Oracle Beats EMC, NetApp

Storage Magazine Quality Awards

#1 in Quality

#1 in Features

#1 in Reliability

#1 in Support

#1 Overall

ORACLE®

oracle.com/storage

Source: *Storage Magazine*, January 2012 issue, network attached storage.

JULY/AUGUST 2012

ORACLE

ORACLE.COM/ORACLEMAGAZINE

ORACLE MAGAZINE

LOCKDOWN

Oracle Database Security solutions deliver layers of enterprise information protection

Future-proof

Choose the best of today's tools to build next-generation Oracle Database applications

Engineered for Backup

Sun ZFS Backup Appliance delivers 20 TB/hour of full backup for Oracle engineered systems

Trusted Future Protection

Fujitsu's SPARC ENTERPRISE delivers Trust and Business Continuity

ORACLE Diamond Partner

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

Please visit : www.fujitsu.com/oracle/

Sun FUJITSU
ORACLE SPARC

shaping tomorrow with you

FUJITSU

ORACLE MAGAZINE

LOCKDOWN /30

Companies spend billions of dollars each year securing their IT systems worldwide, including software, services, and support. Despite these massive investments, attacks have continued, and the attackers—with the help of social engineering and sophisticated automated tools—have continued to succeed. Two-thirds of sensitive and regulated data resides in databases, and unless the databases are protected using a multilayered security architecture, that data is at risk. Learn how Oracle Database security solutions, including Oracle Database Firewall, Oracle Advanced Security, and Oracle Database Vault, mitigate that risk by addressing potential threats from outside and inside the database. *—By David Baum*

Cover: I-Hua Chen

Future-Proof / 36

The next platforms for application delivery are here, and users are demanding a rich and portable experience on these rapidly evolving devices. Application developers are now creating rich applications that need to work on everything from the corporate desktop to smartphones and tablets. Development tools must enable rapid application development to reach new devices quickly, and the tools must deliver a full-featured user experience to all devices. Read how three organizations are using Oracle databases and database development tools to deliver feature-rich applications fast to the latest platforms and devices. *—By David A. Kelly*

Up Front / 6

At Oracle / 11

Community / 24

Technology / 41

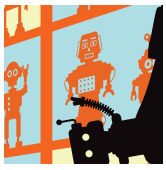

Comment / 67

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

Contributing Editor and Writer
Blair Campbell

Editor in Chief, Oracle Technology Network
Justin Kestelyn justin.kestelyn@oracle.com

Technology Advisor
Tom Kyte

Contributors
Marta Bright, Jeff Erickson, Fred 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, Chris Strach

Production Designers
Sheila Brennan, Kathy Cygnarowicz

EDITORIAL BOARD

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

SUBSCRIPTION INFORMATION

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

MAGAZINE CUSTOMER SERVICE

oracle@halldata.com **Fax** +1.847.763.9638 **Phone** +1.847.763.9635

PRIVACY

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

PUBLISHING

Publisher
Jeff Spicer jeff.spicer@oracle.com

Production Director and Associate Publisher
Jennifer Hamilton jennifer.hamilton@oracle.com +1.650.506.3794

Senior Manager, Audience Development and Operations
Karin Kinnear karin.kinnear@oracle.com +1.650.506.1985

ADVERTISING SALES

Associate Publisher
Kyle Walkenhorst kyle@sprocketmedia.com +1.323.340.8585

Northwest and Central U.S.
Tom Cometa thomas.cometa@sbcglobal.net +1.510.339.2403

Southwest U.S. and LAD
Shaun Mehr shaun@sprocketmedia.com +1.949.923.1660

Northeast U.S. and EMEA/APAC
Mark Makinney mark.makinney@sprocketmedia.com +1.805.709.4745

Recruitment Advertising
Tim Matteson tim@mattesonmail.com +1.310.836.4064

Mailing-List Rentals
Contact your sales representative.

RESOURCES

Oracle Products
+1.800.367.8674 (U.S./Canada)

Oracle Services
+1.888.283.0591

Oracle Press Books
oraclepressbooks.com

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

Printed by Brown Printing

ORACLE PRODUCT INDEX

A Convenient Truth

The Web is stored in tables, and that's good news.

The recent *Computerworld* article "10 Hard Truths for Software Developers" (bit.ly/J8RKme) described "10 aspects of programming developers must learn to live with." The hard truths of the article included information about the process of developing and maintaining code, the people involved in and affected by code development, and software development technologies.

The second item on this list of hard truths was this: most of the Web is just data stored in tables. For me, however, hard truth #2 was anything but difficult, and in fact, it was the best news in that article.

THE TRUTH ABOUT TABLES

The *Computerworld* article gives examples of the many types of information contained in tables on the Web and points out that programming languages have developed features specifically for working with information in tables. The hard part of the "data stored in tables" truth presented in the article appears to be that tables are not new, but rather are years-old constructs that have organized the information of the internet in clearly defined structures. The article also points out that NoSQL *is* new, but the hard truth there is that NoSQL is also about data in tables.

To me, the fact that most of the Web is data stored in tables is fantastic. The downside to that truth is that only *most* of the Web is in tables. And as for the hard truth that even a newer technology like NoSQL works with data in tables, I see only benefit in adding new capabilities for the structures that already contain most of the information of the Web.

Like many people, I've counted on tables—on the Web and in intranet applications—for Web transactions and interactions as well as for storage and access to various types of information for years. For internal projects, I've used relational database tables in situations where I first did not understand how using tables would help, but in each case the truth was that using the tables was convenient and definitely added value to both the process and the result.

IN THIS ISSUE

An important truth about some tables, specifically Oracle Database tables, is that there are multiple layers of security designed specifically to protect table information from both outside and inside threats. "Lockdown" (page 30) describes various Oracle Database security solutions that encrypt, lock, and mask the information in Oracle Database tables. One product, Oracle Database Firewall, monitors activity on the network to help prevent unauthorized access, SQL injections, privilege or role escalation, and other external and internal attacks.

Supporting the excellent truth that the Web is data stored in tables, "Future-Proof" (page 36) describes Oracle database products and database development tools used at three organizations to manage information in database tables for sophisticated Web and mobile device applications. The solutions include a custom application for customer relationship management, a geospatial application, and an e-mail marketing application, working with table data in Oracle Database and MySQL and using Oracle Application Express and Oracle Data Provider for .NET development technologies.

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

NEXT STEPS

READ
"10 Hard Truths for Software Developers"
bit.ly/J8RKme

READ more about
Oracle Database security
oracle.com/us/products/database/security/overview

Oracle developer tools
oracle.com/technetwork/developer-tools

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.

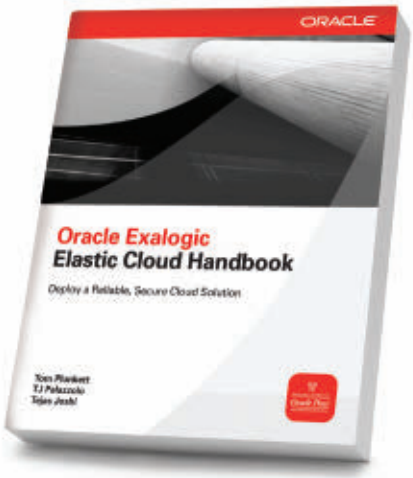

PeopleSoft PeopleTools Data Management and Upgrade Handbook
Paula Dean
Implement proven strategies for the PeopleSoft PeopleTools data management and upgrade process.

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

Oracle Solaris 11 System Administration: The Complete Reference
Michael Jang, Harry Foxwell
Master the skills needed to administer the new release of this powerful operating system.

OCA Java SE 7 Programmer I Study Guide
Edward Finegan, Robert Liguori
Get complete coverage of all objectives for exam 1Z0-803 plus challenging practice exam questions.

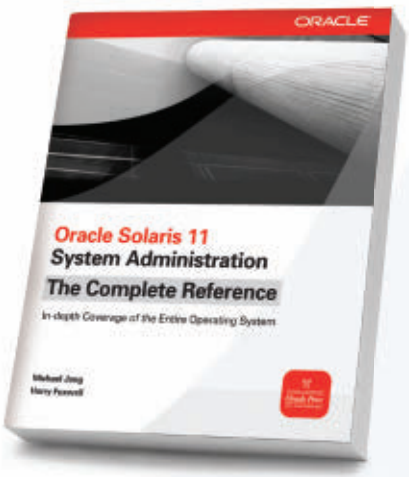

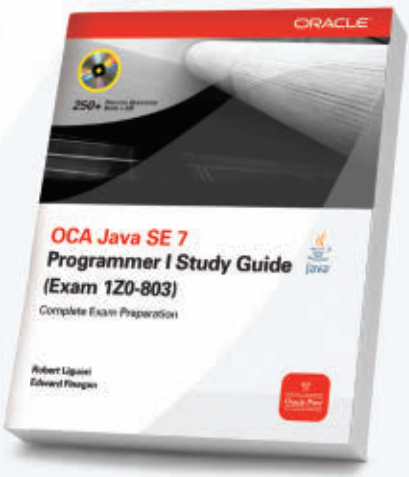

Available in print and eBook formats

Join the Oracle Press Community: www.OraclePressBooks.com

Follow us: @OraclePress

ORIGINAL • AUTHENTIC
Oracle Press
ONLY FROM McGRAW-HILL

CHECK-IN APPS

12 percent[1] of U.S. smartphone owners use check-in apps—geosocial services that allow them to announce where they are. Here are our favorites.

Facebook

Its huge user base makes Facebook's check-in capability, called Places, hard to beat. This app lets you tell people where you've been, where you are right now, and where you are going. You can also check in other Facebook users. Free (iPhone, Android). **facebook.com**

Foodspotting

Don't just check in at a restaurant—check in and snap pictures of the dishes you love. Share and search for favorite foods by location, with maps. Free (iPhone, Android, BlackBerry, Windows Phone). **foodspotting.com**

foursquare

More than 20 million people worldwide use this app to check in at restaurants and other businesses, get recommendations, connect with friends, and score special discounts. Free (iPhone, Android, BlackBerry, Nokia, Palm, Windows Phone). **foursquare.com**

Yelp

With 66 million users and 25 million reviews of local businesses in North America and Europe, Yelp lets users compare and share opinions on the places they go. Free (iPhone, Android, BlackBerry, Palm, Windows Phone). **yelp.com**

UNDER SURVEILLANCE

Setting up a home or office video surveillance system is no small undertaking. Dropcam is trying to change that with Dropcam HD, a small monitoring camera that takes minutes to set up, is Wi-Fi enabled, and streams video to the cloud. With an account on dropcam.com, you can peek at live or recorded video using a computer or smartphone. Dropcam streams 720p HD video and has an 80-degree viewing angle. Cool features include infrared LEDs that let Dropcam HD see in complete darkness, and two-way audio so that you can actually hold a conversation with a person on camera from your smartphone or PC. *US$149; viewing plans range from free to US$29.95/month.* **dropcam.com**

SECURITY TOPS MOBILITY CONCERNS

A survey of 500 business and IT professionals in a variety of industries in the U.S. revealed that for 70 percent of respondents, security considerations are the greatest risk involved in supporting mobility. Respondents cited these concerns:

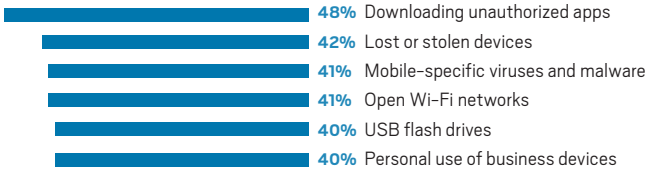

Source: CompTIA's Trends in Enterprise Mobility *study, comptia.org*

YOUR TURN

What mobile-related security concerns keep you up at night? **Take our Facebook poll** at bitly.com/oraclmagfb.

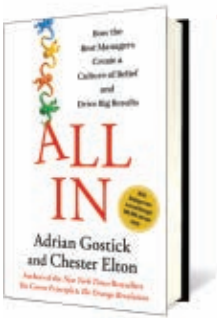

"The average employee spends about 15 hours a month complaining about his or her manager. That's basically 24 days a year, a full month of workdays, grumbling and getting nothing done. Right now employee teams and entire organizational cultures around the world are crumbling from misunderstanding and neglect."

—Adrian Gostick and Chester Elton, authors of *All In* (Free Press, 2012)

STRESSED? YOU'RE NOT ALONE

Ready to quit? **67%** of IT administrators have considered switching careers because of stress. More than 200 IT administrators in U.S. organizations with 10 to 500 employees were surveyed.

Source: *GFI Software and Opinion Matters, bit.ly/KrfnWQ*

Security Threats Rising

83% of business and IT executives surveyed believe information security threats are on the rise. And **70%** rated security as a higher or upper-level priority in 2011, compared to **49%** in 2010. Five hundred U.S. IT and business executives who set or implement information security policies and processes were surveyed online in November and December 2011.

