



Learn more
about Oracle
technology with
Safari Books
Online



Use your QR code reader to access information about a group trial or visit safaribooks.com/oracletrial

SEARCH LESS. DEVELOP MORE.

With the intelligent library in the cloud.

ORACLE MAGAZINE

Oracle Adds RightNow Cloud Service
 Oracle unveils product roadmap for end-to-end customer experience / **25 Role Models**
 Great minds and big ideas shape architectural decisions / **32 Working with Records** Part 7 in a series of articles on understanding and using PL/SQL / **62 An Order of Sorts** Part 5 in a series on the basics of the relational database and SQL / **67 On Connection Pools, Cursor Differentiation, and Optimal Ordering** Our technologist cleans pools, explores cursor types, and looks for order in table creation / **71**

AVENGERS ASSEMBLE

To support the Avengers and its agents, S.H.I.E.L.D. deploys Oracle engineered systems, delivering a new level of mission-critical computing



Storage Stands Tall

Oracle engineered storage systems deliver tiered storage optimized for your I/O

Connected

One telecommunications company depends on Oracle Exadata to cover a nation



Trusted Future Protection



Fujitsu's SPARC ENTERPRISE delivers Trust and Business Continuity

ORACLE Diamond
Partner

Fujitsu and Oracle together are creating UNIX history through progressive delivery of server solutions with SPARC ENTERPRISE. Our customer commitment is reflected in overwhelming performance and excellent reliability that will protect, grow, and expand, the unique value of your business assets.

Please visit : www.fujitsu.com/oracle/



2012 Fujitsu Limited. All rights reserved. Fujitsu, the Fujitsu logo and SPARC Enterprise are registered trademarks or trademarks of Fujitsu Limited in the United States and other countries. UNIX is a registered trademark of open group in the United States and other countries. Oracle and Java are registered trademarks of Oracle and/or its affiliates. All SPARC trademarks are used under license and are registered trademarks of SPARC International, Inc. in the United States and other countries.



shaping tomorrow with you

FUJITSU

ORACLE
MAGAZINE


TM & © 2012 Marvel & Subs.

SUPER HERO IT

/36

Is your data center powerful enough to protect the nations and peoples of Earth from all threats, terrestrial or extraterrestrial? If you're Agent John Smith, it better be. Smith is the chief systems administrator for the Supreme Headquarters International Espionage Law-Enforcement Division (S.H.I.E.L.D.), where he works for Nick Fury. And just as Fury must assemble the Avengers to battle threats to the world that no one else can handle, Smith must assemble the most powerful data center to deliver extreme performance and Earth-saving analytics. See why Smith counts on Oracle engineered systems, storage, and analytics to support the Avengers. —*By Rich Schwerin*

Storage Stands Tall / 42



Oracle's enterprise storage solutions combine the best hardware with software designed to take advantage of that hardware, to meet the requirements of various types of I/O processing and long-term data archiving. Learn why organizations depend on Pillar Axiom 600, Sun ZFS Storage Appliance, and StorageTek modular library systems and drives from Oracle to support their operations and keep information available. —*By David Baum*

Connected / 48



Digicel Haiti, the largest mobile communications operator in Haiti, faced nightly data collection processes that were taking too long on the company's server infrastructure, and as a result, business-critical daily reports were not reliably delivered on time. To overcome that problem, Digicel migrated its 38 TB data warehouse to Oracle Exadata Database Machine. Read how the benefits went far beyond solving the company's immediate reporting problems. —*By David A. Kelly*

**It only takes a second for
an accident to happen.
Why should it take forever to fix it?**

Business responds faster because
data does more in the AT&T network.

In here, intelligent data automates all the right
actions, at all the right points, all at the same time.


Details of the accident are collected and distributed.
Insurance claims are filed and reviewed. Parts are
preordered; paint colors are premixed.

In here, wrecks get repaired faster because the same
data that reports an accident is already at work fixing it.

It's the AT&T network—a network of
possibilities teaching data how to do more.
To learn more, visit att.com/business

WRECKED

REPAIRED



Rethink Possible® 



Download the free scanner app at <http://scan.mobi> and scan this code to learn more.

© 2012 AT&T Intellectual Property. All rights reserved. AT&T, the AT&T logo and all other AT&T marks contained herein are trademarks of AT&T Intellectual Property and/or AT&T affiliated companies.

Up Front / 6

FROM OUR READERS / 6

Readers tell us what they think.

FROM THE EDITOR / 8

Different Coverage

—Tom Haunert

WORK IN PROGRESS / 11

Obsessed with PL/SQL

Steven Feuerstein talks about his PL/SQL journey.
—Tom Haunert

MASHUP / 14

News, views, trends, and tools



At Oracle / 17

EVENTS / 17

Find out about upcoming technology and industry events.

RESOURCES / 19

Your guide to Oracle Webcasts, podcasts, white papers, and more

BRIEFS / 22

The latest product news

NEWS / 25

RightNow Adds Service to Oracle CX

Oracle unveils product roadmap for end-to-end customer experience.
—Leslie Steere

INTERVIEW / 27

Upon Further Analysis

Oracle Advanced Analytics delivers key insights from big data.
—Rich Schwerin

Community / 28

PARTNER NEWS / 28

BOOK BEAT / 28

COMMUNITY BULLETIN / 28

The Magic of Virtualization

Oracle Technology Network embraces all things virtual.
—Justin Kestelyn

ARCHITECT / 32

Role Models

Great minds and big ideas shape architectural decisions. —Bob Rhubart

UP CLOSE / 34


Big Data Boot Camp

An early adopter shares big data expertise at COLLABORATE 12.
—Jeff Erickson

PEER-TO-PEER / 35

Playing Favorites

Gökhan Atıl, Doug Gault, Patrick Barell
—Blair Campbell



Technology / 53

ORACLE ADF / 53

Personalized Dashboards

Extend an Oracle ADF Faces dashboard application to support persistent user personalization.
—Frank Nimphius

BUSINESS INTELLIGENCE / 59

Working with the Summary Advisor

Create in-memory aggregates for better performance on Oracle Exalytics In-Memory Machine. —Mark Rittman

PL/SQL 101 / 62

Working with Records


Part 7 in a series of articles on understanding and using PL/SQL —Steven Feuerstein



SQL 101 / 67

An Order of Sorts

Part 5 in a series on the basics of the relational database and SQL
—Melanie Caffrey



ASK TOM / 71

On Connection Pools, Cursor Differentiation, and Optimal Ordering

Our technologist cleans pools, explores cursor types, and looks for order in table creation. —Tom Kyte

Comment / 77

IN THE FIELD / 77

Virtualized Communities

A virtual community discusses virtualization and delivers real information.
—Andy Flower

ANALYST'S CORNER / 79

Enterprise Servers: Today and Tomorrow

What's driving the server market? —David A. Kelly

ANALYST'S CORNER / 80

Engineered Storage

The present and future of enterprise storage
—David Baum

EDITORIAL

Editor in Chief

Tom Haurert tom.haurert@oracle.com

Senior Editor

Caroline Kvitka caroline.kvitka@oracle.com

Managing Editor

Jan Rogers jan.rogers@oracle.com

Contributing Editor and Writer

Blair Campbell

Editor in Chief, Oracle Technology Network

Justin Kestelyn justin.kestelyn@oracle.com

Technology Advisor

Tom Kyte

Contributors

Marta Bright, Jeff Erickson, Fred Sandmark, Rich Schwerin, Leslie Steere

DESIGN

Senior Creative Director

Francisco G Delgadillo

Senior Design Director

Suemi Lam

Design Director

Richard Merchán

Contributing Designers

Jaime Ferrand, Chris Strach

Production Designers

Sheila Brennan, Kathy Cygnarowicz

EDITORIAL BOARD

Ian Abramson, Karen Cannell, Andrew Clarke, Chris Claterbos, Karthika Devi, Kimberly Floss, Kent Graziano, Taqi Hasan, Tony Jambu, Tony Jedlinski, Ari Kaplan, Val Kavi, John King, Steve Lemme, Carol McGury, Sumit Sengupta, Jonathan Vincenzo, Dan Vlamis

SUBSCRIPTION INFORMATION

Subscriptions are complimentary for qualified individuals who complete the subscription card found in each issue or online at oracle.com/oraclemagazine. For change of address, mail in label with the new address to: *Oracle Magazine*, P.O. Box 1263, Skokie, IL 60076-8263.

MAGAZINE CUSTOMER SERVICE

oracle@halldata.com Fax +1.847.763.9638 Phone +1.847.763.9635

PRIVACY

Oracle Publishing allows sharing of its mailing list with selected third parties. If you prefer that your mailing address or e-mail address not be included in this program, contact Customer Service at oracle@halldata.com.

PUBLISHING

Publisher

Jeff Spicer jeff.spicer@oracle.com

Production Director and Associate Publisher

Jennifer Hamilton jennifer.hamilton@oracle.com +1.650.506.3794

Senior Manager, Audience Development and Operations

Karin Kinnear karin.kinnear@oracle.com +1.650.506.1985

ADVERTISING SALES

Associate Publisher

Kyle Walkenhorst kyle@sprocketmedia.com +1.323.340.8585

Northwest and Central U.S.

Tom Cometa thomas.cometa@sbcglobal.net +1.510.339.2403

Southwest U.S. and LAD

Shaun Mehr shaun@sprocketmedia.com +1.949.923.1660

Northeast U.S. and EMEA/APAC

Mark Makinney mark.makinney@sprocketmedia.com +1.805.709.4745

Recruitment Advertising

Tim Matteson tim@mattesonmail.com +1.310.836.4064

Mailing-List Rentals

Contact your sales representative.

RESOURCES

Oracle Products

+1.800.367.8674 (U.S./Canada)

Oracle Services

+1.888.283.0591

Oracle Press Books

oraclepressbooks.com



Copyright © 2012, Oracle and/or its affiliates. All Rights Reserved. No part of this publication may be reprinted or otherwise reproduced without permission from the editors. *ORACLE MAGAZINE* IS PROVIDED ON AN "AS IS" BASIS. ORACLE EXPRESSLY DISCLAIMS ALL WARRANTIES, WHETHER EXPRESS OR IMPLIED. IN NO EVENT SHALL ORACLE BE LIABLE FOR ANY DAMAGES OF ANY KIND ARISING FROM YOUR USE OF OR RELIANCE ON ANY INFORMATION PROVIDED HEREIN. The information is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle. Oracle and Java are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

Oracle Magazine (ISSN 1065-3171) is published bimonthly with a free subscription price by: Oracle, 500 Oracle Parkway, MS OPL-3C, Redwood City, CA 94065-1600. Periodicals Postage Paid at Redwood City, CA, and additional mailing offices. • POSTMASTER: Send address changes to: *Oracle Magazine*, P.O. Box 1263, Skokie, IL 60076-8263.

Printed by Brown Printing

ORACLE PRODUCT INDEX

PRODUCT NAME	PAGE	PRODUCT NAME	PAGE	PRODUCT NAME	PAGE
JD Edwards EnterpriseOne solutions	22	Oracle Exadata	19, 30, 36, 48	Oracle VM Server for SPARC	22
Ksplice zero-downtime kernel update	22	Oracle Exalogic	36	Oracle VM VirtualBox	31
MySQL	19, 20, 28	Oracle Exalytics In-Memory Machine	29, 36, 59	Oracle WebCenter	19
MySQL Cluster	22	Oracle Fusion Applications	20	Oracle WebLogic	20, 22
Oracle Advanced Analytics	27	Oracle Healthcare Master Person Index	23	Oracle WebLogic Server	31
Oracle Airline Data Model	23	Oracle JDeveloper	53	Pillar Axiom 600	36, 42
Oracle Application Development Framework	22, 31, 53	Oracle Linux	19, 20	SPARC Enterprise M9000	42
Oracle Application Express	22, 28, 35	Oracle Public Cloud	25, 31	SPARC processors	19, 79
Oracle Big Data Appliance	29, 34, 36	Oracle RightNow CX Cloud Service	25	SPARC T4 servers	20
Oracle Business Activity Monitoring	28	Oracle Service Bus	22	StorageTek SL500	42
Oracle Business Intelligence Enterprise Edition	31	Oracle SOA Suite	19	StorageTek SL3000	42
Oracle Business Intelligence Foundation Suite	23	Oracle Solaris	19, 20, 31, 79	StorageTek SL8500	36, 85
Oracle Data Integrator	19	Oracle SQL Developer	22	StorageTek T10000C	42
Oracle Database	11, 28, 35, 42, 62, 67, 71	Oracle User Productivity Kit	23	Sun Blade X6270 M2	29
Oracle Database App Development VM	31	Oracle User Productivity Kit Professional	23	Sun Storage Archive Manager	42
Oracle Database Appliance	20, 29	Oracle Virtual Assembly Builder	23	Sun ZFS Storage Appliance	24, 36
Oracle Enterprise Manager	35	Oracle VM	19, 20, 31, 77	Unbreakable Enterprise Kernel	24
Oracle Enterprise Manager Ops Center	24				



Your corrections, your opinions, and your requests: Here's your forum for telling us what's right and wrong in each issue of *Oracle Magazine*, and for letting us know what you want to read.

DOWNLOADS, PLEASE

Is it possible to download specific technology articles from *Oracle Magazine*?
Shueb Khan

The editors reply: Subscribers to the digital edition of the magazine can download each issue in PDF format. If you would like to switch from a print to a digital subscription, please contact oracle@halldata.com or visit bit.ly/zs8i7R. Another option is to download the free Oracle Magazine app for iPad or iPhone.

MISSING LINKS

In "Working with Dates in PL/SQL" (January/February 2012), Part 5 in the PL/SQL 101 series, you don't have any links to the previous parts. It would be good to add them to the beginning of the article.
Robert Wolf

The editors reply: Our link to the previous parts appeared in that article's "Next Steps" box. All of Steven Feuerstein's past magazine content is available at bit.ly/weeVKY, including the previous parts of his PL/SQL 101 series.

KINDNESS FOR KINDLE

I've been reading *Oracle Magazine* in Portugal since 2005. Is a Kindle edition available?
Rui Madaleno

The editors reply: While a native Kindle edition is not in the works, subscribers to the digital edition can download a PDF that they can view on their second-generation and later Kindle devices. It is

also possible to have a PDF converted to Kindle format using your Send to Kindle e-mail address and Amazon's file conversion service.

SPEAKING OF ERRORS . . .

In "Error Management" (March/April 2012), the examples Steven Feuerstein gives for application-defined exceptions mix exception processing with business logic (and also would hard-code business parameters into the code).

My view is that exceptions should handle events that are, well, exceptional, to the detailed specification of the code. For example, exceptions could be used to report that parameters are outside the specified domain of the input. The specified functionality of the code should be handled entirely using normal conditional logic and return parameters.

Another good use for exceptions would be to communicate failures of assertion tests placed in the code, perhaps basing these assertions on predicate logic techniques defined in the 1970s and subsequently by Dijkstra, Floyd, and Hoare.

Finally, it would be great to explore error handling and reporting in more detail in another article.
Miles Thomas

Steven Feuerstein replies: Miles raises a number of excellent points. I agree that business rules are best placed within their own functions, so that a developer can validate the business rule and then decide if an exception should be raised or a different branch of logic should be taken. When writing "101" articles, it is always a challenge to combine introductory concepts and syntax with best practices. See my blog

(feuerthoughts.blogspot.com) for a more complete discussion of these topics—and Miles' complete critique of my article.

The editors reply: Steven Feuerstein has written about exceptions and error handling numerous times in the many years he has written for *Oracle Magazine*. Refer to the index of his articles at bit.ly/weeVKY for links to his Oracle Magazine articles since 2003.

I am a big fan of Steven Feuerstein and participate in the PL/SQL Challenge (plsqlchallenge.com). I came across his *Oracle Magazine* article "Error Management," and it is a very good article. There is a minor issue, however, in Code Listings 2 and 3: the declaration of `_code` is missing.
Ramesh Thangamani

The editors reply: Thank you for this correction. The declaration

```
_code PLS_INTEGER := SQLCODE;
```

has been added to the code listings in the online article at bit.ly/xdbydd.

SEND MAIL TO THE EDITORS

Send your opinions about what you read in *Oracle Magazine*, and suggestions for possible technical articles, to opubedit_us@oracle.com. You can also follow our [@oraclemagazine](https://twitter.com/oraclemagazine) Twitter feed or join us on Facebook at bit.ly/orclmagfb. Letters may be edited for length and clarity and may be published in any medium. We consider any communications we receive publishable.

Isenberg MBA. Your Way.

Face-to-Face. Online. Blended.
AACSB Accredited.
Adapted to your lifestyle.



Different Coverage

Oracle Magazine always collects content, but this issue also collects covers.



Each issue of *Oracle Magazine* collects different editorial content in familiar sections and a particular order. And just as every editorial piece is designed to have an interesting beginning, a constructive middle, and a conclusive end, the various sections of the magazine are arranged for a similar flow. From Up Front, At Oracle, and Community to Features, Technology, and Comment, the idea is to give the magazine a beginning that makes you want to pick it up (in print), click it (for digital), or tap it open (for mobile distributions); a middle that provides information and awareness about Oracle and the Oracle community; and an end that makes you look forward to the next issue.

The beginning of each issue, the very first section and page, is key to getting readers to pick up or open the magazine. And that beginning is the cover.

CURRENT COVERAGE

In most cases, the editorial team settles on only one cover for each issue of *Oracle Magazine*. In a few cases, however, we've published an issue with two different covers—with some subscribers receiving one version and some receiving the other—and that has consistently caused confusion. We've typically followed these multicover issues with an explanation, but for this issue of *Oracle Magazine*, we'd like to get out in front of any potential confusion with the multicover explanation.

The discussion about what story to put on the cover of this particular issue of the magazine was a brief one. How to put that story—featuring Marvel's *The Avengers*—on the cover was a much longer conversation. The solution: this issue of *Oracle Magazine*

has been published with six different covers, each highlighting one of the Marvel super heroes. Each cover version is numbered—1 through 6—in the "Special Collector's Issue" logo.

Who gets what cover? Our print subscribers will receive a printed issue with one of the six covers. All subscribers—print and digital—and iOS application downloaders will be able to view and collect all six cover images. To view and download the cover images, follow the prompts from your digital subscription or your mobile application. All subscribers—and everyone else reading this—are also invited to view and download the cover images from the *Oracle Magazine* Facebook page at bit.ly/orclmagfb.

I hope whatever cover or covers you view as you begin to explore this issue of *Oracle Magazine* are the prelude to a successful journey into Oracle technology. And if six cover images don't provide enough coverage of these particular characters for you—and quite frankly, they're really not enough for me—*The Avengers* is in theaters on May 4.

Tom Haurert, Editor in Chief
tom.haurert@oracle.com

NEXT STEPS

Marvel's *The Avengers* in Theaters May 4

LEARN more about *The Avengers*
oracle.com/theavengers

GET covered
bit.ly/orclmagfb

The adventures of Harry & Garry

LOOK WHAT ALLROUND AUTOMATIONS SENT ME GARRY



A PL/SQL DEVELOPER-MAN T-SHIRT!



EVERY PROJECT WILL RUN ON SCHEDULE WITH THE HELP OF PL/SQL DEVELOPER-MAN.

WELL, PUT IT ON THEN HARRY.



IT SAYS "DEVELOPE"! YOU'RE TOO SKINNY!



NOT TO WORRY, I JUST NEED TO WORK OUT A BIT.



NOW IT SAYS "SQL DEVELOPER"!



THAT'S TERRIBLE! I WON'T REST UNTIL IT SAYS PL/SQL DEVELOPER-MAN



HOW IS THE PROJECT GOING?

ON SCHEDULE, PL/SQL DEVELOPER HAS BEEN A GREAT HELP.



PL/SQL DEVELOPER-MAN ON THE OTHER HAND...

PL/SQL Developer is the Oracle development tool that gives you maximum productivity, ease of use and all the features you need, for a reasonable price.

Visit our website for additional details:
www.allroundautomations.com/plsqldev



allroundautomations

MY KIND OF ORACLE EDUCATION & TRAINING • APRIL 22-26, 2012

MANDALAY BAY, LAS VEGAS



COLLABORATE 12

TECHNOLOGY AND APPLICATIONS FORUM
FOR THE ORACLE COMMUNITY

▶ EXCLUSIVE IOUG
REGISTRATION BENEFITS



COLLABORATE 12 • REGISTRATION OPENS NOVEMBER 7, 2011

For under \$2,000, COLLABORATE 12 — the IOUG (Independent Oracle Users Group) Forum, offers access to over 1,000 hours of Oracle-related education and training through Hands-on Labs, Boot Camps, full-day Deep Dives, and customer case studies.

With educational tracks devoted to:

- BI/Data Warehousing/EPM
- Big Data
- Database Administration and Development
- Manageability
- Security, Risk and Compliance

On hot topics like:

- Cloud Computing
- Exadata
- High Availability
- Virtualization
- DBA101

<IOUG> EXCLUSIVES:

independent oracle users group

FULL-DAY DEEP DIVES COMPLIMENTARY FOR IOUG REGISTRANTS ONLY!

All other attendees pay the regular rate of \$599.

8 hours of technical training, on topics like:

- Virtualization
- Business Intelligence
- Manageability
- DBA/Developer
- WebCenter
- RAC
- Big Data
- Performance Engineering

EXCLUSIVE ACCESS TO HANDS-ON LABS (A \$350 VALUE!)

Gain actual experience and turn theory into practice at 2-hour hands-on labs, focused on Business Intelligence topics like analytics, warehousing, and OBIEE.



LEARN ABOUT EXCLUSIVE IOUG REGISTRATION BENEFITS & REGISTER:

<http://collaborate12.ioug.org>

Obsessed with PL/SQL

Steven Feuerstein talks about his long and productive PL/SQL journey.

Steven Feuerstein is a PL/SQL evangelist, author, and Oracle Magazine columnist. Oracle Magazine editor in chief Tom Haunert recently sat down with Feuerstein to talk about his past, present, and future with PL/SQL; his advice for PL/SQL developers; and the future of PL/SQL. The following is an excerpt from that interview. Download the full podcast at oracle.com/magcasts.

Oracle Magazine: How did you get started with Oracle technology?


Feuerstein: It was pretty much an accident. I had been working as a programmer, mostly in FORTRAN, when I got out of college, and I started dabbling with databases. I moved into a consulting job at an insurance company and took a step forward with my database programming. That job was incredibly boring. I was basically managing reports.

So, I started looking for a database-related job, and I answered an ad for a company that needed somebody with relational database experience. I had no idea what relational databases were, but I found and memorized E.F. Codd's 12 rules of relational databases and went to my interview. It was with a headhunter looking for people to work for Oracle—which I had barely even heard of—in presales. He didn't ask me a single question about relational databases, though, and the big question for me became, did I want to wear a suit and tie every day?

Fortunately, I decided it was worth it and went from being a programmer to going out and doing presentations on Oracle technology. And at that point—this was 1987—a presentation on Oracle technology consisted of showing canned SQL*Plus scripts and saying, "Check it out—a join of two tables! Check it out—you can update the content of a table with a simple statement!"

Oracle Magazine: What drew you to PL/SQL?

Feuerstein: I've always been interested in creating new things. I joined Oracle in presales,



PL/SQL evangelist Steven Feuerstein

but all the time I was learning how to use the tools and looking for ways to apply them.

Then in 1989, I realized that our sales team had no sales management tools. So, I wrote an application with SQL *Forms 3.0 called Team Sell to help salespeople keep track of activity in their accounts.

At the same time, the head of Oracle USA sales was building a consulting team to create applications on a national level. I was recruited to work with his team and soon found myself spending two-thirds of my time building applications in SQL *Forms with PL/SQL. That was the real grabber for me. When PL/SQL was introduced to the SQL *Forms toolset, it opened up huge vistas in the flexibility of both forms and report building. I was delighted to put SQL *Forms 2.3 triggers behind me and dive into the world of PL/SQL.

Oracle Magazine: Can you give us the high points of your work with PL/SQL and tell us what projects you're working on now?

Feuerstein: I left Oracle in '92 and became a consultant, where my main focus was still on Oracle development tools and PL/SQL. About '94, publishers were starting to ask people if they would like to write books about Oracle7 Database. I wrote a lot of pretty awful science fiction in high school, and teachers liked my essays, so I thought that I should be able to write a technical book. I didn't know Oracle7 Database all that well, but I knew PL/SQL, so I counterproposed. The result was the first book on PL/SQL published independently of Oracle: *Oracle PL/SQL Programming* (O'Reilly Media, 1995).

I wrote it with a sense of humor, and the result seemed to be that developers found that they could actually understand what I wrote—and even enjoy it. Well, that book changed my life, virtually overnight. And from that point on, I focused almost exclusively on PL/SQL—studying it, writing about it, and writing programs in it. I have published 10 different books on the PL/SQL language, which is probably about six or seven too many. At this point, I update the three most popular of these: *Oracle PL/SQL Programming*, *Oracle PL/SQL Language Pocket Reference*, and *Oracle PL/SQL Best Practices* (all O'Reilly Media).

I've also been involved in many tools development efforts over the years. Back in '99, I discovered test-driven development and created utPLSQL, which is basically JUnit for PL/SQL. There are probably several hundred organizations still using that around the world, though it's not an actively maintained project. I was also involved in building another automated testing tool, Quest Code Tester for Oracle, as well as error management and code generation tools.

But for the last several years, my focus has been less on tools and more on strengthening the worldwide community of PL/SQL developers. PL/SQL is a rich, mature language—and there are tens of thousands of program-

mers who have been working with it for a long time. I think one of the key challenges for all of us is to find a way to cohere those individuals into a community in which everybody's expertise can be shared and magnified. Too much of the way that we tend to work in the technology world is that people look for gurus, and basically ask for a "download." I'm pleased to do it, but many people out there know PL/SQL as well as I do in certain areas and have a lot more to offer in their real-world experience. So, what I'm really focusing on these days, primarily through the PL/SQL Challenge Website, is to provide a way for people around the world to not only learn from me, but for me and everybody else to learn from them.

Oracle Magazine: You've said that your days of traveling for work are limited. What's the latest on your PL/SQL roadwork plans?

Feuerstein: In 2010, I finally decided that I had enough of living in airplanes, airports, and hotels. I had been doing 100,000-plus miles a year, and I was tired of being away from home. I was interested in working on building this global community, and the way you build a global community is not necessarily to try to meet everybody personally, but to provide Websites where people can get together and share and promote their expertise. These days, in terms of traveling, I prioritize key conferences—for example, Oracle OpenWorld and ODTUG Kscope.

Oracle Magazine: Do you have any suggestions for aspiring and senior PL/SQL developers to help them manage their own PL/SQL training and experience?

Feuerstein: I get a lot of questions from developers about what they should be doing with their careers. They're trained in PL/SQL. They've got pretty solid experience under their belts, and they want to know what to do to protect and enhance their careers. They also want to know about the future of PL/SQL. Ever since Oracle8i Database, when the Java Virtual Machine was put into the database, people have been asking, "Is PL/SQL going to stick around? Is Oracle switching over to Java? Is PL/SQL going to go away?"

PL/SQL is absolutely not going to go away. It's critical technology in so many parts of the Oracle technology stack. And even if it weren't, there are literally millions of lines of PL/SQL code out there that have been running in

“PL/SQL is absolutely not going to go away. It's critical technology in so many parts of the Oracle technology stack.”

mission-critical applications for decades. So, I think that everybody should be quite confident that PL/SQL is not only going to stick around, but it will be enhanced. The PL/SQL team is still working hard at making it the best possible database programming language.

Having said that, I have heard from a number of developers that they're seeing the opportunities for PL/SQL shrinking. That makes some sense, because it's not the only language you can use to deploy Oracle applications. I think it's really critical that anybody who is a PL/SQL developer today also has a working knowledge of Java—so that they can, at a minimum, leverage Java technology from within their PL/SQL programs.

It's also important that PL/SQL developers are familiar with XML, which will be a critical mechanism for managing and transmitting information between application architecture elements. Also make sure you're really strong on the latest implementations of SQL and the enhancements to SQL in Oracle Database, so you're not writing PL/SQL when you can do the work in SQL instead.

Finally, if you're into PL/SQL and you want to keep it at the top of your technology stack, learn about Oracle Application Express. It's a classic "have your cake and eat it, too" scenario: you get to leverage your PL/SQL skills to build Websites! How cool is that? Oracle Application Express is going to play a bigger and bigger role in the Oracle world in years to come. As more Oracle Application Express-based applications are deployed, I think we'll actually see an upsurge in PL/SQL popularity.

Oracle Magazine: What are the key features that developers should know about PL/SQL?

Feuerstein: One critical datatype of PL/SQL that you've got to be really sharp on is collections—PL/SQL's version of arrays. They form the foundation for critical performance-related features of PL/SQL, including FORALL and BULK COLLECT.

In addition, every PL/SQL developer needs to be aware of the Function Result

Cache feature in Oracle Database 11g. This is a really nice feature that involves putting one keyword into your existing functions, `result_cache`, and potentially speeding up the performance of data retrievals from the database by a significant margin, with minimal impact to your code.

And as a best practice that also enables the use of the Function Result Cache, I recommend that PL/SQL developers encapsulate their implementation details as much as possible. Hide your queries behind functions, hide your SQL statements, hide your formulas, hide your literals, and so on. So encapsulate and learn about the Function Result Cache, and then you'll be able to see a significant boost in performance just in the process of upgrading to Oracle Database 11g and taking advantage of those new features.

Oracle Magazine: Do you have any words of wisdom for PL/SQL newcomers?

