Analyze, Illustrate, Fix and Run It

Troubleshoot and Manage Your PL/SQL

Maintaining legacy applications, documenting old and unfamiliar code or digging into someone else's code? A difficult and time-consuming task. Is there a solution?

Yes! Conquest's *Clear*SQL for Oracle is the solution.

A very powerful and advanced PL/SQL analyzer. The perfect enterprise code review and quality control tool.

WORLD CLASS TOOLS FOR EVERYBODY'S BUDGET

Don't miss ...

Buy One, Get One

Oct. 1 - Dec. 31, 2012

*Clear*SQL 6.5 NEW

Analyzes 1000s of PL/SQL scripts with a single click. Creates Flowcharts, Call Tree diagrams, CRUD matrices from PL/SQL to allow you to instantly understand code.

- quickly detects code errors, code review rule violations and other PL/SQL code anomalies
- creates detailed Code Metrics reports from PL/SQL *(McCab, Halstead, Maintainability Index, and more..)*
- documents PL/SQL creating clickable reports (*HTML*)
- summarizes PL/SQL on a single page *(status overview)*
- offers Command Line feature for batch analysis
- creates PSEUDOCODE from existing PL/SQL
- monitors code complexity trends during development cycles to ensure maintainability
- imports Oracle forms and PL/SQL libraries v. 6 – 10
- ideal when implementing enterprise coding standards *(e.g. code formatting, DB objects naming rules,...)*

***Clear*SQL**
***Clear*DB Documenter**
SQL*Detective*

software solutions
CONQUEST

www.conquest-us.com

ORACLE

ORACLE MAGAZINE

ORACLE.COM/ORACLEMAGAZINE

ORACLE OPEN WORLD

Engineered to Educate, Explore, Engage

2012's Oracle OpenWorld and JavaOne conferences brought technologists to San Francisco for a week of learning, networking, and looking forward

Built for Speed

Oracle unveils new releases of Oracle Exadata Database Machine and Oracle Exalogic Elastic Cloud

APPS ON ORACLE EXADATA

CONSOLIDATE, CUT COSTS, AND SUCCEED WITH ORACLE EXADATA

Reach new heights

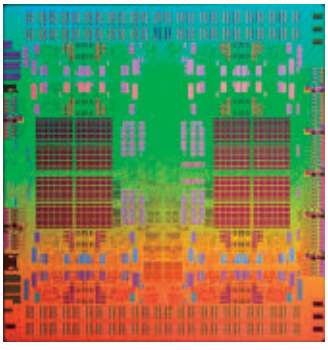

SPARC64 X: Fujitsu's latest SPARC processor carrying the inheritance from the world's top-level supercomputer to provide unprecedented performance and scalability. Fujitsu, together with Oracle, ensures that SPARC systems interoperate well with next generations of Oracle Solaris, Oracle Database, and Oracle software products.

shaping tomorrow with you

FUJITSU

ORACLE MAGAZINE

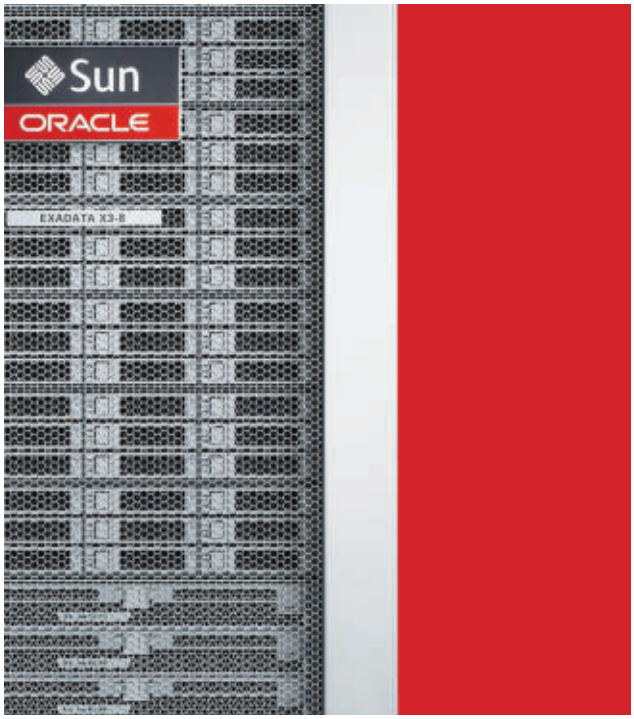

APPS ON ORACLE EXADATA /36

When business growth means that your current business applications platform deployment can no longer support your business, consolidation is certainly worth consideration. By consolidating business applications with Oracle Exadata, organizations are simplifying their information technology; lowering costs; and improving system performance, scalability, and reliability. Find out how Praxair and Alliance Data are benefiting by running their applications on Oracle Exadata. *—By David A. Kelly*

Engineered to Educate, Explore, Engage /32

Oracle OpenWorld, JavaOne, MySQL Connect, and several other events took San Francisco by storm, September 29–October 5. This content-focused week—opening with MySQL Connect on Saturday and closing with the final day of Oracle Customer Experience Summit @ OpenWorld the following Friday—offered a wealth of learning and networking opportunities, with a little rock, blues, ska, and alt-country added in for good measure. Relive the conference—or see what you missed—and start planning your trip to San Francisco next September. *—By Karen Shamban*

Built for Speed /22

At Oracle OpenWorld, Oracle unveiled new releases of Oracle Exadata Database Machine and Oracle Exalogic Elastic Cloud. Learn how these engineered systems deliver hardware and software engineered to work together and in-memory, speed-of-thought compute power to private and public clouds. *—By David Baum*

Up Front / 6

At Oracle / 11

Community / 26

Technology / 41

Comment / 68

EDITORIAL

Editor in Chief
Tom Haunert tom.haunert@oracle.com

Senior Editor
Caroline Kvitka caroline.kvitka@oracle.com

Managing Editor
Jan Rogers jan.rogers@oracle.com

Associate Editor
Patty Waddington

Contributing Editor and Writer
Blair Campbell

Technology Advisor
Tom Kyte

Contributors
Marta Bright, Jeff Erickson, Fred Sandsmark, Rich Schwerin, Leslie Steere

DESIGN

Senior Creative Director
Francisco G Delgadillo

Senior Design Director
Suemi Lam

Design Director
Richard Merchán

Contributing Designers
Jaime Ferrand, Nicholas Pavkovic

Production Designers
Sheila Brennan, Kathy Cygnarowicz

United States Postal Service Statement of Ownership, Management, and Circulation.

Oracle Magazine, ISSN 1065-3171, is published bimonthly (six issues each year) with a free subscription price by Oracle Corporation. The office of publication and the general business office are located at 500 Oracle Parkway, Redwood City, CA 94065. The publisher is Jeffrey Spicer; the Editor is Thomas Haunert; the Managing Editor is Jan Rogers; and the owner is Oracle Corporation, all at the above address. Known stockholders owning or holding 1 percent or more of the total amount of stock is Lawrence J. Ellison at the above address. There are no known Bondholders, Mortgagees, or Other Security Holders Owning or Holding 1 Percent or More of Total Amount of Bonds, Mortgages, or Other Securities. The Issue Date for the following data is September/October 2012. Average No. Copies Each Issue During Preceding 12 months: Total Number of Copies (Net press run): 420,634; Paid/Requested Outside-County Mail Subscriptions stated on Form 3541: 269,324; Sales Through Dealers and Carriers, Street Vendors, Counter Sales, and Other Non-USPS Paid Distribution: 110,862; Total Paid and/or Requested Circulation: 380,186; Free Distribution by Mail: 3,750; Free Distribution Outside the Mail: 35,134; Total Free Distribution: 38,884; Total Distribution: 419,070; Copies Not Distributed: 1,564; Total: 420,634; Percent Paid and/or Requested Circulation: 90.7%. No. Copies of Single Issue Published Nearest to Filing Date: Total Number of Copies (Net press run): 370,040; Paid/Requested Outside-County Mail Subscriptions Stated on Form 3541: 225,674; Sales Through Dealers and Carriers, Street Vendors, Counter Sales, and Other Non-USPS Paid Distribution: 103,459; Total paid and/or Requested Circulation: 329,133; Free Distribution by Mail: 3,721; Free Distribution Outside the Mail: 35,768; Total Free Distribution: 39,489; Total Distribution: 368,622; Copies not Distributed: 1,418; Total: 370,040; Percent Paid and/or Requested Circulation: 89.3%. This Statement of Ownership, Management, and Circulation was filed September 20, 2012 by Karin Kinnear, Director, Audience Development and Operations.

PUBLISHING

Vice President
Jeff Spicer jeff.spicer@oracle.com

Publisher
Jennifer Hamilton jennifer.hamilton@oracle.com +1.650.506.3794

Audience Development and Operations Director
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

Advertising Sales Assistant
Cindy Elhaj cindy@sprocketmedia.com +1.626.396.9440, x201

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

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.

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

Commitment to the Cloud

Cloud is the word at Oracle OpenWorld—and Oracle.

I didn't count the number of conversations I had or heard that included the word, but I'm quite comfortable saying that *cloud* was the word of the week at Oracle's conferences and events in San Francisco, California, September 29 through October 5, 2012. There were announcements focused explicitly on Oracle cloud technologies and solutions, but cloud was also a big part of every Oracle product, technology, and service announcement and discussion.

CLOUD CHECKLIST

Oracle CEO Larry Ellison's Sunday keynote at Oracle OpenWorld 2012 was about cloud computing, including the history of cloud computing going back to NetSuite in 1998 and the beginnings of Oracle Fusion Applications development in 2004. Ellison discussed public and private clouds, the levels of cloud services—software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS)—and their key components, and Oracle's latest offerings in each of these service areas.

Ellison talked about Oracle technologies designed specifically for better cloud computing, from infrastructure and private cloud services to database multitenancy and engineered systems. He also unveiled the latest Oracle engineered systems releases and discussed their infrastructure and platform support for cloud computing.

Oracle's Exadata Database Machine X3-2 and Exadata Database Machine X3-8 are the latest Oracle Exadata engineered systems, featuring increased flash memory—to 22 TB—and support for write I/Os in Exadata Smart Flash Cache. Exalogic Elastic Cloud X3-2 is the latest Oracle Exalogic engineered system; it features updated hardware and uses the Exalogic Elastic Cloud Software 2.0 release. (See "Built for Speed" on page 22 and "Optimized Application Performance" on page 23 for more information.)

Oracle also announced seven Oracle Cloud (cloud.oracle.com) preview services, including Oracle Planning and Budgeting Cloud Service, Oracle Financial Reporting Cloud Service, Oracle Data and Insight Cloud Service, Oracle Social Sites Cloud Service, Oracle Developer Cloud Service, Oracle Storage Cloud Service, and Oracle Messaging Cloud Service. (See "Oracle Expands Cloud Services Portfolio" on page 20 for more information.)

CLOUD CHALLENGES AND SOLUTIONS

While gathering content for this issue of *Oracle Magazine*, I ran across the following phrase in descriptions of cloud computing news at Oracle OpenWorld: "the varied and unpredictable workloads of cloud computing." It looked like a problem statement, but there was no solution with it—at least, not in the first text I saw.

But when I located the phrase in the Oracle press release where it originated, it turns out that the cloud problem statement had a very clear and specific solution: Oracle's new Oracle Exadata X3 engineered systems, the press release said, ". . . can store up to hundreds of terabytes of compressed user data in flash and RAM memory, virtually eliminating the performance overhead of reads and writes to slow disk drives, making Exadata X3 systems the ideal database platforms for the varied and unpredictable workloads of cloud computing."

Tom Haunert, Editor in Chief
tom.haunert@oracle.com

NEXT STEPS

LEARN more about
Oracle Cloud
cloud.oracle.com

Oracle Exadata
oracle.com/exadata
oracle.com/us/corporate/press/1855412

Oracle Exalogic
oracle.com/exalogic

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** 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.

ORACLE

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.

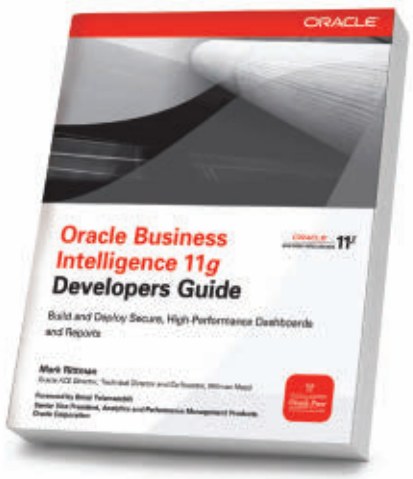

Oracle Business Intelligence 11*g* Developers Guide
Mark Rittman
Deliver meaningful business information to users anytime, anywhere, on any device, with Oracle Business Intelligence 11*g*.

PeopleSoft PeopleTools Data Management and Upgrade Handbook
Paula Dean, Jim J. Marion
Maximize data integrity, maintain peak application performance, and keep your PeopleSoft environment up to date.

Oracle E-Business Suite Financials Handbook, Third Edition
Ben Prusinski, Gustavo Gonzalez
Maintain an integrated, customer-focused financial computing framework that meets global business requirements while lowering total cost of ownership.

Oracle Solaris 11 System Administration: The Complete Reference
Michael Jang, Harry Foxwell
Manage a highly scalable, cloud-based computing platform and deliver unmatched performance levels at every layer of your IT stack.

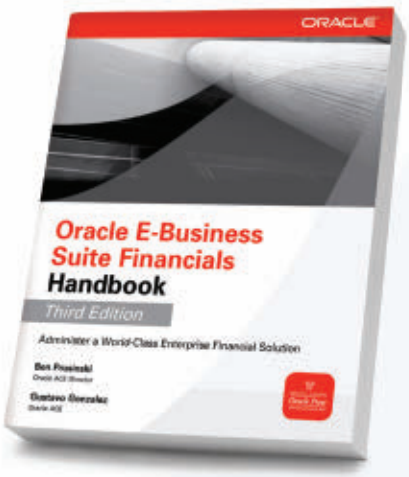

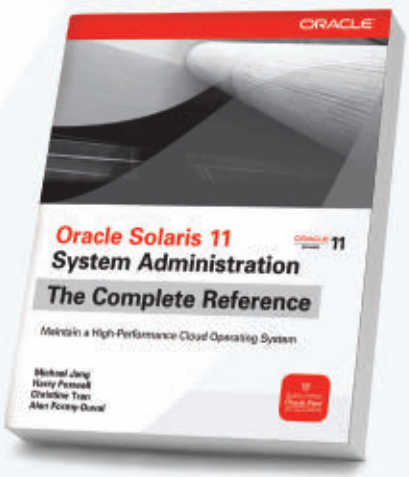

Available in print and eBook formats

Join the Oracle Press Community: www.OraclePressBooks.com

Follow us: @OraclePress

ORIGINAL • AUTHENTIC
Oracle Press
ONLY FROM McGRAW-HILL

NAVIGATION APPS

You're not limited to the navigation apps that come with your phone. Hit the road with these alternatives.

CoPilot
This app's main feature is its ability to chart a course without a data connection. Plus, tap anywhere on the map to navigate there. US$3.99 to US$43.99. **copilotlive.com**

MapQuest
This app puts points of interest on the main screen, provides turn-by-turn spoken directions, and updates traffic info every five minutes. Free. **mapquest.com**

MotionX-GPS Drive
This soup-to-nuts navigation app shows your speed limit, advises you which lane to take at complicated intersections, and connects to social media sites and your address book. US$1. **drive.motionx.com**

Waze
A "social mobile" app, Waze collects traffic, accident, speed trap, and other info in real time from its users and builds it into dynamic maps. The more users in your area, the better. Free. **waze.com**

LIVE STREAM MINICAM

What's better than recording video of your latest adventures on the trail, river, mountain—or in the conference room? Streaming it live via 3G, 4G, and Wi-Fi with the Looxcie HD wearable camera. Simply attach the small camera to a headband, hat, or helmet, wirelessly sync the camera with the Looxcie iOS or Android phone app, and you are ready to broadcast. Looxcie HD also lets you connect directly to Facebook through its companion apps. Applications include live blogging, teaching a workshop, collaborating with a team, and sharing your outdoor exploits. Looxcie HD can also do video capture and sharing with full 1,080p recording. *US$299.* **looxcie.com**

"The impact and almost magic ripple effect of a well-crafted and beautifully delivered three-minute elevator speech in today's changing and challenging market cannot be denied."

—Terri L. Sjodin, author of *Small Message, Big Impact: The Elevator Speech Effect* (Portfolio/Penguin, 2012)

CLOUD CONFUSION

About half of Americans surveyed are hazy on what cloud computing really is. In a survey of more than 1,000 adults, 51 percent of respondents said that stormy weather can interfere with cloud computing.

***Source:** Wakefield Research for Citrix, bit.ly/PJnrGP*

CLOUDS IN THE FORECAST

Of the more than 400 IT professionals responding to a recent survey, 51 percent have or are planning private clouds. Other key findings:

- **99%** of those with private clouds describe their initiatives as successful
- **45%** of nonadopters believe private clouds could deliver a compelling technical advantage
- **38%** of those starting private cloud projects expect to spend 21% or more of their IT budgets on the initiative
- **38%** of respondents will cobble together individual projects despite 31% citing a lack of standards as a roadblock
- **16%** of respondents starting private cloud projects will buy a preconfigured bundle

***Source:** Information Week Reports*, Private Cloud Vision vs. Reality, *bit.ly/Rmi7rL*

Data Transfer Dynamic Duo

If you work with large media files on a Mac, you probably know that external drives can create a bottleneck due to low rpm and low data transfer speeds. Western Digital My Book VelociRaptor Duo solves this problem with an elegantly designed external enclosure packed with two WD VelociRaptor 10,000 RPM drives and dual Thunderbolt ports that move data at transfer rates of up to 10 Gb/sec in both directions. This allows data to travel both ways without compromising bandwidth. VelociRaptor can be set up as RAID 1 (mirroring for data protection), or JBOD ("just a bunch of disks") to use as two individual drives. *US$899.99.* **bit.ly/PX1ILM**

I-HUA CHEN

Data Quality Tools for Oracle

ORACLE Gold Partner

Now, Oracle-based data stewards and information workers can effectively verify, enrich and cleanse contact information assets – from around the country to around the world. Native data quality components make implementation easy and ensure you always have a perfect fit.

- Validate, parse and standardize address, phone, email and name
- Verify, correct and format addresses from 240 countries
- Append rooftop lat/long coordinates for more than 40 countries
- Identify, merge and purge duplicate records with advanced matching
- Update addresses of U.S. & Canadian customers that have moved
- Correctly field and parse contact data from unstructured sources

Available solutions include E-Business Suite plugins, 32/64 bit PL/SQL and Java API toolkits for Linux, Solaris, Windows.

BEFORE

john smith iii phd
melissa data corp.
22382 Empresa 92688
7145895200
john@800miAL.con

AFTER

John Smith III PhD
Melissa Data Corp.
22382 Avenida Empresa Ste 100
Rancho Santa Margarita, CA 92688-2112
949-589-5200
John@melissadata.com
Delivery Indicator: Business

*Highlights indicate added and/or corrected data.

Realtime NCOA^Link Change-of-address Web Service available

Visit us at Oracle OpenWorld, booth #629
Request your free trial at MelissaData.com/myoracle or call 1-800-MELISSA

MELISSA DATA®
Your Partner in Data Quality

TOGETHER ORACLE AND NATIONAL GEOGRAPHIC EDUCATION ARE

Leading the Way in Ocean Education and Marine Research

NATIONAL GEOGRAPHIC education

ORACLE

National Geographic Education supports the mission of the National Geographic Society to inspire people to care about the planet by creating compelling educational materials for young people and the adults who teach them. NG Education provides unique learning experiences to educators and advocates for improved education in geography, the environmental sciences, and other disciplines that are critical to understanding our world.

With support from Oracle, National Geographic Education is engaged in a major project to develop teacher leaders in marine ecology and create materials about ocean science and geography for students, families, the ocean recreation community, and the general public.

Support our work today.
Visit nationalgeographic.org/education.

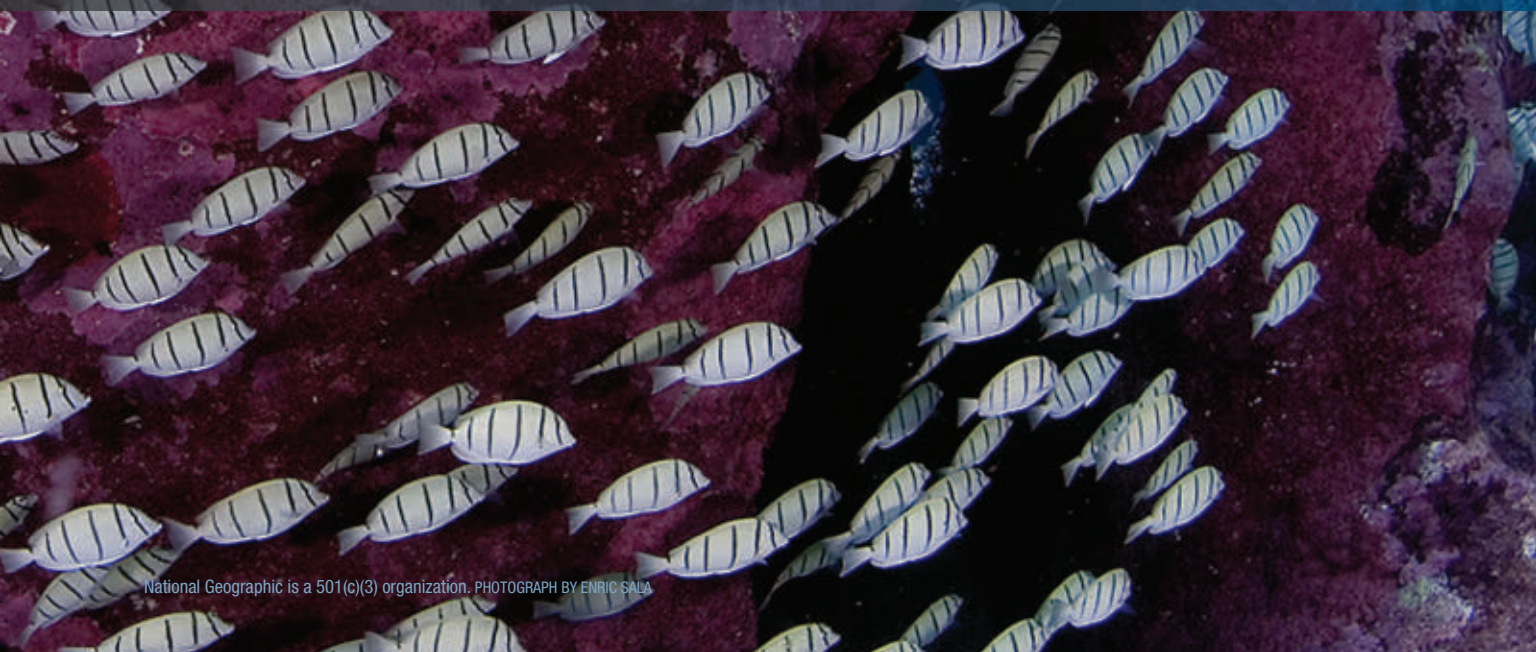

National Geographic is a 501(c)(3) organization. PHOTOGRAPH BY ENRIC SALA

Technology Events

Conferences and sessions to help you stay on the cutting edge

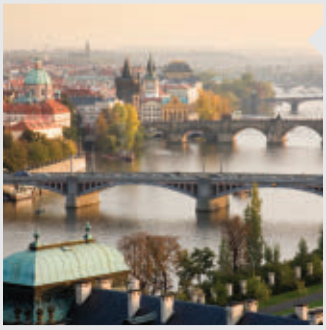

Oracle Day 2012

Through November 28 in various cities throughout Europe, the Middle East, and North Africa.

Learn to absorb disruptive technologies such as cloud, mobile, social, and big data without toppling your infrastructure. With insights from Oracle experts, live demonstrations of new products, customer success stories, and opportunities to brainstorm new concepts and tactics with your peers, Oracle Day provides contacts, ideas, and solutions that drive business value and agility from your Oracle investments. Find the complete list of cities and dates at **oracle.com/oracleday**.

Asia HEUG Conference 2012

November 1–2, Hong Kong, China

▶ **bit.ly/Rh5uP1**

At this new conference featuring Oracle solutions in higher education, participants see presentations by institutions and vendors, hear Oracle strategists discuss Oracle product direction, network with colleagues, and explore complimentary products and services from key vendors.

Sangam '12

November 2–3, Bangalore, India

▶ **aioug.org/sangam12.php**

Cohosted by All India Oracle Users Group, India Oracle Applications User Group, and the Chennai Java User Group, Sangam '12 hosts more than 70 speakers from around the world leading more than 120 sessions. Conference tracks include database, development, Java, Oracle applications, and Linux.

Insurance and Technology Executive Summit

November 4–7, Rancho Palos Verdes, California

▶ **insurancetech.com/summit2012**

Learn about strategies and technologies that are transforming the insurance industry at this summit for senior-level executives. Technology topics include customer experience, big data, core systems, and the role and future of IT.

Oracle Health Sciences User Group (OHSUG) Annual Conference

November 4–7, Monte Carlo, Monaco

▶ **bit.ly/NsHdLO**

This conference features presentations on the technical and functional aspects of Oracle Health Sciences applications for clinical, pharmaceutical, biotech, medical device, contract research, and support companies, along with focus groups covering Oracle Argus Safety, Oracle Clinical, Oracle Remote Data Capture, cloud, SOA, data mining, and metrics and reporting.

OAUG Connection Point—EPM/BI

November 5–6, Orlando, Florida

▶ **connectionpoint.oaug.org/2012/orlando**

Oracle Applications Users Group (OAUG), Florida OAUG, the OAUG Hyperion Special Interest Group (SIG), and Oracle Business Intelligence SIG present a two-day event at which Oracle experts and partners provide information, application roadmaps, and practical solutions in more than 35 educational sessions for accounting, finance, and internal audit teams.

Gartner Symposium/ITxpo

November 5–8, Barcelona, Spain

▶ **bit.ly/OMOiml**

This gathering of CIOs and senior IT executives offers more than 350 analyst sessions, workshops, roundtables, and keynotes over four days.

