme also causes less load on the database because it first stores the audit records in the SGA and writes them to disk as needed.

### **Privileges**

In 12c we have several new system privileges: SYSBACKUP, SYSDBG, and SYSKM. These have been added to support separate users to administer RMAN, Data Guard, and encryption keys, respectively. As always, there are various restrictions with this new feature. A password file is required to make use of these new privileges. Some very specific modifications have been made to the existing privileges SELECT ANY DICTIONARY and UNLIMITED TABLESPACE. A new system privilege, PURGE DBA\_RECYCLEBIN, has been added; users don't need the SYSDBA privilege to purge the DBA recycle bin. For all of this, and the new privilege analysis feature, extensive examples are presented.

### **Oracle Data Redaction**

Data Redaction in 12c uses redaction policies that are set up on columns. This means that all applications accessing the columns will have the same redaction policy enforced. The redaction policies can use full, partial, regular expression, or random redaction. Many aspects of the DBMS\_REDACT package are discussed. Oracle provides several pre-defined character redaction shortcuts, an example being a shortcut to handle Social Security numbers.

### **RMAN and Flashback Data Archive**

In 12c, you can use RMAN to recover an individual table from backups. This process was significantly more complicated in previous versions of RMAN. Another improvement is that you can use SQL commands while in RMAN without needing to use the SQL keyword. The DESCRIBE command also works within RMAN. A new option for cloning a database is active DUPLICATE, which retrieves backup sets over the network to reduce the load on the target database. Another new option is NOOPEN for duplicate, which is used when you need to run upgrade scripts before opening the clone database.

### **Real-Time Database Operation Monitoring**

Oracle automatically monitors a SQL statement that consumes five seconds of CPU or I/O, but you can also define simple or composite database operations that will then be monitored. "Simple" is a single SQL statement and "composite" is the activity between a start and stop time in a database session. Previously, only simple operations were supported. You might want to use this to monitor a batch job or an ETL process. The data dictionary views involved and the function package DBMS\_SQL\_MONITOR are discussed, and an example of the SQL needed to set up this monitoring is reviewed.

### **SQL Tuning**

Adaptive Query Optimization is the formal term for the new features for SQL Tuning in 12c. The two new features are Adaptive Plans and Adaptive Statistics. Adaptive Plans allow the optimizer to watch while SQL is executing and change the execution plan. Adaptive Statistics allows the optimizer to examine the current execution and gather more statistics for use in future executions. Note that both of these features are iterative; that is, they monitor execution and modify plans and gather more statistics to make future plans better. Many aspects of these two new features are discussed, including the new views and packages involved.

### **Emergency Monitoring, Real-Time ADDM, Compare Period ADDM, and Active Session History (ASH) Analytics**

This section covers a lot of new features. Emergency Monitoring is for connecting to a database that is not responding and uses a proprietary mechanism to connect directly to the SGA to collect data that might help resolve the problem without a database restart. Real-Time ADDM is also used to look at databases that are not responding. It uses ASH recent activity gathered from the SGA. You can also use ADDM to compare performance for two time periods. ASH, too, has specific enhancements to help diagnose performance issues. This section is a good example of just how much the reader is assumed to know. Most of the acronyms are not defined; the reader must know about them ahead of time.

### **Resource Manager and Other Performance Enhancements**

Resource Management had to have some new features to deal with a CDB that can have multiple PDBs, all of which are competing for system resources. The author tells us that this has become much more complex in 12c due to multitenant. This works on two levels: for the CDB level and within each PDB. Each of the PDBs is given a number of shares, and this limits the percent-



age of overall resources that each PDB can consume. Within each PDB, these system resources are then managed among the various users of each PDB. Oracle 12c also introduces the multi-threaded architecture, which uses resources more efficiently. Enhancements to the Flash Cache are also discussed.

### Index and Table Enhancements

For indexing, you can now create multiple indexes on the same set of columns. These indexes do need to be different in some way, such as one is B-tree and the other is bitmap, and only one index is visible at a time. The optimizer will only see the one visible index when analyzing how to process SQL using the index. Table columns can also be made invisible such that only SQL that explicitly references the invisible columns will return data from those columns. Online operations have been enhanced so that you can redefine a table while the table remains online. Note that the table will be locked briefly at the end of the redefinition process, but we are told that this is only for a short time.

### ADR and Network Enhancements

There is a new DDL log file that lives in the ADR home and tracks each DDL statement that is executed once the `ENABLE_DDL_LOGGING` parameter is set to `TRUE`. Two files are generated, one in XML format and one in text. The format is the same as the traditional database alert log. A long list is provided showing all the DDL statements that will be logged. I think it might have been better to show what I assume is a shorter list of DDL statements that are not logged. A debug log has also been added for unusual conditions in the database that do not prevent normal operation. This reduces the amount of output that is sent to the alert log.

### Oracle Data Pump, SQL\*Loader, External Tables and Online Operations Enhancements

Data Pump now has an option called full transportable export, which means the entire database—including all the stuff needed to completely re-create the database. To make such an export, you use the `TRANSPORTABLE=ALWAYS` option along with the traditional `FULL` option. Various restrictions on this option are covered. Data Pump can also compress data as it is loaded into a table, and you can control the amount of compression applied during export. SQL\*Loader has a new express mode. This allows loading of data without a control file; only the target table is specified. Since there is no control file, many defaults are applied, such as the name of the data file.

### Partitioning Enhancements

Partitioning can now be set up using Interval-Reference partitioning. The parent table uses interval partitioning, while the child table uses reference partitioning; this gives us better partition modeling. Also new is the cascade functionality for `TRUNCATE PARTITION` and `EXCHANGE PARTITION` operations. Also enhanced are the online partition maintenance operations that can be performed on multiple partitions from one command. This affects the `ADD`, `DROP`, `TRUNCATE`, `MERGE`, and `SPLIT` operations. Partitions can also be moved or redefined while online. To speed up the process, global indexes can be rebuilt asynchronously in the background.

### SQL Enhancements

To help with the migration of databases to the Unicode character set, we now have the Database Migration Assistant for Unicode (DMU). The tools used for this in previous versions were `CSSCAN` and `CSALTER`. The DMU also has a validation mode that can be used to confirm the migration or check databases that are already using a Unicode character set. The row limiting clause has been updated: now you can skip a specified number of results before returning results, and the results will follow any `ORDER BY` clause in the SQL. You can also limit the data returned to a specified percentage of the results.

### Second Part: Core Administration

Note that these topics are also known as “Key DBA Skills” in the official OCP list of exam topics. This section has fewer major sections since the topics are not nearly as well defined as for the new features. Within each section, the discussion is more fluid since each topic is relatively general.

### Core Administration

This section covers a wide range of topics, starting with a discussion of the fundamentals of database architecture. This includes questions like “How does a database differ from the database instance?” The structures within the instance, the SGA, and the background processes are covered in detail as one would expect. Other topics covered include installation, configuration, daily administration tasks, patching, and several other areas that are part of the DBA job description. The description of the daily administration tasks makes it clear just how hard it is to prepare



Helping clients  
get the most out  
of their Oracle  
environments  
since 1999

**TriCore**  
solutions

The Application  
Management  
Experts

[www.TriCoreSolutions.com](http://www.TriCoreSolutions.com)

for the overall topic of core administration. It covers a lot, and many full-time DBAs will handle these tasks very differently. I found this section to be clear and easy to follow, which I hope means that my experience will serve me well for this part of the exam.

### Performance Management

Designing the database to optimize performance is discussed as well as how to monitor performance. How to place different parts of the database across disks is covered, along with how to identify performance issues and memory management. Application tuning and using resource management are also included.