Feuerstein: There are people coming into the language, and there are people who have been around for 20 years. I tell newcomers to PL/SQL: don't be intimidated by people who have been doing it for a long time. We're often stuck in ruts, doing things the "old way." Just make sure that you're keeping up with the latest and leveraging the language fully. You'll have caught up with us in no time. ◀

NEXT STEPS

 **LISTEN to the podcast**
oracle.com/magcasts

VISIT Steven Feuerstein's recommended Websites to learn more about PL/SQL

Steven Feuerstein.com
stevenfeuerstein.com

PL/SQL Obsession
toadworld.com/sf

PL/SQL Challenge
plsqchallenge.com

AskTom
asktom.oracle.com

oracle-developer.net
oracle-developer.net

PROJECT MANAGEMENT

just got **EASIER.** And **FASTER.**



Sure, you could research, buy, host and implement your own Project Management software. But why inflict that misery on yourself when LOADSPRING can make it effortless?

Imagine getting industry-leading PM software, direct from us, hosted by us, implemented by us and made outlandishly simple by us. That's the idea behind our secure turn-key and custom solutions.

And with our even broader range of services, products, training and support, we'll implement your applications faster and get your people up to speed in record time. All while accelerating your IT project executions and boosting ROI.

In the end, we do the heavy lifting, while you make it all look easy.



SOFTWARE



ACCESS



SERVICES



PRODUCTS




SECURITY



SUPPORT





“Innovation is unusual in most businesses and therefore it takes a low priority. Innovation must make a transition from an ‘unusual’ infrequent activity to a more common, consistent capability, regularly applied, to advance a company’s visibility, culture, and profits.”

—Jeffrey Phillips, author of *Relentless Innovation* (McGraw-Hill Professional, 2012)

PRODUCTIVITY APPS

Work smarter with these apps for your smartphone or tablet.

HQ: To Do

A straightforward task and project manager, this app assigns priorities, due dates, interim steps, and color codes to tasks and to-do lists. US\$1.99 (iPhone). sleepinggiantapps.com

LogMeIn Ignition

This app lets you remote-control a PC or Mac over WiFi or 3G so you can move files, run programs, stream audio and video, and more. Price varies from free to US\$100, depending on platform and features (iPhone, iPad, Android). logmein.com

SaveMeeting

Record audio from a meeting to the cloud; have it automatically transcribed; then share the recording, transcription, and any handouts with attendees. Free, including 1,000 minutes of recording (iPhone, iPad, Android). savemeeting.com

Penultimate

If the virtual keyboard is too slow when ideas are flying, you can scribble and sketch using your fingertip. Then save your “notebook” and share it with coworkers. US\$0.99 (iPad; free iPhone version is called Handwriting). cocoabox.com



VIDEO, STARRING YOU

Making a video of yourself using your iPhone is awkward. Enter Swivl, a device that holds an iPhone, or any other video camera under 6 ounces. The base unit holds the camera and has 360-degree horizontal and vertical following capability. The marker is a tiny device that wirelessly communicates with the base so it knows which way to turn. Wear the marker around your neck, or use it like a remote and point it at what you’d like to film (such as a whiteboard). US\$179. swivl.com

Hand Helper

Sitting in front of a computer for eight hours a day can take a toll on your back, neck, eyes, wrists, and hands. You can get a massage for your neck and back, and glasses can relieve eye strain. Now, you can also get a 5- to 15-minute hand massage with the Hand Reflexology Massager from Hammacher Schlemmer. This device applies gentle massage by putting pressure on trigger points on the palm and back of the hand using a mechanical air pressure system that kneads the hand. According to the company, the soothing rhythm helps to relieve tension, stress, and fatigue associated with extended keyboard and mouse use. US\$119.95. hammacher.com



DEADLINE: MISSED


13 percent of IT workers rarely or never deliver IT projects on time, according to a poll of more than 500 business technology professionals. A plurality of 45 percent cited poor requirements planning as a top reason their projects don’t deliver expected results.

Source: InformationWeek 2012 Enterprise Project Management survey, bit.ly/zlSkup

SOCIAL MEDIA AREN'T ALL THE SAME

A survey of nearly 500 business and IT executives in the U.S. found that various social media platforms are being used very differently in business—internally, with customers, and with partners. Respondents were asked, “How do you and/or your company use these social media for business?”

■ Internal to organization ■ Externally with business partners
■ Externally with customers ■ Other



Source: PwC’s Fourth Annual Digital IQ Survey, pwc.to/yaOacZ

CHRIS STRACH



Make your data dance.

Managing your data is a dance between insight and accuracy. Data quality and data integration solutions from Melissa Data make it easy to master each step. We'll help you profile, cleanse, match, enhance and update your customer contact data to build the proper foundation for enterprise business initiatives such as BI, ETL and MDM. Our affordable solutions scale to fit your needs – from thousands of records to millions. With the right partner and the right choreography, your data will really boogie.

Set your data in motion – see it come to life.

Go to MelissaData.com
Call 1-800-MELISSA (635-4772)



MELISSA DATA®
Your Partner in Data Quality

Select Specialized Partners



**70+ Oracle Product
Specializations**

**Skills Certified by
Independent Testing**

**Successful Implementations
Confirmed by Customers**

3,000 Partners with 70,000 Specialists

ORACLE®

oracle.com/specialized

Technology Events

Conferences and sessions to help you stay on the cutting edge



Oracle Cloud Conference

May 3, Helsinki, Finland; May 8, Hamburg, Germany; May 9, Berlin, Germany; May 10, Cologne, Germany

Learn from Oracle's cloud experts how to choose a public or private cloud that fits your enterprise; how to transition to a private cloud that includes self-service, autoscaling, and metering/chargeback; and how to select from Oracle's various cloud services. Other dates and venues to be announced at bit.ly/omagcloud.

CS Week Conference 36

May 1-4, Dallas/Fort Worth, Texas

► csweek.org/conference36

More than 60 workshops at this utility customer service conference cover topics including smart infrastructure, customer engagement, payments, credit and collections, field customer service, billing, and strategy and management.

Forrester's Enterprise Architecture Forum

May 3-4, Las Vegas, Nevada

► bit.ly/xJ6psl

Learn how to infuse enterprise architecture into a business to improve business agility while streamlining costs and minimizing complexity. Sessions include presentations by analysts and industry experts, executive forums, case studies, and networking events.

eHealth Week

May 7-9, Copenhagen, Denmark

► worldofhealthit.org/2012

This gathering of global professionals working in healthcare and health IT includes dedicated sessions for CIOs, a World of Health IT track, presentations on host country Denmark's e-health initiatives, and speakers from healthcare, government, and technology companies.

Metering Billing/CRM Asia 2012

May 8-9, Bangkok, Thailand

► metering-asia.com

This two-day, three-track program for the utili-

ties industry covers smart meters, meter designs and features, meter data management, and customer engagement.

Oracle Utilities Day in Benelux

May 9, Utrecht, Netherlands

► oracle.com/events

Analysts, industry experts, and Oracle executives will discuss the major trends and challenges facing utilities in Benelux. The event features multiple parallel streams and case studies in the morning followed by an afternoon roundtable discussion.

2012 ACORD LOMA Insurance Systems Forum

May 15-17, Orlando, Florida

► acordlomaforum.org

Insurance technology thought leaders will discuss analytics, business strategy, cloud computing, enterprise architecture, mobile technology, operational efficiency, regulatory issues, social media tools, system development, and related topics.

Geecon

May 16-18, Poznan, Poland

► 2012.geecon.org

This conference on Java and Java Virtual Machine-based technologies pays special attention to dynamic languages such as Groovy and Ruby. Attendees will discuss modern software development methodologies, enterprise architectures, software craftsmanship, design patterns, and distributed computing.

ORACLE USER GROUPS

Oracle Performance Firefighting 2012

May 7-9, Melbourne, Australia
bit.ly/wtTB64

Boulder Java User Group Meeting

May 8, Boulder, Colorado
boulderjug.org

DOAG 2012 Applications

May 8-10, Berlin, Germany
doag.org

Advanced Oracle Performance Analysis

May 10-11, Melbourne, Australia
bit.ly/xpB5M9

Bulgarian Oracle User Group Spring Conference

May 11-13, Borovets, Bulgaria
bgoug.org

UKOUG Solaris SIG Meeting

May 16, London, England
ukoug.org

Oracle Conference on the James

May 17, Williamsburg, Virginia
voug.org

Utah Java Users Group Meetings

May 17 and June 21, Murray, Utah
ujug.org

Rocky Mountain Oracle Users Group Quarterly Educational Workshop

May 18, Broomfield, Colorado
rmoug.org

Southern California Oracle Users Group Quarterly Meeting

May 22, Irvine, California
socoug.org

2012 Oracle Spatial User Conference

May 22-23, Washington DC
locationintelligence.net

Greater Cincinnati Oracle User Group Meeting

May 23, Cincinnati, Ohio
gcoug.org

OUG Harmony 2012 Latvia

June 1, Riga, Latvia
jug.lv

NYOUG Summer General Meeting

June 5, New York, New York
nyoug.org

New England Oracle Applications User Group Conference

June 11, Worcester, Massachusetts
neoaug.org

UKOUG Hyperion Planning and Essbase SIG Meeting

June 12, London, England
ukoug.org



Head to San Antonio, Texas, for ODTUG Kscope12, June 24–28.

Gartner PPM and IT Governance Summit

May 21–23, National Harbor, Maryland

► bit.ly/z7w2xD

Learn about IT program and portfolio governance, enterprise program and portfolio management (PPM) issues beyond IT, organizational change management, agile project leadership, and PPM software and solutions at this gathering.

GOTO Conferences

May 21–25, Copenhagen, Denmark

May 24–26, Amsterdam, Netherlands

► gotocon.com

These conferences are designed for software developers, IT architects, and project managers in the Java, mobile, .Net, open source, Lean/agile, architecture, new languages, and process communities. The Copenhagen conference has three days of sessions and two training days; the Amsterdam conference has two days of sessions and one day of training.

IASA 2012 Educational Conference and Business Show

June 3–6, San Diego, California

► bit.ly/zzxP3J

Financial, technology, and operations professionals in the insurance industry can choose from more than 80 technical sessions, listen to speakers, and access demos.

Berlin Buzzwords 2012

June 4–5, Berlin, Germany

► berlinbuzzwords.de

This conference for developers and users of open source software projects focuses on the issues of scalable search, data analysis in the cloud,

and NoSQL databases. More than 30 presentations cover the topics of search, store, and scale.

Gartner IT Infrastructure and Operations Management Summit

June 5–7, Orlando, Florida

► bit.ly/wXwSdd

This summit offers strategic and tactical guidance for navigating critical infrastructure and operations trends and technologies. Topics include cloud-delivered services, virtualization, mobility, and managing data growth.

Cloud Expo

June 11–14, New York, New York

► cloudcomputingexpo.com

Attendees explore cloud and big data solutions and offerings through technical and strategic breakout sessions, general sessions, industry keynotes, expert panels, and an expo floor with two demo theaters.

Gartner Security and Risk Management Summit

June 11–14, National Harbor, Maryland

► bit.ly/ybBVXX

This summit for chief information security officers and security, risk management, and business continuity professionals covers governance, risk, and compliance; cloud computing and recovery; security architecture; mobile applications; and general security threats and vulnerabilities.

OHUG Global Conference

June 18–22, Las Vegas, Nevada

► bit.ly/xA4HAV

The Oracle HCM Users Group's annual gathering includes separate tracks for Oracle's PeopleSoft

applications and Oracle E-Business Suite as well as information on Oracle Fusion Applications and general human capital management (HCM) topics.

Software Systems Best Practices Conference

June 18–22, Anaheim, California

► softwaresystemsbestpractices.com

Learn about software best practices in requirements management, agile development and testing, quality assurance, inspections and review, test design and management, and project management and estimation.

DIA 2012 48th Annual Meeting

June 24–28, Philadelphia, Pennsylvania

► bit.ly/zZgHUo

The nonprofit Drug Information Association (DIA) sponsors this meeting for individuals involved in the development of healthcare products. Topics include the regulatory and public policy environment for pharmaceuticals in the U.S., regulations and economic factors affecting the industry globally, and clinical and statistical systems.

ODTUG Kscope12

June 24–28, San Antonio, Texas

► kscope12.com

The annual conference of the Oracle Development Tools User Group (ODTUG) features four days of in-depth technical sessions for developers, architects, and DBAs who work with Oracle Application Express, Oracle Database, Oracle Hyperion products, Oracle Essbase, Oracle business intelligence solutions, and Oracle Fusion Middleware.

JAZOON'12 International Conference on the Modern Art of Software

June 26–28, Zurich, Switzerland

► jazon.com

This event for Java experts, software developers, solution architects, IT decision-makers, and consultants includes an exhibition of solutions, tools, and continuing education; career opportunities; and networking events.

EVENTS LOCATOR

Oracle Events
oracle.com/events

Locate User Groups
bit.ly/yOFrOc

What's New at Oracle

The latest videos, podcasts, blogs, and more

WEBCASTS

“Big Data Essentials”

▶ oracle.com/goto/bigdata

View presentations from the Big Data Online Forum to get key facts on big data.

“Revolutionizing Economic Impact with Oracle Data Integrator”

▶ bit.ly/ym1MXd

See the potential financial impact of deploying Oracle Data Integrator in your organization.

“Oracle SOA Suite Customer Experiences and Product Overview”

▶ bit.ly/yEIFPP

Learn how customers have justified the migration to Oracle SOA Suite.

“How Oracle Uses Oracle Desktop Virtualization”

▶ bit.ly/xHKmpC

See how Oracle uses desktop virtualization to support developers around the world.

“Leveraging Linux Innovations for Extreme Performance with Oracle Exadata”

▶ bit.ly/Azte8q

Find out about the benefits of using Oracle Linux with Oracle Exadata.

“Simplify Oracle RAC Deployment with Oracle VM”

▶ bit.ly/zSLJzT

Learn how to use Oracle VM prebuilt templates to complete application deployments and updates.

“MySQL Embedded Online Forum”

▶ bit.ly/AscAVQ

See why MySQL is a great embedded database for startups as well as the largest companies.

“Developing a Data Access Layer with Java Persistence API”

▶ oracle.com/goto/weblogicdevcast

Find out what Java Persistence API offers over JDBC.

E-BOOK

Staying Ahead of the Storage Explosion

▶ bit.ly/w3w1OR

Learn how Oracle's storage systems can help you achieve better performance while reducing storage costs.

VIDEOS

I Can't Believe This is Butter!

▶ youtu.be/hxWuaozpe2I

Avi Miller, principal program manager at Oracle, takes you on a tour of the Btrfs file system.

Big Data Insight from Gartner

▶ bit.ly/xLJqAJ

Gartner's Merv Adrian discusses why big data is generating such interest and presents use cases.

Qualcomm Sees 400% Performance Gain from SPARC T4 Processors

▶ bit.ly/x3kvIK

Qualcomm looks to SPARC T4 processors and Oracle Solaris virtualization to solve pressing issues in its mission-critical applications.

Medtronic iPad App Enhances Sales Process with Oracle WebCenter

▶ bit.ly/zu6Pe4

See how Medtronic uses an Oracle WebCenter-powered Apple iPad solution for collateral.

Cloud Computing: An Overview

▶ bit.ly/y6SqzP

Get a handle on cloud computing lingo.

IDC Discusses Oracle Solaris and the UNIX Market

▶ bit.ly/zVsbKh

Matt Eastwood of International Data Corporation (IDC) provides insight on Oracle Solaris.

RESOURCE KITS

Oracle Learning Library

▶ oracle.com/goto/oll

Search for free online learning content, videos, tutorials, articles, demos, step-by-step

instructions, and self-paced interactive learning modules—now with mobile device support.

Data Integration Resource Kit

▶ bit.ly/zE7OVw

See how Oracle Data Integrator leverages the power of Oracle Exadata.

Oracle Linux Resource Kit

▶ bit.ly/yQdWPS

Explore Oracle's complete Linux solution, including one-stop support.

Oracle VM Resource Library

▶ bit.ly/zjwUdD

Find a wealth of information on Oracle VM.

Oracle Direct

For Oracle product information, call the number for your region.

COUNTRY	PHONE NUMBER
U.S. and Canada	1.800.367.8674
Australia	1800.735.467
Austria	0800.29.7626
Belgium	0800.73280
Brazil	0800.901.985
China	800.810.0161
Denmark	8088.1068
Finland	0800.113.573
France	0800.905.805
Germany	0800.1.810.111
Greece	00800.353.12020
Hong Kong	3002.1246
India	1600.44.6725
Indonesia	001.800.1.672.253
Ireland	1850.672253
Italy	8008.74720
Japan	0120.155.096
Malaysia	1800.80.1837
Mexico	01800.221.7321
The Netherlands	0800.0827
New Zealand	0508.555.215
Norway	800.14411
Philippines	811.5831
Portugal	800.853.021
Singapore	1800.6722.531
South Africa	0800.994.225
South Korea	080.2194.114
Spain	900.952900
Sweden	020.798798
Switzerland	0800.55.2574
Taiwan	0800.672.253
Thailand	001800.441.0545
Turkey	00800.353.90014
U.K.	0870.5.332200

OVERHEARD

“When we started testing the SPARC T4 servers, we noticed a huge performance increase . . . a 400 percent increase from our [servers] that we were using on our ERP [enterprise resource planning] systems. Whenever we need an application to be highly reliable, we run it on Oracle Solaris and the SPARC T4 platform.”

—**Matthew Clark**, Senior Director of IT at Qualcomm, in *Qualcomm Sees 400% Performance Gain from SPARC T4 Processors* (bit.ly/x3kv1K)

Oracle WebLogic 12c Resource Kit

▶ oracle.com/weblogic

Access white papers, Webcasts, and more about the release of Oracle WebLogic 12c.

Oracle Database Appliance Resource Kit

▶ bit.ly/z4rUIs

Learn more about Oracle Database Appliance through customer case studies, analyst reports, Webcasts, and demos.

**WHITE PAPERS****“Oracle’s Netra SPARC T4-1 and T4-2 Server Architecture”**

▶ bit.ly/ywSNoq

Read about the new levels of performance that Netra SPARC T4-1 and T4-2 servers offer.

“Oracle Solaris and SPARC T4 Servers—Engineered Together for Enterprise Cloud Deployments”

▶ bit.ly/v09KTU

Learn how the Oracle Solaris/SPARC T4 combination can maximize ROI.

“Guidebook: Oracle Cloud Services”

▶ bit.ly/zzOABZ

This Nucleus Research report explores the benefits of using Oracle’s managed cloud services.

“Oracle StorageTek Tape Storage—A Clear Winner Over IBM”

▶ bit.ly/yro5me

Learn how Oracle’s StorageTek tape offerings outperform IBM’s tape technology.

**CALCULATORS/TOOLS****Database Assessment**

▶ bit.ly/wVSYEf

Take this five-minute survey to gain a better

understanding of your organization’s ability to ensure highly available database services.

Oracle Desktop Virtualization TCO

▶ oracle.com/us/media/calculator/vdi

See how much you can save with Oracle desktop virtualization solutions.

Linux Cost Calculator

▶ bit.ly/mX94rf

Compare the cost of Oracle Linux and Red Hat Enterprise Linux.

Oracle VM Cost Calculator

▶ bit.ly/xfteru

Compare the cost of Oracle VM with VMware vSphere.

**PODCASTS****“Using PHP with MySQL”**

▶ bit.ly/xYCeno

Discover the ways PHP developers can use MySQL to optimize queries.

“EnerNOC Tackles Energy Data Volume Growth with Oracle Exadata”

▶ bit.ly/wtPqKr

Hear how EnerNOC realized a five-times boost in performance with Oracle Exadata.

“NextGen Deploys Oracle Database Appliance in 2.5 Hours”

▶ bit.ly/zSVxLF

Hear about NextGen’s experience deploying Oracle Database Appliance.

“Why Oracle Developed Network Virtualization”

▶ bit.ly/zsB4Fe

Nicholas Droux of the Oracle Solaris core engi-

neering team discusses the development of virtual network interface cards.

“Oracle Solaris 11: Upgrade Your Skills”

▶ bit.ly/xVb5B2

Discover the different ways to upgrade your skills for Oracle Solaris 11.

**BLOGS****BI & Analytics Pulse**

▶ blogs.oracle.com/bi_pulse

Get the latest on the business intelligence market.

**EDUCATION/TRAINING****Oracle Fusion Applications Training**

▶ bit.ly/xHaCPA

Get up to speed quickly on Oracle Fusion Applications with Oracle University training.

Free Sample Live Virtual Class Lesson

▶ bit.ly/z85dJ5

Test-drive Oracle University’s top business intelligence training course with this free lesson.

Oracle Training On Demand

▶ bit.ly/y8BjHo

Get the training you need without waiting for scheduled classes. New courses include Oracle Database: Introduction to SQL, Oracle Database 11g: Administration Workshop I, Oracle Database 11g: Administration Workshop II, and Transition to Oracle Solaris 11.

WEB LOCATOR**Oracle Blog Center**

oracle.com/blogs

Oracle Consulting

oracle.com/consulting

Oracle Events and Webcasts

oracle.com/events

Oracle Podcast Center

oracle.com/podcasts

Oracle Newsletters

oracle.com/newsletters

Oracle University

oracle.com/education

Oracle on Twitter

twitter.com/oracle

Oracle Support

oracle.com/support

My Oracle Support

myoraclesupport.com

My Oracle Support Communities

communities.oracle.com

ORACLE PRESS

YOUR DESTINATION FOR ORACLE AND JAVA EXPERTISE

Written by leading technology professionals, Oracle Press books offer the most definitive, complete, and up-to-date coverage of Oracle products and technologies available.




**Oracle Database 11g
Release 2 Performance
Tuning Tips & Techniques**
Richard Niemiec

Systematically identify and eliminate database performance problems with help from Oracle Certified Master, Oracle ACE Director, Richard Niemiec.

Oracle Exalogic Elastic Cloud Handbook
Tom Plunkett, TJ Palazzolo, and Tejas Joshi
Plan and deploy a reliable, secure, highly available cloud solution.

Effective MySQL: Backup and Recovery
Ronald Bradford
Perform, test, and verify backup and disaster recovery procedures.



Available in print and eBook formats
Join the Oracle Press Community: www.OraclePressBooks.com

 Follow us on Twitter @OraclePress



DEMOS**MySQL Cluster**

In this demo, you'll learn about the capabilities of MySQL Cluster, including write-scalability, real-time performance and 99.999 percent high availability, and the benefits of deploying MySQL Cluster Carrier Grade Edition.

► bit.ly/wqAoyP

Extending Applications with Oracle Application Express

This presentation introduces the Oracle Application Express development environment and demonstrates how to add a calendar page to an existing application in just a few minutes with the help of a wizard and drag-and-drop functionality.

► bit.ly/y8VPjt

TUTORIALS**How to Use an Oracle Service Bus Customization File**

See how to use Oracle Service Bus 11g customization files to change an environment value when propagating a configuration from one domain to another.

► bit.ly/zCEGUt

Oracle Enterprise Manager Cloud Control 12c: Patch an Oracle WebLogic Domain

Learn how to apply multiple patches across all servers within an Oracle WebLogic domain.

► bit.ly/wklEzA

Advanced Oracle ADF Series, Part 1

This self-paced e-course is the first in a series that addresses Oracle Application Development Framework (Oracle ADF) 11g topics through video presentations, quizzes, and practices. You'll find tips, tricks, and best practices from Oracle experts.

► bit.ly/ygLDXP

Getting Started with Oracle SQL Developer 3.0

This tutorial introduces Oracle SQL Developer 3.0 and shows you how to manage your database objects.

► bit.ly/A31AZO

DOWNLOADS**New Downloads**

Oracle SQL Developer 3.1

► bit.ly/wbYNb8

Oracle R Enterprise

► bit.ly/wuuwwd

Oracle Offers Free Trial of Ksplice Zero-Downtime Updates

Oracle has introduced a 30-day free trial of Oracle's Ksplice zero-downtime kernel update technology for Red Hat Enterprise Linux users. The free trial offers organizations the benefit of Ksplice zero-downtime kernel updates for the kernels they have installed today, without any system reconfiguration.

In addition, Ksplice zero-downtime kernel updates for Oracle's Red Hat Compatible Kernel are now available to Oracle Linux Premier Support customers.

"Oracle has been focused on offering the best Linux for enterprise computing, both from a product and support perspective," says Wim Coekaerts, senior vice president of Linux and virtualization engineering at Oracle. "With the innovative zero-downtime update capabilities delivered through Ksplice and the extended support lifecycle for Oracle Linux, Oracle continues to set the industry standard for Linux in the enterprise."

► ksplce.com/rhel-signup

New Mobile Apps for Oracle's JD Edwards EnterpriseOne

Oracle has unveiled new mobile applications that help organizations using Oracle's JD Edwards EnterpriseOne to increase the productivity of employees working in the field.

JD Edwards EnterpriseOne is an integrated applications suite of enterprise resource planning software. With the new mobile applications, JD Edwards EnterpriseOne users can gain real-time insight into sales orders, item availability, and item-based pricing, along with the ability to easily review and approve requisition

self-service orders and purchase orders on mobile devices. The new mobile applications include JD Edwards EnterpriseOne Mobile Requisition Self Service Approval, JD Edwards EnterpriseOne Mobile Purchase Order Approval, and JD Edwards EnterpriseOne Mobile Sales Inquiry.

"Our customers demand real-time access to important business data, regardless of time or location," says Lyle Ekdahl, group vice president and general manager of Oracle's JD Edwards product family. "With the mobile apps, users can closely manage approvals, inventory, and sales orders in the office, in the factory, or on the road."

► bit.ly/xsm7UX

**Oracle VM Server for SPARC 2.2 Released**

Oracle VM Server for SPARC 2.2, the latest release of Oracle's enterprise-class server virtualization solution, is now available. The new release delivers significant enhancements that improve workload agility and performance, maximize the availability of business-critical applications, and increase flexibility in provisioning and deployment. Preinstalled on Oracle's SPARC T-Series servers, Oracle VM Server for SPARC takes advantage of the massive thread scale offered by SPARC T-Series servers and the Oracle Solaris operating system.

New features in Oracle VM Server for SPARC 2.2 allow organizations to benefit from more-flexible control and deployment of server resources as well as workload performance optimization and agility across the SPARC T-Series server family, including SPARC T3 and SPARC T4 servers. These

features include cross-CPU live migration for uninterrupted application availability; single-root I/O virtualization; dynamic performance tuning for business-critical applications; dynamic threading configuration support for SPARC T4 servers; enhanced resource management for SPARC T4-4 engineered systems; and integrated management via Oracle Enterprise Manager Ops Center.


"The latest release of Oracle VM Server for SPARC provides higher performance, greater scalability, and an increased ability to more finely tune the system, making Oracle VM for SPARC, and T-Series servers in general, an excellent platform for an even broader range of enterprise applications," says Wim Coekaerts, senior vice president of Linux and virtualization engineering at Oracle.

► bit.ly/qHLZQN

Oracle Enterprise Healthcare Analytics Enhanced

Addressing healthcare's growing need for actionable business intelligence, Oracle has updated its enterprise healthcare analytics solution and unveiled a new release of Oracle Healthcare Master Person Index.

Oracle enterprise healthcare analytics makes it easier, faster, and more cost effective for healthcare organizations to integrate and analyze clinical, financial, and research



data. It delivers the information they need to drive efficiencies, support evidence-based medicine initiatives, and improve outcomes.

Oracle Healthcare Master Person Index 2.0 provides a single point of reference for information about a patient, provider, or other healthcare entity within and across healthcare organizations. It delivers solutions for patient and provider indexing; enables faster, more flexible deployment; and supports Simplified Chinese characters.

"The need for actionable business intelligence is no longer a 'nice-to-have,' but is now mission-critical as healthcare organizations work to improve outcomes, cut costs, and enable more-accountable care," says Neil de Crescenzo, senior vice president and general manager, Oracle Health Sciences.

bit.ly/xuqSX9

Oracle User Productivity Kit Updated

Oracle User Productivity Kit 11.1 and Oracle User Productivity Kit Professional 11.1 are now available. These content development, deployment, and maintenance platforms streamline the implementation of Oracle, third-party, and custom applications.

Both new releases enable organizations to expand content contribution and collaboration. Testing time is also decreased by

streamlining the capturing and compiling of test step results.

"Providing every user with the ability to easily and rapidly create and share content is exactly what our customers are looking for to enhance collaboration," says Trish Trolley, senior director of Oracle User Productivity Kit product management.

bit.ly/weZbB0

Oracle Launches New Release of Business Intelligence Foundation Suite

The latest release of Oracle Business Intelligence Foundation Suite features an extensive set of capabilities for parallel performance and in-memory analytics that deliver real-time, speed-of-thought visual analysis with subsecond responsiveness. Available now for use in conjunction with

Oracle Ealytics without application changes, the new release features new visualizations, contextual interactions; performance improvements; optimizations for Oracle Ealytics; and simplifications for end users, developers, and administrators.

bit.ly/ydffNyt

New Oracle Airline Data Model Personalizes Travel Experience

Now available, Oracle Airline Data Model is an industry-specific database schema that helps airlines consolidate and analyze passenger data. An option for Oracle Database, Enterprise Edition 11g, Oracle Airline Data Model gives airlines quick insights into passenger and business information to improve customer service and grow their business.

"Airlines need to be able to quickly and easily collect and analyze passenger data in order to deliver the best-possible customer experience," says Cetin Ozbutun, vice

president of data warehouse technologies at Oracle. "Oracle Airline Data Model makes this easier by delivering a comprehensive passenger data model with prebuilt industry-specific key performance indica-



tors, online analytical processing cubes for multidimensional analysis, and data mining for predictive analytics."

bit.ly/AvJaYI

MySQL Cluster 7.2 Unveiled

Now available, MySQL Cluster 7.2 delivers up to a 70-times increase in performance on complex queries and enhanced multi-data center scalability. MySQL Cluster 7.2 allows users to deploy a distributed, highly scalable database with both SQL and NoSQL interfaces and perform complex queries or multitable transactions with ACID guarantees. With SQL and NoSQL access through a new memcached API, MySQL Cluster represents a "best of both worlds" solution that allows key value operations and complex SQL queries within the same database.

