

Spring 15 | Issue 56

**Optimising a Distributed Query**  
A strategy and hint from subject matter expert Jonathan Lewis

**Introduction to Oracle**  
Martin Widlake provides an overview of the Oracle RDBMS architecture

**Getting it Right First Time**  
Pythian's Michael Abbey highlights the importance of planning

# OracleScene

► Serving the Oracle Community

## A User Interface Philosophy

Oracle Alta UI



**UK OUG**  
UK ORACLE USER GROUP

[www.ukoug.org](http://www.ukoug.org)

An independent publication not affiliated with Oracle Corporation

This edition's sponsor:

**HGST**  
a Western Digital company

# Inside this issue

## Oracle Scene Editorial Team

**Editor:** Brendan Tierney  
**Email:** editor@ukoug.org.uk

**UKOUG Contact:** Brigit Wells  
**Email:** brigit@ukoug.org

**Sales:** Kerry Stuart  
**Email:** kerry@ukoug.org

## UKOUG Governance

A full listing of Board members, along with details of how the user group is governed, can be found at:

[www.ukoug.org/about-us/governance](http://www.ukoug.org/about-us/governance)

## UKOUG Office

UK Oracle User Group, User Group House,  
591-593 Kingston Road, Wimbledon  
London, SW20 8SA

**Tel:** +44 (0)20 8545 9670

**Email:** info@ukoug.org

**Web:** [www.ukoug.org](http://www.ukoug.org)

## Produced and Designed by

### Why Creative

**Tel:** +44 (0)7900 246400

**Web:** [www.whycreative.co.uk](http://www.whycreative.co.uk)

## UKOUG Event Photography

Joy Marshall  
Andrew Gourley  
[www.blueskiphotography.com](http://www.blueskiphotography.com)

## Next Oracle Scene Issues

**Issue 57:** Publish month: July 2015

Content deadline: 8th May 2015

**Issue 58:** Publish month: December 2015

Content deadline: 11th September 2015



More than 15,000 people  
follow UKOUG.

Join them now. @UKOUG



## OracleScene® UK Oracle User Group Ltd

The views stated in Oracle Scene are the views of the author and not those of the UK Oracle User Group Ltd. We do not make any warranty for the accuracy of any published information and the UK Oracle User Group will assume no responsibility or liability regarding the use of such information. All articles are published on the understanding that copyright remains with the individual authors. The UK Oracle User Group reserves the right, however, to reproduce an article, in whole or in part, in any other user group publication. The reproduction of this publication by any third party, in whole or in part, is strictly prohibited without the express written consent of the UK Oracle User Group. Oracle is a registered trademark of Oracle Corporation and/or its affiliates, used under license. This publication is an independent publication, not affiliated or otherwise associated with Oracle Corporation. The opinions, statements, positions and views stated herein are those of the author(s) or publisher and are not intended to be the opinions, statements, positions, or views of Oracle Corporation.



## TECHNOLOGY

- Busting Some Myths Around Platform-as-a-Service by Dermot O’Kelly 26
- Distributed Hinting by Jonathan Lewis 28
- Do It Right; Do It Once by Michael Abbey 38
- The Changing World of High Performance Enterprise Database Solutions by Martin Power 40

## EVENTS

- Your Invitation to Submit: Call for Papers is Open! 18
- We’re So Much More Than a Solaris SIG 06

## FOCUS ON

- Next Gen Update by Debra Lilley 07
- Meet a Volunteer: Andrew Watkins 17
- Why You Should Join the Oracle ACE Program by Jennifer Nicholson 36

## REGULAR FEATURES

- News and Reviews 04
- NEW: Partner Showcase: eBiz Answers 37

OracleScene  
DIGITAL

View additional content and get more from your  
Oracle Scene at: [www.ukoug.org/os](http://www.ukoug.org/os)

- Cloud Applications - Introducing a New Kind of User by Debra Lilley 48
- Fusion Middleware Provisioning with Cloud Control by Andreas Chatziantoniou 49
- Creating the Platform of the Future with Fusion Middleware 12c by Antonis Antoniou & Ronald van Luttikhuisen 53
- Oracle APEX in the CERN Java Cloud by Luis Rodriguez Fernandez 60
- Implementing Access Manager by Peter de Vaal & Nicolay Moot 63
- Oracle SQL Developer... for the DBA? by Jeff Smith 67



# First word

**Spring is all about freshness...** and nice weather (hopefully)! That's what we aim to do with every edition – keep it fresh for our readers and I think we've managed it again this time...

**In this issue, Grant Ronald provides some interesting insight into the principles on which the Alta UI is built and how this can be used to improve your existing applications. I'm delighted that Martin Widlake gave in to my arm twisting at Tech14 and has written the first of a series of introductory articles based on his Oracle in 90 minutes conference session.**

A concept we should all know about is the importance of planning, but how many of us actually take the time to do it in practise? In his article, Michael Abbey provides a guide to follow to help us all aim to get it right first time. For the more experienced among you we have an article from Jonathan Lewis on optimising a distributed query. There's some useful information to help you understand licensing Oracle on VMware from dbi Insight's CFO, Gregory Steulet and another good 'how to' article from Greg's colleague Franck Pachot, this time covering workload statistics gathering... and a lot more great content in the extended digital edition too.

We have news and updates from UKOUG galore, special mentions to Next Gen, the mass call for papers for Apps15, Tech15 & JDE15 – get your abstracts in nice and early, I will be – and this time we have a feature on the Solaris SIG... So lots of events for you to get involved with... and maybe even begin your journey towards becoming an Oracle ACE, find out more on p36.

If you've not been to a UKOUG event before or if you haven't attended one in a while you really should, especially if you're a member... make the most of your inclusive SIG places to brush

up on your Oracle knowledge.

Through active participation in community events you get to know a great number of Oracle people, partners and customers. Some of the people you meet even become good friends and there are opportunities galore to develop and improve yourself and your organisation. What's more, Oracle supports the SIGs by having staff involved with each one, which means there's added potential to pick brains in the know for free!

As one member put it recently, UKOUG events are "gold mines of useful information... the sessions cover a range of topics and the presentations from customers, partners as well as Oracle are typically very conversational, so you pick up insight into a lot of practical aspects not just theory like you'd commonly get in say a training session".

Whatever your speciality there is something on offer, check out the calendar and I hope to see you at an event soon:

[www.ukoug.org/events](http://www.ukoug.org/events)

Our call for articles for the summer edition closes on 8th May, send your articles to [articles@ukoug.org.uk](mailto:articles@ukoug.org.uk)

More information of submitting an article can be found online at: [www.ukoug.org/oraclescene](http://www.ukoug.org/oraclescene)



## ABOUT THE EDITOR

**Brendan Tierney**   
Consultant, Oralytics.com

Brendan is an Oracle ACE Director, independent consultant and lectures on Data Mining and Advanced Databases in DIT in Ireland. Brendan has extensive experience working in the areas of Analytics, Data Mining, Data Warehousing, Data Architecture and Database Design for over 20 years. He has worked on projects in Ireland, UK, Belgium and USA. He started working with the Oracle 5 Database, Forms 2.3 and ReportWriter 1.1, and has worked with all versions since then. Brendan is the editor of the UKOUG Oracle Scene magazine and is the deputy chair of the OUG Ireland BI SIG. Brendan is a regular presenter at conferences in Ireland, UK, Norway, Brazil and USA, including Oracle Open World.

Contact Brendan at: [editor@ukoug.org.uk](mailto:editor@ukoug.org.uk)

# 2015 Strategy Update

Last October we published our 2015 Strategy and promised to review a number of areas in January. Fiona Martin, UKOUG Members Advocate Chair, shares the updates with you all.

**Conference Strategy:** In January we completed our post event reviews from our December Conferences. As always we had lots of feedback, mostly very positive, as well as many ideas and suggestions for improvement. The feedback has been shared with the project teams planning this year's events and I know many of the ideas are already being considered. For the last 2 years our strategy has been to bring our Applications communities together. They share a common interest to understand their options as they plan their applications journey, whilst balancing the need to continue to understand and share experience around their current Oracle products. Having listened to all the community feedback from Apps14 we've recognised the need to strengthen

the Apps15 committee and improve links back to our SIG committees. In addition, we have slightly modified our 2015 plans. The Hyperion community has a particular challenge so we will expand upon the already proposed June "super" SIG and run a Hyperion one day event. We do still believe that our conference strategy is correct yet we understand why this does not work for Hyperion. Apps15 in December will still contain Hyperion content. This is a strategic product set that Oracle is speaking to all customers about. We will focus on introductory level and product information and will welcome all Hyperion members to Apps15.

**Volunteers:** Following the feedback from our December Volunteers drinks and the

continuing work by Marisa we will keep with the 2015 plan for Volunteers Drinks in June (Agenda Planning Day/AGM) and December (conferences).

**Membership:** Following positive results from the membership team we will continue to focus on tactical membership initiatives around our up and coming SIGs and events.

I hope this article shows we continue to listen to member feedback and review our products and services strategy based upon what you tell us.

If you have any comments or questions, please do email me:  
[fiona.martin@ukoug.org.uk](mailto:fiona.martin@ukoug.org.uk)



## Interim UKOUG CEO Appointment

We're pleased to announce that Pauline Drummond has been appointed Interim CEO of UKOUG while the Board assess what is required from the leadership position going forward.

## Summer edition Call for Articles

We're inviting functional and technical content submissions for the next edition.

All content submitted will be reviewed and considered for publication. The submissions deadline is **8th May**. Please send your submissions to [articles@ukoug.org.uk](mailto:articles@ukoug.org.uk), with a short bio, your company name, job title and a high resolution photo.

Users Convene in  
Dublin for

## OUG Ireland 2015

Over 200 delegates came together to learn and share experiences at sessions on Applications, Cloud, BI - Data Integration, BI - Business Analytics, Development, Technology and Database.

Opening keynote Jon Paul, Regional Director for Oracle, discussed the developments in Ireland and strategic updates from Oracle before delegates delved deeper into the applications and technology in sessions with Oracle ACEs and end users.

Throughout the day, delegates had the chance to network and visit the lively exhibition of Oracle partners to compare products and services and take part in hosted prize draws and activities. The day concluded with two fantastic keynotes;

Nadia Bendjedou from Oracle discussed co-existing Oracle E-Business Suite investment with Oracle Cloud solutions and Oracle's Maria Colgan explained Oracle's motivation for developing the Oracle Database In-Memory technology.

Thank you to everyone who attended OUG Ireland and we look forward to seeing you at our autumn Special Interest Group Meetings:

24th September 2015, Jurys Inn Dublin Custom House

- OUG Ireland BA & Big Data
- OUG Ireland HCM
- OUG Ireland Technology

Find out more at [www.ukoug.org/events](http://www.ukoug.org/events)



Article submissions deadline:  
**8th May 2015**  
Publish month:  
**July 2015**



UKOUG Partner of the Year Awards 2015

# Call for Nominations

Back in 2007, we introduced the UKOUG Partner of the Year Awards as a way of recognising the value that the partner network contributes to the Oracle user community in the UK. But awards are commonplace; everyone is doing them, so why should you care about ours?

Because they really count! Our awards are unique because they're independent of any third party. This makes them valuable and sought after – they're truly credible and they're won on merit. No one other than the customer can influence who wins. What's more they're the only Oracle-related awards globally that are selected on customer votes alone, so they're a great way of empowering end users.

Last year we saw record numbers of customers voting for their preferred partners - we plan to beat that record this year. But first we need our partner shortlist, so get your nominations in today!

Nominations are open until 18th May. Find out more at [www.ukoug.org/pya](http://www.ukoug.org/pya)

15 OCTOBER 2015  
THE WALDORF HILTON | LONDON

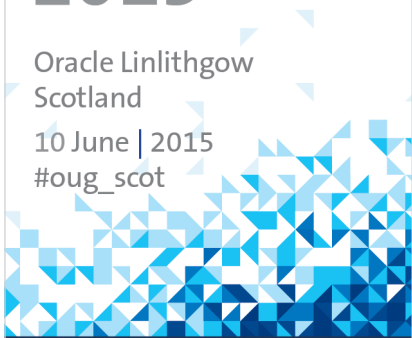
partneroftheyear  
*Awards*  
2015/16

**OUG SCOTLAND**  
ORACLE USER GROUP

**OUG Scotland**  
**2015**

Oracle Linlithgow  
Scotland  
10 June | 2015  
#oug\_scot

Register at  
[www.oug.org/scotland](http://www.oug.org/scotland)



## APPLICATIONS INNOVATION INITIATIVE

# Understanding Your Applications Journey

At UKOUG we understand that applications customers want to know about their options for the future and this goes beyond just Fusion Applications; there are new acquisitions, enhancements in technology and Cloud adoption to also be considered.

To help you get to grips with all the options available to you we introduced our Applications Innovation initiative. This initiative aims to help you access all the information you need to make confident decisions about the next stage of your applications journey.

The next event in the calendar focusing on Applications Innovation is the HCM Community Applications Innovation Day on Thursday 4th June.

Find out more and register at:  
[www.ukoug.org/events](http://www.ukoug.org/events)



## Goodbye & Good Luck

James Haslam left his role of UKOUG CEO on 6th March 2015. We thank James for an outstanding 19 years service with the user group. He is leaving UKOUG in a strong financial position for the experienced team at UKOUG to deliver the services to satisfy the members needs in 2015. We wish James all the very best for the future.



We're so much more than a



# Solaris



OK, you're a DBA or an architect, analyst or maybe you're an IT manager, so why should you be interested in Solaris SIG events?

Peel away the upper layers and you'll very quickly see that the Solaris Special Interest Group (SIG) is so much more than a bunch of Solaris System Administrators and enthusiasts. Our membership has always been diverse and there's a good reason for this. We don't limit ourselves solely to Solaris-focused activities and talks at our meetings; we aim to cater for the wider Oracle community.

Over the years, we've covered subjects as diverse as Getting the most out of MOS; Enterprise Cloud with OpenStack on Solaris 11.2; An Introduction to SuperCluster; Oracle Enterprise Manager Ops Center 12c; and Solaris 11 & Network Virtualisation at our evening meetings – yes, we meet in the evenings, which may make it easier for you to come along after work rather than cutting into your working day.

Not only that, but our evening events are also free to attend. If you're already a UKOUG member, you don't need to worry about using up any of your SIG places, just register and come along.

## Introducing the New Solaris SIG Committee

We've kicked off 2015 with a new committee following elections last year. In addition to committee veteran Andrew Watkins, who stays with us for his second term, we're pleased to welcome Bob Hoekstra and Phil Harman to the team. Our new committee members bring not only many years of experience working with and for Oracle customers but also plenty of fresh ideas to help revitalise the SIG and broaden our appeal to other UKOUG members even further.

## New Look, New Venue

Expect new and exciting things at the Solaris SIG in 2015! Following on from the re-launch of the Solaris SIG in February at its new home in Bonhill Street, London, we're planning a great selection of talks, discussions and networking opportunities throughout the rest of the year including an in depth session on Oracle Software in Silicon in April.

Along with the content, the structure of the SIG also varies, so you may expect to see anything from snappy chalk & talk style sessions to technical deep-dives, lively roundtable or panel discussions, show & tell quick demo's and the occasional hands-on offering or a mixture of all of these.

## Testimonials...

But don't take our word for it, let some of our long time Solaris SIG delegates tell you what's in it for you in their own words.

"The Solaris SIG is a great place to hear about Solaris-related technologies, both Open Source and commercial based. But even more importantly it is an opportunity to hear from others on how and why they use different technologies in their environments. There are often Solaris engineers giving talks or attending the meetings, so it's a good opportunity to talk to them directly. It is a great place to share your experiences, ideas and to learn something." Robert Milkowski, VP in Unix Engineering, Morgan Stanley

"The topics are wide ranging, the presenters are knowledgeable and approachable. It is a very practical way to keep in touch with all that is happening around Solaris."

José Martinez, Independent

Visit [www.ukoug.org/our-communities/oracle-technology/solaris](http://www.ukoug.org/our-communities/oracle-technology/solaris) for more information or find us on Facebook, Twitter or LinkedIn. Solaris-related content will also be on the agenda at the upcoming UKOUG Systems Event on 20th May in London. ■

**Committee:** Andrew Watkins (Birkbeck College), Bob Hoekstra (Hoekstra Systems), Phil Harman (Harman Holistix)  
**Oracle liaison:** Joy Marshall **Project manager:** Jake Potter

# Next Gen Update

Debra Lilley, UKOUG Member Advocate 

UKOUG embarked on a new initiative in 2014 around the Next Generation of Oracle professionals. The statistics we see everyday show that our industry needs people and there are not enough available.

IT career initiatives are all around us, just this week at my local airport there were big billboards advertising to youngsters picking their electives at senior school. The Bring IT On initiative is a careers attractiveness programme aimed at encouraging more students to consider a rewarding career in IT. The programme was designed by e-skills UK, a not-for-profit, employer-led organisation licensed by government as the Sector Skills Council for Business and Information Technology.

Our UKOUG initiative was simple enough, we would encourage young people into our industry, show them what to expect with case studies and talks and bring our community together as one prospective employer, so that CVs could be shared with all our member organisations that wanted either to bring in apprentices or graduates. We shared the idea with user groups in EMEA and there was a lot of interest, Oracle loved the idea and gave us a small grant to help with the getting the idea up and running.

We started off with a great event in Birmingham last year and those who took part all said it was a great start, but it was also very difficult to be a new entrant in an area actually swimming in initiatives, we were almost lost in the crowd, and the cost of setting up something new is high. We stepped back a little and looked at how we can work with what is there and mobilise our community to be a positive force in other initiatives, thus being more productive.

The thoughts behind the original idea was that students would know about Oracle, but what about the wider Oracle community who also use those skills, the customers and partners? So we are looking at how we can plug into what Oracle do in this area and expand their reach for our mutual benefit.



**MAKING SUBJECT CHOICES?**

**The demand is on for IT professionals right now!**

Find out more from your **careers adviser** or **nidirect.gov.uk/skills**

**SKILLS TO SUCCEED** | **SKILLS IN DEMAND BRING IT ON**

We have a great relationship with Oracle Academy that helps bring Oracle to education to ensure students are taught the right skills. They have a very mature program for universities, which was introduced to the UK several years ago, and UKOUG helped recruit educators in their outreach.

Oracle Academy talked about their program at our kick off event in Birmingham. We also had two UKOUG speakers Viswa Swaminathan who had been through a graduate program talked about his experiences and Tim Hall talked about the importance of community to aid learning, and one suggestion from Oracle Academy is that we source user group speakers like these to join them on their 'Job Ready' lecture tour around universities.

Oracle take on a lot of graduates and get involved in the annual 'Milk Round' and various job fairs and we are looking at how we can be part of that, a great way to bring our community to students. Like all our initiatives, we have a committee of volunteers who come together, share ideas and shape take them to a project we can be successful with. Most exciting in our industry at the moment is the

resurgence of apprenticeships and Oracle Academy last year did a pilot with RBS, providing the training content at the local college they attended.

The Future Jobs Framework will provide a platform for partnerships between colleges and industry to create a path to education for industry entrants and work experience for college entrants. Oracle held an open day in February to showcase what they had achieved to the local community, and Pauline Drummond, our then President Elect, represented UKOUG. This is to be rolled out further around the country and another idea to consider is would UKOUG member organisations like to give work experience to younger students and for our members to mentor individuals.

All these ideas take volunteers but that is what we do well and the payback of helping someone start their career is amazing. ■

**If you would like to be involved, keep an eye on our Next Gen webpage [www.ukoug.org/nextgen](http://www.ukoug.org/nextgen) or ask about being part of the committee by emailing [info@ukoug.org](mailto:info@ukoug.org)**





## Oracle Alta UI: A User Interface Philosophy

Last year I wrote a UKOUG Oracle Scene article “Mobile: The Enterprise Game Changer” in which I touched on the impact that mobile is having on developers.

Grant Ronald  
Director of Product  
Management  
Oracle

**But the change isn't only about technology: the bits and the bytes, the tools or the infrastructure. It's also about how we work and interact with individuals, businesses and IT systems.**

In this article I'll look at how Oracle has introduced a new user interface philosophy called Oracle Alta UI, which is driving a mobile-first user experience (UX) for both web-based, and on-device mobile applications.

### Why Mobile is Different

So why is mobile different to our traditional “sitting-in-the-office” applications? Surely approving a purchase order is much-of-a-muchness regardless if you are performing it on a PC or a tablet? Well, yes and no: yes it could be, but no it shouldn't.

Yes, you could implement the same user experience for mobile as you do for your PC based users, but in doing so you are developing to the lowest common denominator and not really exploiting the user gestures that have pretty much

become “the norm” for many of your users.

Just take a minute to think about some of the differences between working on a PC versus a tablet. On a PC you might tab through data but on a tablet you swipe through data. On a PC a mouse gives you a precise way of clicking on content; whilst using a finger on a tablet often requires a larger target area. And of course the nature of working with a smaller screen whilst on the move and subject to interruptions means those mobile tasks are often subjected to the “2 minute rule”.

The whole anthropological change of “how we work” with mobile devices has been the subject of extensive research by User Experience (UX) experts, including Oracle's own dedicated UX team. As a result of their research, UX specialists adopted a phrase which summarises one of the principal interaction patterns on mobile devices: “Glance, Scan, Commit”.

The concept of “Glance, Scan, Commit” applies whether you are using Facebook, LinkedIn or your company's HR application and is even broadly analogous to our

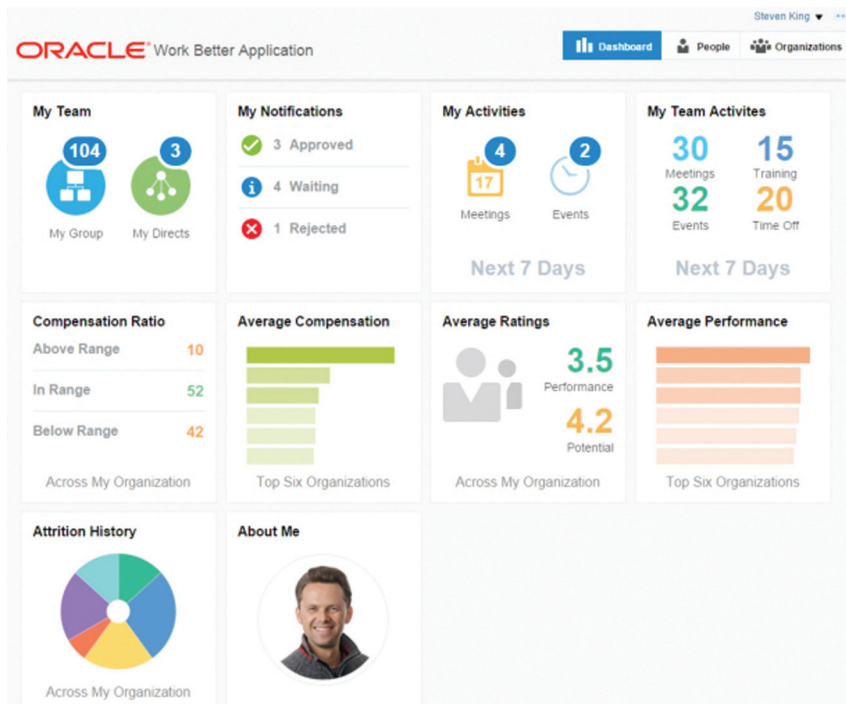


FIGURE 1

everyday lives away from IT, such as out shopping, where we exhibit similar and repeatable habits of interaction. We first glance to ascertain if there is something to action. Did I get a Facebook alert; does the store have a sale on? Then you scan: you scan through your Facebook feed to see if anything interesting is happening; just like scanning the shelves in the store.

**All the time you are making judgment calls, discarding options, ranking possibilities. Then you commit. You post a reply, accept a friend invite, or you buy a new watch at the cash register.**

### Glance

And it is because you are following this pattern of “Glance, Scan, Commit” that the user interface has to evolve to support this mode of operation. You glance at your tablet and immediately need to ascertain if there is something of interest or something you need to action. This means that you need to develop an interface with only essential information and content collated into a top-level launch pad or springboard. Or you might use badging in

order to indicate a particular application feature has pending information.

### Scan

Having ascertained that there is something interesting, the user will now want to drill in and scan through relevant information. Because the brain can process images more quickly than text, you may choose to present information graphically using charts, gauges or images. You could use colour to convey information and even differing font sizes might be used to convey the hierarchy, or importance, of information.

### Commit

And finally you make a decision to action a unit of work: editing information or completing a transaction. And in doing so, the goal is to help the user be as efficient as possible. For example, by presenting “Years in role” rather than “Hire date” the user can more immediately ascertain how long an employee has worked at the company rather than forcing them to perform mental arithmetic to work out if the employees tenure should weight on the decision to give a promotion. And it is the need to embrace these changes in relationship between users and applications that drove Oracle to develop a new design philosophy for user interaction called Oracle Alta UI.

### Introduction to Oracle Alta UI

Oracle Alta UI is much more than a look and feel. Simply putting an Oracle Alta UI skin on an existing application doesn’t “Alta-fy” it. Alta UI also prescribes a design approach as well as a set of guidelines for developing mobile-first user interfaces. Oracle published these design guidelines<sup>i</sup> in 2014<sup>ii</sup> and they cover developing UIs for both web-mobile and on-device mobile applications.

**Of course, a design philosophy sounds all very hippy and Zen, but how does this result in a better user experience for the applications you are developing today?**

The fact is the UX team who developed the Alta UI approach did so in collaboration with the Oracle Development Tools group. That means that the design principles and look & feel that Alta UI defines are baked into the technologies behind Oracle’s cloud and mobile offerings, including Oracle ADF Faces and Oracle Mobile Application Framework.

Let’s first take a look at what an application developed with the Oracle Alta UI looks like.

### Dashboard Style

Figure 1 is a page taken from the publicly available sample “Workbetter” application<sup>iii</sup>. This is a sample application developed by Oracle which is built to demonstrate a modern mobile-first application which is based on the Oracle Alta UI. You can take it for a spin by trying out the hosted version online<sup>iv</sup>.

On viewing this screen you may notice how Oracle Alta UI introduces some key differences when compared with your traditional applications. Firstly, the UI is what we might refer to as very “clean”. The components are large and flat with minimal shading. Space between components is an important factor, and there tends to be very little extraneous visual “chrome”. Of course, there is a reason behind this rather than simply the fashion of minimalistic UIs (also remembering that generally speaking the simpler the UI the less information that needs to be transmitted to render

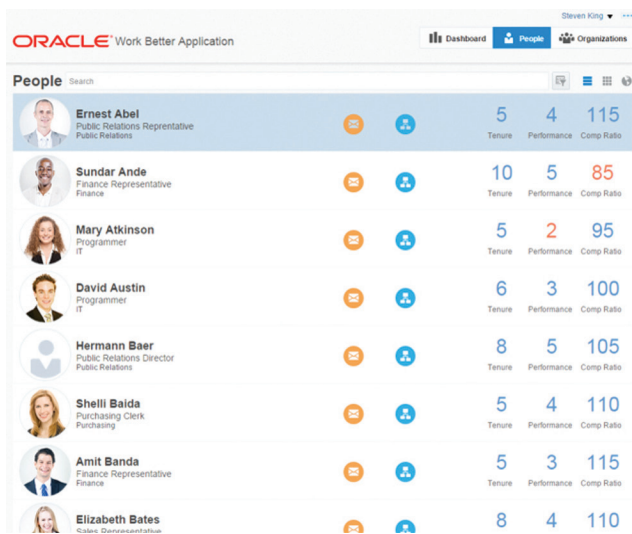


FIGURE 2

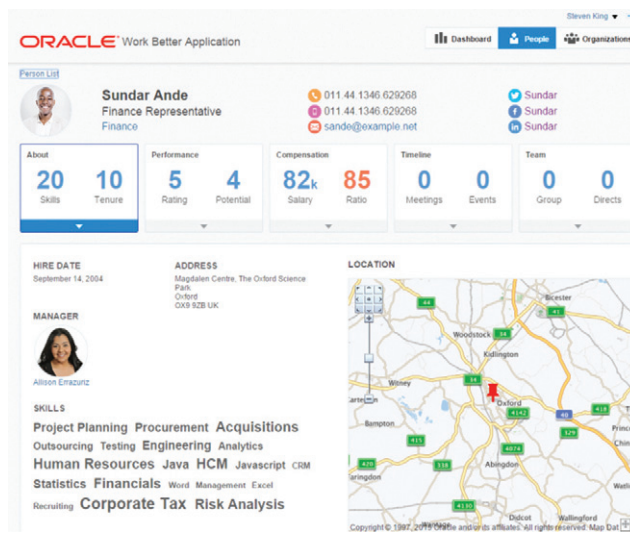


FIGURE 3

it – very important for a mobile device). Because components are larger with clear boundaries they become easier to tap with a finger. And because these components themselves are actionable, you don't generally need extra buttons or menu options.

**Your information also becomes your means of navigation.**

You'll also notice that this example addresses our UX requirement of "Glance". Clear information portlets (sometimes called infolets) and badging give the user an immediate overview of the current application state.

**List Page Style**

The next example, figure 2, is a list of users where you might be scanning (remember the UX edict of "Scan") to decide if any of your employees are currently operating with a low performance rating. What Oracle Alta UI has defined here is that the page is designed "mobile first". Or more accurately, tablet first. The design fits into a width of 1024 pixels, which, lo and behold, is the width of an iPad tablet in landscape mode. Oracle Alta UI defines that all content should be constrained within that width and that content be allowed to flow vertically below the fold of the browser view. And this fits in with the expectations of users. We are used to swiping down to see more information on tablets. However, as soon as we scroll in

two dimensions we start to lose context and work less efficiently.

Each row itself is also the means of navigation; so clicking on a row will change to view the details for that employee. This means selection and navigation becomes a single action. And by having larger rows they are more easily selected using a finger on a tablet.

Figure 2 also demonstrates two other Oracle Alta UI guidelines. Firstly, whilst viewing this page any visual clues the UI can provide will aid speedier scanning. You will notice that colour is being used to convey alert information. Your eye is immediately drawn to the Comp Ratio data in red/orange, which is good, because those are the employees you probably need to action first. Secondly, when designing Oracle Alta UI applications you should be in the mindset of "What information does the user need to complete this task as efficiently as possible". When deciding on an employee's compensation or rating one of the influencing factors is how long the employee has been in the job. In which case, give the user the information they need to do the job (tenure), not just the data which is in the database (hire date).

**Detail Page Style**

Finally, figure 3 is typical of an Oracle Alta UI page where you are committing to undertake some application function. In this example you can see the use of information hierarchy. Not all information is of equal importance and so it can be presented differently. The employee's name is important and is a larger font.

Compare that to the fact he works in the Finance department.

There is also the aim to disclose information only when relevant in order to make best use of space and to keep the user focused on the important facts. Looking at figure 3 the user can immediately see the employee's salary and compensation ratio which are the key facts at this point. Other details such as bonus, share options or benefits can be revealed by selecting the compensation box to disclose further information in the main screen area.

**Implementing Oracle Alta UI**

Having "drunk the Alta Kool-aid" (to borrow a very American metaphor) the next question is how might you implement an Oracle Alta UI in your application today? The Oracle Alta UI guidelines give you some direction, but the real aid to implementation is in the tooling. Currently Oracle ADF Faces and Oracle Mobile Application Framework provide skins and tooling to help you develop an Oracle Alta UI (although you could implement Oracle Alta UI using pretty much any tool or technology).