### Storage

Here we review how to manage database structures, administer ASM, manage VLDB, and set up space management. We are shown screenshots from EM Express for managing tablespaces, undo, redo logs, archive logs, and control files. I was expecting to see a review of the details of storage—including blocks, extents, and segments—but that isn't covered. Given how broad the topic of storage is, I am not surprised that many things had to be eliminated. I assume that what is covered in this section reflects what I will need to know for the exam.

### Security

This section reviews how to set up a security policy, configure and manage auditing, create a password file, and work with encryption.

### Performing Post-Upgrade Tasks

This section covers how to migrate to unified auditing and execute other tasks that come up after an upgrade.

### Migrating Data by Using Oracle Data Pump

Offhand, I wouldn't have expected an entire section dedicated to this topic, but again, I assume this reflects what is on the exam. A detailed explanation follows of all the steps and restrictions that you need to be aware of to use Data Pump to populate a new, empty 12c database.

### Conclusion


This book does exactly what one would expect: it prepares the reader for the OCP exam. Until I take the exam, I can't say from experience how much it helped me, but I'm confident that it will have been helpful. As I have said, this book is not for someone that wants to learn about the new features of 12c. Each section moves quickly, and many acronyms are not defined let alone explained. If you haven't seen a lot of the Oracle database before, I think you will be overwhelmed. For those preparing for the exam, however, I think this book is excellent. ▲

*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/](http://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



### Dr. DR



### by Rich Parsons



Dr. DR is brought to you by Axxana.



# Oracle RAC on Amazon EC2

by Artem Danielov



Artem Danielov

The use of Amazon Elastic Compute Cloud (Amazon EC2) in the Amazon Web Services (AWS) cloud provides IT organizations with the flexibility and elasticity that are not available in the traditional data center. With AWS it is possible to bring new enterprise applications online in hours instead of months. Running the entire IT infrastructure in the public cloud is important in order to fully realize the competitive and cost advantages that the public cloud provides.

Ensuring high availability of backend relational databases is a critical part of the cloud migration strategy. Oracle RAC is the trusted high-availability solution for running mission-critical databases, and many IT organizations prefer to run it in the cloud without switching to cloud-native solutions.

Oracle RAC has the following infrastructure requirements that are not directly available in AWS:

- Shared high-performance storage accessible from all nodes in the cluster
- Multicast-enabled network between all nodes in the cluster
- Separate networks for different types of traffic: client, cluster interconnect, and storage

FlashGrid Storage Fabric and FlashGrid Cloud Area Network™ (CLAN) technologies address these requirements and enable mission-critical Oracle RAC clusters on Amazon EC2. This article provides architectural overview of the solution and can be used for planning and designing Oracle RAC deployments on Amazon EC2.

## Introduction to FlashGrid Software

High-speed shared storage is critical for seamless database infrastructure failure handling with zero downtime and zero data loss. FlashGrid Storage Fabric software enables high-speed shared storage in a variety of infrastructure environments, including bare-metal servers, virtual machines, and extended distance clusters, without the use of proprietary storage arrays. FlashGrid Storage Fabric adds the shared access required by Oracle RAC to the standard storage capabilities of Amazon EC2.

FlashGrid Cloud Area Network software enables migration of mission-critical applications to the AWS by bridging the gap between the standard network capabilities of the Amazon Virtual Private Cloud (VPC) and the networking requirements of Oracle RAC.

## Why Oracle RAC on AWS

AWS provides on-demand computing resources and services in the cloud, with pay-as-you-go pricing. AWS has been disrupting the traditional IT infrastructure world by removing the need to manage dedicated hardware and effectively shifting capital expenditures (CapEx) to operational expenditures (OpEx). After migrating to AWS, IT organizations may expect up to a 70% reduction in the TCO and massive increases in elasticity and agility of development according to an IDC report. By using AWS, customers can free valuable resources used to manage complex and costly datacenters, and repurpose them to focus on other highly strategic areas of the business.

Oracle RAC provides an advanced technology for database high availability. Many organizations, including most financial institutions and telecom operators where high availability and data integrity are of paramount importance, use Oracle RAC for running their mission-critical applications.

Oracle RAC is an active-active distributed architecture with shared database storage. The shared storage plays a central role in enabling automatic failover, zero data loss, and 100% data consistency, and in preventing application downtime. These HA capabilities minimize outages due to unexpected failures as well as during planned maintenance.

Oracle RAC technology is available for both large-scale and entry-level deployments. Oracle RAC Standard Edition 2 provides a cost-efficient alternative to open-source databases while ensuring the same level of high availability that the Enterprise Edition customers enjoy.

FlashGrid software brings the superior economics, flexibility, and agility of AWS to a broad range of Oracle RAC customers. It enables existing enterprise Oracle RAC customers to realize the full benefits of migrating their entire IT infrastructure to AWS. It also lowers entry barriers for new customers starting with small-scale database deployments.

## Supported Cluster Configurations

FlashGrid supports four RAC cluster configurations on Amazon EC2:

- Two RAC nodes in the same Availability Zone
- Three RAC nodes in the same Availability Zone
- Two RAC nodes across different Availability Zones
- Three RAC nodes across different Availability Zones

Support for configurations with more than four RAC nodes is planned in the future.

### Configurations with two RAC nodes

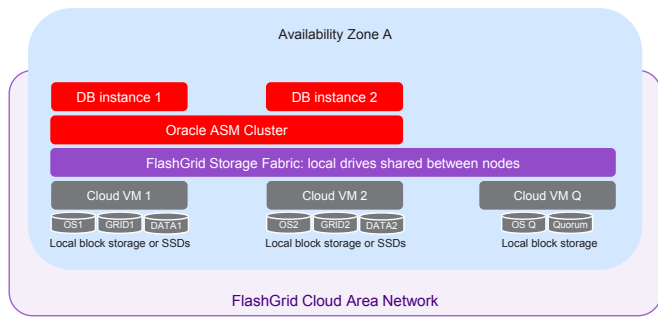


Figure 1. Two RAC nodes in the same Availability Zone

Configurations with two RAC nodes have two-way data mirroring using Normal Redundancy ASM disk groups. An additional EC2 instance is required to host quorum disks. Such clusters can tolerate the loss of any one node without database downtime.

### Configurations with Three RAC Nodes

Configurations with three RAC nodes have three-way data mirroring using high-redundancy ASM disk groups. However, a normal-redundancy ASM disk group is used for clusterware files (the GRID disk group). Such a cluster can tolerate the loss of any one node without database downtime. However, loss of a second node will result in downtime. The main reason for using three (vs. two) RAC nodes is the additional CPU and memory resources, and the additional storage read bandwidth.

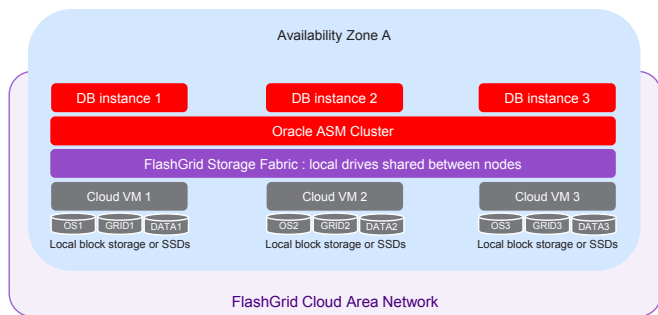


Figure 2. Three RAC nodes in the same Availability Zone

It is possible to have a three-node RAC cluster that can tolerate loss of two RAC nodes without database downtime, but such configurations are beyond the scope of this article.