In addition, MySQL Cluster 7.2 is Oracle VM certified. The combination of MySQL Cluster's elastic, on-demand scalability and self-healing features and Oracle VM support makes MySQL Cluster an appropriate choice for deployments in the cloud.

"The performance and flexibility enhancements in MySQL Cluster 7.2 provide users with a solid foundation for their mission-critical Web workloads, blending the best of SQL and NoSQL technologies to reduce risk, cost, and complexity," says Tomas Ulin, vice president of MySQL engineering at Oracle.

mysql.com/downloads

Oracle Virtual Assembly Builder 11g Released

Oracle Virtual Assembly Builder 11g, now available, helps organizations quickly create and configure entire multitier application topologies and provision them onto virtualized resources. It structures the process of combining software appliances into cohesive, reusable units known as assemblies by making the necessary connections between these appliances and then deploys the entire assembly as a single unit. The product works with Oracle VM 3.0 and Oracle Fusion Middleware 11g with Oracle WebLogic Server.

"Oracle Virtual Assembly Builder accelerates customer efforts toward setting up their own private platform-as-a-service environments by enabling them to rapidly configure and deploy complex multitier applications and middleware platforms in a virtualized resource pool," says Brad Cameron, vice president of product development at Oracle.

bit.ly/b0MMyxv

Oracle Health Sciences Institute Announces First Round of Research Projects

Oracle Health Sciences Institute, in partnership with Oracle Labs, has announced its first group of research projects. Oracle Health Sciences Institute focuses on supporting research that accelerates IT innovation to advance personalized medicine and the delivery of safe and effective treatments and healthcare services to patients around the globe.

The research projects include initiatives with Brigham and Women's Hospital and Harvard Medical School, the



Coalition Against Major Diseases at the Critical Path Institute, and the Human-Computer Interaction Lab at the University of Maryland.

"IT innovation fuels the transformation needed to realize the full potential of the convergence of healthcare and life sciences," says Neil de Crescenzo, senior vice president and general manager, Oracle Health Sciences. "We are pleased to support esteemed research project teams as they seek new ways to apply information technology to advance the discovery and development of next-generation treatments and therapies."

▶ bit.ly/xuqSX9

Unbreakable Enterprise Kernel Release 2 for Oracle Linux Available

Unbreakable Enterprise Kernel Release 2 for Oracle Linux is now generally available. Based on the mainline Linux kernel 3.0.16, the latest release of Unbreakable Enterprise Kernel for Oracle Linux includes performance and scalability enhancements, improved memory and resource management, and greater virtualization support. It also includes one of the first commercially supported production Btrfs, the "next-generation file system" for Linux. With Btrfs included, Oracle Linux supports large files and file systems, snapshots, and checksums for data and metadata; provides integrated RAID and volume management; and simplifies administration.

Unbreakable Enterprise Kernel Release 2 for Oracle Linux is included with Oracle Linux 5 and Oracle Linux 6.

"Oracle Linux continues to deliver timely Linux innovations, backed by real-world testing, providing users a modern, scalable, and reliable platform for their business-critical workload demands," says Wim Coekaerts, senior vice president of Linux and virtualization engineering at Oracle. "Unbreakable Enterprise Kernel Release 2 further demonstrates Oracle's investment in Linux technology and the community and delivers on our goal to make Linux better for everyone."

▶ oracle.com/linux

New Enhancements for Oracle's Sun ZFS Storage Appliance

Oracle has introduced new capabilities for its Sun ZFS Storage Appliance that deliver better performance, simpler management, and higher scalability.

The Sun ZFS Storage Appliance now includes Oracle VM Storage Connect Plug-in for Sun ZFS Storage Appliance, which allows users to seamlessly provision and manage Sun ZFS Storage Appliance logical unit numbers and execute advanced data-protection features. For scalability, Sun ZFS Storage Appliances now have the option of 3 TB hard-disk drives, which allow the Sun ZFS Storage Appliance 7420 to scale to more than 1.7 PB.

"Oracle's Sun ZFS Storage Appliances offer a compelling proposition for customers looking to extend deployment of network-attached storage [NAS] for critical enterprise applications," says Phil Bullinger, senior vice president of storage at Oracle. "With the new enhancements, Oracle is delivering greater value to customers and extending its storage leadership in high-performance, efficient NAS storage."

▶ bit.ly/x5J4ij



Oracle Enterprise Manager Ops Center 12c Unveiled

Now available, Oracle Enterprise Manager Ops Center 12c introduces new capabilities for establishing, managing, and supporting enterprise clouds that enable Oracle systems and server virtualization users to adopt cloud computing faster and with less effort and risk. Oracle Enterprise Manager Ops Center 12c delivers increased server management; support for Oracle VM x86, Oracle Solaris 11, and existing virtualized Oracle Solaris Zones environments; and increased Sun ZFS Storage Appliance management.

"With Oracle Enterprise Manager Ops Center 12c, we have extended our leadership by providing our customers with the best management solution for all of Oracle's systems at no additional license charge," says John Fowler, executive vice president of systems at Oracle. "Oracle Enterprise Manager Ops Center 12c is the only solution our customers need to manage Oracle systems, whether for traditional enterprise IT, private cloud, public cloud, or a hybrid environment."

▶ oracle.com/enterprisemanager

Oracle Acquires Taleo

Oracle has entered into an agreement to acquire Taleo, a leading provider of cloud-based talent management solutions. Taleo's talent management cloud helps organizations attract, develop, motivate, and retain human capital to improve performance and drive growth.

Together, Oracle and Taleo plan to create a comprehensive cloud offering for organizations to manage their human resources operations and employee careers. The combination is expected to empower employees and managers to effectively manage careers throughout their entire employment, enable organizations to retain talent and optimize costs, and improve the employee experience through faster onboarding and better collaboration with team members via social media.

"Human capital management has become a strategic initiative for organizations," says Thomas Kurian, executive vice president of product development at Oracle. "Taleo's industry-leading talent management cloud is an important addition to Oracle Public Cloud."

▶ oracle.com/taleo

RightNow Adds Service to Oracle CX

Oracle unveils product roadmap for end-to-end customer experience.

With the January 2012 acquisition of RightNow, Oracle added the industry's leading cloud-based customer service solution to Oracle Public Cloud, a broad set of best-in-class, integrated services that provide customers with subscription-based, self-service access to Oracle Fusion Applications, Oracle Fusion Middleware, Oracle Database, and more—managed, hosted, and supported by Oracle.

The RightNow solutions join Oracle's already-wide array of customer experience (CX) products—including Web presence, commerce, sales, marketing, search, knowledge, and social networking solutions—to make Oracle's CX offerings "the best, most complete, and most advanced solution in the world," said Oracle President Mark Hurd.

Hurd joined Greg Gianforte, RightNow CEO and founder; Thomas Kurian, executive vice president of product development at Oracle; and Wayne Huyard, senior vice president and general manager of Oracle's new RightNow Global Business Unit, for a January 31 Webcast introducing the RightNow technology and previewing Oracle's CX product roadmap. Oracle RightNow CX Cloud Service delivers "what we consider to be the best customer services technology in the market," Hurd continued. "We have a very aggressive plan to invest in RightNow and build a stronger roadmap than ever before."

THE CX IMPERATIVE

With 86 percent of consumers choosing not to do business with a company because of a bad experience, according to RightNow's annual survey, "CX is today's primary business imperative," said Gianforte following Hurd's introduction. "Not embracing customer experience as a core strategy will leave you unable to compete, whether you sell to consumers

"Not embracing customer experience . . . will leave you unable to compete."

—Greg Gianforte, CEO and Founder, RightNow

or other businesses." RightNow has spent the last 14 years building a cloud-based solution that addresses this CX imperative, he added, with a suite that "empowers organizations to engage directly with customers through great social, Web, and contact center experiences." The benefits, said Gianforte, typically include a 10 to 30 percent reduction in contact-center labor costs, increased customer retention, and accelerated customer acquisition.

CX doesn't begin and end with customer services, Kurian emphasized, turning to the

solution roadmap. "Our fundamental vision is that consumers and business customers move through a closed-loop continuous lifecycle," he explained, from when they perceive need for a product or service through the research, selection, purchase, receive, use, maintain, and recommend stages. "Customers require a simple, consistent, and amazing experience throughout the entire lifecycle, in all the channels through which they perceive the products, solutions, and services you are offering them," said Kurian.

Kurian pointed out that many Oracle products already drive and support the customer lifecycle from needs assessment through the research, select, purchase, and receive phases, including Oracle's Fatwire products (now part of the Oracle WebCenter family); Oracle Social Network; and Oracle's Siebel Marketing, Oracle's Endeca solutions, Oracle ATG Web Commerce, and Oracle Fusion Applications. Oracle RightNow CX Cloud Service completes the process with world-leading capabilities for managing the use, maintain, and recommend phases of the lifecycle, he said.

"Oracle and RightNow are a perfect pair," Huyard concluded, providing "the unique capability of serving the entire customer lifecycle, from beginning to end, with understanding, integration, insights, and full-featured capability." ◀



Thomas Kurian, Executive Vice President of Product Development, Oracle

NEXT STEPS

WATCH the event Webcast
bit.ly/Am5KR4

LEARN more about
Oracle Public Cloud
cloud.oracle.com

RightNow
oracle.com/rightnow



Partner **PRODUCT** WATCH

LoadSpring

Project Portfolio Management Leader

LoadSpring is the solution-centric destination for project management software, helping PMs make business decisions easier. As the leader in hosting and implementation services, we make it easy to select and deploy software. We provide trouble-free access, through our SpringBoard™ portal to projects while offering high levels of service, security, and support. Watch the interview with LoadSpring CEO Eric Leighton from Oracle OpenWorld 2011.

<http://www.bit.ly/ttb9CR> or scan for more info



HiT Software®, Inc.

DBMoto® - Data Replication and Change Data Capture made easy

DBMoto is an efficient, cost-effective software product to extract, migrate and share data among multiple database systems including Oracle, MySQL, IBM DB2 (AS/400, Linux, UNIX, Windows and mainframe), Microsoft SQL Server, Informix, Sybase and others.

Performs automatic, incremental data updates on either batch or real-time scheduling.

Performs data mapping between different data types, and provides graphical data comparison tool to view “before and after” data replication activity. Sets up in minutes, no consulting required. Find out why so many companies around the world depend on DBMoto for tough data integration and data replication jobs! Download Free Trial and White Papers Today!

<http://www.hitsw.com/LP26> or scan for more info

Quest Software

Simplicity and Productivity for DB Professionals

Millions of users trust Toad to provide a simple, consistent way to build, manage, and maintain databases. Whether you're a developer, DBA, or analyst, Toad can dramatically increase your productivity.

<http://www.quest.com/SimplyBetterToad11> or scan for more info



Frontline Consulting and SurfBI

Frontline Consulting is an award-winning Oracle Platinum Specialized Partner offering worldclass solutions on Oracle Business Intelligence Applications, Oracle Business Intelligence Enterprise Edition, Oracle E-Business Suite, and Oracle Fusion Middleware. SurfBI by Enterprise Signal Inc. is an award-winning mobile business intelligence product available natively on multiple platforms. Download by searching “surfbi” on Appstore/Google marketplace.

<http://www.bit.ly/ujPK5d> or scan for more info

To advertise your solutions in Partner Product Watch, contact tom.cometa@oracle.com.

Upon Further Analysis

Oracle Advanced Analytics delivers key insights from big data.

Following the announcement of Oracle Advanced Analytics, Rich Schwerin, Oracle Magazine contributor, sat down with George Lumpkin, vice president of product management for data warehousing at Oracle, to talk about this new option for Oracle Database. The following is an excerpt from that interview. Download the full podcast at oracle.com/magcasts.

Oracle Magazine: How does Oracle define analytics, and why are they so important in relation to big data?

Lumpkin: Today, there is a lot of attention on big data and its potential business benefits. There are lots of new types of data out there that weren't available years ago. To recognize the full potential of this data, organizations have to use deeper analytical techniques, and they need to do more-sophisticated analysis. When Oracle talks about analytics in the context of big data, we really are talking about these more-advanced types of analytic techniques—predictive analytics and data mining, deep statistical analysis, advanced numerical calculations, integrating text, and text mining.

Oracle Magazine: What is Oracle Advanced Analytics?

Lumpkin: Oracle Advanced Analytics is a new option for Oracle Database. It provides analytic capabilities that run in a highly scalable way inside Oracle Database, and it consists of Oracle Data Mining and Oracle R Enterprise, Oracle's implementation of R—an open source programming language and software environment for statistical computing and graphics.

Oracle Magazine: Who uses Oracle Data Mining, and how do they use it?

Lumpkin: Oracle Data Mining is used by analysts and advanced business users, and its most common use is in building predictive models. Oracle Data Mining includes a dozen advanced data mining algorithms



George Lumpkin, Vice President of Product Management for Data Warehousing, Oracle

implemented inside Oracle Database. These algorithms run in parallel in a scalable fashion and allow data analysts to build predictive models. Data mining enables the use of advanced algorithms and the creation of models inside the database that can help characterize the data and provide new insights from the data. In addition to these algorithms, Oracle Data Mining provides a graphical user interface where end users can explore data, build data models, and more.

Oracle Magazine: Who uses Oracle R Enterprise, and how do they use it?


Lumpkin: With Oracle R Enterprise, we allow statisticians to continue to use R, but we execute statistical analysis on the database tier against tables in the database, running statistics in a scalable fashion against millions to billions of records. We also make it much simpler to take advanced statistical programs and make them part of an enterprise production deployment.

Oracle Magazine: Who uses Oracle Advanced Analytics, and how do they benefit from it?

Lumpkin: Oracle Advanced Analytics users are data analysts, statisticians, and data scientists. Oracle provides analytic capabilities in the context of an enterprise database, so users who have lots of data in Oracle data warehouses are now able to get more value from their data through Oracle Advanced Analytics in a way that fits with their current infrastructures.

Oracle Magazine: How does Oracle Advanced Analytics fit in with Oracle Big Data Appliance and Oracle's big data solutions?

Lumpkin: To solve big data problems, you need a complete big data architecture, and Oracle Big Data Appliance is one piece of that. Oracle Big Data Appliance is optimized to run Hadoop and Oracle NoSQL Database and is used in conjunction with Oracle Exadata or Oracle Database, and also in conjunction with Oracle Exalytics and Oracle Business Intelligence solutions. Oracle Advanced Analytics is also part of this ecosystem. Organizations need to be able to do more-sophisticated predictive modeling—targeted marketing, fraud analysis, churn models. Oracle needs to have that as part of our big data ecosystem, so we have Oracle Advanced Analytics built into Oracle Database, which works seamlessly with Oracle Big Data Appliance for an overall big data solution. ◀



Rich Schwerin is a senior manager with Oracle Publishing who focuses on social media.

NEXT STEPS

LEARN more about Oracle Advanced Analytics
bit.ly/oaamag

LISTEN to the podcast
oracle.com/magcasts

Beginning Database Design, Second Edition



By Clare Churcher
Apress
apress.com

Beginning Database Design, Second Edition, provides short, easy-to-read explanations of how to get data-

base design right the first time. It is aimed at novice designers, especially those who are doing ad hoc design as a result of employing end-user tools such as Oracle Application Express, Microsoft Access, and Microsoft Excel. The author offers numerous examples to help readers avoid the many pitfalls faced by new and not-so-new database designers.

The book helps designers ask and answer important questions about their data so they can understand the problems they are trying to solve and create a pragmatic design capturing the essential—while leaving the door open for refinements and extension at a later stage.

Oracle Database 11g Release 2 Performance Tuning Tips and Techniques



By Richard Niemiec
Oracle Press
oraclepressbooks.com

Oracle Certified Master Richard Niemiec shows you how to systematically identify and eliminate data-

base performance problems in this updated Oracle Press book. Filled with real-world case studies and best practices, the book details the latest monitoring, troubleshooting, and optimization methods. Learn how to find and fix bottlenecks, configure storage devices, execute effective queries, and develop bug-free SQL and PL/SQL code. Testing, reporting, and security enhancements are also covered. In addition, Niemiec discusses how to properly index and partition Oracle Database 11g Release 2; work with Oracle Exadata and Oracle Exalogic Elastic Cloud; and efficiently manage disk drives, RAID arrays, and memory.

Agile Oracle Application Express



By Patrick Cimolini and Karen Cannell
Apress
apress.com

Agile Oracle Application Express shows how skilled, motivated, and self-

organizing developers can realize the commercial benefits of Oracle Application Express by coupling it with an agile software development approach to iteratively design,

create, and deliver quality applications on time and within budget. You will learn why Oracle Application Express is well suited to agile processes, with its support for rapid prototyping and team development. The authors also discuss how support for enabling technologies such as SQL, HTML, CSS, and JavaScript in Oracle Application Express allows you to deliver any type of Web application to meet your development needs.

Oracle BAM 11g R1 Handbook



By Pete Wang
Packt
packtpub.com

Oracle BAM 11g R1 Handbook is a practical best-practices tutorial focused entirely on Oracle Business Activity

Monitoring, which enables business service and process monitoring through real-time data streaming and operational reports. Packed with best practices, the book covers business activity monitoring migration, high-availability configuration, and more. Wang dives straight into the architecture of Oracle Business Activity Monitoring 11g before moving on to managing business activity, monitoring server securities, populating data objects, and performing load testing, as well as building an Oracle Application Development Framework-based report. This intermediate-to-advanced guide includes step-by-step instructions and an accompanying demo project that will guide SOA report developers through application development and producing dashboards and reports.

PeopleSoft for the Oracle DBA, Second Edition



By David Kurtz
Apress
apress.com

PeopleSoft for the Oracle DBA, Second Edition, is written for DBAs charged with supporting an Oracle's

PeopleSoft system, and especially those responsible for troubleshooting and resolving PeopleSoft performance issues. In this revised edition, readers will learn about the internet architecture in PeopleSoft applications and its use of Oracle Tuxedo. The book includes full coverage of key database issues such as indexing, connectivity, and tablespace usage as they apply to PeopleSoft applications. The author also provides advice on managing and troubleshooting performance issues in a PeopleSoft environment, including how to pinpoint problems within the code.

Look for other Oracle books at oracle.com/technetwork/community/bookstore

Oracle PartnerNetwork Adds MySQL Specialization

Reflecting its commitment to MySQL innovation, Oracle has launched a MySQL specialization as part of the Oracle PartnerNetwork Specialized program. More than 3,000 independent software vendors, original equipment manufacturers, and value-added resellers already use MySQL in their solutions; those that earn MySQL specialization can spotlight their strengths and expertise, differentiate their offerings, and gain a competitive edge.

Under MySQL specialization, Oracle partners reselling MySQL solutions receive special branding, advanced educational and enablement resources, and easier and faster ordering and reselling. Partners that embed or bundle MySQL with software, hardware, or appliances can access the development license of the MySQL Enterprise Monitor and Query Analyzer features as well as other tools to improve their products' quality, stability, and performance.

▶ bit.ly/zHzE3D

Accenture Garners Oracle Life Sciences Data Hub Specialization

Accenture is the first Oracle Diamond Partner to achieve Specialized status for Oracle Life Sciences Data Hub. The specialization underscores Accenture's ability



to help life sciences companies increase traceability and visualization of clinical data, conduct cross-trial analyses, and improve data analysis. Accenture and Oracle have worked together on successful Oracle Life Sciences Data Hub implementations for five years.

Oracle Life Sciences Data Hub integrates clinical and nonclinical data from multiple sources into a single environment for analysis, visualization, and reporting, providing better insights for more-informed decision-making, improved regulatory compliance, and greater productivity. To achieve Specialized status, Accenture met requirements based on the needs and priorities of the customer and partner communities, including completing online sales, presales, and support assessments and passing an implementation exam.

▶ accenture.com

Oracle PartnerNetwork Launches Oracle Applications and Industry Partner Summit

Oracle PartnerNetwork hosted the first Oracle Applications and Industry Partner Summit January 9–11, 2012, in Scottsdale, Arizona. More than 350 senior sales and practice leaders from 80 leading systems integrators in North America were on hand for presentations from Oracle customers, partners, and executives.

Keith Block, executive vice president of Oracle North America, shared Oracle's strategy and momentum and gave his perspective on the competitive landscape. Thomas Kurian, executive vice president of product development at Oracle, shared roadmaps for Oracle Applications and discussed how partners can capitalize on Oracle Fusion Applications and Oracle's cloud offerings.

Sonny Singh, senior vice president of Oracle's Industries Business Unit, outlined the company's go-to-market strategy, detailed collaboration opportunities, and offered best practices for industry-specific selling.

Attendees learned that more than 47 application- or industry-specific specializations are now available to partners; 1,800 partners have earned nearly 6,000 Oracle Applications or industry specializations; more than 59,000 Application and Industry Boot Camp Training Days have been conducted, including 12,200 in FY 2012 alone; and more than 11,000 certified implementation specialists have achieved application- or industry-specific accreditations.

▶ oracle.com/us/corporate/press/1489378

Acolade Achieves Oracle Enterprise Manager Application Quality Management Specialization

Acolade Consulting, an Oracle Gold Partner based in Sydney, Australia, has achieved Oracle Enterprise Manager Application Quality Management specialization. Acolade is proficient in providing Oracle's application quality management solutions to organizations for use in thoroughly testing their packaged and Web- and SOA-based applications to ensure the quality, scalability, and availability of the applications and their underlying infrastructure.

Oracle Enterprise Manager application quality management solutions consist of three distinct offerings. Oracle Application Testing Suite includes Oracle Test Manager for test management; Oracle Functional Testing for automation testing; and Oracle Load Testing for stress, volume, and per-

formance testing. Oracle Real Application Testing is a change assurance solution that enables businesses to fully assess the outcome of a system change on Oracle Database in a test environment. Oracle Data Masking Pack maintains a high level of data security by scrambling test data and yet maintaining end-to-end data integrity.

Together these solutions enable organizations to manage their entire quality assurance process, including defining test requirements, creating and executing automated and manual test cases, tracking defects, and performing load tests and database tests. This specialization demonstrates Acolade's high level of qualification and project management experience.

▶ acolade.com.au/acolade/partnership.php

ProcessWeaver Achieves Validated Integration

Oracle Gold Partner ProcessWeaver's Enterprise Centralized Shipping (ECS) 3.0 solution has achieved Oracle Validated Integration with Oracle E-Business Suite 12.1.

This solution from ProcessWeaver is used to process shipments, update freight cost and shipment tracking information, and print



shipping documentation such as carrier-compliant labels. The product works with multiple carriers (such as FedEx, UPS, DHL, and USPS) and freight forwarders. Integration of ProcessWeaver's ECS solution with Oracle E-Business Suite increases productivity and improves customer service; validated integration ensures that the two products work together as designed.

▶ processweaver.com

Intel Sets Benchmark Records with Oracle Systems

Oracle Gold Partner Intel and Oracle recently achieved two benchmark records based on the two companies' joint engineering collaboration.

The world's fastest x86 online transaction processing database performance on the TPC-C benchmark was delivered by Oracle's Sun Fire X4800 M2 server based on Intel Xeon E7 processors running Oracle Database, Enterprise Edition 11g with Oracle Linux.¹



The SPECjEnterprise2010 world-record benchmark for Oracle Fusion Middleware 11g was set on a Sun Blade X6270 M2 server module based on Intel Xeon 5600 processors.²

Intel technology is also the basis for several of Oracle's engineered systems:

Oracle Exalytics In-Memory Machine, the first engineered system for analytics, is based on the Intel Xeon E7 processor family.

Oracle Big Data Appliance, an optimized solution for acquiring, organizing, and loading unstructured data into Oracle Database 11g, is based on Intel Xeon 5600 processors.

Oracle Database Appliance, an easy-to-deploy system engineered for midmarket customers and departmental database instances, is based on Intel architecture.

▶ intel.com

ReadSoft PROCESSIT Achieves Several Oracle Ready Statuses

Oracle Platinum Partner ReadSoft AB has earned Oracle Linux Ready, Oracle Database Ready, and Oracle WebLogic Ready designations for PROCESSIT 7.1. ReadSoft software and services address document automation needs, including data entry, information verification, document workflows, and e-invoicing.

ReadSoft has R&D experts dedicated to Oracle-related consulting and development and makes those experts available to customers during implementation of ReadSoft solutions. The company is headquartered in Sweden, operates in 16 countries, and has partners in an additional 70. Some 5,000 customers use ReadSoft software.

▶ readsoft.com

Infosys Adds Oracle Exadata Specialization

Oracle Diamond Partner Infosys has achieved Oracle Exadata specialization and has installed an Oracle Exadata engineered system in the newly opened Shanghai, China, Infosys-Oracle Innovation Center. The Infosys-Oracle Innovation Center located at Oracle headquarters in Redwood Shores, California, also has access to the Oracle Exadata system and the ability to provide



customer demonstrations. To gain Oracle Exadata specialization, Infosys completed Oracle Exadata certification criteria across all categories (sales, pre-sales, implementation, and support).

Infosys is also involved in high-performance testing

for big data and has developed accelerators for data warehouse migration. The company's Value Complete data warehouse migration solution delivers extreme performance and scalability, while significantly reducing costs, time, and risks associated with migrating from any relational database to Oracle Exadata Database Machine. The solution includes workbench frameworks for database; extract, transform, and load; and reporting migration.

The Infosys-Oracle Innovation Center is a state-of-the-art facility jointly established by Infosys and Oracle to demonstrate Oracle Exadata accelerators, such as data warehouse migration and Oracle E-Business Suite migration console, and other solutions. The two companies also use the Innovation Center to develop Oracle-based solutions that maximize customer value and reduce costs.

▶ infosys.com

KPIT Cummins Announces Milestones

Two KPIT Cummins companies recently made Oracle-related announcements.

SYSTIME, an Oracle Platinum Partner, has achieved Oracle Validated Integration for its SYSTIME JD Edwards adapter for Oracle WebCenter 11g Release 1. The adapter optimizes Oracle's JD Edwards EnterpriseOne financial management solutions, using advanced workflow capabilities to accelerate document- and image-intensive business processes such as accounts payable and claims processing.

CPG Solutions, also an Oracle Platinum Partner, has introduced Oracle Accelerate for Life Sciences using Oracle's Agile product lifecycle management solutions. This accelerated implementation of Agile product

lifecycle management for medical device and pharmaceutical manufacturers features pre-configured workflows, data models, and test scripts with optional preconfigured integration to Oracle E-Business Suite.

KPIT Cummins is a worldwide Oracle Platinum Partner that provides consulting, implementation, product development, and 24/7 global services. KPIT Cummins' Oracle solutions-focused companies include CPG Solutions, a leader in manufacturing, supply chain, and engineering; and SYSTIME, the leading Oracle JD Edwards partner. KPIT Cummins companies have achieved five advanced specializations and 20 regular specializations.

▶ kpitcummins.com

ConSol Austria Earns Two Oracle Ready Statuses

ConSol Austria Software GmbH, an Oracle Gold Partner, has earned Oracle Database Ready and Oracle WebLogic Ready status for ConSol*CM version 6.5.5. ConSol*CM manages help desk, service management, customer relationship management (CRM), and sales processes and is built on a SOA platform that allows seamless connection to third-party systems such as enterprise

resource management, CRM, process management, and monitoring systems.

Oracle Database Ready status demonstrates that ConSol*CM has been fully tested on Oracle Database 11g Release 2. Oracle WebLogic Ready status certifies that ConSol*CM is supported on Oracle WebLogic Server 11g Release 1.

▶ consol-software.at

Supply Chain Fusion Earns Demantra Specialization

Supply Chain Fusion, a supply chain consulting firm and an Oracle Gold Partner, has achieved Specialized status for Oracle Value Chain Planning-Demand Management (Demantra). This specialization recognizes Supply Chain Fusion's proficiency in selling and implementing Oracle's Demantra solutions.

Headquartered in Santa Clara, California, Supply Chain Fusion blends supply chain process expertise with comprehensive software skills and adds proprietary Fast Fusion methodology with accelerators to transform clients' supply chains. To earn Specialized status, Supply Chain Fusion met rigorous requirements based on the needs and priorities of the customer and partner community, including completing online sales, presales, and support assessments and passing an implementation exam.

▶ supplychainfusion.com

Delta Energy Solution Achieves Multiple Oracle Ready Statuses



Delta Energy Solution AG, a Swiss company developing gas and electricity trading and scheduling solutions for the European market, has earned Oracle Linux

Ready, Oracle Solaris Ready, and Oracle Database Ready status for its ET3000 version 3.13 solution.

ET3000 provides software modules for the processing of the international energy trading of commodities such as electricity, gas, and CO₂. Additionally, the company's DeltaXE version 4 solution has earned Oracle Linux Ready, Oracle Solaris Ready, Oracle Database Ready, and Oracle WebLogic Ready status. DeltaXE provides time series management and is used in the processing and settlement of trades and energy distribution.

An Oracle Platinum Partner, Delta Energy Solution embeds Oracle software in a way that delivers high availability, high performance, and scalability suited to the fast-paced European utilities market.

▶ delta-energy.ch

The Magic of Virtualization

Oracle Technology Network embraces all things virtual.

The *virtualization* meme has changed our industry forever, and in more ways than you might think.

In its most common context, virtualization has completely rewritten the rules of operational IT. And a range of enabling technology has emerged in the past decade to support that transformation, including technology from Oracle. But if you were to boil down the virtualization concept to its essence, that essence would simply be this: with the right technology, important resources can transcend physical barriers to reach their maximum potential. (It's rare for a technology to meet the famous standard from British writer and scientist Arthur C. Clarke—"Any sufficiently advanced technology is indistinguishable from magic"¹—but this one does.)

So, we know that IT is taking full advantage of virtualization technology for the obvious potential benefits. But similarly, Oracle Technology Network is now using the essence behind it to the advantage of its community—"thinking outside the box," literally, in several ways.