**So let's see what is possible to develop in, say, 20 minutes.**

**Looking at a Simple Use Case**

Those of us of a certain age, development background and lack of artistic talent will have probably spent a considerable



ID	LastName	FirstName	Userid	Email	StartDate	ManagerId	Title	DeptId	Salary
2	Iggas	Ladris	luggas@summi.com	luggas@summi.com	3/8/2010	1	5	41	1450
6	Ungpart	Milly	mungpar	mungpar@summi	1/18/2011	2	8	41	1200
12	Gilgan	Henry	hgilgan@summi.c	hgilgan@summi.c	2/18/2011	3	2	41	1450
16	Mauro	Elena	emauro@summi.c	emauro@summi.c	2/18/2011	6	3	41	1400
17	Smith	George	gsmith@summi.c	gsmith@summi.c	3/8/2011	6	3	41	940

FIGURE 4

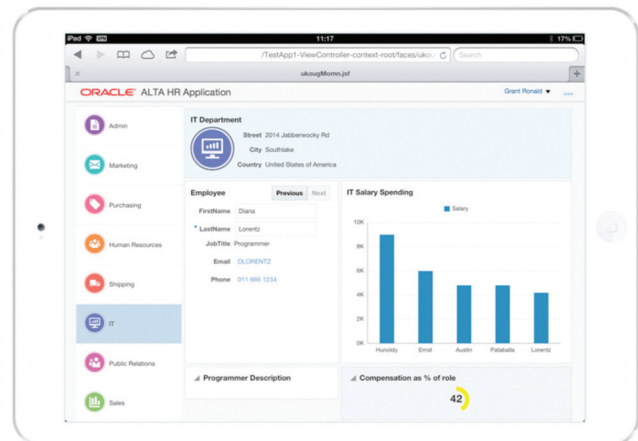


FIGURE 5

part of our lives developing very similar, and decidedly dull, master/detail style of application pages. Figure 4 demonstrates a particularly extreme example. Broadly speaking you have a form layout showing master records, some buttons for navigation and a table of detail rows. Whether you are using Oracle Forms, Oracle ADF or Oracle Apex, you tend to find this is the default layout.

.....

**“What is so wrong with that?” you might be asking. Well firstly, in a straight aesthetic shootout against almost any mobile-specific application, you are never going to be compared favourably.**

.....

Secondly, the whole interaction pattern is so 1980s! Thinking in these “CRUD” terms shows you’re in a database centric mindset. You might argue that it fits the bill of maintaining employee and department records but I’d argue that’s not what your user is really tasked to do. He/she is performing a business function such as hiring a new employee, performing an appraisal or determining if an employee is likely to leave the company. Whilst that may involve a relational database in the background, it doesn’t mean you have to shove your database scheme in their face.

So how might you develop a screen like this in Oracle Alta UI? Of course, you need to understand what are the user’s needs and what are they trying to achieve. But when it comes to implementation you need to familiarise yourself with the Oracle Alta UI patterns. In reading these patterns you find that there is a master/detail interaction pattern<sup>v</sup>. This immediately gives you a couple of examples to aim for, as well as guidance on how to achieve that pattern.

### Skinning

To get the Oracle Alta UI colour scheme and look & feel you need to apply the Alta skin. For an ADF project you set this in the trinidad-config.xml file with the tag <skin-family>alta</skin-family>. This will ensure that when you start building the page that the components you use will have the right colour and style for an Oracle Alta UI.

### Templates

Of course, to ensure consistency there are certain repeatable high-level concepts to each page in Oracle Alta UI. As a result, Oracle JDeveloper offers a tablet-first template which can be used as the basis for your Oracle Alta UI page. When you create a new page, you can choose this template and it will aid you with page sizing as well as including areas for branding, navigation and main content.

### Building the Page

By following the Oracle Alta UI master/details pattern you’ve a well defined target at which you can aim. However, you need to know a bit about the various Oracle ADF

Faces components, especially the layout components, to be able to construct a fully functioning page. Now, a detailed tutorial on the ADF Faces components and how to use them is beyond the scope of what I want to cover here, so we’ll quickly fast forward 20 minutes to see what I was able to create and, broadly speaking, the components I used.

Probably the most immediate take away from looking at figure 5 is that it straight away feels like a “real” mobile application.

The Oracle Alta UI skin gives a clean and clear look to the application and the template ensures the page fits within a landscape iPad and flows vertically.

Oracle ADF Faces provides a component called a listView (as seen on the left displaying departments), which can be used to quickly swipe through the list of departments with each row automatically having a reasonably generous touch point to make them easily selected.

In designing for the tablet I’ve taken the opportunity to very easily hook up icons for each of the departments. This may seem like a purely aesthetic change; however, the icons themselves are representative of the departments and can be an aid to identification.

Focusing on the main content on the screen, I’ve employed a panelGridLayout layout component which has two columns and three rows. This component has a very familiar feel for anyone who has developed with HTML tables allowing you define the number of rows and columns, their size as a percentage and

## Development: Grant Ronald

whether rows span across columns. Here you can see that the top row (light blue displaying the IT Department) is set to span two columns.

Figure 5 also demonstrates the encouraged use of visual components. The user can immediately see from the

bar chart that this employee is the lowest paid in the department and her salary is 42% of the pay grade. Furthermore, the yellow/amber colour of this component conveys that this may not be an immediate concern (green meaning her pay is in the upper percentiles of the job grade, red would show she is in the lower percentiles

and possibly needs action).

You may also notice that by using the panelBox component information is visually grouped and some panel boxes are collapsible if the user decides they don't want to see some particular group of information.

### The Next Step

Of course, I've not finished with this page, it's only a first cut, but with only a template, a few basic components and implementing only a few Oracle Alta UI concepts, I've achieved considerable improvements in only about 20 minutes. I can further enhance my application by following the guidelines for application features such as navigation, branding and even responsive design.

With Oracle Alta UI, along with the associated tooling in JDeveloper, I now have the blueprints and tools to develop a compelling and modern application to address the demands of a very discerning mobile-first audience. ■

<sup>i</sup> <http://bit.ly/oraclealta>

<sup>ii</sup> <https://www.oracle.com/corporate/pressrelease/oracle-alta-093014.html>

<sup>iii</sup> <http://www.oracle.com/technetwork/developer-tools/jdev/index-098948.html#alta>

<sup>iv</sup> <http://jdevadf.oracle.com/workbetter/faces/index.jsf>

<sup>v</sup> <http://www.oracle.com/webfolder/ux/middleware/alta/patterns/MasterDetail.html>



### ABOUT THE AUTHOR

#### Grant Ronald

Director of Product Management, Oracle Application Development Tools

Grant is Director of Product Management in the Oracle Mobility and Development Tools division responsible for Mobile and Cloud technologies including Oracle ADF, Oracle Mobile Application Framework and Mobile Cloud Service. Grant joined Oracle in 1997, working in Oracle support, where he headed up the Forms/Reports/Discover team responsible for the support of the local Oracle Support Centres throughout EMEA. Prior to Oracle, Grant worked for 7 years in various development roles at EDS Defence. He is author of the "Quick Start Guide to Oracle Fusion Development: JDeveloper and Oracle ADF", published by McGraw-Hill.



[uk.linkedin.com/pub/grant-ronald/0/392/ba4/en](http://uk.linkedin.com/pub/grant-ronald/0/392/ba4/en)

[@gwronald](https://twitter.com/gwronald)

# UKOUG brings IT together

UK Oracle User Group is the only UK-based independent user membership organisation, recognised by Oracle, serving the needs of Oracle customers.

From exclusive, focused meetings to our acclaimed national conferences, our membership offerings bring users together in an unbiased, collaborative environment for education, innovation and information.

Call us on +44 (0)20 8545 9670 or email [info@ukoug.org](mailto:info@ukoug.org) today to find out more.



# Demystifying WORKLOAD System Statistics Gathering

You must have system statistics if you want the optimizer to evaluate correctly the cost of full table scans vs. index access. But you have two options to gather them: NOWORKLOAD and WORKLOAD. The latter seems to be the more accurate one because it fits your workload. But do you really know how they are calculated?

Franck Pachot, Senior Consultant, dbi services 

When I recommend gathering system statistics, I always advise checking them in order to avoid weird values (query here: <http://goo.gl/rhyRqA>). Reasons for weird values can be: bug, non-representative workload, filesystem cache effects, any smart optimisation done behind the scene by the storage, etc.

Another reason is the direct-path reads introduced in 11g by serial direct-path reads. They do multiblock reads, but are not taken into account by some system statistics. Let's get deeper and understand how WORKLOAD statistics are gathered.

## The Test Case

On a test database that has no activity, I'm creating a 1000 block table. My goal is to gather WORKLOAD system statistics during a simple table full scan on that table, and see how it calculates SREADTIM, MREADTIM and MBRC.

```
SQL> connect demo/demo
Connected.

SQL> drop table DEMO;
Table dropped.

SQL> create table DEMO pctfree 99 as select rpad('x',1000,'x')
n from dual connect by level <=1000;
Table created.
```

Then I run a simple select between the calls to 'start' and

'stop' procedures of the dbms\_stats WORKLOAD system stats gathering.

```
SQL> exec dbms_stats.gather_system_stats('start');
PL/SQL procedure successfully completed.

SQL> connect demo/demo
Connected.

SQL> select count(*) from DEMO;

COUNT (*)
-----
1000
```

I check the physical read statistics (this is why I have reconnected my session so that I can query v\$mystat without doing the delta).

```
SQL> select name,value from v$mystat join v$statname
using(statistic#) where name like 'phy%' and value>0;

NAME                                VALUE
-----                                -
physical reads                       1000
physical reads direct                 1000
physical read IO requests             23
```

I've read 1000 blocks in 23 i/o calls (thanks to my multiblock read count).



I stop my WORKLOAD statistics gathering:

```
SQL> exec dbms_stats.gather_system_stats('stop');
PL/SQL procedure successfully completed.
```

And check the system statistics that have been set:

```
SQL> select * from sys.aux_stats$;

SNAME          PNAME          PVAL1
-----
SYSSTATS_INFO  STATUS
SYSSTATS_INFO  DSTART
SYSSTATS_INFO  DSTOP
SYSSTATS_INFO  FLAGS          1
SYSSTATS_MAIN  CPUSPEEDNW    2715
SYSSTATS_MAIN  IOSEEKTIM     10
SYSSTATS_MAIN  IOTFRSPEED    4096
SYSSTATS_MAIN  SREADTIM
SYSSTATS_MAIN  MREADTIM     4.708
SYSSTATS_MAIN  CPUSPEED      2715
SYSSTATS_MAIN  MBRC
SYSSTATS_MAIN  MAXTHR
SYSSTATS_MAIN  SLAVETHR
```

I have no SREADTIM because I didn't make any single block read. So it is not set and the CBO will estimate it from IOSEEKTIM+block\_size/IOFRSPEED.

I have MREADTIM because I did enough multiblock reads and they have been timed. 4 milliseconds seems to be accurate for my VM on my laptop.

But MBRC is not set. It should have been set in WORKLOAD mode and I've made multiblock reads. But there is a reason for that, which can be guessed from the above statistics. I've made only 'physical reads direct' and we will see later that system statistics do not account for the multiblock reads that are done in direct-path i/o.

### Tracing dbms\_stats

Yes we can trace dbms\_stats. It's not sql\_trace. It's only some debug messages showing intermediate steps or values. I've seen that referenced for the first time by the Pythian blog ([www.pythian.com/blog/options-for-tracing-oracle-dbms\\_stats/](http://www.pythian.com/blog/options-for-tracing-oracle-dbms_stats/)).

I set the flags '1' in order to have the messages with dbms\_output (and I set serveroutput on) and the '262144' which traces the system statistics gathering (this is new in 12.1.0.2 and was not yet documented in the Pythian blog).

```
SQL> exec dbms_stats.set_global_prefs('TRACE',1+262144);
PL/SQL procedure successfully completed.
```

```
SQL> set serveroutput on
```

And I do the same as before:

```
SQL> exec dbms_stats.gather_system_stats('start');

DBMS_STATS: System Stats Information:
reason is gather_system_stats
start option
target is dict
flag is 51
stattab: statid: statown:
DEMO

Infostats:
```

```
status: MANUALGATHERING, dstart: 01-22-2015 17:21,
dstop: 01-22-2015 17:21, flags: 1
MainStats:
cpuspeednw: , ioseektim: ,
iotfrspeed: , cpuspeed: , sreadtim: , mreadtim: , mbr: ,
maxthr: ,
slavethr:
TempStats:
sblkrds: 15586, sblkrdtim: 28151.772, mblkrds: 2030,
mblkrdtim: 3673.889, mbrtotal: 34254, cpucycles: 76692, cputim:
28245, job: 0, cache_job: 758
```

```
PL/SQL procedure successfully completed.
```

Lots of information came through dbms\_output. I run my full table scan and check that I've the same session statistics about the physical reads I've done:

```
SQL> connect demo/demo
Connected.
```

```
SQL> select count(*) from DEMO;
```

```
COUNT(*)
-----
1000
```

```
SQL> select name,value from v$mystat join v$statname
using(statistic#) where name like 'phy%' and value>0;
```

```
NAME          VALUE
-----
physical reads          1000
physical reads direct  1000
physical read IO requests      23
```

And I stop the WORKLOAD statistics gathering:

```
SQL> set serveroutput on
SQL> exec dbms_stats.gather_system_stats('stop');
```

```
DBMS_STATS: System Stats Information:
reason is gather_system_stats stop
option
target is dict
flag is 31
stattab: statid: statown:
DEMO
```

```
Infostats:
status: COMPLETED, dstart: 01-22-2015 17:21, dstop:
01-22-2015 17:21, flags: 1
MainStats:
cpuspeednw: 2715, ioseektim: 10,
iotfrspeed: 4096, cpuspeed: 2715, sreadtim: , mreadtim: 1.243,
mbr: ,
maxthr: , slavethr:
TempStats:
sblkrds: 15586, sblkrdtim: 28151.772, mblkrds: 2053, mblkrdtim:
3702.479, mbrtotal: 34254, cpucycles: 76692,
cputim: 28245, job: 0, cache_job: 0
```

```
PL/SQL procedure successfully completed.
```

The 'tempstats' are the cumulative statistics that are gathered and dbms\_stat calculates the difference between the 'stop' and 'start'. Let's do it ourselves.

**sblkrds: 15586-15586=0**  
is the number of single block reads (the 'db file sequential read'). We have none of them here because my full table scan is doing only multiblock reads.

**sblkrdtim: 28151.772-28151.772=0**  
is the time spent in single block reads.

**mblkrds:** 2053-2030=23

this matches the 'physical read IO requests' above. Those are my multiblock i/o calls.

**mblkrdtim:** 3702.479-3673.889=28.59 milliseconds

is the time spend in those multiblock reads.

**mbrtotal:** 34254-34254=0

Here the name is misleading. First because the unit is the number of blocks, but we will see that later. However it's not incremented by every kind of multiblock reads, only those that are going through the buffer cache (the 'db file scattered read').

And, in my case, because I'm reading a large table which has no buffers in cache, Oracle is using the serial direct-path reads available from 11g. That's why we have direct-path reads instead of 'db file squattered read'.

So we have a clue here about the reason why MBRC is not set by dbms\_stats WORKLOAD statistics. Only the blocks read to buffer cache are counted and I have none of them. I will verify it below, but first let's check at the statistics gathered:

```
SQL> select * from sys.aux_stats$;
SNAME          PNAME          PVAL1
-----
SYSSTATS_INFO  STATUS
SYSSTATS_INFO  DSTART
SYSSTATS_INFO  DSTOP
SYSSTATS_INFO  FLAGS          1
SYSSTATS_MAIN  CPUSPEEDNW    2715
SYSSTATS_MAIN  IOSEEKTIME    10
SYSSTATS_MAIN  IOTFRSPPEED   4096
SYSSTATS_MAIN  SREADTIME
SYSSTATS_MAIN  MREADTIME    1.243
SYSSTATS_MAIN  CPUSPEED      2715
SYSSTATS_MAIN  MBRC
SYSSTATS_MAIN  MAXTHR
SYSSTATS_MAIN  SLAVETHR
```

Note that we have already seen those numbers from the 'stop' trace 'MainStats'. MREADTIM has been calculated as 28.59/23=1.243

### Without Direct Reads

In order to prove my assumption I disable serial direct path reads:

```
SQL> connect demo/demo
Connected.

SQL> alter session set "_serial_direct_read"=never;
Session altered.
```

I will reproduce here only the relevant lines from the trace:

```
exec dbms_stats.gather_system_stats('start');
...
TempStats:
  sblkrds: 15586, sblkrdtim: 28151.772, mblkrds: 2053,
  mblkrdtim: 3702.479, mbrtotal: 34254, cpucycles: 76692, cputim:
  28245, job: 0, cache_job: 760

PL/SQL procedure successfully completed.

SQL> select count(*) from DEMO;

COUNT(*)
-----
      1000
```

```
SQL> select name,value from v$mystat join v$statname
using(Statistic#) where name like 'phy%' and value>0;

NAME          VALUE
-----
physical reads          1000
physical reads cache    1000
physical read IO requests    30
physical reads cache prefetch    970

8 rows selected.

SQL> set serveroutput on
SQL> exec dbms_stats.gather_system_stats('stop');
...
TempStats:
  sblkrds: 15586, sblkrdtim: 28151.772, mblkrds: 2083,
  mblkrdtim: 3725.163, mbrtotal: 35254, cpucycles: 76692,
  cputim: 28245, job: 0, cache_job: 0

PL/SQL procedure successfully completed.
```

Now that's different. No direct-path reads but still multiblock reads.

MREADTIM is calculated as (3725.163-3702.479)/(2083-2053)=22684/30=0.756

And now we have a value for mbrtotal=35254-34254=1000

My 1000 blocks have been accounted now, in order to calculate the MBRC: 1000/30=33.33

Here are the gathered statistics:

```
SQL> select * from sys.aux_stats$;
SNAME          PNAME          PVAL1
-----
SYSSTATS_INFO  STATUS
SYSSTATS_INFO  DSTART
SYSSTATS_INFO  DSTOP
SYSSTATS_INFO  FLAGS          1
SYSSTATS_MAIN  CPUSPEEDNW    2715
SYSSTATS_MAIN  IOSEEKTIME    10
SYSSTATS_MAIN  IOTFRSPPEED   4096
SYSSTATS_MAIN  SREADTIME
SYSSTATS_MAIN  MREADTIME    .756
SYSSTATS_MAIN  CPUSPEED      2715
SYSSTATS_MAIN  MBRC        33
```

So the formula for MBRC, which can be applied to both cases, is: 'physical read cache'/'physical read' but, from what I've seen (with sql\_trace), it calculates the MBRTOTAL as 'physical\_read'-'physical\_read\_direct' and then MBRC=(MBRTOTAL-SBLKRDS)/MBLKRDS.

## The Formulas

Without going into the detail, here is how the values that are shown are calculated by the tracing of 'start' and 'stop':

```
SQL> select
  singleblkrcs SBLKRDS,
  10*singleblkrdtim SBLKRDTIM,
  (phyrds-singleblkrcs) MBLKRDS,
  10*(readtim-singleblkrdtim) MBLKRDTIM,
  (physical_read-physical_read_direct) MBRTOTAL
from (
  select
    sum(singleblkrcs) singleblkrcs
    ,sum(singleblkrdtim) singleblkrdtim
    ,sum(phyrds) phyrds
    ,sum(readtim) readtim
    ,sum(phyblkrd) phyblkrd
  from v$filestat
),
(select value physical_read from v$sysstat where
name='physical reads'),
(select value physical_read_direct from v$sysstat where
name='physical reads direct')
;
-----
SBLKRDS   SBLKRDTIM   MBLKRDS   MBLKRDTIM   MBRTOTAL
-----
15586      28160        2083      3730        35254
```

Which matches exactly the latest 'TempStats' above.

The '10x' is because v\$filestat times are in 100th of a second. Note that dbms\_stats queries the underlying x\$kcfcio which in 11.2 became in microseconds and this had brought a bug from 11.2 to 11.2.0.2 where the times were 10000 times larger.

According to my tests, here is how the system stats are set from that:

SREADTIM=SBLKRDTIM/SBLKRDS and MREADTIM=MBLKRDTIM/MBLKRDS

Once again, from my tests in 11g and 12c, it seems that SREADTIM and MREADTIM are not set when there are less than 10 i/o calls. That makes sense because it is better to use defaults (from IOSEEKTIM and IOTFRSPEED) rather than non-significant values.

And the observed multiblock read count formula is: MBRC=(MBRTOTAL-SBLKRDS)/MBLKRDS and this one is set only when we have more than 100 multiblocks read (MBRTOTAL-SBLKRDS).

## Single Block Read

Finally, I run the same but doing only single block reads by setting db\_file\_multiblock\_read\_count to 1. As the previous test has read to buffer cache, I need to flush it in order to do some physical i/o.

It will be done with direct-path reads but that doesn't matter because I will have no multiblock reads anyway.

```
SQL> alter system flush buffer_cache;
System altered.

SQL> connect demo/demo
Connected.
SQL> alter session set db_file_multiblock_read_count=1;
Session altered.

SQL> exec dbms_stats.gather_system_stats('start');
...
TempStats:
  sblkrds: 15597, sblkrdtim: 28154.207, mblkrds:
2083, mblkrdtim: 3725.163, mbrtotal: 35265, cpucycles: 76692,
cputim:
28245, job: 0, cache_job: 762

PL/SQL procedure successfully completed.

SQL> select count(*) from DEMO;

COUNT(*)
-----
1000

SQL> select name,value from v$mystat join v$statname
using(Statistic#) where name like 'phy%' and value>0;

NAME                                     VALUE
-----
physical reads                           1018
physical reads cache                       18
physical reads direct                      1000
physical read IO requests                  1018

7 rows selected.

SQL> set serveroutput on
SQL> exec dbms_stats.gather_system_stats('stop');
...
MainStats:
cpuspeednw: 2715, ioseektim: 10,
iotfrspeed: 4096, cpuspeed: 2715, sreadtim: .244, mreadtim: ,
mbrc: ,
maxthr: , slavethr:
TempStats:
  sblkrds: 16706, sblkrdtim: 28425.247,
mblkrds: 2083, mblkrdtim: 3725.163, mbrtotal: 35374, cpucycles:
76692,
cputim: 28245, job: 0, cache_job: 0
```

That time I don't query sys.aux\_stats\$ because I know I've the values above in 'MainStats'.

Because all reads are single block, only SREADTIM has been calculated and the formula above is verified: (28425.247-28154.207)/(16706-15597)=0.244

## Physical Read Total Multiblock Requests

I have deliberately hidden the 'physical read total multiblock requests' statistic in the multiblock read tests above. The full result in the second test was:

```
SQL> select name,value from v$mystat join v$statname
using(Statistic#) where name like 'phy%' and value>0;
```

NAME	VALUE
physical read total IO requests	30
physical read total multi block requests	14
physical read total bytes	8192000
physical reads	1000
physical reads cache	1000
physical read IO requests	30
physical read bytes	8192000
physical reads cache prefetch	970



Do you wonder why I've only 14 multiblock requests accounted when we know that we did 30 'db file sequential read'? That statistic – despite its definition – does not count the multiblocks that are smaller than specified (reason can be extent boundary or blocks already in buffer cache). In my case I've created the

table in an auto allocation extent size where the first 16 extent size is 64k which has only 8 blocks (my block size is 8k). Then only 30-16=14 have been counted. This is the reason why dbms\_stats calculates the number of multiblock reads from v\$filestat rather than from instance statistics.

## Conclusions

The goal was to understand how dbms\_stats gathers WORKLOAD statistics and get the formula which we can use ourselves, for example to get the values from the AWR history in order to put them in a graph and see how they are regular not.

Here is an example of that: [www.dbi-services.com/index.php/blog/entry/when-to-gather-workload-system-statistics](http://www.dbi-services.com/index.php/blog/entry/when-to-gather-workload-system-statistics)

Then you can see the results in a graph and finally set relevant statistics manually. You see how it evolves over time and for all instances (note that when you are in RAC dbms\_stats will gather stats for your instance only).

Final remark: if you are unsure, default noworkload statistics are probably correct for common cases (assuming that you didn't set db\_file\_multiblock\_read\_count). ■



## ABOUT THE AUTHOR

**Franck Pachot**   
Senior Consultant, dbi services

Franck Pachot is a senior consultant at dbi services in Switzerland. He has 20 years of experience in Oracle databases, in all areas from development, data modeling, performance, administration and training. Franck is an Oracle Certified Master and Oracle ACE who tries to leverage knowledge sharing in forums, publications, presentations.



[www.linkedin.com/in/franckpachot](http://www.linkedin.com/in/franckpachot)

@FranckPachot

## MEET A VOLUNTEER: ANDREW

**Name:** Andrew Watkins  
**Job Title:** Assistant Computer Systems Manager  
**Company:** Birkbeck, University of London

### Tell us about yourself in 50 words

I have been a Solaris enthusiast since my University days, so being a co-chair of UKOUG Solaris SIG was a natural progression. My day job means I spend too much time supporting operating systems, keeping up to date with a never-ending list of databases and coding.

### What is your goal in life?

My work goal is to make life easier for our students. My life goal is to win the lottery and spend my days on the beach with no computers in sight.

### What are your interests?

I love sports and, being Welsh, I have a great interest in rugby, but that is not going too well these days. I relish problem solving so I take great pleasure debugging code and applications software.

### When and why did you want to become a UKOUG volunteer?

After a bit of arm twisting I joined the Solaris SIG Committee in 2010. It is my way of giving back having attended the SIG for a few years. It's been a great way of making new contacts, although you can lose friends when you keep badgering them to give talks! I've also volunteered for UKOUG Systems events and UKOUG Tech conference committees. I write for the Oracle Scene magazine and present at the SIG.

### What does a day of volunteering involve for you?

No two days are alike, from organising the next SIG meeting and locating speakers, to preparing a presentation, sending out announcements and the occasional conference call.



### Which volunteer contribution are you most proud of?

It has to be my involvement in the UKOUG Tech conferences where you have the chance to mould a collection of ideas into a full conference. The best thing about volunteering for me is the other volunteers! I get a buzz from their passion for Oracle technology. The Solaris SIGs are rewarding too, you meet a great bunch of people and while sharing a drink you share a good laugh, war stories and, as a bonus, you pick up a lot of technical knowledge.



@awatkins1966

[uk.linkedin.com/in/AndrewWatkins1966](http://uk.linkedin.com/in/AndrewWatkins1966)



**APPS 15**

UKOUG Applications Conference & Exhibition

7-9 DECEMBER 2015

Your Invitation to Submit

# Call for Papers is Open!

Being a member of UKOUG means that you are exposed to the best resources in the industry.

**Hundred of users, speakers from round the globe, Oracle vendors and UKOUG partners flock to our conferences to learn, share knowledge and engage with the best minds in the business in an unbiased, collaborative environment.**

This December, the ICC, Birmingham will play host to the UKOUG Applications, JD Edwards and Technology Conferences & Exhibitions, but before we open our doors we need a world class agenda to offer.

Are you working on an upgrade which will be completed this year? Are you launching a new product to an Oracle audience? Have you a plethora of insider know hows to make an IT departments life easier? Have you got something to say that would be of benefit to the Oracle community? Then we'd love to hear from you.

All we need you to do at this stage is to submit your abstract(s) - up to 2,000 characters about what your proposed 50 minute

session will cover. You will also need to indicate who will benefit from attending the session, what level of experience the audience will need to have in order to enjoy it and what format you'd like to run it as i.e. a roundtable, panel session, presentation etc.

Each conference has its own topic areas to submit under, so explore the relevant conference websites for more details and to submit your abstracts:

**APPS 15** [www.apps15.ukoug.org](http://www.apps15.ukoug.org)  
**JDE 15** [www.jde15.ukoug.org](http://www.jde15.ukoug.org)  
**TECH 15** [www.tech15.ukoug.org](http://www.tech15.ukoug.org)

The call for papers closes on 10th May 2015.



# JDE 15

UKOUG JD Edwards Conference & Exhibition

8-9 DECEMBER 2015



# TECH 15

UKOUG Technology Conference & Exhibition

7-9 DECEMBER 2015



## What happens to your paper(s) after you hit submit?

UKOUG volunteer, Anthony Macey provides some insight based on his experience of being on last year's conference committee: "I first got involved with UKOUG as an ordinary member, then as a speaker and eventually I joined the committee for Tech14. The planning for Tech14 had already started in terms of venue selection well before the kick off meeting in March. We all gathered at the UKOUG offices in London, where we debated and decided some quite fundamental things e.g. the topics, agenda grids, previous year's performance and how we could make it better this year. We took the best part of an afternoon and early evening to churn through the action points for the meeting. After the meeting we continued much of the debate online for the entire year.

Once the call for papers closed, the submitted abstracts were reviewed and scored by even more volunteers and some of the committee.

Essentially the committee's main purpose is to make sure that the right content is available for the paying customer, the right people from Oracle are speaking and that there is a strong coverage of speakers throughout the duration of the conference.

Most of this activity is condensed into agenda planning day which is a long and hard process. It really is quite a difficult nut to crack; trying to line up speakers so that there is good coverage for the entire day for three consecutive days and understanding what people are using or are interested in is quite tricky.

There are hundreds of papers and they have all been scored - some of them won't meet the eligibility criteria, e.g. some papers are blatant product advertising, which can leave some people feeling a little out of pocket. The more tricky papers to select will be very similar to one another so it will be a case of; do you gamble on a new speaker because the content of the talk looks better? Or stick with someone who you know can give a good talk? Trying to strike a balance there is not easy. Eventually late in the evening we had a full agenda grid with gaps for Oracle speakers and reserves for each selected speaker so that we had cover in case of speaker cancellations.

That's it? Not really - the planning, finalising of the grids, a few tweaks here and there were debated online between committee members right up until the agenda hit the press!"

## UKOUG Speaker Awards – Could you be a winner?

If you're lucky enough to be selected to present at our 2015 conferences you're in with a shot at winning one of our 2016 Speaker Awards.

We award Oracle, UK and overseas speakers with awards in the following categories: Best New Speaker; Best Speaker; Best Judged Paper. The results are generated from the call for papers judging scores and the delegate feedback onsite at the event. So not only will presenting raise your profile in the industry, it will also give you the opportunity to be recognised for the great job you did. More about UKOUG Speaker Awards can be found in the Member Activities area of the UKOUG website. ■







# Oracle Database Licensing with VMware

You are running Oracle on VMware and are not 100% sure that you are compliant with the Oracle licensing rules? This article is for you!

Grégory Steulet  
CFO  
dbi services

**As I have been confronted several times with Oracle License reviews with VMware installations, I decided to write this article in order to provide answers to the most frequent questions I have received from customers.**

Licensing is not an easy topic to address and this difficulty is not specific to Oracle. Additionally, this difficulty increases with the number of different software editors and the related software policy you have to manage as a License Manager, IT manager or CTO. In this article, I will focus on Oracle licensing and more specifically on a very frequent use case: Oracle Database installed on a VMware ESX infrastructure. Indeed, there is really few information about this topic in Oracle's documentation and a simple mistake in your setup can easily cost you up to a million dollars. Lots of the information provided in this article originates from the Oracle Software Investment Guide<sup>1</sup>. However, the very first element to read is the Terms and Conditions specified in your Oracle contracts.

## **Oracle on VMware: Supported, but not Certified**

Before diving into the Oracle licensing questions, the very first thing you have to know about any Oracle product, not only the Oracle Database, is that there is no Oracle product certified on a VMware architecture. However, all Oracle products are supported on VMware infrastructure. What does this mean, supported, but not certified? Oracle does not guarantee the proper working of these products on VMware architecture, but will support you in case of incidents not related to the VMware infrastructure. In cases where the VMware could be involved, the Oracle support may ask you to reproduce the incident on a certified environment. Figure 1 shows an extract of the Oracle Support Doc ID 249212.1.

## **Crash Course on Oracle Licensing Terminology**

Since explaining the subtleties of Oracle licensing without sharing some insight about the Oracle licensing terminology is

*Oracle Support will assist customers running Oracle products on VMware in the following manner: Oracle will only provide support for issues that either are known to occur on the native OS, or can be demonstrated not to be as a result of running on VMware.*

*If a problem is a known Oracle issue, Oracle support will recommend the appropriate solution on the native OS. If that solution does not work in the VMware virtualised environment, the customer will be referred to VMware for support. When the customer can demonstrate that the Oracle solution does not work when running on the native OS, Oracle will resume support, including logging a bug with Oracle Development for investigation if required.*

*If the problem is determined not to be a known Oracle issue, we will refer the customer to VMware for support. When the customer can demonstrate that the issue occurs when running on the native OS, Oracle will resume support, including logging a bug with Oracle Development for investigation if required<sup>2</sup>.*

FIGURE 1

very difficult, I will give you some explanation about the most common and basic Oracle licensing terms.