Source: CompTIA's Ninth Annual Information Security Trends *study, comptia.org*

WRITE ON

Typing on the iPad's native keyboard is possible, but it's not the most comfortable thing to do. Plus, the native keyboard covers half the screen, leaving very little real estate for you to work with. Logitech's Ultrathin Keyboard Cover with Bluetooth technology addresses these issues and also doubles as an elegant and sturdy aluminum iPad cover. Thanks to a clever magnetic clip, the keyboard quickly attaches to the iPad. A built-in slot right above the keys lets the iPad stand at the optimal angle for reading or typing. The keyboard has a great feel when you are typing on it, and it makes functions such as text highlighting, copying and pasting, and searching a breeze. *US$99.* **logitech.com**

CHRIS STRACH

1 pewinternet.org/reports/2011/rocation/overview.aspx

Make your data dance.

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

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

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

MELISSA DATA®
Your Partner in Data Quality

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.

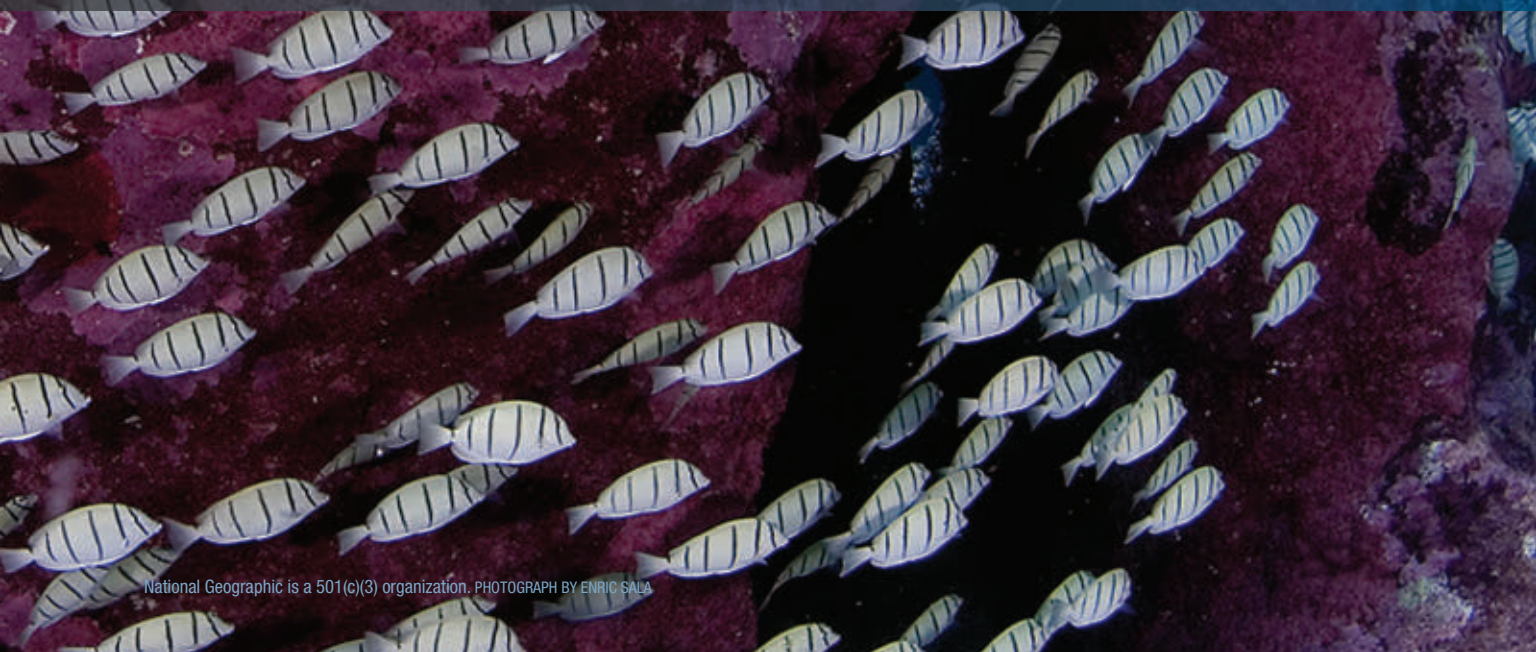

National Geographic is a 501(c)(3) organization. PHOTOGRAPH BY ENRIC SALA

Technology Events

Conferences and sessions to help you stay on the cutting edge

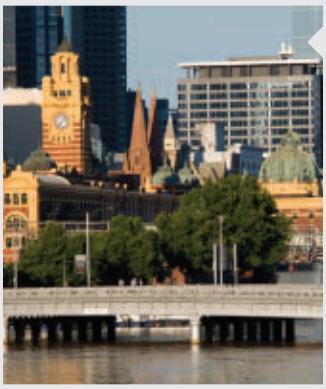

InSync12: Quest Australia and New Zealand Users Conference

August 21–23, Melbourne, Victoria, Australia

The independent Oracle user groups in Australia come together to present a comprehensive conference on Oracle applications, middleware, and Oracle database technology. More than 100 speakers from Australia and abroad, including Oracle employees, developers, and users, will deliver keynotes and sessions.

insync-conference.com.au/blog

AACE International Annual Meeting

July 8–11, San Antonio, Texas

▶ **www.aacei.org/am/currentam**

The American Association of Cost Engineers will gather for technical sessions, exhibits, talks, and networking events.

OR2012: The 7th International Conference on Open Repositories

July 9–13, Edinburgh, Scotland

▶ **or2012.ed.ac.uk**

With the theme "Open Services for Open Content: Local In for Global Out," this conference covers open and augmented content, distributed systems, and data delivery infrastructures. It features general conference sessions and open user group meetings.

TechNet Land Forces—South 2012

July 10–12, Tampa, Florida

▶ **afcea.org/events/tnlf/south12/intro.asp**

At this forum for the U.S. Army and Marine Corps as well as Department of Defense, Navy, Air Force, and Coast Guard security forces, discussions focus on technology issues related to joint and coalition land warfare and U.S. border protection.

OAUG Connection Point

July 11–12, Austin, Texas

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

The Oracle Applications Users Group (OAUG), the South Texas Oracle Applications Users Group (STOAUG), and the OAUG Upgrade Special Interest Group collaborate to present two days of technical sessions on Oracle E-Business Suite Release 12.1 and Oracle Fusion Applications.

The No Fluff Just Stuff Software Symposium Tour

July 13–14, Columbia, Maryland

July 20–22, Austin, Texas

August 3–5, Des Moines, Iowa

▶ **nofluffjuststuff.com**

The No Fluff Just Stuff Software Symposium Tour covers the latest trends within the Java ecosystem and agility space.

Gartner Security and Risk Management Summit

July 16–17, Sydney, Australia

▶ **bit.ly/IcieoK**

This event is for IT and business executives responsible for creating, implementing, and managing a proactive and comprehensive strategy for information security, risk management, governance, business continuity management, and business resiliency.

O'Reilly OSCON Open Source Convention

July 16–20, Portland, Oregon

▶ **oscon.com**

More than 3,000 developers, open source experts, businesspeople, investors, and vendors

GETTY IMAGES

ORACLE USER GROUPS

Sacramento Java Users Group Meetings
July 10 and August 14, Sacramento, California
sacjug.org

Ohio Oracle Users Group Meeting
July 19, Columbus, Ohio
ooug.org

PackerLand Users Group Meeting
July 19, Green Lake, Wisconsin
questdirect.org/events/2464800

Southwest Regional Oracle Applications User Group Conference
July 20, Irvine, California
sroaug.com

Michigan Oracle Users Summit Golf Outing
July 23, Plymouth, Michigan
mous.us/golf.html

Higher Education User Group Regional Conferences
July 25, Arlington, Texas
August 8–9, Naperville, Illinois
heug.org

Austin Java Users Group Meetings
July 31 and August 28, Austin, Texas
austinjug.org

Oracle Transportation Management User Conference
August 12–15, Philadelphia, Pennsylvania
otmconference.com

Credit Research Foundation Oracle Users Group Meeting
Held at the Credit and Accounts Receivable Open Forum and Expo
August 13–15, Chicago, Illinois
crfonline.org/events/august/august.asp

Oregon and Southern Washington Oracle Users Group Meeting
August 14, Portland, Oregon
oswoug.org

North Central Oracle Apps User Group Training Day
August 17, Oakbrook Terrace, Illinois
ncoaug.org

Rocky Mountain Oracle Users Group Quarterly Education Workshop
August 17, Littleton, Colorado
rmoug.org

Southern California Oracle Users Group Quarterly Meeting
August 21, location to be announced
socoug.org

KC R Users Group Meeting
August 25, Kansas City, Missouri
groupspaces.com/kcrusersgroup

Big data is a top agenda topic at TDWI World Conference in San Diego, California, July 29–August 3.

will converge for hundreds of sessions on open source languages and platforms. An exposition, "birds of a feather" sessions, and social events are also on the agenda.

NASPE 2012 Annual Meeting

July 22–25, Jersey City, New Jersey

▶ **bit.ly/JX9sfC**

The National Association of State Personnel Executives conference includes speakers and roundtable discussions, as well as sessions on IT, healthcare, retirement trends, performance measurement, and working with the media.

Sales 2.0 Conference

July 23, Boston, Massachusetts

▶ **sales20conf.com/boston2012**

Leaders of sales, marketing, and related groups will learn to balance people, process, and technology to improve overall organizational effectiveness at this conference.

BriForum

July 24–26, Chicago, Illinois

▶ **briforum.com/US**

This gathering features more than 40 technical sessions on desktop and application virtualization, plus live demos and networking opportunities. Attendance is limited to 325 attendees.

NACUBO 2012 Annual Meeting

July 28–31, Washington DC

▶ **nacuboannualmeeting.org**

The National Association of College and University Business Officers' 50th annual meeting offers a variety of education sessions for professionals in college and university business offices.

NASBO 2012 Annual Meeting

July 29–August 1, Providence, Rhode Island

▶ **bit.ly/ICMpGk**

The National Association of State Budget Officers gathers to hear expert speakers discuss the economy, state revenues, healthcare reform, and other pertinent topics.

TDWI World Conference

July 29–August 3, San Diego, California

▶ **tdwi.org**

With the theme "Big Data Tipping Point," this conference will look at how the internet, social media, and streaming data are fundamentally changing data warehousing and business intelligence as we know it.

TASSCC Annual Conference

August 12–15, Arlington, Texas

▶ **bit.ly/HOuZoc**

Join the Texas Association of State Systems for Computing and Communications for a conference focusing on the challenges and opportunities faced by public sector IT professionals in the Lone Star State.

Enkitec Extreme Exadata Expo (E4)

August 13–14, Dallas, Texas

▶ **enkitec.com**

This two-day gathering is designed for technical directors who already have an Oracle Exadata Database Machine, as well as those who are considering one. A working Oracle Exadata system will be onsite, and experts will explore all facets of working with the engineered system.

XChange

August 19–22, Dallas, Texas

▶ **xchange-events.com/x12**

More than 250 solution providers from the U.S. and Canada will view demonstrations of new products, services, and solutions from leading IT vendors. One-on-one appointments and small group vendor meetings are available.

Medicaid Enterprise Systems Conference

August 19–23, Boston, Massachusetts

▶ **mesconference.org**

This conference brings together thought leaders from the public and private sectors to share ideas and information related to Medicaid systems and initiatives.

Air Force Information Technology Conference

August 27–29, Montgomery, Alabama

▶ **afitc.gunter.af.mil**

More than 5,000 active duty, guard, reserve, and civilian personnel will discuss how hardware, software, and services are used to aid the national defense and intelligence communities.

LinuxCon North America

August 29–31, San Diego, California

▶ **bit.ly/KqjXW7**

This annual conference focuses on collaboration and education for the Linux community. It includes colocated, topic-specific minisummits and workgroups, including the Linux Kernel Summit and the Linux Plumbers Conference.

REGISTER for these upcoming Oracle events

Oracle OpenWorld 2012
September 30–October 4
oracle.com/openworld

JavaOne
September 30–October 4
oracle.com/javaone

MySQL Connect
September 29–30
oracle.com/mysqlconnect

Oracle Customer Experience Summit @ OpenWorld
October 3–5
oracle.com/cxsummit

EVENTS LOCATOR

Oracle Events
oracle.com/events

Locate User Groups
bit.ly/IjxKyc

GETTY IMAGES

What's New at Oracle

The latest videos, podcasts, blogs, and more

WEBCASTS

"2012 OAUG Survey on Cloud Applications"
▶ bit.ly/HZUBic
Hear about the key findings from this recent Oracle Applications Users Group (OAUG) survey.

Webcast Series: *Social Business Thought Leaders*
▶ oracle.com/goto/socialbusiness
This series features perspectives about social tools and the changing workplace.

"Real-Time Data Integration for Next-Generation Data Warehousing"
▶ bit.ly/JYyurX
Discover how Oracle Data Integration products reduce complexity and cost.

"Zero Downtime System Upgrades and Migrations Using Oracle GoldenGate"
▶ bit.ly/I9aOqb
Learn how Oracle GoldenGate's data replication capabilities can eliminate downtime during upgrade, migration, and consolidation projects.

"Maximize ROI for Data Warehousing"
▶ bit.ly/ICyr6A
See how to achieve 20 times higher performance with Oracle Data Integrator on Oracle Exadata.