Prebuilt developer VMs. It has been less than two years since Oracle Technology Network rolled out its first prebuilt developer virtual machine (VM): Oracle Database App Development VM, which was designed for use with Oracle VM VirtualBox. Today, that original VM is available alongside other VM flavors for enterprise Java development, Oracle Solaris 11 administration, Oracle Business Intelligence Enterprise Edition development, Oracle VM management, SOA and business process management development, and more. In all cases, you get a fully capable, preinstalled environment at your fingertips—far better than a demo, it's a genuine test-drive (without the multiple installs or operating system issues).

Virtual Developer Day. Face-to-face gatherings have their place and always will. But today, there is no need to leave the comfort

With the right technology, important resources can transcend physical barriers to reach their maximum potential.

of your desk chair for useful training and enablement. With the Virtual Developer Day platform, Oracle Technology Network offers online hands-on workshops—based on downloadable prebuilt developer VMs—that have a global reach and that can be accessed on demand after the fact. (Of course, attending an in-person Oracle Technology Network Developer Day in your area may be a possibility, if that's your preference. In fact, why not do both?) Developer Days come in various flavors, including Enterprise Java Development with Java EE, Glassfish, and Oracle WebLogic Server; Rich Enterprise Application Development with Oracle ADF; and Java Development with Java SE 7, to name a few.

Cloud computing. With the coming general availability of Oracle Public Cloud, Oracle Technology Network members will have self-service, subscription-based access to Oracle Database Cloud Service and Oracle Java Cloud Service for rapid development and deployment without the burden of installing, configuring, and managing the software and systems.

New virtual communities. Forums.oracle.com, deployed in 2001 and now with 1 million registered users on its own, is the most venerable example of a virtual community brought to you by Oracle Technology Network. In the past year, Oracle has also made major investments in Java.net, the online community and collaborative devel-

opment platform for Java developers around the world that is also the home of Oracle's own open source Java projects.

Java Magazine. Yes, I even put *Java Magazine* in this category. Because the publication is all digital, all the time, its contents can be shared among readers globally much more efficiently than a hard copy. When was the last time you mailed or faxed a paper reprint to someone? Furthermore, the magazine's format allows the editorial team to experiment with rich content and interactivity in new and interesting ways.

In closing, when you think about it, the virtualization concept has applicability far beyond IT operations. On the Oracle Technology Network team, we're putting that thought into action—our specialty.

Do you have any other thinking-outside-the-box ideas for us? Let us know via our community feedback channel at forums.oracle.com. ◀



Justin Kestelyn

(justin.kestelyn@oracle.com) is senior director, Oracle Technology Network and developer programs, as well as

Oracle Technology Network editor in chief.

WEB LOCATOR

Prebuilt Developer VMs

oracle.com/technetwork/developervm

Oracle Technology Network Developer Day (Physical and Virtual)

bit.ly/yjjWSW

Oracle Public Cloud

cloud.oracle.com

Java.net

java.net

Java Magazine

oracle.com/javamagazine

Oracle VM VirtualBox

bit.ly/zEljBx

Role Models

Great minds and big ideas shape architectural decisions.

The decisions that IT architects make are informed by a variety of factors and influences, including an up-to-date knowledge of available technologies, a thorough understanding of how best to apply those technologies to solve the problem at hand, and the input of project stakeholders. These factors are common to every IT architect—or at least they should be. But what other, perhaps less tangible factors come into play? For instance, does the benevolent specter of a role model ever hover just over the shoulders of IT architects as they work toward a solution? How does that role model shape the individual architect? I asked several architects to tell me about their role models.

Randy Stafford, a member of Oracle's Server Technologies A-Team (that's A for *Architecture*), divides his role models into two groups. "The first are seers and teachers, architecturally literary people with a rare gift for perceiving the essence of a relevant topic and having the dedication and skill to clearly convey it in published works for the benefit of the profession," he says. For Stafford, that list of luminaries includes object-oriented programming expert Martin Fowler, engineering and space pioneer Eberhard Rechtin, author and object-oriented programming consultant Rebecca Wirfs-Brock, and several others. "Their contributions to the literature of the software profession have given me, and many other software architects, a body of insight on which to base our daily practice and to develop and communicate new learnings in a hopefully perpetual process," Stafford says.

Stafford describes his second group of role models as "those who have served as guardians of the conceptual integrity of a software system over a long span of time." One such role model from early in Stafford's career is a man who wrote and installed his own assembly language for a CDC Cyber mainframe. "He developed a

When you start a project, what benevolent specter is hovering over your shoulder?

software stack that was critical to the operations of an aerospace company for several decades," Stafford says. Gene Gleyzer, the chief architect of Oracle Coherence, is another of Stafford's role models. "Gene exercised responsibility for the algorithms involved in Oracle Coherence's data management capabilities, among other things, since well before Oracle Coherence 1.0 was released," Stafford says. "These kinds of people exemplify what it means to be a software architect."

Eric Stephens, a director of enterprise architecture at Oracle, mentions enterprise architecture pioneer John Zachman as one of his role models. "Zachman really expanded my thinking of architecture and emphasized the need to engineer the enterprise, like one might engineer an aircraft or ship," Stephens says. Stephens also cites Leonardo da Vinci and Steve Jobs as role models. "Both of these gentlemen leveraged their artistic and creative energy to conceptualize—and in Job's case, implement—great works of artful engineering," Stephens says.

There is also an Apple Inc. connection in Oracle Enterprise Architect Pat Shepherd's choice for role model: venture capitalist and former Apple evangelist Guy Kawasaki. "He has a unique way of breaking a problem down into simple terms," says Shepherd. Shepherd points out that although the term *architect* doesn't appear in Kawasaki's bio, "he has always been at that intersection of business value, vision, and technology." Shepherd believes that Kawasaki's book *Rules*

for *Revolutionaries* (HarperBusiness, 2000) belongs on every enterprise architect's bookshelf, "right next to the TOGAF [the Open Group Architecture Framework] manual."

Dr. Karina Ishkhanova, technical lead for payment systems architecture and design at School-Day Solutions, credits science fiction writer Isaac Asimov as having a profound influence on her career. "When I was 10 years old, I read his short story 'Profession,' and it stuck with me for the years to come," Ishkhanova says. "He inspired me to be different—first to dream and then to build upon that dream." Ishkhanova finds similarities between system architecture and Asimov's futuristic fiction. "On a daily basis, system architects, armed with logic and vision, build a new reality," she says, adding that she looks to Asimov's work for inspiration. "When I start a new project, I try to find an edge that will make it unique and innovative, something that will allow others to grow further and extend."

When you start a project, what benevolent specter is hovering over your shoulder? Whose thoughts influence your decisions? How are those thoughts shaping you as an architect? ◀




Bob Rhubart

(bob.rhubart@oracle.com) is manager of the architect community on Oracle Technology Network, the host of the

Oracle Technology Network ArchBeat podcast series, and the author of the ArchBeat blog (blogs.oracle.com/archbeat).

NEXT STEPS

 LISTEN to ArchBeat podcasts
feeds2.feedburner.com/otnarch2arch
bit.ly/9X6j2T



NetShelter CX enclosure solution includes everything you need to house a highly reliable IT deployment regardless of space limitations.

Expand your IT capabilities, not your real estate.

APC by Schneider Electric helps you deploy your IT whenever and wherever you need it.

Is your lack of IT space a barrier to adopting new technologies?
 Consolidation, virtualization, network convergence, blade servers—these new technologies improve efficiency, cut costs, and allow you to “do more with less.” But they also bring power, cooling, and management challenges, especially when you’re tasked with deploying your IT without dedicated space. You’re relying on guesswork, depending on building air conditioning, or improvising remedies. So, how can you increase the level of reliability and control of your server deployment without spending a fortune?

The availability you need—without the IT room.
 APC™ by Schneider Electric™ understands the challenges of delicately matching IT needs with logistical realities. Fortunately, we have adaptable and flexible solutions that provide everything for your IT deployment: reliable and efficient power, cooling, monitoring, and management. Whatever your logistical or space constraints, we have a total solution to meet your specific needs. All components have been pre-engineered to work together and integrate seamlessly with your existing equipment.

Future-proof your IT deployment.
 There’s no need for confusing cooling configurations or expensive mechanical re-engineering, so our modular, “pay-as-you-grow” designs let you be 100 percent confident that your IT capabilities can keep pace with ever-changing demands. Self-contained cooling, high-density enclosures, rack-level power distribution and monitoring sensors, and integrated management software provide complete remote control and unprecedented visibility into your entire IT system. Simply add power protection (such as undisputed, best-in-class Smart-UPS™ or Symmetra™ units), and you have a total solution for today, tomorrow, and beyond.



Learn how to reduce cooling expenses with our **FREE** Cooling Efficiency kit and enter to win 1 of 5 Smart-UPS units (SMX1000)!

Visit www.apc.com/promo Key Code **m708v** • Call 888-289-APCC x6371 • Fax 401-788-2797

Infra²struxure™

- 1 Enclosures** Vendor-neutral NetShelter™ SX rack design handles high-density airflow and power needs.
- 2 Power** Ultra-reliable Smart-UPS and Symmetra UPS units offer scalable runtime, and PDUs are rack-mounted.
- 3 Environmental monitoring and management** PoE-enabled temperature sensors let you keep an eye on conditions at the rack level, and centralized software gives you real-time insight into the entire system.



The NetShelter CX Office enclosure is available in three sizes: 18U, 24U and 38U.

Expand your IT capabilities without building out
 The unique, soundproof, self-contained NetShelter CX enclosure solution includes everything you need to house a highly reliable IT deployment regardless of space limitations.



Big Data Boot Camp

An early adopter shares big data expertise at COLLABORATE 12.

Ian Abramson, a former president of the Independent Oracle Users Group (IOUG) has a new title: drill sergeant. Abramson, a director of enterprise data at Oracle partner Thoughtcorp, is an early adopter of big data methods and technologies and was tapped by IOUG, the world's largest Oracle technology user group, to lead the Big Data Boot Camp at this year's COLLABORATE conference in Las Vegas, Nevada, April 22–26. Big Data Boot Camp sessions take place on April 23 and 24.

"Most user group members are just starting to look into big data, and these sessions position them well for what's happening in this space," says Abramson. Big data is the mass of data available on public and private networks, such as blogs, social media, e-mail, sensor data, and photographs, that can be mined for useful information.

"When I started to look into big data two years ago, I realized that this collection of massive amounts of information was often being ignored because it seemed too big to go after," Abramson recalls. "I started asking how we could take this unstructured data and start bringing it together with structured data and analyze it. It seemed like a natural extension of data warehousing."

STORE EVERYTHING, KNOW ANYTHING

Abramson says there is a central idea in data warehousing behind the drive toward big data. "If I could store everything, I could answer anything," he says. New tools for big data help organizations get closer to that ideal goal, he says, allowing them to store petabytes of data and look for relationships hidden within that data. "Now that we have the tools, we can capture big data and combine it with structured data so that a retailer, for example, can measure fuzzy things like the sentiment around one of its products. Or a utility company can get a finer understanding of how



Ian Abramson, Director of Enterprise Data, Thoughtcorp

customers use electricity throughout the day, month, and year," says Abramson.

THE WHAT, HOW, AND WHY OF BIG DATA

As drill sergeant of the COLLABORATE 12 Big Data Boot Camp, Abramson aims to give people a firm understanding of the components of big data and how they fit together. "They learn about tools such as NoSQL and Hadoop and what's important about each, and how to connect big data to the structured data in their data warehouse. They will understand why the Oracle Big Data Appliance exists, what's in it, and how to get the most from its potential," he says.

Abramson plans to put big data in a real-world context. "We show attendees who's using it and what they're using it for, in industries such as telecom, retail, and healthcare," he explains.

TRY IT FIRST IN A USER GROUP

Abramson sees a big role for user groups in the uptake of big data. "What a user

group does well is provide a place for people to learn about things they haven't done before," says Abramson. "There are usually early adopters in the group, and they begin to educate the rest of us."

Abramson regards the early stages of big data adoption as being very similar to the first year after the Oracle Exadata launch. "People didn't really understand how to use Oracle Exadata or what it was for," Abramson recalls. "Here we are today in 2012, and Oracle Exadata is huge. People understand it. They get the technology," he observes. "Big data is going to be the same story."

Luckily for IOUG members, there's an experienced leader and big data early adopter offering sessions at COLLABORATE 12. "User group membership has helped me in so many ways over the years," Abramson concludes. "It has helped me stay ahead of the curve on new technologies and provided avenues for growth. I'm looking forward to this opportunity to teach what I've learned about big data and give back to the community." ◀



Jeff Erickson

(jeffrey.x.erickson@oracle.com) is a senior editor with Oracle Publishing.

NEXT STEPS

LEARN more about IOUG
ioug.org

WATCH the interview
oracle.com/oramag/upclose

SEE the kind of information big data initiatives can provide with LinkedIn's visualized network
inmaps.linkedinlabs.com

Playing Favorites

Three peers on the tools and books they rely on most



GÖKHAN ATIL



ORACLE
ACE

Company: Vodafone Net, a subsidiary of mobile phone operator Vodafone Turkey

Job title/description: DBA, responsible for installing, maintaining, and troubleshooting databases; performance; and SQL tuning


Location: Istanbul, Turkey

Oracle credentials: Oracle Certified Professional (Oracle Database 10g, Oracle Database 11g), Oracle Forms Developer Certified Professional, and Oracle Database SQL Certified Expert, with 10 years of experience using Oracle products

What's your favorite tool on the job? Oracle Enterprise Manager 12c. It's absolutely the best tool to manage the Oracle stack—and Oracle Enterprise Manager Cloud Control, which I began implementing last December, provides a comprehensive solution for all layers including hardware, virtualization, database, middleware, and applications.

Which new features in Oracle Database are you currently finding most valuable? Hybrid Columnar Compression. About 10 years ago, compression was an alternative; now it's a must. Hybrid Columnar Compression provides incredible compression ratios, especially for historical data, and will soon be the de facto standard on data warehouse systems.

You've taken Oracle University [OU] classes in the past. What led you to do this? I've taken several classes from OU, including Program with PL/SQL and some Oracle Database 10g workshops. I wanted to improve my knowledge, and I believe the best way to learn Oracle technologies is by taking classes from OU. I also believe in certification, and it has truly helped me in my career.



DOUG GAULT



ORACLE
ACE

Company: Sumneva, a provider of Oracle Application Express training, consulting, and solutions

Job title/description: Director and cofounder, involved in day-to-day operations, client consultations, and the development of software and training materials

Location: Dallas, Texas

Length of time using Oracle products: 25 years

How did you get started in IT? I started off programming on a RadioShack TRS-80 in my high school's computer lab. I then attended a small technical college in Iowa, and my first job out of college was in the IT department of a bank holding company. My second job, for a small consulting and software company, is where I was introduced to Oracle. I haven't looked back.

What are your go-to Oracle reference books?

Pro Oracle Application Express by John Scott and Scott Spendolini [Apress, 2011] and *Oracle PL/SQL Programming* by Steven Feuerstein and Bill Pribyl [O'Reilly Media, 2009]. I have them in both electronic and hard copies—sometimes it's nice to flip through the real thing.

What advice do you have about how to get into Web, database, and application development? Don't be scared of it, and get your hands dirty! Like any other language or technology, the best way to learn is to do.

What's your favorite thing to do outside of work? I have two loves: food and travel. My favorite pastime is traveling wherever I can and sampling the local food at my destination.



PATRICK BAREL



ORACLE
ACE

Company: Employed by AMIS Services, an IT consulting firm, but currently working full time for energy distribution company Alliander

Job title/description: Senior Oracle developer, building SQL and PL/SQL code to update data in a real-time system for Alliander's power network

Location: Haarlem, the Netherlands

Length of time using Oracle products: 15 years

What's your favorite technique on the job?

In SQL, it's analytics—because analytical functions can help you retrieve results faster than doing the same with subqueries. And when I'm doing data manipulation in PL/SQL, one of the best techniques available is using collections combined with bulk processing.

What technology has most changed your life? PL/SQL—because it gets the job done.

PL/SQL provides lots of great constructs, as does Oracle Database, so all the issues posed to me are rather easily solved using this great language. Then I get to let the database work for me, doing the heavy lifting, and I get to go back to the good things in life, such as spending time with my wife and three kids.

How are you using mobile computing in your work? I'm taking baby steps on the jQuery

Mobile side of Oracle Application Express—so far my development is mostly restricted to trying out or rebuilding demos put on the internet by fellow Oracle Application Express enthusiasts such as Peter Raganitsch [oracle-and-apex.com]. Hopefully the upcoming Kscope conference [June 24–28 in San Antonio, Texas] will shed more light on this subject for me. ◀


SUPER HERO

Oracle engineered systems, storage, and analytics support the Avengers.

BY RICH SCHWERIN

Yes, your data center is powerful. But is it powerful enough to protect the nations and peoples of Earth from all threats, terrestrial or extraterrestrial? If you're Agent John Smith, it better be.

Smith is the chief systems administrator for the Supreme Headquarters International Espionage Law-Enforcement Division (S.H.I.E.L.D.), where he works for Nick Fury. And just as Fury must assemble the Avengers to battle threats to the world that no one else can handle, Smith must assemble the most powerful data





TM & © 2012 Marvel & Subs

AT THE CENTER OF EXTREME DATA

With extreme performance in mind, S.H.I.E.L.D. Chief Systems Administrator John Smith assembled an integrated private cloud on a foundation of Oracle engineered systems and business analytics software. Individually, each system component delivers extreme performance. Together, the systems support even more information for deeper analysis and some of the most important decisions in the history of humankind.

Exalogic Elastic Cloud X2-2

Engineered for cloud computing, Oracle Exalogic hardware and software are tuned for Java application performance. S.H.I.E.L.D. agents use Oracle Exalogic to run S.H.I.E.L.D.'s day-to-day operations, enter mission and status reports, and run classified applications.



Exadata Database Machine X2-8

Oracle Exadata provides S.H.I.E.L.D. agents and analysts with extreme database performance for both data warehousing and online transaction processing (OLTP) applications.



Oracle Big Data Appliance

Every day, S.H.I.E.L.D. captures terabytes of information from a variety of sources—surveillance video, satellites, sensors, field reports, network traffic—and all of this high-volume, high-velocity, and high-variety data is processed, filtered, transformed, and sorted in Oracle Big Data Appliance.



Storage


With international jurisdiction comes an international volume of data, all of which needs to be securely stored, backed up, and archived. S.H.I.E.L.D. depends on Oracle's Sun ZFS Storage for hot backup and network-attached storage (NAS), Pillar Axiom for storage-area network (SAN) storage, and StorageTek tape libraries and drives for backup and archiving. Storage includes

- Sun ZFS Storage 7420 appliance
- Pillar Axiom 600 StorageTek SL8500 modular library system



Exalytics In-Memory Machine X2-4

When assessing threats, S.H.I.E.L.D. must quickly draw actionable insights from massive amounts of data. Oracle Exalytics' advanced data visualization and exploration get the job done.



Oracle Business Analytics

For S.H.I.E.L.D., threat assessment begins with daily input—massive amounts of structured and unstructured data, collected from tens of thousands of sources, processed in Oracle engineered systems, and stored intelligently. But while raw information ultimately holds the answers, that information requires analysis before S.H.I.E.L.D. can finalize situation reports and take action. S.H.I.E.L.D. relies on analytics and enterprise performance management from Oracle to discover new ways to strategize, optimize operations, and capture new opportunities for success. Through a custom analytics dashboard, S.H.I.E.L.D. gains insights into all aspects of the operation and can initiate actions that keep humankind safe from threats—terrestrial and extraterrestrial.

center to deliver extreme performance and Earth-saving analytics.

A seven-year S.H.I.E.L.D. veteran, Smith has seen a lot of technology come and go. Some of it is retired, and updated hardware and software take its place. At S.H.I.E.L.D. headquarters, in the Helicarrier, and in field offices around the world, some hardware is retired by force.

"An Asgardian's hammer really does a number on these infrastructures. We're using kevlar and a lot of classified materials in the housings of our server and storage hardware, but superpowers give IT a run for its money," says Smith.

REASSEMBLE

Oracle Magazine caught up with Smith after a classified incident destroyed all of the legacy data center hardware at a previously secret S.H.I.E.L.D. lab. Smith seemed surprisingly upbeat as he looked at scorched pieces of hardware where dozens of racks of state-of-the-art systems once stood.

"Today, we're rolling in Oracle engineered systems, so recovering from this kind of event is easy," Smith says. "The intelligence we deliver enables S.H.I.E.L.D. to complete its missions, so high availability and performance are paramount. The faster we can assemble the data center and deliver vital information, the better."

Smith pauses to watch his trio of tech teams in action. One group rolls in new engineered systems: Oracle Exadata Database Machine, Oracle Exalogic Elastic Cloud, Oracle Exalytics In-Memory Machine, and Oracle Big Data Appliance. Another group works to deploy an arsenal of storage solutions: Oracle's Sun ZFS Storage Appliances, Pillar Axiom 600, and StorageTek tape libraries and drives. And a third group of technicians focuses on analytics, installing Oracle business intelligence (BI) tools and technologies and Oracle Hyperion performance management applications.

ENGINEERED FOR BATTLE

Extraterrestrial security emergencies don't wait, so S.H.I.E.L.D. can't waste time configuring, testing, connecting, and deploying disparate hardware piecemeal—which is why Smith relies on Oracle engineered systems.

"We chose Oracle Exadata because it provides extreme performance for both data warehousing and OLTP [online transaction processing] applications, and its architecture is highly scalable, secure, and redundant," says Smith. "As a result, our agents—and of course, the Avengers—get the right information at the right time and can make better decisions to be more effective in day-to-day operations and in battle."

S.H.I.E.L.D. systems connect thousands of agents and exabytes of classified information in international operations. The frontline data gathering and analysis applications run in a secure private cloud, so building them in Java and running them on Oracle Exalogic machines makes perfect sense to Smith. "Oracle Exalogic is engineered for cloud computing," says Smith. "It's tested and tuned for Java applications, and it delivers extreme performance."

S.H.I.E.L.D. takes in terabytes of information every day, from surveillance video to satellites, threat sensors, field reports, network traffic, and more. This high-volume, high-velocity, and high-variety

data must be processed in batch and in parallel, filtered, transformed, and sorted before it is loaded into an enterprise data warehouse. For that tall task, Smith and S.H.I.E.L.D. rely on Oracle Big Data Appliance, an engineered system optimized for acquiring, organizing, and loading unstructured data into Oracle Database 11g. "Every day we load massive amounts of mission-critical intel into Oracle Big Data Appliance, and it handles everything we give it," says Smith.

Intelligence, planning, and analytics play prominent roles in protecting Earth from all threats, and Oracle Exalytics supports those roles for S.H.I.E.L.D. As an in-memory BI machine, it delivers the extreme performance S.H.I.E.L.D. needs to succeed. "When assessing threats, S.H.I.E.L.D. has to quickly draw actionable insights from massive amounts of data," says Smith. "Oracle Exalytics is our go-to enterprise BI platform, and its advanced data visualization and exploration get the job done."

"The faster we can assemble the data center and deliver vital information, the better."

—Agent John Smith, Chief Systems Administrator, S.H.I.E.L.D.

SUPREME STORAGE

With S.H.I.E.L.D.'s international jurisdiction comes a massive amount of data, both structured and unstructured—all of which needs to be securely stored, backed up, and perpetually available.

So in addition to using all of the storage built into its Oracle engineered systems, S.H.I.E.L.D. uses network-attached storage (NAS), storage-area network (SAN), and tape storage devices—in particular Sun ZFS Storage Appliance, Pillar Axiom storage system, and StorageTek tape libraries and drives. "Our Sun ZFS Storage Appliances provide hot backup for our Oracle Exadata machines, and they deliver NAS storage for the systems that need it," says Smith. "When you're storing as much critical, Earth-in-the-balance data as we are, you simply can't afford any data corruption errors. The ZFS file system's advanced error detection and self-healing capabilities keep S.H.I.E.L.D.'s data accurate and online."

To address SAN storage requirements, Smith turns to the next-generation SAN architecture of the Pillar Axiom 600 storage system.

"S.H.I.E.L.D. agents place pretty high demands on the data center," says Smith. "The storage provided by the Pillar Axiom 600 enables us to meet some very particular and strict service-level agreements [SLAs], prioritize data access across the system, and ensure that all applications get what they need when they need it."

Because the volume and value of S.H.I.E.L.D.'s data grows exponentially, Smith uses Oracle's StorageTek SL8500 modular library system to help him manage complexity, control costs, and deliver on SLAs.

"Considering the massive volume of data we need to back up, the StorageTek SL8500 is well equipped. It's an exabyte storage system with 24/7 high-availability features," says Smith.

Knowing that a superpowered disaster can strike anywhere at any time, S.H.I.E.L.D. operates mirrored Oracle engineered compute and storage systems in multiple locations worldwide in a highly secure private cloud configuration. As Smith's tech teams complete the connections between the Oracle engineered compute and storage systems, they begin data recovery from tape and other systems in the S.H.I.E.L.D. private cloud.

SMART, AGILE, ALIGNED

Oracle engineered systems and storage have an obvious presence in S.H.I.E.L.D.'s rebuilt data center, but S.H.I.E.L.D. also depends on additional software to analyze and manage its own unique operation. For Smith, analytics and enterprise performance management are among the most important software deployed at S.H.I.E.L.D.

"With Oracle Business Intelligence Applications and Oracle Hyperion performance management applications, we're able to manage analytics and optimize operational performance worldwide," says Smith. "With Oracle, we can quickly assemble insights into action anytime, anywhere, on any device."

"When Fury assembles the Avengers, our IT has to be a step ahead," concludes Smith. "And thanks to Oracle's engineered

systems, storage, and analytics, we've assembled a superpowered data center worthy of Earth's Mightiest Heroes." ◀

Rich Schwerin is a senior manager with Oracle Publishing.

NEXT STEPS

Marvel's *The Avengers* in Theaters May 4

READ more about *The Avengers*

oracle.com/theavengers

LEARN more about

Oracle Exadata Database Machine

oracle.com/exadata

Oracle Exalogic Elastic Cloud

oracle.com/exalogic

Oracle Exalytics In-Memory Machine

oracle.com/exalytics

Oracle Big Data Appliance

oracle.com/bigdata

Oracle storage solutions

oracle.com/storage

Oracle business analytics

oracle.com/bi

Oracle Storage Networking

Powered by QLogic, Supports Newly Launched Oracle x86 Servers

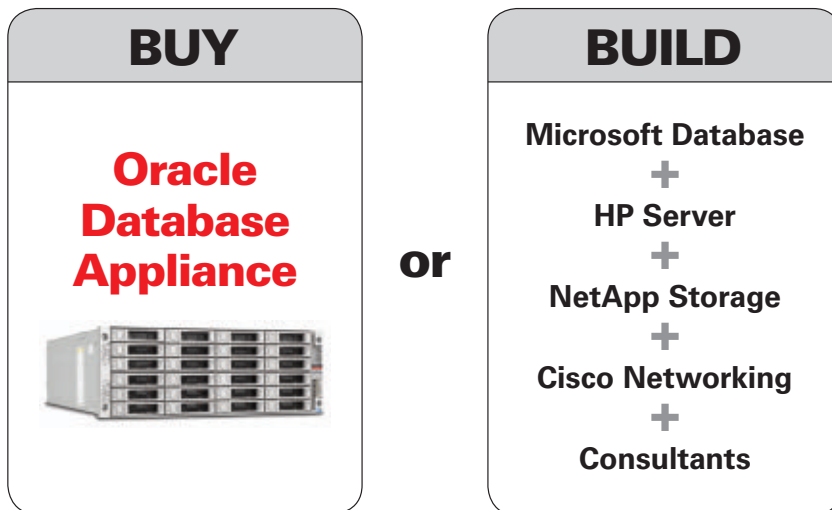
QLOGIC CORPORATION is happy to announce its support of Oracle's new Sun Fire X4170 M3, Sun Fire X4270 M3, and Sun Blade X6270 M3 servers with its Oracle-branded storage networking adapters. QLogic adapters are designed to balance the high compute density of Oracle's Sun Fire M3 servers by providing full hardware offload, enterprise I/O performance, and the highest reliability. Together, these technology solutions from QLogic and Oracle enhance data center efficiencies and increase ROI in IT landscapes where enterprise I/O performance is crucial.

To find out more about Oracle's StorageTek 8 Gb Fibre Channel PCIe HBAs and Sun Storage 10 GbE FCoE Express Module Converged Network Adapters, powered by QLogic, visit us at www.qlogic.com/go/oracle.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.



Database Appliance or Puzzle Pieces?



Database Appliance: Simple, Reliable, Affordable

ORACLE®

oracle.com/databaseappliance
or call 1.800.ORACLE.1

STORAGE S



STANDS TALL

Oracle engineered storage systems deliver tiered storage optimized for your I/O.

BY DAVID BAUM

Oracle's enterprise storage solutions combine the best hardware with software designed to take advantage of that hardware to meet the requirements of various types of I/O processing: database, network-attached storage (NAS), and storage-area network (SAN).

According to Benjamin S. Woo, program vice president for Worldwide Storage Systems at International Data Corporation (IDC), these "purpose-built" storage platforms provide a great deal of control over how I/O can be distributed. With Oracle Exadata storage, for example, "Oracle has advanced database processing





Stephen Schleiger, manager of systems engineering at Navis, depends on the reliability of the Pillar Axiom storage system. "Having a stable data management environment is essential to our customers," he says.

"Pillar Axiom . . . has all of the components we need to maintain exceptional performance."