### Same Availability Zone vs. Separate Availability Zones

Amazon Web Services consists of multiple independent Regions. Each Region is partitioned into several Availability Zones. Availability Zones consist of one or more discrete data centers, each with redundant power, networking and connectivity, housed in separate facilities. Availability Zones are physically separate, such that even extremely uncommon disasters such as fires, tornadoes, or flooding would only affect a single Availability Zone.

Although Availability Zones within a Region are geographically isolated from each other, they have direct low-latency network connectivity between them. The network latency between

Availability Zones is generally lower than 1 ms. This makes the inter-AZ deployments compliant with the extended-distance RAC guidelines.

Placing all nodes in one Availability Zone provides the best performance for write-intensive applications by ensuring network proximity between the nodes. However, in the unlikely event of an entire Availability Zone failure, the cluster will experience downtime.

Placing each node in a separate Availability Zone helps avoid downtime, even when an entire Availability Zone experiences a failure. The trade-off is a somewhat higher network latency and, in some cases, lower network bandwidth, which may reduce write performance. Note that the read performance is not affected, because all reads are served locally.

If you are placing nodes in separate Availability Zones, using a Region with at least three Availability Zones is highly recommended. The current number of Availability Zones for each Region can be found at <https://aws.amazon.com/about-aws/global-infrastructure/>.

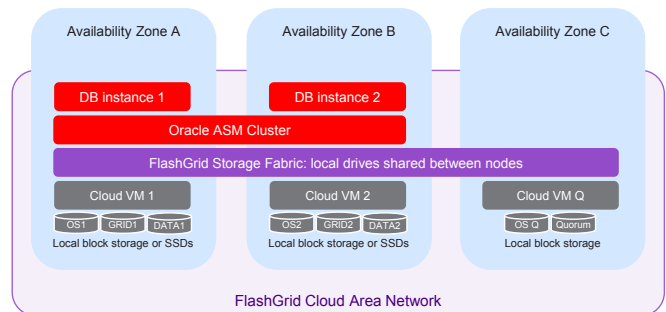


Figure 3. Two RAC nodes in separate Availability Zones

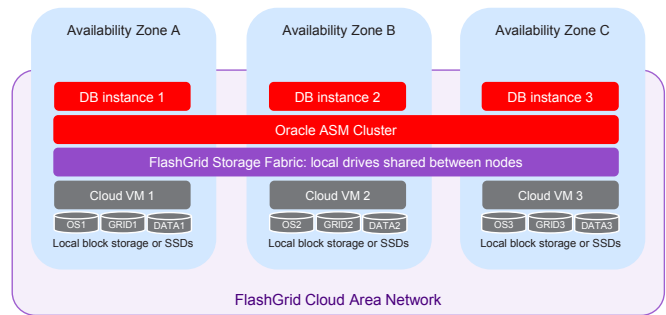


Figure 4. Three RAC nodes in separate Availability Zones

It is possible to deploy in a Region with only two Availability Zones for a two-node RAC cluster. However, in such a case the quorum server must be located in a different Region to prevent network partitioning scenarios. This configuration is beyond the scope of this article.

### How It Works

- FlashGrid Cloud Area Network enables high-speed overlay networks with multicast and bandwidth allocations, leveraging AWS networking features such as cluster placement groups and Elastic Network Adapters rated up to 20 Gbps
- FlashGrid Storage Fabric turns local drives (elastic block storage or local instance-store SSDs) into shared drives accessible from all nodes in the cluster



- FlashGrid Read-Local™ Technology minimizes network overhead by serving reads from local drives
- Two-way or three-way mirroring of data across separate nodes or Availability Zones
- Oracle ASM and Clusterware provide data protection and availability.

## Network

FlashGrid Cloud Area Network (CLAN) enables running high-speed clustered applications in public clouds or multi-datacenter environments with the efficiency and control of a Local Area Network.

The network connecting Amazon EC2 instances is effectively a single IP network with a fixed amount of network bandwidth allocated per instance for all types of network traffic (except for Amazon Elastic Block Storage (EBS) storage traffic on EBS-optimized instances). However, the Oracle RAC architecture requires separate networks for client connectivity and for the private cluster interconnect between the cluster nodes. There are two main reasons for that: 1) the cluster interconnect must have low latency and sufficient bandwidth to ensure adequate performance of the inter-node locking and Cache Fusion, and 2) the cluster interconnect is used for transmitting raw data, and for security reasons it must be accessible by the database nodes only. Also, Oracle RAC requires networks with multicast capability, which is not available in Amazon EC2.

FlashGrid CLAN addresses the limitations described above by creating a set of high-speed virtual LAN networks and ensuring QoS between them.

Network capabilities enabled by FlashGrid CLAN for Oracle RAC in Amazon EC2 offer the following attributes:

- Each type of traffic has its own virtual LAN with a separate virtual NIC, e.g., *fg-pub*, *fg-priv*, *fg-storage*
- Negligible performance overhead compared to the raw network
- Minimum guaranteed bandwidth allocation for each traffic type while accommodating traffic bursts
- Low latency of the cluster interconnect in the presence of large volumes of traffic of other types
- Transparent connectivity across Availability Zones
- Multicast support

## Shared Storage

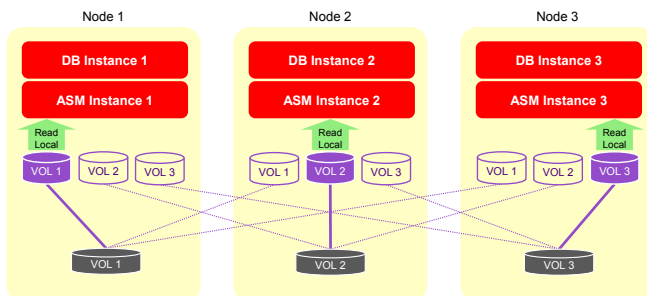


Figure 5. FlashGrid Storage Fabric with FlashGrid Read-Local Technology

FlashGrid Storage Fabric turns local drives into shared drives accessible from all nodes in the cluster. The local drives shared with FlashGrid Storage Fabric can be block devices of any type, including Amazon EBS volumes or LVM volumes. The sharing is done at the block level with concurrent access from all nodes.

Each database node has a full copy of user data stored on Amazon EBS volume(s) attached to that database node. The FlashGrid Read-Local Technology allows serving all read I/Os from the locally attached volumes and increases both read and write I/O performance. Read requests avoid the extra network hop, thus reducing the latency and the amount of the network traffic. As a result, more network bandwidth is available for the write I/O traffic.

The FlashGrid software maintains persistent disk names and sets the required disk permissions. There is no need to configure ASMLib or UDEV rules.

## ASM Disk Group Structure and Data Mirroring

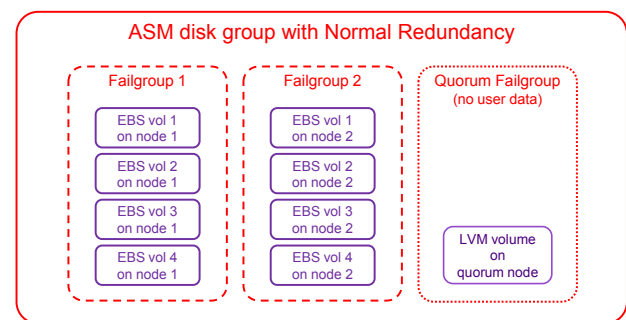


Figure 6. Example of a normal redundancy disk group in a two-node RAC cluster

FlashGrid software leverages proven Oracle ASM capabilities for disk group management, data mirroring, and high availability. In Normal Redundancy mode each block of data has two mirrored copies. In High Redundancy mode each block of data has three mirrored copies. Each ASM disk group is divided into failure groups—one failure group per node. Each disk is configured to be a part of a failure group that corresponds to the node where the disk is located. ASM stores mirrored copies of each block in different failure groups.