***Oracle WebLogic DevCast* Series**
▶ oracle.com/goto/weblogicdevcast
This Webcast series explores the ways that you can accelerate time to market with Java EE, Oracle WebLogic, and Oracle Coherence.

Webcast Series: *Navigating a Pragmatic Path to Your Applications in the Cloud*
▶ bit.ly/IoENnz
Explore best practices for simplifying IT operations with cloud-based deployment.

Oracle's Systems Webcast Center
▶ bit.ly/J7YvDI
Access dozens of Webcasts on Oracle hardware systems and solutions.

CALCULATOR

Database Storage Management ROI Calculator
▶ bit.ly/Iwso4T
See how much you can save with Oracle Database 11*g* and compression technologies.

RESOURCE CENTERS

Big Data Essentials
▶ oracle.com/goto/bigdata
Learn the essentials of big data, from underlying technology to real-world use cases.

Oracle Linux Online Forum
▶ bit.ly/HOGcSz
Oracle Linux experts and Progressive Insurance's John Dome discuss implementation of Oracle Linux to increase database performance.

Automate Invoice Processing Operations
▶ bitly.com/HepsYf
Learn how businesses such as the Home Depot, Costco, and Texas Industries have benefited from Oracle WebCenter.

Oracle WebCenter
▶ bitly.com/HeCs4u
Find out how your organization can increase collaboration with Oracle WebCenter portals.

SOCIAL MEDIA

Oracle Value Chain Execution
▶ on.fb.me/JGV8Yv
Get the latest on Oracle's transportation, trade, and warehouse management solutions.

VIDEOS

Oracle's Pillar Axiom 600 Quality of Service
▶ bit.ly/JBLNUv
See how to consolidate storage-area network–based storage with Oracle's Pillar Axiom 600.

Oracle WebLogic Server YouTube Channel
▶ bit.ly/IoQLO2
Watch videos about Oracle WebLogic Server 12*c* and Java 7.

Server Consolidation with Oracle SPARC SuperCluster
▶ bit.ly/JGVRsH
See why Oracle SPARC SuperCluster provides an ideal platform for server consolidation.

Oracle Fact Finders, Episode 1: Managing Storage Systems
▶ bit.ly/K74ha1
Learn how to simplify your storage with Sun ZFS Storage Appliance.

Oracle SPARC SuperCluster: An Overview
▶ bit.ly/Idyuat
Get an overview of the technical and business advantages of Oracle SPARC SuperCluster.

Oracle's Pillar Axiom 600 Hybrid Columnar Compression
▶ bit.ly/IgY7Zc
Learn how Oracle's Pillar Axiom 600 Hybrid Columnar Compression delivers efficiency and performance for storage-area network solutions.

Oracle Social Network video series
▶ bitly.com/JlV4JL
Discover how to use Oracle Social Network for effective sales interactions, file sharing, and data access through mobile devices.

Orbitz Consolidates OLTP Applications onto Oracle Exadata
▶ bit.ly/ISrrBl
See how Orbitz delivers better hotel search capabilities with Oracle Exadata.

Oracle Exadata: Extra Performance for Oracle Applications
▶ bit.ly/IL9Ax8
AFG, Garmin, IQNavigator, and KPN rely on Oracle Exadata for improved application performance.

Redcats Improves Performance 6x on Oracle Exadata
▶ bit.ly/IEEmIq
Redcats migrated a mixed workload to Oracle

OVERHEARD

"The main reason that we chose Oracle Exadata was that it . . . had proven itself, and therefore we felt more confident in picking this specific platform to move ahead in our consolidation and be able to meet the increased performance and availability requirements that our clients had."

—Eric Zonneveld, Lead Oracle Architect at KPN, in *KPN Consolidates Critical Business Apps on Oracle Exadata* (bit.ly/JqJLlk)

Exadata to dramatically lower cost while improving response time sixfold.

KPN Consolidates Critical Business Apps on Oracle Exadata
▶ **bit.ly/JqJLlk**
KPN uses Oracle Exadata for performance, high availability, consolidation, and quick deployment of critical business applications.

Confidence in the Cloud: Tour Oracle's Global Data Centers
▶ **bit.ly/JBKsgx**
Take a virtual tour of Oracle's data centers with Tom Fisher, Oracle Cloud Services CIO and senior vice president.

WHITE PAPERS

"Oracle's Sun Fire X4170 M3 and Sun Fire X4270 M3 Server Architecture"
▶ **bit.ly/JGeZXP**
Learn why Oracle's x86 Sun Fire M3 systems are the best x86 platform for running Oracle software.

"Sun Blade 6000 Server Module Architecture"
▶ **bit.ly/Ibo7WC**
Read about Oracle's Sun Blade 6000's open modular architecture and multitier blade portfolio for deploying a broad range of applications.

"Enterprise-Grade Cloud Applications: The Benefits of Cloud Plus the Power of Oracle"
▶ **bit.ly/I9HBKN**
Learn how Oracle's cloud strategy delivers the benefits of the cloud, without compromise.

"Integrate for Insight"
▶ **bit.ly/JqHQxl**
Discover how you can combine big data with traditional data management.

"Managing Data with Cloud Computing"
▶ **bit.ly/JzU1eO**
Read how cloud computing can bring flexibility, scalability, and speed to your data center.

PODCASTS

"Immutable Zones in Oracle Solaris 11"
▶ **bit.ly/ImZkcV**
Learn about a new feature in Oracle Solaris 11 that provides read-only root zones for highly secure deployment scenarios.

"EnerNOC Tackles Energy Data Volume Growth with Oracle Exadata"
▶ **bit.ly/JGUOZE**
EnerNOC uses Oracle Exadata to boost online transaction processing and data warehouse performance and maximize I/O efficiency.

"Oracle and Cloudera Working Together on Big Data"
▶ **bit.ly/I66hii**
Read why Oracle is using Cloudera's distribution including Apache Hadoop and how the two companies are working together.

BLOGS

Oracle Supply Chain Management Blog
▶ **blogs.oracle.com/scm**
Get the inside scoop on all things supply chain management.

ORACLE UNIVERSITY

Oracle Fusion Applications Training
▶ **bit.ly/xHaCPA**
Learn how to efficiently and effectively implement and use Oracle Fusion Applications.

Prepare for Your Certification Exam
▶ **bit.ly/Jortnk**
Fast-track your exam preparation with on-demand exam prep training seminars for SQL, Oracle Database 11*g*, and Java certification exams. Save 20 percent when you purchase the training and certification exam voucher together.

E-BOOKS

Archiving for Dummies
▶ **bit.ly/JSaniJ**
Learn how to meet the stringent regulatory mandates that require long-term preservation of critical data.

Server Virtualization for Dummies
▶ **bit.ly/IbEUbt**
Discover how getting virtualization right can help you slash costs, accelerate application deployment, and gain business agility.

Introducing the New SPARC SuperCluster T4-4
▶ **bit.ly/KeSkzd**
Learn how Oracle's SPARC SuperCluster T4-4 is designed, tested, and optimized to handle the most-complex multitier workloads.

Introducing the Next Generation of SPARC Servers
▶ **bit.ly/IsIprV**
Read about the benefits of Oracle's SPARC T4 servers and SPARC SuperCluster T4-4.

WEB LOCATOR

Oracle Blog Center
oracle.com/blogs

Oracle Consulting
oracle.com/consulting

Oracle Events and Webcasts
oracle.com/events

Oracle Podcast Center
oracle.com/podcasts

Oracle Newsletters
oracle.com/newsletters

Oracle University
oracle.com/education

Oracle on Facebook
facebook.com/oracle

Oracle on Twitter
twitter.com/oracle

Oracle Support
oracle.com/support

My Oracle Support
myoraclesupport.com

My Oracle Support Communities
communities.oracle.com

Select Specialized Partners

ORACLE PARTNERNETWORK

70+ Oracle Product Specializations

Skills Certified by Independent Testing

Successful Implementations Confirmed by Customers

3,000 Partners with 70,000 Specialists

ORACLE

oracle.com/specialized

DEMOS

Oracle Linux Animated Overview

Watch this demo of Oracle Linux and Ksplice to understand why Oracle Linux is the best Linux for the enterprise.

▶ bit.ly/Hhlxcc

A Quick Introduction to JavaFX

This demo provides a quick overview of JavaFX, the JavaFX scene graph, and a simple JavaFX application. It is the first demo in the JavaFX Getting Started Demo Series.

▶ bit.ly/KvLYLs

Part 1: Building Your First JavaFX Application

With this demo—Part 1 in a series with Oracle JavaFX evangelist Angela Caicedo—start learning how to create a simple JavaFX application using NetBeans 7.

▶ bit.ly/Kh13D1

Oracle Advanced Analytics and Oracle Data Mining Demonstration

This one-hour-plus session focuses primarily on the data mining component of Oracle Advanced Analytics. Topics include big data and big data analytics, data mining, and typical use cases.

▶ bit.ly/J1vZ9P

TUTORIALS

Oracle Platform Security Services 11*g* Release 1: Lesson 1: Overview

This tutorial provides an introduction to security industry terms and concepts and puts them into context for Oracle Platform Security Services. Get an overview of Oracle Platform Security Services, its architecture, and usage benefits.

▶ bit.ly/J7xvZO

Oracle WebLogic Server 12*c*: Configuring and Using Production Redeployment

This tutorial shows you how to use the Oracle WebLogic Server 12*c* production redeployment feature to simultaneously deploy two versions of the same application.

▶ bit.ly/IKcZPM

Oracle Data Integrator Application Adapter for Hadoop

Learn how to use the Oracle Data Integrator Application Adapter for Hadoop. Topics include an overview of big data and the Oracle Data Integrator application adapter.

▶ bit.ly/Kh6YI4

Oracle Unveils Sun ZFS Backup Appliance for Oracle Engineered Systems

Oracle has introduced the Sun ZFS Backup Appliance for Oracle engineered systems, including Oracle Exadata Database Machine, Oracle Exalogic Elastic Cloud, and Oracle's SPARC SuperCluster T4-4. Tested, validated, and supported with Oracle engineered systems, the Sun ZFS Backup Appliance delivers up to 20 terabytes per hour full backup and up to 9.4 terabytes per hour full restore throughputs—the fastest published recovery rates among general-purpose storage systems for Oracle engineered systems data protection.

The Sun ZFS Backup Appliance helps reduce the risk of failed restores by optimizing availability with predictive self-healing features and a fault management architecture that automatically detects and diagnoses underlying problems. End-to-end data checksumming also automatically checks and repairs corrupted data down to the bit level.

"Business operations are highly dependent on database availability, making fast backup and recovery an enterprise imperative," says Phil Bullinger, senior vice president of storage at Oracle. "Oracle's Sun ZFS Backup Appliance is purpose-built to work with Oracle engineered systems, taking advantage of hardware and software engineered together to deliver the speed, simplicity, and savings customers need to protect their data and meet stringent recovery time objectives."

▶ bit.ly/IsOF3b

New Oracle BI Applications for Manufacturing and Asset Management

Oracle has delivered two new innovations to its business analytics portfolio: Oracle Manufacturing Analytics and Oracle Enterprise Asset Management Analytics. The industry's first and only in-memory analytic applications for manufacturing and asset management analysis, the new offerings enable the manufacturing, energy production, utilities, and other asset-intensive industries to gain better insights and make quick decisions.

Both solutions provide support for large numbers of concurrent users and are certified to run on Oracle Exalytics In-Memory Machine. Native support for tablet and mobile devices, such as the Apple iPad and iPhone, is provided.

"For manufacturing and asset-intensive industries, maintaining the health of critical assets and optimizing production and inventory is the difference between bottom-line business growth and failure," says Gaurav Rewari, vice president of enterprise performance management and business intelligence product development at Oracle. "With Oracle Manufacturing Analytics and Oracle Enterprise Asset Management Analytics, we are enabling customers to get intelligence into the hands of key decision-makers wherever the job may take them."

▶ bit.ly/ItHFjs

New MySQL 5.6 Development Milestone Release

A new Development Milestone Release for MySQL 5.6 delivers enhanced availability, increased performance, and improved manageability for Web, cloud, and embedded applications. MySQL 5.6 is available now for download in the MySQL Developer Zone (bit.ly/JYW1xF).

In addition, Oracle is providing early access to significant features under development for community testing and feedback through labs.mysql.com. These include online operations for ADD index and NoSQL access to InnoDB via the Memchached protocol

"Working with the user community, Oracle continues to lead MySQL innovation, delivering new and anticipated features and enhancements," says Tomas Ulin, vice president of MySQL engineering at Oracle.

▶ bit.ly/JYW1xF

Oracle and Infosys Improve Supply Chain Performance for High Tech

Oracle and Infosys, an Oracle Diamond Partner, have unveiled Oracle's project-driven supply chain for high technology companies to reduce project-related inventory and improve financial forecast accuracy. The solution is designed to increase supply chain performance on a single, unified platform. Oracle's project-driven supply chain for high technology companies also enables organizations to increase installation projection efficiencies by delivering increased visibility into global resources, automate complex billing processes, and integrate project management and supply chain management to provide real-time "what if" scenarios.