GETTY IMAGES

ORACLE USER GROUPS

Ottawa Canada Linux Users Group Meeting
November 1, Ottawa, Ontario, Canada
oclug.on.ca

UKOUG Application Express SIG Meeting
November 6, London, England
ukoug.org

IOUG Virtualization SIG Symposium
November 7–8, online
bit.ly/RKJX7S

Suncoast Oracle Users Group Meetings
November 8 and December 6, Tampa, Florida
soug.org/meetings.php

Michigan Oracle Users Summit
November 14, Livonia, Michigan
mous.us

OUG Scotland APEX SIG
November 14, Linlithgow, Scotland
ukoug.org

Dayton Dynamic Languages SIG Meetings
November 14 and December 12, Dayton, Ohio
dma.org/sigs.shtml#dynamic

Northern California Oracle Users Group Fall Conference
November 15, Pleasanton, California
nocoug.org

OAUG Connection Point—Oracle E-Business Suite Release 12.1/Oracle Fusion Coexistence
November 15–16, San Diego, California
bit.ly/ODmBy2

Calgary Oracle Users Group Meetings
November 15 and December 13, Calgary, Alberta, Canada
coug.ab.ca

OUG Ireland BI and EPM SIG Meeting
November 20, Dublin, Ireland
ukoug.org

Heartland Oracle Users Group Conference and Training Days
November 28–29, Omaha, Nebraska
heartlandoug.wordpress.com

New Zealand Oracle Users Group Meeting
December 5, Wellington, New Zealand
nzoug.org

NYOUG Special Winter General Meeting
December 12, New York, New York
nyoug.org

St. Louis Oracle Users Group Meeting
December 13, Chesterfield, Missouri
sloug.org

Business Intelligence and data warehousing are hot topics at the TDWI World Conference in Orlando, Florida, November 11–16.

With 10 role-based tracks and 6 industry tracks, the agenda targets specific job responsibilities and ways to adapt new ideas and strategies to individual industries.

QCon San Francisco 2012

November 5–9, San Francisco, California

▶ **qconsf.com**

This practitioner-driven conference for development team leads, architects, and project managers includes two tutorial days followed by three conference days covering 18 tracks, including big data and analytics, continuous delivery, and the Java renaissance.

Cyber Security 2012

November 6, London, England

▶ **oraclepublicsector.co.uk/events**

Featuring all aspects of cybersecurity within the public sector, this conference offers the latest strategies from Oracle and leading public sector agencies and opportunities to discuss security with professionals and Oracle experts.

Alliance Down Under

November 6–9, Gold Coast, Queensland, Australia

▶ **heug.org/p/cm/ld/fid=543**

Australian and New Zealand Higher Education User Group presents a program covering all aspects of Oracle's PeopleSoft and other Oracle technologies in higher education, including special sessions for executives and academic registrars.

TDWI World Conference

November 11–16, Orlando, Florida

▶ **bit.ly/Tkwmzh**

The conference features six days of courses covering business intelligence (BI) and data warehousing topics, including BI essentials, business analytics, data analysis and design, data asset management, agile development, cloud computing, text analytics, virtualization, open source, Web 2.0, social media, and mobile BI.

Devoxx

November 12–16, Antwerp, Belgium

▶ **devoxx.com/display/dv12/home**

This annual European Java gathering is historically the world's largest vendor-independent Java conference. It features two days of in-depth university sessions, 30-minute tools-in-action sessions focused on technical tools and APIs, three-hour hands-on labs, three days of hour-long technical presentations, 15-minute quickie sessions during lunch breaks, and informal evening birds-of-a-feather gatherings.

Bulgarian Oracle User Group Conference

November 16–18, Pravets, Bulgaria

▶ **bgoug.org/en/events/details/88.html**

Topics at this conference include database server technology, database administration and tuning, application design and development, middleware, business intelligence, and content management.

DOAG 2012 Conference + Exhibition

November 20–22, Nuremberg, Germany

▶ **bit.ly/NXpbfc**

Approximately 400 technical lectures and displays by 60 companies are highlights of this 24th-annual gathering of Deutsche Oracle Anwendergruppe—also known as DOAG, the German Oracle Users Group. New content this year includes developing mobile applications with Oracle Application Development Framework Mobile Client.

JayDay Munich

December 3, Munich, Germany

▶ **jayday.de**

This new event features JavaOne Rock Stars, Oracle ACEs, and Java Champions presenting a full day of content for Java developers. Sessions on Web, desktop, performance, concurrency, tools, languages, and more are conducted in English and German.

MDM and Data Governance Summit Singapore 2012

December 4–5, Singapore

▶ **bit.ly/T04LpJ**

This is Southeast Asia's premier event focused on master data management (MDM), data governance and quality, and information management. The summit focuses on building a business-driven MDM program and establishing an IT foundation to make MDM work. Content includes best practices and information about leading IT organizations.

Cloud Computing World Forum North America

December 6, New York, New York

▶ **bit.ly/TBSNQx**

This one-day executive-level conference on cloud computing and related topics features content on security, mobility, content delivery networks, applications, software as a service, virtualization, customer relationship management, social, and communications.

Groovy and Grails eXchange 2012

December 13–14, London, England

▶ **bit.ly/OtWWaK**

This two-day conference features two dozen expert-led talks, along with discussion and brainstorming sessions, all focused on learning and sharing ideas, tools, and practices for enterprise Web development with Groovy and Grails.

EVENTS LOCATOR

Oracle Events
oracle.com/events

User Groups
bit.ly/pX7Yob

GETTY IMAGES

Isenberg MBA. Your Way.

Full-Time. Online. Blended.
AACSB Accredited.
Adapted to your lifestyle.

ISENBERG
SCHOOL OF MANAGEMENT

UMassAmherst

Contact our dedicated advising center:
Tel. 413.545.5608 | www.isenberg.umass.edu/mba

Partner Excellence!

Congratulations 2012 Specialized Partners of the Year – Global

accenture	Capgemini CONSULTING.TECHNOLOGY.OUTSOURCING	Deloitte.
Middleware	Oracle on Oracle	Applications
MANNAI CORPORATION QSC	Maplesoft technology™	NEC
Industry	Server & Storage	Database

Specialized. Recognized. Preferred.

ORACLE®

partner.oracle.com
or call 1.800.323.SELL

Oracle
Excellence
Awards

What's New at Oracle

The latest videos, podcasts, blogs, and more

WEBCASTS

"Data Integration Breakthroughs: Bridging the Gap in Big Data, Cloud, and Real-Time Data Integration"

▶ **bit.ly/PH3uAK**

Get up to speed on big data, cloud integration, real-time analytics, and continuous availability.

"Oracle Hardware Systems: The Extreme Performance Online Event"

▶ **bit.ly/Qch7X8**

Learn how Oracle hardware systems deliver extreme performance for competitive advantage.

"Oracle Solaris 11 Online Forum"

▶ **bit.ly/OWXoO3**

John Fowler, executive vice president of systems at Oracle, discusses what's new with Oracle Solaris 11.

"Harness the Power of the New Release of Oracle GoldenGate 11*g*"

▶ **bit.ly/Sx0Gci**

Discover how Oracle GoldenGate 11*g* delivers enhanced, continuous availability and real-time data integration solutions in heterogeneous IT environments.

"Introducing Oracle Identity Management 11*g* Release 2"

▶ **bit.ly/QaaUQI**

See new features and technology in Oracle Identity Management 11*g* Release 2 that expand the reach of comprehensive identity protection to cloud and mobile environments.

2012 Global Database Management Survey Results: "Addressing Your Top Database Management Challenges"

▶ **bit.ly/SwUliF**

In this CIO.com Webcast sponsored by Oracle, get insights into top database management pain points, including the tasks that consume the most DBA time and top challenges in deploying new database services.

Webcast Series: *Accelerate Business-Critical Database Deployments with Oracle Optimized Solutions*

▶ **bit.ly/O9Pcso**

Learn to safely consolidate your business-critical database environment and build a business case based on quantifiable cost savings with Oracle Optimized Solutions.

Webcast Series: *Simplify and Speed Infrastructure Deployments*

▶ **bit.ly/TZUrOH**

Learn about Oracle's solutions with EMC and Cisco that are based on Oracle VM and Oracle Linux.

VIDEOS

Get Faster Results with Oracle SPARC SuperCluster

▶ **bit.ly/RjoNeI**

Get expert information about how Oracle SPARC SuperCluster reduces the time and effort required to install, tune, and administer optimized mission-critical systems.

Reduce Risk with Oracle Advanced Monitoring and Resolution

▶ **bit.ly/QDVXqf**

Learn the unique features of Oracle's 24/7 secure monitoring services and the value of proactive resolutions delivered by advanced support engineers.

Transplace Shares the Value of Oracle Platinum Services

▶ **bit.ly/Q217XK**

Watch the CTO of Transplace describe the unique benefits of Oracle Platinum services.

Overstock.com Achieves 100 Percent Uptime with Oracle Linux and Oracle VM

▶ **bit.ly/NWLr7w**

See how Overstock.com Web operations deliver faster IT response time, 100 percent uptime, and improved customer experience, while reducing hardware costs by 87 percent.

ORACLE UNIVERSITY

Get Your Oracle Training On Demand—More Courses Now Available

▶ **bit.ly/RKlvn2**

Get complete classroom training and access to hands-on labs from your desktop with the latest Training On Demand courses from Oracle University.

RESOURCE CENTERS

Oracle C-Central

▶ **bit.ly/P9apXg**

Get information about industry trends, best practices, and use cases for Oracle solutions through content developed specifically for CFOs and CIOs.

Quantifying the Value of Application-Driven Virtualization

▶ **bit.ly/SKzdUG**

Access an Evaluator Group report about Oracle virtualization, along with other resources such as demos, white papers, calculators, and downloads.

E-BOOKS

Simplify IT Through Hardware Technology Refresh

▶ **bit.ly/U032jO**

Learn how to modernize your infrastructure and optimize your data center with the latest generation of hardware systems.

Reliable, Scalable, Secure: Oracle Solaris 11

▶ **bit.ly/NjhpAL**

Discover how upgrading to Oracle Solaris 11 running on Oracle's new generation of SPARC and x86 systems is helping to transform data centers worldwide.

Oracle's Sun Servers: The Backbone of the Next-Generation Optimized Data Center

▶ **bit.ly/Rh2wOA**

Read about how to achieve new levels of data center optimization using Oracle's SPARC servers, Sun x86 servers, Sun Netra carrier-grade servers, and Sun blade servers.

OVERHEARD

"The SPARC SuperCluster was derived primarily because of such a large SPARC install base. . . . They've invested a lot of time in Oracle Solaris applications, mindset, and training. We felt that delivering a SPARC SuperCluster was a way to continue that SPARC ecosystem and Oracle Solaris ecosystem, yet delivering the performance that we had seen in Oracle Exadata and Oracle Exalogic."

—Gary Combs, Oracle Principal Product Manager, in *Get Faster Results with Oracle SPARC SuperCluster* (bit.ly/RjoNeI)

Oracle's SPARC Systems: What Customers Say
▶ **bit.ly/Qvagb4**
SPARC customers worldwide are minimizing costs and maximizing competitive edge. Discover how.

Easily Unlock the Value of Big Data
▶ **bit.ly/OtH3ig**
Learn how to acquire, organize, integrate, and analyze big data; benchmark your progress against leading companies; and access white papers, analyst videos, and demos.

Advanced Support: Maximize Availability and Performance
▶ **bit.ly/PHq4cB**
Discover how Oracle Advanced Customer Support Services delivers tailored, mission-critical support services.

PODCASTS

"Understanding Big Data Analysis with the R Language"
▶ **bit.ly/NrjTar**
Michael Kane from Yale University talks about big data and the R statistical programming language.

"Linux Security Explained"
▶ **bit.ly/PDm2EA**
Listen to James Morris, consulting member of the Oracle engineering technical staff, discuss the work he is doing to enhance Linux security.

"SQL Server Migration Tool"
▶ **bit.ly/PHx7lz**
Alfredo Kojima, senior software development manager at Oracle, talks about the new tool enabling easy migration from Microsoft SQL Server to MySQL.

"User Interface Changes in Oracle Identity Manager"
▶ **bit.ly/NmhMdL**
Marc Boroditsky, vice president of product management at Oracle, discusses user interface changes to Oracle Identity Manager that improve usability for business users and administrators.

"Oracle Hyperion Planning and Oracle Exalytics In-Memory Machine"
▶ **bit.ly/NPtola**
Oracle product experts discuss the benefits of running the new release of Oracle Hyperion Planning on Oracle Exalytics In-Memory Machine.

"Enhancing Enterprise Planning and Forecasting Through Predictive Modeling"
▶ **bit.ly/MRmbnC**
Hear about best practices in integrating predictive modeling into long-range planning, budgeting, and sales forecasting.

WHITE PAPERS

"SPARC Servers: An Effective Choice for Efficiency in the Data Center"
▶ **bit.ly/LDWfdD**
International Data Corporation analysts discuss how efficiency has become a top priority worldwide for CXOs, IT managers, and data center managers with limited IT budgets and rising operational costs.

"The Business Value of Proactive Support Services"
▶ **bit.ly/PHtAnc**
Read what International Data Corporation has to say about the value of Oracle's proactive support services.

CALCULATOR

Updated Oracle VM vs. VMware vSphere Cost Calculator
▶ **bit.ly/SxVHbs**
The newly updated Oracle VM calculator now includes management costs for VMware vCenter.

BLOGS

Implementing Oracle Analytic Solutions
▶ **blogs.oracle.com/robreynolds**
Read all about implementing successful analytics solutions.

Oracle Enterprise Performance Management Blog
▶ **blogs.oracle.com/epm**
Keep up with the latest key enterprise performance management market trends, recent events, and other enterprise performance management news.

DEMOS

Oracle Linux Animated Overview
▶ **bit.ly/R8JqXp**
Watch this three-minute animated overview and find out why Oracle Linux is the best Linux for enterprise computing.

WEB LOCATOR

Oracle Blog Center
oracle.com/blogs

Oracle Consulting
oracle.com/consulting

Oracle Events and Webcasts
oracle.com/events

Oracle Newsletters
oracle.com/newsletters

Oracle Podcast Center
oracle.com/podcasts

Oracle University
oracle.com/education

Oracle on Facebook
facebook.com/oracle

Oracle on Google+
bit.ly/LdGiMO

Oracle on LinkedIn
linkedin.com/company/oracle

Oracle on Twitter
twitter.com/oracle

Oracle Support
oracle.com/support

My Oracle Support
myoraclesupport.com

My Oracle Support Communities
communities.oracle.com

Only our IT physical infrastructure is as dynamic as your business needs.

"The ability to monitor everything from the Web allows us to catch small problems before they become a big concern."

Andy Koostra
Systems Manager
Consolidated Utility District

"This solution ensures a single integrated power and cooling management system."

Brian R. York – Director
Wireless National Practice
Insight

"APC by Schneider Electric InRow™ cooling architecture... is our design standard going forward."

Mike Mallett – Director of technical operations
Saint Luke's Health System

Protect any IT equipment or deployment with our adaptable APC by Schneider Electric InfraStruxure solutions.

A solution for every IT configuration

Protect your IT system uptime easily where ever your IT is deployed with APC™ by Schneider Electric™ InfraStruxure™ solutions. Our simple, adaptable, and manageable all-in-one physical infrastructure is designed as an easy-to-deploy system to allow for flexible management, physical IT deployment, and by extension, system uptime.

Customers have adapted the solution to all IT configurations — from out-of-the-way network closets to server rooms to data centers. Power protection, cooling, rack systems, and remote management are part of the total architecture for highest availability at all times. With our InfraStruxure solutions working to stave off physical threats, you can focus on more pressing concerns such as network threats, IT hardware failure, and switch hang-ups. When you deploy our solution, it's as if you're getting another IT person to ensure that your IT space or data center will still be at your command when you need it. What's more, our life cycle services enable optimal operations.

Business-wise, Future-driven.™

InfraStruxure

Integrated InfraStruxure solutions include everything for your IT physical infrastructure deployment: backup power and power distribution, cooling, enclosures, and management software. Adaptable solutions scale from the smallest IT spaces up to multi-megawatt data centers.

Mission-critical IT infrastructure without the complexity

- **Simple**
  Solution guides and out-of-the-box installation make deployment easy.
- **Manageable**
  Remote monitoring, management, and reporting simplify IT operations; energy management cuts costs.
- **Adaptable**
  With standardized designs for all types of applications, our solutions can be adapted to fit any IT need at any time for business-minded flexibility.

Make the most of your IT space!

Download our Top 3 solution design guides today and enter to win an iPad® 2.

Visit: www.apc.com/promo Key Code: u263v Call: 888-289-APCC x6459

APC™ by Schneider Electric

TUTORIALS

Using Oracle XML DB to Optimize Performance and Manage Structured XML Data

This tutorial shows you how to store, index, and query XML data in Oracle XML Database (Oracle XML DB). You will learn about Oracle XML DB's capabilities that enable storage, indexing, querying, optimized performance, and management of structured XML data.

▶ bit.ly/SnQW4J

Authenticating and Authorizing Users in Java EE 6 Web Applications Running on WebLogic Server 12*c*

This tutorial covers how to secure a Java EE 6 Web application with form-based authentication and role-based authorization. You will learn how to create a login form with JavaServer Faces as the client to collect user and password data and how to authorize certain roles that access a Servlet. The tutorial also shows you how to add users to Oracle WebLogic Server 12*c* and map them to the Web application roles that will be used in the authentication and authorization processes.

▶ bit.ly/SEBOBr

Oracle Tuxedo 12*c* Release 1: Configuring Extended Authentication and Authorization

Learn how to configure Oracle Tuxedo's extended authentication and authorization feature. This feature allows Oracle Tuxedo applications to use the same centralized LDAP server for authentication that the rest of the enterprise uses, while simultaneously allowing Oracle Tuxedo to provide authorization for protected resources without the need to maintain duplicate identity data.

▶ bit.ly/Q4xo35

ARTICLES

"How I Got Started with the Btrfs File System for Oracle Linux"

Read about the basic capabilities of the Btrfs file system in Oracle Linux, how to create a file system, and more.

▶ bit.ly/LgrQCT

"How I Simplified Oracle Database 11*g* Installation on Oracle Linux 6"

Learn how the oracle-rdbms-server-11gR2-preinstall RPM package can simplify Oracle Database 11*g* installation on Oracle Linux.

▶ bit.ly/UwUpLT

Oracle GoldenGate 11*g* Release 2 Unveiled

Oracle GoldenGate 11*g* Release 2, the latest version of Oracle's comprehensive real-time data integration solution, is now available. Tightly integrated with Oracle technologies, the new release delivers performance enhancements, increased support for business-critical and heterogeneous systems, and expanded management capabilities.

A component of Oracle Fusion Middleware, Oracle GoldenGate 11*g* Release 2 provides real-time data integration to business-critical applications, enabling improved business insight, query offloading to maximize online transaction processing performance, zero-downtime data migration, disaster recovery, and active-active database synchronization for continuous availability.

"To succeed in today's highly competitive global markets, organizations need to have a holistic view of their data landscape," says Brad Adelberg, vice president of development at Oracle. "Oracle GoldenGate 11*g* Release 2 enhances Oracle's industry-leading data integration solution through improved performance, reliability, flexibility, and integration. This release minimizes downtime for our customers' critical systems and supports better decision-making across the enterprise."

▶ bit.ly/Tf2x2o

Oracle Introduces Mobile Point-of-Service for Retailers

Oracle has unveiled Oracle Retail Mobile Point-of-Service, a mobile extension to Oracle Retail Point-of-Service that enables store associates to assist customers and securely complete transactions from anywhere in the store using mobile devices. Fully integrated with Oracle Retail's stores solutions, which ensures consistent business logic, Oracle Retail Mobile Point-of-Service has a scalable and flexible architecture that allows any combination of fixed and mobile checkout points and supports multiple mobile operating systems and devices, so retailers are not locked in to specific vendors.

"Oracle Retail Mobile Point-of-Service provides the flexibility to serve customers on their terms. By providing store associates with a consistent transaction interface and delivering actionable insights at any point of service, we help retailers level the playing field between store associates and digitally enabled customers," says Mike Webster, senior vice president and general manager, Oracle Retail. "Oracle Retail Mobile Point-of-Service provides an integrated and scalable architecture for retailers to create differentiated customer interactions."

▶ bit.ly/QSU9sN

Oracle Application Development Framework Essentials Unveiled

Oracle has released Oracle Application Development Framework (Oracle ADF) Essentials to bring Oracle ADF benefits to the global developer community. Free to use, Oracle ADF Essentials is standards-based and deploys on GlassFish Server Open Source Edition, providing a way for developers to adopt and extend Oracle ADF functionality to new environments.

Oracle ADF Essentials includes Oracle ADF Faces, a set of more than 150 advanced Web user interface (UI) components; Oracle ADF Controller, an extension of the JavaServer Faces controller layer that helps build complete reusable process flows; Oracle ADF Binding, which connects UIs to various business services; and Oracle ADF Business Components, declaratively configured, reusable components that implement common design patterns.

"Oracle ADF is already proven as the strategic framework used to build Oracle Fusion Applications," says Chris Tonas, vice president, Applications Development Tools at Oracle. "With Oracle ADF Essentials, we've combined key tooling from Oracle ADF, Oracle JDeveloper, and Oracle Enterprise Pack for Eclipse in a free offering that should be a great way for many of our customers and partners to jump-start their development efforts."

▶ bit.ly/SHQlew

LINDY GROENING

Oracle Extends Commitment to R for Big Data Analytics

Oracle is offering enhanced support for the R statistical programming language, increasing the accessibility of big data analytics within the enterprise. Oracle's open source distribution of R is available with Oracle Big Data Appliance and can also be downloaded now.

The enhanced support includes new platform ports of R for Oracle Solaris and AIX in addition to Linux and Windows, connectivity to Oracle TimesTen In-Memory Database in addition to Oracle Database, and integration of hardware-specific math libraries for faster performance.

"Big data analytics is a top priority for our customers, and the R statistical programming language is a key tool for performing these analytics," says Andrew Mendelsohn, senior vice president, Oracle Database Server Technologies. "With this announcement, we continue to enhance and expand our industry-leading support for R across the Oracle Database, Oracle TimesTen In-Memory Database, and Hadoop platforms."

oss.oracle.com/ORD

Oracle Launches Two Java Embedded Products

Addressing the strong demand for Java in the embedded market, Oracle has unveiled Oracle Java ME Embedded 3.2, a complete client Java runtime optimized for microcomputers and other resource-constrained devices, and Oracle Java Embedded Suite 7.0, the first Java-based middleware stack for embedded systems.

With Oracle Java ME Embedded 3.2, applications for small embedded devices are no longer tied to a single hardware platform. In addition, Oracle Java ME Embedded 3.2 meets the unique requirements of small embedded, low-power devices, including on-the-fly application downloads and updates, remote operation, and the ability to add new capabilities without affecting existing functions.

Oracle Java Embedded Suite 7.0 helps organizations achieve faster time to market by leveraging proven middleware components with their existing Java skill sets. Oracle Java Embedded Suite 7.0 is now available for Linux on x86 and Linux on ARM. An evaluation implementation is available for download from Oracle Technology Network.

"The rapid growth in the 'internet of things' is driving demand for open and cross-industry platforms that can help decrease time to market and deliver increased capabilities in embedded devices, while retaining tight control on development, production, and support costs," says Nandini Ramani, vice president of engineering, Java Client and Mobile Platforms at Oracle. "Java's uniquely flexible architecture supports these requirements through a highly secure virtual machine designed to support remote applications updates and downloads.

bit.ly/PPx1HD

bit.ly/UzFQwT

Oracle Unveils Oracle Fusion Tap for the iPad

Oracle has introduced Oracle Fusion Tap, a native iPad application that runs off cloud-based enterprise applications and across Oracle Cloud Application Services to improve productivity for mobile users by providing access to key functionalities to keep business moving. Automatically personalized to each user, it enables secure access to day-to-day management of the user's workforce and sales force automation.

"The mobile workforce is a business reality, and Oracle Fusion Tap is an example of how Oracle delivers mobile and cloud innovations that fundamentally improve productivity and how we work," says Chris Leone, senior vice president, Applications Development at Oracle. "With Oracle Fusion Tap, users will have an all-in-one, easily extensible app that puts mission-critical data and colleague connection at their fingertips."

LINDY GROENING

bit.ly/SMAeMP

Oracle Announces Oracle Solaris 11.1

The latest Oracle Solaris release, Oracle Solaris 11.1, delivers more than 300 new performance and feature enhancements.

Oracle Solaris 11.1 features database technology enhancements that deliver the best performance, availability, and I/O throughput of any UNIX platform used to run Oracle Database. These include improving lock latency for Oracle Real Application Clusters, resizing Oracle Database system global area (without a reboot), and observing and understanding database I/O bottlenecks using Oracle Solaris DTrace.

New cloud infrastructure features add to Oracle Solaris 11's highly efficient built-in virtualization capabilities across system, network, and storage resources and include support for the open standard Federated File System and expanded support for software-defined networks.

"Oracle Solaris 11 is the best UNIX operating system to run Oracle applications, deploy mission-critical cloud infrastructure, and protect customer investments," says John Fowler, executive vice president, Systems at Oracle. "The Oracle Solaris engineering team has worked closely with the Oracle Database engineering team to deliver unique value that customers can only get from an Oracle solution."

bit.ly/OOy0gx

Oracle Buys SelectMinds

Oracle has agreed to acquire SelectMinds, a cloud-based talent management company. SelectMinds applications help hiring managers use social connections to manage and market job referrals.

"Recruiting candidates through employee referrals is widely acknowledged as the most effective method to find talent through trusted contacts," says Thomas Kurian, executive vice president of product development at Oracle. "By adding SelectMinds to Oracle's Talent Management Cloud, Oracle can help customers with a complete talent management solution, enabling streamlined recruiting practices, more quality referrals, faster employee on-boarding, and better performance."

oracle.com/selectminds

Oracle Expands Cloud Services Portfolio

Oracle has unveiled seven Oracle Cloud preview services that augment its comprehensive cloud portfolio. These include Oracle Planning and Budgeting Cloud Service, which streamlines financial planning, budgeting, and forecasting processes; Oracle Financial Reporting Cloud Service, which enables the creation and delivery of highly formatted, boardroom-quality management reporting; Oracle Data and Insight Cloud Service, which aggregates insightful and intelligent data from enterprise, social, and external sources; Oracle Social Sites Cloud Service, which enables brands to quickly expose sites to their customers; Oracle Developer Cloud Service, which simplifies collaborative software development by providing a standards-based environment that supports the complete development lifecycle; Oracle Storage Cloud Service, which allows businesses to store and manage digital content in the cloud; and Oracle Messaging Cloud Service, which enables data communications between applications inside and outside of the Oracle Cloud.

"Cloud is a strategic business at Oracle," says Abhay Parasnis, senior vice president of product development at Oracle. "The breadth of opportunities available with the Oracle Cloud already exceeds anything else available in the industry, and with the introduction of these new services, Oracle continues to innovate, shape the market, and define the future of cloud-based services."

cloud.oracle.com

Enhanced Analytics for Oracle Exalytics In-Memory Machine

Oracle has released new enhanced analytics software optimized for Oracle Exalytics In-Memory Machine that enables customers to view and analyze data at the speed of business. New capabilities enhance organizational ability to quickly access and leverage data assets in real time, from information discovery to visual and mobile analytics, through enterprise planning and reporting.