Oracle products are freely available for download on the Oracle website. However, before using it, you need to agree with the Oracle Technology Network (OTN) Developer License Terms. You can find an extract of the LICENSE RIGHTS Chapter below:

*We grant you a nonexclusive, nontransferable limited license to use the programs only for the purpose of developing, testing, prototyping and demonstrating your application, and not for any other purpose. If you use the application you develop under this license for any internal data processing or for any commercial or production purposes, or you want to use the programs for any purpose other than as permitted under this agreement, you must obtain a production release version of the program by contacting us or an Oracle reseller to obtain the appropriate license. You acknowledge that we may not produce a production release version of the program and any development efforts undertaken by you are at your own risk. We may audit your use of the programs<sup>3</sup>.*

A license according to Oracle definition is “the non-exclusive and limited right to use Oracle software under the agreed terms and conditions”<sup>4</sup>. The overarching license rights are described in the Oracle Master Agreement (OMA) and the rights regarding specific products and services are described in the Ordering Document. Specifically, the OMA is the agreement that details the standard rights granted, ownership, restrictions, warranties, disclaimers, confidentiality, etc. The Ordering Document describes the specific products, types of licenses, number of users, level of support, and discounts (if applicable) a customer has ordered and will receive.

In order to quantify and measure the way customers are using Oracle software, Oracle has introduced the notion of License Metric. According to Oracle the “License Metric determines how the usage is being measured, when Oracle licenses its software to a customer”<sup>5</sup>. For the technology related products (i.e. Oracle Database, Application Server products, Business Intelligence Technology products, Identity Management Products, etc.) Oracle uses two metrics:

- a. Processor
- b. NUP (Named User Plus)

The **Processor License Metric** is based on the number of processor cores in the Enterprise Edition or Sockets in the Standard Edition and the Standard Edition One, in the server where the Database or Middleware product is installed and/or running. This metric is used where users are uncountable (i.e. website) or when the number of users changes frequently. It is important to notice that all processors have to be licensed even if one of them is deactivated in the BIOS. In the case of the Enterprise Edition, customers have to count the number of cores multiplied by the core factor in order to determine the number of Oracle processor cores. The core factor determines the coefficient used depending on the processor type. This core factor is defined in the core factor table<sup>6</sup> provided by Oracle at the following address: [www.oracle.com/us/corporate/contracts/processor-core-factor-table-070634.pdf](http://www.oracle.com/us/corporate/contracts/processor-core-factor-table-070634.pdf) (see Figure 2).


Let’s take a simple example in order to demonstrate the calculation method. An Oracle customer runs several Oracle Enterprise Edition Databases on one server equipped with two sockets. Each socket is equipped with an Intel XEON X7560 8 cores processor. The usual Intel core factor is 0.5 (c.f. core factor table). The customer has to license 2 sockets x 8 cores x 0.5 core factor = 8 Oracle processors. If this customer is running the same databases on a Standard Edition or a Standard Edition One, he has to license two Oracle processors since two sockets are occupied.

The NUP Metric is based on users and non-human operated devices accessing the database. This metric is usually used when users can be counted. Customers using this metrics have to license either the number of NUP accessing the Oracle software or the minimum NUP – whichever NUP is higher. The calculation rule for the minimum NUP customers have to license is 25 NUP per Oracle Processor for the Enterprise Edition and 5 NUP for the Standard Edition and the Standard Edition One. Please note that the number of NUP for the Standard Edition and the Standard Edition One is only 5, since the minima in the Standard Edition and the Standard Edition One is not related to the number of processor license.

Let’s come back to the example in the processor section: For the same architecture, the minimum Named User Plus this customer has to license according to the minima rules would be 8 Oracle Processor x 25 = 200 NUP on an Oracle Enterprise Edition and 5 on an Oracle Standard Edition and Standard Edition One.

The same rules described above regarding the Enterprise Edition apply for the Enterprise Edition Options.





Vendor and Processor	Core Processor Licensing Factor
Sun and Fujitsu UltraSPARC T1 processor (1.0 or 1.2 GHz) Only named servers including: Sun Fire T1000 Server, SPARC Enterprise T1000 Server*, with 6 or 8-core 1.0 GHz UltraSPARC T1 processor Sun Fire T2000 Server, SPARC Enterprise T2000 Server*, with 4, 6, or 8-core 1.0 GHz, or 8 core 1.2 GHz UltraSPARC T1 processor	0.25
Sun Netra T2000, 1.0 or 1.2 GHz UltraSPARC T1 processor	0.25
S PARC T3 processor	0.25
Sun and Fujitsu UltraSPARC T1 1.4 GHz Only named servers including: Sun Fire T2000 Server and SPARC Enterprise T2000 Server*, with 8-core, 1.4 GHz UltraSPARC T1 processor	0.5
A MD Opteron Models 13XX, 23XX, 24XX, 32XX, 41XX, 42XX, 43XX, 61XX, 62XX, 63XX, 83XX, 84XX or earlier Multicore chips	0.5
Intel Xeon Series 56XX, Series 65XX, Series 75XX, Series E7-28XX, E7-28XX v2, Series E7-48XX, E7-48XX v2, Series E7-88XX, E7-88XX v2, Series E5-24XX, Series E5-26XX, E5-26XX v2, Series E5-46XX, E5-46XX v2, Series E5-16XX, Series E3-12XX or earlier Multicore chips	0.5
Intel Itanium Series 93XX or earlier Multicore chips (For servers purchased prior to Dec 1st, 2010)	0.5
Intel or AMD Desktop, Laptop/Notebook, or Netbook Multicore chips	0.5
Sun UltraSPARC T2+	0.5
SPARC64 VII+	0.5
SPARC64 X, SPARC64 X+	0.5
SPARC T4 processor	0.5
SPARC T5	0.5
SPARC M5	0.5
SPARC M6	0.5
Sun and Fujitsu SPARC64 VI, VII	0.75
Sun UltraSPARC IV, IV+, or earlier Multicore chips	0.75
Sun UltraSPARC T2	0.75
HP PA-RISC	0.75
IBM POWER5+ or earlier Multicore chips	0.75
All Single Core Chips	1.0
Intel Itanium Series 93XX (For servers purchased on or after Dec 1st, 2010) Intel Itanium Series 95XX	1.0
IBM POWER6	1.0
IBM POWER7, IBM POWER7+	1.0
IBM POWER8	1.0
IBM System z (z10 and earlier)	1.0
All Other Multicore chips	1.0

FIGURE 2 \* SPARC Enterprise T1000 and SPARC Enterprise T2000 Servers may be sold and branded by Oracle, Sun Microsystems, Fujitsu or Fujitsu Siemens

Differences Between the Different Oracle Database Editions  
The main differences between the different Oracle Database Editions are:

- Oracle Database Enterprise Edition has no license restrictions. This is the only edition that allows licensing Database options (e.g. Partitioning, OLAP, Data Mining, Spatial, Enterprise Management Pack, etc.). In addition, Oracle Database Enterprise includes an additional set of features that are not present in other editions (e.g. Data Guard, Transportable Tablespace, Materialised View Query Rewrite, Cross-Platform Backup, Flashback Table, Database, Transaction Query, etc.). In addition this edition is the only one allowed to be monitored by Grid Control or Cloud Control.
- Oracle Database Standard Edition can only be licensed on servers that have a maximum capacity of 4 sockets. This edition encompasses Real Application Cluster, however, the addition of sockets into the cluster nodes cannot be superior to 4. The maximum number of CPUs defined is meant for the entire Real Application Cluster infrastructure; it is not a per node maximum.





- Oracle Database Standard Edition One may only be licensed on servers that have a maximum capacity of 2 sockets. Unlike the Standard Edition, the Standard Edition One does not include Real Application Cluster. In addition, unlike the Standard Edition and the Enterprise Edition, the Standard Edition One does not include Automatic Workload Management.
- Oracle Database Express Edition (XE) can be installed on any size of host machine with any number of CPUs (one database per machine), but XE will store up to 11GB of user data, uses up to 1GB of memory and one CPU on the host machine. Many features integrated in other edition are not present in the Express Edition.

An exhaustive list of differences is provided on the following URL: [www.oracle.com/us/products/database/enterprise-edition/comparisons/index.html](http://www.oracle.com/us/products/database/enterprise-edition/comparisons/index.html)

### Hard and Soft Server Partitioning

According to Oracle's definition: *"Partitioning occurs when the CPUs on a server are separated into individual sections where each section acts as a separate system. Sometimes this is called "segmenting." There are several hardware and software virtualisation technologies available that deliver partitioning capabilities, with varying degree of resource allocation flexibility."*

A partitioning strategy is more and more frequently used by customers. The underlying technologies and products aim to achieve several goals such as:

- Workload balancing by allocating more or less CPU power to specific systems
- Architecture consolidation by running multiple and different operating systems or multiple versions of the same operating system on the same physical server in order to optimise the resource usage
- The deployment of economic models such as "Pay-As-You-Grow" and "Capacity on Demand"

Oracle distinguishes two main types of partitioning:

#### Soft Partitioning

Soft partitioning is defined by Oracle as a flexible way to segment the operating system using OS resource managers, since the CPU capacity allocated to the operating system running an Oracle database can be changed fairly easily. Examples of such partitioning types include (non-exhaustive list): Solaris 9 Resource Containers, AIX workload Manager, Oracle VM, and VMware.

The key message regarding soft partitioning is: *"soft partitioning is not permitted as a means to determine or limit the number of software licenses required for any given server."*<sup>8</sup>

In other terms, customers using such soft partitioning technologies have to license the totality of cores or sockets of the physical systems.

#### Hard Partitioning

Hard partitioning is defined by Oracle as "a way to physically segment a server, by taking a single large server and separating it into distinct smaller systems. Each separated system acts as a physically independent, self-contained server, typically with

its own CPUs, operating system, separate boot area, memory, input/output subsystem and network resources."<sup>9</sup>

Oracle approved hard partitioning technologies as listed in this section of the policy document. They are permitted as a means to limit the number of software licenses required for any given server or a cluster of servers. Approved hard partitioning technologies include: Dynamic System Domains (DSD) -- enabled by Dynamic Reconfiguration (DR), Solaris Zones (also known as Solaris Containers, capped Zones/Containers only), LPAR (adds DLPAR with AIX 5.2), Micro-Partitions (capped partitions only), vPar, nPar, Integrity Virtual Machine (capped partitions only), Secure Resource Partitions (capped partitions only), Fujitsu's PPAR.

Oracle VM Server can be used as a hard partitioning technology if customer bind vCPUs to physical CPU threads or cores as explained in the following paper: [www.oracle.com/technetwork/server-storage/vm/ovm-hardpart-168217.pdf](http://www.oracle.com/technetwork/server-storage/vm/ovm-hardpart-168217.pdf)

### Licensing of VMware on Oracle

As described in the soft partitioning chapter, VMware is considered by Oracle as a soft partitioning technology.

Some white papers provided by VMware such as the one named "Understanding Oracle Certification, Support and Licensing for VMware Environments"<sup>10</sup> state that it is possible to partially license cluster using DRS Host Affinity.

This document has never been approved by Oracle and you won't find anything about such statements in the Oracle documentation, the Oracle licensing terms, or in the Oracle support. Oracle does not allow to partially license an ESX cluster using DRS technology. Up to VMware 5.0, customers running a virtual machine with Oracle installed had to fully license the physical hosts composing the VMware cluster. Since VMware 5.1, if only one Virtual Machine is running Oracle managed within a vCenter Server Instance, you have to fully license all the physical hosts managed by this vCenter Server instance and not only the cluster instances. This decision has probably been taken by Oracle due to the new features coming with VMware vSphere 5.1, particularly one named vSphere vMotion: This new feature leverages the advantages of vMotion (zero-downtime migration) without the need for shared storage configurations. Indeed, this new vMotion capability applies to the entire network.

In addition, in order to license only the vCenter Server Instance (or the cluster up to VMware 5.0) where Oracle products are running, you have to dedicate a separate physical storage to this vCenter (or cluster up to VMware 5.0). Virtual storage separations such as Vplex technology or dedicated LUN are not considered as a physical separation.

Regarding Oracle Standard Edition, it can be licensed if each individual physical server on the vCenter Server instance does not exceed the maximum capacity of 4 sockets. Oracle Standard Edition One can also be licensed if the maximum capacity of each individual physical server in the vCenter Server Instance does not exceed 2 sockets.

#### Practical Case

In the example illustrated in Figure 3 a customer has a VMware infrastructure of version 5.1 with two servers in a cluster and a third separated server. All servers are managed by the same

Vcenter Server Instance.

- The servers in the cluster are equipped with two Intel Xeon with 8 cores each. The third server is equipped with one Intel Xeon with 4 cores.
- Each server on the cluster hosts two Virtual Machines. 8 cores are allocated to each Virtual Machine. The third server has only one Virtual Machine with 4 dedicated cores.
- Oracle Database Enterprise Edition is installed on one Virtual Machine on the cluster (VM1). This VM can only run on server1. The database is used by 300 users.

In this example, the following restrictions apply:

- VMware is counted as soft partitioning and therefore any dedication or restriction of the Oracle programs using it is considered as soft partitioning and hence you will still need to license the whole environment.
- Since no physical server has more than two sockets, Oracle Standard Edition or the Enterprise Edition One can be licensed.
- Oracle will only provide support for issues that either are known to occur on the native OS or can be proven not to be a result of running on VMware.

The customer will have to purchase:

- 450 (36 cores x 0.5 core factor x 25) Named User Plus in Enterprise Edition in order to respect the minima rules.
- Alternatively, the customer can acquire 18 (36 cores X 0.5 Intel core factor) Oracle Processor of Oracle Enterprise Edition.

1 VMware Vcenter with 3 physical hosts and 5 Virtual Machines




FIGURE 3: CUSTOMER VMWARE 5.1 ARCHITECTURE

### Architecture Solutions

Taking into consideration the facts described in the preceding section, what are the possible solutions in order to be compliant with Oracle licensing terms and conditions? Below, you can find a non-exhaustive list of solutions. I voluntarily decided to mostly focus on Oracle propositions:

1. Dedicating and licensing a new VMware vCenter Server Instance to Oracle products with a separated physical storage for Oracle in order to limit the Oracle license footprint to this specific vCenter Server Instance. This solution can be interesting if your company strategy is to install VMware systematically as an “under layer” for each system. The cost of this solution mostly depends on the chosen storage system.

2. Consolidation on a physical host. It is indeed perfectly possible to run several Oracle products of different version on the same physical host. Of course, this solution is not the most flexible, but will cost you close to nothing in terms of license. However, it will be mandatory to migrate all your Oracle products running on Virtual Machine to a physical server and this can be time consuming and tedious. The price of this solution mainly depends on the chosen physical server.

3. Using a hard partitioning technology such as the one presented in the chapter “hard partitioning”. The price and time needed to migrate from VMware to another system varies greatly from one system to another depending on the chosen hard partitioning technology. For instance, Oracle VM (in hard partitioning mode) is capable of loading both VMware and Hyper-V Virtual Machines and converting them automatically to an Oracle VM. An Oracle VM with a one year support will cost you 599 dollars per server with a maximum of 2 CPUs or 1’199 dollars per server with any number of CPUs<sup>11</sup>. This solution also offers the advantage to rapidly deploy the application via the Oracle VM preconfigured templates.

4. Switching on an Oracle Appliance such as the Oracle Database Appliance (ODA X5-2): Indeed, this appliance can be installed in a bare metal or in a virtualised mode. One of the advantages of this technology is that you can license a minimum of 2 cores per node (a total of 4 cores) in a virtualised mode and you can activate up to 72 cores (with increment by 2)! Nowadays, finding a solution of two servers with only 4 overall cores, meaning two Oracle Processor Licenses, is nearly impossible. Unfortunately, one drawback is that you cannot license the Standard Edition on this box. A license revalorisation is however frequently possible, which means you do not have to pay the full conversion price from the Standard Edition to the Enterprise Edition. Finally, Oracle Database Appliance will offer you a simple and highly available solution. The price for the ODA X4-2 according to Oracle’s Engineered Systems price list is 60’000 dollars.<sup>1</sup> The new ODA X5-2 should be around 68’000 dollars and available since April.



FIGURE 4: ORACLE DATABASE APPLIANCE X4-2

5. Using the Server Trusting Partitioning mode. Oracle provides a relatively new third server partitioning mode. It allows the use of Oracle VM Server (OVM) as a means to limit the number of Oracle Processor licenses required in order to license a sub-capacity of total physical cores without the drawback of OVM in hard partitioning mode. This server partitioning mode is named: Trusted partitioning. The usage conditions of trusted server partitioning are the following:

- Using an approved Engineered System (such as Exalogic Elastic Cloud, Exalytics In-Memory Machine, or Oracle Virtual Compute Appliance - OVCA).
- Using an Oracle Enterprise Manager 12.1.0.2 or higher (either in connected or disconnected mode) to monitor VMs participating in Oracle Trusted Partitions.

Customers will get the following advantages when running OVCA with trusted partitioning:

- You do not have to license the totality of cores: Only used physical cores have to be licensed.
- On an OVCA, two Virtual CPUs (vCPU) are counted as equivalent to a physical core. For instance, licensing 4 Oracle processors enables you to use 8 physical cores and 16 vCPUs.
- Licenses are procured in increments of 2 physical cores (1 Oracle Processor), which provide a high level of flexibility.



FIGURE 5: ORACLE VIRTUAL COMPUTE APPLIANCE

- On an OVCA, you can run a Standard Edition, which is not possible on an ODA. Thanks to Oracle Virtualization, you can even run other RDBMS, such as Microsoft SQL Server on Windows, for instance.
- Even if this appliance is an important investment in terms of capital expenditure, such a solution could drastically reduce your operational expenditures.

### Conclusion

Running Oracle products on VMware offers flexibility advantages. However, if you do not have a deep knowledge of Oracle licensing rules, it can expose you to high financial risks. In addition, the probability of facing a license control is quite high.

The core message of this article is not to switch from VMware to another solution. However, if you have any uncertainty, you should ask an Oracle LMS or Oracle sales specialist before designing your Oracle architecture. Of course, dbi services can also help you maintaining your Oracle licenses either through our SLA or through periodic reviews.

I hope that this article has helped you getting a better insight of Oracle licensing on VMware. This is my view of the current rules. Rules are subject to change by Oracle and it is always the responsibility of the customer to make sure that they are correctly licensed. ■

<sup>1</sup> Oracle (March 2014), Oracle Software Investment Guide , [ONLINE] <http://www.oracle.com/us/corporate/pricing/sig-070616.pdf>, [Accessed 13.08.2014]  
<sup>2</sup> My Oracle Support (2014), Support Position for Oracle Products Running on VMWare Virtualized Environments, [ONLINE], Available at: <https://support.oracle.com>, [Accessed 13.08.2014]  
<sup>3</sup> Oracle (2014), Oracle Technology Network Developer License Terms, [ONLINE], Available at : <http://www.oracle.com/technetwork/licenses/standard-license-152015.html> [Accessed 13.08.2014]  
<sup>4</sup> Oracle (2013), Oracle License and Service Agreement, [ONLINE], available at: <http://www.oracle.com/us/corporate/pricing/olsa-ch-e-v111003-070538.pdf> [ACCESSSED 15.08.2014]  
<sup>5</sup> Oracle (March 2014), Oracle Software Investment Guide, [ONLINE] <http://www.oracle.com/us/corporate/pricing/sig-070616.pdf>, [Accessed 13.08.2014]

<sup>6</sup> Oracle (2014), Oracle Processor Core Factor Table, [ONLINE], available at: <http://www.oracle.com/us/corporate/contracts/processor-core-factor-table-070634.pdf> [ACCESSSED 14.08.2014]  
<sup>7</sup> (8,9) Oracle (2013), Oracle Partitioning Policy, [ONLINE], available at : <http://www.oracle.com/us/corporate/pricing/partitioning-070609.pdf> [ACCESSSED 14.08.2014]  
<sup>10</sup> VMware(2011), Understanding Oracle Certification, Support and Licensing for VMware Environments, [ONLINE], available at: <http://www.vmware.com/files/pdf/techpaper/vmw-understanding-oracle-certification-supportlicensing-environments.pdf> [ACCESSSED 15.08.2014]  
<sup>11</sup> Oracle (July 2014), Oracle Linux Support and Oracle VM Support Global Price List, [ONLINE], Available at: <http://www.oracle.com/us/corporate/pricing/els-pricelist-070592.pdf>, [ACCESSSED 15.08.2014]  
<sup>12</sup> Oracle (July 2014), Oracle Engineered System Price List, [ONLINE], Available at: <http://www.oracle.com/us/corporate/pricing/exadata-pricelist-070598.pdf> [ACCESSSED 15.08.2014]



### ABOUT THE AUTHOR

**Grégory Steulet**  
CFO, dbi services

Grégory has more than ten years of experience in database and infrastructure management, engineering. He is an Oracle LMS qualified auditor and has an Executive MBA from the International Institute of Management in Technology, Fribourg (CH).

 [ch.linkedin.com/in/gregorysteulet](http://ch.linkedin.com/in/gregorysteulet)  
 @gsteulet

The views stated in Oracle Scene are the views of the author and not those of the UK Oracle User Group Ltd.







# Busting Some Myths Around Platform-as-a-Service

Interest in cloud platforms is increasing rapidly, but there are some crucial points for organisations to consider before making the jump to ensure that the transition has a positive impact on their bottom line.

**Dermot O’Kelly, Senior Vice President Oracle UK, Ireland and Israel Technology**

In the same way that its “as-a-Service” predecessors have traditionally suffered from (and in some cases still do suffer from) an identity crisis – Platform-as-a-Service (PaaS) cloud offerings are at risk of being so broadly defined that even those familiar with the technology are unclear on where to start.

What is clear, however, is that having access to secure, agile development environments and infrastructure will free IT departments from repetitive maintenance and allow them to focus on responding to the ever-changing needs and requirements of the wider business. Being able to focus IT resource on building new features into existing cloud applications, developing bespoke apps or creating custom dashboards showing real-time data are all compelling possibilities for many businesses.

With the introduction of a cloud platform, IT departments will soon be able to make the fabled switch from cost centre to revenue generator; from maintenance to innovation. However, all cloud platforms aren’t created equal. And, like many other technology investments, deciding which cloud platform is right for the business is an important decision with long term implications.

For those companies that are eager to make the move to cloud, it’s important they are aware of the common myths circulating

around the technology so they can look past the noise and find themselves with a cloud platform that best meets their business needs.

**Myth 1: All cloud platforms offer the same set of services.**

The market is growing steadily as specialist cloud start-ups put forth their value propositions and offer vastly different interpretations of the services that make up a PaaS solution. Most cloud platform providers offer a very limited number of prebuilt services, and many are dedicated to only one type of service, such as storage. Others cobble together a PaaS offering by purchasing smaller companies. These acquisitions are often separate entities and in no way work together easily, although they may appear integrated on their parent company’s PowerPoint slide. Often, the acquiring company has no intention of investing millions in R&D to bring the disparate acquired services together, which means it is unlikely their customers will benefit from a complete service.

If an organisation is spending time and resources on integrating a provider’s clouds, then one of the most significant benefits of cloud technology – simplification – is being overlooked from the start. In addition, IT staff will find themselves spending an inordinate amount of time trying to understand services from multiple cloud provider and expending time and money trying

to connect these different clouds together, something they may not even be able to do at all.

**Myth 2: PaaS is a great catch-all solution to providing capabilities not found in SaaS applications.** Not always. Be on the lookout for SaaS application providers who want organisations to use their legacy cloud platform so they don’t need to improve the functionality of their SaaS apps. They are essentially providing organisations with an old horse and asking them to buy it a new saddle. Some cloud providers even bury expensive cloud platform development services within customers’ monthly SaaS application fees. While bundling is not necessarily a problem, in some cases these providers are really asking organisations to spend more for services that should actually be baked into their basic SaaS applications as basic functionalities to begin with.

**Myth 3: Non-standards-based cloud platforms or open source development languages can bring down costs.** Not necessarily. Finding developers for a programming language that is not well-known can be difficult and time-consuming—not to mention costly. The laws of supply and demand dictate that hard-to-find programmers are more expensive than programmers of industry-standard development languages, such as Java.

**Myth 4: All cloud subscriptions are more or less the same.** Not so. Many cloud providers place numerous, complicated restrictions and limitations on the amount of processing cycles customers can generate in the cloud. In other words, on how much work customers can demand of the platform for their basic subscription fee. Due to these restrictions, customers may be unable to scale their platform to quickly process large amounts of data or perform boiled up analytics that require significant amounts of processing power.

**Myth 5: All clouds are built using the same architecture.** Many providers have designed their cloud platforms to be as cost-effective as possible for them to “run and manage”—rather than to be more convenient for their customers. In practice, this means that the businesses they work with may find themselves forced to share the same database and processing power with other customers on the same platform. This situation is akin to an apartment building in which every tenant shares the same hot water supply—if many neighbours decide to take a long shower at the same time, they’ll all have a chilling experience indeed.

To deal with this and other platform design challenges, many cloud providers have put caps on the amount of processing available with their platform services – the justification being that this way they avoid their entire cloud application going down or underperforming. However, a cloud provider that has all its customers sharing the same application in the cloud exposes every business it works with to poor performance – if one customer hogs all the processing power in any given application, all the others’ response times will slow down, way down. Some cloud providers have so many of these caps and limitations in place that they fill up 50 page word documents. Dealing with these caps actually adds hidden costs and development time to the process of extending or creating new applications.

A modern, personalised cloud platform does not have this problem.

### Best practice for PaaS adoption

Once organisations wrap their heads around the myths clouding the PaaS conversation, they will be in great position to make more informed decisions about how and why to select their ideal PaaS provider.

Those businesses looking to make the leap should bear a few things in mind throughout their search:

- Use a single, well-known, standards-based cloud platform across the entire business. This saves organisations a significant amount of time and money in the long run, and allows them to innovate faster while avoiding unnecessary headaches.
- Use a provider that doesn’t lock customers in to its platform by using lesser-known proprietary languages or databases.
- Look for a cloud platform that connects easily with the business’ SaaS applications, and that makes it easy for customers to move SaaS applications to and from another platform whenever they need to.
- v Finally, choose a provider that has many prebuilt services integrated at every layer. This approach allows customers to innovate quickly while maintaining governance and standards across their entire organisation.

Ultimately, by establishing a foundation today with a secure cloud platform that is built on widely known standards, companies will save themselves both time and money as they make the transition to PaaS. Just as importantly, they will be in an ideal position to pursue their success in the future. ■



## ABOUT THE AUTHOR

### Dermot O’Kelly

Senior Vice President Oracle UK, Ireland and Israel Technology

Dermot O’Kelly is Senior Vice President for Oracle UK, Ireland and Israel region and is responsible for driving Oracle’s operations, growth and profitability across these geographies. He also leads the close alignment of Oracle’s key accounts, and is the Country Leader for the UK.

# Distributed Hinting

It is not easy to optimise a distributed query, but the most important target, usually, is to minimise the traffic between the databases involved. There aren't many steps you can take to achieve this target, but there is a strategy that tends to help and a hint that is particularly relevant.

Jonathan Lewis, Freelance Consultant, JL Computer Consultancy 

## Basic Problem

The biggest problem with distributed queries is that the optimizer doesn't seem to allow for the fact that they are distributed, and once it has acquired whatever statistics it can for all the relevant objects it behaves as if all the objects were in the local database. This statement isn't quite true – there are some indications in trace files from 10g onwards that the optimizer allows some cost for remote execution of distributed queries, and I've seen occasional clues that it has considered executing distributed queries remotely; but the cost adjustments were very small, and the optimizer seemed to ignore, or fail to use, them.

Let's consider an example: I have a query that joins *dist\_home* and *dist\_away*, I collect a few rows from *dist\_home* and for each row in *dist\_home* I find some rows in *dist\_away*. I'm likely to see one of two execution plans: a nested loop join with an indexed access from *dist\_home* to *dist\_away* or a hash join with *dist\_home* as the build table and *dist\_away* as the probe table. Here's an example query and the two plans (nested loop first) when *dist\_away* is a remote table.

There's not a lot of difference between these two plans at first sight, but the cost of the nested loop is significantly higher than the cost of the hash join – so

```
select
  dh.small_vc,
  da.large_vc
from
  dist_home          dh,
  dist_away@orcl@loopback da
where
  dh.small_vc like '1%'
and
  da.id = dh.id;
```

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Inst	IN-OUT
0	SELECT STATEMENT		202	25250	223 (1)		
1	<b>NESTED LOOPS</b>		202	25250	223 (1)		
* 2	TABLE ACCESS FULL	DIST_HOME	202	2020	19 (0)		
3	REMOTE	DIST_AWAY	1	115	1 (0)	ORCL@~	R->S

Predicate Information (identified by operation id):

```
2 - filter("DH"."SMALL_VC" LIKE '1%')
```

Remote SQL Information (identified by operation id):

```
3 - SELECT /*+ USE_NL ("DA") */ "ID","LARGE_VC" FROM "DIST_AWAY" "DA" WHERE "ID"=:1
      (accessing 'ORCL@LOOPBACK' )
```

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Inst	IN-OUT
0	SELECT STATEMENT		202	25250	32 (0)		
* 1	<b>HASH JOIN</b>		202	25250	32 (0)		
* 2	TABLE ACCESS FULL	DIST_HOME	202	2020	19 (0)		
3	REMOTE	DIST_AWAY	2000	224K	13 (0)	ORCL@~	R->S

Predicate Information (identified by operation id):

```
1 - access ("DA"."ID"="DH"."ID")
2 - filter("DH"."SMALL_VC" LIKE '1%')
```

Remote SQL Information (identified by operation id):

```
3 - SELECT /*+ NO_SWAP_JOIN_INPUTS ("DA") USE_HASH ("DA") */ "ID","LARGE_VC" FROM
      "DIST_AWAY" "DA" (accessing 'ORCL@LOOPBACK' )
```

FIGURE 1



```

-----
| Id | Operation          | Name          | Rows | Bytes | Cost (%CPU) | Inst | IN-OUT |
-----
| 0 | SELECT STATEMENT REMOTE |              | 202 | 44844 | 32 (0) |      |        | |
|* 1 | HASH JOIN          |              | 202 | 44844 | 32 (0) |      |        |
| 2 | REMOTE             | DIST_HOME    | 202 | 3434 | 13 (0) |      | ! | R->S |
| 3 | TABLE ACCESS FULL | DIST_AWAY    | 2000 | 400K | 19 (0) | ORCL |        |
-----

Predicate Information (identified by operation id):
-----
 1 - access ("A1"."ID"="A2"."ID")

Remote SQL Information (identified by operation id):
-----
 2 - SELECT "ID","SMALL_VC" FROM "DIST_HOME" "A2" WHERE "SMALL_VC" LIKE '1%'
      (accessing '!' )

```

FIGURE 2

that's the default plan taken in this case. Looking at all the details of operation 3 (REMOTE), though, we can see that really we have to choose the lesser of two evils. (Note that we weren't given much space in the "Instance" column for the database link name for the remote database, and the "IN-OUT" column shows the data flow to be "Remote to Serial").

When we use the nested loop join we have to operate the remote SQL an estimated 200 times – and looking at the Remote SQL Information we see that we're selecting a couple of columns from the remote table for a given ID. That should be a high precision access, so we're looking at a large number of small messages travelling back and forth across the database link.

When we use the hash join we're only going to operate the remote SQL once – but looking at the content of the remote SQL we see that we're selecting all the columns we need from EVERY row in the table. In this case we're going to pull a large volume of (mostly redundant) information across the network (Tip: you can adjust the SQL\*Net parameter SDU\_SIZE and the O/S network buffer size to make this transfer as efficient as possible).

This is often the problem with distributed queries: which option has the smaller impact on your network, a large number of small round trips or a small number of large round-trips. A choice between latency and throughput.