"[Oracle's project-driven supply chain for high technology companies] leverages Infosys' industry-leading technology business and systems integration expertise with Oracle's integrated applications," says Ravi Kumar, global head, Consulting and System Integration Manufacturing Group, at Infosys. "This solution ties together multiple operational layers to help achieve increased precision, financial controls, and cost improvement across the supply chain."

bit.ly/IKqjnh

Oracle Database Now Available on Oracle Linux 6

Oracle Database 11*g* Release 2 and Oracle Fusion Middleware 11*g* Release 1 are now available on Oracle Linux 6 with the Unbreakable Enterprise Kernel as well as on Red Hat Enterprise Linux 6 and Oracle Linux 6 with the Red Hat Compatible Kernel. "The Oracle Database and Oracle Fusion Middleware product availability is the result of stringent certification testing on Oracle Linux 6 with the Unbreakable Enterprise Kernel," says Wim Coekaerts, senior vice president of Linux and virtualization engineering at Oracle.

oracle.com/us/technologies/linux

New Sun x86 Servers Deliver the Industry's Best x86 Platform

Oracle has unveiled a new line of Sun x86 servers based on the Intel Xeon processor E5-2600 product family of CPUs. The servers offer the industry's best x86 platform, with unrivaled performance and reliability for applications and databases, while also offering the tools needed for cloud-enabled deployments including operating system, virtualization, infrastructure provisioning, and system management at no extra cost when combined with Oracle Premier Support services.

The new Sun x86 servers deliver up to 87 percent better computational performance and feature increased memory and network bandwidth over the previous generation, enabling organizations to run additional software and complete tasks faster in the same physical footprint. Running Oracle Solaris or Oracle Linux, Oracle VM, and Oracle Enterprise Manager Ops Center, the new servers reduce costs, streamline data center operations, and offer better performance and reliability. The new servers have set multiple world-record benchmarks, demonstrating superior performance on computational, application, and enterprise workloads, including new world records on SPEC CPU2006, SPECjEnterprise2010, and Oracle E-Business Suite Release 12.1.3 (bit.ly/J7J7vS).

"Oracle's collaboration with Intel provides solutions with great performance and scalability to meet the ever-increasing demands of data center customers," says Doug Fisher, corporate vice president, Software and Services Group, at Intel. "Extensive co-engineering, testing, and validation have gone into creating an excellent balance between processor, memory, and I/O resources in the solution stack. As a result, Oracle's x86 servers based on the Intel Xeon processor E5-2600 offer an unprecedented ability to move large volumes of data efficiently and handle multiple millions of transactions per second."

bit.ly/IMOmSE

I-HUA CHEN

Oracle Releases JD Edwards EnterpriseOne 9.1

The latest release of Oracle's integrated suite of enterprise resource planning software, Oracle's JD Edwards EnterpriseOne 9.1, helps customers further streamline business processes with a personalized user experience, new industry-specific functionality, productivity features, and more support for global operations and compliance mandates. And with the new One View Reporting feature, business users can create interactive reports without IT support.

"This important new release builds on our long-standing commitment to delivering an integrated, standards-based portfolio of innovative applications that can help streamline processes, reduce costs, and drive business value," says Lyle Ekdahl, group vice president and general manager, Oracle's JD Edwards.

bit.ly/xsm7UX

AutoVue 20.2 Supports Better Decision-Making

The latest release of Oracle's family of industry-leading enterprise visualization solutions, Oracle's AutoVue 20.2, is now available. Oracle's AutoVue 20.2 helps create rich and actionable visual decision-making environments on all delivery platforms with augmented business visualization. This framework connects portions of documents with business data from enterprise applications for more-informed decisions.

The new release also supports mobile workers in the utilities and oil and gas industries. AutoVue 20.2 can be integrated with Oracle's enterprise applications to enable a range of mobile devices by leveraging Oracle Virtual Desktop Infrastructure. Mobile support gives field technicians access to work order–related documents and lets them capture a digital record of onsite repair work.

"New capabilities introduced in AutoVue 20.2, such as augmented business visualization, reconcile information from documents and the rest of the Oracle stack, helping companies access critical information in the format they need to make informed business decisions," says Thierry Bonfante, senior director of AutoVue product management at Oracle.

bit.ly/eWCwYO

Oracle Introduces Two Business Intelligence Applications for SAP

To give organizations greater business insight into their SAP systems, Oracle has unveiled Oracle Procurement and Spend Analytics for SAP and Oracle Supply Chain and Order Management Analytics for SAP. Available as part of Oracle Business Intelligence Applications Release 7.9.7.1, these new applications help reduce procurement and spend costs and improve supplier performance and supply chain efficiency.

These business intelligence (BI) applications for SAP are among the 80-plus Oracle BI and enterprise performance management applications available today that do not require any application changes to run on Oracle Exalytics In-Memory Machine.

"Oracle is pleased to deliver the first packaged analytic applications built specifically for SAP customers that run on an engineered system for analytics and are fully enabled for mobile deployment," says Paul Rodwick, vice president of product management, Oracle Business Intelligence.

bit.ly/ItHFjs

RightNow CX Cloud Service and Oracle Fusion Sales Integrated

Oracle RightNow CX Cloud Service is now integrated with Oracle Fusion Sales. This combination allows sales and marketing organizations to leverage service interactions from Oracle RightNow CX Cloud Service, a customer experience suite, and sales prediction and segmentation capabilities from Oracle Fusion Sales. The offering enables organizations to better match products and services to specific customer needs based on customer service history; deliver targeted, personalized customer interactions; and identify new opportunities to increase deal size and conversion rates.

The integration provides organizations with a comprehensive view of all customer interactions across channels and the context and status of these interactions. The offering includes tools that enable sales and service departments to collaborate to proactively solve customer issues, sales representatives to review service history in preparation for sales calls, and agents to understand customer value based on prior buying habits and existing opportunities.

"Every interaction is a relationship opportunity to grow your business. When these interactions are relevant and add value for customers, customers are more likely to trust the relationship and seek purchase advice," says David Vap, group vice president at Oracle.

bit.ly/L5bbgH

New StorageTek Tape Analytics Delivers Intelligent Monitoring

Oracle has introduced Oracle's StorageTek Tape Analytics, intelligent monitoring software for StorageTek tape libraries. StorageTek Tape Analytics simplifies tape management by eliminating manual intervention and helps improve operational efficiency. The new solution captures library, drive, and media performance metrics and presents them in a single console, and can proactively identify potential problems while providing immediate recommendations to prevent data loss. StorageTek Tape Analytics also allows rapid drill-down into detailed drive and media health specifics, enabling customers to proactively make decisions about their tape infrastructure and reduce downtime and data management resources.

"Oracle's StorageTek Tape Analytics software allows customers better insight into the health of their tape environments and arms them with the information to maximize the value of their storage investments," says James Cates, vice president of hardware development at Oracle. "Ideal for backup and archive data residing locally or in the cloud, Oracle's [StorageTek Tape Analytics] is secure, efficient, and cost effective. Its savings for customers can be quite substantial, as it does not require special media but rather works with the customer's existing supported media, regardless of the vendor."

bit.ly/IPqOHD

JDK for Mac OS X Available

Oracle has released its first Java Development Kit (JDK) and JavaFX Software Development Kit for Mac OS X. Java developers can now download Java Platform, Standard Edition 7, Update 4 (Java SE 7 Update 4) and JavaFX Software Development Kit 2.1 for Mac OS X from Oracle Technology Network.

The Java SE 7 Update 4 SDK includes the next-generation Garbage Collection algorithm, Garbage First (G1), which has been highly anticipated by the Java developer community.

JavaFX 2.1 introduces playback support for digital media stored in the MPEG-4 multimedia container format containing H.264/AVC video and Advanced Audio Coding (AAC) audio. It also includes WebView support for JavaScript to Java method calls.

"We look forward to delivering simultaneous releases of the JRE [Java Runtime Environment] across all major operating systems later this year, so all Java users will be able to take advantage of the latest features and security fixes," says Hasan Rizvi, senior vice president of Oracle Fusion Middleware and Java products at Oracle.

bit.ly/qqM4AV

Oracle Acquires ClearTrial

Oracle has entered into an agreement to acquire ClearTrial, a provider of cloud-based clinical trial operations applications that make the planning, sourcing, and tracking of clinical projects and financial performance faster and more accurate. ClearTrial's solutions help life sciences companies manage the complexity of bringing new therapies to trial.

"Adding ClearTrial to the Oracle Health Sciences Cloud will help our customers streamline the clinical development process and help them bring therapies to market with greater predictability and at lower costs," says Neil de Crescenzo, senior vice president and general manager, Oracle Health Sciences.

bit.ly/IrZPmX

I-HUA CHEN

The Business of Analytics

Oracle unveils analytics strategy and releases new EPM and analytics applications, as well as Oracle Endeca.

Oracle President Mark Hurd likes to talk about business intelligence and analytics. Just ask the crowd at Oracle OpenWorld Tokyo, where Hurd discussed those topics—and Oracle's strategy around them—in a keynote address on April 4, 2012. At the same time, several new Oracle products that enable faster, better decision-making were introduced.

To derive benefit from analytics, a company needs systems that can both handle huge data volumes and provide very fast processing—and then they must put that power in the hands of people who need it, Hurd explained. "The key to analytics," says Hurd, "is to get through tons of information, in a fast way, with a lot of performance—to be able to ask really hard questions and get a decision to somebody who can take action."

Hurd added that today's emphasis on business analytics is generational. "I'm used to taking a long time to get information," he said, "but my kids aren't. My kids are used to instant answers to whatever question they've got, whenever they want, and they won't wait."

Joined by Balaji Yelamanchili, senior vice president of analytics and performance management products at Oracle, Hurd outlined the challenges businesses face in implementing analytics—including eight times (or faster) data growth in the last five years, aging infrastructures and applications, and the need to deliver answers to mobile platforms—and the types of solutions Oracle provides to meet these challenges.

Oracle is spending almost US$15 billion in research and development over three years along four lines, Hurd said: Creating best-of-breed hardware, operating systems, databases, and applications, including business intelligence and analytics applications; vertically integrating those components for extreme performance—for example, creating engineered systems built for analytics; creating or acquiring industry-specific solutions; and producing a new generation of applications, built on common middleware, that can be delivered via public or private clouds or on-premises data centers.

Oracle President Mark Hurd

To showcase the early fruits of its R&D investment in business analytics, Oracle announced several new or updated products, including a new release of Oracle's enterprise performance management (EPM) system, general availability of Oracle Endeca Information Discovery, and several new analytic applications. "The next generation of analytic systems is so important because the problem [faced by companies in today's environment] is getting harder," Hurd said. "Because there's more data."

EPM ENHANCED

Release 11.1.1.2 of Oracle's EPM system delivers new modules and capabilities that improve performance and provide companies with greater insight into their projects.

Oracle Hyperion Project Financial Planning, a new product in Release 11.1.1.2 of Oracle's EPM system, enables project-based budgeting, funding, approvals, and tracking. This extends financial planning and forecasting—which in the past typically focused on accounts and cost centers—to activities organized as projects. Project Financial Planning can be used both for projects that occur in all types of business (such as IT, HR, or facilities projects) and for entire operations in project-centric industries, such as aerospace and defense, construction, and manufacturing.

Account Reconciliation Manager, a new feature of Oracle Hyperion Financial Close Management, improves both the speed and accuracy of period-end close activities while reducing reporting risks. Prebuilt components and starter kits help companies comply with International Financial Reporting Standards (IFRS), eXtensible Business Reporting Language (XBRL), and sustainability reporting requirements.

The speed and scalability of Oracle Hyperion Planning—an EPM product that integrates financial and operational planning processes—is significantly boosted in the new EPM system release through support for Oracle Exalytics In-Memory Machine, an engineered system designed for analytics. Usability, performance, and accuracy of Oracle Hyperion Planning are also enhanced through a new Web 2.0–style user interface, built-in predictive analytics based on Oracle Crystal Ball, and built-in integration with Oracle Fusion Financials, Oracle's JD Edwards EnterpriseOne Financials, and the Commitment Control feature of Oracle's PeopleSoft Financials. Additionally, Oracle Hyperion Data Relationship Management, an EPM product used to build and retain consistency within master data, now includes out-of-the-box integration with ERP datasources including

HARTMANN STUDIOS

Oracle E-Business Suite 12.1 and Oracle Fusion Financials.

ORACLE ENDECA UNVEILED

Another product of Oracle's R&D investment in business analytics, Oracle Endeca Information Discovery, is the first Oracle information discovery offering since Oracle acquired Endeca in October 2011. Oracle Endeca Information Discovery enables exploration and analysis of structured, semi-structured, and unstructured data from a variety of sources such as data warehouses, transaction systems, social media, and sensor data. These datasources can reside inside or outside of an organization, and might include sources that have never been modeled and that are constantly changing.

Information discovery, also called data discovery, complements and enhances business intelligence (BI). Whereas BI provides proven answers to known questions, data discovery provides fast answers to new questions and puts the power to ask those questions in the hands of businesspeople who know the information best. Data discovery helps companies create new key performance indicators for BI systems.