—Stephen Schleiger, Manager of Systems Engineering, Navis

and integrity," explains Phil Bullinger, senior vice president of storage at Oracle. "The performance requirements of application workloads, the structure of the data traffic they generate, the architecture of the data center, and the retention lifecycle of the data all influence the optimal implementation of storage solutions for various environments. NAS solutions store data under a file system and service heterogeneous client systems, while SAN solutions function as block storage devices for multiple servers. Tape complements disk storage by providing an

several steps further than other vendors," says Woo. "Rather than just offering an integrated hardware platform, [Oracle] includes technology that optimizes the movement of data to and from the physical media."

Oracle Exadata is just one of the platforms in Oracle's rapidly evolving storage product line. Oracle has realized similar feats of engineering with its Sun ZFS Storage Appliance, which is optimized for NAS storage (see the "Maximizing Options with ZFS" sidebar), as well as with its Pillar Axiom product line, designed primarily for SAN environments. And when it comes to backup, disaster recovery, archiving, and long-term preservation, Oracle continues to lead the industry with modular library systems such as the StorageTek SL8500 and tape drives such as the StorageTek T10000C.

"Storage systems must respond to a variety of workload demands and operational exceptions while always ensuring data availability

order-of-magnitude better capacity economics for large data stores, robust data protection for disaster recovery, and unmatched data retention for long-term archival."

SAN SELECTED

A subsidiary of Cargotec Corporation, Navis provides a software platform that is used at shipping terminals around the world. Its software automates essential functions such as tracking containerized cargo through a port, automating equipment operations, and coordinating cargo among multiple locations. Marine terminals, rail yards, and distribution centers in more than 50 countries rely on Navis' software to automate and optimize many facets of their operations.

Navis storage engineers selected a SAN environment to meet the escalating demands of a global software development team,

Maximizing Options with ZFS

Oracle's Sun ZFS Storage Appliance delivers impressive performance for network-attached storage (NAS)-oriented workloads by combining state-of-the-art hardware and software.

"Hardware provides the capability for speed, but it takes innovative and efficient software to actually deliver the complete potential of the hardware components in application performance," says Phil Bullinger, senior vice president of storage at Oracle. "Sun ZFS Storage Appliance combines the highly scalable Oracle Solaris II OS,

the ZFS file system, and a rich suite of data management features with a massively provisioned Sun server platform that includes multisocket CPUs with multiple cores and large amounts of dynamic random access memory (DRAM) and Flash Cache."

Another unique aspect of the ZFS architecture is Oracle's Hybrid Storage Pool technology, which allows customers to create virtual pools of storage that include a combination of traditional spinning disks and solid-state flash, along with sig-

nificant amounts of DRAM. The storage appliance automatically matches the correct components to designated workloads to maximize performance.

"Our caching algorithms dynamically place frequently and recently used data into the fastest possible storage tier—first DRAM, and then flash, and then onto fast-spinning disks," Bullinger explains. "As data patterns and application workloads change, this system dynamically adjusts to always optimize application performance."

purchasing two Pillar Axiom 600 storage systems and two StorageTek SL3000 modular library systems. Navis is now creating a virtual IT infrastructure to improve the flexibility of its software-development activities.

"We have a very fluid and dynamic environment, with a lot of buildup and teardown in the engineering process," says Stephen Schleiger, manager of systems engineering at Navis. "Being able to logically partition the storage environment to support virtual instances is the secret sauce of the Pillar Axiom storage system. It has all of the components we need to maintain exceptional performance and reliability. It is a solid platform that we can tailor to our individual requirements. With 25 percent of the world's container moves handled using our software, having a stable data management environment is essential to our customers."

In addition to providing components that can be custom tailored to fit an organization, the Pillar Axiom storage system responds to the changing needs of applications.

"The Pillar Axiom storage system has patented technology that adapts and prioritizes all the physical resources of the system to the relative business priorities of the various application workloads it's supporting," explains Oracle's Bullinger. "By dynamically controlling the priority of data in cache, the execution order of I/O activity, and the tiered placement of data on rotating and solid-state disk drives, the Pillar Axiom system can consolidate multiple applications and deliver deterministic levels of performance and priority for each one while linearly scaling to meet growing capacity and performance requirements."

SIZING UP THE OPTIONS

Schleiger first encountered the Pillar Axiom product line when Navis decided to either upgrade or replace an existing storage solution. "We looked at different vendors, but Pillar [Axiom] systems particularly intrigued us," he recalls. "In order to keep up with the demands of our team, we knew we had to virtualize. We saw the Pillar Axiom 600 as the best solution for our software engineering operations."

Navis determined that it would get more for its money with Pillar Axiom, compared to competing products. "From a management standpoint, the Pillar Axiom line is very easy to use," Schleiger notes. "The interface is intuitive. We realized we could come up to speed on the Pillar Axiom faster than we could with any other storage platform."

Working with Oracle Consulting, Navis is implementing the Pillar Axiom 600 systems in its U.S. and India data centers as part of a disaster recovery architecture that uses the Pillar Axiom MaxRep Replication utility to bidirectionally replicate data between these two locations. Net changes to specified volumes are automatically captured via snapshots every 15 minutes and then replicated to the Pillar Axiom 600 volumes in the alternate data center. Navis plans to use Oracle's storage archive management file system (SAM-FS) to

SNAPSHOTS

Navis
navis.com
Headquarters: Oakland, California
Industry: Transportation
Employees: 300
Oracle products and services: Pillar Axiom 600 systems, StorageTek SL3000 modular library systems, Oracle Database, Oracle Consulting

FamilySearch
familysearch.org
Headquarters: Salt Lake City, Utah
Industry: Research
Oracle products: StorageTek SL8500 modular library systems, StorageTek T10000C tape drives, Sun Storage Archive Manager, StorageTek SL500 modular library systems, SPARC Enterprise M9000 servers, Oracle Database

create a tiered storage environment.

"Using SAM-FS in conjunction with MaxRep will ensure that the development team always has the right content in the right place at the right time," explains Schleiger. "SAM-FS provides a great way for us to develop a complete, policy-based storage solution that covers everything from archival storage to the critical data that we're looking at every day."

By offloading some storage functions to tape technology, the tiered storage solution is expected to reduce Navis' costs by more than half. "As we examined the StorageTek SL3000 modular library systems, we realized that tape provides a much more robust and cost-effective disaster recovery solution than we currently have in place," Schleiger adds.

"When you look at the price and the benefits, the Oracle tape drives and libraries are an excellent solution."

A COST-EFFECTIVE ALTERNATIVE

FamilySearch came to the same conclusion when it sought a digital preservation and archiving system for its massive genealogical service. For more than 100 years, FamilySearch has been actively gathering, preserving, and sharing genealogical records with people all over the world—initially through thousands of family history centers in 70 countries, and later through a public Website where any Web visitor can tap into a database that contains billions of genealogical records. Millions of people use this online service each year to piece together their family trees and learn more about their ancestors.

The online service provides access to an Oracle database that contains linking information gathered from individuals building their family trees. Currently 20 terabytes (TB), it's a small amount of data when compared to the FamilySearch preservation sites, where the organization maintains digital images of historic documents. Digital camera crews throughout the world are constantly gathering new data as they photograph vital records from historical archives, churches, hospitals, and government bureaus. At one archival site, FamilySearch has been storing images of historic documents on microfilm since the 1930s and has amassed 3.5 million rolls of film containing 4 billion records.

"Microfilm is a very stable medium, but it is not very accessible," says Randy Stokes, group architect for engineering services, storage

Newly Awarded

Oracle's Sun ZFS Storage Appliance has earned the highest overall scores in the six-year history of *Storage Magazine's* Quality Awards for NAS, for both enterprise and midrange NAS. The Quality Awards identify and recognize products that have proven their quality and reliability in actual use. The results are based on a survey that assesses products in initial product quality, features, reliability, technical support, and sales force competence. Learn more at oracle.com/us/corporate/press/1505421.



FamilySearch uses StorageTek SL8500 systems to store digitized documents. “We amass 10 to 12 terabytes of new image data each day,” says Randy Stokes, group architect for engineering services, storage infrastructure, at FamilySearch.

infrastructure, at FamilySearch. “We knew that if we wanted to make it easier for people to do research and find their ancestors, we would have to put this information into digital form.”

It’s a huge job. Approximately 25 automated scanners are used to scan the microfilm to digital images. Additionally, new photographic images are continually coming in from 185 digital cameras in the field. These operations yield terabytes of new image data each day. The original lossless images are saved as JPEG 2000 images for long-term storage on LTO-4 tapes in three StorageTek SL8500 modular library systems. One copy is written to an in-house preservation system and another to an offsite archival location.

“We decided to use tape libraries mainly because of the cost and long-term stability of this medium,” Stokes says. “Between the scanners and the cameras, we amass 10 to 12 terabytes of new image data each day. Our primary challenge has been to figure out how to effectively manage this massive influx of information. It is difficult to find a system that can deal with this much data around the clock with no downtime.”

Before purchasing the StorageTek SL8500 systems, Stokes and his team spent six months evaluating tape storage vendors, all the way down to the chemical makeup of the various types of tape. After an intense investigation, they decided to go with the StorageTek SL8500 and LTO tapes because they liked the standard format. FamilySearch

Talking Tape

Jim Cates, vice president of storage technology at Oracle, talked with *Oracle Magazine* about the past, present, and future of enterprise tape technology. Listen to the podcast at oracle.com/magcasts.

currently has more than 8,000 LTO tapes preserving more than 6 petabytes (PB) of data. The organization expects to have more than 10 PB of preservation data in storage by the end of 2012.

UPGRADING THE INFRASTRUCTURE

FamilySearch also appreciates the tremendous raw capacity available in the new StorageTek T10000C tape drives—up to 5 TB per tape—as well as being able to write cyclic redundancy check (CRC) checksums on the tapes. “This will save us a lot of money and effort, because it allows us to validate the integrity of the data without having to pull the data off of the tape and onto a server,” Stokes explains. “We simply mount a drive and issue a ‘validate’ command to ensure that all the checksums match the actual data.

Going forward, preservation data will be written to these new StorageTek T10000C tape drives. “Our intent is to preserve the data for at least 1,000 years,” Stokes says. “After comparing the T10000C against competitive products, we decided to continue with the StorageTek technology.”

In addition to using the StorageTek SL8500 library and StorageTek T10000C tape drive storage technology for archiving, FamilySearch runs its Website on Oracle’s SPARC Enterprise M9000 servers with a 20 TB Oracle database. FamilySearch also uses four StorageTek SL500 modular library systems with LTO-4 tape drives for system development and testing as well as tape migration. FamilySearch uses another StorageTek SL8500 as part of its disaster recovery scheme.

“All of the Oracle storage devices work well together and have been very reliable,” Stokes concludes.

That’s just the kind of engineered success that sets Oracle apart as a storage vendor. “Unlike most storage vendors, Oracle has an end-to-end solution—from the applications to the middleware to the database platform to the operating systems to the hypervisor through to the service port and the network—that is unique in the market today,” says IDC’s Woo. “Other vendors integrate various parts of this stack, but Oracle is in the unique position of providing end-to-end, optimized, engineered systems.” ◀

David Baum (david@dbaumcomm.com) is a freelance business writer.

NEXT STEPS

LEARN more about

Oracle storage solutions

oracle.com/storage

Pillar Axiom storage systems

oracle.com/us/products/servers-storage/storage/san

StorageTek products

oracle.com/us/products/servers-storage/storage/tape-storage

Sun ZFS Storage Appliance

oracle.com/us/products/servers-storage/storage/nas

Leaders Run Oracle Storage

10 of the 10

Top Banks

Top Telcos

Top Pharmas

Top Governments

Get Better Results



ORACLE®


oracle.com/storage
or call 1.800.ORACLE.1

CONNECTED

One telecommunications company depends on **Oracle Exadata** to cover a nation.

Sometimes, when you tackle one problem, you solve another. Take the case of Digicel Haiti. Launched in 2006, Digicel Haiti is now the largest mobile communications operator for Haiti, serving 4 million customers across the country as well as several additional subscriber bases around the Caribbean and Pacific, including Jamaica and Trinidad.

BY DAVID A. KELLY







To run nightly data collection processes faster and support subscriber growth, Digicel migrated its 38 TB data warehouse to Oracle Exadata Database Machine. “We needed a system that could handle an increase in the amount of data,” says Marwan Salman, head of IT at Digicel.

In 2011, the organization set out to solve a problem with its data warehouse—the nightly data collection processes were simply taking too long for the company’s existing server infrastructure to process, and the business-critical daily reports were not reliably delivered on time. To overcome that problem and support substantial growth in its subscriber base, Digicel migrated its 38 TB data warehouse to Oracle’s Exadata Database Machine X2-2 Half Rack with high-capacity disks.

For Digicel, the benefits of implementing Oracle Exadata went far beyond solving the company’s immediate reporting problems. Indeed, the new solution significantly improved analysis and report performance across the board, with reports showing response time improvements of up to 16 times. And yes, daily reports are now processed and delivered on time each morning before the employees walk in the door, giving them the tools they need to strengthen and refine sales strategies and marketing plans.

But Oracle Exadata is also helping Digicel address a key strategic requirement for growth and customer retention that the company simply couldn’t address without a robust, highly scalable, and integrated engi-

neered system. On its old systems, Digicel couldn’t run antichurn reports. Now, antichurn reports allow the company to identify situations of potential customer loss and attempt to address them.

“To keep subscribers, we needed to understand their behavior,” says Marwan Salman, head of IT at Digicel. “Oracle Exadata and our ability to run new antichurn reports on the system are key to doing that.”

MEETING NEW CHALLENGES

A relative newcomer to the communications field, Digicel Haiti is part of the privately owned Digicel Group, which was launched in 2001 and currently serves more than 11 million subscribers across

several markets in the Caribbean, Central America, and the Pacific. Digicel Haiti’s goal is to provide customers with reliable cellular service. It’s an important mission, because for many Haitians, their cell phone is the only phone they have.

Digicel knew that its existing data warehouse would not scale effectively enough for the rapid expansion of services it needed to provide. “To be able to grow and accommodate all our new subscribers, we needed a system that could handle an increase in the amount of data,” says Salman.

SNAPSHOT

Digicel Haiti

digicelhaiti.com

Location: Port-au-Prince, Haiti

Industry: Communications

Employees: 1,000

Oracle products and services: Oracle Exadata Database Machine, SPARC Enterprise M5000 servers, Oracle Solaris, Oracle Enterprise Manager, Oracle Consulting Advanced Technology Services, Oracle Advanced Customer Support Services

The company also needed its data warehouse infrastructure to support significantly greater analysis of data by Digicel managers. The company's data warehouse is used to analyze call data records and generate customer-related reports for greater business insight.

STEPPING UP TO AN ENGINEERED SYSTEM

Upon its founding in 2006, Digicel Haiti's data warehouse was deployed on a set of Sun servers. Digicel later scaled up the data warehouse and server cluster to four servers as the company grew, but eventually it became clear that the solution would not be able to scale with Digicel's long-term growth—especially after the influx of new customers over the last two years.

Specifically, while Digicel could continue to scale out its server cluster, the company found that adding new servers didn't provide linear scalability. As they examined expanding the cluster even more, Salman and his team realized that an increasing portion of the server capacity would be consumed in overhead. That's when Digicel's IT group started looking for a solution that could grow more efficiently.

“Oracle Exadata has given us the ability to provide our business units with better and more-dynamic models of our customer base.”

—Marwan Salman, Head of IT, Digicel

After an in-depth evaluation of data warehouse options, the company chose Oracle Exadata, based primarily on the engineered system's performance.

Digicel was able to migrate its existing data warehouse to Oracle Exadata in less than two months, even under very difficult postearthquake conditions. The migration was also completed without interruption of day-to-day business operations. The result was a significant improvement over Digicel's legacy data warehousing infrastructure.

THE ORACLE EXADATA EXPERIENCE

Oracle Exadata has transformed Digicel's report operations. “Oracle Exadata really delivered the benefits we were looking for,” says Salman.

The migration also simplified the management of Digicel's IT environment, which encompassed more than 300 servers and many databases. “Oracle Exadata is much more manageable than our old system,” says Salman. “One reason is that we don't have to have external storage, so everything is managed inside one system.”

With Oracle Exadata, Digicel also improved recoverability in case of any system problems.

ORACLE EXADATA BUSINESS BENEFITS

Oracle Exadata not only helped Digicel reduce report times, save data center power and floor space, consolidate existing systems, increase

Award-Winning Solution

Digicel's Oracle Exadata Database Machine deployment was so effective that at Oracle OpenWorld San Francisco 2011, Digicel Haiti was recognized as the Data Warehouse Leader of the Year for Latin America and the Caribbean. Andy Mendelsohn, senior vice president of database server technologies at Oracle, presented the award.

In 2011 the Data Warehouse Leader of the Year awards recognized a leading organization in each of the following regions: Asia Pacific, Europe, Middle East, and Africa, Latin America/Caribbean, and North America.

The Data Warehouse Leader of the Year awards are part of the Oracle Excellence Awards and highlight organizations that have demonstrated excellent technical knowledge and innovative use of Oracle data warehouse technologies to the benefit of their businesses. For information about the Oracle Excellence Awards program, including descriptions of the different award categories and instructions for nominating candidates, see oracle.com/us/corporate/awards.

manageability, and provide greater recoverability; it also helped the company generate business growth.

“Oracle Exadata has given us the ability to provide our business units with better and more-dynamic models of our customer base, so they can create business offers that have the biggest positive impact,” says Salman.

Oracle Exadata also gives Salman and his team a way to ensure that critical corporate reports and key performance indicators can be put on managers' desks first thing in the morning.

And Oracle Exadata has made it possible for Digicel's IT group to be more responsive to the business. “It has enabled us to provide business units with solutions and reports that meet their needs and plans,” says Salman. “Now we can build new data mining procedures using the power of Oracle Exadata that would have been too complicated for our previous systems to handle. It's a huge benefit for our company.”

A PLATFORM FOR THE FUTURE

Digicel continues to grow. Over the past year, the company has expanded its customer base by 25 percent. Supporting that type of expansion is exactly why Digicel chose Oracle Exadata, with its ability to move from a half-rack implementation to a full-rack implementation.

Digicel is also in the planning stages of creating a secondary Oracle Exadata site for recoverability.

“The earthquake taught us a lesson on the importance of having a redundant site as well as being able to recover more quickly in the case of a major catastrophe,” says Salman. ◀

David A. Kelly (davidakelly.com) is a business, technology, and travel writer who lives in West Newton, Massachusetts.

NEXT STEPS

READ more about Oracle Exadata

oracle.com/exadata
bit.ly/AggUjv
bit.ly/zZ9z78



Great minds don't think alike.

To continue to lead the way in technology and science, we must develop the next generation of innovators. Steps we take today can empower young people to develop and test their own ideas and approach science, technology, engineering, and math not just as fields of study in school but as ways of understanding our world.

Join the Lawrence Hall of Science in inspiring and preparing kids from all backgrounds to do science.

lawrencehallofscience.org/brilliant



**THE LAWRENCE
HALL OF SCIENCE**
UNIVERSITY OF CALIFORNIA, BERKELEY

ORACLE JDEVELOPER AND ORACLE APPLICATION DEVELOPMENT FRAMEWORK

Personalized Dashboards

Extend an Oracle ADF Faces dashboard application to support persistent user personalization.

Many enterprise applications being built today require flexible, dashboard-style presentation of visual data—bar charts, graphs, and the like—showing various levels of information at a glance. Flexibility these days also assumes a certain amount of personalization. Users want to be able to rearrange how the information appears in any dashboard to best suit their needs.

The panelDashboard component of the Oracle Application Development Framework (Oracle ADF) Faces feature is a JavaServer Faces (JSF) layout container that Oracle ADF developers can use to implement such information dashboards. The panelDashboard component supports various types of child components that can be arranged into columns and rows of equal sizes.

This column will teach you how to change the Oracle ADF Faces panelDashboard component default layout behavior so that users can reorganize information areas according to their needs. You'll also learn how to persist the user's changes to the UI beyond the

current session by configuring Oracle Fusion Middleware's Oracle Metadata Services (MDS) feature in conjunction with the Oracle ADF Security feature so that the user's personalized UI is displayed whenever that person logs in. This column steps you through the process of making the necessary changes to a baseline application that ultimately supports personalization across application restarts.

The JSPX language was used for building the sample application's page definitions, but the techniques detailed in this column will also work for projects that use Facelets as the view declaration language rather than JSPX. Facelets, a feature of JSF 2.0, are fully supported in Oracle JDeveloper 11g Release 2. To learn more about Facelets, see bit.ly/AORpBT.

SAMPLE APPLICATION OVERVIEW

The sample application workspace contains a panel dashboard page that uses graphs and gauges in the Oracle ADF Faces Data Visualization Tools feature to display salary information in a static grid.

Instances of the Oracle ADF Faces af:panelBox component are used as immediate child components of the af:panelDashboard to form a two-row, two-column grid.

To follow the instructions provided in this column, download the o32adf-1549375.zip sample zip file at bit.ly/xfq2gK. The zip file contains two folders:

- **AdfDashboardCompleted** contains the Oracle JDeveloper 11g Release 2 (11.1.2.2) workspace for the completed application.
- **AdfDashboardHandsOn** contains the starter workspace.

GETTING STARTED

To step through the example, you need the studio edition of Oracle JDeveloper 11g Release 2 (11.1.2.2) or a later release, available as a free download from Oracle Technology Network. You also need an Oracle Database instance with the unlocked HR schema.

To set up and run the example application

1. Unzip the o32adf-1549375.zip file into a directory on your local machine.
2. Launch Oracle JDeveloper 11g Release 2. From the Oracle JDeveloper menu, select **File -> Open** and navigate to the directory containing the unpacked sample application.
3. Open the **AdfDashboardHandsOn** folder and then the **PanelDashboard** folder, and select the **PanelDashboard.jws** file.
4. Click **Open**. The workspace opens in Oracle JDeveloper.

Before attempting to run the application, you must change the database connection information to point to your HR database schema.

5. From the Oracle JDeveloper menu, select **View -> Database -> Database Navigator**.
6. Expand the **PanelDashboard** node.
7. Right-click the **hrconn** node and select **Properties** from the context menu. Edit



Figure 1: Dashboard starter application

the database connection information to work with your setup. Apply changes, and click OK to close the connection properties.

With a working connection to your database, you can now run the sample application.

8. In the Application Navigator, expand the ViewController project node to display the Application Sources and Web Content folders.
9. Expand the Web Content folder, click the DashBoardPage.jspx file within it to select it, right-click it, and select Run from the context menu.

If this is the first time you're running an application in Oracle JDeveloper, the Create Default Domain dialog box will appear, prompting you to create a password for the default Oracle WebLogic Server domain ("weblogic") associated with Oracle JDeveloper. Choose a password, enter it in the dialog box in the appropriate fields, and click OK. The application opens in the browser, as shown in Figure 1.

The dashboard application consists of the base panelDashboard component with four panelBox components. The topmost component (the "Salaries in Departments" graph) is not contained within the panelDashboard

component and thus remains fixed at the top of the UI.

Explore the sample application by clicking any bars in the "Salaries in Departments" graph, and you'll see that the two panel boxes in the middle of the page are updated to reflect the data change. Click a bar in the "Employee Salaries in selected Department" panel box to change the gauge graph data displayed in the "Employee Salary in Department Salary Range" panel box. This shows the selected employee salary in relation to the low-, medium-, and high-salary range for the selected department.

The two panel boxes at the bottom of the page both contain bar graphs that are not dynamically updated.

Select one of the four panel boxes and try to move it to another location in the dashboard, and you'll see that it cannot be moved. However, that's one of the first things you will enable as you step through this column.

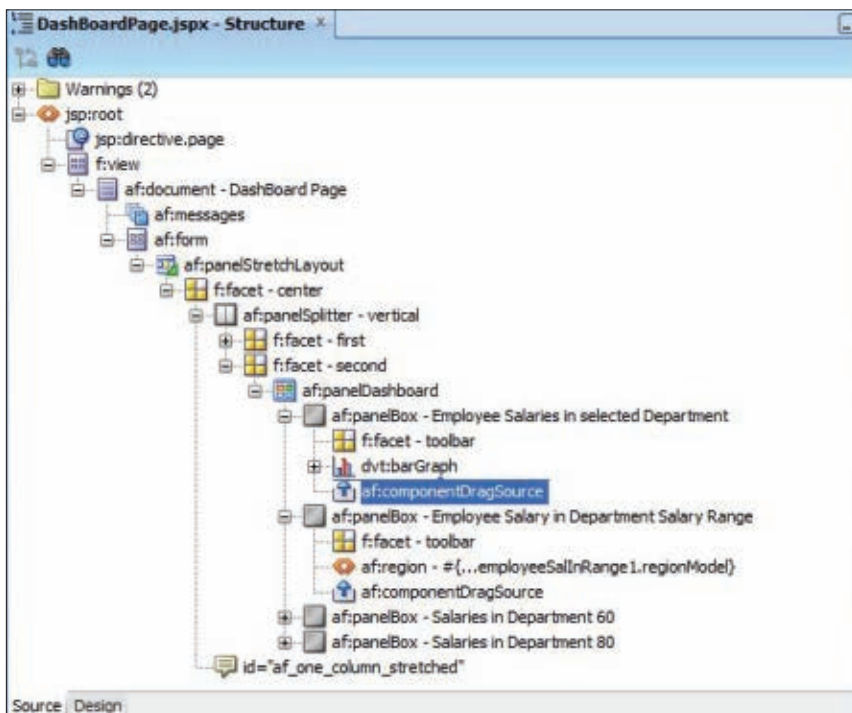


Figure 2: Oracle JDeveloper Structure window, showing component tree



Figure 3: Dragging a panel box at runtime

EXTENDING THE SAMPLE APPLICATION TO SUPPORT PERSISTENT PERSONALIZATIONS

The starter application is a one-size-fits-all dashboard UI. You'll extend the base functionality to enable users to drag and drop the charts and graphs. You'll also modify the application so that user changes are persisted, resulting in a personalized UI presented each time the application is launched. To implement drag-and-drop for the panelDashboard child components and to persist user-applied changes, application developers must make changes to the base application in several different layers, as follows:

- Enable drag-and-drop on the Oracle ADF Faces panelBox child components in the panelDashboard.
- Enable personalization by configuring MDS for the ViewController project and the panelDashboard component.
- Enable detection of panelBox component reordering and notification when such changes occur.
- Configure Security for user authentication so that any changes made to the UI are associated with the user who made them and will be presented to that user even after application restart.

Most of these tasks are accomplished with Oracle JDeveloper's declarative and visual tools or with a few changes to configuration files, as detailed in the remainder of this column.

ENABLING BASIC DRAG-AND-DROP FUNCTIONALITY

To enable panelBox components to be moved within the panelDashboard com-

ponent, the panelBox must be marked as a drag source. You can do this declaratively in Oracle JDeveloper by simply adding to each panelBox the necessary tag—an `af:componentDragSource` tag in this case. The `af:componentDragSource` tag is an Oracle ADF Faces behavior tag that encapsulates the JavaScript code (which you would otherwise have to write) for implementing drag-and-drop in HTML pages. Behavior

tags are just one example of how Oracle ADF Faces improves developer productivity by providing functionality that typically would require custom code.

To mark the panel boxes as drag sources, you drag the `componentDragSource` tag from the Component Palette and drop it onto each of the panelBox objects of the application. This gesture (dragging items from the Component Palette onto a page) is easy to do in the Structure window:

1. From the Oracle JDeveloper menu, select **View -> Structure** to open the Structure window, and select **View -> Component Palette** to open the Component Palette.
2. In the Application Navigator, double-click the `DashBoardPage.jspx` file to launch the Oracle JDeveloper visual editor.
3. In the opened editor, click in one of the panel boxes, such as the "Employee Salaries in selected Department" box.

In the Structure window, you should see the `af:panelDashboard` component and its contained `af:panelBox` children. Now move to the Component Palette:

1. Select **ADF Faces** from the menu.
2. In the search field above the menu, type `component drag source` and press Enter (or click the right-arrow search button). The search results return the **Component Drag Source (ADF Faces .Operations)** tag.
3. Drag the **Component Drag Source** tag from the Component Palette to the `af:panelDashboard`, and drop it onto the first `af:panelBox` child component in the `af:panelDashboard`.

Repeat this process for the other three panel boxes in the dashboard. When you finish, every panel box will have an `af:componentDragSource` as a child component (see Figure 2).

Right-click `DashBoardPage.jspx`, and select **Run** from the context menu to run the sample application again. When the application appears in the browser, select a panel box and move it within the dashboard, as shown in Figure 3. When you release the mouse button, Oracle ADF Faces will drop the panel box into position.