But there is a third way – if we push the selected `dist_home` data to the remote site it will be a bulk transfer of a fairly small amount of data; then we can do the join remotely and push the (relatively) small result set back. In this way we have kept both the volume of traffic and the number of messages to a minimum. Choosing which site actually

runs the query can make a big difference to performance – and that's what the **driving\_site()** hint gives us. The hint should reference a table at the site where you want the query to operate; In this case I would add `/*+ driving_site(da) */` to my query, which would change the hash join plan into the following (see Figure 2):

Note how the first line of the plan now says "select statement **REMOTE**". This is the execution plan as seen from the perspective of the remote database. This is why the reference to the `DIST_HOME` table (operation 2) has become REMOTE; the odd "Instance" identifier of "!" is the name given to the local database when viewed from the perspective of the remote, and notice how the IN-OUT columns now says that it's the `DIST_HOME` table that is transferred "Remote to Serial".

```

select
  sale_date,product, site, qty, profit
from
  sales@&m_target          sal,
  sites                    sit,
  products@&m_target      prd
where
  sit.id = sal.site
and
  prd.id = sal.product
and
  prd.promoted > date'2014-06-17'
;

```

```

-----
| Id | Operation          | Name          | Rows | Bytes | Cost (%CPU) | Inst | IN-OUT |
-----
| 0 | SELECT STATEMENT |              | 3627 | 162K | 10030 (1) |      |        |
| 1 | NESTED LOOPS     |              | 3627 | 162K | 10030 (1) |      |        |
| 2 | NESTED LOOPS     |              | 10000 | 322K | 11 (28) |      |        |
| 3 | REMOTE           | SALES        | 10000 | 263K | 9 (12) | ORCL@~ | R->S |
|* 4 | INDEX UNIQUE SCAN| SI_PK        | 1 | 6 | 0 (0) |      |        |
| 5 | REMOTE           | PRODUCTS    | 1 | 13 | 1 (0) | ORCL@~ | R->S |
-----

Predicate Information (identified by operation id):
-----
 4 - access ("SIT"."ID"="SAL"."SITE")

Remote SQL Information (identified by operation id):
-----
 3 - SELECT "SALE_DATE","SITE","PRODUCT","QTY","PROFIT" FROM "SALES" "SAL"
      (accessing 'ORCL@LOOPBACK' )

 5 - SELECT "ID","PROMOTED" FROM "PRODUCTS" "PRD" WHERE
      "PROMOTED">TO_DATE(' 2013-06-17 00:00:00', 'syyyy-mm-dd hh24:mi:ss')
      AND "ID"=1 (accessing 'ORCL@LOOPBACK' )

```

FIGURE 3

If we check the numbers in this plan we can see, first of all, that the cost hasn't changed from that of the original hash join (which helps to explain why the optimizer hasn't chosen to make this switch between local and remote). But according to the Bytes column we're going to pull 3,434 bytes to the remote database, do the join, then send 44,844 bytes back – compared to pulling 224KB from the remote database when we executed from the local database. That looks like a potential benefit to me (albeit a small one with this little example).

### Join Order

Unfortunately it's not just the choice of **where** to execute the query that matters, and it's possible to run a query from the "right" database but still cause too much network traffic by accessing the tables in the wrong order. Here's a sample query with execution plan (see Figure 3):

We're joining remote tables **sales** and **products** with local table **sites**; that being the case we might consider using replication technology to replicate the sites tables (which sounds as if it shouldn't be subject to much change) to the remote database so that we can do a three-table **remote** (as opposed to **distributed**) join. But we're stuck with what we've got at present and what we've got is a plan where Oracle gets

every row from the sales table (in array fetches) and checks each row against the sites index then, for each result row individually, accesses the products table. As always there's the question of balancing the number of roundtrips and the volume of data to find the best strategy – in this case I'll suggest that I want to operate locally, get a remote join to take place between products and sales (since this eliminates a lot of data as early as possible), pull the result back in a bulk transfer to the local database then join to sites. In this case all I have to do is put in the hint `/*+ leading (prd sal sit) */` (with an optional and currently redundant `driving_site(sit)` to get the following plan (see Figure 4).

The cardinality estimate at operation 2 (REMOTE) is clearly wrong, but the optimizer still manages to get a better estimate at the subsequent nested loop operation; and we can see from the Remote SQL Information that the two-table join we wanted to take place at the remote site has indeed occurred as expected.

Strangely, when I had hinted the order as `/*+ leading(sal prd sit) */` – just swapping the order of sal and prd – Oracle used two remote operations to fetch the data from sales and products separately, then joined them locally with a hash join. It doesn't seem reasonable that the optimizer should arrive at this plan, but that's the sort of surprise you can get with distributed queries – even in 12.1.0.2 which is the version I've been using throughout this article.

### Distributed DML

When we move from “select” to “create as select” we make a horrid discovery: for no obvious reason the `driving_site()` hint is not valid, so we need to find a different way of dealing with the problem of controlling the query. (Note: this behaviour is not a bug it's deliberate; MoS Bug note 5517609 states: “*This is not a bug. A distributed DML statement must execute on the database where the DML target resides. The DRIVING\_SITE hint cannot override this.*”)

There are two well-known ways of working around this problem – sometimes it is possible to make the query “efficient enough” by creating suitable join views at the remote site and then querying the view; the alternative, which I will demonstrate here, is to create a pipelined function to hide the select

```
-----
| Id | Operation          | Name | Rows | Bytes | Cost (%CPU) | Inst | IN-OUT |
-----
| 0 | SELECT STATEMENT  |      | 3627 | 162K | 27 (8)      |      |      |
| 1 | NESTED LOOPS      |      | 3627 | 162K | 27 (8)      |      |      |
| 2 | REMOTE            |      | 45   | 585  | 17 (0)      | ORCL@~ | R->S |
|* 3 | INDEX UNIQUE SCAN | SI_PK | 1    | 6    | 0 (0)      |      |      |
-----
Predicate Information (identified by operation id):
-----
   3 - access("SIT"."ID"="SAL"."SITE")

Remote SQL Information (identified by operation id):
-----
   2 - SELECT /*+ LEADING ("PRD" "SQL" "SIT") */ "A1"."ID",
        "A1"."PROMOTED", "A2"."SALE_DATE", "A2"."SITE", "A2"."PRODUCT", "A2"."QTY",
        "A2"."PROFIT" FROM "PRODUCTS" "A1", "SALES" "A2" WHERE "A1"."ID"="A2"."PRODUCT"
        AND "A1"."PROMOTED">TO_DATE(' 2013-06-17 00:00:00', 'syyyy-mm-dd hh24:mi:ss')
        (accessing 'ORCL@LOOPBACK' )
```

FIGURE 4

statement. Unfortunately this mechanism can't work with CTAS, or with “insert `/*+ append */`”, so you still suffer a penalty for using distributed queries.

I'm going to use the first query I discussed in the article (joining `dist_home` and `dist_away`) to populate a table in the local database. First I create a table, a scalar type and an array type, then I create a pipelined function that “pipes” rows of the scalar type, then I can write a select statement to insert from the pipelined function into the table (see Figure 5):

You'll notice I've added the “hint” “FIND ME” to the embedded SQL- this was to allow me to search the library cache (`v$sql`) for the statement so that I could find its `sql_id` and `child_number` and check its execution plan and some of the execution statistics.

Unfortunately the execution plan seemed to disappear the moment the insert completed, so I had to fall back on the extended SQL trace (event 10046 at level 8) to see where each part of the code ran. This allowed me to see that the

```
create table dist_join (small_vc varchar2(10), large_vc varchar2(200));

create or replace type myScalarType as object (
  small_vc      varchar2(10),
  large_vc      varchar2(200)
)
/

create or replace type myArrayType as table of myScalarType
/

create or replace function pipe_fun
return myArrayType pipelined
as
begin
  for r1 in (
    select
      /*+
        driving_site (da) FIND ME
      */
      dh.small_vc,
      da.large_vc
    from
      dist_home      dh,
      dist_away@&m_target da
    where
      dh.small_vc like '1%'
      and
      da.id = dh.id
  ) loop
    pipe row (myScalarType(r1.small_vc, r1.large_vc));
  end loop;
  return;
end;
/

insert into dist_join
select *
from   table(pipe_fun)
;

commit;
```

FIGURE 5

```

Rows (1st) Rows (avg) Rows (max) Row Source Operation
-----
1111      1111      1111 HASH JOIN   (cr=65 pr=0 time=8342 us)
1111      1111      1111 REMOTE  DIST_HOME (cr=0 pr=0 time=6285 us)
2000      2000      2000 TABLE ACCESS FULL DIST_AWAY (cr=65 pr=0 time=888 us)

Elapsed times include waiting on following events:
Event waited on                      Times    Max. Wait     Total Waited
-----
SQL*Net message to client              15         0.00           0.00
SQL*Net message from client            15        68.18          68.19
    
```

FIGURE 6

embedded query was executed by the remote database, which sent a request for rows from the *dist\_home* table to the local database and pulled them to the remote database in just two fetch

calls. The trace files also showed that the embedded select statement was fetching 100 rows at a time thanks to the standard PL/SQL optimisation of “cursor for loops” – so the execution time was suitably close

to standard array processing time and not the row-by-row processing that the PL/SQL appears to be.

Running *tkprof* against the trace file from the remote session this (with a few cosmetic cuts) is the information I got for the critical query (see Figure 6).

As you can see, the plan shows that the remote database is, indeed, the driving site for this query. The 15 SQL\*Net round-trips also give you some idea of the array-processing efficiency, although you really need to see the details in the trace file to understand exactly where they come from (and why one of them – the last one in the trace file - is 68 seconds).

### Summary

It's possible that the optimizer has some code that is supposed to allow distributed queries to execute at a remote site and some code that should recognise that remote access is more expensive than local access, but at present that code doesn't seem to be functioning properly. Because of this we have to tell Oracle when it would be appropriate to execute a query at a remote site.

We may also need to force Oracle into a particular join order to ensure that when two tables are located at the same database the join between them takes place at that database. The *driving\_site()* hint dictates where the work is done, and the *leading()* hint (possibly with the help of a *no\_merge()* hint) can dictate locality of joins. Unfortunately the *driving\_site()* hint is not valid as part of either “create as select (CTAS)” or “insert as select”, so we have to find an alternative mechanism that allows us to control the distributed query. ■

### Footnote

There are other limitations and problems with distributed joins and I've written several articles about the topic on my blog: <http://jonathanlewis.wordpress.com/category/oracle/distributed/> For example, one of the common strategies to make remotely joinable tables join remotely is to create a remote view joining them and then query the view; but when you query a remote view, the local optimizer isn't able to move inside the view to discover the statistics of the underlying tables – so joins involving remote views can produce very bad execution plans.

That's not the only feature that results in the optimizer losing information and producing bad execution plans, another is that it does not collect histogram information from remote objects, it only collects the simple column-level statistics. I've even got an example where a four table join does a remote join of the first three tables if the cardinality for the first table is less than 1,000 and switches to three separate remote operations – doing effectively the same sequence of nested loop joins – when the estimated cardinality hits 1,000.

All in all, distributed queries lead to lots of traps and, going back to my opening comment on the difficulties; perhaps there are some bugs in the optimizer that have spent the last few years hiding away waiting for someone to raise the SR that will introduce the fixes that will change everything.



## ABOUT THE AUTHOR

Jonathan Lewis

Freelance Consultant, JL Computer Consultancy

Jonathan's experience with Oracle goes back more than 25 years. He specialises in physical database design, the strategic use of the Oracle database engine and solving performance issues. Jonathan is the author of 'Oracle Core' and 'Cost Based Oracle – Fundamentals' both published by Apress and 'Practical Oracle 8i – Designing Efficient Databases' published by Addison-Wesley, and has contributed to three other books about Oracle. He is one of the best-known speakers on the UK Oracle circuit, as well as being very popular on the international scene – having worked or lectured in 50 different countries. Further details of his published papers, presentations and tutorials can be found through his blog: [jonathanlewis.wordpress.com](http://jonathanlewis.wordpress.com).

# UKOUG Systems Event 2015

20 MAY 2015

CAVENDISH CONFERENCE CENTRE | LONDON

**UKOUG**  
UK ORACLE USER GROUP

#ukoug\_sys

**Register now**



## Back to Basics | Series

A series of articles focusing on an introduction to Oracle-related topics and themes

# An introduction to the Oracle RDBMS Architecture

This article (hopefully the first of three) is an overview of “How Oracle Works”. That is, the activities and process that underlie how the core RDBMS functions. Although Oracle is complex, what it does at the core is a set of easy-to-understand tasks. This is not one of the usual two introductory explanations of either how SQL works as a language or how a database is like a filing cabinet (who still uses filing cabinets?). Rather, it is about how the Oracle software moves information around in order to achieve being a robust and reliable store of information.

Martin Widlake, Database Architect & Performance Specialist, ORA600 🐘

Often even experienced Oracle developers and DBAs have never considered the database in this way, but then they don't need to in order to do their jobs - but understanding it makes some key aspects of Oracle make sense. For example, what actually happens when you COMMIT a change to the database?

We won't go into deep technical details as, frankly, we don't need to know them. Indeed, the first thing I will explain is that we can not only think of the Oracle database as a black box that does stuff for us - but in fact, this is the true situation. Look at Figure 1.

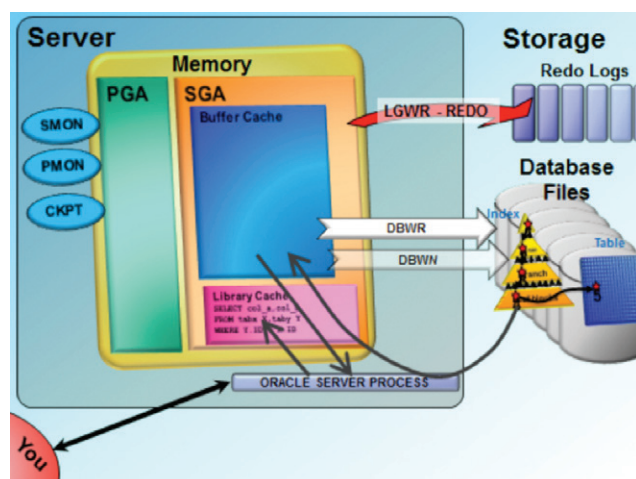


FIGURE 1: ORACLE ARCHITECTURE OVERVIEW

The Oracle instance consists of several parts (which I will go over soon) but notice that you, the user of the database, are represented as the red blob at the bottom left of Figure 1 - outside the server and all you do is talk to an Oracle Server Process. That actually is *all you can do*. When you interact with Oracle, be it via Java or SQL\*Plus or a management reporting tool, you communicate with an Oracle Server Process and make requests of it. It does everything for you and simply feeds back to you the results of your request – your SQL select, update, delete and insert statements. As a user (and usually even as a DBA) you have no direct access to the files of the database or background processes.

**Normally, one Oracle Server Process exists for each connection to the database.**

### Oracle is your Data Waiter

This situation is a little like a restaurant, where all your interactions are via your waiter. You place your orders via the waiter, get your results via them (your prawn salad, steak & chips and chocolate parfait) and they organise for things to be cleared up after you. The Oracle Server Process is your waiter. In a restaurant there are many staff you never see, and may not even know about, that do specific tasks. They interact to move stuff out of the store rooms and fridges (Storage) into the working kitchen (the computer server), process the ingredients into what you want and then get the waiter to take the results back to the customer. The restaurant staff equate to the various internal oracle processes like SMON, PMON, LGWR etc.

That's enough analogy; I want to cover the main parts of the Oracle database architecture.

### What is an Oracle Instance? What is the Database?

An Oracle Instance is what does the work, the thing that processes data. It consists of a chunk of shared memory and a set of server processes, often termed the background processes. Many of us recognise the names of the key background processes, SMON, PMON, CKPT, LGWR etc. With each new version of Oracle there seem to be more background processes, often supporting a new feature of the database such as Data Guard Broker, Flashback Database or Automatic Memory Management. You will probably never need to know what each process does and you can easily look each one up. But it is important to realise that each oracle instance needs a set of processes and the number is increasing with newer versions. This is why it is not sensible to have dozens (and I have come across hundreds) of instances running on one server. Even if the server has several multi-core CPUs, all these processes require resource and time on the CPUs to run.

**A CPU (or rather, a core) can only run one program, one process, at a time – they just swap between them so quickly it seems to us that several things happen at the same time.**

The allocated memory is the other half of an instance and it consists of two main components:

1. The Shared Global Area or System Global Area. Everyone just calls it the SGA. This holds the stuff that is shared between sessions, mainly consisting of the Database Buffer Cache and the Library cache. The Buffer Cache holds data blocks from the tables and indexes that are held in the database. Oracle does nothing directly to the data held on disc, data is read from disc into the Buffer Cache and then worked on. In fact, in many ways Oracle assumes the data is in the Buffer Cache and simply goes there for it. If it is not there, it goes to the storage to get it. The core Oracle RDBMS software has a lot of "intelligence" about trying to keep the most often used data in the Buffer Cache as fetching from memory is several

orders of magnitude faster than fetching from storage. The library cache holds the SQL and PL/SQL that has been run against the database. Most systems run the same SQL and PL/SQL statements over and over again and parsing the SQL statements takes a lot of work, so it is eminently sensible that Oracle caches it and re-uses it.

2. The Process Global Area. This holds the working data set for each session; the result sets coming back from queries, data that is sorted, information about what SQL the user is running and where the statements have got up to. Information in the PGA is not shared between users and that is an important factor. Strictly speaking, the PGA is not part of the Instance, but it is a core part of how Oracle works.

The Database is actually the set of files on storage that hold the tables and indexes. Oracle does not care what the storage is – it can be traditional hard discs, it can be solid state storage, I've even seen a database where some of it was held on CD (It did not work that well! It was a long time ago).

**So long as the storage is visible to the computer server, Oracle does not care. It is just files.**

Without going into the details I'll cover in part two of this series, tables and indexes are stored in blocks of data, usually 8K. Oracle defaults to 8K and 95% of the time it is best to stick with that. Oracle copies data into the Buffer Cache in these 8K blocks. They are read in and written out as whole blocks and sometimes read in as sets of 8K blocks. This sometimes leads to tension with the storage or network teams. The DBA says the storage is slow and the storage team say the database is not reading anything like the volume the storage can provide – but Oracle is doing a lot of its work in 8K chunks, so it is the NUMBER of reads that is the problem, the reads per second or IOPS.

Until the very latest version of Oracle 12c, an instance can only talk to one database (all the users and objects described in one Data Dictionary, if you know something about being a DBA). You can set up Oracle so that several instances (memory-plus-background-processes) each running on a different server can reference the one database, held on shared storage. This is Oracle Real Application Cluster or RAC. I'm going to ignore RAC for now as it only complicates matters – and at the core it does the same few things as a single-instance database.

### The Core of Commit and Recovery – Redo

The main item on Figure 1 that has not been discussed to date is the Redo Logs. The Redo Logs are actually the most critical part of your database. By default **every change made to a database block is recorded in the online redo logs**. It is written as a "redo vector", which holds the information to apply the change. I think of this as the "Golden Change Record" – it is the information that has to be saved to storage for a commit to complete. Until the redo vector has been written and marked as complete, Oracle does not count the change as permanent and the

session issuing the commit will not get the “commit complete” message. Once the change has been written to the online redo log and the commit completed then Oracle (as in the software, not the company) will ensure the change remains.

The redo logs are written to by one and only one background process – the Log Writer, known as LGWR. And as the tag line for the film “Highlander” put it, “There can be only one!”. There can be only one LGWR process running for an Oracle instance. Why? My assumption is that because it is so vital a step it has to be kept as fast as possible and, in computers, fast means simple. One is simple.

So, when a change is made to an Oracle block several steps occur {this is not the complete process}; The Oracle Server Process looks for the block in the Buffer Cache and, if it is not there it fetches it into the Buffer Cache; the change required is prepared; The change is written to the redo log; The change is applied to the block in memory; on commit Oracle ensures the change has been completed to the redo log and only then marks the change as committed and informs the session – i.e. you.

Note, **the changed block has not been written from the Buffer Cache to storage.** The changed blocks are written to disc in a slightly lazy manner, by the Database Writer process, known as DBWR. Unlike the LGWR, there can be several DBWR processes, called DBW0, DBW1, DBW2 etc. On a large computer with many CPU cores more than one database writer will be started at instance startup and more are started automatically as the Oracle software sees fit.

**Why does Oracle do it this way?**  
Again, speed. Writing a block to disc means writing 8K to disc and writing it to a correct place for that block.

A lot of the data written may not have changed, other rows in the block. If the change effects many rows over many blocks, that soon becomes a lot of data to write and a lot of it may not be significant to the change – those other rows in the block. Only the change vector is written to the online redo logs, in 512 byte chunks, and in a simple manner to the end of the log.

### How Oracle Processes Redo Rapidly?

As there is only one log writer, on a busy Oracle database doing lots and lots of changes and commits, the log writer can be the bottleneck. To help with this, Oracle has a set of structures and activities to make redo writing as fast as possible. This is represented in figure 2.

In figure 2, the first thing to notice is the Log Buffer. This is a special area of memory in the SGA and, being in the SGA, it is shared by all Oracle Server Processes. They all write their change vectors to this circular buffer and if need be, prompt the LGWR process to start handling them. LGWR works around this

memory buffer and writes the change vectors to an online redo log. There must be at least two online redo log files, the DBA will often configure a database with more – 3 and 5 seem common. The LGWR writes to one file and, once it is full, starts writing to the next file. This filled redo log is held until all the changes it covers have been written to disc by the DBWR process(es) and the file can be marked as clear. Optionally, it is also held until archived (see below).

The use of the in-memory log buffer and the set of online redo logs mean that the LGWR can process a significant volume of change data in a simple manner, whilst other processes (DBWR and ARC) process the database blocks and copy the filled redo logs to the archive destination. Database activity has peaks and troughs, the buffering allows for higher peaks of activity to be catered for and caught up during troughs. Of course, if the workload is high for a long period, the log buffer may fill or the archiver/database writers may struggle to keep up and the log writer architecture will become the bottleneck.

You will have noticed that there are two online redo log groups in figure 2. **This is not mandatory but it is very sensible.** The redo logs are the golden copy of the change – Changes have to be written to them to be committed and if the instance crashes, the online redo logs are used to recover the database up to the point of the crash, to ensure that all changes that were committed at that point are re-done, irrespective of the data blocks having been written from the Buffer Cache to disc. If the online redo log is corrupted, lost or accidentally deleted (and over the years I have seen all three) then you would be in real trouble. Thus you have a second set of online redo logs - they are duplexed. LGWR will write simultaneously to both sets. I am not aware of any significant performance impact of having duplexed or even triplexed online redo logs.

The final part of the redo log architecture is also not mandatory but is also very sensible, which is that you enable archivelog mode and automatic archiving, i.e. the archiver process(es) copy completed online redo logs to an archive destination. This process is called ARCn and there can be several. This is a simple copy of the whole archive log and is quick. Strictly speaking these archive logs

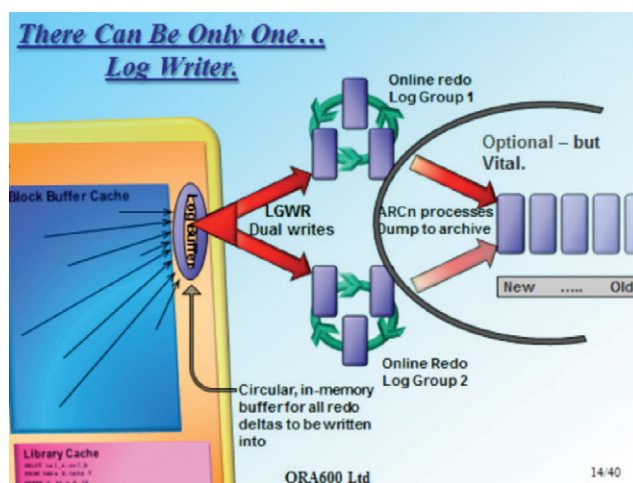


FIGURE 2: THE REDO LOG WRITING MECHANISM



are not classed as part of the database, but they contain all the changes made to the database. They are needed if you want to do hot backups of the database or be able to recover a database to a point in time – or any time later than the last full backup. 99% of businesses want this.

---

**As I said, all changes go into the online redo logs and then the archive redo logs by default. In theory, if you had all the archived redo logs since your created the database, you should be able to recover it to any point in time in the past.**

---

The archive redo logs hold the whole history of the changes to your database and you can mine them with Oracle LogMiner. They are also vital for Standby Database/Dataguard.

Some of you will be aware that you can carry out some actions against the database in nologging or unrecoverable mode. This stops Oracle from recording the changes into the online redo logs and are only allowed for bulk operations, such as SQL\*Loader in direct loader mode. In theory, you could repeat these actions after a database recovery. In practice, people are not very good at being able to identify what tasks have been done like this and may need recovering. I always use nologging with the greatest care.

Many of you will have heard of the new Oracle 12c Multitenant option that allows you to have several databases, pluggable database, under one container database, i.e. under one instance. With the above architecture in mind three things spring to mind. The first is that this could be an improvement over many separate instances on one server as you only have one set of background processes, saving resources. The second is, you still only have one log writer. I'm going to be mindful of not consolidating several busy systems under one container database as that one LGWR process could get very busy! The third is that although there is still only one LGWR process, 12c introduces redo workers – a look at the processes on a Linux 12c database shows lg00 and lg01. I can't say I know enough about these to make sensible comments but spreading the log writing effort is probably a benefit for systems with a high level of change. A little bird has mentioned to me that this "only one LGWR" situation may change in future versions...

So, in summary:

- The Oracle Server Processes read data from the database into the Buffer Cache
- The single Log Writer writes redo data to the online redo logs and that is the key information for commit and recovery
- The Database Writers copy modified data blocks to disc in a more lazy manner
- Separate Archiver process handle the filled redo logs
- Your only interaction with the database is via the Oracle Server Process waiter

Next time I'll give an overview of tablespaces, tables, indexes and how Oracle reads and writes data to them. ■



### ABOUT THE AUTHOR

**Martin Widlake** 

Database Architect & Performance Specialist, ORA600

An independent consultant specialising in Oracle database design, performance and making systems work better. Martin has been working with Oracle technology for half his life. Despite this he is passionate about user groups, sharing knowledge and explaining how Oracle works. He is a regular conference presenter both in the UK and internationally. Martin is an Oracle ACE and a member of the OakTable Network. He has a blog at <https://mwidlake.wordpress.com> which is part technical, part management and part just musing on working in I.T. His real passion is genetics. And cats.

**Blog:** [mwidlake.wordpress.com](https://mwidlake.wordpress.com)



[uk.linkedin.com/pub/martin-widlake/2/7a2/8b/en](https://uk.linkedin.com/pub/martin-widlake/2/7a2/8b/en)



@MDWidlake

# Why You Should Join the Oracle ACE Program

By Jennifer Nicholson, Oracle

The Oracle ACE Program is a community of some of the foremost Oracle technology and product mavens in the world. It recognises not just domain expertise and professional achievement, but also community advocacy: whether they're a database administrator, solution architect or in another technical role, ACE Program members have been vital in helping others succeed when it comes to using and deriving business innovations from Oracle technologies and applications.

If you've been to an Oracle conference or read your fair share of Oracle how-to articles, chances are you've learned directly from an Oracle ACE. Prolific contributions are a common thread amongst ACEs. Birmingham-based ACE Director Tim Hall, for instance, has published more than 700 articles to date on his resource website [ORACLE-BASE.com](http://ORACLE-BASE.com), covering everything from Oracle Database 12c to Oracle PL/SQL.

The ACE Program provides a new level of access and resources to its members, specifically in terms of ACE-only networking and professional development opportunities. Debra Lilley, a member of UKOUG's Governing & Operating Board, describes her tenure as an Oracle ACE Director:

"Once part of the program, you are part of the greatest network of experts – and you get more opportunities to share, creating a perpetual circle of learning and sharing with ever increasing benefits. As an ACE Director, my network covers all Oracle disciplines: I learn from the very best – including development and product managers – and have been able to share my discoveries not just in my local user



ACE DIRECTOR

ACE

ACE ASSOCIATE

group community, but also worldwide through Oracle Technology Network ACE Tours and as a regular contributor to Oracle blogs and Oracle's Profit Magazine."

Today, there are more than 500 Oracle ACE Program inductees across three membership tiers: ACE Associate, ACE, and ACE Director. The program enjoys global representation across 57 countries, with UK and European leaders comprising a big part of the total: currently 200 members live in an EMEA country. (Browse the entire ACE directory at: <http://bit.ly/ace-search>.)

Hopefully by now, we've piqued your interest: we hope you'll consider working toward or applying to become an Oracle ACE! In assessing your own candidacy, the first step is to gauge your community involvement to date: in accepting new ACEs, we look for ways in which you've added value to your own Oracle networks – that

might mean actively volunteering in an Oracle user group, sharing advice regularly through your blog or platforms such as [CommunityOracle.com](http://CommunityOracle.com), and/or presenting at conferences or user group meetings. These and other community-focused activities constitute the requirements for obtaining your ACE membership.

As the top tier of the program, ACE Directors get additional valuable benefits, including:

- An annual exclusive briefing at Oracle headquarters, allowing you to share feedback with Oracle executives and product management;
- Fast-track publication of your articles and papers on Oracle Technology Network;
- Speaking opportunities at Oracle and third-party events and conferences;
- Access to Oracle product development and strategy teams;
- Ongoing support from the Oracle ACE Program office.

We hope you'll take the next step: you can learn more about the Oracle ACE Program, membership requirements and the overall application review process at <http://bit.ly/join-ace-program>. ■



**Organic growth** from modest beginnings to a strong position as the only truly global Oracle eBTax specialist: **eBiz Answers** was founded by Andrew Bohnet, a leading Oracle eBTax consultant.



**THIS ISSUE'S PARTNER SHOWCASE**

# eBiz Answers

**Membership with UKOUG has been instrumental in nurturing growth and development as networking opportunities have fostered flourishing business relationships. "Each event that we attend and are invited to present at, provides a successful channel for demonstrating the benefits of our solutions," comments Andrew.**

"Without UKOUG we would not have known where to start, where to advertise or how to promote the business effectively. We always receive fast and helpful responses to our enquiries and without their support eBiz Answers would not be where it is today."

It was after resolving one client's tax issues within a week, when three eBTax 'experts' had failed in nine months, that Andrew set up the business with the aim of specialising in global solutions using standard Oracle functionality. Providing a 'third party solution without it being a third party product' underlines the consultancy's dynamic approach to offering advice, support, configuration, reporting, system audits, testing and training around Oracle eBTax.

"Due to our level of expertise it proved impossible to find good resources in eBTax, prompting us to develop the eBTax Rapid Install tool™," explains Andrew. "As its name suggests, it allows us to fast track installation; a procedure which could historically take up to 100 days of physical configuration for 40 tax regimes can now be achieved within a few hours."

Forward-thinking innovation which is recognised by the eBTax industry: awarded

OAUG GEO/SIG Certificate of Distinction – Outstanding Achievements Award 2014, Andrew is current chair of the OAUG Tax Management SIG. Winner of three Gold UKOUG Partner of the Year Awards in 2014 for Innovative (Product/Service), Small Business/General Business and Independent Partner of the Year for the contribution made and services offered to Oracle users, eBiz Answers regularly present at leading international events including Oracle Open World and Oracle Collaborate.

"Andrew Bohnet is one of the strongest Functional Oracle eBusiness tax consultants in Europe," comments a senior VAT manager of compliance at RBS. "He has built a superb rapport with the implementation team and users and brought a wealth of valuable Oracle knowledge to the table. I would not hesitate to recommend Andrew to my clients."

What makes eBiz Answers stand out from other consultancies? Andrew and his team offer automated tax solutions that optimise the rich functionality of the R12 and Fusion tax modules. With the development of automated solutions they demonstrate their commitment to focusing on global efficiency and compliance in tax management. Solutions are built in Oracle then tailored to fit each client's specific requirements.

A powerful example of the advantages of this groundbreaking approach is the experience of one of the top five international banks.

"Unable to obtain a sign-off for the indirect tax solution from consultancies

including the 'big 4', we were brought in and took the original 15 day installation time taken by the previous consultancy down to just 2 hours using our eBTax Rapid Install™ tool" explains Andrew.

"In three weeks we installed the additional functionality specified and used eBTax Rapid Install™ to migrate the newly created solution from one instance to another with a total implementation time on a new instance of just two hours. Because our tool allows us to clone the solution as many times as needed we ensured every possible combination of tax elements was covered."