Oracle Endeca Information Discovery uses a faceted analytical model, similar to that used in online shopping. A faceted model encourages data exploration by presenting all relevant data in a straightforward click-to-select manner. (Compare it to a shoe-shopping Website that lets users narrow a selection: first they might choose men's shoes, then sneakers, then white leather, then size 9. Now extend that concept to searches for commonalities and relevant themes in multiple corporate datasources, such as ERP, customer relationship management, and HR management systems, as well as unmodeled and external datasources such as social networking sites.)

Oracle Endeca Information Discovery "takes you from traditional reporting and analysis to interactive discovery and exploration—where you don't even know all the questions you're going to ask," Yelamanchili said at the Oracle OpenWorld Tokyo launch. "You're going to start with a question, and then move very quickly into understanding, connecting the dots, finding the patterns, and finding the outliers. And very quickly, from that point on, zero in on the real issues and real problems."

> **"Business analytics—getting the right data to the right person at the right time to make the optimal decision—is the most important thing you can get from your data."**
>
> *—Mark Hurd, President, Oracle*

Oracle Endeca Information Discovery's three-tier architecture works like this: After various datasources are combined and enriched in the Integration Suite feature of Oracle Endeca Information Discovery (the first tier), Oracle Endeca Server (the second tier) organizes that information into a unified index, deriving a model from the data if necessary—a bottom-up approach that is the core of the Endeca application line. The Studio feature of Oracle Endeca Information Discovery (the third tier), with its drag-and-drop interface, creates applications for exploring and analyzing that information. In the process of information discovery, Oracle Endeca's optimization algorithms exploit the server's full memory hierarchy, from on-CPU cache to disk, to optimize system performance.

ANALYTICS APPLICATIONS EXTENDED

Oracle has also extended its current family of applications to support the need for analytics. The Tokyo announcement highlighted new offerings in this area, including Oracle Procurement and Spend Analytics for SAP, Oracle Supply Chain and Order Management Analytics for SAP, enhancements to Oracle Financial Analytics for SAP, Oracle Manufacturing Analytics, and Oracle Enterprise Asset Management Analytics. (See Briefs, page 16, for details.)

These new applications, as well as Release 11.1.1.2 of Oracle's EPM products and Oracle Endeca Information Discovery, support Oracle's overall business analytics strategy. As described by Hurd and Yelamanchili, the strategy recognizes trends such as big data, mobility, in-memory analytics, and cloud, and is built on four foundational principles:

Any data, any source. Deliver the ability to tap into diverse data sets—structured, unstructured, social media, and machine-generated data—to improve analytics.

A full range of analytics. Provide software and engineered systems purpose-built for the full range of analytics—including reporting, modeling, planning, and predictive analysis—that democratize analytics and allow anyone in an organization to be a data artisan.

Integrated analytic applications. Support those customers who want to create their own analytic applications, but provide out-of-the-box analytic capabilities tied to business processes that allow people to act on what they discover.

On premises, on cloud, on mobile. Deliver analytics via flexible models in any way that customers—both IT professionals and end users—want them. This includes traditional data center solutions and software-as-a-service (SaaS)/cloud platforms, and on both traditional screens and mobile devices, without compromising the usability or power of the analytics.

"Business analytics—getting the right data to the right person at the right time to make the optimal decision—is the most important thing you can get from your data," Hurd told the Oracle OpenWorld Tokyo audience. "Analytics is in Oracle's DNA and is a vital part of everything we do—database, middleware, applications, and engineered systems." ◀

Fred Sandsmark is a freelance technology and business writer.

NEXT STEPS

WATCH the event Webcast
bit.ly/JzXgPY

LEARN more about enterprise performance management
bit.ly/KbtjK0

Oracle Endeca Information Discovery
bit.ly/IRbkqX

Virtually Everywhere

Oracle Virtual Desktop Infrastructure 3.4 enables secure desktop access from almost anywhere.

Following the availability of Oracle Virtual Desktop Infrastructure 3.4, Rich Schwerin, Oracle Magazine *contributor, sat down with Andy Hall, product management director at Oracle, to talk about desktop virtualization. The following is an excerpt from that interview. Download the full podcast at oracle.com/magcasts.*

Oracle Magazine: How does Oracle desktop virtualization deliver enterprise desktops?

Hall: Oracle desktop virtualization offers a different approach to delivering enterprise desktops to end users. Don't run the desktop on the client machines, where it's insecure, where it's tough to manage, where basically you're tied to a device. Instead, why don't you use virtualization to run your desktop in a virtual machine in the data center? It's more secure, it's easier to manage, and you can access it from more places and from more devices.

Oracle Magazine: Which products comprise Oracle's desktop virtualization solutions?

Hall: We have two desktop virtualization classes: client-side and server-side. A product called Oracle VM VirtualBox handles client-side virtualization. Oracle VM VirtualBox is one of the most popular client-side hypervisors in the world today. There are more than 2 million downloads of Oracle VM VirtualBox happening from Oracle Technology Network every month, and people are using Oracle VM VirtualBox on their Microsoft Windows PCs, on their Mac desktops, and on their Linux machines as well. The big brother to Oracle VM VirtualBox is Oracle Virtual Desktop Infrastructure—which is server-side desktop virtualization—running desktop operating systems such as Windows in individual VMs on servers in the data center. And if you're running your desktops from the data center, you can access them from your Windows PC, from your Mac, from your iPad, or from specific devices, such as Oracle's Sun Ray Clients. Sun Ray Clients are an endpoint

JOHN BLYTHE

Andy Hall, Product Management Director, Oracle

device that are specifically designed for accessing virtual desktops in the data center, so they have no operating system, no moving parts, and are really secure. Finally, we have Oracle Secure Global Desktop, which provides secure browser-based access to your applications and data running in the data center from wherever you are in the world. The idea is you start with a browser, type in a URL, log in, and you get access to the Oracle applications and other enterprise applications that are in the data center.

Oracle Magazine: How do system administrators benefit from Oracle desktop virtualization?

Hall: If you've ever managed a fleet of desktops, you've realized that you're in a very reactive mode. You're forever chasing around and looking to update people's PCs, you're fixing security problems, you're helping end users who have gotten themselves into a bit of a predicament somehow—and that's because the device that you're trying to help them with is out there in the wild. Effectively many users are their own administrators in a traditional PC environment. But if you move to a virtual desktop environment, the administrator is able to deliver virtual desktops from the data center. And the way this works is that the administrator would create one or more pools of desktops with a golden template—a repeatable, reusable reference model. So you can create, for instance, the golden master Windows template for your sales team, and another golden master template for your marketing people or your engineers, and they can be completely different golden images. Oracle Virtual Desktop Infrastructure can clone these golden master images into desktop pools that you can assign as desktops to end users. So when I as your end user log in, I am assigned a particular desktop that the system administrator wants for me. And when I log out, Oracle Virtual Desktop Infrastructure has a selection of features that allows the administrator to either keep that virtual desktop for me, repurpose or reassign it to someone else, or potentially just throw it away altogether and clone a brand-new one. This means that you can always re-create and deploy pristine desktop images to your end users, so you don't let them get into a mess.

Oracle Magazine: How do developers and DBAs benefit from Oracle desktop virtualization?

Hall: Most developers I know tend to use Oracle VM VirtualBox on the client side, because one of the things developers really want to do is to be able to work from anywhere and to be able to do everything. And being able to work from anywhere means that you want a complete development environment of clients, a middle-tier server, and a back-end server, so you can create multiple VMs on your desktop device, whatever that might be. Let's say you're a Web developer or you're developing something in HTML5. You need to test that against all the different

browsers, so you create multiple VMs to test against Mozilla Firefox, Microsoft Internet Explorer, or others, and on all the different platforms that you expect your customers to be using as well. By using virtualization in that way, you really can simulate almost every environment in which you expect your application to be used.

Oracle Magazine: What role does Oracle desktop virtualization play when employees bring their own tablet devices to work?

Hall: One of the big problems that system administrators are having is that not just any old end user, but the C-level executive is coming to work and saying, "Hey, I'm the chief financial officer, and I expect to be able to access my enterprise applications from this device." How on earth can you do that? Typically the applications haven't been designed for that device, so how can you deliver your standard enterprise applications to such a device? Well, it's very easy to do with Oracle Virtual Desktop Infrastructure. You simply assign your CFO a virtual desktop that runs in the data center, and then install the Oracle Virtual Desktop Client for iPad. You start that up, and that is able to find your virtual desktop server, authenticate, and log in, and the CFO can access the desktop on his iPad. And we've gone further than that, because we realize that the gestures you use when you're using a tablet device are very often different from the ones you use when you are using a keyboard and mouse. So we've mapped a certain number of those gestures to make it more natural to be able to interact with this traditional desktop environment on your tablet. Plus, you can hop between devices and locations and still get access to the same virtual desktop. So let's say you're sitting in the office and using your virtual desktop, and all of a sudden you get a telephone call that says you have to come home for dinner now. So you dash home, and when you're at home you pick up your iPad, connect back into your virtual desktop server, and pick up the session from exactly where you left it.

Oracle Magazine: How does Oracle address desktop virtualization security issues?

Hall: The good news is that none of the data you're interacting with from your Oracle Virtual Desktop Client on your PC, Mac, iPad, or Sun Ray Client exists on the client device—you're really just seeing a view onto that data. So if someone steals the device, they really haven't stolen anything. Your desktop is still running in the data center, where it's being looked after; where it's tightly secured; and where you can keep audit trails of who accessed it, from where, when, and for how long.

> **"One of the things developers really want to do is to be able to work from anywhere. . . . And [that] means that [they] want a complete development environment."**
>
> *—Andy Hall, Product Management Director, Oracle*

Oracle Magazine: What's new in Oracle Virtual Desktop Infrastructure 3.4, and what makes it important?

Hall: Traditionally, putting together a virtual desktop infrastructure solution from any vendor is a little bit hairy. You need quite a lot of capital expenditure: you have to get some shared storage devices, such as a NAS [network-attached storage] appliance; you need some servers on which you're going to run your virtual desktops; you need some software that's going to virtualize the desktop; and you need some software that's going to allow you to get remote access from it securely over a protocol that has been designed specifically for accessing remote desktops.

One of the things that we're doing with Oracle Virtual Desktop Infrastructure 3.4 is to address some of those challenges by offering more and simpler deployment options. For example, we're offering a greater range of storage options, so you don't necessarily need to go out and buy a big NAS device. You can use either local storage or shared storage such as NFS [Network File System] or SMB [Server Message Block] to set up the virtual desktop infrastructure. The other area that we've improved on greatly is the delivery of a richer user experience, because if you're using your virtual desktop, you really want it to be as good as the experience you have when you use a traditional PC. For example, if you go to YouTube, you'll see rich media types inside of your browser, and you expect them to be displayed from your virtual desktop in a way that is comparable to a traditional desktop deployment. To achieve that, we've put in acceleration techniques at various places in the stack. So, for instance, when Oracle Virtual Desktop Infrastructure 3.4 sees video being displayed in a virtual desktop, it uses a different codec to transmit it down to the client. And as you jump from device to device, we've made the *hotdesking*, or session mobility, a lot smoother and smarter. Oracle Virtual Desktop Infrastructure 3.4 also passes the location of the endpoint up the wire to the virtual desktop, and that's very important in deployments such as healthcare.

Oracle Magazine: What makes Oracle's desktop virtualization solutions unique?

Hall: First, Oracle is the only virtual desktop vendor that can provide the hardware for your virtualization layer, the storage service, the server virtualization software, the desktop virtualization software, the remote access, and the endpoint device. So we deliver the whole stack—hardware and software, engineered to work together. Second, when you're using Oracle Applications, Oracle Virtual Desktop Infrastructure is the only virtual desktop environment certified for use with those enterprise applications. ◀

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

NEXT STEPS

LEARN more about Oracle desktop virtualization
bit.ly/ovdimag

LISTEN to the podcast
oracle.com/magcasts

Database Appliance or Puzzle Pieces?

BUY		BUILD
Oracle Database Appliance 	**or**	Microsoft Database + HP Server + NetApp Storage + Cisco Networking + Consultants

Database Appliance: Simple, Reliable, Affordable

ORACLE®

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

Oracle BPM Suite 11g Developer's Cookbook

By Vivek Acharya

Packt
packtpub.com

Oracle BPM Suite 11g Developer's Cookbook will help readers learn how to develop rich, interactive business processes using Oracle Business Process Management Suite 11*g*. A common process model based on Business Process Model and Notation is presented to the specific role assigned to readers in each chapter. The book is divided into four sections: modeling; implementation; measuring; and deployment, migration, and runtime. Each chapter offers practical, task-based recipes.