Nice! But we're not done yet. If the user quits the application at this point, the changes to the UI won't be retained. To

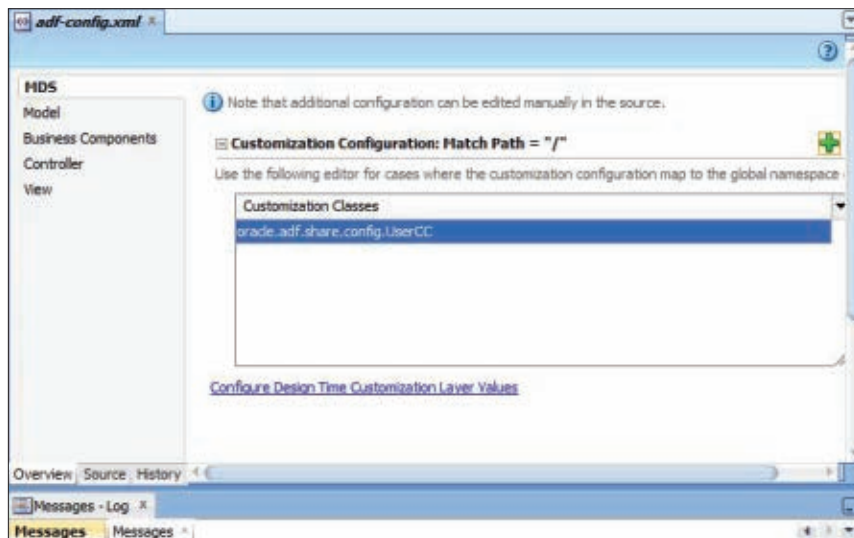


Figure 4: Editing the MDS customization layer for the UserCC class

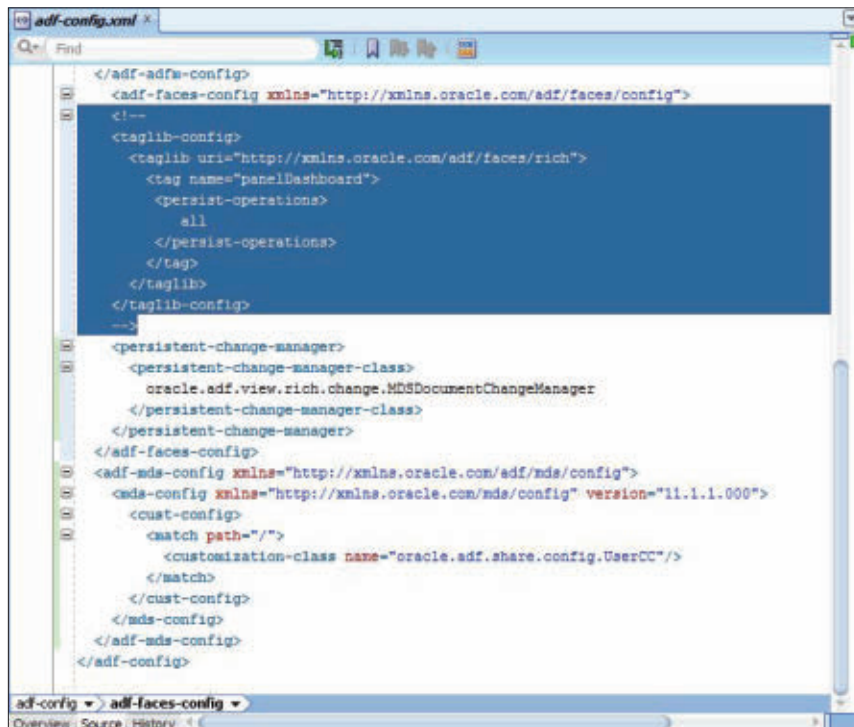


Figure 5: MDS configuration for panelDashboard in adf-config.xml

persist changes made by any user, we must configure personalization with MDS for this ViewController project.

ENABLING AND CONFIGURING PERSONALIZATION

MDS, the personalization and customization engine at the heart of the sample dashboard application, is a repository-based customization framework that is tightly integrated with the Oracle ADF request lifecycle. Two levels of personalization are available:

- “Seeded” personalization, in which different UI components and data are presented to different users, according to group membership, role, location, or other information
- Per-user personalization

For this sample application, we’ll set up MDS for per-user personalization. Incorporating MDS functionality into an Oracle ADF Faces application starts by enabling the functionality in the ViewController project layer. Follow these steps:

1. Click the **ViewController** project node in the Application Navigator.
2. Right-click **ViewController** and select **Project Properties** from the context menu.
3. In the Project Properties dialog box click the **ADF View** menu item. The ADF View panel appears.
4. Click the **Enable User Customizations** checkbox.
5. Click **Across Sessions using MDS**.
6. Click **OK** to save the settings and close the Project Properties page.

When you save these settings, the dashboard application’s web.xml deployment descriptor is (transparently) updated with the necessary context parameters for MDS. The application’s adf-config.xml is also updated with the change manager configuration.

However, the adf-config.xml file still requires a few other changes to support the personalization mechanism on a per-user basis.

In MDS, information is saved in layers identified by a Java object, the customization class. The user layer is defined by an existing Java object, the UserCC class. Internally, the UserCC class accesses the Oracle ADF Security context for the authen-

	Name	Password	Display name	Description
Test user 1	sKing	welcome1	Steven King	Leave this field empty.
Test user 2	dfaviet	welcome1	Daniel Faviet	Leave this field empty.

Table 1: Test data for sample application

ticated username (you’ll configure this later). The name of the authenticated user is a unique ID that MDS uses to store and retrieve data for a particular person.

For any application for which you want changes individually saved for an application user, you can configure the UserCC class.

To select and configure the UserCC class for the sample, follow the next few steps.

From the Application Navigator,

1. Open the **Application Resources** accordion panel.
2. Expand the **Descriptors** and **ADF META-INF** nodes.
3. Click **adf-config.xml** to select the file, right-click, and select **Open** from the context menu. adf-config.xml opens in the editor; the **Overview** tab appears by default.
4. In the editor, select **MDS** and click the green plus (+) icon above the **Customization Classes** entry (see Figure 4).
5. In the **Edit Customization Class** dialog box, type **UserCC** in the **Match Class or Package Name** field.
6. Select the **UserCC** class from the **Matching Classes and Packages** list, and click **OK**.

These steps add metadata for the UserCC class to the adf-config.xml file.

For our sample dashboard application, only changes to the child component order of the af:panelDashboard component will be persisted, so let’s configure that now. Unlike other Oracle ADF Faces components, the af:panelDashboard component does not have a menu option for configuration, so you must manually modify the adf-config.xml file and add the appropriate metadata to identify allowable customization. As a convenience, the application’s adf-config.xml file already contains the required metadata, nested inside comment tags (see Figure 5), so you need only remove the tags.

To manually edit the adf-config.xml file,

1. Click the **Source** tab at the bottom of

the editor where the adf-config.xml file is open

2. Remove the comments around the taglib-config element
3. Click the save icon in the Oracle JDeveloper tool bar to save all changes, and close adf-config.xml by clicking the x icon on the tab

DETECTING PERSONALIZATIONS

For many Oracle ADF Faces components, the component’s renderer notifies the framework when an end user personalizes a page, and then it uses the Oracle ADF Faces change persistence framework to render the page. However, because the panelDashboard’s drag-and-drop behavior is a feature of the af:componentDrag behavior tag rather than of the component itself, the panelDashboard component does not automatically notify the framework about user-applied changes. Thus, for this sample application, you must use Java and the change persistence API directly to create a bean that notifies the framework. That’s been done for you in the DndHandlerBean, which you can now associate with the listener on the panelDashboard.

The panelDashboard DropListener component receives an event notification about any drag-and-drop action performed on the child components. To respond to component drop events in the application, the DropListener property must be configured to point to a managed bean method:

1. Double-click the **DashBoardPage.jspx** file in the Application Navigator to open and display the page in the visual editor.
2. Click the binoculars (search) icon at the top of the Structure window to toggle open the search input field.
3. Type **panelDashboard** in the input field, and press Enter. Oracle JDeveloper will find and highlight the af:panelDashboard component (its node) in the page structure tree.
4. With the panelDashboard component

selected in the Structure window, open the Property Inspector by selecting **View -> Property Inspector** from the Oracle JDeveloper menu.

5. In the Panel Dashboard—Property Inspector pane, click in the Find field (next to the magnifier icon) to activate the search input field.
6. Type `DropListener` in the search field. As you type, the `DropListener` attribute becomes highlighted (in the Behavior section of the Property Inspector pane).
7. Select **Edit** from the `DropListener` property menu. The `Edit Property:DropListener` dialog box appears.
 - Select `DndHandlerBean` from the menu as the managed bean for the listener.
 - Select `dashBoardDropHandler` as the method.
8. Click **OK** to save the setting.

Upon drop event notification, the managed bean identifies the change in the `panelDashboard` child component order and writes this change to the Oracle ADF Faces change persistence framework, which passes this information on to MDS as per the configuration in the `adf-conf.xml` file.

For a better understanding of using Java to access the Oracle ADF Faces change persistence API, take a look at the code that constitutes the `DndHandlerBean`. From the Application Navigator, go to the **ViewController -> Application Sources** folder, in the `oramag.sample.dashboard.view.beans.dashboard` package.

At this point, you've implemented drag-and-drop for the `panelDashboard` component, so your application users can reorder child components at runtime. You've also set up MDS to persist changes to specific UI components. The one task remaining is to set up Oracle ADF Security to enforce user authentication so that UI changes can be saved for each user.

SETTING UP ORACLE ADF SECURITY AND TESTING THE COMPLETED APPLICATION

When users access the dashboard from their browser, you want them to be prompted for a username and password. Implementing this in the sample requires configuring Oracle ADF Security in a simple wizard-driven process. From the Oracle JDeveloper menu,

1. Select **Application -> Secure -> Configure ADF Security** to launch the Oracle ADF Security wizard. The `Enable ADF Security` pane (step 1 of 4) appears, enabling you to choose the type of authentication.
2. Click **ADF Authentication** (in the Security Model section), and click **Next** to continue. By default, `ADF Authentication` has `Http Basic Authentication` selected as its type on the next pane (step 2 of 4) of the wizard.
3. Leave `Http Basic Authentication` selected, and click **Finish** to save the settings and complete the wizard. A "Security Infrastructure Created" message appears.
4. Click **OK** to dismiss the message.

With Oracle ADF Security enabled for authentication, you can now create some test users. From the Oracle JDeveloper menu, select **Application -> Secure -> Users**. Click the green plus (+) icon in the Users section of the dialog box for adding users, and create the test user accounts listed in Table 1.

When you have finished adding these test users, click the save icon in the Oracle JDeveloper tool bar to save your work.

You can now run the completed dashboard application with some real-world personalization scenarios. Run the application by right-clicking the `DashBoardPage.jspx` file in the Application Navigator pane and selecting **Run** from the context menu. In the browser's login dialog box, authenticate as test user `sking` (`sking/welcome1`).

Rearrange the panel boxes by dragging and dropping them to new locations within the dashboard. (You'll notice that the dashboard functionality itself hasn't changed—the graph and the gauge are updated as they should be, and the result data in the two graphs at the bottom of the UI remains static.) When you've sufficiently modified the order of the UI items, close the browser.

Now rerun the application. Note that you can quickly restart the application from the Oracle JDeveloper log window, which contains the runtime URL, similar to the following:

```
http://localhost:7101/PanelDashboard-
ViewController-context-root/faces/
DashBoardPage.jspx
```

Open the log window if necessary by selecting **View -> Log Window** from the

menu, and click the URL to rerun the dashboard application.

When the browser login dialog box appears, authenticate as `sking/welcome1`. This time, the `panelDashboard` child component order is read from MDS and the UI is populated in the order in which test user `sking` left it.

Close the browser, and click the URL in the Oracle JDeveloper log window to relaunch the dashboard. When prompted for user credentials, authenticate as test user `dfaviet`. The application displays the dashboard components in their default positions, because test user `dfaviet` has not yet personalized the UI.

CONCLUSION

This column has stepped you through the process of using Oracle JDeveloper's declarative and visual tools to support persistent personalization of UI elements in a dashboard application. You can see how the combination of Oracle ADF Security for user authentication, Oracle ADF Faces behavior tags for implementing drag-and-drop, and the integrated change persistence framework in Oracle ADF Faces enable application developers to build personalizable application user interfaces with almost no need to write code. ◀

Frank Nimphius is a senior principal product manager for Oracle JDeveloper and Oracle Application Development Framework. He is a coauthor of *Oracle Fusion Developer Guide: Building Rich Internet Applications with Oracle ADF Business Components and Oracle ADF Faces* (McGraw-Hill, 2010).

NEXT STEPS

READ more about Oracle ADF

"Building Customizable Oracle ADF Business Applications with Oracle Metadata Services (MDS)"
bit.ly/wgLDfE

"Security for Everyone"
bit.ly/xgsD9R

LEARN more about Oracle ADF and Facelets
bit.ly/AORpBT

 **DOWNLOAD the Personalized Dashboards project**
bit.ly/xfq2gK



**Deliver
database results
with the complete
set of tools from Altova®**

The Altova MissionKit® is an integrated suite of database, XML, and data integration tools with powerful support for working with all major relational databases.

NEW in Version 2012:

- Support for JDBC drivers (in addition to ADO and ODBC options)
- Creation and editing of Oracle® PL/SQL packages
- Support for logical files in IBM® iSeries
- Streaming reading/writing of files for large-scale ETL applications

All MissionKit database tools support all of the following:

- Microsoft® SQL Server®
- Oracle®
- IBM DB2®
- Sybase®
- MySQL®
- Postgre SQL
- Microsoft Access™
- Cloud database systems

The Altova MissionKit includes multiple tools for working with databases:

DatabaseSpy®

SQL editor + database query, design & comparison tool

MapForce®

graphical data mapping, transformation & conversion tool

XMLSpy®

industry-leading XML editor with strong database integration

StyleVision®

database report and eForms design tool



 **Download a 30 day free trial!**

Try before you buy with a free, fully functional, 30-day trial from www.altova.com.



Scan with your smart-phone to learn more about these products

	DEPTNO	DEPTNAME	
1	A00	SPIFFY COMPUTER SERVICE DIV	0000
2	B01	PLANNING	0000
3	C01	INFORMATION CENTER	0000

ORACLE EXALYTICS IN-MEMORY MACHINE

Working with the Summary Advisor

Create in-memory aggregates for better performance on Oracle Exalytics In-Memory Machine.

If you create analyses and dashboards with Oracle Business Intelligence Enterprise Edition 11g, you know that users expect short response times even when querying datasources many terabytes in size. How do you provide short, consistent response times when querying such large amounts of data?

Oracle Exalytics In-Memory Machine, one of Oracle's latest additions to its engineered systems products, works alongside Oracle Exadata and Oracle Exalogic to provide a "speed of thought" business intelligence platform. Oracle Exalytics brings in-memory analysis to Oracle Business Intelligence solutions and provides a platform for analyzing terabytes of information with lightning-fast query responses.

Under the covers, Oracle Exalytics works by gathering statistics on your query workload and then recommending in-memory aggregates to improve query response time. To make this possible, Oracle Business Intelligence Enterprise Edition provides a set of cache management tools that are available only when you license Oracle Exalytics. One such management tool is the Summary Advisor, a utility within the Oracle Business Intelligence administration tool. Let's take a look now at how you can use the Summary Advisor with your Oracle Exalytics system to first analyze your query statistics and then to recommend in-memory aggregates, which you will then implement to improve the performance of queries on your system.

USING THE SUMMARY ADVISOR TO GENERATE RECOMMENDATIONS

The following example uses the SampleAppLite repository and catalog that come preinstalled with Oracle Business Intelligence Enterprise Edition 11g to demonstrate the Summary Advisor at work. This

article assumes that your Oracle Exalytics system has been installed and set up with a standard configuration: query caching is disabled, usage tracking is set up for your system, the required tables are registered in your repository, and the standard TT_AGGR_STORE Oracle TimesTen In-Memory Database connection is enabled.

1. Let's start by looking at the performance of the SampleAppLite repository and catalog. Using your Web browser, navigate to the Oracle Business Intelligence Web page, typically at [http://\[machine_name\]:9704/analytics](http://[machine_name]:9704/analytics), and log in with your user ID and password—biadmin/welcome1, for example.

From the Dashboards menu in the header area, select **Sample Lite -> QuickStart**. When the dashboard opens, click the **Overview** tab and then locate the **History Overview** pivot table on the dashboard page, as shown in Figure 1.

In this pivot table, drill into the **All Products** hierarchy to get a feel for the responsiveness of the analysis.

2. Now let's get ready to use the Summary Advisor to recommend aggregates to place in Oracle TimesTen In-Memory Database. Start by opening the Oracle Business Intelligence administration tool, and then connect to the repository running on your Oracle Exalytics server (**File -> Open -> Online**). When prompted, enter the credentials for the SampleAppLite repository and your Oracle Exalytics server. For example

```
Repository Password: Admin123
User:                biadmin
Password:            welcome1
```

3. To start using the Summary Advisor, select **Tools -> Utilities -> Oracle BI Summary Advisor** from the Oracle Business

	2008	2009	2010	Grand Total
All Products	1,500,000	1,700,000	1,800,000	5,000,000
BizTech	658,692	821,626	1,019,482	2,500,000
Communication	348,509	439,807	515,163	1,303,478
Cell Phones	169,997	233,361	263,727	667,085
Smart Phones	178,511	206,446	251,436	636,393
Touch-Screen	102,482	118,062	140,899	361,443
KeyMax S-Phone	76,029	88,383	110,537	274,950
Electronics	310,183	382,019	504,320	1,196,522
FunPod	542,613	556,666	400,721	1,500,000
HomeView	296,695	321,508	379,797	1,000,000

Figure 1: The History Overview pivot table

Intelligence administration tool. Note that this option will be listed on the Utilities menu only if you are using Oracle Exalytics.

4. The Summary Advisor first uses log activity to present aggregate recommendations.

Because you want to consider all activity on your system for these recommendations, click **Next** on the first three screens—Filter Logs - Logical Fact Tables, Filter Logs - Time Window, and Filter Logs - Execution Time Threshold—to accept the default settings, which will instruct the Summary Advisor to consider all log activity when making its recommendations.

5. The Targets screen enables you to select which physical database in your repository will be used to store the aggregates recommended by the Summary Advisor. Because you want to use Oracle TimesTen In-Memory Database, select the details that correspond to your Oracle TimesTen In-Memory Database instance, such as

Database Schema: "20 - RCU"..
"BIRCU_BIPLATFORM"

Connection Pool: "20 - RCU"."RCU CP"
Capacity (MB): 200000

Click **Add Target** to register the target database, and then click **Next**.

6. On the Select File Location screen, click **Browse** to specify a filename and a file system location for the script the Summary Advisor will create for you—`C:\TEMP\summary_advisor.sql`, for example—and click **Next**.
7. The Stopping Criteria screen enables you to place limits on how long the Summary Advisor will spend generating recommendations. Again, click **Next** to accept the default values, which place no limits on the recommendation generation process.
8. On the Miscellaneous screen, leave the **Max size of any single aggregate** setting at its default maximum value and click **Browse** to specify a location for the script that can be used to rerun the recommendation process—`C:\TEMP\rerun_advisor.xml`, for example—and click **Next**.
9. On the Run screen, click **Run** to generate these recommendations.

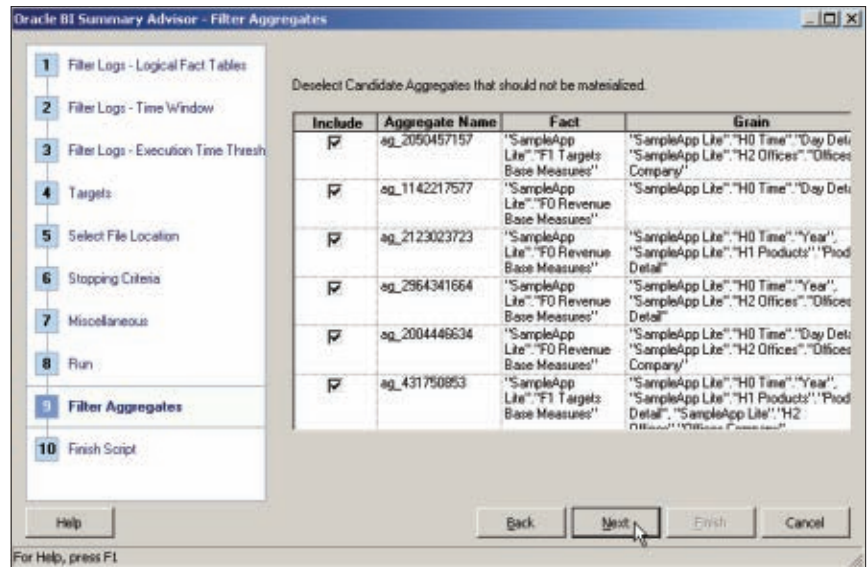


Figure 2: The aggregate recommendations

Depending on the settings you chose previously, the recommendation process will take a few moments to run, and then you will be presented with a list of aggregate recommendations, as shown in Figure 2.

10. On the Filter Aggregates screen, you can select which aggregates will be included in the script that the Summary Advisor will generate.

For this example, leave all the aggregates selected and then click **Next** to proceed to the final step.

11. On the final screen, Finish Script, click **Finish** to direct the Summary Advisor to create the specified scripts and close the wizard.

CREATING THE AGGREGATES

At this point, the Summary Advisor has created two scripts for you:

- `C:\TEMP\summary_advisor.sql` contains the aggregate definitions that you will need to run through the `nqcmd` utility to create or re-create the aggregate definitions in Oracle TimesTen In-Memory Database and the Oracle Business Intelligence Enterprise Edition 11g repository.
- `C:\TEMP\rerun_advisor.xml` contains the parameters you specified on the various Summary Advisor screens.

Let's now use the first script created by the Summary Advisor to create the aggre-

gates it has recommended. To do this, close the Oracle Business Intelligence administration tool; open a command-prompt session (`Start -> Run -> cmd.exe`); and run the `nqcmd` (`nqcmd.exe`) utility, using the following parameters:

- `-d`: an Open Database Connectivity (ODBC) datasource name for the Oracle Business Intelligence Server to which you want to connect
- `-u`: username of an administration user, such as `biadmin`
- `-p`: password of the user account
- `-s`: path to the aggregate definitions script created by the Summary Advisor

For example, for a default installation and the file, account, and password naming suggested in this article, the command-line entry is

```
C:\>"C:\Program Files\Oracle Business Intelligence Enterprise Edition Plus Client\oraclebi\orahome\bifoundation\server\bin\nqcmd.exe" -d exalytics -u biadmin -p welcome1 -s C:\TEMP\summary_advisor.sql
```

After the script runs, you should see a message indicating that several statements have been executed successfully. You can now reopen the Oracle Business Intelligence administration tool, connect to your repository, and navigate to the **Business Model**

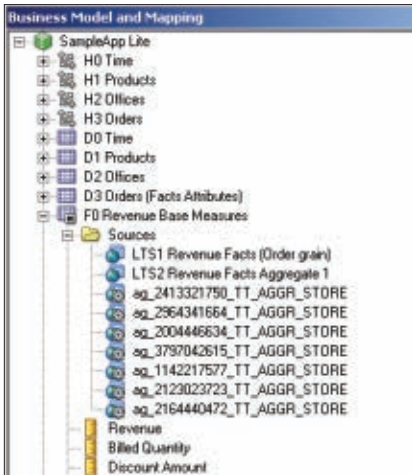


Figure 3: Aggregate tables mapped into the Business Model and Mapping layer

and Mapping layer to check out the new aggregate logical table sources that will have been automatically mapped in for you. For example, Figure 3 shows the FO Revenue

Base Measures logical table with seven new logical table sources mapped to it, representing the seven in-memory aggregates that were created by the Summary Advisor script.

Later on, if the volume of data loaded into your source database changes or new dashboards and analyses are created, you can generate a fresh set of aggregate recommendations by using the same Summary Advisor settings. Click Load Parameters from File on the first Summary Advisor screen to load your settings from the c:\TEMP\rerun_advisor.xml file you created in the previous steps.

TESTING THE IN-MEMORY AGGREGATES

So now that you have generated a set of in-memory aggregate recommendations and implemented them for your repository, you can test the impact on your dashboards and analyses. Using your Web browser, log in to the Oracle Business Intelligence Web page

again, navigate to the Overview dashboard, and try drilling into the History Overview analysis. You should now notice a much shorter response time as you drill into the hierarchies in your data. ◀



Mark Rittman is an Oracle ACE Director and cofounder and technical director of Rittman Mead. He writes for the Rittman

Mead blog, at rittmanmead.com/blog.

NEXT STEPS

READ more about
Oracle Exalytics In-Memory Machine
bit.ly/yzyjvQ

READ more Rittman
bit.ly/ywdgTP



When you provide a hot meal to a disaster victim, or give blood to someone you will never meet, train in first aid, or help a member of our military, you join the American Red Cross.

Your support makes the difference. Because of you, the Red Cross can respond to nearly 200 neighborhood emergencies every day.

Click, text or call to join today!



1-800-RED CROSS | redcross.org

Text REDCROSS to 90999 to give \$10

Text REDCROSS to 90999 to make a \$10 donation to the Red Cross. Charges will appear on your wireless bill, or be deducted from your prepaid balance. Msg & Data rates may apply. Reply STOP to 90999 to STOP. Reply HELP to 90999 for HELP. Full terms and privacy policy: redcross.org/m

ORACLE DATABASE

Working with Records

Part 7 in a series of articles on understanding and using PL/SQL



The Oracle PL/SQL language was designed to be a portable, high-performance transaction processing language that is tightly integrated with the SQL language. It is rare, indeed, to find a PL/SQL program that does not either read from or make changes to tables in a database. Tables are made up of rows of data, each consisting of one or more columns, so it stands to reason that Oracle Database would make it as easy as possible to work with those rows of data inside a PL/SQL program. And it does precisely that through its implementation of the *record*.

A record is a *composite* datatype, which means that it can hold more than one piece of information, as compared to a *scalar* datatype, such as a number or string. It's rare, in fact, that the data with which you are working is just a single value, so records and other composite datatypes are likely to figure prominently in your PL/SQL programs.

This article explores how you declare records, populate them with rows from a table, and even insert or change an entire

row in a table by using a record. It also takes a look at *user-defined* record types, which enable you to work with records that are not necessarily related to a relational table.

DECLARE A RECORD WITH %ROWTYPE

PL/SQL makes it very easy to declare records that have the same structure as a table, a view, or the result set of a cursor by offering the `%ROWTYPE` attribute.

Suppose I have an `employees` table in an application that looks like this:

```
SQL> DESCRIBE omag_employees
Name          Null?     Type
-----
EMPLOYEE_ID   NOT NULL  NUMBER(38)
LAST_NAME     VCHAR2(100)
SALARY        NUMBER
```

Each row in the table consists of three columns, and each column has its own datatype. The following query retrieves all the columns in all the rows in the table:

```
SELECT employee_id, last_name, salary
FROM omag_employees
```

I want to write a block of code that retrieves a single row of data from `omag_employees` for an employee ID and then work with the column values in that row. I *could* declare a variable for each column and then fetch into those variables, as follows:

```
CREATE PROCEDURE process_employee (
    employee_id_in IN
    omag_employees.employee_id%TYPE)
IS
    l_employee_id
        omag_employees.employee_id%TYPE;
    l_last_name
        omag_employees.last_name%TYPE;
    l_salary
        omag_employees.salary%TYPE;
BEGIN
    SELECT employee_id,
           last_name,
           salary
```

Pseudorecords in Database Triggers

Row-level triggers defined on tables can reference *pseudorecords* named `NEW` and `OLD` (you can override these default names with the `REFERENCING` clause of the trigger). They are called *pseudorecords* because they are similar in structure to a record defined on a table with `%ROWTYPE` but are restricted in their usage.

Both of the pseudorecords contain a field for every column in the table on which the trigger is defined. When you execute an `INSERT` or `UPDATE` statement, the `NEW` pseudorecord's fields contain the "post" values of the columns (the values after the `INSERT` or `UPDATE` has taken place).

When you execute a `DELETE` or `UPDATE` statement, the `OLD` pseudorecord's fields contain the "pre" values of the columns—how the row looks before the statement executes.

I can, for example, use pseudorecords to validate business rules, determine whether a column value has changed, and more. In the following trigger, I enforce a salary freeze; no one is allowed to get a raise during these tough economic times:

```
CREATE OR REPLACE TRIGGER
    omag_employees_freeze_trg
```

```
BEFORE INSERT
ON omag_employees
FOR EACH ROW
DECLARE
BEGIN
    IF :NEW.salary > :OLD.salary
    THEN
        RAISE_APPLICATION_ERROR (
            -20000,
            'Salary freeze in effect: ||
            ' no increases allowed!');
    END IF;
END omag_employees_freeze_trg;
```

There are, however, record features that do not apply to pseudorecords. I cannot, for example, pass a pseudorecord as an argument to a subprogram, even if the parameter for that subprogram is defined as `tablename%ROWTYPE`, where `tablename` is the name of the table that causes the trigger to be fired.

```

    INTO l_employee_id,
        l_last_name,
        l_salary
    FROM omag_employees
    WHERE employee_id = employee_id_in;
END;
```

(Note that I use suffixes in my parameters to indicate their mode. Here `_in` indicates an IN parameter.)

That is, however, an awful lot of code to write, read, and maintain. A much better approach is to fetch that row of data into a record, and the best way to declare that record is as follows:

```

CREATE PROCEDURE process_employee (
    employee_id_in IN
        omag_employees.employee_id%TYPE)
IS
    l_employee omag_employees%ROWTYPE;
BEGIN
    SELECT employee_id,
        last_name,
        salary
    INTO l_employee
    FROM omag_employees
    WHERE employee_id = employee_id_in;
END;
```

When this procedure is compiled, PL/SQL looks up the structure of the `omag_employees` table and defines a record that has a field for each column in the table, with the same name and datatype. By using `%ROWTYPE` to declare the record, I also tell Oracle Database that this procedure *depends on* the `omag_employees` table. If the database administrator changes the maximum length of the `last_name` column to 200, for instance, this procedure's status will be changed to `INVALID`. When the procedure is recompiled, the compiler will update the definition of the record in this procedure to match the table's new structure.

I can even shorten things further and write

```

CREATE PROCEDURE process_employee (
    employee_id_in IN
        omag_employees.employee_id%TYPE)
IS
    l_employee omag_employees%ROWTYPE;
BEGIN
    SELECT *
```

Answers to the Challenge

Here are the answers to the PL/SQL Challenge in last issue's "Error Management" article:

Answer 1:

Choices (a), (c), and (d) all raise `ORA-00001`, also known as `DUP_VAL_ON_INDEX` inside PL/SQL code.

Answer 2:

To make it possible for the `plch_proc` procedure to compile without error, change "AND" to "OR" in the WHEN clause of the exception section. Because only one exception can be raised at a time in a session, it doesn't make any sense to allow you to check for two (or more) exceptions in a single handler.

```

    INTO l_employee
    FROM omag_employees
    WHERE employee_id = employee_id_in;
END;
```

The `SELECT *` syntax tells Oracle Database to fetch all the columns in the table.