"As we developed our enhanced solution we realised that the end tax rate was only one part of it. We soon recognised the importance of giving the power back to the tax department so they could have real time access to data via our suite of tax reporting tools and prevent tax issues from occurring rather than the traditional route of manual corrections after the financial period has been long closed."

eBiz Answers continues to improve their solution in a quest to change the way companies deal with indirect tax. ■

**Connect with Andrew Bohnet at eBiz Answers:**

-  @ebizanswers
- W** www.ebizanswers.net
- E** andrew@ebizanswers.net
- T** UK +44 78999 44474
- T** US +1 818 435 6435



# PLAN

## Do It Right; Do It Once

You may think that based on the title of this article, it is from a motivational therapist who lives in a van down by the river; not quite. How many times have you embarked on a journey through IT land just to find yourself hopelessly un-prepared for the work at hand? I pray your answer is “not often” but alas in way too many cases, some of us have.

Michael Abbey, Manager; Internal Principal Consultants, Pythian 🍀

We all learned lessons the hard way. How about the DBA that once completely wiped \$ORACLE\_HOME underneath a running database? How about the hapless DBA that wiped \$ORACLE\_HOME/dbs with one *rm* command? There is nothing like planning and planning and, oh ya, did I mention planning? Throughout the balance of this article, we will cover the first few steps in a detailed workplan to rebuild a set of online redo logs on top of a running database – running database you say! Even more reason to practice what I preach – *do it right; do it once.*

### In the Beginning

Upon first examining the requirements for an online redo log rebuild, one easily thinks “piece of cake”. I have done that a bazillion times. I know the nuances and idiosyncrasies of the work at hand and nothing could possibly go wrong/go wrong/go wrong ... Based on that adage and one’s confidence since the work has been done before, let’s now look at the first cut at a workplan to perform the task at hand.

```
Time:      8:00
Date:      12/4/13

Database:  prod
Server    :  margaret

[ ] connect to server
[ ] change to owner of software binaries
[ ] set oracle environment
[ ] log into sqlplus
[ ] switch logfile
[ ] drop old group
[ ] add new group
[ ] drop another old group
[ ] add another new group
[ ] erase files for old groups
```

The sad part about the above list of steps is they are indeed significantly more detailed than some DBAs get for such an important task on a production environment. Let’s pick apart the above steps that conjure up a wealth of questions that accompany each.

#### connect to server

As a seasoned DBA the following questions spring to mind:

1. How do I connect? Is this a Windows box of Linux as I would look foolish trying to ssh to most Windows servers
2. What login do I use to make my first connection?
3. If someone has never connected to this server and knows none of its metadata, where could that information be found?

#### change to the owner of the software binaries

More questions:

1. Who indeed might that be as there are 4 copies of the binaries on the machine?
2. What command should I use to change to that user?

#### set oracle environment

This is pretty straightforward since we already know the desired ORACLE\_SID:

1. What happens when I run *oraenv* and after entering the desired ORACLE\_SID I am prompted for ORACLE\_HOME?
2. If I find the desired ORACLE\_SID missing from *oratab* how do I proceed?

#### log into sqlplus

Easier said than done:

1. What login credentials should I use?
2. What if I am told *insufficient privileges* when using SYSDBA?
3. What if the O/S does not understand the command *sqlplus*?

#### The next 5 commands

These are the heart of the work and the place where, if one does not do it right, may get to do it many many more times than just once:

1. What group should I drop as I believe I need a group number?
2. What query should I used to ascertain the group numbers?
3. What if the group is in use?
4. What if the group is not in use but Oracle is not done with it?
5. ...
6. ...

We could go on for days. We are now ready to move this into the next step. While doing just that some of the perplexing questions from above will be answered...

## Version 2

The following plan is the expanded and much more usable than its predecessor; we will only concentrate on the first three steps to get across the point I am trying to make in this section.

```
Time: 8:00am pacific
Date: Dec 4, 2013

[ ] connect to server
    server name = Margaret
    IP address = 10.21.11.121
    login = pythian
    location of login metadata = client sheet
    connection method = ssh
    COMMAND = ssh pythian@margaret

[ ] change to owner of software binaries
    obtain oracle O/S password
    location of login metadata = client sheet
    COMMAND = sudo su - oracle

[ ] set oracle environment
    ORACLE_SID = PROD
    COMMAND = . oraenv
    REPSONSE = PROD
```

We have made massive amounts of progress since the first version. I commonly use the baseball diamond analogy to explain just how far we have progressed; in this case we are somewhere between first and second. In another land, we are only half way down the pitch. This is what we are missing and is the heart of *do it right, do it once*. Let's build the next version.

## Version 3

```
Time: 8:00am pacific
Date: Dec 4, 2013

Client contact information:
Primary   : Network Ops (NOC)
Technical : Ben Racine

Escalation tree:
Primary   : Jenn Leachman
Secondary : Mac Henderson

Please await the go-ahead from one of the client contact
personnel to begin the work.

[ ] connect to server
    server name = Margaret
    IP address = 10.21.11.121
    login = pythian
    location of login metadata = client sheet
    connection method = ssh
    COMMAND = ssh pythian@margaret
    CHECK-UP = are we on the right server
    COMMAND = hostname
    EXPECTED RESULT = margaret.pythian.com

[ ] change to owner of software binaries
    obtain oracle O/S password
    location of login metadata = client sheet
    COMMAND = sudo su - oracle
    CHECK-UP = are we really the oracle user
    COMMAND = id
    EXPECTED RESULT = uid=501(oracle) gid=20(orainst)

[ ] set oracle environment
    ORACLE_SID = PROD
    COMMAND = . oraenv
    REPSONSE = PROD
    CHECK-UP = is the ORACLE_SID as expected
    COMMAND = echo $ORACLE_SID
    EXPECTED_RESULT = PROD
```

I think we all get the gist of the above and recognise the importance of planning and formulation of the intimate details

of work to be performed. This meticulous planning will protect us against the Achilles heel when performing maintenance ... *human error*. Where's a good place to start—have a look at the following video ([www.pythian.com/blog/how-to-reduce-mistakes-when-managing-a-database](http://www.pythian.com/blog/how-to-reduce-mistakes-when-managing-a-database)) that outlines the direction detailed planning can take to mitigate risk and reduce the likelihood of human error.

## Guidelines

This meticulous planning benefits the consumer by reducing the potential of something going wrong. The most important steps in preparation of a workplan for carrying out maintenance no matter how seemingly trivial are as follows:

- Document the environment
- Assemble contact list
- Final checkpoint on date and time of work
- Prepare detailed list of commands
- Insert checkpoint information after each step
- Share material with second resource to elicit feedback
- Incorporate feedback as/if required
- Circulate final plan to colleagues and client personnel

Following the above will ensure the work is done properly leading to a successful implementation and allow you to *do it right; do it once*. ■



## ABOUT THE AUTHOR

### Michael Abbey

Manager; Internal Principal Consultants, Pythian

Michael is a seasoned presenter in the Oracle database space, having appeared at a wealth of user group and vendor events for the last 25 years. He honed his skills on the core database technology starting with Oracle version 3 in 1986. He was on the team that produced the first offering in the Oracle Press series in 1994. He is Manager of the Internal Principal Consultants at Pythian who offer managed services that cross legacy as well as some of the latest big data offerings such as MongoDB, Cassandra and Hadoop.



[ca.linkedin.com/pub/michael-abbey/1/3a7/5a0](https://ca.linkedin.com/pub/michael-abbey/1/3a7/5a0)

@MichaelAbbeyCAN

This article was first published in RMOUG SQL>UPDATE, Vol 71, Summer 2013.



# The Changing World of High Performance Enterprise Database Solutions

The birth of the current generations of high performance database systems started when Oracle and HP introduced Exadata to the world in 2008, at Oracle OpenWorld. It began with an internal project called SAGE, an acronym for Storage Appliance for Grid Environments. This was to be an open hardware stack solution, i.e. with no proprietary hardware.

Martin Power, Consultant and Director, Parish Crest Solutions 

**The competition was Teradata and Netezza (now owned by IBM), who relied much more on defined or proprietary hardware solutions, so Oracle's aim was to compete through a) openness, and b) Moore's law. The latter meant that as commodity hardware improved with technological advances, the SAGE software running on it would benefit from these advances whilst the competition struggled to test and release newer appliance-based solutions. The idea of openness is also important, because it could be argued that much of Oracle's success in becoming the dominant relational database vendor over the past few decades came from the policy of offering a wide variety of ports and platform choices.**

The headline from the OpenWorld release read as, "Oracle Introduces The HP Oracle Database Machine: Delivering 10x Faster Performance Than Current Oracle Data Warehouses".

In September 2009, Oracle introduces Exadata V2, this time on top of Sun – Intel based hardware utilising its 11g R2 database platform and a significantly upgraded Cluster and ASM extension that allowed individual ASM disks to go offline and return without system outages and long disk synchronisation recovery periods. This new feature is not restricted to ASM, its operation is only available from 11.2.0.2 of the database, with this database compatibility setting needed in the ASM instance for instantiation. So if you have ever wondered why Oracle does not allow Exadata's to run pre 11g R2 databases, this is the reason, a reason that will become more evident as we cover high speed PCIe based flash cards later in this article.

Exadata V2 was a turning point. It was the first Oracle Database appliance architected as no compromise engineered system to be released to the market. Its architecture set new standards and ignored the classic wisdom promoted by vendors like IBM, EMC, HDS and Cisco. Technologies that were beyond consideration, just too new, found their way into this platform. The creation of Storage Cells that ran the database within the

storage device have previously never been seen and today form a major part of Oracle's strategy.

## The Challenge of Today's High Performance Architecture

As a result of Oracle's alignment with open standards and the view that inclusion is paramount in extending market share we are provided with the opportunity to create high performance database systems with performance characteristics that rewrite the rules book on system performance. In the early part of the last decade tablespaces were tuned with data files spread across multiple disks or jbods using combinations of RAID and individual disks to resolve IO bottlenecks. Disks were attached by SCSI interfaces, the fore runners to SAN technologies and then Fibre Channel.

If we spend a minute looking at how this world has changed we see InfiniBand 40 Gigabit network connections used to access both high speed SAN disks, commonly comprising SSDs and background block transition to SAS and finally SATA based on storage over time. These arrays are capable of achieving sustained rates of IO well above 30,000 IO/Per Second. However, if we are to match Exadata then we need another quantum leap in performance, not an incremental increase. One answer to this problem comes in the form of PCIe Flash Cards that reside within the server and provide 1 to 1.5 Million IO/Per Second with ASM Failure Groups used to provide redundancy and storage capacity above 30 Tera Bytes in a single server. The InfiniBand networks compares to Fibre Channel and its 8 Gigbit capacity with 5 times this capacity, and although it was esoteric when Oracle first used it with Exadata V2, today it is readily in reach of any enterprise architecture. 10 Gigabit ethernet provides an alternative lower cost option to InfiniBand, however its performance is less impressive and given network virtualisation options very much the second choice.

Any Exadata DBA will now arch up and tell me there is a



great deal more to Exadata than networks and servers, and this is absolutely true, but for this article I will focus on the foundations of achieving system performance and leave the management and virtualisation until the next article, and I have no illusions, Oracle are well advanced on the next leap, but until then we can explore how we optimise our existing systems to meet the changing performance expectation our business and customers impose on us all.

### The Example System using this Technology

To take the architectural theory into a practice, I have used an AUSOUG partner's technology, HP, and assembled an Oracle certified system configuration using DL380 servers, InfiniBand fabric and switching, HGST or as they are now known, HGST, Flash Storage and built a three node 12.1.0.2 RAC environment on Oracle Enterprise Linux 6.3 with the UEK Kernel to see how things go and yes, No SAN. I have done this as the combination establishes a pattern for blocks of compute and storage that can be used to build scalable RAC environments without the need for SAN storage. So we can now build modular high performance compute systems by racking and stacking servers alone with no compromise in performance, redundancy or scalability. For AUSOUG members there is also a chance for you to test these systems and see just how they work through a series of Webinars and workshops, including an option to carry our a POC of your own to validate the systems and their performance.

The test system comprises,

- 3 x DL380 servers with 64 Gigbytes of memory and 2 x 8 core Intel CPUs.
- 6 x HSGT (HGST) 1 terabyte PCIe cards
- 3 x Dual InfiniBand network cards
- 2 x InfiniBand network switches
- Oracle Enterprise Linux 6.3
- Oracle 12.1.0.2 Enterprise Edition Database and Clusterware

Diagram 1 illustrates our system including its network topology.

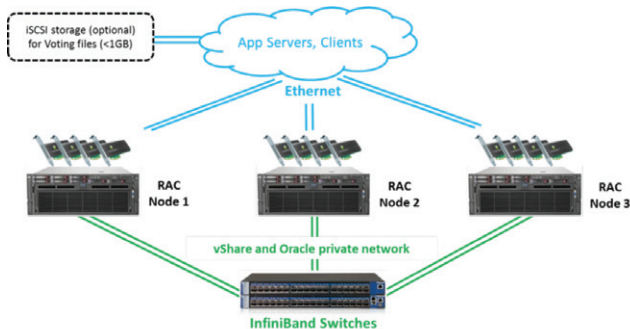


FIGURE 1: NETWORK TOPOLOGY WITH 3 OR MORE NODES

By using InfiniBand we are able to assemble mirrored storage across all three nodes with lower latency and improved performance than we can achieve from fibre channel connectivity with a shared storage array like a SAN. We also have the option to use Oracle's new network virtualisation, previously Xsigo, to establish load aggregation across multiple cards and traffic separation with quality of service standards for key traffic. The layout employed uses ASM failure groups, configured when establishing the initial Disk Groups by setting the redundancy to "High", providing 3 way mirroring of data across the three nodes. The PCIe Flash Cards include RPMs for

Linux to enable these block devices to be exposed to all nodes in the cluster across the network. They therefore appear as shared disks across all nodes and are treated like LUNs with the capacity to name them with meaningful aliases that identify the node, card and logical unit. These devices are then enabled using udev rules and grouped within ASM disk groups. Diagram 2 illustrates this at a logical level, with data striped across all three nodes in our cluster.

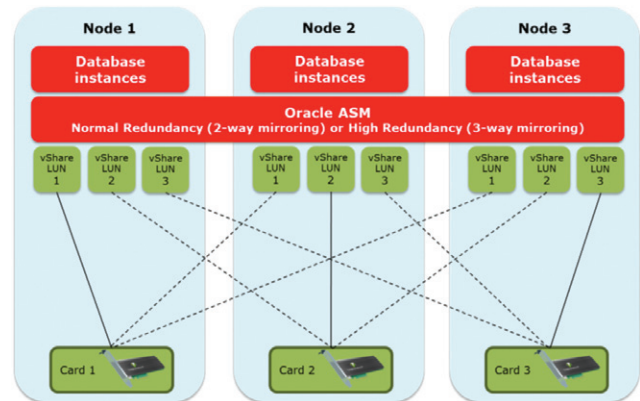


FIGURE 2: LOGICAL PRESENTATION OF LUNs ACROSS ALL NODES AS SHARED DISK STORAGE DEVICES

To obtain maximum performance, dynamic CPU frequency scaling should be disabled. This is especially important when measuring I/O latencies. To disable dynamic CPU frequency scaling run the commands.

```
# service cpuspeed stop
# chkconfig cpuspeed off
```

One of the reasons I selected OEL was that Linux UEKR3 kernel has Open Fabric Enterprise Distribution (OFED) driver 2.0 embedded and does not require additional OFED driver installation. However, installation of InfiniBand software is still required. Use the following commands to install and enable InfiniBand software on each node:

#### To install and enable InfiniBand software on each node:

1. Run the following commands to install InfiniBand software:

```
# yum install infiniband-diags.x86_64
# yum install libibverbs-utils
# yum install rdma
# yum install opensm
```

2. Add the following line to the /etc/rdma/rdma.conf file:

```
MLX4_LOAD=yes
```

3. Enable and start rdma service:

```
# chkconfig rdma on
# service rdma start
```

#### Configuring the Subnet Manager for Two InfiniBand Links

Each InfiniBand link must be a part of a subnet. It is necessary to configure, enable and start the opensmd (InfiniBand subnet manager) service on each server if:

- A back-to-back (no switch) InfiniBand connection is used
- Unmanaged (no embedded subnet manager) InfiniBand switches are used – not recommended for production deployments with two or more InfiniBand ports per server.

If each server has only one InfiniBand link then it is not necessary to customise the opensmd configuration.

If each server has two InfiniBand links then each server must be configured to manage one subnet assigned to one of the InfiniBand ports. This can be done by adding the IB Port GUID of the corresponding port to the opensmd configuration file. For example in a 2-node cluster with two InfiniBand ports (assuming that port 1 on node A is connected to port 1 on node B):

```

• On node A add the line "guid<local port 1 GUID>" to /etc/opensm/
opensm.conf file
• On node B add the line "guid<local port 2GUID>" to /etc/opensm/
opensm.conf file

# cat /etc/opensm/opensm.conf
guid 0x0002c90300e7cd18

IB Port GUIDs are available in the ibstat -p command output as shown below.
# ibstat -p
0x0002c90300e7cd10
0x0002c90300e7cd18

```

### Enabling and Starting the opensmd Service

To enable and start the opensmd service. Run the following commands:

```

# chkconfig opensmd on
# service opensmd start
Starting IB Subnet Manager. [ OK ]

```

This service should be started and enabled on all servers connected either back-to-back or to an unmanaged switch.

### Validating the InfiniBand Connectivity

To validate the InfiniBand connectivity

On each server use the **ibstat** command to verify that the InfiniBand driver is loaded and the InfiniBand links are established.

```

# ibstat
CA 'mlx5_0'
CA type: MT4113
Number of ports: 2
Firmware version: 10.10.1000
Hardware version: 0
Node GUID: 0x0002c90300e7cd10
System image GUID: 0x0002c90300e7cd10
Port 1:
State: Active
Physical state: LinkUp
Rate: 56
Base lid: 4
LMC: 0
SM lid: 1
Capability mask: 0x06514848
Port GUID: 0x0002c90300e7cd10
Link layer: InfiniBand
Port 2:
State: Active
Physical state: LinkUp
Rate: 56
Base lid: 2
LMC: 0
SM lid: 2
Capability mask: 0x0651484a
Port GUID: 0x0002c90300e7cd18
Link layer: InfiniBand

```

If *Physical* state of any link is not *LinkUp* then check whether the InfiniBand cables are connected correctly. If *State* of any link is not *Active* this indicates that the link does not have a subnet configured. If you are using InfiniBand switches that have an embedded Subnet Manager then refer to the administration manual for the switch used to configure subnets. If you are using back-to-back (no switch) connections or InfiniBand switches that do not have an embedded Subnet Manager, see the next section.

Do not stop InfiniBand drivers or openibd service while FlashMAX Connect drivers are loaded (vgcd service is running). This can cause kernel panic.

### Validating Accessibility of All Cluster Nodes via InfiniBand

To verify that all cluster nodes are accessible via InfiniBand. Run the **ibhosts** command as shown below.

```

# ibhosts
Ca : 0x0002c9030036da30 ports 1 "server02 HCA-1"
Ca : 0x0002c9030036d970 ports 1 "server01 HCA-1"

```

### Blacklisting Devices in a Multipath Driver Configuration

If multipath driver is installed it must be configured to have all HGST devices in its blacklist. Otherwise the multipath driver will lock the HGST devices and will prevent normal operation. The multipath driver is typically used when FibreChannel cards are used, but may be present in other cases too.

To add devices to multipath driver blacklist: Add the following to the "/etc/multipath.conf" file (assuming there are no other devices already in the blacklist):

```

blacklist {
devnode "^vgc*"
devnode "^vshare*"
}

```

The above format assumes that vShare devices will be configured with the recommended "vshare" prefix in their names.

### Verifying the Installation of FlashMAX II Devices

Installation of the FlashMAX II devices must be verified when the system is powered on for the first time after installation.

To verify the device installation:

Run the **lspci** command as shown below. This displays a listing for every HGST device installed. It verifies that the device is detected and initialised properly by the PCI subsystem.

```

# lspci -d 1a78: 05:00.0 FLASH memory: Virident Systems Inc.
Device0040 (rev 01)

```

### Installing FlashMAX Connect Drivers and Utilities To install the drivers and utilities

1. Download the RPMs listed below from the HGST website. It is necessary to select the RPMs corresponding to your distribution and kernel version.

- Base/vStore and vCache driver RPM:

```
kmod-vgc-drivers-<distribution or kernel version>-1.2.FC.<driver version>-<driver build>.x86_64.rpm
```

- Utilities RPM:

```
vgc-utils-<distribution name and version>-1.2.FC.<driver version>-<driver build>.x86_64.rpm
```

- RDMA transport driver RPM (for vShare and vHA only):

```
vgc_rdma-<distribution or kernel version>-1.2.FC.<driver version>-<driver build>.x86_64.rpm
```

- RDMA Utilities RPM (for vShare and vHA only):

```
vgc-rdma-utils-<distribution name and version>-1.2.FC.<driver version>-<driver build>.x86_64.rpm
```

## 2. Confirm the running kernel version with the `uname-a` command.

```
# uname -a
Linux hostname 2.6.32-220.el6.x86_64 #1 SMP Tue Dec 6 19:48:22 GMT 2011 x86_64 x86_64 x86_64 GNU/Linux
```

## 3. Install the RPMs using the standard `rpm` command as shown below.

```
# rpm -ivh kmod-vgc-redhat6.1+1.2.FC-65244.V5A.x86_64.rpm
Preparing... ##### [100%]
1:kmod-vgc-redhat6.1+ ##### [100%]
# rpm -ivh vgc_rdma-2.6.32-220.el6.x86_64-1.2.FC-65244.V5A.x86_64.rpm
Preparing... ##### [100%]
1:vgc_rdma-2.6.32-220.el6 ##### [100%]
# rpm -ivh vgc-utils-1.2.FC-65244.V5A.x86_64.rpm
Preparing... ##### [100%]
1:vgc-utils ##### [100%]
# rpm -ivh vgc-rdma-utils-1.2.FC-65244.V5A.x86_64.rpm
Preparing... ##### [100%]
1:vgc-rdma-utils ##### [100%]
```

## Starting the Driver

After installing the drivers and utilities RPMs, the driver will load automatically on every system boot.

To start the driver without rebooting the system:

### 1. Run the **service vgcd** command as shown below.

```
# service vgcd start
Loading kernel modules... [ OK ]
Rescanning SW RAID volumes... [ OK ]
Rescanning LVM volumes... [ OK ]
Enabling swap devices... [ OK ]
Rescanning mount points... [ OK ]
```

### 2. After the driver starts successfully, run the **vgc-monitor** command to confirm that the status is Good.

```
# vgc-monitor
vgc-monitor: FlashMAX Connect Software Suite 1.2(65244.V5A)
Driver Uptime: 13:57
Card_Name Num_Partitions Card_Type Status
/dev/vgca 1 VIR-M2-LP-550-1A Good

Partition Usable_Capacity RAID vCache vHA vShare
/dev/vgca0 555 GB enabled disabled disabled disabled
```

## Formatting and Enabling vShare

Before configuring vShare target, it is necessary to enable them on all FlashMAX II devices. This requires formatting the physical

partition on the FlashMAX II device using the **vgc-config** utility. The formatting mode can be changed while formatting the physical partition. The modes available are:

- **Max Capacity:** This is the default mode. It makes the full advertised capacity of the card available for user data.
- **Max Performance:** This mode is useful for write-intensive applications. It provides twice the sustained random write performance of “Max Capacity” while reducing available user capacity of the device by 15%. Only workloads that generate significant amount of random write I/O will benefit from using the “Max Performance” mode.

Read performance and sequential write performance are the same in both modes.

Note: Oracle ASM version 11.2 has 2TB disk size limit. When using FlashMAX II cards of 2.2TB capacity there are three possible ways to configure them:

- Use Max Performance mode, which reduces user capacity of the card to 1.85TB
- Configure 1999 GB or smaller size in vShare auto-configuration file (See the following section)
- Split each card into two 1.1 TB physical partitions. Contact HGST support for help if this is your preferred option. Do not use `fdisk` or parted utilities. This limitation was removed in Oracle ASM 12c.

## Configuring a vShare Cluster Using a Configuration File

vShare devices can be configured automatically using the **vgc-vshare-auto-config** utility. This utility offers a series of commands that can be used to configure, monitor, and maintain vShare target vSpaces, and initiators in a cluster. This utility reduces the manual effort of configuring servers/hosts individually and makes the vShare configuration easier and faster. Automatic configuration requires password-protected, ssh-based remote root access to the hosts in the cluster. The utility can be used with:

- ssh-setups where the passwords are all the same and setups where different passwords are used on different servers.
- ssh-setups that require shared, password-protected keys for authentication.

The passwords, however, are neither stored persistently nor logged during the configuration.

Configuring vShare automatically can be done by creating a configuration file and then using it as an input to configure the target vSpaces and initiators. The configuration file contains information like cluster name, vShare name, hosts, backing device path, size of the vShare vSpaces, and IB port GUIDs.

To configure vShare automatically with a configuration file:

1. Run the **ibstat-p** command on each server node to get the IB port GUIDs. This step is optional and required only if you want to enable multipathing with two InfiniBand ports in each server.

```
# ibstat -p
0x0002c90300e7cd10
0x0002c90300e7cd18
```



2. Create a configuration file with any name (e.g. example.conf) on any node of the cluster.

```
[cluster]
name: myrac
[vShare:vshare-rac1-a]
host: rac1
backing-dev: /dev/vgca0
size: 300
initiators: rac2 rac3
[vShare:vshare-rac1-b]
host: rac1
backing-dev: /dev/vgcb0
size: 300
initiators: rac2 rac3
[vShare:vshare-rac2-a]
host: rac2
backing-dev: /dev/vgca0
size: 300
initiators: rac1 rac3
[vShare:vshare-rac2-b]
host: rac2
backing-dev: /dev/vgcb0
size: 300
initiators: rac1 rac3
[vShare:vshare-rac3-a]
host: rac3
backing-dev: /dev/vgca0
size: 300
initiators: rac1 rac2
[vShare:vshare-rac3-b]
host: rac3
backing-dev: /dev/vgcb0
size: 300
initiators: rac1 rac2
[ib:path1]
rac1: 0x002590ffff4813fd
rac2: 0x002590ffff481401
rac3: 0x002590ffff481267
[ib:path2]
rac1: 0x002590ffff4823eb
rac2: 0x002590ffff482413
rac3: 0x002590ffff4822d4
```

These configuration files create one vShare target on each of the FlashMAX II cards on all nodes and configure initiators for each target on all other nodes. The targets and initiators are configured so that they will use the specified IB Ports to communicate with each other.

The first section is [cluster] with the cluster name. It can be any name but the first character in the name must be alphanumeric ([a-zA-Z0-9]) or '\_' and can contain alphanumeric characters or any of the characters: '+', '=', '-', or '\_'.

The [vShare:vshare-rac1-a] section specifies parameters for a vShare target named 'vshare-rac1-a'. It is recommended that the name consists of the following:

- A fixed word in the beginning (e.g. "vshare") to make it easy to filter vShare devices by their names
- Name of the server node to make it easy to identify the location of the device
- A letter corresponding to the FlashMAX II device, for example, 'a' for /dev/vgca0

This name is currently limited to 24 characters. The first character in the name must be alphanumeric ([a-zA-Z0-9]) or '\_' and can contain alphanumeric characters or any of the characters: '+', '=', '-', or '\_'.

- "host" is the name of the target host/server
- "backing-dev" is the name of the existing FlashMAX II backing device, /dev/vgc[a-z][0-1] on the target server
- "size" is the size of the target vSpace in GB
- "initiators" is the list of one or more initiator host(s)/server(s) for this target, separated by a space. Typically the initiators list includes all nodes, except for the target node itself

The [ib:path1] and [ib:path2] sections are optional and used for configuring InfiniBand multipathing. Each of these sections provides a list of IB port GUIDs that are part of the same subnet (connected to the same IB switch or, in case of back-to-back connection, connected directly to each other). Each of the two sections must list host names for all nodes with corresponding IB port GUIDs.

3. Run the **vgc-vshare-auto-config-configure** command, as shown below. This creates the target vSpace, grants access, and creates the initiator, with multipathing enabled. The prerequisites, IB connectivity and configuration are verified. The cluster configuration is stored on all cluster hosts as "/var/lib/vshare/\*\_cluster.conf". The steps involved in the configuration are displayed one by one until completion. The process ends with disconnecting the SSH sessions.

- The cluster configuration is persisted on the cluster hosts for use by the other management commands.
- If the command is run on a host that is not a member of the cluster, the cluster configuration file is not persisted locally. Future management commands must be run from one of the configured cluster hosts.

```
# vgc-vshare-auto-config --configure example.conf
Is the same root password used for all hosts [y/N]: y
Enter root password:
Step: Verify Passwords
[rac1, rac2]: Opening SSH sessions
Succeeded.
...
Step: Completion
[rac1, rac2]: Disconnecting SSH sessions
Succeeded.
```

### Verifying the vShare Devices

After the vShare cluster is successfully configured verify that all vShare devices are visible on all cluster nodes with the **ls-l** command, as shown below.

```
# ls -l /dev/vshare*

brw-rw----. 1 root disk 250, 32 Dec 25 22:16 /dev/vshare-rac1-a
brw-rw----. 1 root disk 250, 48 Dec 25 22:16 /dev/vshare-rac2-a
```

### Enabling Synchronised Boot Service

In an Oracle RAC environment **vgc-vshare-wait** service must be enabled immediately after configuring vShare and before disk groups are created in Oracle ASM. During system boot the **vgc-vshare-wait** service pauses the system boot process until all vShare initiators are in a connected state. The service then exits and the Oracle **ohasd** service starts. Without this service enabled when bringing up a vShare/RAC cluster ASM may find some of its disks missing and will fail to start. This might lead to unexpected startup failures if any of the cluster nodes boots with a significant delay compared to other nodes. The **vgc-vshare-wait** service is only for environments running Oracle RAC. If Oracle is not installed on the node and the **ohasd** service is not present then the **vgc-vshare-wait** service exits without waiting.

#### To enable the service

1. On each server node that has vShare initiators configured, run the **chkconfig vgc-vshare-wait on** command. This enables the service.
2. Verify that the service has been enabled, with the command **chkconfig --list | grep vshare**.

```
[@initiatorhost]# chkconfig vgc-vshare-wait on
[@initiatorhost]# chkconfig --list | grep vshare
vgc-vshare-wait 0:off 1:off 2:on 3:on 4:on 5:on 6:off
```

The **vgc-vshare-wait** utility has configuration options that the administrator can customise. They are found in the “vgc-vshare-wait.conf” file.

### To customise the configuration options

1. On the target server, run the command **vim /etc/sysconfig/vgc-vshare-wait.conf**. This opens the configuration file in the vim editor. You can use any available text editor. The options **TIMEOUT**, **SLEEP\_INTERVAL**, **MSG\_THROTTLE**, and **IGNORE\_OHASD\_AUTOSTART** are displayed with the default/recommended values along with comments explaining about the accepted values.

```
# Copyright (C) 2013 Virident Systems, Inc.
# TIMEOUT < 0 Wait until all devices are connected.
# TIMEOUT = 0 Do nothing (no check) and exit. disables the
script.
# TIMEOUT > 0 Check for vShare devices for this number of
seconds.
# TIMEOUT > 0 always checks at least once.
TIMEOUT=-1
# SLEEP_INTERVAL the number of seconds to wait between checking
vShare status.
# SLEEP_INTERVAL will be adjusted to be at least 0 and at least
TIMEOUT seconds.
SLEEP_INTERVAL=5
# MSG_THROTTLE controls how often this script prints runtime
details of the disconnected vShares.
# Note MSG_THROTTLE < SLEEP_INTERVAL implies MSG_THROTTLE =
SLEEP_INTERVAL.
# Status messages will always be printed when TIMEOUT expires or
all vShares are connected.
MSG_THROTTLE=5
# IGNORE_OHASD_AUTOSTART controls whether this script will
change the current ohasd autostart value to disable if vShares
aren't connected at timeout.

# IGNORE_OHASD_AUTOSTART = 0 Allow changing.
# IGNORE_OHASD_AUTOSTART = 1 Don't change anything.
IGNORE_OHASD_AUTOSTART=0
```

The **TIMEOUT** value specifies the amount of time in seconds until which the **vgc-vshare-wait** service will wait for vShare initiators to be connected. By default the **vgc-vshare-wait** service will wait indefinitely until all vShare initiators are in a connected state. This value can be set as follows:

- A **TIMEOUT** value of ‘-1’ causes the service to wait indefinitely until all devices are connected. This is the default value.
- A **TIMEOUT** value of ‘0’ disables the script and no checking of vShare status is performed.
- A **TIMEOUT** value greater than ‘0’ will check the connection of the vShare devices for that amount of time.