In addition, Oracle Endeca Information Discovery is now certified with Oracle Exalytics, allowing business users to easily explore diverse and unstructured information to find answers to questions that cannot be easily answered by more-traditional, quantitative business intelligence.

"Oracle Exalytics enables organizations to make decisions faster in the context of rapidly shifting business conditions while broadening user adoption of business intelligence through the introduction of interactive visualization capabilities that make every user an analyst," says Balaji Yelamanchili, senior vice president, Oracle Analytics and Performance Management Products. "The latest software enhancements and certifications further increase the extreme performance and value of Oracle Exalytics—allowing users of all skill sets to gain the actionable intelligence they need to make better, more informed decisions."

bit.ly/R7vEIO

New Commercial Extensions for MySQL Enterprise Edition

Oracle has unveiled new commercial extensions for MySQL Enterprise Edition. Available at no additional cost as part of MySQL Enterprise Edition and MySQL Cluster Carrier Grade Edition, the new commercial extensions deliver advanced-security and high-availability options.

One of the new extensions, MySQL Enterprise Audit, provides an out-of-the-box, easy-to-use auditing and compliance solution that helps organizations conform to industry best practices and satisfy regulatory requirements. Other extensions provide capabilities to dynamically enable and disable audit streams, implement policies that log all or selected login or query-based activities, and automatically rotate audit log files based on size.

"MySQL customers want to reduce the risk, cost, and time to deploy and manage business-critical applications," says Tomas Ulin, vice president, MySQL Engineering at Oracle. "Our new security and high-availability commercial extensions specifically address those needs."

mysql.com/products/enterprise

Oracle Unveils Platform for Social

Oracle has launched Oracle's social relationship management solution, an integrated enterprise service that enables companies to listen, engage, create, market, and analyze interactions across multiple social platforms in real time. The solution is integrated with Oracle's enterprise applications, including Oracle Fusion Marketing, Oracle Fusion Sales Catalog, and Oracle ATG Web Commerce, allowing organizations to use social capabilities to transform their corporate business processes and systems. It is also integrated with Oracle Platform Services, including Oracle Java Cloud Service and Oracle Database Cloud Service, enabling marketing teams to integrate social technology with their custom Web pages, landing pages, and marketing tools.

Oracle's social relationship management solution includes Oracle Social Network Cloud Service, which provides a secure collaboration platform; Oracle Social Marketing Cloud Service, which enables marketers to centrally create, publish, moderate, manage, measure, and report across multiple social campaigns and platforms; Oracle Social Engagement and Monitoring Cloud Service, which enables organizations to analyze social media interactions while also empowering customer service and sales teams to effectively engage with customers and prospects; Oracle Social Sites Cloud Service, which provides a rich editing experience that end users can leverage to dynamically develop and launch social sites; and Oracle Data and Insight Cloud Service, which provides information and insights about common business entities.

"By fundamentally changing the way organizations connect with their different stakeholders, social is changing the rules of business," says Thomas Kurian, executive vice president of product development at Oracle. "With Oracle's social relationship management solution, we are empowering our customers to embrace this change by integrating the tools required to listen, engage, create, market, and analyze social interactions into existing applications and services."

cloud.oracle.com

LINDY GROENING

Oracle Optimized Solution for Oracle E-Business Suite Unveiled

Oracle has launched a new hardware offering, Oracle Optimized Solution for Oracle E-Business Suite. Based on Oracle's SPARC SuperCluster T4-4, Oracle Solaris, and Oracle VM Server for SPARC, Oracle Optimized Solution for Oracle E-Business Suite can deliver significant cost savings versus both new and existing hardware infrastructure.

"Oracle Optimized Solution for Oracle E-Business Suite delivers high-volume batch and online transaction processing performance with mission-critical availability and disaster recovery for Oracle E-Business Suite environments, while reducing risk and enabling significant savings in TCO," says Ganesh Ramamurthy, vice president of engineering at Oracle. "Built on the SPARC SuperCluster, our newest solution takes advantage of Oracle hardware and software engineered together to deliver the high performance and availability needed for business-critical applications."

bit.ly/SYt18L

Oracle VM VirtualBox 4.2 Now Available

Oracle VM VirtualBox 4.2, the new release of Oracle's open source virtualization software, delivers a smarter user interface; advanced networking capabilities; and Windows 8, Mac OS X 10.8, and Oracle Linux 6.3 support.

"As the only free, open source virtualization software that supports Windows, Mac, Linux, and Oracle Solaris platforms, users can install Oracle VM VirtualBox 4.2 on their preferred host platforms and run a huge variety of guest operating systems in virtual machines," says Wim Coekaerts, senior vice president, Linux and virtual engineering at Oracle.

bit.ly/R7mZGd

Oracle Retail Data Model Upgrades Analytics

Oracle has updated Oracle Retail Data Model, its standards-based, enterprise-class data warehouse for retailers. The latest enhancements deliver advanced analytics for unique retail formats including consumer goods, wholesale and private label, multi-channel, franchise control, and quick-serve restaurants and fine dining.

It has also been optimized for Oracle Database, Oracle Exadata Database Machine, and Oracle Big Data Appliance to deliver extreme scalability and performance.

"The enhancements to Oracle Retail Data Model provide deeper and broader insight for a wider range of retailers," says Mike Webster, senior vice president and general manager, Oracle Retail. "By taking advantage of its new capabilities, retailers will have unparalleled access to information that is relevant, timely, and actionable."

bit.ly/PPzH8f

Built for Speed

Oracle unveils new releases of Oracle Exadata Database Machine and Oracle Exalogic Elastic Cloud.

Pausing for waves of applause in the packed keynote hall on the first evening of Oracle OpenWorld, Oracle CEO Larry Ellison announced a new generation of engineered systems anchored by Oracle Exadata X3.

"If you thought the old Exadata systems were fast, you ain't seen nothing yet," he said.

One of the major performance drivers in the Exadata Database Machine X3-2 and Exadata Database Machine X3-8 is their mass memory system, which permits these engineered systems to store hundreds of terabytes of compressed user data in flash and RAM memory, virtually eliminating the performance overhead of reads and writes to disk drives. The Exadata flash components are part of a memory hierarchy that includes 4 TB of dynamic random-access memory (DRAM) for Exadata Database Machine X3-8, 22 TB of flash cache, and 500 TB of disk storage.

Intelligent storage optimization in Oracle Exadata Smart Flash Cache automatically keeps all active data in flash memory and RAM, while keeping less-active data on low-cost disks. Exadata Smart Flash Cache now also sends write I/Os to flash in addition to read I/Os, accelerating write performance by 20 times.

"As we move more and more of our databases onto Exadata machines, those databases aren't going to go onto disk drives," he predicted. "They're going to go into flash memory where the information is available instantaneously, where you can ask a question and get an answer at the speed of thought, where the new generation of clouds delivers answers faster than you ever imagined was possible because of Oracle Exadata X3, the world's fastest computer for business."

Oracle CEO Larry Ellison

UPGRADE ADVANTAGE

Organizations using Oracle Exadata X2 can also take advantage of the new write I/O support simply by upgrading their Exadata Storage Server Software. "When you push the button and upgrade your Exadata software, you will suddenly speed up write I/Os by someplace between a factor of 10 and a factor of 20," Ellison estimated.

With the Oracle Exadata X3, speed is easy. "Everything's faster with Exadata X3, and they consume less power," Ellison said. "We now can do a million writes per second in a single rack of Exadata. You would need 10,000 disk drives for that kind of write capacity, or 100 racks compared to a single Exadata rack."

In addition to previously available Oracle Exadata full-rack, half-rack, and quarter-rack configurations, the new Oracle Exadata X3 eighth-rack configuration provides an entry point for smaller workloads, including testing, development, and disaster recovery systems.

FASTER PERFORMANCE, SAME COST

Also announced at Oracle OpenWorld, Exalogic Elastic Cloud X3-2, the second hardware generation of Oracle's engineered system for running business applications, offers similar boosts in functionality and performance, thanks in part to the eight-core Intel Xeon E5-2600 series of processors, which are resident in the new Oracle Exadata and Oracle Exalogic systems. The new Exalogic platform builds on the recent release of Oracle's Exalogic Elastic Cloud Software 2.0 (see "Optimized Application Performance," page 23) to provide extreme performance, reliability, and scalability for Java, Oracle Fusion Middleware, Oracle Applications, rehosted CICS/IMS TP, and other business applications.

Oracle Exadata and Oracle Exalogic are key components of Oracle's new infrastructure-as-a-service cloud offering, which Ellison announced during the same keynote presentation.

"What we're offering is OS, VM [virtual machine], compute services, and storage services on the fastest, most reliable, most secure machines in the world—our engineered systems, Exadata, Exalogic, SuperClusters, and Exalytics—all networked together with a modern InfiniBand network," Ellison concluded. "The reason we're making these systems faster is to improve not just their peak performance, but to improve their cost performance. You can save a lot of money by using these engineered systems, Oracle Exadata and Oracle Exalogic." ◀

David Baum is a freelance technology writer based in Santa Barbara, California.

NEXT STEPS

LEARN more about
Oracle Exadata
oracle.com/exadata

Oracle Exalogic
oracle.com/exalogic

HARTMANN STUDIOS

Optimized Application Performance

Oracle's Exalogic Elastic Cloud Software runs applications in isolation, provides traffic metering, and delivers self-service–based infrastructure as a service.

Following the recent release of Oracle's Exalogic Elastic Cloud Software 2.0, Tom Haunert, editor in chief of Oracle Magazine, *sat down with Mohamad Afshar, vice president of product management at Oracle, to talk about Oracle Exalogic Elastic Cloud, the new software release, and what's next for Oracle Exalogic. The following is an excerpt from that interview. Download the full podcast at oracle.com/magcasts.*

Oracle Magazine: What is the history of Oracle Exalogic, what does it do, and how does it integrate with other Oracle engineered systems?

Afshar: Larry Ellison launched Oracle Exalogic Elastic Cloud at Oracle OpenWorld in 2010. The vision for the product was to deliver a complete infrastructure to data centers for running their middleware and packaged applications. Oracle's engineered system for the database tier, Oracle Exadata, was already available, and the idea behind Exalogic was to deliver a machine that would work in tandem with Oracle Exadata to run applications and midtier workloads. Effectively, the combination of the two products becomes the foundation for the modern data center on which applications are consolidated.

Oracle Magazine: Other Oracle engineered systems are optimized to run specific installed Oracle software, but Oracle Exalogic Elastic Cloud is optimized for the installed Oracle software and a variety of Oracle and other applications. How is this possible?

Afshar: When we started looking at building Oracle Exalogic, we looked at where the bottlenecks are in the midtier and the application tier, and we noticed that the primary bottleneck that hinders performance and throughput tends to be in the network. So we built Exabus technology, which encompasses

> **"We enable customers to have a true IaaS environment."**

technologies spanning InfiniBand switches, gateways, host channel adapters, firmware, device drivers, operating system extensions, and software libraries. Exabus provides an interoperability layer that enables any application that runs on Oracle Linux on x86 to run on top of it with no code changes and with significantly improved performance and throughput. For Oracle middleware and applications we also built native integration with Exabus to enable breakthrough performance and throughput—and that helps deliver the 5 to 10 times improvements that you see us talk about for Oracle WebLogic Server, Oracle Tuxedo, Oracle Coherence, and Oracle SOA Suite, as well as our applications such as Oracle E-Business Suite, Oracle ATG applications, and more.

BOB ADLER

Mohamad Afshar, Vice President of Product Management, Oracle

Oracle Magazine: How does Oracle Exalogic provide optimized support for Java?

Afshar: Oracle has a significant investment in Java and the Java community and is committed to ensure ever-greater adoption of Java. Hence we looked at how we could improve response times and throughput for Java applications by leveraging the Exabus technology within the Java Virtual Machine [JVM], Oracle WebLogic Server, and Oracle Coherence. This engineering effort involved delivering capabilities such as JDBC over SDP [Sockets Direct Protocol] and RDMA [Remote Direct Memory Access]–backed state replication.

Oracle Magazine: Oracle recently announced Exalogic Elastic Cloud Software 2.0. What is the background behind this release, and what are the key features?

Afshar: Many of our customers have been running multiple Oracle applications on Oracle Exalogic—Oracle E-Business Suite, Siebel, Oracle ATG, and PeopleSoft, as well as Oracle Fusion Middleware—and they've typically been isolating applications by running one application per compute node. For example, a customer running an Exalogic quarter rack system could dedicate two Exalogic compute nodes for Oracle ATG, two compute nodes for Oracle E-Business Suite, and then two to four compute nodes for integration and Oracle WebCenter components.

With Exalogic Elastic Cloud Software 2.0, we've delivered a number of capabilities that enable customers to run more applications

on Exalogic by delivering application isolation through virtualization. We built Exabus integration into the virtualization layer in Exalogic Elastic Cloud Software 2.0 to enable near native (nonvirtualized) application performance, and incorporated cloud management capabilities around that virtualized environment based on an infrastructure-as-a-service (IaaS) model. As part of this we also delivered a self-service user interface that provides full IaaS self-service capabilities for cloud users and cloud administrators. With Exalogic Elastic Cloud Software 2.0, we enable customers to have a true IaaS environment where they can deploy a self-service model or provision for different users within groups, and enable those users to go in and spin up environments and start running applications.

Oracle Traffic Director is another important part of Exalogic Elastic Cloud Software 2.0. It is an integrated application delivery controller capable of doing everything from basic load balancing to complex traffic shaping, traffic metering, and security enforcement.

Oracle Magazine: How does Exalogic Elastic Cloud Software 2.0 fit into Oracle's cloud applications and Java strategies?

Afshar: If you look at Oracle Cloud, Oracle has announced application services such as Oracle Sales and Marketing Cloud Service, Oracle Human Capital Management Cloud Service, and Oracle Social Network Cloud Service, in addition to Oracle Database Cloud Service and Oracle Java Cloud Service. Specifically in the context of Java, Oracle Java Cloud Service runs on Oracle Exalogic, and so we can deliver 100 percent portability between customers' systems and Oracle Cloud. As far as Oracle's Java strategy is concerned, you can expect to see continued investment in optimizing Java applications on Oracle Exalogic across the JVM, the virtualization layer, networking, and storage.

Oracle Magazine: What's next for Oracle Exalogic?

Afshar: We are committed to taking the latest industry-standard hardware and fully testing it within the engineered system environments, certifying it, and delivering it to customers. And that's exactly what we are doing with the introduction of the Exalogic X3 product series, which brings Intel's Sandy Bridge chip set to the Exalogic platform, delivering ever-better density, performance, and compute power. ◀

NEXT STEPS

LEARN more about
Oracle Exalogic Elastic Cloud
oracle.com/exalogic

Oracle's Exalogic Elastic Cloud Software 2.0
oracle.com/us/corporate/press/1715394

LISTEN to the podcast
oracle.com/magcasts

#1 Middleware

- ✓ #1 in Application Servers
- ✓ #1 in Service-Oriented Architecture
- ✓ #1 in Identity & Access Management
- ✓ #1 in Enterprise Performance Management

Oracle Middleware

Trusted by 100,000 Customers Worldwide

ORACLE®

oracle.com/goto/middleware
or call 1.800.ORACLE.1

Oracle 11g Anti-Hacker's Cookbook

By Adrian Neagu

Packt
packtpub.com

This book covers the most-important security measures as well as tips and tricks that can be used for better Oracle Database 11*g* security. The author uses real-world scenarios to show you how to secure the Oracle Database server from different perspectives and against different attack scenarios. Almost every chapter has a possible threats section, which describes the major dangers that can be confronted. The initial chapters cover how to defend your operating system, network, data, and users. The defense scenarios are linked and designed to prevent these attacks. The later chapters cover Oracle Database Vault, Oracle Virtual Private Database, Oracle Label Security, and Oracle Audit Vault. The author also demonstrates how to perform a security assessment against the operating system and the database.

Oracle Business Intelligence 11g Developers Guide

By Mark Rittman

Oracle Press
oraclepressbooks.com

This Oracle Press guide explains how to build business intelligence (BI) applications using Oracle Business Intelligence 11*g*. It covers how to execute effective queries, build accurate models, use scorecards and key performance indicators, create dynamic reports, set up dashboards, and publish to smartphones and wireless devices. The author also describes how to integrate BI applications with Oracle Fusion Middleware security, Oracle WebLogic Server, and Web services. Aimed at developers creating BI solutions using Oracle's platform, this book also covers Oracle Exalytics In-Memory Machine, a hardware appliance that combines Oracle Business Intelligence and Oracle TimesTen In-Memory Database, for "speed of thought" analytics.

Oracle APEX Best Practices

By Alex Nuijten, Iloon Ellen-Wolff, and Learco Brizzi

Packt
packtpub.com

Oracle APEX Best Practices guides you through the development of real-world applications and gives you a broader view of Oracle Application Express to help you maximize its capabilities. The book covers setting up an Oracle Application Express environment, testing and debugging, security, and getting the most out of SQL and PL/SQL. It is filled with best practices on how to best use Oracle Application Express. Both new and experienced application developers will benefit from *Oracle APEX Best Practices*.

Effective MySQL: Replication Techniques in Depth

By Ronald Bradford and Chris Schneider

Oracle Press
oraclepressbooks.com

Effective MySQL: Replication Techniques in Depth describes what is needed to understand and implement MySQL replication to build scalable solutions. This book includes detailed syntax examples to demonstrate the features, options, and limitations of native MySQL replication. It provides an evaluation of new replication features and additional third-party product implementations, and helps you to ensure that your MySQL environment can support the various high-availability needs of your business. This guide helps you to master the strengths and limitations of native asynchronous replication in a MySQL topology and identify the important features to improve replication for growing business requirements. You will also learn how to recognize the key business factors to determine your optimal high-availability needs and gain an understanding of the benefits of using MySQL replication for failover scenarios.

Oracle SOA Suite 11g Administrator's Handbook

By Ahmed Aboulnaga and Arun Pareek

Packt
packtpub.com

This book provides detailed explanations of the core administrative and management activities associated with Oracle SOA Suite. It includes information for end-to-end administration of Oracle SOA Suite 11*g* and also delves into advanced topics such as silent installs, cloning, backup and recovery, and high-availability installations. The book begins with managing composite applications, deployments, and lifecycles, and then moves on to monitoring instances, service engines, Oracle WebLogic Server, and composite applications. Toward the end, the authors take you through identifying faults and exceptions, troubleshooting approaches, and securing various components. The book explains core concepts while providing real-world implementation specifics.

Look for other Oracle books at bit.ly/NjG3KM.

Partners Achieve Oracle Database Appliance Specialization

Two members of Oracle PartnerNetwork (OPN) have achieved OPN Specialized status for Oracle Database Appliance.

Cintra Software & Services, a Global Platinum Partner and the OPN Specialized Global Partner of the Year for Oracle Database in 2011, has created a blueprinted delivery model for Oracle Database Appliance to fully leverage the appliance's strengths. Cintra has completed hundreds of successful Oracle Database implementations; its database infrastructure services include design, implementation, and ongoing remote support of Oracle Database architectures. This marks Cintra's ninth specialization.

Mythics, an Oracle Platinum Partner, Oracle Product GSA Schedule Holder, and OPN Specialized Consulting Partner, has demonstrated success in selling, deploying, and managing Oracle Database Appliance in both public sector and commercial environments. An OPN member since 2000, Mythics provides Oracle licensing, hardware, training, and premier consulting services across the Oracle product portfolio, and was the first partner to achieve specialization on Oracle Exadata Database Machine. Among its thousands of customers are the U.S. Departments of Defense and Homeland Security, federal civilian agencies, intelligence agencies, state and local governments, and healthcare and commercial organizations worldwide.

To achieve Oracle Database Appliance Specialized status, Cintra and Mythics met stringent requirements based on the needs and priorities of the customer and partner community. Those requirements measured the companies' technical competency and business results. Specialization gives Cintra and Mythics access to exclusive tools and resources to help them market and sell Oracle Database Appliance.

cintra.com

mythics.com

Higher One CASHNet Achieves Validated Integration with Oracle's PeopleSoft Campus Solutions

Higher One Holdings (Higher One), an Oracle Gold Partner, has achieved Oracle Validated Integration of its CASHNet payment processing suite 2012.2 with Oracle's PeopleSoft Campus Solutions 9.0. With this integration, colleges and universities—as well as students and their parents—can use CASHNet to process any payment, anytime, anywhere, and be assured that transactions occur in real time. Higher One CASHNet, which is mobile-friendly, met stringent requirements based on customer needs and priorities to achieve Oracle Validated Integration.

Higher One offers solutions and services to help college business offices manage operations and provide enhanced service to students.

higherone.com

Kewill Flagship Achieves Validated Integration with Oracle's JD Edwards EnterpriseOne

Oracle Gold Partner Kewill, a provider of shipping solutions for global trade and logistics, has achieved Oracle Validated Integration of Kewill Flagship with Oracle's JD Edwards EnterpriseOne 9.0 and 9.1. By achieving validated integration with two versions of JD Edwards EnterpriseOne, Kewill has enabled an easy upgrade path from JD Edwards EnterpriseOne 9.0 to 9.1 for Flagship users.

Kewill Flagship is a multicarrier shipping management system for moderate and high-volume shipping operations and complex distribution environments. It is available in on-premises and hosted models.

kewill.com

Leeyo Software and SKG are Oracle Ready

Two Oracle partners have achieved Oracle Ready status across multiple product lines through Oracle PartnerNetwork and the Oracle Exastack Ready program.

Leeyo Software, a provider of software for revenue recognition automation and management, has achieved Oracle Database Ready, Oracle WebLogic Ready, and Oracle Linux Ready status for RevPro 2.0. This achievement demonstrates that Oracle Gold Partner Leeyo's RevPro 2.0 is fully tested and supported on Oracle Database 11*g* Release 2, Oracle WebLogic Server, and Oracle Linux.

RevPro automates and manages every process in an organization's revenue cycle. It integrates with the quote-to-cash processes of any enterprise resource planning system, eliminating time-consuming and error-prone manual systems. RevPro can simplify revenue and cost-of-goods-sold management, monitor revenue-related activities in real time, and accelerate period-end close and financial reporting cycles. RevPro can be deployed on premises or in the cloud.

I-HUA CHEN

Oracle Gold Partner SKG S.A. has achieved Oracle Database Ready, Oracle WebLogic Ready, Oracle Solaris Ready, Oracle Linux Ready, and Oracle VM Ready status for SKG eSAMBO 1.0. This achievement demonstrates that eSAMBO 1.0 is fully tested and supported on Oracle Database 11*g* Release 2, Oracle WebLogic Server, Oracle Solaris 11, Oracle Linux, and Oracle VM 3.0.

eSAMBO is a modern, centralized retail chain management system. It manages heterogeneous, sophisticated retail chains from a central system instance. All data can be managed from the system owner's central level, system owner's chain stores, franchisee subchain level, and franchisee stores.

Oracle Database Ready, Oracle WebLogic Ready, Oracle Solaris Ready, Oracle Linux Ready, and Oracle VM Ready are part of the Oracle Exastack Ready program, which recognizes partners for developing, testing, and tuning their applications on the latest component products of Oracle Exadata Database Machine and Oracle Exalogic Elastic Cloud engineered systems.

leeyo.com

skg.pl/english/skg_products.aspx

RxLogix Achieves Oracle Argus Safety Specialization

RxLogix, an Oracle Gold Partner, has achieved OPN Specialized status for Oracle Argus Safety, Oracle's integrated platform of end-to-end pharmacovigilance solutions for regulatory compliance and comprehensive product stewardship.

RxLogix was recognized for its safety, risk management, and regulatory global practices in implementing Oracle Argus Safety at major life sciences organizations. Specialization also acknowledges the company's ability to complement those practices with value-added strategic software solutions. RxLogix, with a global staff of more than 70 employees, has completed more than 50 implementations of Oracle Argus Safety from its offices in the United States, Europe, and Japan.

rxlogix.com

Peloton Tests Oracle Exalytics, Launches Oracle Exalytics Practice

Oracle Platinum Partner Peloton, a U.S.-based professional services firm specializing in analytics and enterprise performance management, has installed Oracle Exalytics in its data center and tested the in-memory analytics platform on the roll-up of an Oracle Hyperion Planning application, using actual financial performance data for a multibillion-dollar company with multiple business units. The process, which usually takes more than five hours to complete, was finished in less than two hours—a 274 percent reduction in processing time compared to the baseline solution (a Microsoft Hyper-V virtualized server running release 11.1.2.1 of Oracle's enterprise performance management products on Windows 2008 R2 SP1 64 bit, with 4 virtual processors and 4 GB of RAM; details at bit.ly/OFpx9z).

Peloton has launched an Oracle Exalytics program that supports organizations with complex analytical processing needs—ranging from early-stage evaluation through implementation. Peloton serves organizations in the aerospace and defense, financial, healthcare, high technology, life sciences, manufacturing, and not-for-profit sectors.

pelotongroup.com

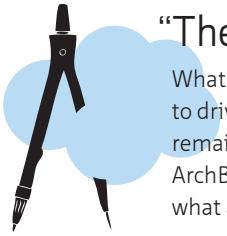

PODCAST

"The Role of the Cloud Architect"

What does cloud mean to IT architects? Cloud computing continues to drive the evolution of IT architecture, and IT architects who want to remain relevant need to up their game. In this Oracle Technology Network ArchBeat podcast, cloud architects Ron Batra and Dr. James Baty discuss what architects need to do to take their talents into the cloud. **bit.ly/Socws4**

Sangam Brings Oracle and Java Users Together

How are Oracle database, middleware, and applications users and Java developers collaborating in your community? In **Bangalore, India,** they are gathering under one roof at the All India Oracle Users Group's (AIOUG) fourth annual Sangam conference, November 2–3, 2012. Oracle ACE Director and AIOUG President Murali Vallath organized the event, which is cohosted by the India Oracle Applications Users Group and JUG (Java user group) Chennai. The largest independent Oracle user event in India, Sangam features more than 70 speakers from all over the world, provides 120-plus sessions, and attracts more than 600 delegates. **sangam12.info**

Oracle Solaris 11 Services Administration

Does taming the behavior of your operating systems services manually make you want to run for the hills? Help is here. The Service Management Facility feature, introduced in Oracle Solaris 10 and extended in Oracle Solaris 11, provides the discipline those unruly services need. In his two-article series, Oracle Solaris Principal Product Manager **Glynn Foster** helps you get the most out of the Service Management Facility. First, in "Introducing the Basics of the Service Management Facility in Oracle Solaris 11," Foster explains what the Service Management Facility does and how to perform basic services administration with it, including how to use four specific commands to get information about and manage your system services. Then, in "Advanced Administration with the Service Management Facility in Oracle Solaris 11," he describes how to use some of the more advanced features—including service bundles, which you can use to deliver custom configuration across systems. **bit.ly/Q80KN7**

Virtual Developer Day

Have you started looking at how **Oracle Fusion** will affect your own development projects? Jumpstart your education in Oracle Fusion development with this on-demand Virtual Developer Day that covers Oracle Application Development Framework (Oracle ADF), Oracle WebCenter Portal, Oracle Business Process Management, Oracle Business Intelligence, and more. Tracks include Introduction to Fusion Development and What's New in Fusion Development. Session topics range from Oracle Business Intelligence and Oracle ADF integration to a comparison of Oracle ADF to Oracle Forms, Oracle Application Express, and Microsoft .NET. Hands-on lab documentation that shows you how to develop rich Web applications with Oracle ADF and Oracle WebCenter Portal is also available. **bit.ly/TtdcKp**

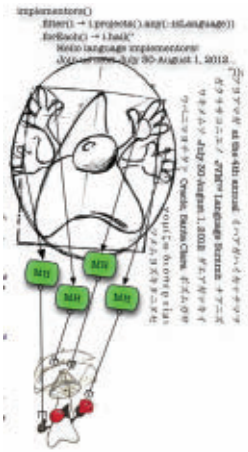

JVM LANGUAGE SUMMIT 2012

Are you a language designer, compiler writer, tool builder, runtime engineer, or VM architect? Did you miss the 2012 JVM Language Summit? The 2012 JVM Language Summit was held at Oracle's Santa Clara, California, campus July 30–August 1. If you missed it, or maybe you were there and want to revisit it, videos and presentations are available at **bit.ly/Or8AAR**

DINODIA PHOTOS/GETTY IMAGES; DUKE ILLUSTRATION/DAVE FONTAINE

Architects Matter

Making sense of the people who make sense of enterprise IT

Why do architects matter? Oracle Enterprise Architect Eric Stephens suggests that you ask yourself this question the next time you take the elevator to the Oracle offices on the 45th floor of the Willis Tower in Chicago, Illinois (or any other skyscraper, for that matter). If you had to take the stairs to get to those offices, who would you blame? "You get the picture," he says. "Architecture is essential for any necessarily complex structure, be it a building or an enterprise."