Effective MySQL: Backup and Recovery

By Ronald Bradford and Patrick Galbraith

Oracle Press
oraclepressbooks.com

Effective MySQL: Backup and Recovery offers a no-nonsense approach to one of the most critical tasks of an operational DBA—performing, testing, and verifying backup and disaster recovery procedures. It helps DBAs understand the limitations and quirks associated with MySQL backups. This is key to ensuring that the important business information that has been backed up and the method used to recover it will meet business needs. The book drills down into real-world problems and provides succinct, proven formulas for achieving faster backups and gaining greater overall understanding of the systems—regardless of the DBA's experience level. The book includes a "Five Minute DBA" chapter that presents a real-world scenario and the correct way to manage it.

Getting Started with Oracle Data Integrator 11g: A Hands-On Tutorial

By David Hecksel, Bernard Wheeler, Peter C. Boyd-Bowman, and Julien Testut

Packt
packtpub.com

Getting Started with Oracle Data Integrator 11g: A Hands-On Tutorial offers readers all the information they require to get up and running with Oracle Data Integrator, Oracle's data integration platform for high-speed data transformation and movement between different systems. Following an example scenario, this book covers essential information about the Oracle Data Integrator architecture. It provides details about using the product across different databases and file types such as XML. Other topics include orchestrating data integration workflows, error management, operational management and monitoring, and more. This book is a resource for extract, transform, and load developers; software developers; and DBAs who want to get hands-on experience with Oracle Data Integrator quickly.

Developing Essbase Applications: Advanced Techniques for Finance and IT Professionals

Edited by Cameron Lackpour

Auerbach Publications
crcpress.com

Learn advanced techniques for building analytical models, reporting systems, and forecasting applications using Oracle Essbase. This book, written by some of the most experienced Oracle Essbase practitioners around the world, covers infrastructure, datasourcing and transformation, database design, calculations, automation, APIs, reporting, and project implementation. The book includes practical cases that stem from the authors' decades of combined experience using Oracle Essbase tools and techniques. The authors are committed to showing readers the right way to work with Oracle Essbase. Their motto: "We love Essbase and hate to see it done wrong."

Oracle Advanced PL/SQL Developer Professional Guide

By Saurabh Gupta

Packt
packtpub.com

Master PL/SQL concepts with *Oracle Advanced PL/SQL Developer Professional Guide*. The book starts with an overview of PL/SQL and outlines the characteristics and benefits of the language. It then covers advanced PL/SQL features that include code writing using collections, tuning recommendations using result caching, and enforcing row-level security. The book also dives deep into the use of Oracle SQL Developer, employing best practices in database environments and safeguarding the vulnerable areas in PL/SQL code to avoid code injection.

The book can serve as a study guide for the Oracle Advanced PL/SQL Developer Certified Professional certification. Each major certification topic is covered in a separate chapter, with practice questions included at the end of every chapter.

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

Apps Associates Offers Oracle Exalytics Service and Support

Oracle Platinum Partner Apps Associates has launched an initiative to deliver services and solutions for Oracle Exalytics In-Memory Machine. Apps Associates is committed to bringing Oracle Exalytics to the thousands of customers using Oracle Business Intelligence, Oracle Essbase, and Oracle Hyperion Planning. The company will provide early stage proof-of-concept evaluation, implementation, deployment, and support.

"Apps Associates' commitment to the development and rollout of a program

centered on leveraging the power of Oracle Exalytics underscores the inherent ability of the platform to address the growing need for high-speed, real-time analytics in the most-demanding environments," says Tyler Prince, group vice president, North America Sales, Alliances and Channels, at Oracle.

appsassociates.com/oracle-exalytics

Heiler Earns Validated Integration with Oracle ATG Web Commerce

Heiler Software, an Oracle Gold Partner, has achieved Oracle Validated Integration between Heiler Product Manager 5.3 and Oracle ATG Web Commerce 10. Heiler Product Manager 5.3 is a scalable central master data platform for e-commerce and multichannel businesses. It provides distributors and manufacturers with master data management for all product data in a central datasource.

Heiler's adapter between Heiler Product Manager 5.3 and Oracle ATG Web Commerce 10 enables transmission of information about items, products, structure groups, media assets, and more. The validated integration includes preconfigured export templates, export utilities, and data field validations.

heiler.com

LINDY GROENING

QLogic Supports Newly Launched Oracle x86 Servers

QLogic, an Oracle Gold Partner, supports Oracle's new Sun Fire X4170 M3 server, Sun Fire X4270 M3 server, and Sun Blade X6270 M3 server with its line of Oracle networking products. This includes an 8 Gb FC PCI Express (PCIe) host bus adapter and 10 Gb/sec Ethernet FC over Ethernet (FCoE) PCIe converged network adapters that are also available in ExpressModules for Sun Blade X6270 M3 server.

With more than 1 million ports shipped to Oracle customers, QLogic's Oracle-branded networking adapters provide Sun Fire M3 customers proven FC performance and full hardware offload for FCoE and iSCSI. These performance attributes support the high compute density of Oracle's Sun Fire and Sun Blade servers by delivering superior application performance and I/O scalability in virtual and nonvirtual environments. And, with time-tested QLogic FC, FCoE, and network interface card driver stacks, the highest level of SAN reliability is achieved to ensure the robustness and continuous uptime of Oracle's hardware infrastructure and operating environments.

qlogic.com/go/oracle

Kewill Flagship Achieves Validated Integration with Oracle E-Business Suite

Kewill, a solution provider in global trade and logistics and an Oracle Gold Partner, has achieved Oracle Validated Integration of Kewill Flagship with Oracle E-Business Suite 12.1. Kewill Flagship is a multicarrier enterprise parcel shipping solution for high-performance and high-volume shipping. The validated Kewill Flagship Adapter is a PL/SQL package that enables standard integration between Kewill Flagship 5.0 and Oracle E-Business Suite 12.1.

Kewill Flagship can manage multiple shipping stations and support hundreds of thousands of shipments per day with a single software deployment. Oracle has also developed and markets a standard integration between Oracle Transportation Management and Kewill Flagship, and the two companies have a cooperative marketing agreement.

kewill.com

Megabyte Earns Multiple Specializations

Oracle Gold Partner Megabyte has achieved Specialized status in Oracle Business Intelligence foundation, Oracle Database, and Oracle Financial Management.

Founded in 1979, Megabyte is one of Malta's first business computer systems companies. It offers IT project management, systems analysis, software design, customization, modification, and support, and has many years of experience in the development and implementation of IT solutions.

megabyte.net

Elemica Introduces Transportation Management Solution

Elemica, a provider of supply chain integration and optimization services to the chemical, tire and rubber, energy, and selected manufacturing industries, has introduced the Elemica Transportation Management solution, a BPO Powered by Oracle service based on Oracle Transportation Management. Elemica is an Oracle Gold Partner.

LINDY GROENING

The Elemica Transportation Management solution provides control over, and visibility into, the end-to-end logistics order lifecycle. It incorporates chemical industry and logistics best practices, and its preconfigured processes support both common and anticipated needs of the chemical industry.

elemica.com

Omnix Achieves Platinum Partner Status

Omnix International, a distributor and systems integrator for digital software imaging solutions based in Dubai, United Arab Emirates, has become an Oracle Platinum Partner. The status recognizes Omnix for its expertise in delivering Oracle's Primavera P6 Enterprise Project Portfolio Management solutions. To earn Platinum status, partners must achieve any combination of five specializations, Oracle Exastack Ready status, or Oracle Validated Integrations.

With its Platinum status, Omnix International will now be able to broaden specialization areas; receive priority placement in the Oracle PartnerNetwork solutions catalog; and enjoy joint marketing and sales opportunities, discounted training, and more.

omnix.ae

BIAS Achieves Oracle Exalogic Elastic Cloud Specialization

BIAS Corporation, an Oracle Platinum Partner, is the first Oracle partner to achieve Oracle PartnerNetwork specialization on Oracle Exalogic Elastic Cloud. BIAS designs, acquires, implements, and manages comprehensive IT solutions based on Oracle technologies for companies worldwide.

In achieving a Specialized distinction, BIAS has been recognized by Oracle for its expertise in delivering services specifically around Oracle Exalogic Elastic Cloud through competency development, business results, and proven success.

BIAS also holds the following specializations: Oracle Exadata, Oracle GoldenGate, Oracle Business Intelligence Foundation, Oracle Enterprise Manager, Oracle Application Grid, Service-Oriented Architecture, Oracle WebCenter Portal, Data Warehousing, Oracle Database Performance Tuning, Oracle Linux, Oracle Database, and Oracle Real Application Clusters.

biascorp.com

Transcepta Achieves Validated Integration with Oracle E-Business Suite

Oracle Gold Partner Transcepta has attained Oracle Validated Integration for the Transcepta Procure to Pay (P2P) network with Oracle E-Business Suite 12.1. The Transcepta P2P network enables enterprise customers to electronically collaborate with suppliers through the exchange of invoices, purchase orders, and other documents. Validated integration between the Transcepta P2P network and Oracle E-Business Suite eases integration challenges for IT departments while enabling accounts payable and procurement teams to leverage Transcepta's network.

"A rigorous technical review is required for Oracle Validated Integration," says Ognjen Pavlovic, vice president of product strategy at Oracle. "Since Transcepta's supplier messaging through Oracle Supplier Network has been tested and validated by Oracle, customers can be confident that the Transcepta Network integrates with Oracle E-Business Suite as designed, allowing accelerated time to value for our joint customers."

transcepta.com

LoadSpring Becomes Oracle Gold Partner

LoadSpring Solutions has achieved Gold status in Oracle PartnerNetwork in recognition of its expertise in Oracle's Primavera P6 Enterprise Project Portfolio Management solutions and for addressing the challenges of joint Oracle/LoadSpring customers. LoadSpring resells project management software and offers hosting and implementation services to its customers.

With its Gold status, LoadSpring receives the benefit of being able to start developing specializations that will allow it to grow its business, increase its expertise, reach higher levels of customer retention, and create differentiation in the marketplace. Gold members also become eligible to resell all Oracle technology products and can apply to resell Oracle applications and industry solutions. In addition, they receive access to Oracle account representatives and My Oracle Support updates for all products, discounts on training, limited free assessment/exam vouchers, reduced rates on the purchase of Oracle licenses for internal use, discounts on advanced customer services, and more.

loadspring.com

INFOPRO's eICBA Achieves Two Oracle Ready Statuses

INFOPRO, an Oracle Platinum Partner headquartered in Malaysia, has achieved Oracle Database Ready and Oracle WebLogic Ready status for its eICBA 9 software. Oracle Ready status demonstrates that eICBA 9—a comprehensive, integrated, modular banking solution—has been tested on Oracle Database 11*g* Release 2 and Oracle WebLogic Server 11*g* Release 1. eICBA provides a full spectrum of conventional and Islamic banking operations at retail, commercial, midsize enterprise, and corporate levels, and incorporates a full complement of transactional services.

Oracle Database Ready and Oracle WebLogic Ready are components of the Oracle Exastack Ready program. That program recognizes partners who have developed, tested, and tuned their applications on component parts of the Oracle Exadata Database Machine (pictured above) or Oracle Exalogic Elastic Cloud engineered systems.

infopro.com.my

Cintra Achieves Three Specializations

Cintra Software & Services, an Oracle Platinum Partner, has earned three Oracle PartnerNetwork specializations: Oracle Linux, Oracle Data Warehousing, and Oracle Exadata Database Machine. Cintra also holds Oracle Database 11*g*, Oracle Real Application Clusters, Oracle Database Performance Tuning, Oracle GoldenGate 10, and Oracle Business Intelligence Foundation Suite specializations.

Cintra's achievement of multiple specializations demonstrates its understanding of how various Oracle solutions can work together to create value across an organization. Cintra provides database infrastructure services including design, implementation, and ongoing remote DBA support.

cintra.com

Avnet Distributes Oracle Healthcare Information Exchange and Analytics in U.S. and Canada

Avnet Technology Solutions is now authorized to distribute Oracle Health Sciences solutions in the U.S. and Canada. Those solutions include Oracle's healthcare information exchange and Oracle enterprise healthcare analytics. Oracle Health Sciences has authorized Avnet to distribute these specialized solutions because of Avnet's proven ability to develop and foster an ecosystem of highly skilled healthcare IT solution provider partners through its SolutionsPath methodology and HealthPath practice.

Under this agreement , Oracle Platinum Partner Avnet will support a team of healthcare IT solution providers—Avnet partners—who assist midsize hospitals with issues related to accountable care, electronic health records, and healthcare information exchanges using Oracle Health Sciences solutions. Those Avnet partners participate in Avnet's SolutionsPath practices, which include strategic planning consultation on high-growth vertical markets and data center technology areas.

Avnet helps its partners to obtain specialization in Oracle Health Sciences solutions through its "Under Our Wing" program that includes curricula, study guides, exam syllabi, and more.

avnet.com

SAVE THE DATE

Oracle PartnerNetwork Exchange @ OpenWorld

September 30–October 4
oracle.com/opnexchange

LINDY GROENING

Bring the Noise

Passion and participation drive the Oracle Technology Network Architect Community.

In this issue, Bob Rhubart, manager of the architect community on Oracle Technology Network, provides an update for and about that community.