A typical Oracle RAC set up in Amazon EC2 will have three Oracle ASM disk groups: GRID, DATA, and FRA.

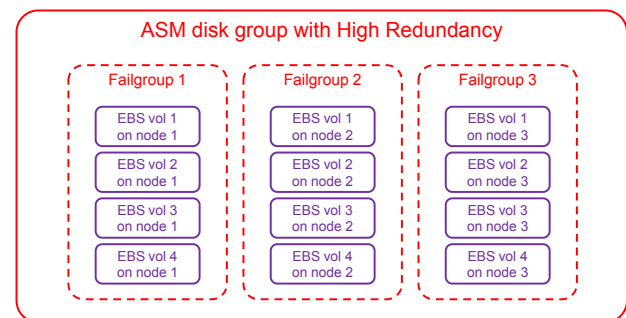


Figure 7. Example of a high-redundancy disk group in a three-node RAC cluster

In a two-node RAC cluster all disk groups must have normal redundancy. The GRID disk group containing voting files is required to have a quorum disk for storing a third copy of the voting files. Other disk groups also benefit from having quorum

disks as they store a third copy of ASM metadata and improve failure handling.

In a three-node cluster all disk groups, except the GRID disk group, must have high redundancy in order to enable full Read-Local capability. In a three-node RAC cluster the GRID disk group would typically have normal redundancy. Note that in such a three-node RAC cluster, loss of no more than one node is tolerated without causing downtime.

If a three-node RAC cluster must tolerate simultaneous failure of two nodes without causing downtime, the GRID disk group must have high redundancy, and an additional two quorum nodes must be provisioned to accommodate five copies of voting files. Details of such configuration are not covered in this article.

### High Availability Considerations

FlashGrid Storage Fabric and FlashGrid Cloud Area Network have a fully distributed architecture with no single point of failure. The architecture leverages HA capabilities built into Oracle Clusterware, ASM, and Database.

***Oracle RAC has the following infrastructure requirements that are not directly available in AWS: Shared high-performance storage accessible from all nodes in the cluster, multicast-enabled network between all nodes in the cluster, and separate networks for client, cluster interconnect, and storage. FlashGrid Storage Fabric and FlashGrid Cloud Area Network address these requirements.***

### Node Availability

Because all instances are virtual, failure of a physical host causes only a short outage for the affected node. The node instance will automatically restart on another physical host. This significantly reduces the risk of double failures.

A single Availability Zone configuration provides protection against loss of a database node. It is an efficient way to accommodate planned maintenance (e.g., patching database or OS) without causing database downtime. However, a potential failure of a resource shared by multiple instances in the same Availability Zone, such as network, power, or cooling, may cause database downtime. Additionally, when using instance types that do not occupy an entire physical host, there is a chance that two instances may be running on the same physical host. A failure of that physical host will cause database downtime. There is no such risk for large instance types occupying an entire physical host, e.g., m4.10xlarge, m4.16xlarge, or r4.16xlarge.

Placing instances in different Availability Zones virtually eliminates the risk of simultaneous node failures, except for the unlikely event of a disaster affecting multiple data center facilities

in a region. The trade-off is higher network latencies and, in certain cases, lower network bandwidth between the nodes.

### Data Availability

An Amazon EBS volume provides persistent storage that survives a failure of node instance where the volume is attached. After the failed instance restarts on a new physical node all its volumes are attached with no data loss.

Amazon EBS volumes have built-in redundancy that protects data from failures of the underlying physical media. The mirroring by ASM is done on top of the built-in protection of Amazon EBS. Together Amazon EBS plus ASM mirroring provide durable storage with two layers of data protection, which exceeds the typical level of data protection in on-premises deployments.

### Performance Considerations

An instance type must meet the following criteria:

- At least two vCPUs
- Enhanced networking—direct access to the physical network adapter
- EBS optimized—dedicated I/O path for Amazon EBS, not shared with the main network

The following instance type families satisfy the above criteria and are optimal for database workloads:

- M4: optimal memory-to-CPU ratio
- R4: high memory-to-CPU ratio, high peak network bandwidth
- X1: large memory size; large number of CPU cores

Oracle Database Standard Edition 2 customers can use two-node RAC clusters with up to four vCPUs per node, e.g., m4.xlarge or r4.xlarge instance types. More details about licensing Oracle software in the cloud computing environment are available at [www.oracle.com/us/corporate/pricing/cloud-licensing-070579.pdf](http://www.oracle.com/us/corporate/pricing/cloud-licensing-070579.pdf).

Quorum servers require fewer resources than RAC. However, the above criteria are still important to ensure stable cluster operation. c4.large or m4.large instances can be used as quorum servers. Using T2 family for quorum servers is not recommended. Note that there is no Oracle Database software installed on the quorum servers; hence, the quorum servers do not increase the number of licensed CPUs.

### Single vs. Multiple Availability Zones

Using multiple Availability Zones provides substantial availability advantages. However, it does affect network performance in the following ways:

- Network bandwidth between instances in different Availability Zones is limited to 5 Gb/s. Network speeds above 5 Gb/s are available only within a single Placement Group. A Placement Group is limited to a single Availability Zone. This constraint must be taken into account when using instances that normally have network bandwidth higher than 5 Gb/s.
- Latency between instances is increased. In the US-West-2 region for 8 KB transfers we measured 0.6 ms between



Availability Zones compared to 0.1 ms within a single Availability Zone.

The impact of inter-AZ configurations may be significant for the applications that have high ratios of data updates. However, read-heavy applications will experience little impact because all read traffic is served locally and does not use the network.

## EBS Volume Types

Use of Provisioned IOPS SSD (io1) volumes is recommended for data in order to achieve best performance. Each io1 volume can be configured with up to 20,000 IOPS.

General Purpose SSD (gp2) volumes are recommended for OS, GRID disk group, and quorum volumes. Normally, use of gp2 volumes for data is not recommended because of variability in their performance. However, for a volume of 1 TB or larger the performance will be 3000 IOPS, guaranteed. If this level of performance is sufficient, 1 TB gp2 volumes can be used. Note that using N x 1 TB volumes has N times performance advantage over using one N TB volume, while the cost is the same.

## Local SSDs

Use of local SSDs as the primary storage offers higher bandwidth and lower cost compared to Amazon EBS volumes. For example, the newly announced i3 instance type includes NVMe SSDs with up to 16 GB/s and up to 3.3 mln IOPS. The use of local SSDs is not covered in this revision of this document; it is planned for future revisions.

## Reference Performance Results

The main performance-related concern when moving database workloads to the cloud tends to be around storage and network I/O performance. There is a very small to zero overhead related to the CPU performance between bare metal and EC2. Therefore, in this article we focus on the storage I/O and RAC interconnect I/O.

## Calibrate\_IO

The CALIBRATE\_IO procedure provides an easy way to measure storage performance, including maximum bandwidth, random IOPS, and latency. The CALIBRATE\_IO procedure generates I/O through the database stack on actual database files. The test is read-only and it is safe to run it on any existing database. It is also a good tool for directly comparing the performance of two storage systems because the CALIBRATE\_IO results do not depend on any non-storage factors, such as memory size or the number of CPU cores.