Complexity in enterprise IT is unavoidable—to a point. "You need a role to manage out the unnecessary complexity and prevent architectural entropy from taking hold," Stephens says. "I've seen a number of companies with little to no architectural presence, and it shows. The duplication of IT assets, as well as inefficient and often ossifying integration between systems, leads to the business's inability to respond to change."

Stephens' colleague Pat Shepherd, also an Oracle enterprise architect, shares a similar opinion. "Enterprise architects matter because history has demonstrated that project-based IT development does not lead to reuse, consolidation, or any of the economies of scale that come from enterprise architecture," he says. "Enterprise architecture is the only real hope an organization has for delivering IT as a strategic asset."

That awesome responsibility filters downward from the enterprise level to the level of the solution or IT architect.

"Architects provide the counterweight to the business," says Oracle ACE Director Lucas Jellema, CTO at AMIS Services. It's a matter of achieving balance between the interests of business and IT stakeholders. According to Jellema, the IT architect's focus is on preparing the IT environment to meet key business demands. "Therefore the business should not regard architects as opponents, the frustrators of new developments," he says. "The business should recognize the architect as an enabler."

> **"Architecture matters because it optimizes the use of resources in order to achieve the most-important qualities in each IT system."**
>
> *—Manuel Ricca, IT Architect*

Development teams need to see the architect in a similarly positive light. "The architect is also the guide, coach, and sparring partner for the development team, and perhaps the administrators as well," Jellema says. "The architect will instruct and monitor the development team in how they design and implement the software artifacts and how they make use of the various tiers of underlying infrastructure."

That focus on efficient, intelligent use of resources is key to defining the architect's value.

"Architecture matters because it optimizes the use of resources in order to achieve the most-important qualities in each IT system," says Manuel Ricca, an IT architect with a large European financial institution. Ricca points out that stakeholders in security, operations, and finance and even the CEO may not see beyond their respective concerns. It's the architect's job to relieve them of their tunnel vision. "The architect is responsible for getting agreement on the relative priorities of those different concerns, and for designing the system structures that allow the organization to achieve the desired qualities," says Ricca.

The ability to see and understand the interconnection and interdependencies between disparate systems is another key value proposition for architects, and that value is increasing in the evolving IT landscape.

"When software could work in silos, architects weren't overly useful," says Derek Sharpe, director of the Oracle Fusion Middleware A-Team (that's *A* for *architecture*). But times have changed. "In today's world, business users expect their systems to communicate with each other, share information, and be highly available."

A good architect knows how the pieces fit together and can turn that information into a cohesive strategy that spans a company's IT infrastructure and forms the design basis for a reliable platform that satisfies IT and business stakeholders. "But most importantly," Sharpe says, "the architect is management's technical communicator, someone who can articulate to both the development teams and the business what is necessary to meet current and future requirements."

To put it another way, if the elevators don't work, nothing works. ◀

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
bit.ly/9X6j2T

GET more architect information
oracle.com/technetwork/architect

A Path to Leadership

New ODTUG program grooms leaders for user group success.

Monty Latiolais' first attempts to get involved with his Oracle user group came up dry. Latiolais is the president of Oracle Development Tools User Group (ODTUG) but recalls a time when his submissions to the user group's events were rejected and his bid for a committee seat flopped. "The user group meant a lot to me, and I wanted to get more involved," he says. "But I couldn't find a way to get over the hurdle." Now that he's president, he wants to give enthusiastic ODTUG members a more direct path to leadership. In June 2012 he announced the ODTUG Leadership Program to do just that.

The six-month program will run January through June 2013, culminating at ODTUG's Kscope13 conference in New Orleans, Louisiana, June 23–27. It will cover a range of leadership topics including effective communication, mentoring, ethics, goal setting, and accountability. Training will consist of a series of classes and conference calls with the current executive committee. The program is limited to a small group of participants, because each one will be assigned an ODTUG board member as a mentor and will receive a complimentary pass to ODTUG Kscope13 upon graduation.

The program is designed to give ODTUG a new group of trained leaders, and gives participants visibility when it comes to filling committee positions and other leadership opportunities, says Latiolais. "This is a way for members to, in a sense, raise their hands and say, 'Pick me, pick me!'" he says.

Latiolais has been saying, "Pick me!" since 2008, when he made the decision to pursue speaking and leadership roles in the user group. "I started out as an audience member who put speakers on a pedestal," Latiolais says. "Later I learned that they are just people who know something about the technology that they want to share. I approached the people I considered rock stars at the time and found that they were very willing to engage me and encouraged me to get involved."

ODTUG President Monty Latiolais

Thus began Latiolais' journey to leadership. He was appointed to ODTUG's conference committee in 2009 and then campaigned unsuccessfully for the ODTUG board of directors. He ran again the following year and succeeded. In 2011 he was elected to the executive committee as vice president. At ODTUG Kscope12 he was elevated to president. "It's gratifying to have a goal and work your way up and achieve that goal," says Latiolais. "Professionally it's done wonders for me in the workplace."

Latiolais attributes promotions and new leadership positions at work to skills he gained from his ODTUG experience. He wants to give a new crop of leaders the same opportunity. "If ODTUG can better equip a member to become a leader in his or her workplace, that's real value," says Latiolais. "The employer benefits because they have someone who is better equipped and more confident. They now have a person with leadership skills, communication skills, ethics training, and mentoring skills."

Another benefit to employers, says Latiolais, is connections. "My company is gaining visibility in the IT community because of my involvement in ODTUG," he says.

With 25,000 members worldwide and growing, ODTUG is the model of a successful user group. The group has opened up to new and acquired Oracle technologies, and blends an ethic of service with hard-nosed technical content delivery. "ODTUG's success thus far has been the result of outstanding leaders," says Latiolais. He credits former president Mike Riley and former executive director Kathleen McCasland with leaving him a strong community to lead. "The real challenge is to continue what these outstanding leaders have done up to this point," he says.

Latiolais believes the emphasis on service is a key to the group's success. "ODTUG is all about giving back," he says. The ODTUG Leadership Program will be, Latiolais hopes, another way to serve the Oracle community. "We give employers an employee with new skills and confidence," he says. "For members with a sincere desire to give back to the community, we provide a path from the audience to the podium." ◀

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

NEXT STEPS

LEARN more about ODTUG
odtug.com

WATCH the interview
oracle.com/oramag/upclose

PAUL S. HOWELL

Tech Destiny

The power of MySQL, life choices, and staying engaged with Oracle Database

WAGNER BIANCHI

ORACLE ACE

Company: WBConsulting, an IT consulting firm

Job title/description: Senior principal consultant, running all processes involving MySQL technologies

Location: Belo Horizonte, Brazil

Oracle credentials: Oracle Certified Associate (MySQL 5, MySQL 5.1, MySQL 5.5), Oracle Certified Professional (MySQL 5 Developer, MySQL 5 Database Administrator), and Oracle Certified Expert (MySQL 5.1 Cluster Database Administrator), with nine years of experience using Oracle products

What are your favorite tools on the job? MySQL Enterprise Backup, MySQL Workbench, and the Enterprise Monitor feature of MySQL Enterprise Edition. Those tools have made it possible to have all MySQL areas under control with no necessity to access mysqld through the command line.
What technologies have most changed your life? MySQL and Linux. MySQL was my passport to many of the companies I've worked with. The software gave me great perspective and opened my mind to many other things about database servers. It led me to start studying Linux, which showed me how to monitor services and configure them for greater performance.
How are you using MySQL and cloud computing? Some time ago we developed a system in which MySQL was running in the cloud to support our company systems. This was an important advance, and it has held an important place in our architecture plans. MySQL is easy to use, easy to configure, and easy to drive into the cloud. There is a great opportunity for MySQL to dominate this field.

EMRE BARANSEL

ORACLE ACE

Company: Turkcell, a communications and technology company

Job title/description: Oracle DBA, responsible for managing the standby databases in Turkcell's IT environment

Location: Ankara, Turkey

Oracle credentials: Oracle Certified Professional (Oracle Database 10g), with seven years of experience using Oracle products

How did you get started in IT? When I was 14, I passed the very competitive entrance exam for my high school, and my dad decided to give me a reward. He asked me if I wanted a computer or a pair of skis—I guess I made my life's choice when I chose the computer.
How are you using social media in your work these days? I generally use it to keep myself up to date and share valuable information about Oracle technologies. I mostly use Twitter to get and share Oracle-related information, but I do use Facebook and LinkedIn for information sharing inside our local Oracle user group, TROUG [Turkish Oracle User Group, troug.org]. There are also some personal Oracle blogs, OTN forums, and Oracle mail groups that I follow.
What would you like to see Oracle, as a company, do more of? I'd like to see Oracle with products in end-user technology. Oracle is a company that always changes the standards and achieves the best. It would be great to use Oracle smartphones or Oracle computers with Oracle operating systems.
What is your favorite thing to do outside of work? I enjoy exercising and spending time with my family.

JEFF JACOBS

ORACLE ACE

Company: PayPal, a global online payment company

Job title/description: Senior data architect, modeling, designing, and troubleshooting site databases

Location: San Jose, California

Length of time using Oracle products: More than 20 years

What's your favorite tool on the job? Quest Software's Toad for Oracle is still one of my favorite tools, and I also find reports generated by the Automatic Workload Repository feature of Oracle Enterprise Manager incredibly useful. Most of my performance work is after the fact, and it most often involves correcting application issues. Automatic Workload Repository reports provide the necessary data to identify the cause of performance issues.
What technology has most changed your life? Definitely Oracle Database. It has not only provided me a technology that I enjoy and find constantly engaging but has also given me a wide variety of opportunities to engage in user group activities, conferences, presentations, and writing. I am honored to serve on the board of directors of the Independent Oracle Users Group Exadata special interest group, and privileged to have served on the ODTUG board of directors for 15 years.
What's your go-to Oracle reference book? My favorite authors are Tom Kyte, Jonathan Lewis, Richard Niemiec, and Steven Feuerstein. I still prefer real books that I can mark, flag, and read in an easy chair. ◀

Learn more about the Oracle ACE program at oracle.com/technetwork/community/oracle-ace.

Oracle CEO Larry Ellison discussed Oracle's comprehensive cloud offerings in one of two featured keynotes.

Oracle President Safra Catz shared views on business strategy and vision at the Executive Edge @ OpenWorld event.

More than 50,000 attendees participated in seven events across Oracle's unprecedented week of conferences.

Oracle President Mark Hurd appeared live from Oracle OpenWorld on CNBC's *Closing Bell* with host Maria Bartiromo.

BY KAREN SHAMBAN

ENGINEERED TO EDUCATE, EXPLORE, ENGAGE

2012's Oracle OpenWorld and JavaOne conferences brought technologists to San Francisco for a week of learning, networking, and looking forward.

From September 29 through October 5, 2012, Oracle's week of conferences in San Francisco, California, brought together a wide range of attendees interested and invested in Oracle technologies, announcements, and expertise. In addition to the long-standing flagship conferences Oracle OpenWorld and JavaOne, new this year were conferences focused on customer experience (Oracle Customer Experience Summit @ OpenWorld), partners (Oracle PartnerNetwork Exchange @ OpenWorld), MySQL (MySQL Connect), C-level and line-of-business executives (the Executive Edge @ OpenWorld), and Java embedded technology (Java Embedded @ JavaOne).

With this diverse content offering, the numbers were impressive. More than 50,000 people from 144 countries attended. More than 3,500 speakers conducted presentations and workshops over seven days. More than 700 exhibits and demos from Oracle and its partners filled four exhibition halls. Of course, not everyone could attend in person, but between live-streaming and on-demand video, there were more than 1 million views on oracle.com and YouTube (youtube.com/oracle), virtually increasing attendance by 20 times.

Press and customers—including attendees from 100 percent of the Fortune 500—were eager to hear the latest announcements from Oracle, and Oracle did not disappoint. Significant technology

HARTMANN STUDIOS

The thousands of conference sessions required the use of all of San Francisco's Moscone Center, plus the conference space of 10 hotels, and numerous other venues.

announcements included Oracle Exadata X3, a lightning-fast database in-memory machine; Oracle Solaris 11.1, which is optimized for cloud computing deployments; Oracle's social relationship management solutions, designed to take advantage of big data to help create improved experiences for customers; and expanded Oracle Cloud services.

Oracle conferences are known for creating unique opportunities for networking and community building, and this year was no exception. Not only did tens of thousands attend the renowned Oracle Appreciation Event on Treasure Island featuring Pearl Jam and Kings of Leon, but the Oracle OpenWorld Musical Festival—new this year—brought 30 performers (including Grammy winners Joss Stone and Macy Gray) to multiple venues for five nights of rock, blues, ska, and alt-country. Other notable happenings included a conversation with *Moneyball* author Michael Lewis—and the conferences captured national attention at very close range when the set of CNBC's *Closing Bell* moved to Moscone Center for two days, where host Maria Bartiromo conducted live interviews with Oracle CEO Larry Ellison and Oracle President Mark Hurd.

From the opening keynote and Welcome Reception to the final It's a Wrap! celebration, there was much to see, do, and learn at Oracle's week of conferences. Mark your calendar: next year's events run September 22 through 26. ◀

Karen Shamban is a senior director of communications and editorial at Oracle and has run dozens of events during her corporate career.

Oracle introduced its next-generation Oracle Exadata X3 systems.

Attendees flocked to four exhibition halls, which featured more than 700 exhibits and Oracle demos.

Oracle tented two streets—Howard and Taylor—to accommodate attendee meals, meetups, and lounging.

Java creator James Gosling made a surprise appearance at the JavaOne community keynote to talk about Liquid Robotics' Wave Glider.

Networking is always an attraction at the conferences, and attendees had numerous formal and informal opportunities, such as the Oracle OpenWorld Welcome Reception, throughout the week to do just that.

Joss Stone was one of 30 performers at the debut Oracle OpenWorld Music Festival, which gave attendees the chance to get their groove on after sessions.

Oracle Executive Vice President Thomas Kurian's keynote focused on how Oracle's cloud platform and application services are transforming business.

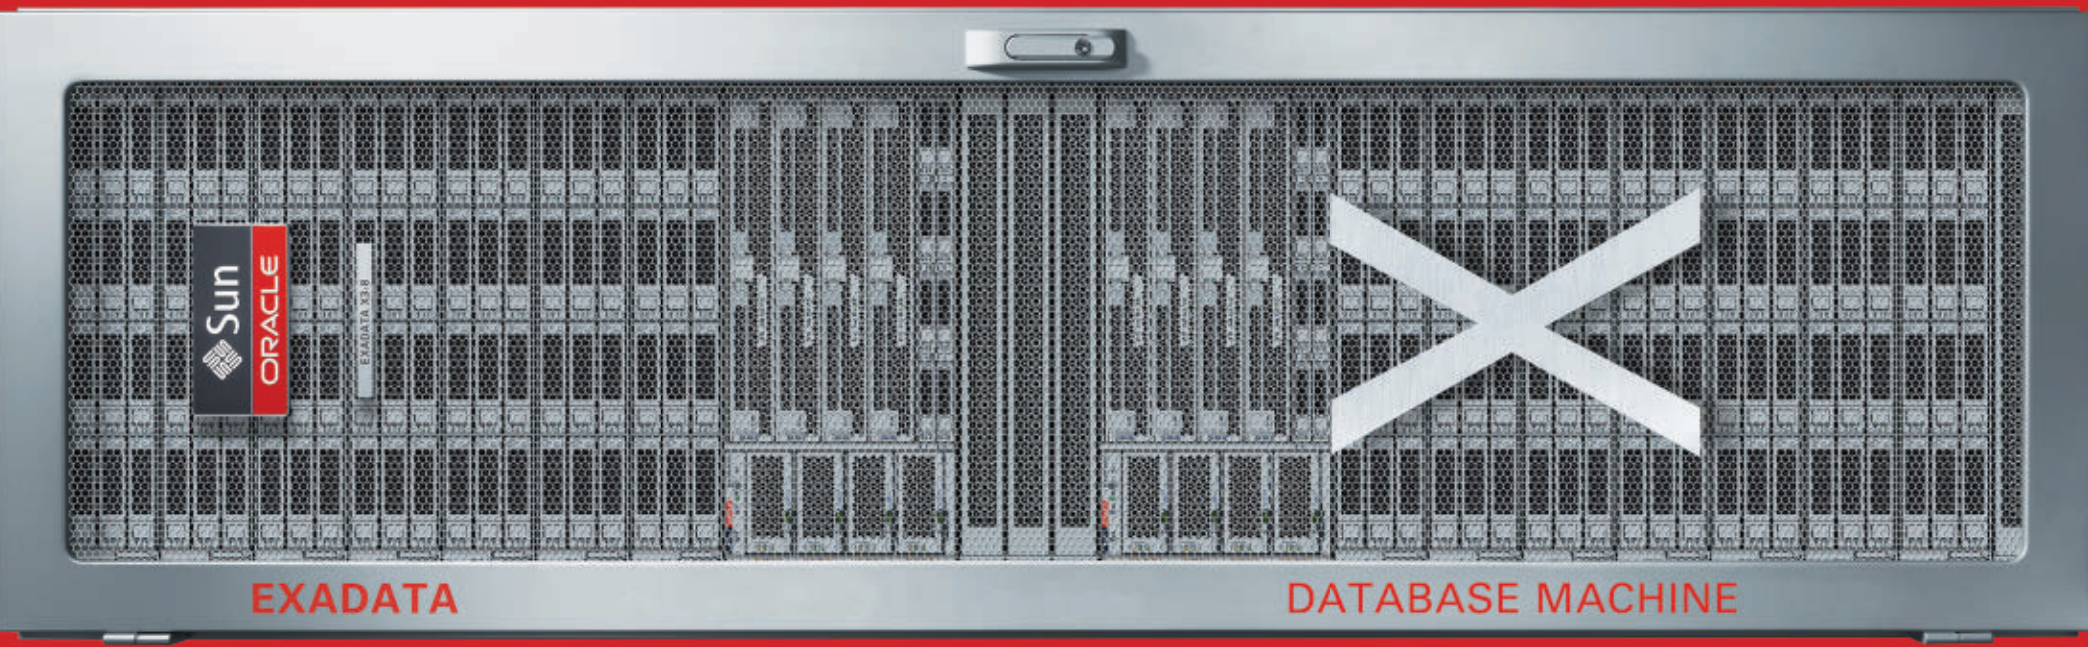
Sun
ORACLE
EXADATA X3-8
EXADATA
DATABASE MACHINE

BY DAVID A. KELLY

APPS ON ORACLE EXADATA

CONSOLIDATE BUSINESS APPLICATIONS, LOWER IT COSTS, BOOST PERFORMANCE, AND SUCCEED WITH ORACLE EXADATA.

Consolidation may not be the first word that comes to mind when you think of strategies for managing business growth, but when business growth means that your current business applications platform can no longer support your business, consolidation is certainly worth consideration.

Rather than adding separate hardware and software—configuring, testing, reconfiguring, retesting, and reconfiguring and retesting again to update a business applications platform—organizations are transforming their business technology as they are transforming their businesses. By consolidating business applications with Oracle Exadata, they are simplifying their information technology; lowering costs; and improving system performance, scalability, and reliability as part of that transformation.

INDUSTRIAL TRANSFORMATION

Praxair is one company that's deriving extreme value from the combination of Oracle Applications and Oracle Exadata.

With more than 26,000 employees in 50 countries, the Fortune 300 company, located in Danbury, Connecticut, is one of the largest industrial gases companies in the world and the largest in North and South America. The company has grown from about US$7.5 billion in 2005 to US$11 billion in 2011.

Key factors in Praxair's growth have been its use of shared services and the global replication of best practices. The organization has standardized on a globally consistent Oracle's JD Edwards EnterpriseOne enterprise resource planning (ERP) platform supported by an Oracle Exadata technology stack running in its global data center.

Although the company has different business units around the world, it has deployed a standardized JD Edwards EnterpriseOne configuration across them. Previously, Praxair had deployed an older JD Edwards EnterpriseOne footprint, set largely on top of a Microsoft technology stack. Now, the JD Edwards EnterpriseOne ERP platform is supported by four Oracle Exadata machines: two Exadata Database Machine X2-8 Full Racks for production and disaster recovery, an Exadata Database Machine X2-2 Half Rack for development, and an Exadata Database Machine X2-2 Quarter Rack for testing. All of the company's global environments are being deployed on those four racks of Oracle Exadata.

"We've had a pretty aggressive data center consolidation strategy over the past couple of years," says Praxair CIO Marc Franciosa. "But the transition to JD Edwards [EnterpriseOne] running on Oracle Exadata has definitely saved us significant time and money."

Deploying Oracle Applications on Oracle Exadata has also changed the way that Praxair does business. With its previous system, scaling out or up required more infrastructure, with a high-speed switching fabric, and more systems engineering as the system grew.

"Our Oracle Exadata deployment has removed this added burden," says Franciosa. "With Oracle Exadata, we can spend our time understanding the business processes and on parts of the solution where we need to integrate other systems."

In addition to its transactional systems, Praxair is consolidating its data warehouse on Oracle Exalytics and Oracle Exadata.

"There are definitely inherent benefits to bringing our transactional systems and data warehouse systems together on Oracle Exalytics and Oracle Exadata," says Franciosa. "Having real-time and near-real-time access to every type of operational report we need enables us to efficiently drive production, distribution, and route planning processes."

A few years ago, when Franciosa looked around Praxair's IT environment, he saw an IT group that was supporting every country with its own set of servers and applications. "We had no way to scale our IT environment to meet the types of business changes we wanted to make," says Franciosa. "Even if we moved everything into a global data center, we would have ended up with a lot of capacity that was significantly underutilized."

At full deployment, Franciosa expects to have about 13,000 employees using the JD Edwards EnterpriseOne and Oracle

> **"With Oracle Exadata, we can spend our time understanding the business processes."**
>
> *—Marc Franciosa, CIO, Praxair*

As part of an aggressive data center consolidation strategy, Praxair's transition to JD Edwards EnterpriseOne running on Oracle Exadata saved both time and money, says Marc Franciosa, CIO at Praxair.

CHRISTOPHER LANE/GETTY IMAGES

SNAPSHOT

Praxair
praxair.com
Location: Danbury, Connecticut
Industry: Industrial gases
Size: US$11 billion in 2011
Oracle products: JD Edwards EnterpriseOne applications, Oracle Hyperion Financial Management, Oracle Hyperion Planning, Oracle E-Business Suite process manufacturing applications, Oracle Product Information Management, Oracle Business Intelligence solutions, Oracle Exadata, Oracle Database

SNAPSHOT

Alliance Data
alliancedata.com
Location: Plano, Texas
Industry: Loyalty marketing solutions, consulting, and programs
Size: US$3 billion in 2011
Oracle products: PeopleSoft financials and human resources applications, Oracle E-Business Suite, Oracle Essbase, Oracle Enterprise Manager, Oracle Real Application Clusters, Oracle Exadata, Oracle WebCenter, Oracle Data Warehouse, Oracle Database

Running Oracle Applications on Oracle Exadata enables Alliance Data to deliver additional IT capacity and keep up with increased application needs, says Collin Harrison, vice president of finance and IT at Alliance Data.

Exadata system, with an additional 6,000 to 10,000 using it indirectly. For example, the drivers in the company's trucks are connected through onboard computers that interface indirectly into the JD Edwards EnterpriseOne environment so Praxair can calculate things such as the cost of the sale and the cost of the trip. As drivers scan cylinders at customer sites, the data goes into a tracking system and then straight into JD Edwards EnterpriseOne. "Almost every Praxair employee will touch these systems at some point," notes Franciosa.

Praxair made the transition in steps. Initially, it simply moved a lot of its legacy environments to a more complete Oracle technology stack. But Praxair didn't focus just on migration: the company made business transformation the goal. "The real opportunity for us as we move to a global environment and deploy the shared services around the world is that ability to enable some pretty big business transformations, which will give us even more benefits," says Franciosa.

Praxair was an early adopter of Oracle Exadata, and its deployment has been successful. "I don't think we would have done anything differently, because so far we've had a very smooth experience," he says.

> **"Running applications on Oracle Exadata saves us significant productivity costs."**
>
> *—Collin Harrison, Vice President of Finance and IT, Alliance Data*

ADVENTURES IN GROWTH

Keeping up with growth is a big part of the business transformation process, and it has been a big part of business and technology at Alliance Data.

"As our company continues in high-growth mode, we have to make sure we can deliver additional IT capacity and keep up with increased application needs," says Collin Harrison, vice president of finance and IT at Alliance Data. "Running our Oracle Applications on Oracle Exadata enables us to do that. The platform is tried and true, thoroughly tested, and stable."

Based in Plano, Texas, Alliance Data is a US$3 billion business-to-business company. It operates three independent, high-growth, high-performance Fortune 1000–size companies—Alliance Data Retail Services, Epsilon, and LoyaltyOne—that specialize in loyalty marketing solutions. Across each of the businesses, Alliance Data's services enable its clients to create deeper and longer-term relationships with their best customers.