Oracle Technology Network is organized around four pillars: Java developers, database administrators and developers, system administrators and developers, and architects. The boundaries between those community segments are fuzzy and porous, reflecting the practical reality of enterprise IT, where individual players often straddle multiple roles. Each role makes a specific contribution to the creation, implementation, and maintenance of IT solutions, and each contribution is an integral piece of the larger IT puzzle. While each role is essential to the overall process, it is the architect's job to see that all of those pieces fit into a coherent, fully functioning whole.

That's not to say that architects stand above the other roles. But architecture is, by necessity, the context in which all of the individual roles must operate. This is increasingly true as service orientation, cloud computing, the wildly expanding use of mobile devices, social computing, and the explosion of data continue to reshape not just the IT landscape but the very environment in which businesses operate.

Architects must be able to clearly communicate how all the pieces fit together within the solution, and how the solution fits into the larger enterprise picture. That's no small task, which is exactly why Oracle recognizes the role of the architect.

In supporting the architect community, Oracle Technology Network draws on the community itself, especially the leaders within that varied and dynamic network: experienced IT professionals whose interests center on Oracle technologies and how they fit into the enterprise IT ecosystem. Oracle's own experts are part of that community, but they are by

It is the architect's job to see that all of those pieces fit into a coherent, fully functioning whole.

no means the sole source of the architecture content on Oracle Technology Network. Members of the Oracle ACE and Oracle ACE Director programs, whose collective expertise spans a variety of specialties, are high on the list of contributors. The resulting content stream draws from articles, blog posts, social network updates, podcasts, and videos.

Just as architecture itself focuses on ensuring that the various elements of a solution fit together and act as a cohesive whole, the architect pillar strives to funnel and filter the community information stream into a coherent channel that represents the interests and serves the needs of the architect community. The focal point of that effort is the Oracle Technology Network Architect home page, a portal to a wide variety of resources for architects.

Among those resources is **IT Strategies from Oracle**, an evolving library of reference architectures and best practices. Another is the **OTN ArchBeat podcast**, featuring an ongoing series of panel discussions in which subject matter experts and community leaders share insight in an open forum.

Speaking of open forums, the Architect home page also includes links to companion social channels on Facebook and Twitter where you can track and participate in the community conversation as it happens.

Moving out of social networks and into the real world, the events listings on the Architect home page will keep you up to date on opportunities to connect with community leaders and technical experts at a wide variety of live, in-person events, including the ongoing series of Oracle Technology Network Architect Days.

Of course, conversation is a two-way process. So while we invite you to take advantage of the stream of architecture content on Oracle Technology Network as an information consumer, it is equally important that you contribute to the stream. The aforementioned social channels offer an easy, informal way to participate in the conversation. But those interested in a more in-depth, formal approach are encouraged to submit articles for publication in order to demonstrate and publicize their architecture expertise and practical knowledge.

In the end, a community is about more than shared interests. It is a collection of individual voices and opinions. True, that can often be a cacophonous chorus. But that's when you know that the participants are passionate about the topic. Oracle Technology Network is listening, always working to gather and distill the architecture conversation in a way that makes sense and provides the greatest value to the community. That's our end of the bargain. It's up to you to add your voice and your passion. ◀

WEB LOCATOR

Oracle Technology Network Architect Home Page
oracle.com/technetwork/architect

Article Submission Guidelines
bit.ly/KqsNmM

OTN ArchBeat Podcast
feeds2.feedburner.com/OtnArch2Arch

ArchBeat Blog
blogs.oracle.com/archbeat/

OTN ArchBeat on Twitter
twitter.com/otnarchbeat

OTN ArchBeat Facebook Page
facebook.com/otnarchbeat

Catching Up to Mobile Computing

Mobile computing presents challenges and opportunities for architects.

If mobile computing isn't already on your architect radar, it will be. Apple has sold 55 million iPads since 2010. Gartner expects a 98 percent increase in tablet sales in 2012, to 118 million. Nielsen reports that smartphones now account for nearly half of all mobile phones in the U.S., a 38 percent increase over 2011. And the mobile juggernaut is just getting started.

Karina Ishkhanova, technical lead for payment systems architecture and design at School-Day Solutions, sums up her take on the mobile challenge for architects in one word: *uncertainty*. That uncertainty, according to Ishkhanova, applies to the evolution of the operating systems for mobile apps, the availability of APIs and other necessary components, and the stability of data streams that can be interrupted at any time by lost connections or dead batteries.

"Properly designed, thought-through architecture takes center stage and becomes a matter of survival not only for single applications but for entire companies," Ishkhanova says. "Mobile computing demands data loss–tolerant architectures with multiplatform adaptability and insanely optimized resource usage."

Those mobile computing demands place an even greater importance on interoperability and standards, according to Oracle ACE Director Ronald van Luttikhuizen, managing partner at Vennster. "New developments happen so quickly nowadays that your architecture needs to be able to support the rapid development of new functionality and services," van Luttikhuizen says. "Mashups, cloud computing, and service orientation are a few things that can help."

Mike van Alst, an architect with MShift, believes that service orientation especially is key to mobile success. "Without service-oriented architecture (SOA), building mobile applications may be a bridge too far," he says. "We need to make sure that mobile computing converges with a SOA environment by correlating service capabilities with the needs of the mobile applications."

Smartphones now account for nearly half of all mobile phones in the U.S.

Addressing the needs of mobile applications will require a head-on approach. "The growth in mobile computing forces architects to think about multichannel delivery from the ground up rather than as an afterthought," says Anbu Krishnaswamy, an enterprise architect at Oracle. "The challenge for architects is to identify the right content for the right channel and reuse the plumbing underneath to provide consistent access to the underlying business processes."

An organization's ability to reuse existing "plumbing" and meet the increasing demands of mobile computing depends a great deal on the viability and integrity of the organization's existing architecture. "The introduction of multichannel support should have no major impact on an architecture that has always followed best practices and core architecture principles of abstraction, modularity, and reuse," he says. "In theory, mobile devices become just another delivery channel."

Even so, adding and supporting that mobile delivery channel won't be without its challenges. "The scalability of enterprise systems will be put to the test," says Oracle ACE Director Lucas Jellema, CTO at AMIS Services. "To alleviate that anticipated load, smart caching solutions that leverage a memory grid will play an important role."

Jellema also cites user expectations and the unique characteristics of mobile apps as significant considerations. "Mobile apps are a mashup of various social networks, multiple types of rich data, and technical aspects such as communication protocols and types of data," he says. "And the expectation of immediate response and constantly fresh data will have serious and interesting consequences with regard to event-driven architecture."

Folding mobile computing into the enterprise IT mix seems inevitable as consumers' enthusiastic embrace of mobile computing continues to change how they interact with commercial, governmental, and social institutions. But as with any disruptive change in the enterprise IT landscape, new challenges mean new opportunities for architects. ◀

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

GET more architect information
oracle.com/technetwork/architect

LISTEN to ArchBeat podcasts on mobile challenges
bit.ly/la758G

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

GET the Oracle mobile app for iPhone and iPad
bit.ly/lqFhvL

Immersion Learning

Peers delve into social media, virtualization, and "extreme BI" to gain insight.

JOHANNES AHRENDS

ORACLE ACE

Company: CarajanDB, a tech consulting firm specializing in Oracle solutions

Job title/description: Manager, focusing on the development of private cloud infrastructures, especially for Oracle databases

Location: Erftstadt (near Cologne), Germany

Length of time using Oracle products: 20 years

What's your favorite tool or technique on the job? The whole concept of managing storage in a grid infrastructure with Oracle Automatic Storage Management. With that very reliable infrastructure, it's much easier to maintain an Oracle database, migrate it to Oracle Real Application Clusters, or move it on the fly to a new SAN [storage-area network] box or other disk structures.

What technology has most changed your life? I'd say first computers, next the internet, and finally virtualization. I've been working with virtualization solutions for about four years now, and these solutions have allowed me to easily set up Oracle Real Application Clusters with Oracle Database 10*g* and Oracle Database 11*g*, and to create test databases with Oracle Data Guard and Oracle GoldenGate. Virtualization is also helping us help our clients to save energy.

How are you using social media in your work life? I blog and use Twitter and Facebook, but Xing [xing.com, a professional business network]—which in Germany is more popular than LinkedIn—is my main platform for connections to peers, looking for contacts, and exchanging basic knowledge through the site's Oracle group.

WISSEM ELKHLIFI

ORACLE ACE

Company: Motorola Mobility, a communications corporation

Job title/description: Senior DBA, responsible for Oracle Database installation, upgrades, configuration, tuning, and PL/SQL development

Location: Barcelona, Spain

Oracle credentials: Oracle Certified Professional (Oracle Database 10*g*) with 12 years experience using Oracle products

What advice do you have for those just getting into database development? I see a lot of programmers who, after taking on a project, start programming immediately after reading a specification. In my opinion, it's better to lean back and think about the task and the best way to approach it. Also, take time to understand the concepts and best-practice examples. You'll benefit in the long run by avoiding performance problems, bad programming, and extensive bug-fixing.

How are you using social media in your work life? Lately I've become addicted to social media tools such as Twitter and Facebook to connect with Oracle experts around the globe. I'm now connected with many experts, including Hans Forbrich, Eddie Awad, Syed Jaffar Hussain, Kai Yu, Ben Prusinski, and Tanel Poder.

What would you like to see Oracle, as a company, do more of? Provide more Oracle documentation in Spanish, French, and German. Certainly English is the number one IT language, but there are a lot of DBAs—especially here in Spain—who face difficulties understanding Oracle concepts due to the language barrier.

STEWART BRYSON

ORACLE ACE

Company: Rittman Mead, a technology firm focused on Oracle business intelligence solutions

Job title/description: U.S. managing director, overseeing projects, staff, and sales while doing consulting and training

Location: Atlanta, Georgia

Length of time using Oracle products: 15 years

How did you get started in IT? After receiving my master's degree in philosophy, I realized that I needed to go out and make a living—that I wouldn't be able to generate a job through a syllogism. After working at a marketing job, I decided my best bet was to get into technology. So I pounded the bricks and convinced a company to give me a break in the application support group.

Which Oracle technologies are you currently finding most valuable? I've been working with the combination of Oracle Exadata Database Machine and Oracle Business Intelligence Enterprise Edition 11*g*, delivered with an agile methodology. I like to call this setup "extreme BI." These two technologies are a natural fit in terms of performance and functionality.

What advice do you have for those just getting into business intelligence development? Most developers get their feet wet with Oracle Business Intelligence Enterprise Edition doing front-end development: delivering analytic reports and dashboards. I'd recommend getting started with Oracle Business Intelligence Enterprise Edition metadata development first. This type of development provides a deeper respect for what the product can really deliver. ◀

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

LOCKDOWN

I-HUA CHEN

Oracle Database security solutions deliver layers of enterprise information protection.

BY DAVID BAUM

Hackers who manage to break in to government and corporate enterprise systems make the news—sometimes very big news. An organization's most privileged IT users may not get the same news attention, but their ability to access sensitive information as well as to configure systems,

Andrew Meade, senior database administrator at TransUnion Interactive, leads a team of DBAs charged with keeping the company's data available, secure, and compliant. "We have a lot of sensitive information that we must protect, and Oracle gives us many different ways to defend our data against attacks," he says.

modify databases, and grant privileges makes internal system administrators and DBAs a security risk to the enterprise and a target for hackers. And they aren't the only potential liabilities: in some cases, sensitive production data moves through development and test environments, where any developer can see it.

"In most organizations, two-thirds of sensitive and regulated data resides in databases," points out Vipin Samar, vice president of database security technologies at Oracle. "Unless the databases are protected using a multilayered security architecture, that data is at risk to be read or changed by administrators of the operating system, databases, or network, or hackers who use stolen passwords to pose as administrators. Further, hackers can exploit legitimate access to the database by using SQL injection attacks from the Web. Organizations need to mitigate all types of risks and craft a security architecture that protects their assets from attacks coming from different sources."

Companies spend billions of dollars each year securing their IT systems worldwide, including software, services, and support, Samar adds. Despite these massive investments, attacks have continued, and the attackers—with the help of social engineering and sophisticated automated tools—have continued to succeed. This is partly because traditional security strategies have focused on protecting the network perimeter and desktop and laptop machines, rather than internal server assets. Ultimately, what is most important is protecting the databases themselves, yet this is the area that many companies overlook. According to Forrester Research, while 70 percent of companies have an information security plan, only 20 percent of them have a database security plan.

Samar thinks security professionals can learn a lesson from history. In medieval Europe, castles had multiple defenses—wide moats, high walls, iron doors, and even counterattacking archers to repel different types of attackers. "Similarly, in the IT world, you have to defend your databases from casual onlookers, opportunistic insiders, and state-sponsored hackers," he says. "Data is your king, but if your defensive moat is a firewall of pawns, it is easy for an enemy knight to jump across and checkmate your king."

DATABASE ENCRYPTION

Oracle Advanced Security, an Oracle Database 11*g* option, helps organizations protect sensitive data on the network, on storage media, and within the database. It addresses privacy and regulatory requirements including the Payment Card Industry Data Security Standard (PCI DSS), Health Insurance Portability and Accountability Act (HIPAA), and numerous breach notification laws.