Test configuration:

- Two database nodes, M4.16xlarge
- Four io1 20000 IOPS 400GB volumes per node

Test script:

```
SET SERVEROUTPUT ON;
DECLARE
  lat INTEGER;
  iops INTEGER;
  mbps INTEGER;
BEGIN DBMS_RESOURCE_MANAGER.CALIBRATE_IO (8, 10, iops,
  mbps, lat);
  DBMS_OUTPUT.PUT_LINE ('max_iops = ' || iops);
```

```
DBMS_OUTPUT.PUT_LINE ('latency = ' || lat);
DBMS_OUTPUT.PUT_LINE ('max_mbps = ' || mbps);
end;
/
```

Our results:

```
max_iops = 154864
latency = 0
max_mbps = 2219
```

Note that the Calibrate\_IO results do not depend on whether the database nodes are in the same or different Availability Zones.

## SLOB

SLOB is an open-source tool for generating I/O-intensive Oracle workloads. SLOB generates database SELECTs and UPDATEs with minimal computational overhead. It complements Calibrate\_IO by generating mixed (read+write) I/O load. AWR reports generated during the SLOB test runs provide various performance metrics. For the purposes of this article, we focus on the I/O performance numbers and the private interconnect latency.

Test configuration:

- Two database nodes, M4.16xlarge
- Four io1 20000 IOPS 400GB volumes per node
- SGA size: 2.6 GB (small size selected to minimize caching effects and maximize physical I/O)
- 8 KB database block size
- Schemas: 30 x 240 MB
- UPDATE\_PCT= 20

The table below shows our results for tests performed in the same configuration (provided above) in one Availability Zone and different Availability Zones.

	Same AZ	Different AZs
Read+Write Database Requests (both nodes combined)	121,237 IOPS	111,546 IOPS
Read Database Requests (both nodes combined)	100,539 IOPS	92,081 IOPS
Write Database Requests (both nodes combined)	20,697 IOPS	19,465 IOPS
Interconnect Ping Avg Latency 500 B msg	0.23 ms	0.69 ms
Interconnect Ping Latency Stddev 500 B msg	0.05 ms	0.06 ms
Interconnect Ping Avg Latency 8 K msg	0.27 ms	0.77 ms
Interconnect Ping Latency Stddev 8 K msg	0.07 ms	0.15 ms

While there is a roughly 3X increase on the interconnect latency between the Availability Zones (and similar increase on the write latency), the difference in throughput is less than 10%. With over 100 K IOPS in both cases, the performance results are comparable to having a dedicated all-Flash storage array.

## Compatibility

The following versions of software are supported by FlashGrid as part of the solution:

- ▶ Oracle Database: ver. 12.1.0.2 or 11.2.0.4 with the latest PSU
- ▶ Oracle Grid Infrastructure: ver. 12.1.0.2 with the latest PSU
- ▶ Operating System: Oracle Linux 7.3, Red Hat Enterprise Linux 7.3, CentOS 7.3
- ▶ FlashGrid software: ver. 16.11

The solution can be deployed on any EBS-optimized Amazon EC2 instance type with enhanced networking.

## Deployment Process

Below is a brief overview of the steps for deploying an Oracle RAC cluster on Amazon EC2.

- 1) Optionally customize an AMI provided by FlashGrid or create your own AMI
  - a. Download and install Oracle Grid Infrastructure software
  - b. Download and apply the latest Grid Infrastructure patch set update
- 2) Customize the cluster configuration file: network parameters, instance types, storage volumes, etc.
- 3) Run automatic cluster provisioning using the *flashgrid-aws-cluster-setup* tool
- 4) Run Grid Infrastructure configuration
- 5) Create DATA disk group(s) using the *flashgrid-create-dg* assistant tool



## Dr. DR



Dr. DR is brought to you by Axxana.

***Amazon Web Services consists of multiple independent Regions. Each Region is partitioned into several Availability Zones. Availability Zones consist of one or more discrete data centers, each with redundant power, networking and connectivity, housed in separate facilities.***

- 6) Download, install, and patch Oracle Database software
- 7) Create a database

## Conclusion

Running Oracle RAC clusters on Amazon EC2 out of the box has historically been challenging due to storage and network constraints. FlashGrid Cloud Area Network and FlashGrid Storage Fabric remove those constraints and enable a wide range of highly available database cluster solutions ranging from small cost-efficient and easy-to-deploy Oracle RAC Standard Edition 2 clusters to high-performance mission-critical Oracle RAC Enterprise Edition clusters with high availability characteristics exceeding those of the traditional on-premises deployments. ▲

*Artem Danielov is CTO at FlashGrid.*

Copyright © 2017 FlashGrid


## By Rich Parsons





# Oracle Licensing in the Cloud

by Mohammad Inamullah



Mohammad Inamullah

*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 before performing any actions based on the information provided.*

## Introduction

Cloud computing offers extraordinary opportunities for getting the best value out of your Oracle license investment. Understanding the licensing implications for the cloud is crucial to maintaining license compliance while optimizing your Oracle spend.

Whether you are planning a move to the cloud, considering a major Oracle purchase, performing a re-architecture, or concerned about Oracle license compliance and optimization, having a clear understanding of the licensing implications in cloud environments is crucial for your Oracle-in-the-cloud strategy.

## Why the Cloud?

*“You can’t stop change any more than you can stop the suns from setting.”—Shmi Skywalker—Star Wars: Episode I—The Phantom Menace*

The emergence and rapid growth of cloud computing has been one of the defining tectonic shifts in the IT world over the last decade. Cloud adoption by early adopters has already given way to the laggards. IT and engineering departments are being told by their CIOs to move to the cloud in rising numbers.

The promise and potential of the cloud are not lost on many. This includes companies that have significant investments in Oracle software licenses.

Customers switching to PaaS and SaaS services have less to worry about. Billing is typically simple, and compliance risks and issues are greatly reduced. In fact, customers wishing to switch to Oracle’s own PaaS and SaaS cloud offerings may be able to trade in their existing Processor and NUP licenses for partial Oracle Cloud credit. However, such a transaction likely won’t be free and will require good contract negotiation from the customer.

The area of most licensing complexity is moving existing Oracle on-premise licenses (Processors and Named Use Plus (NUP)) into public cloud IaaS environments. This is the specific aspect of Oracle licensing in the cloud that we will be dealing with in the remainder of this article. We will be addressing the

issue of taking existing on-premise Oracle licenses and applying them to IaaS server instances. We will answer questions like “How many Processor licenses do I need for an AWS x4-2xlarge instance with eight vCPUs?” and “How do I correctly count cores in Azure and Google Cloud, and should I worry about hyper-threading?” Consequently, we will not be touching on PaaS and SaaS alternatives, or any other cloud options that may good alternatives. Furthermore, we will be limiting the coverage to the following cloud vendors: **Amazon (AWS), Microsoft (Azure), Google (Google Cloud Platform), and Oracle (Oracle Cloud).**

The key concepts and fundamentals, though, can be applied to other cloud vendors as well. Let’s start with the basics of Oracle licensing.

## The Fundamentals

*“Obi-Wan has taught you well.”—Darth Vader—Star Wars: Episode VI—Return of the Jedi*

When it comes to Oracle’s database and middleware products, Oracle’s primary licensing metrics are the “Processor” and “Named User Plus (NUP)” metrics. The Processor metric is dependent on the number and type of physical processor cores, which is then scaled by an Oracle-defined factor. For most x86 chips with more than one core, that factor is 0.5.