Alliance Data's corporate IT department runs an enterprise stack of applications that are used at the company's corporate office as well as by many of its 9,000 associates across North America, at Alliance Data Retail Services, Epsilon, and LoyaltyOne. The applications stack is based on Oracle solutions that include Oracle's PeopleSoft financials and human resources applications as well as Oracle Hyperion Planning for forecasting, budgets, and models; Oracle WebCenter for the corporate portal; and Oracle Data Warehouse and Oracle Business Intelligence analytics solutions.

In 2011, Alliance Data chose to upgrade its existing servers to support the company's growth—nearly 23 percent over the last three years. The company's IT team investigated several alternatives, including combinations of distributed systems. However, analysis showed that under traditional approaches, Alliance Data's data center footprint would continue to grow too rapidly, resulting in higher overhead as well as upfront capital expenditures.

Instead, Alliance Data decided to support its applications with Oracle Exadata. By running its Oracle Applications with Oracle Exadata, the corporate IT team found they could dramatically reduce the amount of hardware they would have to deploy. The solution also

PAUL S. HOWELL

> **"When we scope out a new project using applications running with Oracle Exadata, we're already three to four weeks ahead of where we used to be with our old approach."**
>
> *—Collin Harrison, Vice President of Finance and IT, Alliance Data*

increased flexibility and reliability, while reducing the data center footprint. The company soon migrated its PeopleSoft and Oracle WebCenter systems to Oracle Exadata.

"Our time to uptime with our first Oracle Exadata Database Machine was substantially quicker than we had planned because so much of the solution was pre-engineered and preconfigured," says Harrison. "We had an Oracle database up and running in a very short period of time—compared to weeks of time with other solutions—and we could immediately get to the process of installing our applications on it."

Ultimately, Alliance Data consolidated seven database servers onto two Oracle Exadata quarter racks, one for development and one for production. Many of the company's employees took notice. "Things that used to have a little bit of a delay are now much faster," Harrison says. "For the first few weeks, we received a surprising amount of user feedback, all noting how fast the system was responding."

Harrison also appreciates Oracle Exadata's ability to stay ahead of the company's continued growth. And while increased speed, reduced complexity, and a solid growth path are compelling benefits from Alliance Data's move to Oracle Exadata, the company is also excited about greater efficiencies as a result of running its applications on Oracle Exadata.

For example, the move to Oracle Exadata delivered increased developer efficiency that reduces the time it takes to deploy updates and new applications. "Running applications on Oracle Exadata saves us significant productivity costs," says Harrison. "Our development refreshes can be ready over a lunch hour now, instead of taking overnight."

But that's not all. "Just the simple fact that everything runs faster allows us to schedule more intelligently—based on business needs, not on IT limitations," says Harrison. For example, the company previously had to balance scheduling some operations to avoid the company's month-end close process or other periods where workload exceeded normal capacity. By running Oracle Applications with Oracle Exadata, Alliance Data now is able to schedule system updates and new application rollouts based on business need.

Running applications at Alliance Data with Oracle Exadata, Harrison has also noticed other benefits. "Things are not only running faster, but they're breaking down less. That means less troubleshooting time, fewer all-hands-on-deck incidents, and more stability."

Just as important, with applications running with Oracle Exadata, Alliance Data's corporate developers and other IT teams can remain focused on their strategic goals, such as adding new capabilities to support the company's continued growth, and providing the highest levels of service and support.

"When we scope out a new project using applications running with Oracle Exadata, we're already three to four weeks ahead of where we used to be with our old approach," says Harrison. "Previously, new applications required new hardware, which typically took three to four weeks to requisition, configure, deploy, and test."

Enabling business transformation requires an agile IT infrastructure—one in which an organization's hardware and software work together to enable rapid business change. Alliance Data has achieved that with its combination of Oracle Applications and Oracle Exadata.

"One of the big benefits we've seen from running Oracle Applications with Oracle Exadata is that it's much closer to plug-and-play," says Harrison. "The Oracle hardware and software are engineered to work better together, especially when it comes to running PeopleSoft and many other Oracle Applications with Oracle Exadata." ◀

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

Oracle Exadata in the Cloud

For many organizations, a good place to start consolidating and creating a highly scalable Oracle Exadata–based global IT environment is with Oracle Managed Cloud Services.

Oracle Managed Cloud Services offers comprehensive end-to-end services across applications and technology including a broad range of preconfigured, pre-integrated, and pretested certified configurations of Oracle products, making it easier to deploy and manage Oracle technologies such as Oracle Applications and Oracle Exadata.

"If an organization wants to deploy Oracle E-Business Suite or Oracle's PeopleSoft on Oracle Exadata, Oracle Managed Cloud Services provide an integrated 'apps-to-disk' management framework for deployment in the cloud or at the customer's location," says Guru Shashikumar, director of product management for Oracle Cloud Services at Oracle.

For Shashikumar, consolidating systems isn't just about removing things. It's also about enabling new opportunities. "The ultimate benefit of Oracle Managed Cloud Services for Oracle Applications and Oracle Exadata is to accelerate business transformation and maximize the value of the transformation with faster uptake of the latest Oracle technologies and lower, more predictable costs," says Shashikumar. "Oracle Exadata enables organizations to consolidate legacy systems and rationalize a legacy infrastructure into a very high-performing engineered system, with integrated hardware and software designed to deliver optimal performance for online transaction processing and data warehousing applications."

NEXT STEPS

LEARN more about
Oracle Exadata
oracle.com/exadata
bit.ly/VG2lj1

Oracle Cloud
cloud.oracle.com

ORACLE APPLICATION DEVELOPMENT FRAMEWORK AND ORACLE JDEVELOPER

Master and Commander

Implement a reusable command pattern template strategy using contextual events for unlimited parent–to–region interaction.

I-HUA CHEN

An Oracle Application Development Framework (Oracle ADF) region is an aspect of the Oracle ADF Controller that executes bounded task flows in defined layout areas within a view, page, or page fragment, without forcing a refresh of the entire view. Think *portlets* and you have the right idea.

There are many ways to establish communication between a view and a bounded task flow exposed in an Oracle ADF Region. These Oracle ADF Region interaction strategies include contextual events, data exchange via input parameters, managed bean injection, shared Oracle ADF data control, and many others. Developers should choose an interaction strategy that best suits the use case they want to implement.

One common Oracle ADF Region use case that frequently appears on the Oracle Technology Network Oracle JDeveloper forum is a global toolbar or menu bar within a parent view. The goal is to have the global toolbar perform actions, such as data iteration, create/update/delete operations, or context-sensitive help display, on the view displayed in a separate region.

This article will show how to implement this use case by borrowing the idea of the command pattern from object-oriented programming. A command pattern implementation, deployed as an Oracle ADF library, is provided for you to use within your custom development projects.

AN ADVANCED TOPIC

Unlike previous *Oracle Magazine* articles, which were more suitable for beginners, this article covers an advanced topic in Oracle ADF task flow communication. It's assumed that you have an understanding of core Oracle ADF technology concepts such as bounded task flows, task flow templates, Oracle ADF Regions, page templates, and the Oracle ADF binding layer.

THE PROBLEM STATEMENT

The use case for this article is comparable to a universal remote control that, independent of the receiving device, always issues the same set of commands.

Using this analogy, the receiving device in Oracle ADF is the current view displayed in an Oracle ADF Region. The remote control in this scenario is a toolbar or menu bar, in the parent view, that sends commands into the Oracle ADF Region to process in the context of the current view. The commands passed from the parent view to the region continue working even when the current view changes due to users navigating within the task flows. The interpretation of a given command, however, may be different; it is up to each receiving region to react properly to the signals from the parent view.

A COMMAND PATTERN SOLUTION

The command pattern is a well-known design pattern in object-oriented programming that uses a command object to encapsulate all the information needed to invoke methods on other objects. The command pattern describes three acting parts: a *client* that creates the command object, including information about the actions to invoke on a target object; an *invoker* that dispatches the command defined in the command object; and a *receiver* that represents the target object on which the command is executed.

With slight variations applied to the invoker and receiver, the command pattern defined for object-oriented programming can be used with Oracle ADF and task flows to implement the Oracle ADF Region interaction use case covered in this article.

In Oracle ADF, the invoker object is replaced by contextual events. They take a payload (command object) defined by the client and broadcast it as an event to registered listeners (receiver). The event receiver does not execute commands directly; instead, it further dispatches the execution to a managed bean loosely coupled with the current view. For example, in Figure 1, Action 1 and Action 2 are managed beans whose roles are to interpret incoming action requests and to execute the appropriate methods for the view.

Figure 1 shows a simplified architecture of the command pattern, as implemented with Oracle ADF. In this architecture, the buttons in the parent view represent the command pattern client. When a button is pressed, it creates a payload object (command object) that includes the name of the action to invoke in the context of the view displayed in the region, along with optional additional information developers may want to provide to the receiving end. The event receiver in the Oracle ADF Region is a data control method that is referenced in the PageDef file (binding container definition) of each view.

The solution in Figure 1 retains task flow encapsulation and does not expose objects located within the task flow to the outside consumer. One change to the command pattern in object-oriented programming is that the Oracle ADF Region container allows developers to check for commands supported by the current view to let them disable or hide command buttons accordingly.

THE COMMAND PATTERN IMPLEMENTATION

The code sample for this article (at bit.ly/OHuJLe) is packaged by implementing the

command pattern design for Oracle ADF Region interaction in a separate Oracle JDeveloper workspace. This code is deployed into an Oracle ADF library and can be used in the sample application provided as well as in your own custom applications.

Figure 2 provides an overview of the command pattern implementation. Both sides of the command pattern implementation, the client (producer) and the receiver, are developed as page templates. Each page template has its own Oracle ADF binding file (PageDef) defined. The PageDef file of the producer page template contains the contextual event metadata that defines the event name, the action type to listen for, and a reference to the location of the payload object.

The PageDef file associated with the receiver page template contains the event handler definition and a mapping that binds the producer event name with it. A lifecycle configuration in the form of an Oracle ADF Region Controller is used to let the calling side (producer) know the action commands supported by the current view (receiver).

A task flow template is provided for any implementing receiver task flow to consume. The task flow template ensures that the required managed beans and input parameters are configured on the receiving task flow and exposed to the Oracle ADF application developer.

The public APIs exposed to the application developer by the task flow command pattern template consist of two page templates, a task flow template, and a set of abstract classes for the developer to extend.

With the provided sample implementation, the assumption is that the view action handlers shown in Figure 1 are implemented as JavaServer Faces (JSF) managed beans in backing bean scope within the bounded task flow.

The name of the managed bean should be the same as the view activity, along with a prefix of *tfcp*, which stands for *task flow command pattern*. So if the view activity name is DepartmentsView, then the associated backing bean name configured in the task flow configuration file should be tfcpDepartmentsView.

A third pillar in this architecture is a set of implementation classes containing logic to invoke, dispatch, and handle the defined contextual events. Following the Oracle ADF framework naming convention, these classes are stored in a package structure starting with "internal." Application developers don't need to know about these classes and should not have any need to use them.

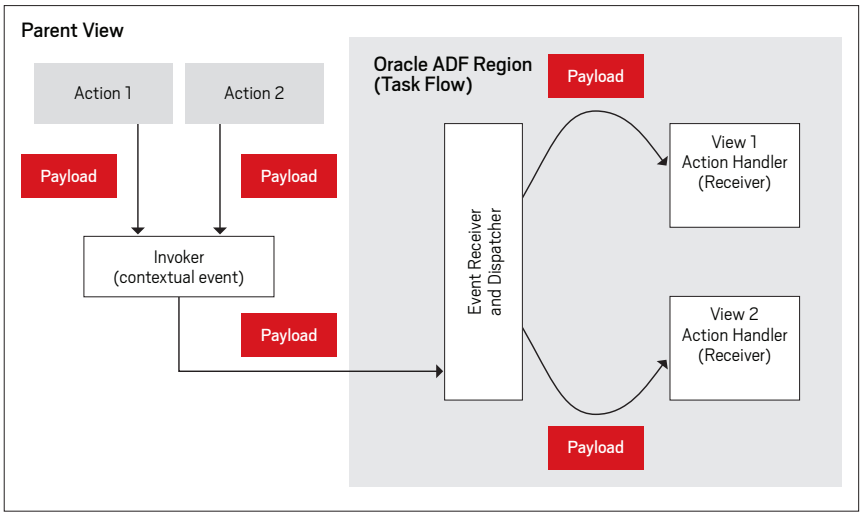

Figure 1: Technical implementation outline

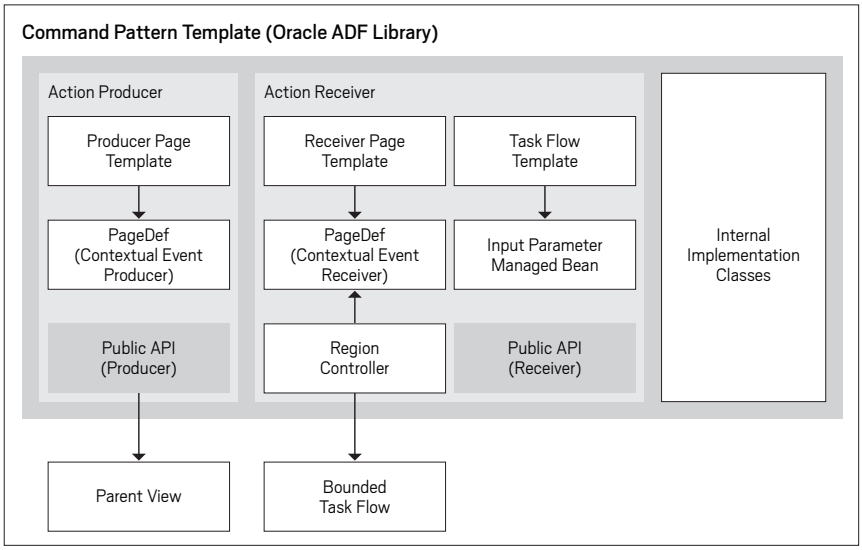

Figure 2: Task flow command pattern template architecture

GETTING STARTED

Before exploring how the task flow command pattern solution works, you should first download the command pattern sample application at bit.ly/OHuJLe and unzip the o62adf-1841823.zip file to a local folder on your computer. To open the sample application, you need the studio edition of Oracle JDeveloper 11*g* Release 2 (11.1.2.2.0), available as a free download on Oracle Technology Network. You also need an Oracle database instance with an unlocked HR schema.

The sample application contains two folders:

- **CommandStrategyTemplates** contains the task flow command pattern template source code. The command pattern template is deployed to an Oracle ADF library, TaskFlowCommandStrategyLib.jar, which is referenced in the OraMagSample sample application ViewController project properties.

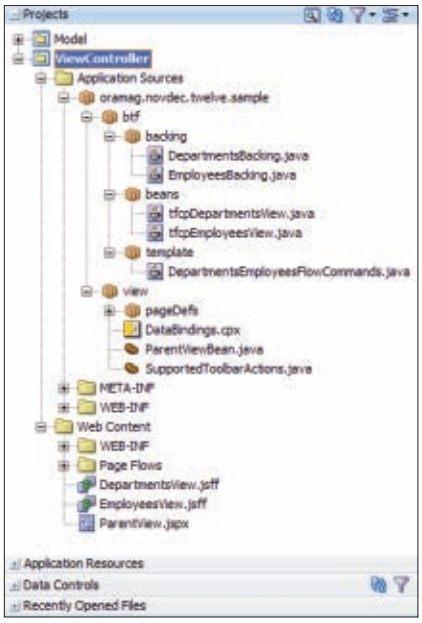

Figure 3: Sample application ViewController project

You need to open this folder only if you are interested in the task flow command pattern implementation, or if you want to modify the library to suit your needs. The source code is well documented, using Java comments.

- **OraMagSample** holds a sample application that implements the task flow command pattern template. We will use this application to further explain how to use the template in custom projects. The sample application uses a copy of the TaskFlowCommandStrategyLib.jar file in its Lib folder and does not depend on the CommandStrategyTemplates folder.

To run the sample application, you also need to change the database connection to point to your HR database schema:

1. Launch Oracle JDeveloper 11*g* Release 2. Select **File** -> **Open,** and then navigate to the directory containing the unpacked sample application.
2. Open the **OraMagSample** folder, select the **OraMagSample.jws** file, and click **Open**. The workspace opens in Oracle JDeveloper.
3. Select **View** -> **Database** -> **Database Navigator**, and expand the **OraMagSample** node to display the **hrconn** node.
4. Right-click the **hrconn** node, and select **Properties** from the context menu. Edit the database connection information to work with your database configuration. Test the changes and then click **OK** to close the Connection Properties dialog box.

As another preliminary task, start the Oracle WebLogic Server instance integrated with Oracle JDeveloper. To start the Oracle WebLogic Server, select **Run** -> **Start Server Instance (IntegratedWebLogicServer)**.

If this is the first time you've run the integrated Oracle WebLogic Server, a Create Default Domain dialog box will open. Create a password for the default Oracle WebLogic Server domain, and select an address from those listed for Listen Address. For example, choose **localhost** rather than leaving the address empty.

Click **OK** to save the change and create and configure the default domain for Oracle ADF.

SAMPLE APPLICATION OVERVIEW

With the OraMagSample workspace open, expand the **ViewController** node to see a project view similar to what is shown in Figure 3.

Figure 4 shows the architecture of the sample application, with the task flow command pattern template library parts colored in green.

The sample application contains a single JSF document, ParentView.jspx, that is based on the event-producer-template.jspx page template defined in the TaskFlowCommandStrategyLib.jar library.

In Oracle ADF, views that are based on a template not only adopt the template layout definition but also maintain a reference to the template Oracle ADF binding file (if the template has its own Oracle ADF binding).

The event-producer-template.jspx file, upon which the ParentView.jspx document is based, has a PageDef file defined and configured as a contextual event producer. Because the ParentView.jspx document maintains a reference to the template binding, it can broadcast contextual events as indicated in Figure 4.

The event-producer-template.jspx file does not contain layout definitions; it only helps to automate the registration of views as a contextual event producer. Luckily, Oracle JDeveloper 11*g* Release 2 lets application developers nest page templates. As a result, an Oracle ADF Faces view built based on a functional template can contain additional page templates to improve look and feel.

In the OraMagSample application, the ParentView.jspx page uses a nested template reference to the default Oracle three-column template.

Also shown in Figure 4, the ParentView.jspx template uses two managed beans in viewScope that are configured in the adfc-config.xml file located in the Web Context ->public_html->WEB-INF folder.

The SupportedToolbarActions bean is passed as an input parameter value to the bounded task flow, also shown in Figure 4. The bounded task flow uses this bean reference to call back into the parent view. In the sample application, the callback is used to change the button bar disabled state depending on the view currently displayed in the bounded task flow.

Let's have a detailed look at each of the two beans:

- **ParentViewBean** contains action methods referenced from the **ActionListener** property of the toolbar buttons. Each time a user clicks a button, a referenced method is invoked on ParentViewBean.

  ParentViewBean extends the ActionProducer class of the task flow command pattern template library and can raise contextual events via calls to invokeTaskFlowAction. Information passed with the contextual event include the method name to invoke on the current display view in the bounded task flow, a discriminator for the event receiver to decide whether or not to respond to the request, and an optional payload object to pass extra information if needed.
- **SupportedToolbarActions** allows the parent view to show a button as disabled if a method is not supported by the current view in the task flow.

To enable this functionality, the Oracle ADF Region tag (af:region) in the ParentView.jspx page has its RegionNavigationListener property configured to reference a method in ParentViewBean.

The region navigation listener receives a notification each time the view displayed in the bounded task flow changes. This

notification is used to refresh the button bar by disabling a button if its action is not supported by the new view.

The command receiver in the sample application consists of a bounded task flow that is created from the library command-receiver-flow-template.xml, as well as views that are built based on the event-receiver-template.jspx page template in the TaskFlowCommandStrategyLib.jar.

Similar to event-producer-template.jspx for the parent view, event-receiver-template.jspx automatically configures the views in the bounded task flow for contextual events. This time, however, the views are configured as contextual event receivers, as shown in Figure 4.

The departments-employees-flow.xml bounded task flow definition located in the Web Content -> WEB-INF folder defines two managed beans used as view action handlers to receive contextual event notification, as well as two backing beans used to access and launch a dialog box contained in the view.

Let's look at two representative beans in more depth:

- **tfcpDepartmentsView** is the view action handler (as shown in Figure 1) for the DepartmentsView activity in the bounded task flow.

  Following the naming convention used by the task flow command pattern template, the managed bean filename has a prefix of *tfcp*. When the contextual event handler receives a contextual event, it identifies the view action handler on which to invoke the invokeAction method using the name of the current viewId and the *tfcp* prefix.

  The invokeAction method in a view action handler bean interprets the contextual event payload and calls the method that handles the request. View action handler beans extend DepEmpBaseCommands, an abstract class that contains method signatures for the global methods supported by the application. In addition, the DepEmpBaseCommands class provides access to the current displayed view binding container. With contextual events, developers cannot use the currentBindingsEntry method on the BindingContext object to gain this access.

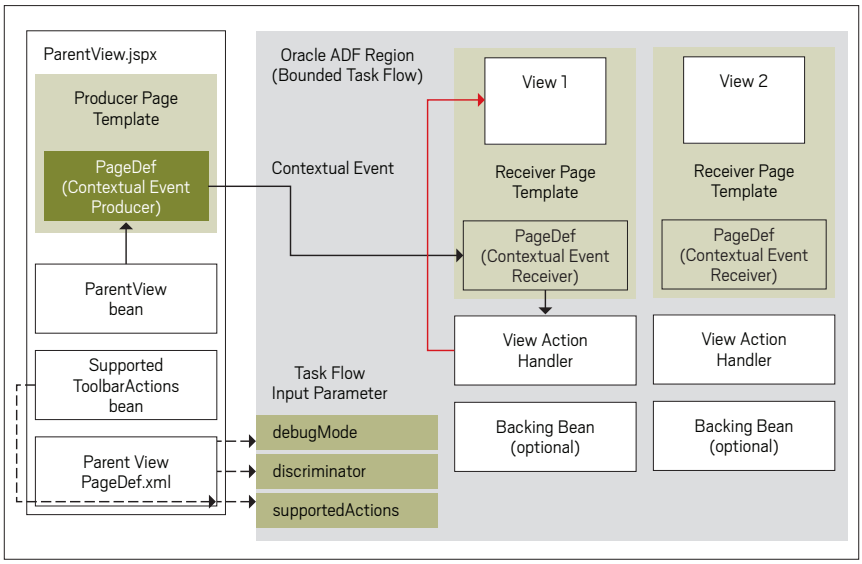

Figure 4: Sample application architecture

- **DepartmentsBacking.** This is a backing bean referenced from the showHelp method in tfcpDepartmentsView to launch a help dialog box.

  When dragging a bounded task flow built on the command-receiver-flow-template.xml task flow template to a page, three input parameter values must be provided.

- **debugMode.** When set to true, this parameter prints debug messages at runtime.
- **discriminator.** Contextual events are broadcast to all listening binding containers. To use the library with multiple regions on a page, a discriminator key can be configured as a task flow input. A receiver will ignore contextual events if the discriminator string defined on the region configuration does not match the discriminator string passed with the producer event.
- **supportedActions.** This parameter points to the managed bean in the parent view that contains the list of commands supported by the views in the bounded task flow.

To see this configuration defined for the ParentView.jspx file in the sample application, expand the **Application Sources** node in the **ViewController** project and navigate to the **ParentViewPageDef.xml** file located in the **oramag.novdec.twelve.sample.view** pageDefs package. To open the page definition file, right-click the **ParentViewPageDef.xml** file and select **Open** from the context menu. In the binding editor, click the **Source** tab at the bottom to switch to the source view.

RUNNING THE SAMPLE

To run the sample, right-click the **ParentView.jspx** file in the Application Navigator and select **Run** from the context menu.

The application opens in a browser showing a global button bar and an edit form for the HR Departments table. The button bar resides in the ParentView.jspx file, while the edit form is defined in the DepartmentsView.jsff page fragment in the bounded task flow (Oracle ADF Region). Note that the departments edit form does not implement query functionality. As a result, the **Execute Query** button is disabled, as shown in Figure 5.

Clicking any of the buttons in the button bar issues a contextual event that is broadcast to the view in the bounded task flow. The contextual event handler checks the discriminator value passed with the event request. If the value matches the value configured in the task flow, it invokes the command on the view action handler.

To navigate to the employees form, click the **Navigate to Employees** button. Upon task flow navigation, the region navigation listener on the af:region tag is invoked for the parent view to refresh the button bar with the capabilities of the employee view. As a result, the **Execute Query** button

Figure 5: Runtime view of sample application

becomes enabled. As you would expect from a command pattern implementation, the same set of buttons now operate on the employees view.

SUMMARY

This article outlines a solution for a frequently asked-about use case in Oracle ADF Region interaction by providing an implementation that can be reused in many projects.

The task flow command pattern template library contains page templates, a task flow template, and a public API that simplifies the configuration of this solution in custom projects. The contextual event implementation detail in the template is hidden from the application developer.

In addition, this article has shown how to adopt the command design pattern from object-oriented programming for use within an Oracle ADF development environment. ◀

Frank Nimphius is a senior principal product manager for Oracle JDeveloper and Oracle ADF. 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

LEARN more about Oracle ADF
oracle.com/adf

DOWNLOAD
Oracle JDeveloper 11*g*
bit.ly/LRlg7k

the sample application for this article
bit.ly/OHuJLe

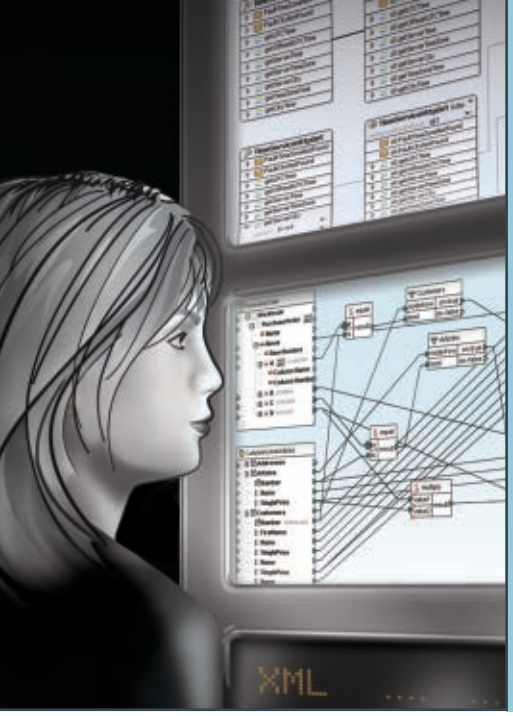

Connect legacy technologies affordably with the complete set of data integration tools from Altova®

ALTOVA® missionkit®

Experience how the Altova MissionKit®, the integrated suite of XML, data mapping, and database tools, can help you leverage existing technology and business software investments while integrating modern technologies – without breaking your budget.