I can also use `%ROWTYPE` to declare a record that has the same structure as a `SELECT` statement in a cursor. This is especially helpful for fetching either a subset of columns from a table or columns from multiple tables. Here's an example:

```

DECLARE
    CURSOR no_ids_cur
    IS
        SELECT last_name, salary
        FROM omag_employees;
    l_employee no_ids_cur%ROWTYPE;
```

(Note that I usually add a "`_cur`" suffix to the names of my explicitly declared cursors.)

Whenever you are fetching data from a cursor into PL/SQL variables, you should declare a record based on that cursor with `%ROWTYPE` and fetch into that record. This way, when and if the `SELECT` list of the cursor changes, the number and type of fields in the record will change accordingly and everything will stay in sync.

WORKING WITH RECORD VARIABLES

Once you have declared a record in your block, you can both read and change the record's value. You can do this at the record level or by referencing individual fields of that record, with the same dot notation used

in SQL to refer to the column of a table.

So if I declare a record as follows,

```

DECLARE
    l_employee omag_employees%ROWTYPE;
```

I will be able to display the value of the `last_name` field of `l_employee` in the executable section of the block as follows:

```

DBMS_OUTPUT.put_line (
    l_employee.last_name);
```

I can *change* the value of a field, using an assignment operator:

```
l_employee.last_name := 'Picasso';
```

I can also perform the following record-level operations:

1. Set a record to `NULL`. This simple assignment will set the values of all fields to `NULL`.

```
l_employee := NULL;
```

2. Assign one record to another.

```

DECLARE
    l_employee1 omag_employees%ROWTYPE;
    l_employee2 omag_employees%ROWTYPE;
BEGIN
    l_employee1 := l_employee2;
END;
```

POPULATING RECORDS WITH DATA

Most of the time when you work with records, you will be assigning a row from a table to a record. You can also, however, assign values directly to individual fields or even to the record as a whole by using the PL/SQL assignment operator (`:=`). Let's look at examples of the ways to populate a record.

1. Declare a record with the same structures as those of the `omag_employees` table and then fill the record with the contents of one row from that table:

```

DECLARE
    l_employee omag_employees%ROWTYPE;
BEGIN
    SELECT *
    INTO l_employee
```

```

    FROM omag_employees
    WHERE employee_id = 100;
END;
```

2. Declare a cursor that fetches the last name and salary of all employees. Then use %ROWTYPE to declare a record that contains two fields: l_employee.last_name and l_employee.salary. Finally, open the cursor, fetch one row into the record, and close the cursor.

```

DECLARE
    CURSOR no_ids_cur
    IS
        SELECT last_name, salary
           FROM omag_employees;

    l_employee omag_employees%ROWTYPE;
BEGIN
    OPEN no_ids_cur;
    FETCH no_ids_cur INTO l_employee;
    CLOSE no_ids_cur;
END;
```

3. Populate a record by using native dynamic SQL. (Note: the SELECT statement is not dynamic; this is just to show that it is possible to populate a record with an EXECUTE IMMEDIATE ... INTO statement).

```

DECLARE
    l_employee omag_employees%ROWTYPE;
BEGIN
    EXECUTE IMMEDIATE
        'SELECT * FROM omag_employees'
        INTO l_employee;
END;
```

4. Populate the fields of a record by using assignments.

```

DECLARE
    l_employee omag_employees%ROWTYPE;
BEGIN
    l_employee.last_name := 'Renoir';
    l_employee.salary := 1500;
END;
```

Note that even though I define the record based on the table, I do not *have* to set the values of the record's fields from

the table. I might, for example, want to insert a *new* row into the employees table by using the record (see "Inserting and Updating with Records" for details).

5. Assign one record to another. Oracle Database supports record-level assignments, even the assignment of NULL to a record.

```

DECLARE
    l_old_employee omag_
employees%ROWTYPE;
    l_new_employee omag_
employees%ROWTYPE;
BEGIN
    l_new_employee := l_old_employee;
    l_old_employee := NULL;
END;
```

CURSOR FOR LOOPS AND IMPLICITLY DECLARED RECORDS

Suppose I want to write a program to display the last names of all employees. An elegant and simple way to do this in PL/SQL is to take advantage of the cursor FOR loop (which I discussed in part 2 of this PL/SQL 101 series). The cursor FOR loop is a variation on the numeric FOR loop, which looks like this:

```

FOR index IN low_value .. high_value
LOOP
    loop_body_statements
END LOOP;
```

The index is implicitly declared by Oracle Database as an integer and can be referenced only inside the body of this loop.

A cursor FOR loop has a similar structure but replaces a numeric range with a query:

```

FOR index IN ( SELECT_statement )
LOOP
    loop_body_statements
END LOOP;
```

Oracle Database also implicitly declares *this* loop index as well, but in the case of a cursor FOR loop, it declares the index as a record by using %ROWTYPE against the query in the loop header.

The following block uses a cursor FOR loop to fetch only the last name of each employee, deposit that name into a record,

and then display the value of the last_name field of that record:

```

BEGIN
    FOR employee_rec
    IN (SELECT last_name
        FROM omag_employees
        ORDER BY last_name)
    LOOP
        DBMS_OUTPUT.put_line (
            employee_rec.last_name);
    END LOOP;
END;
```

PASSING RECORDS AS PARAMETERS

You can define parameters based on record types, and you can therefore pass records as arguments to subprograms. Suppose I need to write a procedure that displays an employee. I could implement it as follows:

```

CREATE PROCEDURE show_employee (
    employee_id_in IN
        omag_employees.employee_id%TYPE,
    last_name_in IN
        omag_employees.last_name%TYPE,
    salary_in IN
        omag_employees.salary%TYPE)
IS
BEGIN
    DBMS_OUTPUT.put_line (
        employee_id_in
        || '-'
        || last_name_in
        || '-'
        || salary_in);
END;
```

But I can also avoid having to declare each of those individual parameters (imagine a 100-column table!) by passing a record:

```

CREATE PROCEDURE show_employee (
    employee_in IN
        omag_employees%ROWTYPE)
IS
BEGIN
    DBMS_OUTPUT.put_line (
        employee_in.employee_id
        || '-'
        || employee_in.last_name
        || '-')
```

```

    || employee_in.salary);
END;
/

```

Of course, as new columns are added to the table, their contents will not automatically be displayed by this procedure. So when you use %ROWTYPE to pass arguments to subprograms, make sure to review the subprogram logic after any change to the table.

INSERTING AND UPDATING WITH RECORDS

As you have seen, PL/SQL makes it very easy to populate a record from a row in a table. But what if you want to change the contents of a row in a table by using a record? PL/SQL offers special syntax in both the INSERT and UPDATE statements so that you can easily use records to perform those data manipulation language (DML) operations as well.

The most common form of an INSERT statement is

```

INSERT INTO table_name (column_list)
VALUES (expression_list)

```

where column_list is the list of columns that

will be populated on insert and expression_list is the list of expressions that will be assigned to their respective columns.

If I want to provide a value for each column in a table that has, say, 500 columns, writing and managing that code can become quite tedious. Inserting with a record comes in very handy in such a scenario.

To perform a record-level insert, simply leave off the parentheses around the record in the VALUES clause. Listing 1 demonstrates an insert of a single row into the omag_employees table that specifies each column individually. The following demonstrates the same insert, using a record:

```

DECLARE
    l_employee omag_employees%ROWTYPE;
BEGIN
    l_employee.employee_id := 500;
    l_employee.last_name := 'Mondrian';
    l_employee.salary := 2000;

    INSERT
        INTO omag_employees
        VALUES l_employee;
END;
/

```

So if you ever find yourself typing what feels like an endless list of variables in the VALUES clause of your INSERT statement, try using a record instead.

For updates, use SET ROW to update all the columns in a row from the record:

```

DECLARE
    l_employee omag_employees%ROWTYPE;
BEGIN
    l_employee.employee_id := 500;
    l_employee.last_name := 'Mondrian';
    l_employee.salary := 2000;

    UPDATE omag_employees
        SET ROW = l_employee
        WHERE employee_id = 100;
END;

```

Remember: this UPDATE sets the value of every column in the table, including your primary key, so you should use the SET ROW syntax with great care.

USER-DEFINED RECORD TYPES

So far you've seen how to declare a record variable based on a table or a cursor by using the %ROWTYPE attribute. You can also

Take the Challenge!

Each PL/SQL 101 article offers a quiz to test your knowledge of the information provided in it. The quiz questions appear below and also at PL/SQL Challenge (plsqlchallenge.com), a Website that offers online quizzes on the PL/SQL language. You can read and take the quiz here in *Oracle Magazine* and then check your answers in the next issue. If, however, you take the quiz at PL/SQL Challenge, you will be entered into a raffle to win an e-book from O'Reilly Media (oreilly.com).

I create and populate this table:

```

CREATE TABLE plch_parts
(
    partnum    INTEGER PRIMARY KEY,
    partname   VARCHAR2 (100) UNIQUE
)
/

BEGIN
    INSERT INTO plch_parts
        VALUES (100, 'Keyboard');

    COMMIT;
END;
/

```

Question

Which choices contain code I can use in place of the /*DECLARE*/ comment in the following block so that after the resulting block executes, "Keyboard" will be displayed?

```

DECLARE
    /*DECLARE*/
BEGIN
    SELECT *
        INTO l_part
        FROM plch_parts
        WHERE partnum = 100;

    DBMS_OUTPUT.put_line
        (l_part.partname);
END;
/

a.
l_part plch_parts%TYPE;

b.
l_part plch_parts;

c.
l_part plch_parts%ROWTYPE;

d.
CURSOR parts_cur
IS
    SELECT * FROM plch_parts;
l_part parts_cur%ROWTYPE;

```

Code Listing 7: Insert of a single row with each column specified

```

DECLARE
  l_employee_id  omag_employees.employee_id%TYPE
                := 500;
  l_last_name    omag_employees.last_name%TYPE
                := 'Mondrian';
  l_salary       omag_employees.salary%TYPE
                := 2000;
BEGIN
  INSERT
    INTO omag_employees (employee_id,
                        last_name,
                        salary)

  VALUES (
    l_employee_id,
    l_last_name,
    l_salary);
END;

```

declare your own, user-defined record types by using the TYPE... RECORD statement.

User-defined record types come in handy when you find yourself declaring “sets” of individual variables, such as

```

DECLARE
  l_name1          VARCHAR2 (100);
  l_total_sales1   NUMBER;
  l_deliver_pref1  VARCHAR2 (10);
  --
  l_name2          VARCHAR2 (100);
  l_total_sales2   NUMBER;
  l_deliver_pref2  VARCHAR2 (10);

```

Instead, why not create your own record type and then declare two records:

```

DECLARE
  TYPE customer_info_rt IS RECORD
  (
    name          VARCHAR2 (100),
    total_sales   NUMBER,
    deliver_pref  VARCHAR2 (10)
  );

  l_customer1    customer_info_rt;
  l_customer2    customer_info_rt;

```

(Note that when I declare types, I use a root “t” suffix and then add the “type of type.” Here I added “_rt” for *record type*.)

With this approach, you do more than avoid writing repetitive statements. You also document that those three pieces of information are all related to a customer. And once you’ve “moved up” to using a record,

you can pass that record as an argument or perform record-level operations, further reducing the volume of code needed to implement your requirements.

Another excellent time to use the TYPE... RECORD statement to create your own record type is when a field of your record needs to be a PL/SQL-specific type, such as BOOLEAN. If you use %ROWTYPE, the datatypes of all the fields will be constrained to SQL types.

Here’s an example of a record type that contains two Boolean fields:

```

DECLARE
  TYPE user_preferences_rt IS RECORD
  (
    show_full_name  BOOLEAN,
    autologin       BOOLEAN
  );

  l_user            user_preferences_rt;

```

Records are, themselves, PL/SQL-specific datatypes, so another nice feature of user-defined record types is that you can define a record type as a field in *another* record type. In the declaration section below, I have created one record type that holds the different numeric elements that make up a telephone number. I then create another record to hold the various telephone numbers for a salesperson:

```

DECLARE
  TYPE phone_rt IS RECORD
  (

```

```

    area_code      PLS_INTEGER,
    exchange       PLS_INTEGER,
    phn_number     PLS_INTEGER,
    extension      PLS_INTEGER
  );

  TYPE contact_rt IS RECORD
  (
    day_phone#     phone_rt,
    eve_phone#     phone_rt,
    cell_phone#    phone_rt
  );

  l_sales_rep      contact_rt;

```

COMPOSITE DATATYPES TO THE RESCUE!

PL/SQL’s support for records, one of several composite datatypes, enables you to write code that is simple, clean, and easy to maintain. Rather than work with long lists of variables or parameters, you can work with a record that *contains* all that information. User-defined records offer the flexibility to construct your own composite datatype, reflecting program-specific requirements that may not be represented by a relational table.

In the next article in this PL/SQL 101 series, I will explore another key composite datatype, the collection. Collections, PL/SQL’s implementation of arraylike structures, are used in some of the most important performance-related PL/SQL features, including FORALL and BULK COLLECT. ◀



Steven Feuerstein

(steven.feuerstein@quest.com) is Quest Software’s PL/SQL evangelist. He has published 10 books on

Oracle PL/SQL (O’Reilly Media) and is an Oracle ACE Director. More information is available at stevenfeuerstein.com.

NEXT STEPS

 **DOWNLOAD Oracle Database 11g**
bit.ly/fherki


TEST your PL/SQL knowledge
plsqlchallenge.com

READ PL/SQL 101, Parts 1–6
bit.ly/fc0u0j

ORACLE DATABASE

An Order of Sorts

Part 5 in a series on the basics of the relational database and SQL



Part 4 in this SQL 101 series, “Why WHERE Matters” (*Oracle Magazine*, March/April 2012), introduced readers to the WHERE clause of a SQL SELECT statement (a *query*) and the importance of filtering your data. The WHERE clause and the SELECT list tell the database *which* rows you want your SELECT statement to retrieve. Now that you know how to narrow the scope of the data a query fetches, you’re ready to learn how to *sort* (or *order*) the data. This article focuses on the SQL ORDER BY clause and how it behaves in conjunction with certain options and keywords to tell the database *how* you want retrieved rows to be sorted.

To try out the examples in this and subsequent articles in the series, you need access to an Oracle Database instance. If necessary, download and install an Oracle Database edition for your operating system from bit.ly/fherki. I recommend installing Oracle Database, Express Edition.

If you install the Oracle Database software, choose the installation option that enables you to create and configure a database. A new database, including sample user accounts and their associated schemas, will be created for you. SQL_101 is the user account you’ll use for the examples in this series; it’s also the schema in which you’ll create database tables and other objects. When the installation process prompts you to specify schema passwords, enter and confirm passwords for SYS and SYSTEM and make a note of them.

Whether you installed the database software from scratch or have access to an existing Oracle Database instance, download and unzip the SQL script available at bit.ly/yrEVIT and execute the script to create the example tables for the SQL_101 schema. (View the script in a text editor to get instructions on how to execute the script and information on what it does.)

The SQL queries in this article are executed against tables in the SQL_101 schema with Oracle’s SQL*Plus tool. In addition to discussing the ORDER BY clause, this article provides a closer look at SQL*Plus.

MAKING ORDER OUT OF DISARRAY

Oracle Database table data isn’t stored in any specific order, regardless of the order in which it was inserted into the database. To retrieve rows in either ascending or descending order by column, you must *tell* the database that you want to do so. For example, you might want to list all employees in the order they were hired,

display all employees in order of highest to lowest annual salary, or list the last names of all employees in the accounting department in alphabetical order. You retrieve sorted data by adding an ORDER BY clause to your SELECT statement. ORDER BY is always the last clause in a query.

Listing 1 shows a simple query of the EMPLOYEE table that doesn’t filter or order its result set. Compare Listing 1’s result set with the one in Listing 2. When you use the ORDER BY clause, the result set is in ascending order by default. Listing 2 displays the employees in the EMPLOYEE table sorted by last name in default ascending alphabetical order.

Code Listing 1: Simple query for listing all rows in the EMPLOYEE table

```
SQL> set linesize 32000
SQL> set feedback on
SQL> select first_name, last_name, hire_date, salary
       2   from employee;
```

FIRST_NAME	LAST_NAME	HIRE_DATE	SALARY
Roger	Friedli	16-MAY-07	60000
Betsy	James	16-MAY-07	60000
Emily	Eckhardt	07-JUL-04	100000
Frances	Newton	14-SEP-05	75000
Donald	Newton	24-SEP-06	80000
Matthew	Michaels	16-MAY-07	70000

6 rows selected.

Code Listing 2: Query that lists all rows in ascending alphabetical order by last name

```
SQL> select first_name, last_name, hire_date, salary
       2   from employee
       3  ORDER BY last_name;
```

FIRST_NAME	LAST_NAME	HIRE_DATE	SALARY
Emily	Eckhardt	07-JUL-04	100000
Roger	Friedli	16-MAY-07	60000
Betsy	James	16-MAY-07	60000
Matthew	Michaels	16-MAY-07	70000
Frances	Newton	14-SEP-05	75000
Donald	Newton	24-SEP-06	80000

6 rows selected.

You can obtain a result set in descending order by adding the DESC keyword immediately after the column name in the ORDER BY clause. The query in Listing 3 retrieves all employees from the most recent to the least recent date of hire. Note the DESC keyword in the ORDER BY clause. (You can use the ASC keyword to explicitly request ascending order, but it isn't necessary, because ascending order is the default.)

NAMES, NUMBERS, AND ARRANGEMENTS

Your ORDER BY clause does not need to explicitly name the column(s) by which you want to order the data. If you prefer, you can use the number of the column's position in the query's SELECT list. Listing 4 shows a query that retrieves all employees ordered from highest to lowest salary, using the sequence number (4) of the salary column in the query's SELECT list.

Code Listing 3: Query that retrieves and displays all employees in descending order by date of hire

```
SQL> select first_name, last_name, hire_date, salary
 2   from employee
 3   ORDER BY hire_date DESC;
```

FIRST_NAME	LAST_NAME	HIRE_DATE	SALARY
Roger	Friedli	16-MAY-07	60000
Betsy	James	16-MAY-07	60000
Matthew	Michaels	16-MAY-07	70000
Donald	Newton	24-SEP-06	80000
Frances	Newton	14-SEP-05	75000
Emily	Eckhardt	07-JUL-04	100000

6 rows selected.

Code Listing 4: Query that retrieves and displays all employees in descending order by column 4

```
SQL> select first_name, last_name, hire_date, salary
 2   from employee
 3   ORDER BY 4 DESC;
```

FIRST_NAME	LAST_NAME	HIRE_DATE	SALARY
Emily	Eckhardt	07-JUL-04	100000
Donald	Newton	24-SEP-06	80000
Frances	Newton	14-SEP-05	75000
Matthew	Michaels	16-MAY-07	70000
Betsy	James	16-MAY-07	60000
Roger	Friedli	16-MAY-07	60000

6 rows selected.

Code Listing 5: Query that retrieves and displays all employees, using multiple ORDER BY criteria

```
SQL> select first_name, last_name, hire_date, salary
 2   from employee
 3   ORDER BY hire_date DESC, 4, last_name ASC;
```

FIRST_NAME	LAST_NAME	HIRE_DATE	SALARY
Roger	Friedli	16-MAY-07	60000
Betsy	James	16-MAY-07	60000
Matthew	Michaels	16-MAY-07	70000
Donald	Newton	24-SEP-06	80000
Frances	Newton	14-SEP-05	75000
Emily	Eckhardt	07-JUL-04	100000

6 rows selected.

A query can sort on multiple columns, using multiple ascension and descension requests. You list the columns (or SELECT list column sequence numbers) in the ORDER BY clause, delimited by commas. The results are ordered by the first column, then the second, and so on for as many columns as the ORDER BY clause includes. If you want any results sorted in descending order, your ORDER BY clause must use the DESC keyword directly after the name or the number of the relevant column.

Listing 5 shows a result set that displays all employees in descending order of hire date (most recent to least recent), within which the employees are further sorted from lowest to highest salary and then by last name. Because ascending order is the default, the second column in Listing 5's ORDER BY clause doesn't need to include the ASC keyword; for the same reason, the ASC keyword associated with the last_name column is superfluous.

ENSURING THAT YOU ARE SET

Whenever you log in to the database with your username and password, you're creating a *session* in the database. You can change certain environment settings for your session that have no effect on other connected sessions (logged-in users). In Listing 1, note that a couple of SQL*Plus set commands appear before the SQL statement. These commands set system variables to customize the SQL*Plus environment settings for the current session. For example, in Listing 1, the following command sets the number of characters that SQL*Plus displays on a line before beginning a new line:

```
set linesize 32000
```

The shorter notation for this command is `set lines n`. This command is helpful if you want to ensure that the lines of your SQL query results do not wrap.

The other set command used in Listing 1 is

```
set feedback on
```

This command directs SQL*Plus to display a final count of the number of rows returned

in your result set. The shorter notation for this command is `set feed on`.

The last line displayed in Listing 2's result set is

```
6 rows selected.
```

This line appears because the SQL*Plus feedback setting was turned on (in Listing 1). If you do not want to see this final count of

the number of rows returned in your result set, you can turn this setting off with the `set feed off` command.

Your environment settings will apply to all of your current session's subsequent query execution results.

ORDERING THE UNKNOWN

Recall that a *null* value is one that is not known. Listing 6, for example, lists all

employees from the EMPLOYEE table with their manager values. Two of the six returned records have null values in the manager column.

When an ORDER BY clause sorts results in ascending order, any null values are displayed last by default. Conversely, if an ORDER BY clause specifies descending order for a column containing null values, as in Listing 7, the null values are displayed first by default. By using the NULLS FIRST or NULLS LAST option in the ORDER BY clause, you can override the defaults and explicitly specify how you want null values to be sorted. The example in Listing 8 uses the NULLS FIRST option to override the default display-nulls-last behavior of an ORDER BY clause.

Code Listing 6: Query that displays all employees with their manager values

```
SQL> select employee_id, first_name, last_name, manager
2   from employee
3  ORDER BY manager, last_name;
```

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	MANAGER
6567	Roger	Friedli	28
6568	Betsy	James	28
7895	Matthew	Michaels	28
1234	Donald	Newton	28
28	Emily	Eckhardt	
37	Frances	Newton	

```
6 rows selected.
```

Code Listing 7: Query that orders a column containing null values in descending order

```
SQL> select employee_id, first_name, last_name, manager
2   from employee
3  ORDER BY manager DESC, last_name;
```

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	MANAGER
28	Emily	Eckhardt	
37	Frances	Newton	
6567	Roger	Friedli	28
6568	Betsy	James	28
7895	Matthew	Michaels	28
1234	Donald	Newton	28

```
6 rows selected.
```

Code Listing 8: Query that orders a column containing null values with the NULLS FIRST option

```
SQL> select employee_id, first_name, last_name, manager
2   from employee
3  ORDER BY manager NULLS FIRST, last_name;
```

EMPLOYEE_ID	FIRST_NAME	LAST_NAME	MANAGER
28	Emily	Eckhardt	
37	Frances	Newton	
6567	Roger	Friedli	28
6568	Betsy	James	28
7895	Matthew	Michaels	28
1234	Donald	Newton	28

```
6 rows selected.
```

SORTING WITH DISTINCTION

When including an ORDER BY clause in a SQL SELECT statement, you will usually choose to sort by a column or an expression that's in the statement's SELECT list. However, you can also order by columns or expressions that are not in the SELECT list. Listing 9 displays a list of employees ordered by the most recent to the least recent date of hire, within which the employees are sorted alphabetically by last name. Although the sort occurs and displays correctly, only the employees' first and last names appear in the output, because `hire_date` is not in the SELECT list.

If you include the DISTINCT keyword in the SELECT list, only columns or expressions in the SELECT list may be included in the ORDER BY clause. As Listing 10 shows, an error will occur if a query using the DISTINCT keyword tries to order by a column not included in the SELECT list.

THE ERRORS OF OUR WAYS

You will inevitably make mistakes while learning to write SQL statements. Being able to interpret the Oracle Database error messages you receive is key to your understanding of SQL. Some error messages make it easy to understand what you've done wrong, whereas others are not so straightforward. The best approach is to try to resolve one error message at a time (a process called *debugging*).

Oracle Database tells you on which line of a query an error has occurred. Listing 10, for example, displays the following error message:

Code Listing 9: Query that orders by a column not included in the SELECT list

```
SQL> select first_name, last_name
2   from employee
3   ORDER BY hire_date DESC, last_name;
```

FIRST_NAME	LAST_NAME
Roger	Friedli
Betsy	James
Matthew	Michaels
Donald	Newton
Frances	Newton
Emily	Eckhardt

6 rows selected.

Code Listing 10: Query with DISTINCT fails because ORDER BY column is not in the SELECT list

```
SQL> select DISTINCT hire_date
2   from employee
3   ORDER BY manager NULLS FIRST;
ORDER BY manager NULLS FIRST
*
ERROR at line 3:
ORA-01791: not a SELECTed expression
```

Code Listing 11: Query that uses column aliases

```
SQL> select first_name first, last_name "Last", hire_date "Start Dt", salary "sal"
2   from employee
3   ORDER BY manager NULLS FIRST, first;
```

FIRST	Last	Start Dt	sal
Emily	Eckhardt	07-JUL-04	100000
Frances	Newton	14-SEP-05	75000
Betsy	James	16-MAY-07	60000
Donald	Newton	24-SEP-06	80000
Matthew	Michaels	16-MAY-07	70000
Roger	Friedli	16-MAY-07	60000

6 rows selected.

ERROR at line 3:
ORA-01791: not a SELECTed expression

Now you know that the database program is having difficulty with the following line:

```
3 ORDER BY manager NULLS FIRST;
```

If you add the MANAGER column to the query's SELECT list, as the error message implies, you will be able to rerun the statement successfully (assuming that the query contains no other errors).

Syntax errors will probably be the most common errors you make while learning SQL. The importance of carefully reading

(and rereading) your statements while debugging cannot be overemphasized. Simple typos, misplaced or missing commas, and unpaired single quotation marks (to name a few common mistakes) can cause a myriad of problems to which the solution might not be readily apparent.

USING ALIASES AND FORMAT MODELS

Sometimes you might want your query output to display meaningful headings for specific columns or expressions. You can make this happen by adding a column alias to any of the columns or expressions in your SQL statement's SELECT list. Listing 11 shows examples of the types of column aliases you

can use. Note that if a column alias contains more than one word—or you want it to appear in exact case (uppercase is otherwise the default)—you must enclose the alias in double quotation marks. As Listing 11 shows, you may use a column's alias in a query's ORDER BY clause, provided that it is not enclosed in double quotation marks in the SELECT list.

The online version of this article at bit.ly/xEN8dU includes information on SQL*Plus format commands and running SQL statements in SQL*Plus.

CONCLUSION

This article has shown you how to expand on simple SQL SELECT statements via the ORDER BY clause to order the data you display in a more meaningful way. You've seen how the DESC, NULLS FIRST, and NULLS LAST options behave and how null values are handled by default in an ORDER BY clause. You've also seen how the presence or the absence of the DISTINCT keyword in a SELECT list affects query execution if the ORDER BY clause includes a column that's not in the SELECT list.

The next installment in the SQL 101 series will look at character functions. ◀



Melanie Caffrey is a senior development manager at Oracle. She coauthored *Expert PL/SQL Practices for Oracle Developers and DBAs* (Apress, 2011) and *Expert Oracle Practices: Oracle Database Administration from the Oak Table* (Apress, 2010).

and DBAs (Apress, 2011) and *Expert Oracle Practices: Oracle Database Administration from the Oak Table* (Apress, 2010).

NEXT STEPS

READ online-only article content
bit.ly/xEN8dU

READ more about relational database design and concepts

Oracle Database Concepts 11g Release 2 (11.2)
bit.ly/aonqPP

Oracle Database SQL Language Reference 11g Release 2 (11.2)
bit.ly/yWtbz1

DOWNLOAD the script for this article
bit.ly/yrEVIT

ORACLE DATABASE

On Connection Pools, Cursor Differentiation, and Optimal Ordering

Our technologist cleans pools, explores cursor types, and looks for order in table creation.

We are facing severe performance issues on our server running Oracle Database 10g Release 2 (with 2 CPUs and 12 GB RAM). Automatic Workload Repository reports state that we have to tune our SQL statements, but application-level tuning is not an option, because we do not have the resources. Please refer to the top part of the Automatic Workload Repository report and suggest a way to improve performance (any way except SQL tuning). We are at our wit's end.

Note: Due to its size, I'm not including the submitted report from the Oracle Database Automatic Workload Repository tool. If you are interested in reviewing the report, it is available online at bit.ly/zNOvK3.

My initial answer to this question follows. The people who posted the question sent a follow-up comment, and I've also included that here. It is always nice to get feedback (good, bad, or indifferent) to see how things are going. This time, the feedback was good and reinforced the points I made below.

The questioners wrote, "... application-level tuning is not an option" and that they do not have the resources to do it, but I maintain that they have to do it. "Tuning a database" will deliver single- or even double-digit percentage decreases in application runtime, but these performance issues need orders-of-magnitude application runtime improvement. There is no magic here. If getting the response you wanted didn't require looking at the application and its overall design and implementation, *no one* would ever look at the application, its design, or its implementation.

If you want orders-of-magnitude decreases in runtimes, you will be looking at the application and its design and imple-

mentation. If you want things to run 10 percent faster, you *might* (and I stress the word *might*) get lucky and be able to tune the database without touching the application. It is exceedingly rare—and getting rarer as the database becomes more self-tuning.