For example, consider licensing Oracle Database Enterprise Edition using the Processor metric. As such, a quad-core Intel Xeon E5 chip will require: 4 (# of cores) x 0.5 (Oracle-defined core factor) = 2 Processor licenses of Oracle DB EE.

As its name suggests, the NUP metric is dependent on the number of users that directly—or through front-end applications—access the underlying Oracle database or middleware. That said, NUPs do maintain a connection to the physical cores. NUP licensing requires a minimum number of NUP licenses per processor. For example, for Database Enterprise Edition, that minimum is 25 NUPs per processor. For Internet Application Server Enterprise Edition, it’s 10 NUPs per processor.

While the Processor metric was suitable for licensing in bare-metal server environments, it can lead to some ambiguity and confusion in virtualized and cloud environments. Oracle has not introduced any changes to the definition in response to the emergence of cloud computing. Therefore, having a firm understanding of your Oracle contracts is crucial as you plan to move your on-premise licenses to a cloud environment of your choice.

Oracle has published guidance on licensing Oracle software in cloud environments in the form of the document titled

**“Licensing Oracle Software in the Cloud Computing Environment.”** In the absence of any contractual revisions or the introduction of new Oracle licensing metrics for cloud computing environments, this document is the best guidance we have from Oracle. It sheds useful light on Oracle’s thinking on counting processors. The key points in this document, along with a brief discussion for each, are as follows:

➤ **The document is for educational purposes only.**

First, it should be noted that Oracle has identified this document as being for “educational purposes only” and “may not be incorporated into any contract.” Customers often make the mistake of treating this document as binding or as having some contractual value. It does not. Oracle states that at the bottom of the document. Rather, we feel customers should treat this document as guidance on a new topic. **Remember: the only authority on Oracle licensing is the agreement you signed with Oracle.**

➤ **Approved vendors.**

Oracle has recognized Amazon Web Services and Microsoft Azure as “approved” cloud environments. This is interesting; misunderstanding and misinterpreting this is common and unfortunate. What this means is that the policy document, which is educational and not binding anyway, is specifically applicable to AWS and Azure. The document does not say anything about, or refer to, any other cloud vendors. There is no reason or basis to assume that other vendors like Google Cloud Platform are unapproved or otherwise entail any contractual or compliance issues. Oracle has not stated anything to this effect in this document.

➤ **Virtual cores have been correlated with physical cores.**

Now this is some useful insight. The document states:

For the purposes of licensing Oracle programs in an Authorized Cloud Environment, customers are required to count each virtual core as equivalent to a physical core. This policy applies to all Oracle programs available on a processor metric.

This statement about counting virtual cores raises some important questions: *How do you count a virtual core? What is a virtual core? How does it vary from vendor to vendor?* We will be diving deeper into this topic later in this article, so hold your thoughts on that. The document goes on to state that for purposes of counting sockets (as in DBSE), each set of four virtual cores will be counted as one socket.

➤ **ULA certifications cannot include virtual core counts from cloud environments.**

This statement seems straightforward—and not too inviting for moving your ULA deployments to the cloud! This implies that you can deploy your ULA products in AWS or Azure but cannot include them in your ULA certification, and thus you lose out on a lot of value. For example, suppose you have a ULA that includes Database Enterprise Edition. Further suppose that at the end of the ULA peri-

od, you end up deploying 50 Processors of DBEE on-premise and 25 virtual cores in AWS. According to this policy document, you may count the 50 Processors toward your final certification count but not the 25 virtual cores in AWS. Instead, for the 25 virtual cores in AWS, you must now purchase regular licenses from Oracle to cover that usage. *It’s at this point that I’d remind the reader that the document is a policy document that has no enforceability. This position has no contractual basis.*

So, now that we have some insight on counting cores in the cloud, let’s take a close look at the key concepts for correctly counting virtual cores and how they apply to the leading cloud vendors.

## AWS

*“Traveling through hyperspace ain’t like dusting crops, farm boy!”—Han Solo—Star Wars: Episode IV—A New Hope*

Amazon’s AWS EC2 instances are defined in terms of vCPUs. What exactly is a “vCPU”? How does it correlate with virtual and physical cores? Per Amazon:

*“Each vCPU is a hyperthread of an Intel Xeon core except for T2 and m3.medium.”* (<https://aws.amazon.com/ec2/instance-types/>)

This implies that a vCPU correlates to a thread and not to a full core. An Intel Xeon core with hyper-threading means that two threads are executing on a given core. So, counting vCPUs in AWS for purposes of licensing Oracle software is potentially misleading and incorrect. Almost as if to address this issue, Amazon has a (somewhat obscure) page titled “Virtual Cores by Amazon EC2 and RDS DB Instance Type” (<https://aws.amazon.com/ec2/virtualcores/>). The page states:

*Some independent software vendor (ISV) licensing is based on the number of virtual cores an instance provides. To assist you with virtual core licensing calculations for ISV software, the following tables shows the virtual cores provided by Amazon EC2 Instances and Amazon RDS DB Instances.*

For calculating Oracle Processor license requirements, customers should use this page and not rely on vCPUs, as they are incorrect and lead to overestimating the licensing requirement.

For example, suppose we want to calculate the Processor licensing requirement for Oracle Database Enterprise Edition on an Amazon EC2 **m4.2xlarge** instance. This has eight vCPUs—correlating to eight hyper-threads:

Select	Name	vCPU	Memory (GiB)	Instance Storage (GiB)	I/O	EBS Opt.	On-Demand Hourly Cost	Reserved Effective Hourly Cost (Savings %)
<input type="radio"/>	t1.micro	1	0.6	--	Very Low	--	\$0.020	\$0.008 (60%)
<input type="radio"/>	t2.nano	1	0.5	--	Low	--	\$0.006	\$0.003 (49%)
<input type="radio"/>	t2.micro	1	1.0	--	Low to Moderate	--	\$0.012	\$0.005 (58%)
<input type="radio"/>	t2.small	1	2.0	--	Low to Moderate	--	\$0.023	\$0.010 (57%)
<input type="radio"/>	t2.medium	2	4.0	--	Low to Moderate	--	\$0.047	\$0.021 (55%)
<input type="radio"/>	t2.large	2	8.0	--	Low to Moderate	--	\$0.094	\$0.041 (56%)
<input type="radio"/>	t2.xlarge	4	16.0	--	Low to Moderate	--	\$0.188	\$0.083 (56%)
<input type="radio"/>	t2.2xlarge	8	32.0	--	Low to Moderate	--	\$0.376	\$0.167 (56%)
<input type="radio"/>	m4.large	2	8.0	--	Moderate	Yes	\$0.108	\$0.040 (63%)
<input type="radio"/>	m4.xlarge	4	16.0	--	High	Yes	\$0.215	\$0.081 (62%)
<input checked="" type="radio"/>	m4.2xlarge	8	32.0	--	High	Yes	\$0.431	\$0.161 (63%)
<input type="radio"/>	m4.4xlarge	16	64.0	--	High	Yes	\$0.862	\$0.321 (63%)
<input type="radio"/>	m4.10xlarge	40	160.0	--	10 Gigabit	Yes	\$2.155	\$0.804 (63%)



The actual number of virtual cores for the **m4.2xlarge** instance is four, as seen below in the table of virtual cores for EC2 instance types:

EC2 Instance Type	Virtual Core Count
t1.micro	1
t2.micro	1
t2.small	1
t2.medium	2

Scrolling down, we find the following:

EC2 Instance Type	Virtual Core Count
m3.medium	1
m3.large	1
m3.xlarge	2
m3.2xlarge	4
m4.large	1
m4.xlarge	2
<b>m4.2xlarge</b>	<b>4</b>
m4.4xlarge	8
m4.10xlarge	20

Based on the virtual cores above, the licensing requirement for Oracle Database Enterprise Edition on a AWS **m4.2xlarge** EC2 instance would be: 4 virtual cores \* 0.5 core factor = 2 Processor licenses.