NEW in Version 2012:

- Streaming reading/writing of files for large-scale ETL applications
- Sorting of data mapping results
- New data processing functions
- JDBC drivers (in addition to ADO and ODBC)
- New, native Java APIs

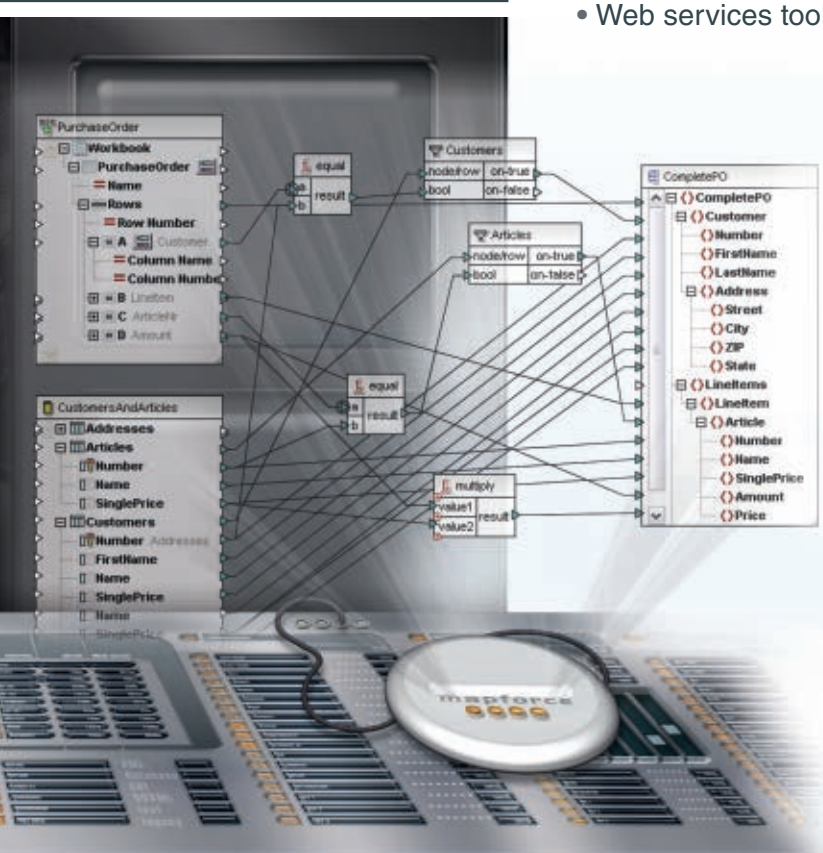

The Altova MissionKit includes multiple intelligent tools for data integration:

MapForce® – Graphical data mapping, transformation, & conversion tool

- Drag-and-drop data conversion with instant transformation & code gen
- Support for mapping XML, DBs, EDI, Excel®, XBRL, flat files & Web services

XMLSpy® – XML editor and Web services tool

- XML editor with strong database integration
- Web services tool, JSON <> XML converter

DatabaseSpy® – multi-database query, design & comparison tool

- Support for all major relational databases and translation between DB types
- SQL editor, graphical database design & content editor

Download a 30 day free trial!

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

Scan to learn more about data integration with the MissionKit.

ORACLE DATABASE

The Data Dictionary: Make Views Work for You

Part 10 in a series of articles on understanding and using PL/SQL

If you're reading this article, there's a really good chance that you write PL/SQL code. Lots of it. Which means that you also will at least occasionally need to analyze that code, answering questions such as

- On which database objects does my program depend?
- Which of my packages contain calls to a subprogram in another package or a reference to a global variable?
- Do any of my subprograms contain parameters whose datatypes should not be used?
- Are all of my subprograms compiled with a sufficiently high level of optimization?

You can, of course, always use the search feature of your editor or integrated development environment to look through multiple database objects and files to find specific chunks of text. But that won't be enough to answer all of the above questions and many more you will encounter.

Don't despair! One of the most lovely aspects of writing PL/SQL code and compiling that code into the database is that Oracle Database offers a set of views—known collectively as the *data dictionary*—that enable you to use the SQL *and* PL/SQL languages to get answers to just about any question you have about your code. Table 1 offers a high-level overview of the data dictionary views most often used to manage PL/SQL code.

This article explores many of the views in the table, describing the most useful columns in the views and offering examples of how you can put those views to use.

DATA DICTIONARY FUNDAMENTALS

The data dictionary consists of numerous tables and views created by the database instance. User schemas generally have no privileges on these tables; Oracle Database grants only SELECT access on the views.

Most data dictionary views come in three versions:

1. The USER view: information about database objects *owned* by the schema to which you are connected
2. The ALL view: information about database objects to which the currently connected schema has access
3. The DBA view: unrestricted information about all the database objects in a database instance (non-DBA schemas usually have no authority to query DBA views)

Let's look at an example. Suppose I want to obtain a list of the objects—tables, views, packages, and so on—defined in the database.

The following query returns all the objects defined in my schema:

```
SELECT * FROM user_objects
```

This query returns all the objects that are defined in my schema or for which I have been granted the privilege to use those objects in some way:

```
SELECT * FROM all_objects
```

Finally, the following query returns a list of all the objects defined in the database instance—if I have the authority to select from the view:

USER_ARGUMENTS	The arguments (parameters) in all the procedures and functions in your schema.
USER_DEPENDENCIES	The dependencies to and from objects you own. This view is used mostly by Oracle Database to invalidate the status of database objects when an object on which they depend changes.
USER_ERRORS	The current set of compilation errors for all stored objects (including triggers) you own. This view is accessed by the SHOW ERRORS SQL*Plus command. You can, however, write your own queries against it as well.
USER_IDENTIFIERS	Introduced in Oracle Database 11*g* and populated by the PL/Scope compiler utility. Once populated, this view provides you with information about all the identifiers—program names, variables, and so on—in your code base.
USER_OBJECT_SIZE	The size of the objects you own. Actually, this view shows you the source, parsed, and compile sizes for your code. Although it is used mainly by the compiler and runtime engine, you can use it to identify the large programs in your environment.
USER_OBJECTS	The objects you own. You can, for instance, use this view to see if an object is marked INVALID, find all the packages that have EMP in their names, and so on.
USER_PLSQL_OBJECT_SETTINGS	Information about the characteristics—such as the optimization level and debug settings—of a PL/SQL object that can be modified through the ALTER and SET DDL commands.
USER_PROCEDURES	Information about stored programs, such as the AUTHID setting, whether the program was defined as DETERMINISTIC, and so on.
USER_SOURCE	The text source code for all objects you own (in Oracle9*i* Database and above, including database triggers and Java source). This is a very handy view, because you can run all sorts of analyses of the source code against it with SQL and, in particular, Oracle Text.
USER_STORED_SETTINGS	PL/SQL compiler flags. Use this view to discover which programs have been compiled via native compilation.
USER_TRIGGERS and USER_TRIGGER_COLS	The database triggers you own (including the source code and a description of the triggering event) and any columns identified with the triggers, respectively. You can write programs against USER_TRIGGERS to enable or disable triggers for a particular table.

Table 1: Useful views for PL/SQL programmers

```
SELECT * FROM dba_objects
```

Usually the only difference between the USER view and the ALL view is that the latter contains one extra column, OWNER, that shows which schema owns the object.

The remainder of this article provides examples based on the USER view.

DISPLAY INFORMATION ABOUT STORED OBJECTS

The USER_OBJECTS view contains a row for every database object owned by your schema. The most commonly used columns are

- OBJECT_NAME: Name of the object
- OBJECT_TYPE: Type of the object, such as PACKAGE, FUNCTION, or TRIGGER
- STATUS: Status of the object—VALID or INVALID
- LAST_DDL_TIME: Time stamp indicating the last time this object was changed

Here are some examples of queries against USER_OBJECTS.

- Show the names of all tables in my schema:

```
SELECT object_name
  FROM user_objects
 WHERE object_type = 'TABLE'
 ORDER BY object_name
```

Show the names of all objects whose status is invalid:

```
SELECT object_type, object_name
  FROM user_objects
 WHERE status = 'INVALID'
 ORDER BY object_type, object_name
```

The status of a program unit (PL/SQL package, procedure, or function) is set to INVALID if a database object on which it depends is changed. That program unit must then be recompiled (which Oracle Database will often do automatically the next time you try to use that program unit).

- Show all objects that have been changed today:

```
SELECT object_type, object_name,
       last_ddl_time
  FROM user_objects
 WHERE last_ddl_time >= TRUNC (SYSDATE)
 ORDER BY object_type, object_name
```

DISPLAY AND SEARCH SOURCE CODE

All the program unit source code you've compiled into the database is accessible through the USER_SOURCE view, whose columns are

- NAME: Name of the object
- TYPE: Type of the object (ranging from PL/SQL program units to Java source and trigger source)
- LINE: Number of the line of the source code
- TEXT: Text of the source code

You can write queries against USER_SOURCE to

- Find all the program units that call a particular subprogram of a package
- Verify that coding standards are being followed
- Find all occurrences of a literal value that needs to be changed

Here is an example: I need to change the parameter list and code of a procedure named CALC_TOTALS in the SALES_MGR package. I'd like to find out where this procedure is called, outside of the SALES_MGR package itself.

```
SELECT name, line, text
  FROM user_source
 WHERE UPPER (text)
  LIKE '%SALES_MGR.CALC_TOTALS%'
 ORDER BY name, line
```

Of course, this query will also find comments that contain this string, and there *could* be invocations of CALC_TOTALS that are *not* found, such as

```
SALES_MGR.
CALC_TOTALS
```

Assuming, however, that you don't write or format your code to break up subprogram calls like that, the query will do a pretty good job of identifying the places in your code you need to review.

And for an Oracle Database 11*g* instance, you could use the PL/Scope feature. See the "A Better USER_SOURCE" sidebar in the online version of this article at bit.ly/SamK9l for more information.

COMPILER SETTINGS OF STORED CODE

The USER_PLSQL_OBJECT_SETTINGS view provides information about compiler settings of stored PL/SQL objects. Key columns are

- PLSQL_OPTIMIZE_LEVEL: Optimization level that was used to compile the object
- PLSQL_CODE_TYPE: Compilation mode for the object
- PLSQL_DEBUG: Whether or not the object was compiled for debugging
- PLSQL_WARNINGS: Compiler warning settings that were used to compile the object
- NLS_LENGTH_SEMANTICS: NLS length semantics that were used to compile the object

Here are some examples of queries against USER_PLSQL_OBJECT_SETTINGS.

- Find all the program units that are not taking sufficient advantage of compile time optimization in Oracle Database:

```
SELECT name
  FROM user_plsql_object_settings
 WHERE plsql_optimize_level < 2
```

An optimization level of 0 means no optimization at all. An optimization level of 1 means a minimal amount of optimization. Neither of these levels should be seen in a production environment.

- Identify all programs for which compile time warnings (which provide feedback on the quality of your code) are disabled:

```
SELECT name, plsql_warnings
  FROM user_plsql_object_settings
 WHERE plsql_warnings LIKE '%DISABLE%';
```

DETAILED INFORMATION ABOUT PROCEDURES AND FUNCTIONS

The USER_PROCEDURES view provides information about all functions and procedures, both schema-level and those defined within packages, in your schema. Columns of this view are

- AUTHID: Shows whether a procedure or a function is defined as an invoker rights (CURRENT_USER) or definer rights (DEFINER) program unit
- DETERMINISTIC: Set to YES if the function is defined to be deterministic, which theoretically means that the value returned by the function is determined completely by the function's argument values
- PIPELINED: Set to YES if the function is

defined as a pipelined function, which means that it can be executed in parallel as part of a parallel query
- OVERLOAD: Set to a positive number if this subprogram is overloaded, which means that there are at least two subprograms with this name in the same package

Here are some examples of queries against USER_PROCEDURES.
- Find all the procedures and functions that will run under invoker rights (the privileges of the invoker of the program are used at runtime to resolve references to database objects such as tables):

```
SELECT   object_name
       , procedure_name
    FROM user_procedures
   WHERE authid = 'CURRENT_USER'
ORDER BY object_name, procedure_name
```

- Show all the functions declared to be deterministic:

```
SELECT   object_name
       , procedure_name
    FROM user_procedures
   WHERE deterministic = 'YES'
ORDER BY object_name, procedure_name
```

ANALYZE AND MODIFY TRIGGER STATE

If you work with database triggers, USER_TRIGGERS, which contains a row for each trigger defined in your schema, will come in handy. Key columns are
- TRIGGER_NAME: The name of the trigger
- TRIGGER_TYPE: A string that shows if this is a BEFORE or AFTER trigger and whether it is a row- or statement-level trigger (in a trigger that is fired before an INSERT statement, for example, the value of this column is BEFORE STATEMENT)
- TRIGGERING_EVENT: The type of SQL operation—such as INSERT, INSERT OR UPDATE, DELETE OR UPDATE—that will cause the trigger to fire
- TABLE_NAME: The name of the table on which the trigger is defined
- STATUS: The status of the trigger—ENABLED or DISABLED
- WHEN_CLAUSE: An optional clause you can use to avoid unnecessary execution of the trigger body
- TRIGGER_BODY: The code executed when the trigger fires

Here are some examples of queries against USER_TRIGGERS.
- Find all disabled triggers:

```
SELECT *
  FROM user_triggers
 WHERE status = 'DISABLED'
```

- Find all row-level triggers defined on the EMPLOYEES table:

```
SELECT *
  FROM user_triggers
 WHERE table_name = 'EMPLOYEES'
   AND trigger_type LIKE '%EACH ROW'
```

- Find all triggers that fire when an UPDATE operation is performed:

```
SELECT *
  FROM user_triggers
 WHERE triggering_event LIKE '%UPDATE%'
```

One limitation in the USER_TRIGGERS view is that the TRIGGER_BODY column type is LONG, which means that it cannot be used in a SQL comparison.

Suppose, for example, that I want to find all the triggers whose trigger body contains the string "emp". The following query, unfortunately, fails and produces an ORA-00932 error:

```
SELECT *
  FROM user_triggers
 WHERE trigger_body LIKE '%emp%'
```

So if you do want to search the contents of trigger bodies, you will need to use PL/SQL, in a block like this:

```
BEGIN
  FOR rec IN (SELECT *
                FROM user_triggers)
  LOOP
    IF rec.trigger_body LIKE '%emp%'
    THEN
      DBMS_OUTPUT.put_line (
        'Found in ' || rec.trigger_
name);
    END IF;
  END LOOP;
END;
```

Note that the USER_TRIGGER_COLS view keeps track of the columns that are referenced inside a trigger body.

OBJECT DEPENDENCY ANALYSIS

The USER_DEPENDENCIES view describes the dependencies between the procedures, packages, functions, package bodies, and triggers accessible to the current user. You can use it to perform impact analysis on your code, as in: How many programs will need to be changed if I change this table?

Key columns in this view are
- NAME: Name of the object
- TYPE: Type of the object
- REFERENCED_OWNER: Owner of the referenced object
- REFERENCED_NAME: Name of the referenced object
- REFERENCED_TYPE: Type of the referenced object

Here are some examples of queries against USER_DEPENDENCIES.
- Find all the objects that depend on (reference) the EMPLOYEES table:

```
SELECT type, name
   FROM user_dependencies
  WHERE  referenced_name = 'EMPLOYEES'
ORDER BY type, name
```

- Find all the objects in the current schema on which the ORDER_MGR package depends:

```
SELECT referenced_type
     , referenced_name
  FROM user_dependencies
 WHERE name = 'ORDER_MGR'
   AND referenced_owner = USER
ORDER BY referenced_type,
         referenced_name
```

A best practice that I, and others, strongly recommend is to avoid repeating SQL statements by "hiding" those statements inside a procedure or a function. Let's look at an example and then at how the USER_DEPENDENCIES view can help us identify *violations* of this best practice.

Code Listing 1: Find functions that have an OUT or an IN OUT argument

```
 1  SELECT ua.object_name,
 2         ua.package_name,
 3         ua.argument_name,
 4         ua.in_out
 5    FROM (SELECT *
 6            FROM user_arguments
 7           WHERE position = 0) funcs,
 8         user_arguments ua
 9   WHERE     ua.in_out IN ('OUT', 'IN OUT')
10         AND ua.position > 0
11         AND ua.data_level = 0
12         AND funcs.object_name = ua.object_name
13         AND funcs.package_name = ua.package_name
14         AND (   funcs.overload = ua.overload
15              OR (funcs.overload IS NULL
16                  AND ua.overload IS NULL))
```

Lines	Description
5–7	I use an inline view in the FROM clause to identify all those rows in USER_ARGUMENTS that are RETURN clauses (and therefore identify functions).
9–11	I look for OUT or IN OUT arguments that are not in RETURN clauses and are not "nested" information, such as fields of a record argument.
12–16	I use this rather lengthy join condition between the inline view (abbreviated as "funcs") and USER_ARGUMENTS. The object names and package names must match, and the overload value must be the same or both must be NULL. The overload column is not NULL if the package has two or more subprograms with the same name.

It is very common in PL/SQL code to find many queries that retrieve a single row for a primary key. Here's a PL/SQL example with a query that uses the standard Oracle Database EMPLOYEES table:

```
PROCEDURE process_employee (
   employee_id_in IN INTEGER)
IS
   l_name   VARCHAR2 (100);
BEGIN
   SELECT last_name
     INTO l_name
     FROM employees
    WHERE employee_id = employee_id_in;
END;
```

Instead of writing this query each time, I suggest writing a function *once* that contains this query and returns the desired value. Then you can call the function as needed. Assuming that I have created a package named EMPLOYEES_API with a function named LAST_NAME, the above procedure can be changed to

```
PROCEDURE process_employee (
   employee_id_in IN INTEGER)
IS
   l_name   VARCHAR2 (100);
BEGIN
   l_name :=
      employees_api.
        last_name (employee_id_in);
END;
```

Now if I ever need to change the query for any reason (such as to take advantage of Oracle Database 11*g*'s function result cache feature), I'll be able to make the change in *one place*, rather than having to find all occurrences of the query in my application code.

So suppose my development team has added this best practice to its coding standards: the only PL/SQL program units that should contain SQL statements are packages that end with the suffix _API.

I can then write a query against USER_DEPENDENCIES that identifies all program units that violate this rule:

```
SELECT name,
       TYPE,
       referenced_owner,
       referenced_name
  FROM user_dependencies
 WHERE     TYPE IN
              ('PACKAGE',
               'PACKAGE BODY',
               'PROCEDURE',
               'FUNCTION',
               'TRIGGER',
               'TYPE')
   AND referenced_type = 'TABLE'
   AND name NOT LIKE '%\_API' ESCAPE '\'
ORDER BY name
       , referenced_owner
       , referenced_name
```

The online version of this article includes information on the USER_ARGUMENTS view and its columns as well as this issue's PL/SQL Challenge question.

IT'S A GOLD MINE IN THERE

This article merely scratches the surface of the application information that can be mined from the data dictionary views in Oracle Database. PL/SQL editors such as Oracle SQL Developer provide user interfaces to many of these views, making it easier to browse their contents. ◀

Answer to Last Issue's Challenge

The PL/SQL Challenge question in last issue's "Bulk Processing with BULK COLLECT and FORALL" article asked, "Which of these blocks will uppercase the last names of all employees in the table?" Choices b, c, and d are correct.

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 11*g*
bit.ly/fherki

TEST your PL/SQL knowledge
plsqlchallenge.com

READ
PL/SQL 101, Parts 1–9
bit.ly/omagplsql

online-only article content
bit.ly/SamK9l

READ more about the data dictionary
bit.ly/OR14hA

ORACLE DATABASE 11*g*

Beginning Performance Tuning: Diagnose the Past

Take the next step in diagnosing Oracle Database performance issues.

In "Beginning Performance Tuning," in the July/August 2012 issue of *Oracle Magazine*, I demonstrated the basics of diagnosing performance issues in Oracle Database. In this article, I will show how to take performance tuning to the next level—identifying patterns and measuring resource utilization. As was the case with the previous article, you can create this article's examples by executing the setup script—setup.sql—in the online version of this article, at bit.ly/U8ml8P.

HISTORICAL WAITS

Recall from "Beginning Performance Tuning" that you can identify why a session is slow by looking at what event it is waiting for right now in the V$SESSION view. This technique is effective when you are able to peek into the view while the issue is occurring, but what if a user reports that there *was* an increase in the response time of a session but it is back to normal *now*? Because the issue is no longer present, the session may appear to be working productively with the STATE column value of WAITED KNOWN TIME. What's worse, the current wait event shown in the V$SESSION view may be something unrelated to what the user encountered earlier. To diagnose this issue, you need to know *all* the different events the session had to wait for—not just what it is waiting for *now*. How difficult is it to get that information?

Fortunately, it's quite trivial. There is another view—V$SESSION_EVENT—that shows the details of *all* the wait events the session has ever waited for. Here are the important columns in that view:

- **SID:** the session ID
- **EVENT:** the name of the wait event experienced by the session
- **TOTAL_WAITS:** the total number of times the session had to wait for that wait event
- **TIME_WAITED:** the total length of time the session had to wait for that event, recorded in cs (centiseconds [hundredths of a second])
- **AVERAGE_WAIT:** the average length of time (also in cs) the session waited for that event
- **MAX_WAIT:** the maximum length of time (in cs) the session waited for that event

Let's examine the usage of the view with an example. Make sure you have set up the example as described in the online version of this article. Connect to the Oracle Database instance as ARUP, and execute the test1.sql script, available in the online version. (Note that the script will take some time to complete). While that script is running, connect as SYS from a different session and find the SID of the ARUP session as follows:

```
select sid from v$session
where username = 'ARUP';

 SID
————
  37
```

Now display the session events of SID 37 by executing the script in Listing 1. Because the V$SESSION_EVENT columns report time in centiseconds, I have multiplied them by 10 to show them in milliseconds (ms), a unit of measure that is a bit more familiar. Study the output carefully—it shows the various events the session waited for earlier,

Code Listing 1: History of wait events in a specific session

```
set lines 120 trimspool on
col event head "Waited for" format a30
col total_waits head "Total|Waits" format 999,999
col tw_ms head "Waited|for (ms)" format 999,999.99
col aw_ms head "Average|Wait (ms)" format 999,999.99
col mw_ms head "Max|Wait (ms)" format 999,999.99
select event, total_waits, time_waited*10 tw_ms,
       average_wait*10 aw_ms, max_wait*10 mw_ms
from v$session_event
where sid = 37
/
                                  Total      Waited     Average         Max
Waited for                        Waits    for (ms)   Wait (ms)   Wait (ms)
——————————————————————————— ————————— ———————————— ———————————— ————————————
Disk file operations I/O              8          .00         .10         .00
KSV master wait                       2       350.00      173.20      340.00
os thread startup                     1        20.00       19.30       20.00
db file sequential read               5       160.00       32.10       70.00
direct path read                  1,521    51,010.00       33.50      120.00
direct path read temp           463,035   513,810.00        1.10      120.00
direct path write temp               20       370.00       18.70       50.00
resmgr:cpu quantum                   21       520.00       24.60      110.00
utl_file I/O                          8          .00          .00         .00
SQL*Net message to client            20          .00          .00         .00
SQL*Net message from client          20     9,620.00       481.20    9,619.00
kfk: async disk IO               904,818     3,050.00          .00         .00
events in waitclass Other             35        20.00          .70       20.00
```

I-HUA CHEN

not what it is experiencing now. You can see that the session has waited for one event in particular: "kfk: async disk IO." The session has waited 904,818 times for this event, but it has waited for a total of only 3,050 ms. The average wait shows zero, but it is simply because the average wait was so short that it couldn't be shown within the two digits after the decimal point. So although the total number of waits for this event was high, it added only 3,050 ms to the overall time of the session—an insignificant number. You can thus rule this wait event out as a cause of delay for the session.

In this case, you should focus your attention on the event that caused the session to wait for the maximum time. From the output in Listing 1, you can see that session 37 (SID = 37) waited 513,810 ms, or more than 8.5 minutes, for the "direct path read temp" event. Every time the session waited for this event, it waited 1.1 ms on average, so if you can reduce the time for this event, you can reduce the overall time for the session. Looking at the session event history enables you to identify the biggest contributors to the delay in the session, whether they are currently affecting the session or have already affected it.

Did you notice the column for the maximum time waited: "Max Wait (ms)"? Why is that information useful? You see, the average wait time does not tell the whole story. Consider the "SQL*Net message from client" event in the output in Listing 1. The session waited 20 times with an average wait time of 481 ms for that event. Does that mean that the session waited for approximately 481 ms at each of the 20 occurrences or that the session waited a very short time for most of the event instances and a very long time for one event? The latter will skew the average to a high value but will indicate an isolated issue rather than a persistent one—the two possibilities lead to very different conclusions.

The "Max Wait (ms)" column shows the maximum time the session had to wait for one occurrence of this event. The value here is 9,619 ms, and because the total wait time was 9,620 ms, it appears that the session waited 9,619 ms on one occasion, leaving a 1 ms combined total for the other 19 occasions—a very small wait each time.

Code Listing 2: Output of the top command

```
top - 10:56:49 up 18 days, 18:48,  4 users,  load average: 1.02, 0.92, 0.48
Tasks: 180 total,   2 running, 178 sleeping,   0 stopped,   0 zombie
Cpu(s): 49.8%us,  0.5%sy,  0.0%ni, 49.2%id,  0.5%wa,  0.0%hi,  0.0%si,  0.0%st
Mem:   1815256k total,  1771772k used,    43484k free,    66120k buffers
Swap:  2031608k total,   734380k used,  1297228k free,   747740k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
 5946 oracle    25   0  706m 177m 159m R  100 10.0   9:20.26 oracle
 6104 oracle    15   0 2324 1060  800 R    1  0.1   0:00.12 top
31446 oracle    15   0  688m 135m 129m S    0  7.7   0:08.24 oracle
... output truncated ...
```

Considering the single incidence of a large wait, this event should not be a general cause of concern. On the other hand, had you seen a maximum time close to the average time, you could surmise that all occurrences had to wait about the same amount of time. In such a case, reducing the time for this event would likely apply uniformly to all occurrences and consequently reduce the overall elapsed time.

Although the V$SESSION_EVENT view shows *what* the session waited for earlier, it doesn't show *when*. That information is visible in another view—V$ACTIVE_SESSION_HISTORY (part of the extra-cost Oracle Diagnostics Pack)—which is beyond the scope of this article.