The **SLEEP\_INTERVAL** value is the number of seconds that the service waits between checking the vShare connection status. By default the **vgc-vshare-wait** service will check the vShare connection status, once in every 5 seconds.

While waiting for vShare devices to connect, the **vgc-vshare-wait** service prints periodic status messages. The **MSG\_THROTTLE** value controls how often the service prints runtime details of the disconnected vShares.

The **IGNORE\_OHASD\_AUTOSTART** value specifies whether or not **vgc-vshare-wait** should attempt to change the automatic startup of Oracle Clusterware High Availability Services. If the **TIMEOUT** value is set and expires before all vShare devices are in a connected state, the **vgc-vshare-wait** will change the current Oracle Clusterware High Availability Services autostart value in order to disable its startup. To disable changing the

ohasd autostart value, this value should be set to 1.

### To disable the service

1. On the target or initiator server, run the command **chkconfig vgc-vshare-wait off**. This disables the service.
2. Verify that the service has been disabled with the command **chkconfig --list | grep vshare**.

```
[@initiatorhost]# chkconfig vgc-vshare-wait off
[@initiatorhost]# chkconfig --list | grep vshare
vgc-vshare-wait 0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

All devices used by Oracle ASM must have correct permissions set. If the permissions are not set correctly ASM will not allow using those devices. The permissions can be set using a UDEV rule or using ASMLib. HGST recommends using a UDEV rule as described below.

1. Create a file named **/etc/udev/rules.d/99-vgc-oracle.rules**
2. Add the following line to the file and save it. Change the owner and group names as needed.

```
KERNEL=="vshare*", OWNER="grid", GROUP="dba", MODE=660
```

3. Run the following command to apply the rule.

```
# udevadm trigger
```

4. Verify that the permissions are correct.

```
# ls -l /dev/vshare*
brw-rw---- 1 grid dba 252, 0 Dec 4 15:03 /dev/vshareXXXX
brw-rw---- 1 grid dba 252, 16 Dec 4 14:59 /dev/vshareXXXX
```

### Measuring Storage Performance

The easiest way to assess performance of an Oracle storage subsystem is using Oracle’s standard PL/SQL procedure called **CALIBRATE\_IO**. This procedure uses real Oracle database processes accessing the actual blocks in the database files. It can be used with any storage configuration including Oracle ASM and including various redundancy modes. Its use is dependent on **async\_io** being set to true, the default with ASM.

The following query will display the current state of your environment and the enabling of **async\_io**.

```
SELECT d.name, i.async_io
FROM v$datafile d, v$iostat_file i
WHERE d.file# = i.file_no
AND i.filetype_name = 'Data File'
/

/u01/oradata/CDB1/CDB1/EAf64EB7BEE6102EE04500000000001/datafile/
o1_mf_martin_98 3n9kxz_.dbf
ASYNC_ON
```

If the tablespaces and their matching data files are not displayed with **ASYNC\_ON** then run the following command and restart the database.

```
ALTER SYSTEM SET filesystemio_
options=setall SCOPE=SPFILE;
```

Syntax: DBMS\_RESOURCE\_MANAGER.CALIBRATE\_IO (<DISKS>, <MAX\_LATENCY>, iops, mbps, lat);

The CALIBRATE\_IO procedure has two input parameters:

- **DISKS:** This parameter affects the increment and the maximum number of outstanding I/Os used during the test. In case of HDDs, user has to set this parameter equal to the number of physical spindles. However, due to much higher concurrency of FlashMAX II, we recommend setting it to 4 per card. For example, if you have 6 cards in the cluster, set the parameter to 24.
- **MAX\_LATENCY:** This parameter sets the limit for acceptable latency in milliseconds. The procedure will stop increasing the amount of outstanding I/Os when latencies become higher than this value. We recommend setting this parameter to the minimum allowed value of 10 milliseconds.

The CALIBRATE\_IO procedure has three output values:

- **iops:** This value returns the best measured number of I/Os per second achieved with latencies lower than MAX\_LATENCY.
- **mbps:** This value returns the best measured bandwidth in MB/s using large blocks. Note that the mbps value is not the iops value multiplied by block size. A separate test using larger block sizes is used for measuring the mbps value.
- **lat:** This value returns the actual latencies in milliseconds measured during the iops test. On vShare, the iops value is expected to be zero in most cases, as the latencies are <1ms.

#### To execute the CALIBRATE\_IO procedure

Run the following PL/SQL script in SQL\*Plus. It takes 10 minutes per node (30 minutes on a 3-node cluster) to complete. Make sure to modify the DISKS parameter corresponding to the number of FlashMAX cards being used (4 per card).

```
SET SERVEROUTPUT ON
DECLARE
  lat INTEGER;
  iops INTEGER;
  mbps INTEGER;
BEGIN
  DBMS_RESOURCE_MANAGER.CALIBRATE_IO (24, 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;
/
```

```
oracle@tm01:~
File Edit View Search Terminal Help
SQL> select name from V$databases;
NAME
-----
CAL4K

SQL> SET SERVEROUTPUT ON
SQL> DECLARE
  2 lat INTEGER;
  3 iops INTEGER;
  4 mbps INTEGER;
  5 BEGIN DBMS_RESOURCE_MANAGER.CALIBRATE_IO (24, 10, iops, mbps, lat);
  6 DBMS_OUTPUT.PUT_LINE('max_iops=' || iops);
  7 DBMS_OUTPUT.PUT_LINE('latency=' || lat);
  8 DBMS_OUTPUT.PUT_LINE('max_mbps=' || mbps);
  9 end;
 10 /
max_iops=1434268
latency=0
max_mbps=12310

PL/SQL procedure successfully completed.
SQL>
```

## Summary and Findings

The system we assembled delivered a sustained 1.43 Million IOPs with immeasurable latency using 4.8Tb Flash Cards. We can compare this to systems of three years ago that typically achieved less than 20,000 IOPS based on enterprise grade SAN technologies and fibre channel networks. Its assembly comprised off the shelf components from recognised and trusted suppliers like HP and HGST. The combination is an Oracle certified configuration, illustrating just what can be done to engineer cost competitive high performance solutions. ■

## References

Exadata history details, <http://flashdba.com/history-of-exadata/>

Test Systems setup guidelines, Oracle RAC 11.2 with FlashMAX Connect™ vShare ver. 1.2 Deployment Guide rev1.01



## ABOUT THE AUTHOR

### Martin Power


Consultant and Director, Parish Crest Solutions

Martin Power is an Oracle ACE, member of Oracle's Beta Test Teams for Database, Virtualisation and Linux operating systems. He is the current president of the Australia Oracle User Group and a regular contributor to the Oracle community's publications and as a speaker. He works with Parish Crest, an Australian Oracle Gold Partner and professional services company in the Infrastructure, database administration, architecture and middleware. He can be contacted via email at [president@ausoug.org.au](mailto:president@ausoug.org.au)

This article was first published in AUSOUG's Foresight magazine in 2014. Reprinted with permission.



# Accelerate Oracle® Databases with Smarter, Faster and More Reliable Storage



FlashMAX® PCIe SSDs







Ultrastar® 12Gb/s SAS SSDs

**6x**

**Six times better performance at a fraction  
of the cost of typical SAN for Oracle RAC**

Traditional architectures cannot keep up with growing application performance needs. HGST FlashMAX® PCIe SSDs, HGST Share software and Oracle® ASM together offer a new cost-effective approach to implementing high-performance, scalable Flash storage for Oracle and Oracle RAC databases.

## Key benefits include:

-  Consistently high performance with linear scalability
-  DBAs fully control storage resources with Oracle ASM
-  0.5TB to 72TB of Flash in a standard x86 server
-  One-third the cost of a SAN or a Flash array

HGST offers a full breadth of solutions tailored specifically for Oracle environments, including single-instance caching and RAC caching.

Learn more at: [www.hgst.com/solutions/Oracle-RAC-Solutions](http://www.hgst.com/solutions/Oracle-RAC-Solutions).

OracleScene **additional**  
DIGITAL **content**

# Cloud Applications- Introducing a New Kind of User

**As an ACE Director I speak at a lot of Oracle conferences about Cloud Applications (formerly known as Fusion), and the audience almost exclusively falls into two camps...**

Debra Lilley, VP Certus Cloud Services, Certus Solutions 

**Partners who have an insatiable appetite for all things 'Fusion'. The world is changing but with all products there is a lifecycle and not all partners have made the move yet, but know they will.**

**Existing Oracle applications customers, who are thinking about the value to them or simply want to know what Oracle is doing.**

However, Oracle Cloud Application customers are not all from this second group, Cloud applications have opened up Oracle to a wider audience, the mid market.

Oracle and partners also talk to the line of business more now in the sales cycle. I joined Certus in September, they specialise in HCM Cloud, and in November I joined them at the The Chartered Institute of Personnel and Development (CIPD) conference in Manchester; having conversations with HR specialists from organisations of all sizes. These HR professionals know that their most precious assets are also their most complex; their employees. This was a new challenge for me. I normally talk about the

value the technology can deliver from a business perspective, but here they didn't care about the technology, they were only interested in the functionality.

HCM systems are normally a manual process, until an organisation reaches the stage of needing a process that cannot be handled by any single person. At that point, a business may be limited in what it can afford. Traditionally an organisation would then have to keep moving to a bigger and better HCM IT system as it grows. They couldn't automate everything they needed, but it would be a step in the right direction. When it reached a critical size they would start looking at serious IT players. However, these best of breed applications may still have been beyond their reach. The cost of hardware, licensing, implementation and on-going maintenance could simply be too expensive.

Cloud can be the solution, Cloud changes all this. Cloud is about access to a system, not ownership. The vendor owns the hardware and software and simply leases you the right to use the application. The vendor is responsible for maintenance, patching, security; everything.

So nowadays, much smaller organisations are able to use systems that can transform their employee experience, simple self-service that can empower them, mobile access to manager functions and a truly world class set of processes for HR professionals, enabling their people to be as effective as possible in delivering success for their organisation.

You hear the Cloud might be more of a lock-in to a vendor yet in the mid market I think it is easier to walk away. Cloud to them is about Choice. They have a smaller footprint and can move, if the service excellence they demand is not delivered, they will not renew.

Oracle HCM Cloud is growing rapidly in this area, with many net new customers. This new segment of customers for Oracle, and there are some great stories going to be references later this year, can be a challenge. With the deal size being smaller they will have less of a voice with Oracle and the need for a common voice will be greater, they will need a user group.

It's exciting times in the applications world, I'm loving it. ■



## ABOUT THE AUTHOR

Debra Lilley 

VP Certus Cloud Services, Certus Solutions

Debra Lilley, VP Certus Cloud Services at Certus Solutions, ACE Director and Board Member Advocate UKOUG, Debra has worked with Oracle Applications for 18 years.

**Blog:** [debrasoracle.blogspot.co.uk](http://debrasoracle.blogspot.co.uk)



@debralilley

# Fusion Middleware Provisioning with Cloud Control

Andreas Chatziantoniou, Foxglove-IT BV  & Ingo Reisky, OPITZ CONSULTING

Oracle Enterprise Manager Cloud Control 12c (EMCC12c) is an excellent tool for monitoring Oracle environments, especially those containing Fusion Middleware (FMW) components. Setting up of such environments in the real world often uses bespoke scripting solutions or just Oracle Configuration Tools (config.sh). Cloud Control offers possibilities in the provisioning arena to assist in creating and administrating Oracle software and FMW domains. The added value is obviously the fact that such an environment can be provided consistently and repeatedly. This is completely embedded in the framework of Cloud Control employing concepts of privileges and roles, but also offering reporting capabilities and configuration auditing.

In this article we would like to demonstrate how such a solution can be designed and implemented. Our main focus lies on the quality of the provisioned environments because variations and discrepancies in the various project phases (DEV, TST, INT, PRD) are avoided.

## Introduction

An Oracle Fusion Middleware (FMW) environment can relatively easy evolve into a very complex entity. Often more than 100 (sub-) activities need to be carried out just to set up one environment only once!

Very soon such provisioning needs to be executed more than once and, in projects with several lifecycles (DEV, TST, INT, PRD), the problem arises to ensure and guarantee that a provisioning is consistent with its predecessor and successor. Furthermore, each provisioning can be very time consuming.

Additionally, each provisioning operation (possibly) needs adaptations according to the project phase, e.g. a production environment has RAC databases while previous stages used single instances, and so on.

Many organisations employ a script-based approach (often its historic roots are buried in time) or execute the tasks manually. Here a number of well-known issues arise: traceability and governance (Who built the INT environment? Which changes were carried out in DEV?).

## What is provisioning anyway?

Provisioning knows a number of aspects that typically cross the boundaries of classical system management teams and realms. Especially in Oracle FMW environments various components are to be addressed:

- Hardware
- Network
- Storage
- OS
- Oracle Software (DB, FMW)

After providing these components they need additional attention - such as a configuration effort, setup of database instances and, last but not least, the application that brings another level of complexity to the party.

The main focus of this article is on the Oracle FMW components. The reason for this is the fact that in each organisation different underlying infrastructures and platforms are present. This does not however rule out an extension of the described concepts of this article towards these components.

What does provisioning of FMW components mean in particular? We need to face the reality and accept that the day to day practice depends on an organisation's choice of components. A classic J2EE application serving the needs of an individual department needs less effort and attention than a highly-available enterprise application that crosses various



network zones, deploys the SOA Suite and Service Bus and needs to be secured with certificates. Furthermore, some Oracle FMW products are not yet integrated into the provisioning mechanisms, which leads to an extended usage of script-based solutions.

However, this does not imply to abandon the idea of using EMCC12c for provisioning. The more components of a system architecture that are provisioned in an integrated way using EMCC12c the better the overall quality of the complete solution will become. Additionally, the number of outages decreases. The necessary steps for quality assurance benefit from this approach because the huge effort to carry out such tasks can be reduced or eliminated. A system which was checked OK once in a previous lifecycle requires less checks when everybody knows that it is an (almost) exact replica of its predecessor.

### Steps for Oracle components

Starting as always one needs to provide and install the Oracle software. Possibly a technical database (Metadata Repository) is required to harbour configuration data and even actual run-time data. Finally, the FMW software is configured (creation of a domain and possibly an extension of domain templates) and then the application is deployed.

### Provisioning with EM Cloud Control 12c

The provisioning process, as stated before, should be based as much as possible on available Oracle tools. Using EM Cloud Control 12c (EMCC12c) is the logical choice. In addition to provisioning, there is the integration of governance aspects (who is allowed to execute this provisioning operation?) and the subsequent monitoring of the components, which is mandatory for a seamless system administration. The new FMW targets are discovered in a post-provisioning step in the EMCC12c. This offers such goodies as an automated QA of EMCC12c standard metrics (e.g. application deployment in status failed, data source in status suspended) as an integral part of the provisioning workflow. If required, such an automated QA can be expanded individually by using "Configuration Extensions". Likewise the "Configuration Management" in EMCC12c offers afterwards a quick insight into configuration changes after the provisioning; which changes happened after completion of the provisioning workflow? (Many configuration parameter of an environment are collected by EMCC12c every 24 hours by default).

Depending on your version of OEM FMW plugin, the following components can be provisioned with the EMCC12c: WebLogic Server, Java EE applications, Coherence nodes and clusters, SOA artifacts, Oracle Service Bus resources, Oracle BPEL processes and Oracle Application Server.

This shows that the complete width of the Oracle FMW components is not yet covered (see also the official "Cloud Control Lifecycle Guide" for updates).

The complete process is embedded in the framework of the Oracle Enterprise Manager, so that the creation of a software library in the EMCC12c, granting of privileges and roles to users,

the creation of provisioning profiles and the creation/extension of WebLogic domains need to be executed in that particular order. The transition of lifecycles (DEV, TST, INT, PRD) calls for a decreased list of entitled users (larger user group or self-service in DEV <--> Super Administrator in PRD).

Application of provisioning profiles to target systems requires an OEM agent on those systems. Therefore the target systems need to be "discovered" by the OEM as "Host Targets" before provisioning takes place.

How is a provisioning of the FMW components executed? Globally this can be divided into two areas: *Profiles* and *Deployment Procedures*. Profiles describe the Oracle software (installation media), the Oracle Home and the WebLogic domain. A profile is created in the Profile Wizard and will then be stored in the Software Library. The steps depend primarily on the desired task, but typically it is a combination of the following activities:

- Creation of an Installation Media Profile
- Creation of a Middleware Home Profile
- Creation of a WebLogic Domain Profile
- Execute the Provisioning Profile

The *Deployment Procedures* describe the provisioning of the FMW domains and Oracle Homes, scaling out of SOA, Oracle Service Bus, and WebLogic Server domains, the deployment of J2EE applications up to the provisioning of BPEL processes and OSB resources. Adapted procedures (Perl, WLST or shell scripts) can be stored in Cloud Control (inside the EMCC12c Software Library) and executed from the EMCC12c Console (GUI) as well as by EMCLI (command line).

### Execution of a Provisioning Profile

Comparably to the set up without EMCC12c, a number of prerequisites need to be met before the successful provisioning can happen. For example, the target machines need to be known as targets in Cloud Control. Numerous read- and write permissions on those machines are to be set for the user executing the commands. Furthermore, the ports for the Administration Server, Managed Server and Node Manager are not to be occupied by other processes.

In an interactive session, a defined sequence of dialogs is to be executed to provide the necessary data.

The first activity is to call the FMW Provisioning Page. From here a *Profile* or a *Deployment Procedure* is initiated. This is followed by the *Middleware Target Home Page*. The targets are selected here (e.g. a WebLogic domain). This is then cloned from an existing WebLogic domain or from a Middleware home. Further dialogs ask for *Source Environment Details*, *Destination Environment Details*, *Domain Configuration Details*, passwords and a start time.

Such an interactive approach obviously suits only the preparation of a provisioning. Repeated execution of such provisionings will therefore use the "EMCLI" (Enterprise Manager

Command-Line Interface).

The following example shows the creation of a profile:

```
emcli create_fmws_domain_profile
-name="Profile Name"
-ref_target="Reference Target Name"
[-description="Profile Description"]
[-oh_cred="Oracle Home Owner
Credentials"]
[-includeOh]
[-schedule=start_time:yyyy/MM/dd HH:mm];
```

Selecting the *Profiles* or the *Deployments Procedure* as starting point requires the GUID of the selection:

```
emcli get_procedures
```

This command provides a unique global ID (GUID) of the existing *Profiles/Deployment Procedures*. These have been created manually in the EMCC12c Console. With the GUID a *Properties File Template* is created.

```
emcli describe_procedure_input -procedure=<proc_guid> -
name=<proc_name>
```

The output can then be redirected into a file and will be edited. Such a *Properties File Template* contains a large number of parameters - which can be adapted for the more specific needs and parameters that were not set by the interactive dialog.

```
CREATE_DOMAIN=true
DEST_ADMIN_HOST.0.ADMIN_SQL_HOME=
DEST_ADMIN_HOST.0.DATASOURCE_PROPERTY_FILE_DIR=
DEST_ADMIN_HOST.0.DATASOURCE_PROPERTY_FILE_NAME=
DEST_ADMIN_HOST.0.DOMAIN_HOME_DEST_ADMIN_HOST=/opt/oracle/
domains/MY_DOMAIN
DEST_ADMIN_HOST.0.DOMAIN_NAME_DEST_ADMIN_HOST=MY_DOMAIN
...
DEST_ADMIN_HOST.0.PORT_DETAILS_DEST_ADMIN_HOST=7001:Listen
Port,7002:SSL Listen Port
...
DEST_ADMIN_HOST.0.WLS_PASSWORD_DEST_ADMIN_HOST=welcome1
DEST_ADMIN_HOST.0.WLS_USERNAME_DEST_ADMIN_HOST=weblogic
...
DEST_ADMIN_HOST.0.name=as01.mydomain.com
...
DEST_FMWS_HOST.0.name=srv01.mydomain.com
...
DEST_MANAGED_SERVERS.0.NM_LISTEN_PORT=5556
```

Handing over the *Properties File Template* to the provisioning tool is achieved by the following command:

```
emcli submit_procedure -input_file=data:<input_properties_file>
-procedure=<proc_guid> -instance_name=<optnl_DP_Instance_Name>
```

This will set up the environment like any interactive session of the EMCC12c Console would do.

The above command "**emcli submit\_procedure ...**" delivers (when successful) a so-called Instance ID (IID), which can be

used to check the status of the set-up activities (comparable to the overview present in the "Procedure Activity" of the EMCC12c Console (GUI)):

```
emcli get_instance_status -instance=<IID> -details
-showJobOutput
```

The above mentioned properties-file can also serve as input for different environments. This makes transition between the different lifecycles easier. If, for example, an environment contains a larger number of hosts, the file can be simply edited and adapted by multiplying the number of blocks describing a target host.

The properties file can also be saved or updated in the EMCC12c Procedure Library:

```
emcli save_procedure_input -name=<...> -procedure=<...> -input_
file=<MYFILE>
emcli update_procedure_input -name=<...> -procedure=<...> -input_
file=<MYFILE>
```

## Deployment

Provisioning the environment (platform) is a big step forward. Now the application should also be provided automatically. Such a deployment can take different forms depending on the used FMW components. Take the example of a WebLogic environment with a J2EE application. This can be compared to the **emcli get\_procedures** approach of the provisioning. The main difference is the usage of Procedures that employ a *J2EE Application Provisioning Template*. Even in this case a parameter file is generated with **emcli describe\_procedure\_input**. The structure of this file is comparable to the one seen before.