Had we relied on vCPUs, the incorrect result would have been: 8 vCPUs \* 0.5 core factor = 4 Processor licenses.

## Azure

While not explicitly documented, Azure does not use hyper-threading. Their documentation makes no references to hyper-threading and refers to its instances as having CPU cores from Xeon processors. Our own technical investigation confirms that Azure does not use hyper-threading. Let's look at Azure's documentation for their **Dv2-series** of instances:

Sizes for web and worker role instances

There are multiple standard sizes to choose from on Azure. Considerations for some of these sizes include:

- D-series VMs are designed to run applications that demand higher compute power and temporary disk performance. D-series VMs provide faster processors, a higher memory-to-core ratio, and a solid-state drive (SSD) for the temporary disk. For details, see the announcement on the Azure blog, [New D-Series Virtual Machine Sizes](#).
- Dv2-series, a follow-on to the original D-series, features a more powerful CPU. The Dv2-series CPU is about 35% faster than the D-series CPU. It is based on the latest generation 2.4 GHz Intel Xeon® E5-2673 v3 (Haswell) processor, and with the Intel Turbo Boost Technology 2.0, can go up to 3.1 GHz. The Dv2-series has the same memory and disk configurations as the D-series.
- G-series VMs offer the most memory and run on hosts that have Intel Xeon E5 V3 family processors.

Dv2-series

Size	CPU cores	Memory, GiB	Local SSD, GiB	Max data disks	Max data disk throughput, IOPS	Max NICs / Network bandwidth
Standard_D1_v2	1	3.5	50	2	2x500	1 / moderate
Standard_D2_v2	2	7	100	4	4x500	2 / high
Standard_D3_v2	4	14	200	8	8x500	4 / high
Standard_D4_v2	8	28	400	16	16x500	8 / high
Standard_D5_v2	16	56	800	32	32x500	8 / extremely high

Based on the above, with an Azure CPU core correlating to a Xeon core, the licensing requirement for Oracle Database Enterprise Edition on an Azure **Standard\_D3\_v2** instance with four CPU cores would be: 4 CPU cores \* 0.5 core factor = 2 Processor licenses.

## Google Cloud Platform (GCP)

GCP documentation is clear and transparent on this topic. Like the other vendors, GCP offers several classes of instances with varying allocations of virtual CPUs. GCP documentation footnotes clearly define a virtual CPU as a "single hardware hyper-thread" on a Xeon chip. Consider their "**n1-highmem**" class of instances and the footnote below the table:

Machine Types

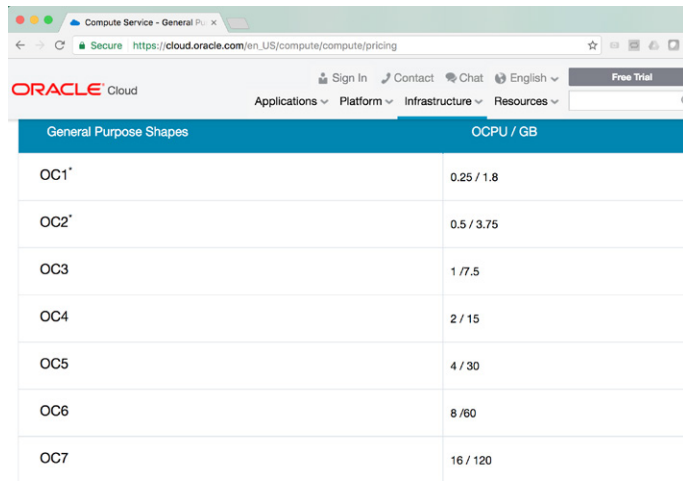
Machine Name	Description	Virtual CPUs <sup>1</sup>	Memory (GB)	Max number of persistent disks (PDs) <sup>2</sup>	Max total PD size (TB)
n1-highmem-2	High memory 2 CPU machine type with 2 virtual CPUs and 13 GB of memory.	2	13	16 (64 in Beta)	64
n1-highmem-4	High memory 4 CPU machine type with 4 virtual CPUs, and 26 GB of memory.	4	26	16 (64 in Beta)	64
n1-highmem-8	High memory 8 CPU machine type with 8 virtual CPUs and 52 GB of memory.	8	52	16 (128 in Beta)	64
n1-highmem-16	High memory 16 CPU machine type with 16 virtual CPUs and 104 GB of memory.	16	104	16 (128 in Beta)	64
n1-highmem-32 <sup>4</sup>	High memory 32 CPU machine type with 32 virtual CPUs and 208 GB of memory.	32	208	16 (128 in Beta)	64

<sup>1</sup>For the n1 series of machine types, a virtual CPU is implemented as a single hardware hyper-thread on a 2.6 GHz Intel Xeon E5 (Sandy Bridge), 2.5 GHz Intel Xeon E5 v2 (Ivy Bridge), 2.3 GHz Intel Xeon E5 v3 (Haswell), or 2.2 GHz Intel Xeon E5 v4 (Broadwell).

Therefore, the licensing requirement for Oracle Database Enterprise Edition on an **n1-highmem-8** instance with eight virtual CPUs would be: (8 Virtual CPUs / 2 hyper-threads per core) \* 0.5 core factor = 2 Processor licenses.

## Oracle Cloud

Oracle Cloud offers a variety of instance options with different OCPU allocations. A quick look at their General Purpose Compute shows clear and unambiguous documentation on what an OCPU is— equivalent to “one physical core of an Intel Xeon processor with hyper-threading enabled. Each OCPU corresponds to two hardware execution threads . . .”



The screenshot shows the Oracle Cloud website's pricing page for General Purpose Shapes. The table lists various shapes (OC1 through OC7) and their corresponding OCPU and GB specifications. The OC1 and OC2 shapes are noted as being supported in a future release.

General Purpose Shapes	OCPU / GB
OC1*	0.25 / 1.8
OC2*	0.5 / 3.75
OC3	1 / 7.5
OC4	2 / 15
OC5	4 / 30
OC6	8 / 60
OC7	16 / 120

\*OC1 and OC2 shapes will be supported in a future release.  
An OCPU provides CPU capacity equivalent of one physical core of an Intel Xeon processor with hyper threading enabled. Each OCPU corresponds to two hardware execution threads, known as vCPUs.  
Each 1 OCPU General Purpose shape can be provisioned with up to 15 GB of RAM.  
Monthly Computation based on 744 hours per month.

As an example of licensing Oracle Database Enterprise Edition, for an **OC5** instance with four OCPUs, we would need: 4 CPU cores \* 0.5 core factor = 2 Processor licenses.

While the above discussion has focused on AWS, Azure, GCP, and Oracle Cloud, the concept is the same for other IaaS vendors as well: Oracle licensing is done at physical cores and not virtual CPUs, which typically represent a single hardware execution thread; the underlying hardware (commonly Intel Xeon-based) may or may not have hyper-threading enabled.

## Named User Plus (NUP)

As previously mentioned, while NUPs are based on the number of users, they nonetheless have processor-related minimums. This calculation must always be done after the correct Oracle Processor requirements have been calculated for an instance type. For example, if we are considering Oracle Database Enterprise Edition with NUP licensing for a GCP **n1-highmem-8** instance, once we have calculated the Processor requirement of two in the example above, we can calculate the NUP minimum: 2 Processors \* (25 NUP/Processor minimum) = 50 NUPs.