The question ends, "... suggest a way to improve performance (any way except SQL tuning). We are at our wit's end." I can imagine that you are at your wit's end. A request like this is sort of like being asked to go from Europe to America on foot, blindfolded, in one day. Tuning an application *without being able to touch the application* is an impossible job. I suggest that you start by rejecting that tuning proposition and realizing that you are almost certainly going to have to touch the application.

In looking at your Automatic Workload Repository report, I see that the longest-running events at the system level are latch-related: cache buffers chains and library cache. Additionally, your CPU time was way up there. Concurrency-based waits are caused by one thing and one thing only: having many concurrently active sessions. If you had fewer concurrently active sessions, you would by definition have fewer concurrency-based waits (fewer people contending with each other). I see that you had 134 sessions in the database running on a total of 4 CPU cores. Because it is not really possible for 4 CPU cores to allow 134 sessions to be concurrently active, I recommend that you decrease the number of concurrent sessions by reducing the size of your connection pool—radically. Cut your connection pool to something like 32. It is just not possible for 4 cores to run 134 things at the same time; 4 cores can run only 4 things

at exactly the same time. If you reduce the concurrent user load, you'll see those waits go down, response time go down, and your transaction rates improve. In short, you'll find out that you can actually do more work with *fewer* sessions than you can with *more*.

I know that this fewer-does-more suggestion sounds counterintuitive, but I encourage you to watch this narrated Real World Performance video: bit.ly/AqMOvy. In this video, you'll see what happens in a test of a 12-core machine running transactions and decreasing the size of a connection pool from an unreasonable number (in the thousands) to a reasonable number: 96. At 96 connections on this particular machine, the machine was able to do 150 percent the number of transactions per second and took the response time for these transactions from ~100 milliseconds down to ~5 milliseconds.

Short of reducing your connection pool size (and therefore changing the way the application is using the database by queuing in the middle-tier connection pool instead of queuing hundreds of active sessions in the database), you would have to change your queries to make them request cache buffers chains latches less often. In short: tune the queries and the algorithms in the application. There is literally no magic here. Tweaking things at the system level might not be an option. Touching the application might have to be an option.

That was my original answer, and shortly after I posted it, the people who'd posed the question wrote back:

Thanks a lot for your response. . . . We reduced the connection pool, and we are already seeing better performance!

That drives home the fact that managing resources contributes hugely to your runtime performance. Far too often, I see databases with hundreds or thousands of concurrently connected sessions. Unless you have hundreds or thousands of CPUs, this just isn't reasonable in a three-tier architecture. You have connection pools in the middle tier, and you should be using them to limit the amount of concurrent workload you ask the database to do. It comes down to simple math in the end: if you have N cores, you can truly have only N things concurrently "active." If some of the things that are trying to be active are blocked, waiting for I/O, you might be able to have some number of active sessions larger than the number of cores on your machine, but the multiple would be in the single digits. For example, in the above-referenced video, the demo ran 8 connections \times 12 cores = 96 total connections. In many cases, I see people running 20, 30, 50, or more (many more, in some cases) connections per core. The only thing that can happen if you do that: your machine will end up in a state requiring you to pull the plug to regain control of it.

Watch your connections! Keep them reasonable. It makes a huge difference.

CURSOR VERSUS REF CURSOR

In an interview for an Oracle PL/SQL developer job position, the interviewer asked what the difference between a cursor and a ref cursor was and when I would appropriately use each of these. Could you please tell me how I could have effectively answered this question?

Technically, under the covers, at the most basic level, they are the same. They are both cursors and can be processed in the same fashion. In the following answer, I'll discuss only ref cursors and "regular" PL/SQL cursors (not DBMS_SQL cursors, for example).

In short, the salient differences between regular cursors and ref cursors are as follows:

- A ref cursor can be dynamically opened and is defined at runtime; a regular cursor is static and defined at compile time.
- A ref cursor can be passed to another PL/SQL routine or returned to a client. A regular cursor must be directly addressed (not passed) and cannot be returned to a client application.
- A regular cursor will be cached "open"

Code Listing 1: Creating a ref cursor and a regular cursor

```
declare
  type rc is ref cursor;
  cursor c is select * from dual;
  l_cursor rc;
begin
  -- here we have a ref cursor whose definition is determined at runtime.
  if ( to_char(sysdate,'dd') = 30 ) then
    open l_cursor for 'select * from emp';
  elsif ( to_char(sysdate,'dd') = 29 ) then
    open l_cursor for select * from dept;
  else
    open l_cursor for select * from dual;
  end if;

  -- Here we have a regular cursor. Its definition was decided
  -- at compile time and cannot change
  open c;
end;
/
```

Code Listing 2: Opening, passing, and returning a ref cursor

```
SQL> declare
  2   l_cursor sys_refcursor;
  3   l_rec   all_users%rowtype;
  4   l_query varchar2(300) :=
  5     'select *
  6       from all_users
  7       where rownum <= 5';
  8 begin
  9   open l_cursor for l_query;
 10   p( l_cursor );
 11   :x := l_cursor;
 12 end;
 13 /

data in procedure = SYS

PL/SQL procedure successfully completed.

SQL> print x

USERNAME          USER_ID   CREATED
-----
SYSTEM              5    05-SEP-10
OUTLN               9    05-SEP-10
DIP                 14    05-SEP-10
ORACLE_OCM         21    05-SEP-10
```

by PL/SQL; a ref cursor cannot be. In other words, using a ref cursor incurs a parsing penalty.

- A regular cursor can implicitly array-fetch 100 rows at a time. A ref cursor must use explicit array fetching. In other words, a regular cursor can more efficiently retrieve data by using significantly less code.
- A regular cursor can be defined outside of a procedure or a function as a global package variable. A ref cursor cannot be; it must be local in scope to a block of PL/SQL code.

Let's look at these points one by one.

First, a ref cursor can be dynamically opened and defined at runtime, in contrast to a regular cursor, which is static and defined at compile time. Listing 1 creates a ref cursor and a regular cursor.

Given the block of code in Listing 1, you see perhaps the most salient difference between cursors: no matter how many times you run that block, cursor C will always be SELECT * FROM DUAL. The ref cursor (L_CURSOR), however, can be anything. Another thing

that block of code demonstrates is that a ref cursor can be opened with either a query you construct at runtime or a query that is predetermined at compile time. The statement

```
open l_cursor for 'select * from emp';
```

demonstrates that the query defining the ref cursor L_CURSOR can be truly dynamic—constructed entirely at runtime. You could replace 'select * from emp' with any string—any PL/SQL variable of a string type—and be able to open the ref cursor at runtime.

The other two open calls for L_CURSOR in Listing 1 demonstrate that the query associated with the ref cursor can be dynamic at runtime. Depending on what day of the month it is, the ref cursor opens the static query SELECT * FROM DEPT or SELECT * FROM DUAL. The ref cursor is not bound to a fixed definition at compile time as the regular cursor C is. No matter how many times you run the code, the query associated with cursor C will be static and constant.

Now, looking at the second point, a ref cursor can be passed to another PL/SQL routine or returned to a client, whereas a regular cursor must be directly addressed (not passed) and cannot be returned to a client application. You can easily see what this means for a ref cursor. The following is a procedure that accepts a ref cursor as input, fetches a row from it, and prints out the row.

```
create or replace
procedure
p( l_cursor in out sys_refcursor )
is
    l_rec all_users%rowtype;
begin
    fetch l_cursor into l_rec;
    if (sql%notfound)
    then
        return;
    end if;
    dbms_output.put_line
    ( 'data in procedure = ' ||
      l_rec.username );
end;
/
```

Then, using SQL*Plus, I'll define a ref

Code Listing 3: Report comparing performance of regular and ref cursors

```
SQL> declare
2   l_rec      t%rowtype;
3   l_cursor  sys_refcursor;
4   begin
5       for x in
6         ( select * from t regular )
7         loop
8             null;
9         end loop;
10
11      open l_cursor
12      for
13      select *
14      from t refcur;
15
16      loop
17          fetch l_cursor
18          into l_rec;
19          exit when
20            l_cursor%notfound;
21      end loop;
22      close l_cursor;
23  end;
24  /
```

PL/SQL procedure successfully completed.

TKPROF report

SELECT * FROM T REGULAR

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	722	0.23	0.23	0	1748	0	72198
total	724	0.23	0.23	0	1748	0	72198

SELECT * FROM T REFCUR

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	72199	0.40	0.42	0	72203	0	72198
total	72201	0.40	0.42	0	72203	0	72198

cursor host variable in the "program":

```
SQL> variable x refcursor
```

And last I'll run the code in Listing 2, which will dynamically open a ref cursor, pass it to a PL/SQL subroutine, and then return it to the client program for further processing.

As you can see, the first row was fetched by the PL/SQL routine and printed out, and the remaining four rows were fetched and printed by the client program, not PL/SQL.

Those two differentiators show the flex-

ibility of a ref cursor over a regular cursor, so that might beg the question "Why use regular cursors at all?" The answer is twofold: performance and ease of programming. Let's look at performance first, starting with the facts that a regular cursor will be cached open by PL/SQL and that a ref cursor cannot and will not be. In other words, a parsing penalty is involved in using a ref cursor. Again, I can observe this easily by coding a simple routine that will open a regular cursor, fetch from it, close it, and then do the same to a ref cursor. Here is a routine to do that:

```

SQL> create or replace
2 procedure p
3 is
4   l_cursor sys_refcursor;
5   l_rec    dual%rowtype;
6   cursor c is select *
7     from dual d1;
8 begin
9   open c;
10  fetch c into l_rec;
11  close c;
12  open l_cursor for select *
13    from dual d2;
14  fetch l_cursor into l_rec;
15  close l_cursor;
16 end;
/

```

Procedure created.

Now, if I run this procedure in SQL*Plus 10 times, using “exec p;” with SQL tracing enabled, my TKPROF report will show the following:

```
SELECT * FROM DUAL D1
```

call	count
Parse	1
Execute	10
Fetch	10
total	21

```
SELECT * FROM DUAL D2
```

call	count
Parse	10
Execute	10
Fetch	10
total	30

Note how both SQL statements were executed 10 times apiece, which was expected, because I ran the code 10 times, but the query associated with the ref cursor was also parsed 10 times, as compared to the regular cursor’s query, which was parsed only once. In PL/SQL, regular cursors are cached in an open state in

Code Listing 4: 1,000-column table created, queried, and reported

```

declare
  l_create long := 'create table t ( c1 number';
begin
  for i in 2 .. 1000
  loop
    l_create := l_create || ',c'||i||' number default ' || i;
  end loop;
  execute immediate l_create || ');
end;
/

insert into t (c1, c1000 ) select rownum, rownum from all_objects;

exec dbms_stats.gather_table_stats( user, 'T' );

SELECT C1 FROM T

call      count          cpu    elapsed        disk    query    current    rows
-----
Parse         1         0.00         0.00         0         0         0         0
Execute      10         0.00         0.00         0         0         0         0
Fetch       7230         6.41        15.72        414610       420920         0       722790
total       7241         6.41        15.72        414610       420920         0       722790

SELECT C1000 FROM T

call      count          cpu    elapsed        disk    query    current    rows
-----
Parse         1         0.00         0.00         0         0         0         0
Execute      10         0.00         0.00         0         0         0         0
Fetch       7230         8.66        17.93        421260       3304860         0       722790
total       7241         8.66        17.94        421260       3304860         0       722790

```

the hope that you’ll execute them again. PL/SQL can then skip the entire parse process, which can lead to a significant reduction in CPU utilization. (If you don’t want to take my word for it, watch this video and be convinced: bit.ly/zPMeVw.) This PL/SQL cursor cache is controlled by the SESSION_CACHED_CURSORS database initialization parameter and is 100 percent transparent to you and completely nonintrusive. By that, I mean that PL/SQL’s caching of cursors open will not affect your programming at all. If the number of open cursors (not PL/SQL cached cursors) needs to exceed the number of currently free open cursor slots, PL/SQL will silently start closing the cached open cursors. PL/SQL is not taking a chunk of open cursor slots away from you—it just transparently uses the ones you are not currently using.

The next performance-related advantage of a regular cursor is also an ease-of-use feature. A regular cursor can implicitly array-fetch 100 rows at a time, whereas a ref cursor must use explicit array fetching. In other words, a regular

cursor can more efficiently retrieve data by using significantly less code. For example, if you were to copy the ALL_OBJECTS view into a table T and run the code in Listing 3 with tracing enabled, you would find something similar to that listing’s TKPROF report.

Note how many times the regular cursor was fetched from—722—about a 100th of the number of rows. PL/SQL uses an implicit array fetch of 100 rows for that implicit, regular cursor. For the ref cursor, you can see that the number of fetches is one more than the number of rows (the fetches must go one past the last row to verify that the last row has been fetched). That materially affects the CPU utilization of the query and may also affect the number of I/Os (the number in the query column in the TKPROF report) performed, as it did in this example. (If you are curious about why I/O was reduced, go to bit.ly/wwOWkI, which shows why with examples.)

Last, consider that a regular cursor can be defined outside of a procedure or function as a global package variable. A ref cursor cannot

Code Listing 5: Report for table with only 250 columns

```

SELECT C1 FROM T

```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	10	0.00	0.00	0	0	0	0
Fetch	7230	0.62	0.62	1117	94520	0	722790
total	7241	0.62	0.62	1117	94520	0	722790

```

SELECT C250 FROM T

```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	10	0.00	0.00	0	0	0	0
Fetch	7230	0.96	0.97	7	94520	0	722790
total	7241	0.96	0.97	7	94520	0	722790

be—it must be local in scope to a block of PL/SQL code. I consider this to be neither an advantage nor a disadvantage for either type of cursor. I am not a huge fan of global variables to begin with—I consider them to be a bad practice in general—so I personally never really take advantage of a regular cursor’s ability to be global. To show what this particular cursor difference means, I need just a small snippet of code:

```

SQL> create or replace package my_pkg
  2 as
  3     cursor global_cursor is
        select * from dual;
  4 end;
  5 /
Package created.

```

```

SQL> create or replace package my_pkg
  2 as
  3     global_cursor sys_refcursor;
  4 end;
  5 /
Warning: Package created with
compilation errors.

```

```
SQL> show err
```

```
Errors for PACKAGE MY_PKG:
```

```
LINE/COL ERROR
```

```

3/16 PL/SQL: Declaration ignored
3/16 PLS-00994: Cursor Variables

```

cannot be declared as part
of a package

See how a package with a regular cursor defined outside of a procedure or a function compiles successfully but a package with a ref cursor defined that way will not compile.

OPTIMAL ORDERING

Is there an optimal order for creating columns in a table? The DBA at my shop enforces a standard of putting VARCHAR2 columns last, and the likelihood of nulls is not considered. Is there a real-world benefit to this order?

Here’s a little background first. A row is stored in a manner similar to this in the database block:

```
[null flag][length][data][null flag]
[length][data]. . .
```

In order to get to the third column in a table, the database must parse through the first two columns (not pointers—the database must read the row byte by byte)—even if the columns are numbers, dates, whatever. The columns will be stored as a length field followed by the data, and the database must parse through them to get to the next column.

So, in general you should put the most frequently accessed columns first for the best performance.

However, there is something to be said for putting the column(s) most likely to be null last, because they will consume zero bytes.

If the database hits the end of a row before finding the Nth column, you know that that column is NULL or the default column value (for not-null columns added to the table with fast add in Oracle Database 11g).

However, there is a convention to put primary key columns first in all cases.

However . . . there are too many *however*s! And *in general*, it really isn’t going to matter too much. The bit about putting the most frequently accessed columns first would have the most weight in a data warehouse in which you are scanning over many millions of rows to build the answer (the repeated overhead of parsing over a lot of columns that are not interesting would add up). In an online transaction processing (OLTP) system, where you are processing tens of rows, though, it wouldn’t be as big a deal.

In a strange case involving dozens of columns that are almost always null, putting them last in a table with lots and lots of rows could save you quite a few bytes, but the table would have to be large and the columns would have to be all null and frequently all null.

So, my recommendation is to put the columns in the order that makes the most sense from an aesthetic and design perspective.

As for the DBA’s rule above (VARCHAR2 columns last and the likelihood of nulls not considered), it would do nothing for performance or manageability. In many cases, it could negatively affect performance. Just remember that pretty much all database data is stored in a stringlike fashion:

- A string will have a null flag and a length byte or length bytes (<=250 characters [1 byte], >250 characters [2 bytes]) followed by the data.
- A number will have a null flag and a length byte or length bytes (numbers are 0–22 bytes in length) followed by a varying amount of data.
- A binary_float will consume 5 bytes—a leading null flag and a length byte followed by 4 bytes of floating-point data.
- A binary_double will consume 9 bytes—a leading null flag and a length byte followed by 8 bytes of floating-point data.

The database must read the length byte to figure out how many bytes of that row con-

stitute that column, read over that column's data to get to the next length byte, and so on. You might want columns you frequently access to be first in a table, because it will take less CPU to access them. But you'd need to be accessing a lot of rows for this to be true!

Suppose you have a table with 1,000 columns and then you query the first column and the last column and compare the TKPROF reports for each. Listing 4 creates the table and shows the reports for the queries.

In this case, the CPU overhead was partially from the parsing of 1,000 columns and the chasing down of the chained-row piece (because any row with more than 254 columns will be stored in multiple pieces). You can see this row-chaining side effect in the report in the increased query column values, which resulted from the number of buffer gets processed.

If you change the 1,000 columns in Listing 4 to 250 columns to avoid having to

chase down the chained rows, you might see the result in Listing 5.

So, the farther down the CREATE list a column is, the longer it will take to retrieve that column. Having a rule that says, "The longer a column is, the farther down the create table statement it must be" doesn't seem to make sense. There are arguments for putting nullable columns last (for space savings) and for putting infrequently accessed columns last (for performance) but none that I'm aware of for putting longer columns last. ◀



Tom Kyte is a database evangelist in Oracle's Server Technologies division and has worked for Oracle since 1993. He is the author of *Expert*

Oracle Database Architecture (Apress, 2005, 2010) and *Effective Oracle by Design* (Oracle Press, 2003), among other books.

NEXT STEPS


ASK Tom

Tom Kyte answers your most difficult technology questions. Highlights from that forum appear in this column.
asktom.oracle.com

READ more Tom

Oracle Database Concepts 11g Release 2 (11.2)
bit.ly/aonqPP

Expert Oracle Database Architecture: Oracle Database 9i, 10g, and 11g Programming Techniques and Solutions, Second Edition
amzn.to/ckGXaR

 **DOWNLOAD Oracle Database 11g Release 2**

bit.ly/epBiUG

LEARN more about limiting connection pools

bit.ly/xR4I2T

decreasing connection pools to improve performance

bit.ly/AqM0vy

parsing

bit.ly/zPMeVw



Polar Bears International is a trusted voice focused on funding scientific research for the survival of this magnificent animal. Polar Bears International also funds educational programs that are inspiring people to reduce their carbon emissions.

Conservation through Research and Education
www.polarbearsinternational.org



Help Us Help the Polar Bear

PHOTO © R&C BUCHANAN

Virtualized Communities

A virtual community discusses virtualization and delivers real information.

Oracle's launch of Oracle VM 3.0 in 2011 underscored the increasingly important virtualization layer in data and application architecture. Nearly all computing platforms are being virtualized, and indeed 62 percent of Independent Oracle Users Group (IOUG) members who responded to a 2009 survey have deployed virtualized database servers to support development, testing, and production requirements. Without a more recent survey of the same IOUG members, we cannot be absolutely sure, but the number of IOUG members who have deployed virtualized database servers through 2012 has almost certainly increased.

OPTIMIZE PHYSICAL RESOURCES

A physical machine can be partitioned as one or more virtual servers. To the users and to applications running on a virtual server, operations appear to run on an independent physical server. Virtual servers enable us to optimize the use of our physical resources by sharing them and running a mix of applications, which distributes and balances the workload to get greater utilization of those physical resources.

No longer do we need a dedicated machine for each database or application server. We can use just enough of a machine's capacity to optimally handle the workload. We now have a greater capacity to have more servers running more applications without a linear increase in server hardware. Along with this ability to maximize resources comes a greater capacity to support high availability and disaster recovery.

Configuring a physical server is a time-consuming process. From the day it arrives on your loading dock to the day it is fully provisioned and usable can take days or weeks. Virtualization doesn't eliminate this initial provisioning of physical hardware, but once the hardware is provisioned and the virtu-

Join the IOUG Virtualization SIG and the associated LinkedIn group.

alization software is installed, provisioning a new virtual server happens in minutes or hours instead of days or weeks.

According to the same 2009 IOUG survey, 81 percent of respondents see an annual increase in the number of servers needed to support their businesses. Beyond the resource flexibility and time savings, virtualization can also address the need to increase server utilization while minimizing hardware costs. The benefits of virtualization were clear in 2009, but the experience and information needed to take full advantage of virtualization were not.

A VIRTUAL COMMUNITY GROWS

For several years, IOUG members have sought information on how best to run their Oracle software on virtual servers. More recently, perhaps as a consequence of engineered systems such as Oracle Exadata or the recent improvements in Oracle VM, our members have become more interested in Oracle VM. Not surprisingly, as has often been the case when a topic is being vigorously discussed and debated, some have wanted to formalize the discussion, idea sharing, and peer networking.

Thus, in the summer of 2011, the IOUG Virtualization SIG was born, and it is one of IOUG's fastest-growing special interest groups (SIGs). It includes proponents of both Oracle VM and other vendors' virtualization software. The SIG has a regular schedule of Webinars and a dedicated educational track in the IOUG Forum at COLLABORATE.

I am a virtualization novice, as you can tell from this column. But, as an IOUG member, when I needed to find out more about virtualization I did the three things any community member should do:

1. Join the IOUG Virtualization SIG and the associated LinkedIn group.
2. Go to the IOUG Virtualization SIG Website and access the Webinar archives to find and watch recorded Webinar presentations, and add the next Webinar to your calendar.
3. Contact your peers that are actively involved in the IOUG Virtualization SIG to benefit from their expertise.

Some of IOUG's most tenured and well-known members are active contributors to this SIG, underscoring its importance to our community. If understanding everything from the basics of virtualization to advanced implementations is important to you, then do what I did: join the IOUG Virtualization SIG.

While you are joining, have a look at the other IOUG SIGs. Perhaps you will find another growing virtual Oracle community to join and begin taking advantage of an IOUG SIG. ◀



Andy Flower

(andy_flower@ioug.org) is president of IOUG and has been an active volunteer with the organization since 1998.

In his day job, he is an information management and business intelligence consultant with Right Triangle Consulting.

NEXT STEPS

LEARN more about the IOUG Virtualization SIG and other SIGs
ioug.org

Print. Digital. Mobile.

Read Your Favorite Technology Magazines on Your Favorite Device.



Download the most recent issue **FREE** on Apple iTunes.

Or get your **FREE** print or digital subscription.

Go to: oracle.com/us/corporate/publishing/subscribe

ORACLE[®]

Enterprise Servers: Today and Tomorrow

Workload and new trends drive markets for commodity and high-end servers.

Oracle Magazine spoke with Rich Partridge, senior server analyst for Ideas International, about the ongoing changes in the server market and the future of enterprise servers.

Oracle Magazine: What's the state of the enterprise server market today?

Partridge: Organizations have different options depending on the type of workload. Clearly, for routine work, they are looking for servers that will provide performance at an attractive price. That's where x86 systems have become the dominant force. With x86 systems, you get a lot of performance at a very attractive price. That's why they've become the standard computing solution for doing things that are routine.

But there's another area of the market where companies are clearly looking for a competitive edge. They have unique requirements and are willing to spend a little bit more to get well-integrated software, servers, and storage. That's where SPARC processors and servers fit in very nicely.

Oracle Magazine: What's important to businesses at the higher end of the server market?

Partridge: At the very high end, everyone is looking for a bulletproof, highly reliable, and resilient server. When you think about it, those are the things that had been hallmarks of mainframes for decades. High-end servers need to provide resiliency—that ability to keep on computing even if there is some kind of a hiccup.

Oracle Magazine: What's the value of an integrated server?

Partridge: With higher-end systems there is a lot of data coming from all of the different business processes, from managing inventories to analyzing data for trends for future products. So in these systems, there is a real need to integrate a lot of different

“Oracle can identify where the bottlenecks are and address them from both a hardware and software perspective.”

applications—a lot of different usages of the same massive amounts of data—and that requires someone to think about how all of the pieces go together.

At the high end of the server market, many companies recognize the value of an integrated system that can exploit the interaction between different business processes. Oracle has a real advantage there, since it provides both hardware and software. With these types of high-end solutions, Oracle can identify where the bottlenecks are and address them from both a hardware and software perspective. For example, one bottleneck is cryptographic work, and in response, Oracle has determined that it's really important to put cryptographic processing into its SPARC T-Series processors.

Oracle Magazine: What's your current impression of the Oracle roadmap for SPARC and Oracle Solaris?

Partridge: There was some skepticism when Oracle first took over Sun, but the success of the first few products to be released after the combination of Oracle and Sun—particularly the SPARC T4 processor—has shown that Oracle is committed. What they've done so far with SPARC and Oracle Solaris is pretty impressive.

Of course, Oracle's ability to deliver and complete the roadmap will be the real test. Oracle clearly wants to be a leading player in the high-end systems using SPARC, and it's got a map that is aggressive, which is going to be inspiring to its customers.

Oracle Magazine: What do you think the future holds for enterprise servers?

Partridge: You are going to see continued improvement in the price performance and the raw performance of x86-based systems. At the same time, as the cost of computing continues to come down, there will be so many more computing problems that people never thought they could afford to attack before. Some of these problems will involve huge amounts of data to analyze so that customers can identify salient opportunities in ways never before possible. Solving these problems will be the future role of high-end enterprise servers.

The high end will continue to grow because we are not going to be doing only the kind of computing that was done 40 years ago, much of which was based on straightforward automation. Organizations will increasingly be exploring and exploiting trends within the data that were simply not visible without high-end computers. ◀

David A. Kelly (davidakelly.com) is a business, technology, and travel writer who lives in West Newton, Massachusetts.

Ideas International (ideasinternational.com) provides enterprise research, insight, analysis, and tools to the IT industry.

NEXT STEPS

READ more about Oracle's Sun servers

oracle.com/us/products/servers-storage/servers

Engineered Storage

Preintegrated storage platforms combine the best hardware with software to simplify operations and improve performance.

Oracle Magazine spoke with Benjamin S. Woo, program vice president for Worldwide Storage Systems at International Data Corporation (IDC), about how organizations are looking at the present and future of enterprise storage with engineered storage solutions.

Oracle Magazine: What are the top priorities and solutions in enterprise storage today?

Woo: From a technology perspective, today's storage solutions are a lot more focused on performance. Additionally, there are a lot more types of media available, from tape to multiple types of magnetic disks such as large-capacity SATA [serial ATA] drives to very fast solid-state devices. Today's intelligent storage platforms understand the types of data being stored and the performance characteristics that organizations require from the storage system.

Oracle Magazine: How does the type of I/O affect the choice, cost, and operation of enterprise storage solutions?

Woo: Customers often focus on cost per gigabyte, but they also need to consider cost per I/O operation. IDC research reveals that solid-state technology can be 10 times less costly per I/O operation than traditional magnetic media. Storage vendors provide different tiers of storage corresponding to different types of media, along with built-in software to predict and analyze which types of I/O operations equate with which types of media.

Oracle Magazine: How are today's engineered storage solutions addressing these I/O types?

Woo: Engineered storage systems have different degrees of sophistication to match the storage media to the optimum I/O, performance, and capacity needs of each application. Engineering takes different forms, from controlling the basic array technology that is integrated into any storage system, to

“Today's intelligent storage platforms understand the types of data being stored.”

managing virtualization or thin provisioning, to provisioning storage capacity based on certain needs, to organizing data into storage pools. These levels of sophistication make the system progressively more complex. However, they also provide more granularity and more control over how I/O capabilities can be distributed to the servers.

Oracle Magazine: How are storage administrators managing today's different workloads?

Woo: No human can keep up with the complexity of managing various application workloads, particularly in large virtual environments. Even a single application can have several I/O profiles, from transaction processing to archiving. Automated storage environments prioritize these workloads based on the needs and policies that the IT organization establishes. They also remove any possibility of human error by applying the right techniques and policies at the right time, while handling millions of operations per second.

Oracle Magazine: What is the status of tape in enterprise storage, and what are the leading advancements in tape storage technology?

Woo: The prediction about the death of tape is about as long in the tooth as the death of the mainframe. Tape is a critical part of any organization because it's still the cheapest medium on which to store data for the long term. Secondly, tape is unique in that, unlike online backups and snapshots, data stored

on tape is truly offline and the media can be stored offsite. One significant advancement in tape technology is LTO [Linear Tape-Open] version five and Linear Tape File System [LTFS], which enables files to be written to tape in a self-describing format. LTFS presents a tape in the form of a disk. You can search on it and follow the directory structure, which enables you to be more precise with what you want to restore or how you want to take action based on what you have stored. Thus, for the first time ever, we can refer to data on tape just by referencing data on disk.

Oracle Magazine: What are the most interesting storage changes and advancements on the horizon?

Woo: The greatest change is in your pocket. The smartphone is changing the way data is created, consumed, and stored. For the first time in computing history, we no longer just have to manage what's in the data center but also what's *not* in the data center—and at some point that set of data will be the greater of the two. That's frightening for a lot of organizations, and it makes the ability to understand the flow of data and its persistence and its usage and so forth absolutely critical to any evolving storage strategy. ◀

David Baum (david@dbaumcomm.com) is a freelance business writer based in Santa Barbara, California.

IDC (idc.com) is a global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets.


NEXT STEPS

READ more about Oracle storage
oracle.com/storage

Exadata

20x Faster

Replaces IBM Again



**Giant European Retailer
Moves Databases from
IBM Power to Exadata
Runs 20 Times Faster**

ORACLE®

**For more details
oracle.com/EuroRetailer**