```
deployMode=true
domains.0.continueOnDeployError=true
domains.0.domainName=MY_DOMAIN
domains.0.javaeeApps.0.appName=MY_APPLICATION
...
domains.0.javaeeApps.0.archivePath=/tmp/MY_APPLICATION.ear
domains.0.javaeeApps.0.copyComponents=true
domains.0.javaeeApps.0.copiedComponentsList.0.componentPath=sw
lib/MY_APPLICATION
...
domains.0.javaeeApps.0.copiedComponentsList.0.name=srv01.
mydomain.com
...
domains.0.javaeeApps.0.deplMode=Deploy
...
domains.0.javaeeApps.0.stageMode=DEFAULT
domains.0.javaeeApps.0.startMode=full
domains.0.javaeeApps.0.targets="MS_001_APPLICATIONS"
domains.0.javaeeApps.0.type=host
domains.0.javaeeApps.0.wlsAdminURL=t3s://as01.mydomain.com:7002
...
undeployMode=false
```

Modifying manually the properties-file happens as described above when the application should be deployed in the various lifecycle environments (more Managed Server, different ports, etc.).

## Conclusion

Employing EMCC12c for the provisioning it most certainly useful when addressing the ever-growing complexity of environments. One should keep in mind that the preparation is not trivial as it “executes” the manual provisioning steps in order to automate them. Furthermore, a number of FMW products is not yet supported. However when regarding FMW products or J2EE applications that need for example a database, the advantages become obvious. Although this article has not shown the database and schema provisioning, it is very similar to the concepts depicted above. Combining the FMW and database activities in one Provisioning Workflow will take a large burden off the shoulders of system administrators as a large number of repetitive tasks can be automated with EMCC12c. ■



## ABOUT THE AUTHORS

### Andreas Chatziantoniou

Fusion Middleware Specialist, Foxglove-IT

Andreas (Diplom-Informatiker TU Berlin, OCA, RHCT, Oracle ACE, Oracle SOA Suite specialised, Oracle Enterprise Manager specialised) is a freelance Oracle Fusion Middleware specialist with his own company Foxglove-IT. His focus is the technology area of all Oracle Fusion Middleware products (except BI), WebLogic Server, Enterprise Manager and the Oracle Database. Years of hands-on experience in literally 100s of environments ensure a backpack full of tricks to get systems running. Andreas is a frequent speaker at various Oracle events.

**Blog:** [achatzia.blogspot.com](http://achatzia.blogspot.com)



[www.linkedin.com/in/andreaschatziantoniou](http://www.linkedin.com/in/andreaschatziantoniou)



@MagicChatzi



### Ingo Reisky

Senior Consultant, OPITZ CONSULTING Deutschland GmbH

Ingo (Diplom-Wirtschaftsingenieur TH Karlsruhe, OCP, RHCT, LPIC-2) is a Senior Consultant for the Infrastructure Consulting division of OPITZ CONSULTING Deutschland GmbH. He is specialised for the Oracle Fusion Middleware, with a focus on WebLogic Server, SOA Suite and Enterprise Manager. Broad practical experience and years of administration of very complex system environments (databases, application servers, system administration of Linux, as well as large VM farms) are the foundation of his knowledge. Alongside this he works as trainer for Oracle University delivering WebLogic Server Administration courses.

**Blog:** [thecattlecrew.wordpress.com](http://thecattlecrew.wordpress.com)



[www.linkedin.com/in/ireisky](http://www.linkedin.com/in/ireisky)



@IngoReisky



# Creating the platform of the future with Oracle Fusion Middleware 12c

This article shows how Oracle Fusion Middleware 12c offers one complete, modern, open and integrated stack that allows you to create a future-proof platform. Oracle Fusion Middleware provides mobile and web access from any device based on user and customer experience best practices. It enables applications to be responsive and to create valuable insights by incorporating big and fast data capabilities; and gives you control over both structured as well as unstructured knowledge-driven processes. It also has best-in-class integration capabilities, both between the products in the Fusion Middleware stack, as well as with other applications and data either running in the Cloud or on-premise.

Antonis Antoniou  & Ronald van Luttikhuizen  eProseed

## Introduction

The advancement of technology and the changing and disruptive ways we use technology provides both opportunities and threats. In particular the adoption of social, mobile, Cloud, and big and fast data (Internet of Things), together termed the “Nexus of Forces” by Gartner, have introduced new business models. These models change the way we interact and collaborate with customers, employees and partners.

Enterprises strive to leverage this disruption to boost their customer intimacy and operating profits by optimising time-to-market and create solutions that better fit the needs of customers to deliver quicker return on investment and faster innovation cycles.

In order to support these new business models we need an IT-platform with specific characteristics to enable us in the key areas we just identified:

- The User: Easy to use and secure mobile and web-access based on user and customer experience best practices and patterns. Availability on all devices;
- The Activities: Activity-based applications that give insight and support both structured processes aimed for efficiency as well as dynamic, adaptive and knowledge-driven processes;
- The Events: Tapping into the vast amount of streaming data out there and being able to extract valuable information out if it to use and analyse in real-time;
- The Integration: Connecting the services and data from on-premise systems and Cloud-based applications, integrating both modern as well as legacy systems.

## The Case

We know how annoying it can be to lose

your luggage when you're travelling. We use this real-life example as a showcase to demonstrate how a fictitious airliner builds its platform of the future using Oracle Fusion Middleware. This platform increases customer-intimacy by detecting and responding in near real-time to customer's complaints and by making it easy for customers to track and trace the progress of their luggage. This is done through the user's channel of choice such as social media, mobile and web, and by proactively notifying customers through push-notifications. The platform also streamlines the airline's lost luggage process and provides its employees with real-time data and analytics to speed-up recovery of the luggage and to help the organisation keep track of their operational performance.

## Introducing Oracle Fusion Middleware

Oracle Fusion Middleware with its innovative platform empowers enterprises to truly transform into a

## Technology: Antonis Antoniou & Ronald van Luttikhuisen

“digital business” by blurring the digital and physical worlds, enabling customers, partners and employees to interact more efficiently and more productively, delivering tangible business value.

As can be seen in figure 1, Oracle Fusion Middleware is a vital part of the complete portfolio of hardware and software offered by Oracle, which ranges from storage and servers to applications. Besides being widely used by a vast range of enterprises, Oracle itself uses Fusion Middleware to build and deploy its applications.

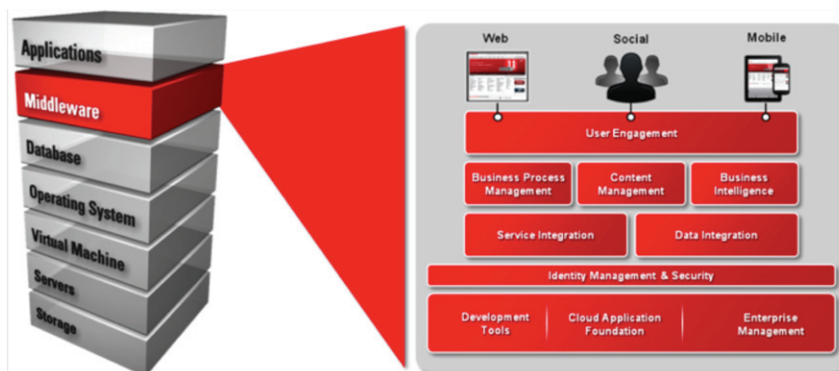


FIGURE 1: ORACLE FUSION MIDDLEWARE

This article focuses on the main Oracle Fusion Middleware 12c building blocks to support the four key areas: Mobile Application Framework (MAF) to enable “The User”, Business Process Management (BPM) and Adaptive Case Management (ACM) to orchestrate and support “The Activities”, Oracle Event Processing (OEP) to harness the information available through the Internet of Things (IoT), and SOA Suite and Data Integration to facilitate “The Integration”.

### The User

The omnipresence of tablets, smartphones and wearables has changed the way users interact with technology and what they expect from enterprises. If we look at the lost luggage case from a user’s perspective and identify the customer touch points (see Figure 2 – User Lost Luggage Activities) you can see how these have changed over the years.

Before boarding the plane you check-in online from your mobile, tablet or desktop computer. Assuming you bring carry-on luggage, you drop it off at a check-in counter or baggage drop off point at the airport. You tweet that you’re all set and possibly post a photo on Facebook of you at the airport. Once you arrived at your destination it would be really nice if you could look up the carousel number on your mobile device or even better, get a push notification with the carousel number rather than wandering around the arrivals hall. While your fellow travellers pick up their suitcases, you become more and more anxious and impatient since your suitcase isn’t there. Once you realise your luggage is lost you tweet about it (#fail), post a Facebook message and report the missing bag through the airline’s mobile app. After an apology from the airline’s representative you go to your hotel. Over the next hours

you receive push-notifications informing you that your suitcase has been located and will be delivered to your hotel the same evening. You also check the track and trace section of the mobile app that confirms your suitcase has been located.

Mobile Application Framework (MAF) is Oracle’s rapid development framework for creating mobile applications that you can use to build mobile functionalities such as described in the lost luggage example. MAF is a hybrid platform leveraging Java, HTML5 and JavaScript to provide a declarative user interface definition and integration capabilities with device services to develop single source and highly secure applications. These applications can run both on iOS and Android devices, always using the same code base. See figure 3.

We can leverage MAF’s rich component set of over 80 layout, input, selection and operation components to build the “Look Up Carousel”, “Report Lost Luggage” and “Trace Suitcase” screens of our luggage

app. And of course we can use MAF’s rich data visualisation components such as graphs, charts, maps and many other components to provide simplified and meaningful data visualisations; for example display the suitcase current location on a map. Figure 4 - MAF UI Components shows some of these UI components that are available.

All three screens require interaction with services and APIs that provide functionality and data to our app. Oracle MAF provides a declarative binding layer to streamline the connection of such services to user interfaces using POJO, REST or SOAP based interfaces. For example, the RESTful TrackSuitCase service can be exposed as data control and visually dragged and dropped on the mobile user interface to hook it up with an input form.

Oracle MAF enables you to quickly and declaratively access native device services such as phone, camera, SMS, GPS, contacts and others via the same

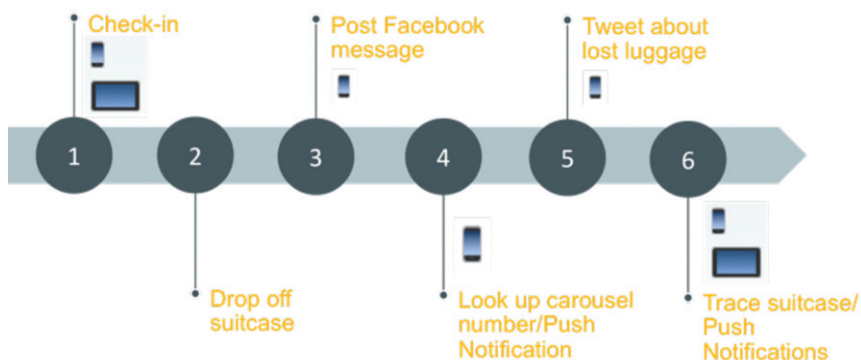


FIGURE 2: USER LOST LUGGAGE ACTIVITIES



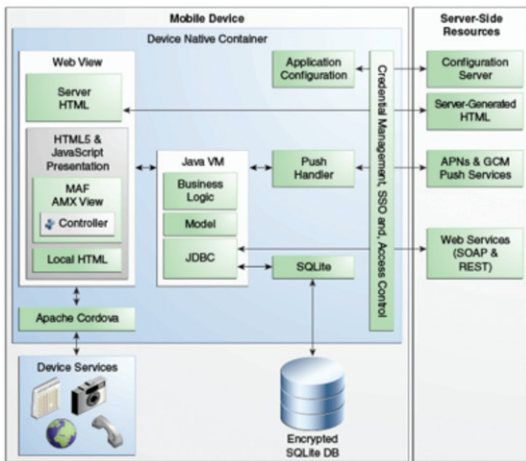


FIGURE 3: ORACLE MAF ARCHITECTURE



FIGURE 4: MAF UI COMPONENTS

declarative binding layer. In the “Report Lost Luggage” screen, instead of having the user typing in the luggage tag we use QR to scan the tag number by using the device’s camera.

The airline’s mobile app needs to be secured and support both real-time and offline data access. MAF uses a local encrypted SQLite database to enable an application to run both online as well as in offline mode. Furthermore, MAF supports authentication and access control for fine-grained security, interfacing with any OAuth supported server and enforcing encryption at all levels: database, on-device and communication with other apps.

Besides its Mobile Application Framework, Oracle also offers a rapid application development framework for developing web-based interfaces and desktop applications called ADF (Application

Development Framework). Both frameworks shares similar visual and declarative development approaches thus leveraging existing skillsets for developing ADF and MAF based applications.

**The Activities**

The transition to new business models and the changing way we use IT also impacts the way we collaborate together and the support we need for executing our activities and providing insight into them.

Traditional Business Process Management (BPM) has focused on predictable and structured processes to streamline our activities. The process path is predetermined at design-time and the goal is to maximise efficiency. Even though you can inject some flexibility in a process using for example business rules, the adaptability remains limited. While some processes or parts of processes are in fact predictable and can be optimised by traditional BPM approaches, we also need to support other types of processes. Adaptive Case Management (ACM) is used to denote flexible and knowledge-driven

business processes around the concept of a “case”. The execution path of these processes varies for different instances and is determined by knowledge workers at runtime based on the context and data surrounding the case. Business processes are often a mix of these types: some parts are highly predictable and designed for efficiency using traditional BPM, while other parts lean more to a case-oriented approach. Standard baggage handling activities fall in the first category, while handling lost luggage falls in the latter category. The way the airline handles lost luggage may depend on a wide variety of facts: length of stay at your destination, frequent flyer status, travel insurance, whether the luggage reported missing at the final destination or at a stop-over, the current location of your luggage, and so on.

Oracle BPM Suite (figure 5) is an integrated suite that supports both traditional BPM processes modelled in BPMN, as well as case-oriented processes using the Adaptive Case Management component. Both BPMN and ACM components run in the same SCA-fabric and can therefore be easily used together for mixed processes while incorporating other components such as Rules, Human Workflow and automated steps. Oracle BPM Suite provides functionality for all aspects of the process lifecycle: modeling, implementation, simulation, management, monitoring and optimisation. Moreover, Oracle BPM Suite empowers IT and business users through various business-friendly web and mobile environments to participate more effectively, providing a unified lifecycle management of BPM projects between business and IT users.

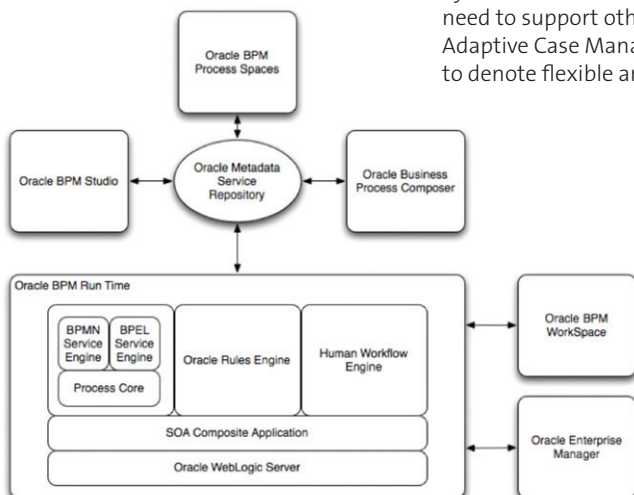


FIGURE 5: ORACLE BPM SUITE COMPONENTS



human-centric and knowledge-intensive processes that require context, relevant data and insight to execute. In ACM the available human and automated activities are modelled together with a set of constraints and milestones at design-time. At runtime, business users then decide what activities to execute and in what order; as long as this is within the boundaries defined in the case. ACM heavily uses Business Rules to capture and guide every case.

If we look again at the lost luggage process, we can see that there is a great deal of unpredictability involved. Luggage can be delayed, damaged or even lost. The airline might offer to repair the suitcase, provide the customer with a replacement or provide a reimbursement. We simply cannot model all these various scenarios using BPMN, the process would become too cluttered, error-prone and difficult to change over time. To manage this unpredictability we need the capabilities of Adaptive Case Management offered by Oracle BPM (figure 6).

The lost luggage case consists of a collection of data, processes (both dynamic and structured), human workflow, business rules and automated services. During execution the process is entirely driven by an airline investigator who is making decisions based on information that is available and linked to the case: documents, social media, structured data, geographical data, and so on. The lost luggage process consists of various checkpoints to indicate the progress of the case, referred to as “milestones”, such as “Initial Investigation Started” and “Initial Investigation Completed”. Furthermore, the lost luggage case can have various activities, like a human task activity or a BPMN process promoted as a case activity to define the ad-hoc work that can be performed in the context of a case. For example, “Reimburse Customer” and “Contact Customer”. During the execution of a case various events are raised (lifecycle, milestone, activity and custom events) that can be used to define the governance of a case; for example to withdraw the “Reimburse Luggage” activity when the luggage is found.

Business Activity Monitoring (BAM) completes the BPM-lifecycle by providing real-time process analytics for both structured, as well as unstructured, processes and thereby empowering business users with contextual insight



FIGURE 6: ORACLE BPM FOR ADAPTIVE CASE MANAGEMENT

to make better decisions. With a set of out-of-the-box dashboards and analytics and with an embedded Oracle Event Processing engine you can do complex and advanced pattern matching on data streams, time-based queries and trend analysis to measure your SLAs and get alerted if you are likely to violate your SLAs or not meet your KPIs so that you can take corrective actions on the spot. For example, to be alerted when there is an increasing number of lost suitcases for specific journey segments due to short layovers.

**The Events**

The next key ingredient of our platform of the future is that of fast data – the ability to process continuous streams of high-volume data from many devices, social media and other sources, to detect meaningful business events out of these seemingly irrelevant streams of data and to respond in near real-time to these discovered events by analysing them and passing them on to interested parties. Challenges in this arena are:

- the exponentially growing volume and throughput of data;
- time-constraints on the processing time of the data since it loses business value fast;

- variation in the data contents and formats: structured, unstructured, media, text and so on;
- the complexity of discovering relevant business value from these events; since meaning is often only discovered by combining raw data with low meaning in context: location, time, and occurrences of other events.

The sooner a relevant business event is detected, the sooner an enterprise can act on it. This provides companies with opportunities to increase customer satisfaction and customer intimacy (e.g. by informing and helping customers proactively), to increase revenue (e.g. by making targeted offers based on customer interest), and to cut costs and risk (e.g. through fraud detection).

In the case of the lost luggage example, we want to detect new complaints from customers about luggage handling based on tweets and Facebook messages. We want to continuously update customers with the latest status of their luggage by tapping into the sensor data of the baggage handling systems, check-in counters and aircraft cargo hold. These are just a few examples of fast data. In the future the airport wants to be able to filter social media and news feeds to



FIGURE 7: FAST DATA CAPABILITIES OFFERED BY ORACLE FUSION MIDDLEWARE

detect possible threats to the airport, to process sensor data to guard capacity in their bag store and to process car park events and inform customers in advance about the best place to park.

Figure 7 shows the relevant capabilities we need for fast data and how this maps to products in the Oracle Middleware stack.

One of the key products in Oracle Fusion Middleware for fast data is Oracle Event Processing (OEP). This is a lightweight event-processing engine that is able to handle millions of events per second with low-latency throughput. It can be deployed standalone, as part of your SOA stack or in an embedded Java environment running, for example, on a Raspberry Pi.

Event-processing applications in OEP consist of the following building blocks:

- **Adapters.** Components that are able to receive events from the outside world (event sources) or send events to the outside world (event sinks). OEP is bundled with various standard adapters for HTTP, JMS, CSV files besides offering an API to create your own adapters.
- **Channels.** Used to wire the various building blocks together; for example to connect an inbound adapter to a processor.
- **Processors.** Processors receive events through channels and query them using the Continuous Querying Language (CQL). This is a querying language that extends SQL and is optimised to detect relevant events in context. CQL gives you the ability to specify queries that take into account the time-period in which events occur, to easily correlate between various events, to do pattern matching, and to incorporate missing events in queries (the non-occurrence of something can be relevant as well). The results of those queries - events with more business value - are passed upstream to the application through other channels.
- **Beans.** Components that contain Java code for processing, receiving events or sending events. These beans can be Spring beans or OEP-specific event beans.
- **UI** to graphically design event-processing applications and flows.

In our lost luggage case, we create an adapter to receive tweets through

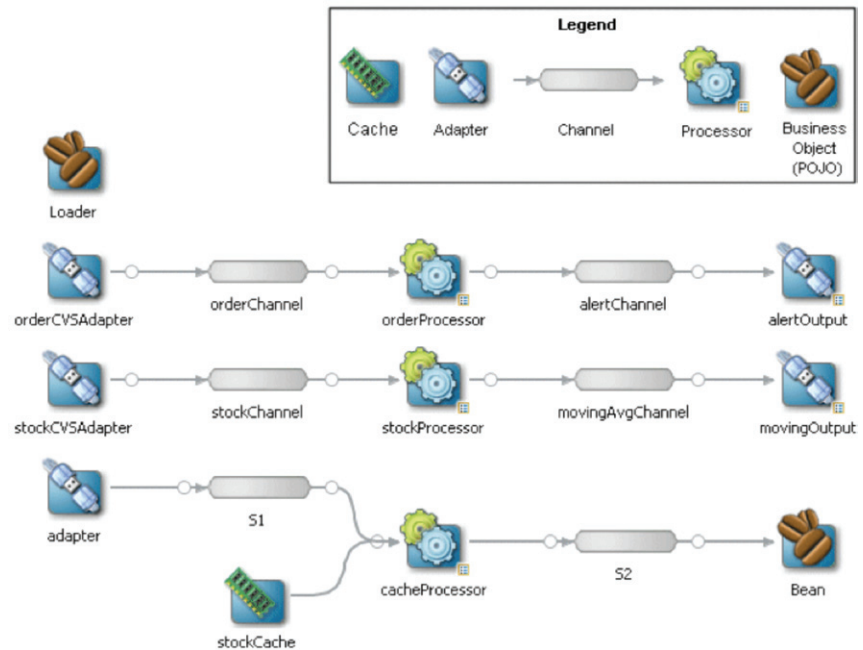


FIGURE 8: EXAMPLE OF AN EVENT-PROCESSING APPLICATION IN THE OEP DESIGN-TIME

available Twitter APIs. We can then use beans and processors using CQL to detect tweets with certain keywords, hashtags and user references that indicate complaints about lost luggage. Finally, the filtered tweets are used to publish a business event on the Event-Delivery Network (EDN) indicating a new complaint. This event is picked up by a service running on SOA Suite which starts a new lost luggage case. We can have the same business event fed into Oracle Business Activity Monitoring for analytics using JMS and reply to the original tweet that we have noticed the complaint and try our best to resolve it.

Besides OEP, Oracle offers various other big and fast data solutions in its Fusion Middleware platform. To name a few: Oracle Data Integrator and Golden Gate for both real-time data replication and bulk data loading, various business analytics and business intelligence solutions in Oracle's Business Analytics, and Oracle Coherence that provides a scalable, robust and highly-available in-memory data grid. As we will see next, all these various building blocks are designed to work together (integrated), as well as being easy to use from other platforms and applications (open).

### The Integration

We have various integration needs in our lost luggage scenario to connect all

the components together in a coherent solution:

- **Mobile integration:** the mobile apps from the airline and airport need to be tied to the backend applications that provide the actual data and functionality;
- **Integration between on-premise and Cloud:** various applications, such as baggage handling systems, CRM systems, social media apps such as Facebook and Twitter, and self-service portals and so on are involved when handling lost luggage; these systems can run either on-premise as well as offered as SaaS in the Cloud and these systems can be legacy, as well as modern applications that offer APIs;
- **B2B integration:** the airline needs to integrate with its insurance company through a B2B channel;
- **Data integration:** data on passengers, their travel plans and their baggage is stored in several databases and needs to be replicated, some of the data is provided as bulk files by legacy systems and needs to be imported into a database;
- **Internet of Things:** as we saw from the previous section we need to integrate with our event-processing system to receive business events and pass them on to interested parties such as our business processes and mobile apps.

Oracle's key products in the integration




FIGURE 9: ORACLE'S OFFERING FOR SERVICE- AND DATA INTEGRATION

arena are SOA Suite, including Service Bus for service-oriented integration, and Oracle Data Integrator, including GoldenGate for data-oriented integration.

### Mobile integration

The de facto standard for accessing backend systems from mobile applications is through RESTful services using JSON payloads. Oracle SOA Suite natively supports REST/JSON through its REST adapter: both by exposing APIs and services through REST/JSON interfaces as well as being able to consume REST/JSON interfaces. XML payloads and SOAP Web Services are of course also supported.

Oracle WebLogic Server, which is used to run SOA Suite as well as to run Java applications, implements the Java EE standards for REST such as JAX-RS.

Note that Oracle Database 12c natively supports JSON besides relational data and XML. This is particular useful in a document-centric approach where we want to store and retrieve documents instead of highly-structured relational data. Almost all of Oracle's database features such as indexes, SQL constructs, high-availability, etc. can be used regardless of the type of data stored: relational, XML, JSON, text, etc.

### Connectivity

Oracle SOA Suite ships with out-of-the-box adapters for popular SaaS applications such as Salesforce.com, as well as a Cloud Adapter SDK to develop your own Cloud adapters. These Cloud

adapters hide the complexity and intricacies of the APIs of such SaaS applications.

Besides these Cloud adapters, SOA Suite also offers various technology and application adapters for integration with on-premise applications and legacy systems. To name just a few of these adapters: JMS, EJB, File, FTP, SAP, JD Edwards, and E-Business Suite Adapter.

Oracle B2B supports B2B integration for specific industry standards such as HL7, UN/EDIFACT, and ebXML. SOA Suite and B2B are easily integrated through a dedicated B2B adapter or messaging standards such as JMS.

For data integration between various data sources, Oracle provides Oracle Data Integrator (ODI) which operates based on an ELT (Extract, Load and Transform) pattern. ODI has extensive transformation capabilities so it can transform between a huge variety of different data sources and formats such as various databases, flat files, XML files, etc. It also ships with a Web Service API to easily integrate it with services and processes through SOA Suite. Oracle Golden Gate is the product to use for (real-time) data replication. It is non-intrusive and uses log-based data capture, possibly transforming data before replicating it into the target database(s).

### Internet of Things

The business events generated by our fast data platform need to be fed to all consumers that are interested in these

business events: the ACM process for lost luggage, social media to alert customers, the mobile apps and so on. Fusion Middleware offers, among others, the following out-of-the-box functionality for integrating our event platform and consumers:

- OEP can publish business events on the Event-Delivery Network (EDN), a declarative eventing framework, on which all SOA and BPM composites can subscribe;
- SOA Suite has a dedicated adapter to integrate with Coherence in-memory grids to which data and events are published by OEP.

Besides covering the required integration needs we just discussed, cross-product integration between the components in the stack itself is also very important. Compared to stacks based on heterogeneous best-of-breed products, the Fusion Middleware stack is tightly integrated and designed to work together. Just a few examples to showcase this are the out-of-the-box Coherence adapter in SOA Suite, optimised data sources and connection pools in WebLogic Server for Oracle Databases and RAC, declaratively publishing events from SOA Suite to Oracle BAM, integrated monitoring and administration through Enterprise Manager providing insight into message flows spanning various products, etc.



### Conclusion

We have seen how technology is changing the way we live and changing the expectations and demands that we as consumers, employees and partners have regarding enterprises and each other. Big drivers for these changes are mobile-enablement, constant connectivity, the Internet of Things and the need to easily and quickly integrate various sources of data and functionality in real-time. We also know that these expectations and demands will keep changing in the future and very likely at a rapid pace.

To support this changing way of living and working we need a platform that is built for disruptive change rather than a platform that breaks down from it. A platform that provides us with all necessary modern building blocks to create the business solutions we need.

This article has shown that Oracle Fusion Middleware 12c provides us with all the ingredients for creating such a platform as we have seen from its mobile, process, fast data and integration capabilities. Oracle supports several deployment models for the platform including on-premise, private Clouds using engineered systems and appliances, and IaaS, PaaS, and SaaS offerings so you can choose the deployment that best suits your needs.

### In summary:

- The technology we need is here today. Mobile Application Framework, Fast data solutions using OEP and Coherence, process and case management using BPM and ACM, service and data integration through SOA Suite, Data Integrator, and Golden Gate are all part of Oracle Fusion Middleware 12c.
- General and industry-specific trends such as mobile-first customer experience, cross-channel experience, personalisation, insight in performance, and even further integration with social media and a higher-degree of self-service are easily implemented using Fusion Middleware.
- Fusion Middleware is both best-of-breed through acquisitions and home-grown products as well as integrated into one platform. This integration is notable through several factors including a consistent and recognisable user experience over the various components of the platform, by offering cross-platform monitoring and tracking and tracing, and coherent security features for the whole stack.
- Business focus and a user-centred design are key to success when building your platform of the future. Oracle and its ecosystem offer a rich set of patterns and prebuilt templates for this purpose. ■



### ABOUT THE AUTHORS

#### Antonis Antoniou

Technical Director, eProseed

Antonis is a Fusion Middleware Expert and a specialist and certified professional in the areas of Enterprise 2.0, BPM and SOA. Antonis has long-time experience as developer, coach, trainer and architect and has led multiple complex projects on Oracle Fusion Middleware worldwide. Antonis is an Oracle ACE Associate for Middleware & SOA.

**Blog:** <http://antonis-antoniou.blogspot.com>

 <https://cy.linkedin.com/in/antonisantoniou>

 @AntoniuAntonis



#### Ronald van Luttikhuizen

Managing Partner, eProseed NL

Ronald has experience in IT in various roles such as coach, architect, lead designer/developer, analyst and team lead. Ronald focuses on information and solution architectures in SOA and BPM environments. Ronald is a frequent blogger, author and presenter at (inter)national conferences. Ronald is an Oracle ACE Director for Middleware & SOA.

**Blog:** <http://blog.vennster.nl>

 <https://nl.linkedin.com/in/soaarchitecture>

 @rluttikhuizen

# Oracle APEX in the CERN Java Cloud

Usage of Oracle APEX at CERN is continuously evolving. CERN first started using the tool in 2005 with the first installations of the HTML DB. Since then, it has been used in all the areas of the organisation, from administrative applications to accelerators control systems. In the past few years we have accomplished two big milestones: the development of an application that centrally manages the Database Access Descriptors and most recently the integration of the APEX applications into the CERN SSO. The challenge now is to use Oracle APEX to move from a centrally managed infrastructure to a true 'Platform as a Service' in which the user becomes the owner of the service.

Luis Rodríguez Fernández, CERN

## Introduction

The purpose of this article is not to be an installation guide or a "how-to" note. Rather, its aim is to show how the CERN Oracle APEX platform has evolved from a simple setup to a cloud model. You might identify your own installation and decide to move towards a more "cloudish" platform. If this is the case, be sure to keep reading!

## What? Is this another Databases On Demand Service?

Short answer: no. This system lets the APEX developers fully manage the websites of their applications. We can do this thanks to the integration of the Oracle Rest Data Services web application (ords.war) with our CERN Middleware On Demand Java PaaS. Basically, the approach of our solution is that our Java PaaS will bind a set of APEX applications with one ords.war deployment. In this way each group of applications will be independent from the other ones. But before entering into details let's take a look at how the Oracle APEX service has evolved at CERN.

## The evolution of the Oracle APEX service at CERN

Before the first version of APEX (formerly Oracle HTML DB) was released, we were already using the Oracle PL/SQL Gateway. Developers were very excited by the ease with which they could build web applications. Also, the administrators were delighted with the fact that they did not have anything to deploy: the html code was cooked in the database and served fresh to the client! Oracle APEX was also adopted early on by the CERN Web developer community. The reasons behind this were not only its easy accessibility and agility, but also the fact that most of our data was, and still is, stored in Oracle databases.

Our first APEX service initially consisted of Oracle HTTP Server with the mod\_plsql. 99% of the administrators reading this article are probably familiar with this kind of setup. If you still have architecture like this one, be aware that mod\_plsql is no longer actively developed and it is not supported by some new Oracle products, such as E-Business Suite R12. If you start an installation from scratch, my recommendation is to use the Oracle REST Data Services, formerly known as APEX listener.



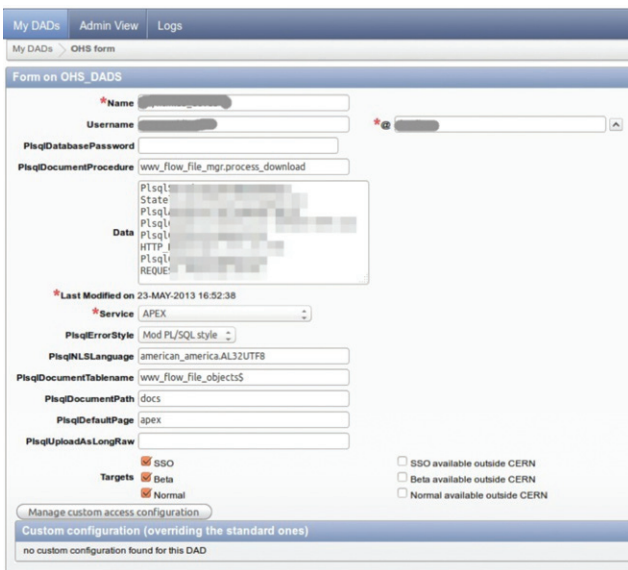


FIGURE 1: DADEDIT. DAD FORM

As usual, an increase in the number of applications and developers makes the management of the platform more complex. One sensitive and critical point is the configuration of the Database Access Descriptors (DADs). The challenge faced was how to automatically configure each user's DADs. For this purpose we developed the DadEdit tool. This application links the database schemas with the different DADs, so the owner can update the DAD information online.

The DadEdit tool takes all the information for each DAD and generates the configuration file for the mod\_plsql, the dads.conf. This procedure, which is executed every 60 minutes, compares the current dads.conf with the new one. If there are changes, it applies the new one and gracefully restarts the OHS, making the process transparent to the users.

Once the issue relating to the management of DADs was solved, we faced another challenge: integration with our SSO (luckily for me, SSO was the subject of my first project at CERN). This work involved the applications deployed in both Oracle WebLogic and the APEX ones. Fortunately, we were running Oracle WebLogic 10.3 at the time, and that version included a brand new SAML2 module! Additionally, the CERN SSO implementation for the Identity Provider was Microsoft ADFS2, which was SAML2 compatible (n.b. SAML stands for Security Assertion Markup Language and is an XML-based protocol for exchanging authentication and authorisation information between security domains).

But what about the APEX applications? No need to panic on that front either; just put Oracle WebLogic Server in front of OHS. WebLogic will authenticate the user requests and proxy them to OHS. It's not necessary to reinvent the wheel, as WebLogic also provides the HttpProxyServlet for proxying the user requests. Also, APEX offers the HTTP Header Variable among its scheme types. This authentication schema looks for the HTTP REMOTE\_USER header that contains the username. At this stage, there was only one question remaining: how to extract the username from the original authentication request and add it to the proxy one? For this we had to get our hands dirty and write a couple of Java libraries (see <https://github.com/cerndb/wls-cern-ss0>).

The purpose of these libraries is not only to add the username to the request but also fill some gaps in the WebLogic implementation, such as the single logout. The next diagram summarises the final result:




FIGURE 2: THE SSO INTEGRATION

### Why a cloud? One APEX platform is not enough to implement all scenarios?

No, unfortunately not. Soon we realised that only one APEX installation was not enough to fulfil all the requirements. We were especially concerned about the next two issues:

1. We are using WebLogic both as a proxy and for the integration with our SSO Identity Provider. This means that all the APEX applications must be SSO protected. But what happens with the ones that do not require SSO integration?
2. How to lockdown the URLs of specific applications such as the APEX development environment (f?p=4550)?

We solved the first one adding a little bit of extra configuration in the OHS: the apex-ss0.cern.ch domain serves the SSO protected applications and apex.cern.ch serves the rest.

We faced the second one in 2013. CERN opened its doors to the general public for a full weekend. To control the visits to the experiments, a ticketing system was developed in APEX. Users were able to authenticate using an external account, such as Google, Facebook or Microsoft, as well as others. The challenge here was to prevent users from changing the value of the application id parameter (p) of the URL and thus gaining access to other applications. Our solution was to set up a parallel infrastructure using the Oracle REST Data Services, aka APEX Listener, deployed in Oracle WebLogic Server. One of the most useful functionalities offered by this component is the possibility of having a request validation function (apex.security.requestValidationFunction). This enables the binding of some URL parameters, like the application id with the parameters of one PL/SQL function. Through this, we are able to allow or forbid access to any application in of our database. You can find more information about this in a very useful post by Kris Rice, APEX Listener PL/SQL Validations (<http://krisrice.blogspot.fr/2013/01/apex-listener-plsql-validations.html>).

So the need for an APEX cloud became clear:

- We realised that manually setting up ad hoc platforms wouldn't be feasible over the long term.
- We had to provide a tool where the developers independently managed the access (authentication and authorisation) to their applications.

### Mw On Demand and Oracle REST Data Services: the path towards the APEX Cloud

All the ingredients are ready for the cloud: a system generating the data access descriptors, integration with the SSO and an application that can be easily deployed and configured in Java containers, namely the Oracle REST Data Services. Now it is time for the final stage in the process. This will be the duty of our Platform as a

Service for Java Applications, the Middleware On Demand (MWOD).

The objective of MWOD is to fulfil all the requirements that a Java web application has for running. Among other functionalities, it handles (a) registration of the URL in the DNS, (b) creation of the virtual host in the web server, (c) integration with the SSO, (d) creation of the container and (e) deployment of the application.

For each different application server version there is one template that contains the metadata (ports, certificates, connectors, etc.) and some custom scripts and packages, such as the monitor agent. During the creation stage, the custom configuration values are filled, the agent is deployed and symbolic links to the original container binaries are created. Similarly, there is a template for APEX that contains the Oracle REST Data Services web application (ORDS). What the MWOD has to do is just set the path of the configuration directory. This value is set via a configuration parameter (config.dir) in the web.xml deployment descriptor of the ORDS application.

The most important file inside the configuration directory is the apex-config.xml. This archive is the equivalent of the dads.conf in the mod\_plsql and is automatically generated by the DadEdit tool.

The procedure is now close to an end. There is only one final important topic to review: the security.

- Note that the default settings allow any user with an account in the database to execute the f procedure, which is the entry point of any APEX application. In our installations only the APEX\_PUBLIC\_USER has the right to execute this procedure, because applications using the HTTP Header Variable authentication schema would otherwise be vulnerable. Any user can set up his or her own ORDS installation and inject any REMOTE\_USER header.
- Each new installation should only access certain permitted applications. For this purpose, we developed the CERN APEX Management system, through which the APEX developer is able to declare which applications should be “deployed” in the cloud. Once validated, the application invokes the MWOD web service that triggers the creation of the site and generates the code for the apex.security.requestValidationFunction.

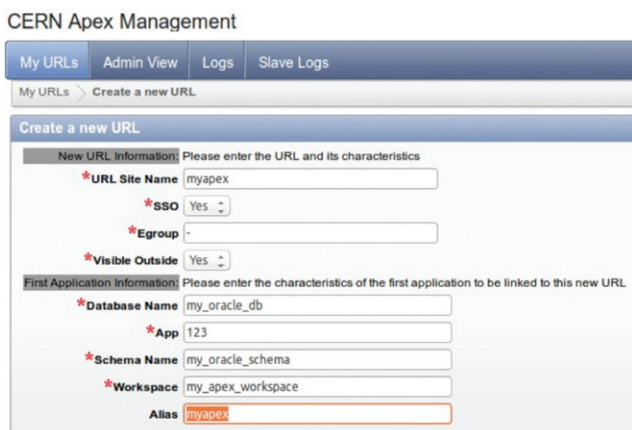


FIGURE 3: THE CERN APEX MANAGEMENT APPLICATION

A little bit lost? No worries, the diagram below gives a clear overview:

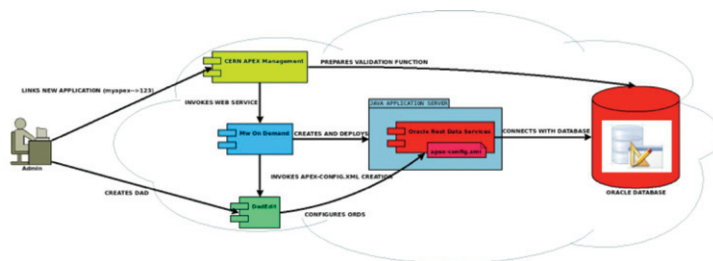


FIGURE 4: ORACLE APEX IN THE CERN JAVA CLOUD: THE BIG PICTURE

Finally the users access the application through the new URL created in the DNS by our cloud:




FIGURE 5: ORACLE APEX IN THE CERN JAVA CLOUD: USER VIEW

### Wrap up

This article has taken you through all the steps to build an Oracle APEX infrastructure in the cloud: (a) an application for managing the database configurations, (b) a system to lockdown the applications URLs, (c) SSO integration, (d) a Java Platform as a Service and (e) the Oracle Rest Data Services. As you can see, we have not invented anything new. We have just deployed a standard Oracle product, the ORDS, in a Java PaaS, automated the configuration and hardened the security.



## ABOUT THE AUTHOR

Luis Rodríguez Fernández  
Software Engineer  
CERN

Luis Rodríguez Fernández is a Java developer and middleware specialist with over 10 years' experience in building portals, web applications and systems integrations. He loves to find simple solutions to complex problems. Luis has been a speaker at UKOUG and JavaOne. He currently works for CERN, in Switzerland, and is a member of CERN openlab.



# Implementing Access Manager the Easy Way

Oracle Identity Management 10g (IdM 10g) has been widely used in the Oracle community in combination with Oracle Application Server 10g deployed applications such as Oracle Forms applications. When migrating Application Server middle-tiers to Fusion Middleware 11g it is still possible to continue using the 10g identity management services. The complexity of the new products Access Manager and Identity Manager 11g may often have been a reason to stay with the reliable 10g stack. Support on IdM 10g, however, has stopped some time ago, and it is high time to migrate Identity Management as well. In this article we describe how to do this for a Forms-Reports environment without drowning in the details of all the functionality you are not going to use (initially or ever).

Peter de Vaal & Nicolay Moot, Transfer Solutions

## Differences between IdM 10g and Identity & Access Management 11g

Oracle Identity Management up to version 10g consisted of a directory server, Oracle Internet Directory (OID), a light weight Single Sign On server (OSSO), and a handy Self Service Console (OIDDAS) for managing users and groups in the directory. Also included was a synchronisation and provisioning tool (DIP) to synchronise the directory with other directories (such as MS Active Directory) or other data sources. All these tools provided an authentication service for applications.

In 11g the Identity and Access Management suite contains (among several other products) Access Manager, Identity Manager and OID. Access Manager 11g takes care of the Single Sign On functionality, Identity Manager 11g is the intended tool for managing identities. They replace OSSO and OIDDAS which are not included in the 11g stack.

Access Manager and Identity Manager add a huge amount of extra functionality compared to IdM 10g, but most likely your intention is not to make use of that when you are merely upgrading your technology stack to stay supported. Moreover, the products are very expensive while IdM 10g, as you used it, was part of the Application Server EE licence. The latter, however, should not keep you from migrating to 11g, because a limited license of Access Manager is still included in the application

server licence, and even very limited use of Identity Manager, for the purpose of user and password policy management, is included.

## Why do we need to migrate to 11g?

Of course a valid reason to migrate is that IdM 10g is out of support. However, as long as you do not replace the server running it and there is no need to upgrade the Operating System, the risk of a failure is low, and it is unlikely that you need a patch or other support for a system that may be running without change for many years. Security patches might be an absolute requirement in a large organisation or when the system is exposed to the internet, but for a small scale intranet deployment these might be skipped.

Waiting too long with a migration, however, may impact other future changes. We have had a case where IdM 10g was integrated with Windows native authentication (WNA). Another organisation was responsible for the Microsoft environment, and decided to upgrade Active Directory (AD) to the 2008 release. This broke the WNA integration because SSO 10g did not support the stronger cipher suites required by the new AD version. We had to remove the Single Sign On functionality from the applications and could implement it again only after migration to Access Manager 11g.

## Impact of migration to Identity & Access Management 11g

Three aspects exist that are important to consider when determining the impact of the migration from IdM 10g to Identity & Access 11g:

- Complexity of installation
- Increased requirements for hardware resources
- Desired functionality

### Installation complexity

Although it was possible in IdM 10g to set up a distributed, high available architecture that needed expert knowledge to set up and manage, it was also possible to perform an almost black box installation with all components included, even the repository database. Such an installation was often performed for an intranet deployment, where only a logon service for Oracle Portal or for Forms or ADF applications was needed.

To obtain the same functionality in an 11g architecture you have to install, configure and integrate many components as we will discuss in this paper. Note that we will restrict the discussion to a non-high available, non-SSL and non-distributed architecture, the one you installed in 10g just with a few mouse clicks.

### Hardware resources

First of all re-using the server which ran the 10g IdM is hardly an option. The 10g software ran exclusively 32-bit, while in 11g you would certainly opt for 64-bit. The 10g stack, including the repository database, could run on a single core server with 500 MB memory. You will now require at least 8 GB of memory.

### Desired functionality

If you really only need a replacement for the single box 10g IdM then you can do the following:

- Do not install Identity Manager  
You don't need OIM if you can manage users in another way, e.g. from an application
- Use stand-alone OID  
If you do not need to manage users with ODSM. This saves a complete WebLogic domain

### Steps to implement the minimum required architecture

As you are likely going to deploy the new components on a 64-bit server but the 10g stack is still running on a 32-bit machine a simple upgrade using upgrade assistants is not possible, because such tools only work if you do an in-place upgrade on the same machine.

The steps to set-up the new environment will include those described in the following paragraphs. This is only an outline, detailed steps (e.g. all installation screens, how to start and stop WebLogic domains and Oracle instances) can be found in the regular product documentation.

### Install a supported JDK

Use the latest update of JDK 7.

### Install WebLogic version 10.3.6 software on the new machine

In case you are going to use Access Manager 11g R2 (11.1.2.2) and you want to manage OID from a WebLogic domain using Enterprise Manager and/or Oracle Directory Service Manager, you need to do this twice, creating 2 different Middleware Homes, one for OID software, the other for Access Manager software. You must have the different Middleware Homes because the versions of the Fusion Middleware Common Components of the OID (11.1.1.7) and OAM (11.1.2.2) installations are not compatible.

At the end of the installation do not run the quick start to create a WebLogic domain.

### Install Identity Management 11g (11.1.1.7) software

This is an operating system specific distribution. If you installed a 64-bit JDK then you must use the 64-bit version of this distribution. You must install this software in the Middleware Home, which is created by the (first) WebLogic software installation, unless you want to install a stand-alone OID, in which case you can specify to install in a new Middleware Home.

Use the Install software only, do not configure mode of the installer. After installation download and apply the latest patch bundle for OID, e.g. OID 11.1.1.7.5, using the OPatch utility.

### Install Identity & Access Management 11g R2 (11.1.2.2) software

This software distribution is generic for all platforms. Install it in the second Middleware Home. After installation download and apply the latest patchbundle (11.1.2.2.3) for Oracle Access Manager. See MOS note 736372.1 for details.

### Install the Repository Creation Utility (RCU) software

You must download and install the proper versions of the utility. The OID schemas will be created using the RCU 11.1.1.7. For the OAM schemas you need version 11.1.2.2.

### Create the metadata repository

We suppose you already have an Oracle database available for the metadata repository. If not, create one first on the same or a different machine. See the FMW 11g certification matrix for supported versions.

Run RCU 11.1.1.7 according to the installation guide to create the OID schemas.

Run RCU 11.1.2.2 to create the schemas for Metadata Services (MDS), Audit Services (IAU), Oracle Platform Security (OPSS) and, of course, Oracle Access Manager (OAM).

### Configure the OID domain or instance

Start the Identity Management configuration wizard (config.sh or config.bat) from the <Oracle Home>/bin directory.

In case you have chosen to install OID in a WebLogic Middleware Home, you will be asked to select a domain. Choose Create a New Domain and specify a domain name. Note that you cannot specify the location of the domain directory. The domain directory will be created in <Middleware Home>/user\_projects/domain.

Select the products to configure: OID and optionally DIP. In case you have specified a WebLogic domain ODSM and Fusion Middleware Control will automatically be selected. Do not select OVD (Oracle Virtual Directory) nor Identity Federation components.

**Migrate the relevant OID content**

Although a migration assistant exists that can be used to upgrade an existing OID 10g instance to 11g including the database schemas, this can only be used when the 10g and 11g installations are on the same machine. This is unlucky enough not a realistic situation as pointed out earlier. So skip the frustrating part of getting the upgrade done as described by the Identity Management Upgrade Guide and use LDAP tools to migrate the OID content from 10g to 11g.

If you have configured it you can use the Directory Service Manager (ODSM) to export the users and groups from OID 10g, and import them in OID 11g. Otherwise you need to use the OID commandline tools to do the Idif export and import. You also have to migrate the Resource Access Descriptor (RAD) entries as well. This is the trickiest thing to do, but you may already have a script to create them from scratch in the new OID.

**Configure the Access Manager domain**

Start the Domain Configuration wizard config.sh/cmd. Only configure the components: Oracle Access Management, Oracle Enterprise Manager, Oracle Platform Security Service, Oracle JRF. Finish the configuration as described in the Installation Guide for Oracle Identity and Access Management 11.1.2.2. See chapter Configuring Oracle Access Management for details. Do not forget to read the sections Upgrade OPSS schema Using Patch Set Assistant and Configuring Database Security Store, and follow the steps described there.

**Associate Access Manager and OID**

Access Manager is not associated to a directory service out of the box, so we have to connect it to the OID. We use the OAM Console for this, which is accessed from: <oamhost>:7001/oamconsole. Logon to the console as weblogic. The first page is a launchpad from which you can access all the configuration pages, see Figure 1. The amount of options is quite intimidating, so we will keep the number of visited pages to a minimum here. The link to the User Identity Stores page is found in the section Configuration. When you open this page you will find a number of sections.

We are first going to add OID as an identity store. Define OID as an IDS Profile. Specify all the details of the OID and apply the data. You can also test the OID connection from this page.

Now we set the OID profile as the default store by selecting it in the corresponding poplist in the top section of the User Identity Stores page. Leave the system store as it is, it is the identity store for system management (such as we are doing now) and it uses the embedded LDAP of WebLogic by default (that is why we can logon as weblogic).

**Associate Forms and Reports applications with OID 11g**

As your Forms and Reports applications are currently associated with OID 10g you have to change this to the 11g OID. This is done using Fusion Middleware Control of the Forms-Reports server (Forms menu Associate/Disassociate OID, Reports menu Reports OID Association).

After you have specified the details for the 11g OID you will be asked which Access Control method you want to use: Access Manager or Oracle Single Sign On. Select Access Manager and specify the details for this server. Note that such details need

not specified if you selected SSO 10g. The reason is that the SSO 10g details were stored in OID during the SSO installation, but the Access Manager details are not available in that way.

**Register Forms and Reports applications with Access Manager**

Access Manager uses so called access clients to protect applications with single sign on and access policies. The access client is deployed on the web server hosting the application, typically the Oracle HTTP Server that front-ends the Oracle Forms or Reports applications. Access Manager has an own access client called WebGate, but can also use the OSSO 10g client, i.e. the mod\_osso module of Oracle HTTP Server. Because we have the latter already configured on the Forms servers we will continue to use

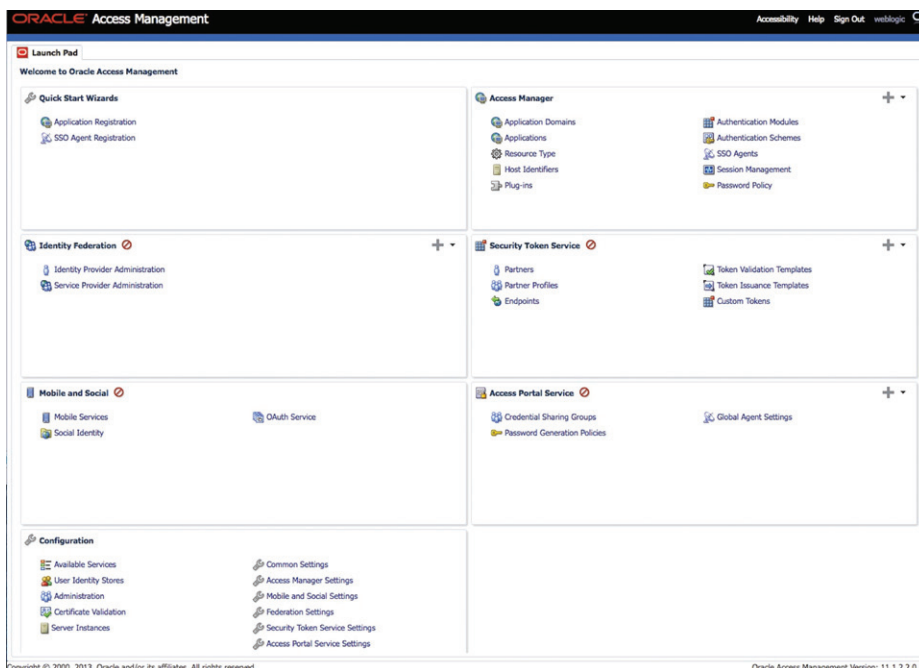


FIGURE 1

this OSSO agent. Register each Forms service by selecting SSO Agent Registration from the Access Manager Console launchpad. Choose OSSO-agent as the type, specify a comprehensive name for the agent and specify the base URL to your Forms-Reports server, e.g. <http://formsprod.mycompany.com:8888>. This should be the URL that a user of the application will use. If you have a load-balanced Forms application then this URL will point to the load-balancer, and you have to register only one SSO Agent.

After you finish each registration you will find 3 new objects in the Access Manager configuration: The SSO Agent definition, a host identifier, and an Application Domain, all three with the name you specified for the SSO Agent. Don't worry about what all these things are, they only have meaning if Access Manager is used for anything else then Forms-Reports.

The SSO Agent registration also produced a file `osso.conf` on the Access Manager server. It can be found in the output/<SSO Agent name> subdirectory of the domain directory. This file should replace the original `osso.conf` on the Forms-Reports server (in the HTTP Server configuration directory) that was produced during the registration of the 10g Single Sign On server.

After restart of the HTTP Server and the Form and Reports managed servers the (SSO enabled) applications will now be protected with Access Manager instead of OSSO 10g.

### Implementing password policies with Access Manager 11g R2

In Identity Management 10g it was possible to define password policies, such as expiration rules or special requirements for password syntax, in OID. SSO 10g then used these rules and would present a "Change Password" screen when a password expired. In OID 11g it is still possible to define such rules, but Access Manager 11g will ignore them. If you are not aware of this then users will not be able to logon after the expiration date.

If you want to have this functionality back then you may disable it in OID and implement it in Access Manager 11g release 2 (11g release 1 did not have this functionality). To set up password policies in Access Manager 11gR2 read the MOS note 1533134.1.

### A final word

Maybe the word 'easy' is not the best choice for this installation after all. The recipe we have described here, however, makes it possible to configure the 11g identity management environment within an acceptable amount of time, with all the features needed for your Forms and Reports applications.



## ABOUT THE AUTHORS

### Peter de Vaal

Principal Consultant, Transfer Solutions

Peter de Vaal is Principal Consultant at Transfer Solutions, an Oracle Technology partner based in Leerdam, the Netherlands. With 20+ years' experience in Oracle technologies he now focuses mainly on Fusion Middleware. He is lead of the Fusion Middleware Cloud Application Foundation SIG of Oracle User group Holland (OGH). He is a frequent presenter at conferences such as UKOUG TECH and ODTUG.



### Nicolay Moot

Middleware & Database Administrator, Transfer Solutions

Nicolay Moot is Middleware and Database Administrator at Transfer Solutions. He joined Transfer after other administration jobs in IT and picked up experience in Oracle technology quickly. His daily work is Remote Administration, Monitoring and Installation of customer's Oracle environments. Recently he worked on various projects involving Identity Management components OID, OIM and OAM.



# Oracle SQL Developer... for the DBA?

Short Answer: YES!

Jeff Smith, Product Manager, Oracle

## Introduction

Oracle SQL Developer is the graphical user interface (GUI) for Oracle Database developers, administrators, and IT support staff. With more than 4 million active users, it is very well-known for making it easier to execute queries, develop stored procedures and design database structures.



FIGURE 1: SQL DEVELOPER DEBUTED AS PROJECT RAPTOR IN 2005.

When I give my well-rehearsed 30 second elevator speech, I always describe SQL Developer as the IDE for SQL & PL/SQL. Or I'll describe it as a graphical user interface alternative to SQL\*Plus.

What sometimes gets left out or overlooked is that SQL Developer also hosts an entire suite of features for the database administrator. Yes, just because it has the word 'Developer' in its product name, doesn't mean we can't help make the DBA's life better too!

I recently put in a request to change the product name to Oracle SQL Developer/DBA, but I don't think that's going very far.

Author's Note: I was joking about changing the product name.

So now that it has been established that SQL Developer is for DBAs too, the rest of this article will highlight the DBA features available in version 4.0 of SQL Developer.

## Some Historical Perspective

SQL Developer was first released ten years ago. For the first few

years, the PL/SQL developer and general Oracle user's workflow and perspective were definitely the primary focus.

## Release Dates

- Project Raptor – 2005
- Version 1.1 – 2006
- Version 1.2 – 2007
- Version 1.5 - 2008
- Version 2.1 – 2009
- Version 3.0 – 2011 ← The DBA Panel is introduced
- Version 3.1 – 2012
- Version 4.0 – 2013

To be fair, we weren't completely ignoring the administrator or power user prior to 2011. SQL Developer has always shipped with some reports that assist with managing sessions, resources and users. And if you right-click on a connection, you can get a brief overview of storage and memory allocation.

TABLESPACE_NAME	PERCENT_USED	PCT_USED	ALLOCATED	USED	FREE	DATAFILES
1 SYSTEM	99.94	820	819.5	0.5	1	
2 USERS	94.97	232.5	220.81	11.69	1	
3 SYSAUX	90.18	1210	1091.13	118.88	1	
4 EXAMPLE	89.73	345.625	310.13	35.5	1	
5 APEX_1334411028004709	86.78	7.5625	6.56	1	1	
6 PART1	42.04	862	362.38	499.63	3	
7 PART2	7	25	1.75	23.25	1	
8 UNDOTBS1	6.63	260	17.25	242.75	1	
9 PART3	6	25	1.5	23.5	1	
10 PART4	4	25	1	24	1	
11 TEMP	(null)	(null)	(null)	0	(null)	

FIGURE 2: MANAGE DATABASE SCREEN, AVAILABLE OFF A CONNECTION CONTEXT MENU

This feature, as well as all of the reports, is 'read only.' That is, the DBA could see what was happening, but not directly interact

with the database and do things like unlock a user or resize a tablespace. The one exception to this would be the 'Sessions' report, which we also have under the Tools menu as 'Monitor Sessions.' A user with the appropriate privileges can kill or trace sessions from that interface.

### Introducing the DBA Panel

The name says it all – a collection of features designed with the DBA in mind. We also assume any user accessing this panel has the DBA role or equivalent set of privileges. Trying to access a feature without these privileges will raise the appropriate error messages. These features were first introduced in version 3.0, but a few of the ones discussed below require at least version 4.0.

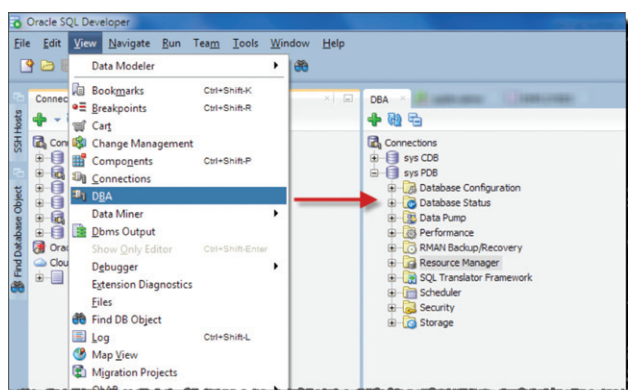


FIGURE 3: THE DBA PANEL

### Getting Started

The panel will be blank when you first open it. You will need to add connections from the main connections panel using the '+' button in the panel toolbar as needed.

Opening a connection in the DBA panel has the same effect of opening it from the main connection panel. If the connection is already open, expanding the connection tree node will show the feature areas as shown above.

I do not have the required space and time here to go over every single feature. Instead, I will try to focus on some areas that will give you a good idea of what the tool is capable of.

### Database Configuration

Unlike the reports, these screens are definitely interactive. Want to change a startup parameter? You can do that.




FIGURE 4: INITIALISATION PARAMETERS

There are a lot of parameters – so be sure to take advantage of the filter mechanism. Hit Ctrl+Space in the filter dialog to get a list of columns to filter on. Double-click on the 'Value' cell, make your changes, set your commit strategy, and post your updates. I'm assuming you know to test these changes in your QA environments before ever doing anything in production! Another area of interest in the Database Configuration tree is the Database Feature Usage report. It will show you how many times the partitioning feature in the database has been used for example.

### Data Pump

This may be one of the few features in the DBA panel that will also be of interest and value to your developers. The main view displays currently running Data Pump jobs, and more interestingly, provides a wizard to create new export and import jobs.

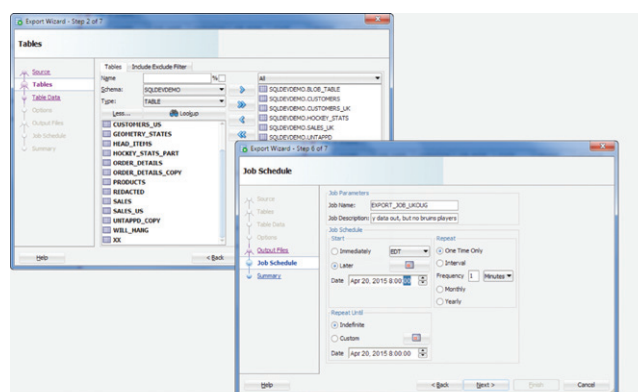


FIGURE 5: DATA PUMP EXPORT JOB WIZARD

In a matter of just a few clicks, I can quickly define the type of export, the objects to be included (or filtered), what data to filter on, all of the job details and much more.

At the end of the process, I'm given the 40 line PL/SQL anonymous block that will be fired to kick off the job. So whether you need the GUI to help you with the syntax, or you're just looking to save some keystrokes, the wizard can launch the job for you or just generate the code for you.

### Performance, Also Known as the Diagnostic Pack

A quick word regarding licensing in SQL Developer: SQL Developer is included with your database and is also separately licensed via the Oracle Technology Network when you agree to download it. For all intents and purposes, SQL Developer is free. However, using it to interact with additional cost database features requires said features to be licensed. SQL Developer will prompt you the first time you attempt to use a Diagnostic Pack feature. You can also go into the preferences and disable pack features for specific databases.

So, what can you do here? It turns out, quite a bit:

- Manage snapshots, baselines and baseline templates
- Manage Automatic Database Diagnostic Monitor (ADDM)

Tasks, drill into findings and view reports

- Run Active Session History (ASH) reports
- Generate Automatic Workload Repository (AWR) Reports
  - AWR Report Viewer
  - Difference Report Viewer
  - SQL Report Viewer

Instead of showing everything, I'll just demonstrate how to step into a ADDM report from a Snapshot.

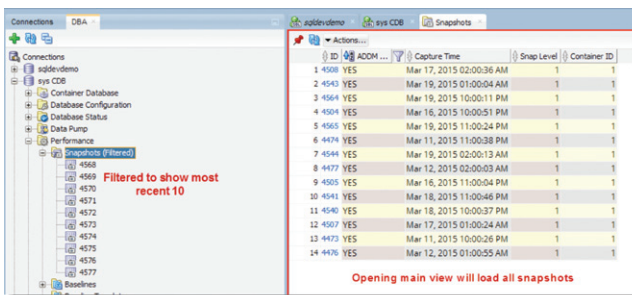


FIGURE 6: SNAPSHOTS. NOTE THAT IN A 12c MULTITENANT ARCHITECTURE DATABASE, THE SNAPSHOTS ARE COLLECTED AT THE CONTAINER DATABASE (CDB) LEVEL.

Connect to the CDB – the tree will immediately show the most recent 10 snapshots. I could force a new snapshot to be taken by right-clicking on the Snapshots node. I can also open an individual snapshot by clicking on it in the tree. Instead, I'm opening all of the snapshots by double-clicking on the main Snapshots tree node.

Once it has opened, I'm applying another filter – only show me snapshots where there are ADDM findings. I can then click on the ADDM ID to open that task.

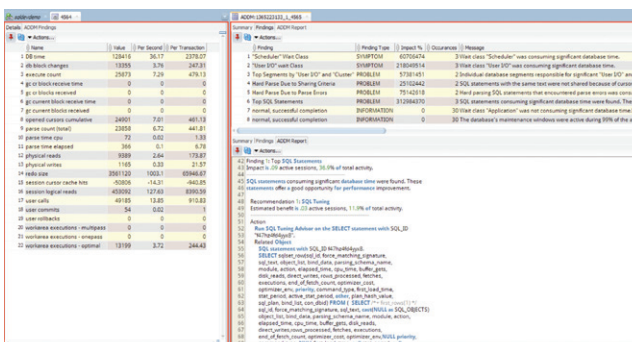


FIGURE 7: ADDM TASK FINDINGS AND REPORT

I have used a couple of my favorite SQL Developer user interface tricks in the screenshot above. I've cloned the ADDM item to a separate document. I then moved the copy of that ADDM item to a new document tab group – this allows for the screens to be viewed side-by-side. Finally, I used a split horizontal editor to see both the findings and the report in the same editor. Let's spend a few additional moments talking about the Active Session History reports since those are one of the most popular features of the Diagnostic Pack.

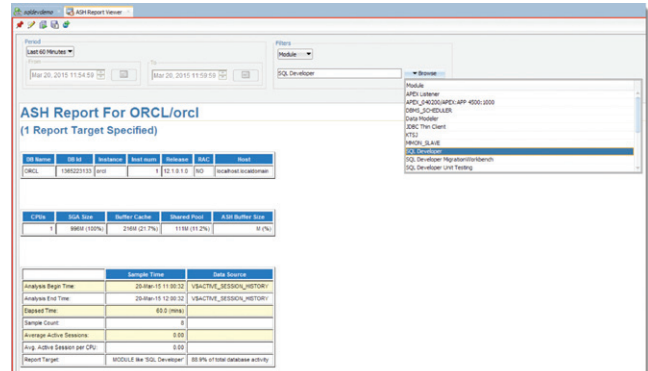


FIGURE 8: ACTIVE SESSION HISTORY (ASH) REPORT

I have run this ASH report from my pluggable database. I can set predefined filters on the time period or also set a custom start and end time. I can also optionally set a filter on something like 'Module.' SQL Developer will look for available values to filter on for you. You then run the report.

The expected HTML report is generated and you can navigate the hyperlinks directly in SQL Developer. You can also choose to send the report to your browser as a temp file or save it as an HTML file.

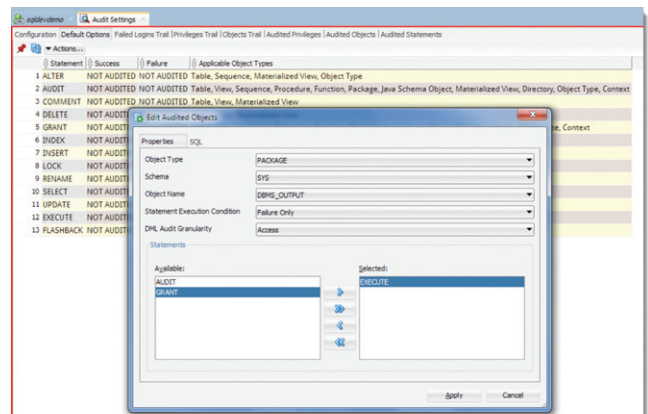
I can re-run the report with different parameters by using the 'Toggle Inputs' button on the toolbar, re-setting my filters and hitting the 'Generate Report' button again.

## Security

Can I use SQL Developer to reset a user's password? Yes.

The user interface for managing users, roles and privileges is very self-explanatory. So instead of showing that workflow, I would beg for your indulgence and would like to use this space instead to discuss our Database Auditing support.

I think many of our users completely miss that we have a complete Auditing interface in the Security tree. You can configure the Oracle Auditing settings and view the Audit trails.





The trails are standard report/grids – meaning you can filter and export them to your favorite file formats. This is a fancy way of saying that yes, you can export the data to Excel.

### Storage

Can I use SQL Developer to add space to a tablespace? Yes.

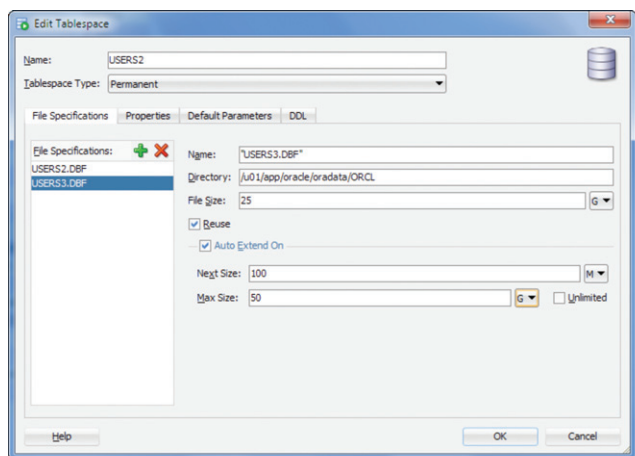


FIGURE 9: TABLESPACE EDITOR – USING THE EDIT DIALOG TO ADD A DATA FILE

You can edit a tablespace or a data file. As with all of our other wizards, you can take the generated code on the DDL page and copy that out for testing or adding to your scripts if you're not going to be doing this live in your database.

### Multitenant

You have seen me mention our multitenant architecture a few times. When connected to a 12c Container database, you'll additionally be able to manage your pluggable databases. You can create a new pluggable, drop an existing, clone an existing and plug or unplug a database. For a pluggable database, you can create a startup trigger, modify its state, and coming soon – clone it to our public Oracle Database Cloud Service!

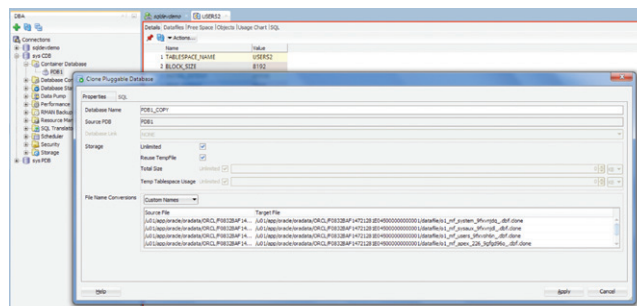


FIGURE 10: CLONING A PDB WITH SQL DEVELOPER

### Summary

Oracle SQL Developer isn't just for the developer or business user. Since version 3.0, we have offered specific features for the DBA including interfaces for Data Pump, Recovery Manager, Auditing and the Diagnostic Pack via the DBA panel.

The features assume that you are a DBA and have the DBA role, or equivalent set of privileges.

All of the wizards allow you to just generate the code required to implement the defined changes so you can copy that out to your SQL\*Plus scripts or your testing environments.

We are continuing to expand and add new features for the DBA. Oracle SQL Developer version 4.1 is now available as an Early Adopter on OTN for your evaluation and feedback.

You can download SQL Developer, read the release notes, and watch video tutorials of our major features at the Oracle Technology Network.

I presented this topic at Oracle OpenWorld 2014 and it was recorded. You can watch this session on demand, for free!

I also have the slides available in the form of Slideshare on my blog. ■



## ABOUT THE AUTHOR

**Jeff Smith**  
Product Manager, Oracle

Jeff is a Product Manager in the Database Development Tools Group at Oracle, and has been obsessing over saving people clicks and keystrokes for the last decade.

**Blog:** [www.thatjeffsmith.com](http://www.thatjeffsmith.com)

[@thatjeffsmith](https://twitter.com/thatjeffsmith)

# Creating opportunities to connect Creating connections that count

## APPS 15

UKOUG Applications Conference & Exhibition

7-9 DEC 2015 | BIRMINGHAM

The largest event in the UK for the Oracle applications community, attracting over 800 delegates. Apps15 will feature a variety of presentations ranging from users telling their stories to high level strategic presentations covering product roadmaps.

[www.apps15.ukoug.org](http://www.apps15.ukoug.org) #ukoug\_apps15

## JDE 15

UKOUG JD Edwards Conference & Exhibition

8-9 DEC 2015 | BIRMINGHAM

This respected Oracle community event attracts over 200 attendees with a variety of skills and experience. It has a consistently cutting edge agenda, captivating speakers and an exhibition packed with JD Edwards partners.

[www.jde15.ukoug.org](http://www.jde15.ukoug.org) #ukoug\_jde15

## TECH 15

UKOUG Technology Conference & Exhibition

7-9 DEC 2015 | BIRMINGHAM

The largest event in the UK for the Oracle Technology community with over 850 delegates, 200+ internationally renowned speakers, a highly targeted agenda and an exhibition hall filled with suppliers offering cutting-edge solutions and products.

[www.tech15.ukoug.org](http://www.tech15.ukoug.org) #ukoug\_tech15

## OracleScene

### Summer Edition Call for Content

# Now open - Get Writing!

Users all around the world really do benefit from the knowledge and experience you share in your Oracle Scene articles.

"Broad coverage, well-written articles. Good editorial practices applied."

"Excellent magazine with lots to read and gives new ideas"

"An informative read, with strong technical content delivered in a concise easy to digest manner. A must read for Oracle professionals."

"Great for wider issues as well as the "hot" topics"

Article submissions to [articles@ukoug.org](mailto:articles@ukoug.org) by **8th May 2015**.  
See [www.ukoug.org/oraclescene](http://www.ukoug.org/oraclescene) for guidelines.

# Oracle **Cloud** Applications

## **HCM**

Human Capital  
Recruiting  
Talent

## **CRM**

Sales  
Service  
Marketing

## **ERP**

Financials  
Procurement  
Projects  
Supply Chain

**More Enterprise **SaaS** Applications  
Than Any Other **Cloud** Services Provider**

**ORACLE®**

**JOIN US: MODERN BUSINESS SUMMIT | April 28-30, 2015**  
Hotel Pullman London St Pancras | [oracle.com/goto/ModernBusinessSummit2015](http://oracle.com/goto/ModernBusinessSummit2015)