STATISTICS

Although wait events are great for helping with understanding the speed bumps the sessions experience, they do not show another important attribute of sessions: the use of resources such as CPU, I/O, and memory. A resource-hogging session deprives other sessions of the same resources, thus causing performance issues. When the root of the problem is that the session is consuming too much CPU, you should look at resource consumption—not the events waited for—by a session. Fortunately, finding that information is quite simple: it's visible in a view named V$SESSTAT, which has three columns:

- **SID:** the session ID
- **STATISTIC#:** the ID of the statistic being captured (You can get the name of the statistic from the NAME column in another view—V$STATNAME—by using this ID.)
- **VALUE:** the value of the statistic

Let's see how to use this information with a performance problem example.

CPU SPIKE

Suppose you've heard from several users that the performance is terrible across the board. Further, the UNIX system administrator has reported that both CPU and memory consumption are very high on the server and most of the consumption is by Oracle Database–owned processes. Sound familiar? From my experience, it's the second-most-common performance issue in Oracle Database systems. You can reproduce this problem by running the test1.sql script, available in this article's example scripts, as user ARUP.

To diagnose this resource issue, review the top resource-consuming processes by issuing the `top` command at the UNIX command prompt. Listing 2 shows the output of the `top` command.

In the output in Listing 2, you can see that the process with ID 5946 consumes the most CPU (100 percent) and memory (10 percent) and therefore should be the focus of your attention. To find out more about the process, enter the following command at the UNIX prompt:

```
$ ps -aef|grep 5946

oracle    5946  5945 63 10:59 ?
00:01:52 oracleD112D2
(DESCRIPTION=(LOCAL=YES)
(ADDRESS=(PROTOCOL=beq)))
```

The output shows the entire description of the process, which is clearly an Oracle "server process"—a process that is created by Oracle Database when a session is established—

and that the process has been running for 1 minute and 52 seconds. The next question, then, is which Oracle Database session this process was created for. For that, you should look into another view—V$PROCESS—where the SPID column shows the server process ID. However, this view does not show the session information, so you need to join this view with the familiar V$SESSION view, as follows:

```
select sid
from v$session s, v$process p
where p.spid = 5946
and s.paddr = p.addr;

SID
———
37
```

Once you know the SID, you can get everything you need to know about the session—the user who established the session, the machine it came from, the operating system user, the SQL it is executing, and so on—from the V$SESSION view. This script for getting the information was described in Listing 5 in "Beginning Performance Tuning," in the July/August 2012 issue of *Oracle Magazine*. To find the SQL being run by session 37, use this query:

```
select sql_fulltext
from v$sql l, v$session s
where s.sid = 37
and l.sql_id = s.sql_id;
```

Here is the output:

```
select max(test1.owner)
from test1, test2, test2, test2,
     test2, test2, test2, test2,
     test2, test2, test2, test2,
     test2, test2, test2, test2,
     test2, test2, test2, test2,
     test2, test2, test2, test2;
```

This SQL is performing multiple Cartesian joins, so it's no wonder it's consuming so much CPU and memory.

Now that you have found the culprit, should you just go ahead and kill the session to release the CPU consumption? Not so fast. First, you need to know if the CPU consumption was recent or if the session has been chewing it up since the beginning. This is where the V$SESSTAT view comes in very handy—it shows the resource consumption (CPU in this case) by a specific session. To find out the CPU used by session 37, you would use the following query:

```
select s.value
from v$sesstat s, v$statname n
where s.sid = 37
and n.statistic# = s.statistic#
and n.name = 'CPU used by this session';

VALUE
———
47379
```

The output shows the number of CPU "ticks" that have been consumed by this session since it started. Considering that this session has been running for about two minutes, the CPU consumption is pretty high, so it is likely that this session has been consuming CPU all the time. Again, checking the session's other details, such as the SQL it is executing, makes it fairly easy to understand why this is the case: the session is performing a multiple-table Cartesian join, which is bound to consume a lot of CPU. At this point, you may decide to kill the session to stop the high CPU consumption, or you may decide to let it run for now and fix the SQL later to avoid the Cartesian product.

Code Listing 3: All session statistics

```
select name, value
from v$sesstat s, v$statname n
where sid = 37
and n.statistic# = s.statistic#
order by value desc
/

NAME                                              VALUE
————————————————————————————————————————— ——————————
table scan rows gotten                       1.0236E+10
session logical reads                          25898547
consistent gets                                25898543
table scan blocks gotten                       25325165
session pga memory max                         21250020
session pga memory                             21250020
session uga memory max                         20156552
session uga memory                             20156552
bytes sent via SQL*Net to client                 878760
recursive calls                                  576848
opened cursors cumulative                        143367
parse count (total)                              143292
parse count (hard)                               143118
table scans (short tables)                       143086
sql area evicted                                 141996
DB time                                           70007
CPU used by this session                          69724
... output truncated ...
```

ALL STATISTICS

Let's revisit the current problem by checking the CPU consumption once again with this:

```
select s.valuez
from v$sesstat s, v$statname n
where s.sid = 37
and n.statistic# = s.statistic#
and n.name = 'CPU used by this session';

VALUE
———
69724
```

Now the result—the CPU used—is 69,724. Note that this number is larger than the number the last time I checked CPU usage (47,379). This is because the statistic value increases over time. The first time I checked CPU usage, I surmised that the multitable Cartesian product was to blame for that

CPU consumption—but can I prove it? The answer is a resounding "Yes, I can, using other statistics."

Listing 3 shows the query for gathering all the statistics for session 37. In the output, note the "table scan rows gotten" statistic value: 1.0236E+10—about 10 billion rows! This is indeed a very high number of rows to be accessed by one session in two minutes. The value for the "consistent gets" statistic is 25,898,543—about 25.9 million blocks read from the buffer cache. The high number of buffer gets takes up a considerable amount of CPU.

Another cause of CPU consumption is the parsing of SQL statements. In the output in Listing 3, note the "parse count (total)" statistic, a very high number at 143,292. In plain English, it means that the session had to parse—not just execute—SQL statements that many times in about two minutes, which is quite unusual. Examine the SQL statement executed by the session (shown in the script test2.sql script in the example setup in the online version of this article). You can see that it creates distinct literal SQL statements. Each literal SQL statement needs to be parsed, which is something you have confirmed by looking at the parse count. Therefore, you surmise from the Listing 3 output that there are two causes of high CPU usage for this session: a high number of buffer gets and a high number of parses.

In Listing 3, also note two other sets of statistics: "session pga memory max" and "session uga memory max," which indicate the total memory consumed by the session. The very high numbers explain the high memory consumption by the Oracle server process that you noticed in the output of the operating system `top` command earlier. If you want to reduce the CPU and memory consumption of the server, you need to make sure the session consumes fewer of these resources, by appropriately modifying the SQL statement it issues.

REDO SPIKE

Occasionally you may have a performance issue that will not appear as clearly at the OS level as CPU and memory consumption. One such case is redo generation by the database instance, which, in turn, increases both the rapid switching of redo logs and the creation rate and number of archived logs. This may cause an increase in overall I/O on file systems (or Oracle Automatic Storage Management disk groups), causing a systemwide performance issue. To alleviate this type of issue, you need to locate the session or sessions that caused the generation of high amounts of redo, but looking at OS metrics will not provide any insights into the offending session. In this case, you need to look at the sessions responsible for most of the load: the sessions generating maximum redo. Again, this information is available quite easily in the same V$SESSTAT view. The following query shows the sessions generating the most redo:

```
select sid, value
from v$sesstat s, v$statname n
where n.statistic# = s.statistic#
and n.name = 'redo size'
order by value desc;
```

SID	VALUE
13	11982752
10	3372240
17	964912
26	571324

... output truncated ...

It's clear from the output that SID 13 produced most of the redo, followed by SID 10, and so on.

The online version of this article at bit.ly/U8ml8P describes how to find the SQL statement that is generating the redo, provides a replacement SQL statement, and tests that new statement for redo generation.

OTHER STATISTICS

I've shown you how to use the redo size, session pga memory max, and CPU statistics used by a session. Here are some other useful statistics visible in the V$SESSTAT view:

- **physical reads:** the number of database blocks retrieved from disk
- **db block changes:** the number of database blocks changed in the session
- **bytes sent via SQL*Net to client:** the bytes received from the client over the network, which is used to determine the data traffic from the client

These are just a few of the 604 such statistics available in the V$SESSTAT view. It is not possible to describe each one of them, but I hope you get the general idea of how to use statistics to see the various resources used by a session and focus on performance issues. (Note that there is another view—V$SYSSTAT—that shows the statistics of the entire instance.)

CONCLUSION

In this article, you learned about two very important sources of performance tuning information in Oracle Database: (1) the history of wait events experienced by sessions—visible in V$SESSION_EVENT, and (2) the resource consumption by session—visible in V$SESSTAT. From the history, you can find out why a session has waited in the past and for how long—very useful information for diagnosing performance issues when they are no longer present. The resource statistics show the consumption of various resources such as CPU, memory, and redo by a specific session—very useful for focusing on resource-hogging sessions. Using these two views and scripts mentioned in this article, you can resolve many performance issues encountered in Oracle Database instances. Happy tuning! ◀

Arup Nanda (arup@proligence.com) has been an Oracle DBA since 1993, handling all aspects of database administration, from performance tuning to security and disaster recovery. He was *Oracle Magazine*'s DBA of the Year in 2003 and Enterprise Architect of the Year in 2012.

NEXT STEPS

READ
online-only article content
bit.ly/U8ml8P

"Beginning Performance Tuning"
bit.ly/SdU9j8

LEARN more about performance tuning
Oracle Database 2 Day + Performance Tuning Guide 11g Release 2 (11.2)
bit.ly/OOxhbH

Exadata

Replaces IBM Again

Giant European Retailer Moves Databases from IBM Power to Exadata

IBM
Building Smarter Planets

Oracle
Building Faster Computers

ORACLE®

For more details oracle.com/EuroRetailer

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 DATABASE

Selecting a Type That Is Right for You

Part 8 in a series on the basics of the relational database and SQL

Part 7 in this series, "From Floor to Ceiling and Other Functional Cases" (*Oracle Magazine*, September/October 2012), introduced common SQL number functions and showed how your queries can use them to modify the appearance of numeric result set data. It also introduced SQL substitution functions and showed how you can use them to manipulate result set data to convey more-meaningful results. Similarly, you can use SQL *date* functions and *datatype conversion* functions to manipulate data so that it displays differently from how it is stored in the database. This article introduces you to some of the more commonly used SQL date functions, along with some useful datatype conversion functions.

To try out the examples in this 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 11*g* Release 2.

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. (Note that SQL_101 is the user account to 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.

Finally—whether you installed the database software from scratch or have access to an existing Oracle Database instance—download, unzip, and execute

Last Two Digits of Current Year	Two-Digit Year Specified in Query	Century Referred To
Between 00 and 49	Between 00 and 49	Current
Between 00 and 49	Between 50 and 99	Previous
Between 50 and 99	Between 00 and 49	Next
Between 50 and 99	Between 50 and 99	Current

Table 1: Relationship among current year, two-digit year specified, and the century referred to as a result

Code Listing 1: Display date data in the Oracle Database default date format

```
SQL> set feedback on
SQL> select first_name, last_name, hire_date
  2    from employee
  3  order by hire_date desc, last_name, first_name;

FIRST_NAME                     LAST_NAME                      HIRE_DATE
------------------------------ ------------------------------ -----------
Theresa                        Wong                           27-FEB-10
Thomas                         Jeffrey                        27-FEB-10
mark                           leblanc                        06-MAR-09
michael                        peterson                       03-NOV-08
Roger                          Friedli                        16-MAY-07
Betsy                          James                          16-MAY-07
Matthew                        Michaels                       16-MAY-07
Donald                         Newton                         24-SEP-06
Frances                        Newton                         14-SEP-05
Emily                          Eckhardt                       07-JUL-04

10 rows selected.
```

Code Listing 2: Display date data in a different format by using TO_CHAR with a format mask

```
SQL> select first_name, last_name, TO_CHAR(hire_date, 'YYYY-MM-DD') hire_date
  2    from employee
  3  order by hire_date desc, last_name, first_name;

FIRST_NAME                     LAST_NAME                      HIRE_DATE
------------------------------ ------------------------------ ----------
Thomas                         Jeffrey                        2010-02-27
Theresa                        Wong                           2010-02-27
mark                           leblanc                        2009-03-06
michael                        peterson                       2008-11-03
Roger                          Friedli                        2007-05-16
Betsy                          James                          2007-05-16
Matthew                        Michaels                       2007-05-16
Donald                         Newton                         2006-09-24
Frances                        Newton                         2005-09-14
Emily                          Eckhardt                       2004-07-07

10 rows selected.
```

the SQL script available at bit.ly/OzGgwj to create the tables for the SQL_101 schema that is required for this article's examples. (View the script in a text editor for execution instructions.) Some of the examples also use the DUAL table. Recall that DUAL is an Oracle system table owned by the SYS user, not the SQL_101 schema. DUAL contains no meaningful data itself, but it is useful to query it as a way to experiment with functions that work on literals.

THE PERFECT FORMAT FOR YOUR DATE

The DATE datatype is stored in Oracle Database in an internal format that consists of both date and time information: the century, year, month, day, hour, minute, and second. For input and output of dates, every Oracle Database instance has a default date *format model* (also called a *mask*) that is set by the NLS_DATE_FORMAT *initialization parameter*. (Initialization parameters determine the default settings for Oracle Database instances. Users who have appropriate permissions can change some of these parameters on a per-database, per-instance, or per-session basis.) When you first query the data stored in a table column with a DATE datatype, Oracle Database displays it with a format mask of either DD-MON-YYYY or DD-MON-RR, depending on which is set as the default.

The RR format mask, which represents a two-digit year, was introduced to deal with end-of-century issues such as the Y2K problem. With RR, a two-digit year can refer to a year in the previous, current, or next century—depending on the current year and the two-digit year specified in the query. Table 1 shows the relationship between the current year, the range of two-digit year combinations, and the corresponding century referred to as a result.

For example, the last two digits of the current year (2012) are 12, which falls between 00 and 49. A SQL query issued during 2012 that specifies an RR year value of 15, therefore, refers to the year ending in 15 (2015) in the current century (the twenty-first), because 15 is between 0 and 49. A query issued in 2012 that specifies an RR year value of 98 refers to the year ending in

Code Listing 3: Default date format mask is used when optional parameter is not provided

```
SQL> select first_name, last_name, TO_CHAR(hire_date) hire_date_formatted
  2    from employee
  3  order by hire_date_formatted desc, last_name, first_name;

FIRST_NAME                     LAST_NAME                      HIRE_DATE
------------------------------ ------------------------------ -----------
Thomas                         Jeffrey                        27-FEB-10
Theresa                        Wong                           27-FEB-10
Donald                         Newton                         24-SEP-06
Roger                          Friedli                        16-MAY-07
Betsy                          James                          16-MAY-07
Matthew                        Michaels                       16-MAY-07
Frances                        Newton                         14-SEP-05
Emily                          Eckhardt                       07-JUL-04
mark                           leblanc                        06-MAR-09
michael                        peterson                       03-NOV-08

10 rows selected.
```

Code Listing 4: Use the TO_DATE conversion function in a WHERE clause

```
SQL> select first_name, last_name, TO_CHAR(hire_date, 'DD-MON-YYYY') hire_date
  2    from employee
  3   where hire_date > TO_DATE('01-JAN-2008', 'DD-MON-YYYY')
  4  order by hire_date desc, last_name, first_name;

FIRST_NAME                     LAST_NAME                      HIRE_DATE
------------------------------ ------------------------------ -----------
Thomas                         Jeffrey                        27-FEB-2010
Theresa                        Wong                           27-FEB-2010
mark                           leblanc                        06-MAR-2009
michael                        peterson                       03-NOV-2008

4 rows selected.
```

Code Listing 5: Error when the format mask does not match the provided string literal

```
SQL> select first_name, last_name, TO_CHAR(hire_date, 'DD-MON-YYYY') hire_date
  2    from employee
  3   where hire_date > TO_DATE('01-JAN-2008', 'MM/DD/RR')
  4  order by hire_date desc, last_name, first_name;
 where hire_date > TO_DATE('01-JAN-2008', 'MM/DD/RR')
                   *
ERROR at line 3:
ORA-01858: a non-numeric character was found where a numeric was expected
```

Code Listing 6: Implicit date conversion (not recommended) returns a result set

```
SQL> select first_name, last_name, TO_CHAR(hire_date, 'DD-MON-YYYY') hire_date
  2    from employee
  3   where hire_date > '01-JAN-2008'
  4  order by hire_date desc, last_name, first_name;

FIRST_NAME                     LAST_NAME                      HIRE_DATE
------------------------------ ------------------------------ -----------
Thomas                         Jeffrey                        27-FEB-2010
Theresa                        Wong                           27-FEB-2010
mark                           leblanc                        06-MAR-2009
michael                        peterson                       03-NOV-2008

4 rows selected.
```

98 (1998) in the previous century (the twentieth), because 98 is between 50 and 99.

The query in Listing 1 uses the EMPLOYEE table in the sample schema for this article. The query displays employees sorted from most recent to least recent date of hire. As you can see, the hire date data is displayed in DD-MON-RR format. For example, it shows that Roger Friedli was hired on 16-MAY-07. To change the way this data is displayed, you use the TO_CHAR conversion function in conjunction with a format model of your choosing. (You had a brief introduction to TO_CHAR in the last installment, where you saw that it can be used to convert a number to a text string.)

The query in Listing 2 modifies the way the date data from Listing 1 is displayed. To convert data of DATE datatype to a specific date format model, TO_CHAR takes one required parameter and one optional parameter. The required parameter is data of DATE datatype from a column, expression, or literal. The optional parameter is a textual format-mask representation of the date to be displayed. In Listing 2, the default format mask of DD-MON-RR is changed to display as YYYY-MM-DD.

Listing 3 demonstrates that the second parameter for TO_CHAR is optional. If it is left off, the format mask of the date data returned will simply be the default format mask. Note also that the datatype of the date returned is VARCHAR2. The output from Listing 3 is sorted by HIRE_DATE in descending order, but in *character*, not *date*, descending order. So, be aware that when you apply the TO_CHAR conversion function, your data is returned as character strings; you should plan and sort accordingly.

Code Listing 7: Attempted implicit date conversion fails

```
SQL> select first_name, last_name, TO_CHAR(hire_date, 'DD-MON-YYYY') hire_date
  2    from employee
  3   where hire_date > '01/01/2008'
  4  order by hire_date desc, last_name, first_name;
 where hire_date > '01/01/2008'
                   *
ERROR at line 3:
ORA-01843: not a valid month
```

Code Listing 8: Find the default date format for your current session

```
SQL> select sys_context ('USERENV', 'NLS_DATE_FORMAT')
  2    from dual;

SYS_CONTEXT('USERENV','NLS_DATE_FORMAT')
----------------------------------------------------------------------------------------------------
DD-MON-RR

1 row selected.
```

Code Listing 9: Display the time component of a value with a DATE datatype

```
SQL> set lines 32000
SQL> select first_name, last_name, TO_CHAR(hire_date, 'DD-MON-YYYY HH:MI:SS')
hire_date
  2    from employee
  3  order by hire_date desc, last_name, first_name;

FIRST_NAME                     LAST_NAME                      HIRE_DATE
------------------------------ ------------------------------ --------------------
Thomas                         Jeffrey                        27-FEB-2010 12:00:00
Theresa                        Wong                           27-FEB-2010 09:02:45
Donald                         Newton                         24-SEP-2006 12:00:00
Roger                          Friedli                        16-MAY-2007 12:00:00
Betsy                          James                          16-MAY-2007 12:00:00
Matthew                        Michaels                       16-MAY-2007 12:00:00
Frances                        Newton                         14-SEP-2005 12:00:00
Emily                          Eckhardt                       07-JUL-2004 12:00:00
mark                           leblanc                        06-MAR-2009 12:00:00
michael                        peterson                       03-NOV-2008 12:00:00

10 rows selected.
```

DATES WITH STRINGS ATTACHED

Just as you can convert a date to a string, you can convert a string literal to a date. The resulting expression can be compared with any other column's data of DATE datatype or another date expression. You perform the conversion by applying the TO_DATE conversion function to a text string, as shown in Listing 4. The query in Listing 4 not only returns all employees whose HIRE_DATE value is found to be greater than the date value 01-JAN-2008; it also demonstrates that the TO_DATE conversion function can be used in WHERE clauses as well as SELECT lists. The TO_DATE function is applied to the string literal 01-JAN-2008, with a format mask that helps the database interpret the supplied literal as a date.

When you provide a format mask to the TO_DATE function, the mask you choose must be the same as the one used in the string literal you supply. If the two do not agree, you will receive an error message similar to the one shown in Listing 5. When you convert a text literal, it is good practice to use the TO_DATE conversion function and explicitly specify an appropriate format mask. This way, your statement can be interpreted independently of any database, instance, or session default date settings.

Oracle Database will perform implicit date conversion where it can, if (and only if) the literal is already in the default date format. However, I do not recommend that you allow it to do so, because your code will be more fragile and less likely to perform well long-term. Listing 6 shows a query that relies on the default date format in Oracle Database and its ability to perform implicit date conversion on a string literal. Compare the result

in Listing 6 with that in Listing 7, which also attempts to perform an implicit date conversion. The query in Listing 7 fails because the database cannot interpret the date format mask of the literal value being compared with the values in the HIRE_DATE column of the EMPLOYEE table.

Because the default date format can be changed, it is best not to allow your queries to rely on an expected default format. Instead, always use the TO_DATE function on date string literals. One way to find out which default date format your current session is using is to execute the query shown in Listing 8. The SYS_CONTEXT function can be used by any session (and, therefore, any user) to see current session attributes.

Code Listing 10: WHERE clause using TO_DATE might not capture all possible values

```
SQL> select first_name, last_name, TO_CHAR(hire_date, 'DD-MON-YYYY HH:MI:SS')
hire_date
  2    from employee
  3   where hire_date = TO_DATE('27-FEB-2010', 'DD-MON-YYYY')
  4  order by last_name, first_name;

FIRST_NAME                     LAST_NAME                      HIRE_DATE
------------------------------ ------------------------------ --------------------
Thomas                         Jeffrey                        27-FEB-2010 12:00:00

1 row selected.
```

Code Listing 11: Truncate the time from a DATE value to return all records for a particular day

```
SQL> select first_name, last_name, TO_CHAR(hire_date, 'DD-MON-YYYY HH:MI:SS')
hire_date
  2    from employee
  3   where TRUNC(hire_date) = TO_DATE('27-FEB-2010', 'DD-MON-YYYY')
  4  order by last_name, first_name;

FIRST_NAME                     LAST_NAME                      HIRE_DATE
------------------------------ ------------------------------ --------------------
Thomas                         Jeffrey                        27-FEB-2010 12:00:00
Theresa                        Wong                           27-FEB-2010 09:02:45

2 rows selected.
```

Code Listing 12: Date range that returns records for a particular day

```
SQL> select first_name, last_name, TO_CHAR(hire_date, 'DD-MON-YYYY HH:MI:SS')
hire_date
  2    from employee
  3   where hire_date >= TO_DATE('27-FEB-2010', 'DD-MON-YYYY')
  4     and hire_date <  TO_DATE('28-FEB-2010', 'DD-MON-YYYY')
  5  order by last_name, first_name;

FIRST_NAME                     LAST_NAME                      HIRE_DATE
------------------------------ ------------------------------ --------------------
Thomas                         Jeffrey                        27-FEB-2010 12:00:00
Theresa                        Wong                           27-FEB-2010 09:02:45

2 rows selected.
```

Code Listing 13: The SYSDATE function

```
SQL> select SYSDATE, TO_CHAR(SYSDATE, 'DD-MON-YYYY HH24:MI:SS') sysdate_with_time
  2    from dual;

SYSDATE   SYSDATE_WITH_TIME
--------- --------------------
08-AUG-12 08-AUG-2012 14:25:08

1 row selected.
```

TAKING TIME WITH YOUR DATES

Recall that the Oracle DATE datatype includes a time component. You can either ignore the time component, as the examples in this article have done so far, or you can include it for display or comparison purposes. Listing 9 shows a query that includes the time component from each HIRE_DATE value for every employee listed in the EMPLOYEE table. Note that all the employee records except the one for Theresa Wong show a time value of 12:00:00. If you do not include a time when inserting a value into a column with a DATE datatype, the time will default to midnight (12:00:00 a.m. or 00:00:00 military time). To display or compare a date value in military time, use the HH24 format mask instead of HH.

Unless you know the exact time of the date values on which you'd like to filter—or unless all the time portions for your date values are already set to midnight—using date values in your WHERE clauses can produce unexpected results. Consider the query in Listing 10. You know from the results in the previous listings that two employees were hired on February 27, 2010, yet only one is returned in Listing 10's result set. The reason is that the TO_DATE function in the WHERE clause does not specify an exact time, so Oracle Database assumes that the time is midnight and returns only those records that contain the specified date value *and* midnight as the time component.

CUTTING YOUR DATE SHORT

When you would like to be able to filter on a certain date but do not want to have to include each individual time component, you can use a couple of different methods. One method is to include the TRUNC function (introduced in the previous installment in this series). It, like the TO_CHAR function, works not only on numbers but also on

date values. The TRUNC function helps cut off the time portion of a date if no optional format parameter is passed to it. This can be useful for date comparison purposes. Listing 11 shows a revised version of the query from Listing 10. As you can see, eliminating the time portion of the values in the HIRE_DATE column enables the comparison against the date value 27-FEB-2010 to retrieve all records with a HIRE_DATE value of 27-FEB-2010, irrespective of the time. The truncated HIRE_DATE value is made into a *date only* value to be compared with the corresponding *date only* value returned from the result of applying the TO_DATE function on the literal string 27-FEB-2010 with a date-only format.

Be aware, however, that you might sacrifice performance by applying a function to your table column values in a WHERE clause. Indexes (used to assist with data access efficiency—and not discussed in this series) can improve query performance in certain situations. Applying a function to a table column has the effect of ensuring that an index on the column might never be used. Also, this function would be applied to every value in that column for every row. Both actions are extreme performance inhibitors. Therefore, another method you can use is to specify a date range outside of the date(s) you would actually prefer to filter on. The query in Listing 12 retrieves the same result set as the query in Listing 11. The difference between the two is that the query in Listing 12 does not apply a function to the HIRE_DATE column data. Instead, it chooses a range just outside of the desired date(s) and encloses the filtered date data inside this range of values.