Of course, this is the minimum. The licensing requirement will be the higher of this or the actual number of users.

## Risks and Opportunities

*“Great, kid. Don’t get cocky.”—Han Solo—Star Wars: Episode IV—A New Hope*

Now that we’ve talked about correctly sizing your cloud environment for Oracle licensing purposes, it’s worth spending some

time talking about the Oracle license compliance risks and opportunities inherent in moving to the cloud.

While the overall license compliance risk level does not seem to change when moving from on-premise to public cloud, the specifics do change. In other words, some risks diminish while new ones pop up.

Oracle licensing risk factors specific to public cloud environments include the following:

1. The ease of starting new instances makes it easy to rapidly deploy Oracle software without the need of provisioning hardware in the first place.
2. Related to the above point, it’s very easy to deploy Oracle software on very large instances. A junior engineer may fire up a 16-Processor instance with Database Enterprise Edition and incur a \$760k licensing requirement in only a few minutes.
3. Seasonal spikes in processing demand may launch new instances in automated fashion and trigger spikes in Oracle licensing needs. Careful control of the environment setup is key.
4. Related to the above point, cluster setups need to be carefully planned and set up to avoid license compliance issues.
5. Instances may be shut down but not fully terminated in a timely manner. Given that Oracle software requires licensing anywhere it’s “installed and/or running,” an instance that is shut down but the Oracle software is not uninstalled requires licensing. To fully free the corresponding Oracle licenses, instances must be fully terminated.

Moving your Oracle deployments to the cloud has many obvious benefits. Some areas of reduction in license compliance risk and optimization opportunities include the following:

1. Depending on how companies handle things, tracking their deployments in cloud environments may be easier due to the centralized nature of cloud management, thereby reducing the likelihood of unplanned or unauthorized Oracle deployments.
2. Oracle licenses may be more readily reused due to the ease of shutting down and terminating instances, effectively uninstalling Oracle software very quickly and re-purposing the licenses.
3. Using instance sizes that are specific to the need can help reduce the overall licensing need, rather than having to acquire specific hardware each time specific needs arise or deploying Oracle software on machines that are bigger than those needed. Some providers, like Google Cloud, allow customers to define their own custom instance sizes in terms of CPUs and memory.
4. Customers moving to cloud environments can use the move as an opportunity to assess their usage of Oracle software and reduce their footprint through consolidation—potentially reducing their Oracle support costs by trimming excess Oracle licenses.
5. Companies can readily standardize and monitor the deployment of Oracle software. Launching machine images with pre-installed Oracle software may reduce the risk of deploying the incorrect Oracle program editions and/or



versions, as well as options and features that may require separate licensing.

An exhaustive list of risks and opportunities of moving your Oracle installation to the cloud is beyond the scope of this article, and specific factors depending on each company's individual circumstances can add a lot of variability.

### Planning Your Move to the Cloud—Practical Steps

***“Do. Or do not. There is no try.”—Yoda—Star Wars: Episode V—The Empire Strikes Back***

Okay—now that you have your vCPUs and threads sorted out and your NUPs calculated, it's time to move, right?

Before you “lift and shift,” some licensing due diligence is highly advisable. Involving third-party Oracle licensing experts may be a good idea. Here are steps I would recommend before lifting and shifting:

1. Review your Oracle contracts and know exactly what product licenses you own and the respective quantities. Identify any limitations or restrictions. If you have Unlimited License Agreements (ULAs), make sure you know exactly which products are included in the ULAs and any upper caps.
2. Perform an internal audit of your Oracle deployment. Use the opportunity to establish a clear and firm baseline of your Oracle deployments. Make sure you have a clear enumeration of any licensing shortfalls and surpluses so you know exactly what compliance gaps need to be addressed, and strategize around how best to use your surpluses (be it optimized deployments, support cost reduction, and/or license termination).

3. Perform a value assessment of moving to the cloud. This involves calculating and comparing the cost of staying on-premise against moving to the cloud over the next several years. The value of moving to the cloud should be clearly demonstrated before moving (I know this sounds obvious, but it is worth repeating).
4. Perform a value assessment of PaaS and SaaS alternatives. When considering a move, it's worthwhile to evaluate alternative services like storage options and other PaaS options that may replace the need for traditionally installed software in IaaS instances. The value of these options should include the reduction in any administrative costs as well.
5. Calculate the cost of the move. Whether you have an in-house team or expert consultants coming in to facilitate the lift and shift, make sure you include this cost in your overall calculus of moving to the cloud.

It is crucial to note that the above steps are highly summarized, and that each step is a significant but important undertaking that must not be underestimated. The overall process will take time to plan and execute but will go a long way toward helping you move your Oracle deployments to the cloud with confidence and certainty. ▲

***“Remember...the Force will be with you, always.”—Obi Wan Kenobi—Star Wars: Episode IV—A New Hope***

*Mohammad Inamullah is the principal at Redwood Compliance in Palo Alto, California. He can be contacted at [mohammad@redwoodcompliance.com](mailto:mohammad@redwoodcompliance.com).*

© 2017 Mohammad Inamullah

***Editor's Note: As a testament to the fluid nature of the topic, just as we were about to go to print, Oracle updated its “Licensing Oracle Software in the Cloud Computing Environment” policy document (<http://www.oracle.com/us/corporate/pricing/cloud-licensing-070579.pdf>). The author provided the following comments:***

The key modification in the updated policy document is that the Processor Core Factor table is no longer applicable when calculating licensing requirements in Authorized Cloud Environments. The licensing requirement for AWS and Azure is explicitly stated as follows:

- **Amazon EC2 and RDS** - count two vCPUs as equivalent to one Oracle Processor license if hyper-threading is enabled, and one vCPU as equivalent to one Oracle Processor license if hyper-threading is not enabled.
- **Microsoft Azure** - count one Azure CPU Core as equivalent to one Oracle Processor license.

*When counting Oracle Processor license requirements in Authorized Cloud Environments, the Oracle Processor Core Factor Table is not applicable.*

Given that the typical core factor for multi-core x86 processors is 0.5, the net effect of this change is that it doubles the cost of Processor-based licensing in cloud environments. The policy also revises the approach to socket-based licenses for DB SE, SE1, and SE2, effectively raising the licensing requirement for these products too in cloud environments.

Keep in mind that this is a policy document which has no contractual value. Oracle states this very explicitly in the footer of the document: *“This document is for educational purposes only and provides guidelines regarding Oracle's policies in effect as of January 23rd, 2017. It may not be incorporated into any contract and does not constitute a contract or a commitment to any specific terms.”*

The only authority on correctly calculating licensing requirements is the contract signed between the customer and Oracle. Oracle's agreements (Oracle Master Agreement, OLSA, etc.) almost universally require Processor license calculations to be done using Core Factors. Oracle's policy document cannot negate any portion of the signed agreement. The methodology that I used in my article for calculating the Processor requirements in cloud environments is based on contractual terms between customers and Oracle. It is therefore unaffected by the updated policy document.


# 30th Anniversary Celebration *Sponsored by Google Cloud Platform*



*Happiness is a breakfast donut.*



*Database Specialists is now TriCore Solutions.*








*Post-conference celebration at Casino M8trix.*




*Included with your NoCOUG membership.*



*Also included with your NoCOUG membership.*






*Knock three times on the ceiling if you want me.*



*First-ever life-time member of NoCOUG.*



*Golden voices.*



# 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](http://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](mailto:support@orapub.com).



Find out more at [orapub.com](http://orapub.com).