Code Listing 14: SYSDATE used in date arithmetic

```
SQL> select SYSDATE, (TO_DATE('01-JAN-2013', 'DD-MON-YYYY') - TRUNC(SYSDATE))
Days_till_2013
  2    from dual;

SYSDATE   DAYS_TILL_2013
--------- ---------------
08-AUG-12             146

1 row selected.
```

Code Listing 15: SYSDATE and date arithmetic combined with DATE data

```
SQL> select substr(last_name, 1, 10) last_name, substr(first_name, 1, 10)
first_name, hire_date, ROUND(MONTHS_BETWEEN(TRUNC(SYSDATE), TRUNC(HIRE_
DATE))/12, 2) YEARS_OF_SERVICE
  2    from employee
  3  order by years_of_service desc, last_name, first_name;

LAST_NAME  FIRST_NAME HIRE_DATE YEARS_OF_SERVICE
---------- ---------- --------- ------------------
Eckhardt   Emily      07-JUL-04              8.09
Newton     Frances    14-SEP-05               6.9
Newton     Donald     24-SEP-06              5.88
Friedli    Roger      16-MAY-07              5.23
James      Betsy      16-MAY-07              5.23
Michaels   Matthew    16-MAY-07              5.23
peterson   michael    03-NOV-08              3.77
leblanc    mark       06-MAR-09              3.42
Jeffrey    Thomas     27-FEB-10              2.45
Wong       Theresa    27-FEB-10              2.45

10 rows selected.
```

Code Listing 16: BETWEEN operator uses midnight in a date range comparison

```
SQL> select last_name, first_name, hire_date
  2    from employee
  3   where hire_date BETWEEN TO_DATE('26-FEB-2010', 'DD-MON-YYYY')
  4                       AND TO_DATE('27-FEB-2010', 'DD-MON-YYYY');

LAST_NAME                      FIRST_NAME                     HIRE_DATE
------------------------------ ------------------------------ ---------
Jeffrey                        Thomas                         27-FEB-10

1 row selected.
```

A SYSTEM FOR GETTING YOUR DATES RIGHT

You will often need to perform date arithmetic. A useful *built-in* function (one already built into Oracle Database) is SYSDATE. This function returns the current date and time that are set on the operating system of the computer on which the database resides. It takes no parameters. Listing 13 shows an example of using the SYSDATE function to return and display the current date and time.

SYSDATE can be extremely useful in date arithmetic. Listing 14 shows how many days are left in 2012 from the current date (August 8, 2012, in the example). Note that if the SYSDATE value were not truncated, the returned DAYS_TILL_2013 value would include some fraction of the SYSDATE value (to account for the time component). Because it is truncated, however, the entire current date is subtracted from January 1, 2013, to arrive at the result of 146 days left in the year. Listing 15 uses SYSDATE and date arithmetic (using a date function called MONTHS_BETWEEN) against the HIRE_DATE column of the EMPLOYEE table, to show the number of years of service for each employee.

Another method for performing date arithmetic is to use the BETWEEN operator, as demonstrated by the query in Listing 16. Be aware, however, that the BETWEEN operator uses the midnight (or 00:00:00) time component of the upper-range value in a date-range comparison. To include all

possible values for the date specified in the upper range of the date comparison, ensure that the date includes the full time component of your upper range. In the example in Listing 16, an upper-range date value of 27-FEB-2010 23:59:59 would have allowed both employee records with a HIRE_DATE value of 27-FEB-2010 to be included in the result set.

CONCLUSION

This article has shown you a few of the most common date functions and how they can be used to manipulate the way data is displayed. You've seen how to use the TO_CHAR and TO_DATE conversion functions and have learned the differences between them. You now know that dates all contain a time component that can be used or truncated according to your needs. You've been introduced to the SYSDATE function and date arithmetic. Last but not least, you now know the pitfalls to be aware of when you use DATE comparisons in WHERE clauses with TO_DATE and BETWEEN—and what you can do to avoid unexpected results. By no means has this article provided an exhaustive list of the Oracle Database date and datatype conversion functions. You can review the documentation for more details at bit.ly/PR7GQh and bit.ly/NOgf01.

The next installment of SQL 101 will discuss aggregate functions. ◀

Melanie Caffrey is a senior development manager at Oracle. She is a coauthor of *Expert PL/SQL Practices for Oracle Developers and DBAs* (Apress, 2011) and *Expert Oracle Practices: Oracle Database Administration from the Oak Table* (Apress, 2010).

NEXT STEPS

READ
SQL 101, Parts 1–7
bit.ly/omagsql101

LEARN more about date and datatype conversion functions
bit.ly/PR7GQh
bit.ly/NOgf01

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

Oracle SQL Developer User's Guide Release 3.1
bit.ly/I9w52z

Oracle Database development essentials
Oracle Database 2 Day Developer's Guide 11g Release 2 (11.2)
bit.ly/OzNal3

DOWNLOAD the sample script for this article
bit.ly/OzGgwj

ORACLE DATABASE

On External Table Queries, Data Consistency, and Nothing

Our technologist queries the operating system, locks manually, and uses the right NULL.

In some of my customer's databases, the DBAs are using datafiles with automatic extension but with many datafiles sharing the same file system, such as

tablespace A, datafiles /u01/oradata/ts_A_file01.dbf autoextend unlimited
tablespace B, datafiles /u01/oradata/ts_B_file01.dbf autoextend unlimited

and so on. The requested extension is that all datafiles must be capable of growing by at least 20 percent of their current size, so if, for example, ts_A_file01.dbf currently is 100 GB and ts_B_file01.dbf is 200 GB, we must ensure that at least 20 + 40 = 60 GB is free in the /u01/oradata file system.

The question is: how can we monitor this in a single query inside the database? Right now we have a complex script gathering free space from the df command in a text file, opening a cursor, calculating the current allocated space from DBA_DATA_FILES, and reading the df data via an external table.

This can be done in a single SQL query. To accomplish this, I'll need to be able to query disk free (df) interactively—without a complex set of maneuvers such as running a script or redirecting output. I'll start by making it so that the df output can be queried as if it were in a table. I'll be relying on a feature first added in Oracle Database 11*g* Release 2 but subsequently back-ported to version 10.2.0.5 as well (so this works in Oracle Database 10*g* Release 2 and above). This feature is the preprocessor directive. (I've written about this previously—go to bit.ly/P9jdGT.)

To start I'll create a directory where I can place a small shell script that will produce the df output:

```
SQL> create or replace
  2  directory exec_dir
  3  as '/home/tkyte/df'
  4  /
Directory created.
```

Next, I'll create a shell script named run_df.sh in that directory. This script will contain only the following:

```
#!/bin/bash
/bin/df -Pl
```

And the output of that script will look something like Listing 1.

Note that in the run_df.sh script, I used explicit pathnames to run `df`—I did not rely on the environment and on the path environment variable in particular. This is very important: when coding scripts for external tables—when coding scripts in general!—you always want to use explicit paths to run the program you actually intend to run. You don't have any real control over the environment this script will run in, so relying on the environment being set a certain way is a recipe for disaster.

So, now that I have the script and the directory, I am ready to create the external table. As you can see from the output in Listing 1, all I need to do is have the external table skip the first record and then parse each subsequent line, using white space as a delimiter. This is something an external table can do easily, as shown in Listing 2.

With the df external table created, I can now review the df output easily in a query, as shown in Listing 3.

Note: Hopefully you see how this approach could work easily for ps, ls, du, and so on—all the UNIX utilities could easily be considered "tables" now!

With this data available to me in the df external table, it is easy to start working on the query. All I need to do is join df to DBA_DATA_FILES with a join condition that matches the longest mount point possible to each filename. Before I do this, however, I'm going to change my df output for testing purposes. Because I have a rather small file system with just one mount point and I want to test my logic, I'm going to "fake" my df data by changing the script as shown in Listing 4.

Instead of running `df`, I'll fake the output by echoing two possible mount points for the datafiles. Also, I'll be able to change the available amount of data to test for the 20 percent datafile extension condition.

My single-query solution to the issue in this question is in Listing 5. Here's what's

Code Listing 1: Output of the df shell script

```
SQL> !./run_df.sh
Filesystem          1024-blocks      Used  Available   Capacity    Mounted on
/dev/mapper/VolGr...   18156292  10827600    6391528        63%    /
/dev/sda1                101086     12062      83805        13%    /boot
tmpfs                    517520         0     517520         0%    /dev/shm
```

happening on some of the lines:

On lines 3 and 4, I query the df external table. I purposely use a materialize hint to force the optimizer to load the df data into the equivalent of a global temporary table, because the query would tend to read and reread the external table over and over and the results of the df table could change as the query runs. This provides the logical equivalent of a consistent read on the df data. Also, if the query plan did involve rereading the external table, I would receive an error message at runtime:

```
KUP-04108 unable to reread file string
```

The documentation (bit.ly/OQ3PBT) explains the error:

Cause: The query that is executing requires that a datasource for an external table be read multiple times. However, the datasource is a sequential device that cannot be reread. Examples of this type of datasource are a tape or pipe.

Action: There are a few ways around this problem. One is to rewrite the query so that the external table is referenced only once. Another option is to move the datasource to a rereadable device such as a disk file. A third option is to load the data for the external table into a temporary table and change the query so that it references the temporary table.

On lines 27–30, I join DBA_DATA_FILES to df data with a WHERE clause, using LIKE. This will join every file in DBA_DATA_FILES to every possible mount point in the df output. I know that the goal, however, is to find the "longest" matching mount point, so to accomplish that I assign—on lines 24–26—a ROW_NUMBER to each row. This ROW_NUMBER will be sequentially assigned to each duplicated row in DBA_DATA_FILES, so if the FILE_NAME matches more than one MOUNT, each FILE_NAME occurrence will be assigned a unique, increasing ROW_NUMBER value. This ROW_NUMBER will be assigned after the data is sorted by the length of the MOUNT, from big to small.

Once I have that data, I apply a WHERE clause to save only the first entry for each FILE_NAME value—that predicate is WHERE rn = 1, on line 32. At the same time, I've added another column—TOT_BYTES—on lines 18-19. This will enable me to verify the 20 percent threshold.

The last step is to format and output the data. I print the columns I'm interested in and add a CASE statement on lines 11–15 to verify that 20 percent of the total bytes of storage allocated on a given mount point does not exceed the remaining available bytes of free storage.

So, now you see how to use external tables to query operating system output such as df, ps, and ls. Additionally, you can use them to

Code Listing 2: Creating the df external table

```
SQL> create table df
  2  (
  3   fsname   varchar2(100),
  4   blocks   number,
  5   used     number,
  6   avail    number,
  7   capacity varchar2(10),
  8   mount    varchar2(100)
  9  )
 10  organization external
 11  (
 12    type oracle_loader
 13    default directory exec_dir
 14    access parameters
 15    (
 16      records delimited
 17      by newline
 18      preprocessor
 19      exec_dir:'run_df.sh'
 20      skip 1
 21      fields terminated by
 22      whitespace ldrtrim
 23    )
 24    location
 25    (
 26      exec_dir:'run_df.sh'
 27    )
 28  )
 29  /
Table created.
```

Code Listing 3: Querying the df external table

```
SQL> select * from df;

FSNAME                                  BLOCKS      USED     AVAIL CAPACITY  MOUNT
--------------------------------------- ---------- ---------- -------- --------- ---------
/dev/mapper/VolGroup00-LogVol00  18156292  10827600  6391528 63%       /
/dev/sda1                          101086     12062    83805 13%       /boot
tmpfs                              517520         0   517520 0%        /dev/shm
```

Code Listing 4: Fake data for df (for testing purposes)

```
SQL> !cat run_df.sh
#!/bin/bash
#/bin/df -Pl

echo Filesystem         1024-blocks      Used  Available  Capacity  Mounted on
echo /dev/mapper/VolG...   18156292  10827600    6391528       63%  /
echo /dev/mapper/VolG...   18156292  10827600    6391528       63%  /home/
ora11gr2/app/ora11gr2/oradata/ora11gr2/ORA11GR2
echo /dev/sda1               101086     12062     83805       13% /boot
echo tmpfs                   517520         0    517520        0% /dev/shm
```

Code Listing 5: Single-query monitoring solution

```
SQL> with fs_data
  2  as
  3  (select /*+ materialize */ *
  4     from df
  5  )
  6  select mount,
  7         file_name,
  8         bytes,
  9         tot_bytes,
 10         avail_bytes,
 11         case
 12         when 0.2 * tot_bytes < avail_bytes
 13         then 'OK'
 14         else 'Short on disk space'
 15          end status
 16    from (
 17  select file_name, mount, avail_bytes, bytes,
 18         sum(bytes) over
 19           (partition by mount) tot_bytes
 20    from (
 21  select a.file_name,
 22         b.mount,
 23         b.avail*1024 avail_bytes, a.bytes,
 24         row_number() over
 25           (partition by a.file_name
 26            order by length(b.mount) DESC) rn
 27    from dba_data_files a,
 28         fs_data b
 29   where a.file_name
 30             like b.mount || '%'
 31         )
 32   where rn = 1
 33         )
 34   order by mount, file_name
 35  /

MOUNT    FILE_NAME                              BYTES  TOT_BYTES AVAIL_BYTES STATUS
-------- ------------------------------------ ---------- ---------- ------------ -------
/        /home/ora11gr2/app/ora11gr2/    360710144 2410283008  6544924672 OK
         oradata/ora11gr2/example01.dbf

/        /home/ora11gr2/app/ora11gr2/   1101004800 2410283008  6544924672 OK
         oradata/ora11gr2/sysaux01.dbf

/        /home/ora11gr2/app/ora11gr2/    924844032 2410283008  6544924672 OK
         oradata/ora11gr2/system01.dbf

/        /home/ora11gr2/app/ora11gr2/       131072 2410283008  6544924672 OK
         oradata/ora11gr2/system02.dbf

/        /home/ora11gr2/app/ora11gr2/     23592960 2410283008  6544924672 OK
         oradata/ora11gr2/users.dbf

/home/or /home/ora11gr2/app/ora11gr2/    144703488  376438784  6544924672 OK
a11gr2/a oradata/ora11gr2/ORA11GR2/
pp/ora11 datafile/o1_mf_big_tabl_
gr2/orad 7y3thv78_.dbf
ata/ora1
1gr2/ORA
11GR2

/home/or /home/ora11gr2/app/ora11gr2/    231735296  376438784  6544924672 OK
a11gr2/a oradata/ora11gr2/ORA11GR2/
pp/ora11 datafile/o1_mf_undotbs_
gr2/orad 78w1hprj_.dbf
ata/ora1
1gr2/ORA
11GR2

7 rows selected.
```

query anything that writes to standard out, including gunzip, sed, and so on. If you'd like to see an interesting example of using an external table to query multiple gzipped compressed files, take a look at bit.ly/PQgOol.

QUESTION ON LOCKING

If Oracle Database provides automatic locking and resolves deadlock situations automatically, why are there manual locking and deadlock solutions?

There are many reasons. The automatic locking falls into the 80/20 rule: 80 percent of the time, it is sufficient. But developers need to understand when locking is something they must be concerned about. Here is one of the simplest reasons: loss of update detection.

Suppose you have a table containing a primary key ID, names, addresses, and phone numbers. Further, there is an application that can read a row from this table, put it onscreen, and let end users update the values. When a user clicks a button, the application will issue UPDATE TABLE t SET NAME = :name, address = :address, phone = :phone WHERE id = :id.

Now, think about what would happen if two people pulled up your record at about the same time and one of them changed the address and the other changed your phone number. What would happen if they both clicked the button now?

One update would succeed—say the one that changed your address. All would be well and good so far.

The second update, to change your phone number, would end up updating the entire row again, reverting your address to what it was before, and the address update would be lost.

In a client/server application, you might employ what is known as pessimistic concurrency control. Before the application lets you type over the address or the phone number, it will issue a SELECT * FROM t WHERE id = :id AND name = :name AND address = :address AND phone = :phone FOR UPDATE NOWAIT to manually lock the record as long as the record still exists. If it does, great—you can modify the columns and be assured that the update will succeed and not overwrite anyone else's changes. If

the SELECT returns zero rows, you'll know that someone else modified the record and that you have to requery the record before you can modify it. If the query returns an error message, you'll know that someone else is currently working on it.

That is one reason for manual locking. Here is another situation. Suppose you have a business rule that says, "The total number of employees in a department cannot exceed 100." How do you enforce that? You have to make sure that when you make an insert into the EMP table, no one else is inserting into that same department. You have to serialize. You might use a DBMS_LOCK (user-defined lock) to do this, you might use the LOCK TABLE command, you might use a SELECT * FROM dept WHERE deptno=:x FOR UPDATE NOWAIT to serialize at the department level, and so on. But you need a manual lock.

Automatic locking covers 80 percent of everything you need. You need manual locking for the rest to ensure logical data consistency.

NULLS AND CARDINALITY

A few issues ago (bit.ly/NROO1a), I wrote up part of a presentation I delivered at Oracle OpenWorld in 2011. It was from "Five Things You Probably Didn't Know about SQL." In that article, I said I'd be writing up other sections of that presentation, and that is what I'm doing here. This time I'll take a look at NULLs and their possible effect on cardinality estimates and index use.

I've written many times (see bit.ly/Pjrbyr, for example) about how important cardinality estimates are to the cost-based optimizer (CBO). In short, if the optimizer guesses incorrectly how many rows will flow through steps in a query plan, you'll probably get a bad plan and a poorly performing query. For example, if the optimizer believes that it is going to retrieve 2 rows out of 1,000,000, it will likely choose to use an index if one is available. But what if, at runtime, you discover you are not getting just 2 rows but 500,000? Then using the index will be the worst-possible approach.

Getting the right estimated cardinality values is important to the optimizer, and anything you do to throw that off—to make it harder to get the right estimate—is bad.

Code Listing 6: Query of NULLs and a good plan

```
SQL> set autotrace traceonly explain
SQL> select *
  2    from t
  3   where end_date
  4         between to_date( '01-sep-2010', 'dd-mon-yyyy' )
  5             and to_date( '30-sep-2010', 'dd-mon-yyyy' );

Execution Plan
----------------------------------------------------------------------

Plan hash value: 1601196873
-------------------------------------------------------------------------------------
| Id  | Operation         | Name | Rows  | Bytes | Cost (%CPU)| Time     |
-------------------------------------------------------------------------------------
|   0 | SELECT STATEMENT  |      | 36024 |  3588K|   339   (1)| 00:00:05 |
|*  1 |  TABLE ACCESS FULL| T    | 36024 |  3588K|   339   (1)| 00:00:05 |
-------------------------------------------------------------------------------------
Predicate Information (identified by operation id):
---------------------------------------------------

   1 - filter("END_DATE"<=TO_DATE(' 2010-09-30 00:00:00', 'syyyy-mm-dd
              hh24:mi:ss') AND "END_DATE">=TO_DATE(' 2010-09-01 00:00:00',
              'syyyy-mm-dd hh24:mi:ss'))
```

But what does this have to do with NULL values? Nothing really—it has to do with what can happen when developers do not use a NULL value when they should have. Many times developers fear using NULLs: they do not understand them, and they do not believe they can be indexed, so they avoid them. They will use a "fake" value—such as 01-JAN-9999—to represent a missing date value. This is a bad idea for many reasons; the first I'll show you is how it can throw off cardinality estimates.

I'll start with a table that uses a NULL value. This table will represent "effective-dated" records—that is, each record has an effective start date and end date. If a record is "current"—that is, it doesn't have an end date yet—the end date will be NULL. Here is the CREATE TABLE statement for that table:

```
SQL> create table t
  2  pctfree 20
  3  as
  4  select a.*,
  5       case when mod(rownum,100) <= 50
  6            then last_ddl_time
  7       end end_date
  8    from all_objects a;

Table created.
```

In this table, about half of the rows have an END_DATE and half are NULL. Next, I'll create an index on the END_DATE column for retrieval purposes:

```
SQL> create index t_idx
  2  on t(end_date);

Index created.
```

If I look at the data, I'll discover that many of the values are clumped together in one month:

```
SQL> select count(*)
  2    from t
  3   where end_date
  4   between to_date('01-sep-2010',
                      'dd-mon-yyyy')
  5       and to_date('30-sep-2010',
                      'dd-mon-yyyy');

  COUNT(*)
----------
     36267
```

My ALL_OBJECTS view has about 72,000 records in it, so table T has about 72,000 records, and half of them fall in the month of September 2010. Next I'll gather statistics on this table, and these statistics will include histograms on the END_DATE column. If you are curious about how or why they have histograms, refer to the "Why Does My Plan Change?" section of bit.ly/PQnpPB.

Code Listing 7: Result of rerunning the Listing 6 query against fake data—and a bad plan

```
-------------------------------------------------------------------------------
| Id  | Operation                    | Name  | Rows  | Bytes | Cost (%CPU)|
-------------------------------------------------------------------------------
|   0 | SELECT STATEMENT             |       |   175 | 18375 |    10   (0)|
|   1 |  TABLE ACCESS BY INDEX ROWID| T     |   175 | 18375 |    10   (0)|
|*  2 |   INDEX RANGE SCAN           | T_IDX |   175 |       |     2   (0)|
-------------------------------------------------------------------------------
Predicate Information (identified by operation id):
---------------------------------------------------
   1 - filter("END_DATE"<=TO_DATE(' 2010-09-30 00:00:00', 'syyyy-mm-dd
              hh24:mi:ss') AND "END_DATE">=TO_DATE(' 2010-09-01 00:00:00',
              'syyyy-mm-dd hh24:mi:ss'))
```

```
SQL> begin
  2     dbms_stats.gather_table_
stats(user, 'T');
  3  end;
  4  /
PL/SQL procedure successfully completed.
```

Now I'll look at the data to understand it before I proceed:

```
SQL> select count(*),
  2         count(distinct end_date),
  3         count(end_date),
  4         min(end_date),
  5         max(end_date)
  6    from t;

CNT    CNTD  CNT2   MIN        MAX
-----  ----  -----  ---------- ----------
72228  703   36850  01-OCT-02 30-SEP-11
```

As you can see, there are 72,228 rows in the table, with a low value of October 1, 2002 and a high value of September 30, 2011. I know that about half of these rows fall into September 2010 (there is data skew). Also, it has 703 distinct dates, which will affect the type of histogram the optimizer can store. Because there are more than 255 distinct values, the optimizer will have access to a height-balanced, rather than frequency-based, histogram.

Now if I execute a query against this table, asking for the rows in September 2010, the optimizer will be able to intelligently come up with a query plan, as shown in Listing 6.

Because the optimizer correctly estimated the cardinality of this query at 36,024, it knew to do a full table scan. Using an index would have been a bad idea.

However, what if the developer decides to use a fake date instead of NULL? What will happen then? I'll update the data and gather statistics again:

```
SQL> update t
  2    set end_date =
  3        to_date( '01-jan-9999',
'dd-mon-yyyy' )
  4   where end_date is null;
35378 rows updated.

SQL> commit;
Commit complete.

SQL> begin
  2    dbms_stats.gather_table_
stats(user, 'T');
  3 end;
  4 /
PL/SQL procedure successfully completed.
```

Now when I run the query in Listing 6 again, I see the results in Listing 7.

The optimizer has incorrectly estimated the cardinality for this predicate, leading it to think an index would be appropriate when it is not. This is all due to the fake values.

Here is another reason to avoid fake values. In Washington DC, a car owner decided to get a vanity (personalized) license plate for his car. The license plate the owner chose was NO TAGS. The owner chose it as a joke, but guess what value the programmers at the Department of Motor Vehicles (DMV) had decided to use (instead of NULL) to represent a missing license plate value? The title of the news article (yhoo.it/NRUisL) about this is "DC Man's 'NO TAGS' Vanity Plate Earns Him $20,000 in Tickets."

Easy to fix, you might say: just don't get NO TAGS as a license plate. Well, not so fast. Another motorist, Nick Vautier, wanted his initials—NV—as his tag. In California, where he lives, the DMV had decided to use NV, for "not visible," as its "no tags" value.

Alabama uses XXXXXXX as its "no tags" value, and a news article (bit.ly/RNuqUT) about the car owner—Scottie Roberson—who got that license plate, mentioned a possible solution: "Birmingham Traffic Engineer Gregory Dawkins says the city may change the system to keep Roberson from receiving more tickets. He says 'maybe we just need to leave that part blank altogether.'"

See bit.ly/OgmnJq for other examples. There is no good default value to represent data that is missing. Just use NULL. ◀

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

FOLLOW Tom on Twitter
@OracleAskTom

READ more Tom
bit.ly/omagasktomk

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 11*g* Release 2
bit.ly/epBiUG

LEARN more about preprocessor directive
bit.ly/P9jdGT

external tables that query files
bit.ly/PQgOol

cardinality and the cost-based optimizer
bit.ly/Pjrbyr

1981

Moscone Center

Encompassing nearly 1 million square feet today, Moscone Center—San Francisco's largest convention center—opened in **1981** and played host to the 1984 Democratic Convention.

2002

Taxi! Take me to OracleWorld?

No, you're not hallucinating and that's not a typo. For two years ('02 and '03), the annual conference was simply called OracleWorld. Have you heard people talking about OracleWorld since 2003?

1982 INTERNATIONAL ORACLE USERS WEEK (IOUW)

Imagine this: When Oracle was still called Relational Software Inc. (RSI), about 50 attendees from a dozen or so companies gathered in a small, windowless room at the San Francisco Grand Hyatt for the first International Oracle Users Week (actually just three days, August 23–25).

After the inaugural event in **San Francisco**, the International Oracle Users Week (IOUW) went on tour; our ears are still ringing.

1986
Two Exhibitors

Oracle hosted IOUW at the Hyatt Regency in San Francisco with nearly 500 attendees and two exhibitors. During the conference, a general user meeting was held, and the International Oracle Users Group (IOUG) was formed.

1995
Going to California

During IOUW in Philadelphia, Oracle CEO Larry Ellison announced via satellite that beginning in 1996 Oracle would take control of the conference, call it Oracle OpenWorld, and host it in California.

1996

Open Your Golden Gate

Oracle OpenWorld '96 featured 420 sessions and 17,200 visitors at San Francisco's Moscone Center.

International Perspective

In addition to San Francisco, many other world-class cities have hosted their own Oracle OpenWorld conferences. **Beijing, Copenhagen, Paris, São Paulo, Tokyo**, and more have welcomed Oracle OpenWorld and added their international flavor to it.

REVISIONIST HISTORY

We asked our Facebook fans where and when the first Oracle OpenWorld was held:

San Francisco, 1986	54%
San Francisco, 1998	30%
Los Angeles, 1997	8%
Boston, 1983	8%

Your Turn

From **Elton John** and **Billy Joel** to **Journey** and **Aerosmith**, the Oracle Appreciation Event has hosted a number of major musical artists over the years. Who is your favorite? Anyone you would like to see in the future? Visit Facebook/OracleMagazine and let us know. **bit.ly/orclmagfb**

BETTMANN/CORBIS; HARTMANN STUDIOS

See ORACLE TEAM USA

World-Class Sailing, Cutting-Edge Technology, and the Best Viewing in Italy

Photography: Guilain Grenier/ ORACLE TEAM USA

Join fans from around the world as they gather in Italy to watch the America's Cup World Series events, featuring ORACLE TEAM USA.

Venice: April 16–21, 2013
Naples: May 14–19, 2013

It will be—quite literally—the biggest thing in sports.

Find Out More
oracleteamusa.com

Stay Connected

#oracleteamusa

ORACLE TEAM USA