A need to comply with PCI DSS regulations motivated TransUnion Interactive to adopt Oracle Advanced Security's Transparent Data Encryption feature and Oracle Database Firewall.

"We have a lot of sensitive information that we must protect, and Oracle gives us many different ways to defend our data against

PETER STEMBER

attacks," says Andrew Meade, senior database administrator at TransUnion Interactive. "We don't want that information stored in plain text on a disk. Oracle Advanced Security's Transparent Data Encryption ensures that any time information is written to disk or backed up to another location, it is completely encrypted."

SNAPSHOT

TransUnion Interactive
transunion.com
Headquarters: Chicago, Illinois
Industry: Financial services
Oracle products: Oracle Database 11*g*, Oracle Advanced Security, Oracle Database Firewall

Meade leads a team of database administrators charged with keeping TransUnion Interactive's data available, secure, and compliant. It's a high-volume, nonstop operation that processes thousands of transactions per second.

TransUnion Interactive is the consumer subsidiary of TransUnion, providing credit-report, credit-monitoring, and alert services for consumers along with educational tools to help them stay on top of their finances and avoid identity theft.

"We used tablespace encryption in Oracle Database 11*g* to protect our databases and address regulatory compliance issues," says Ramdas Kenjale, director of architecture and infrastructure at TransUnion Interactive. "This method allowed us to encrypt our data very quickly, without changing our applications or modifying our infrastructure. Transparent Data Encryption encrypts data when written to disk and decrypts it after a user has been successfully authenticated and authorized."

Kenjale says Oracle's approach with Transparent Data Encryption shields his team from the details of encrypting specific columns in each database table. It fulfills PCI DSS requirements by encrypting data in storage, in transit, and on backup media. All access controls that are enforced by Oracle Database remain in effect, including object grants, roles, virtual private database, and Oracle Database Vault. Oracle's two-tier system includes a master encryption key that protects data encryption keys.

WEIGHING THE ALTERNATIVES

TransUnion Interactive considered alternatives from other security vendors such as *full-disk encryption*, in which data is encrypted at the hardware level, and *tokenization*, in which a token represents the actual data. "Tokenization would have meant changing all of our applications and parts of our architecture, which would have been time consuming and costly," says Meade. "For us, tokenization is not really a viable solution."

Full-disk encryption wasn't viable either because it would have required TransUnion Interactive to take key databases offline whenever encryption keys are rotated. "Neither solution satisfied our needs for a zero-downtime implementation," Meade says.

TransUnion Interactive chose Transpararent Data Encryption in part because of the flexibility it provides. "Oracle Advanced Security with Transparent Data Encryption is the perfect solution for us," Meade says. "It lets us encrypt all of our data without any application or infrastructure changes. It's fully integrated with Oracle Database. And key management is built in. No downtime is required to create or rotate the keys, so it works well for us. It is easy to use, easy to implement, and easy to maintain."

TransUnion Interactive recently passed a PCI audit that focused on encryption and key rotation, validating the effectiveness of its Oracle solution. Users see little or no difference in the level of service. "The performance impact of Transparent Data Encryption is negligible," adds Meade. "In our case, it is less than 1 percent."

ORACLE DATABASE FIREWALL

TransUnion Interactive is now implementing Oracle Database Firewall to complement its existing network security strategy. Most security experts see *database firewalls* as an important adjunct to *network firewalls*, which protect a data center from unauthorized access from the outside.

To guard against unauthorized database access, Oracle Database Firewall monitors the SQL network traffic going to the database, and provides a first line of defense against threats originating from both outside and inside the organization. It monitors data access, enforces access policies, highlights anomalies, and protects against network-based attacks.

"Oracle Database Firewall reveals precisely what types of queries are hitting our database, who is submitting them, and where they come from," says Meade. "All that information is exposed based on our preferences, which we specify via a graphical user interface."

Database Security: The Big Picture

Database security is an essential component of a complete IT security program. According to Vipin Samar, vice president of database security technologies at Oracle, it often begins with the three *A*s: authentication, authorization, and auditing.

Many Oracle customers implement Oracle Identity Management to enable centralized access control, along with granular role-based controls and provisioning capabilities.

In addition, Oracle Database Vault limits the activities of privileged users by placing sensitive database tables and applications in a protective realm. Oracle Audit Vault provides robust monitoring and auditing of these privileged users.

"Some organizations still use generic shared accounts, which doesn't allow them to track which users made which changes," says Samar. "Rather than granting excessive privileges for the sake of convenience, it is better to assign the least privileges for each user to do his or her job."

When considering their overall security architecture, Samar believes that customers should adopt a multitiered perspective by viewing potential threats from outside and inside the database. He suggests that organizations follow these guidelines:

1. Use a network-based database firewall to reduce threats from Web users. Deploy a database firewall that can analyze SQL and prevent rogue statements from hitting the database.
2. Use data encryption to secure primary database and backup data when it is written to disk.
3. Encrypt network traffic to reduce the risk of someone sniffing data on the wire.
4. Empower DBAs to tune and manage databases, but restrict them from viewing or modifying sensitive data.
5. Audit sensitive operations and the activities of privileged users.
6. Mask sensitive data before sharing it with development and testing teams.

Meade is in the process of developing a white list and a black list of various types of SQL statements. "Anything that is on the white list gets through, and everything on the black list is blocked," he explains. "The database firewall analyzes SQL traffic. Based on policies we establish, it chooses to block, substitute, log, or send an alert about each suspicious statement."

In addition to evaluating the legitimacy of SQL statements, Oracle Database Firewall can consider factors such as the requester's IP address, time, and program name. TransUnion Interactive can choose to deploy it in blocking mode as a database policy enforcement system as well as for supplemental auditing and compliance purposes.

A COMPREHENSIVE VIEW

Martin Kuppinger, founder and principal analyst at KuppingerCole, a leading analyst company for identity-focused information security, explains the importance of this type of defense. "While a network firewall controls access to IT resources at the IP level, it looks at packets, so it doesn't have a very deep understanding of what happens at the SQL level. Database firewalls provide in-depth protection for communication to databases by monitoring and enforcing normal application behavior. They prevent SQL injection attacks and unauthorized SQL commands."

Kuppinger sees database encryption and database firewalls as an important part of a complete database security strategy. Other essential technologies include data masking for test environments and controlling access to databases. "It's important to have solutions for every aspect, starting with strong authentication and granular access control to databases, ensuring that operators and database administrators can't abuse their privileges," Kuppinger concludes. "If you trust only one database security solution, you will fail to address all of these different aspects. Leaving some doors open doesn't really solve your enterprise security issues." ◀

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

NEXT STEPS

READ more about database security

"Effective Data Leak Prevention Programs: Start by Protecting Data at the Source—Your Databases"
oracle.com/us/products/database/039434.pdf

Oracle Database Security
oracle.com/us/products/database/security/overview

FUTURE

-PROOF

Choose the best of today's tools to build next-generation Oracle Database applications.

BY DAVID A. KELLY

Trends in application development have come and gone over the years. Remember when client/server was new? Rich, thick-client applications ran on UNIX-based workstations and later on PCs. Remember the rise of the thin client? A Web browser on a PC became the delivery method for portable—across different desktops—thin-client applications. The next platforms for application delivery are here, and users are demanding a rich and portable experience on these rapidly evolving devices.

I-HUA CHEN

Purdue Pharma's Sayee Natarajan, director of sales systems (left), and Larry Pickett, vice president and CIO, rely on the development speed and agility of Oracle Application Express for rapid iteration. "We've been able to make changes based on business needs very quickly and very effectively," Pickett says.

"These days, application developers are creating rich, browser-based applications that need to work on everything from the corporate desktop to smartphones and tablets," says Michael Hichwa, vice president of software development at Oracle. Hichwa is responsible for several Oracle developer products aimed at database application development, including Oracle SQL Developer (see "Database Development Choices" sidebar).

According to Hichwa, development tools for Oracle database products are leveraging HTML5/CSS3 technology at the same time that they're becoming more declarative, to facilitate faster development. "We want to enable our customers to build out applications faster, and with all the features they need," Hichwa says.

Faster application development enables organizations to see lower development costs and realize a faster return on investment. It also makes it possible to keep up with the latest IT drivers, such as cloud, social networking, and mobile computing.

THE EXPRESS PATH

Pharmaceutical company Purdue Pharma is one company with an eye on trends as it deploys its state-of-the-art custom-built customer relationship management (CRM) database application, originally built for the desktop, to 1,000 field users on multiple platforms, including the iPad.

Purdue is a privately held pharmaceutical company based in Stamford, Connecticut. The company—and independent associated companies—engage in the research, development, production, and marketing of prescription and over-the-counter medicines and healthcare products. Purdue is an industry leader in pain management and is introducing important products in other therapeutic areas such as insomnia.

Since a restructuring in 2005, the company has focused on bringing new, innovative medicines to the marketplace while simultaneously managing costs and creating an agile IT infrastructure that can rapidly adapt to meet changing business needs, all the while staying ahead of technology.

"The rapid pace of technology change is something that we have to stay on top of," says Larry Pickett, vice president and CIO at Purdue. "Over the past several years, we've focused on using newer technologies and approaches, to support the growth of the company."

For example, after the restructuring, Pickett and his team looked for creative ways to drive down costs while still meeting Purdue's business needs. "That's where Oracle Application Express first came into play," Pickett says. Sayee Natarajan, Purdue's director of sales systems, used Oracle Application Express to

CATHERINE GIBBONS

SNAPSHOTS

Purdue Pharma L.P.
purduepharma.com
Location: Stamford, Connecticut
Employees: More than 1,500
Oracle products: Oracle Application Express, Oracle Database

Lyris
lyris.com
Location: Emeryville, California
Employees: 150
Revenue: US$40 million in 2011
Oracle products: MySQL

Spatial Eye
spatial-eye.com
Location: Culemborg, the Netherlands
Oracle products: Oracle Data Provider for .NET, Oracle Database 11*g* Release 2, Oracle Spatial, Oracle Workspaces

prototype a new CRM solution to manage Purdue's sales territories.

"The new CRM system worked so well, we rolled it out to our entire field sales force," says Pickett. "We're continuously improving it with new capabilities, while continuing to receive rave reviews from users."

The application is also providing Purdue with significant cost savings. Comparing it to the system it replaced, Pickett estimates savings of about US$2 million annually, "so over the life of the system so far, we've saved the company roughly [US]$12 million," he says.

Another benefit is the agility and development speed Purdue gets from Oracle Application Express, says Natarajan. "Oracle Application Express allows us to do rapid iteration and rapidly respond to feedback from our users, proving that the agile development idea really works."

Pickett concurs. "We've been able to make changes based on business needs very quickly and very effectively," he says. Based on input from users in the field, Natarajan's team "has been able to iterate six major versions a year along with six minor versions."

According to Pickett, Purdue's salespeople are very excited about the solution—especially new hires coming from competitors. "Our salespeople feel that our Oracle Application Express–based CRM system is better than anything on the market," says Pickett. "It's adding a lot of value."

Over the years, Natarajan and his team have deployed the CRM application to Microsoft Windows, Red Hat Linux, Oracle Linux, Mac OS, and most recently, the Apple iPad. The port to the iPad took only 90 days. "The fact that we could port our application to the iPad in 90 days speaks to the power of Oracle Application Express and our ability to rapidly prototype," says Natarajan.

The company uses Oracle Application Express for more than 40 production applications, including its eCaseFile system for adverse events and the product R&D pipeline project-tracking system, both of which are mission-critical. The project-tracking system, says Pickett, works on the iPad as well as the company's portal and "captures all of the decisions and decision criteria for advancing projects through our R&D pipeline to our product portfolio, so the system is the tool we use to get new products to market—managing the cost, the timing, critical-path decisions, and so on. It's a really innovative application that our R&D executives are using."

BUILDING CLOUD-READY DATABASE APPLICATIONS

When you're building a cloud-based product that must be able to scale exponentially, you need a solid, reliable database solution. "That's where MySQL comes in," says Navdeep Singh, vice president of IT and operations at Emeryville, California–based Lyris. "We needed a very solid database that could easily scale and seamlessly integrate with other components—such as [Apache] Hadoop—to enable the rapid processing of massive amounts of data and automate digital marketing activities." Like many other companies Singh had worked for in the past, Lyris developed its current application offering using MySQL.

Lyris needed a solid database that could scale for cloud deployments, says Navdeep Singh, vice president of IT and operations. MySQL was the solution.

Lyris has been providing digital marketing automation solutions and services to its customers—sophisticated marketers who want to deploy intelligent multichannel campaigns across e-mail, social, and mobile channels—for more than 10 years. "Our software helps organizations automate digital messaging campaigns using rules based on the specific customer segments they're targeting," says Singh. For example, using the Lyris platform, marketers can define rules that specify precise message sequences to selected subsegments, such as sending follow-up e-mails or social notifications to a group of customers only after an initial message has been viewed, and not sending subsequent messages to customers who have not yet taken action on earlier communications.

Lyris ensures that its products deliver top-notch performance by scaling horizontally with multiple instances of MySQL Database and then taking advantage of MySQL's replication feature. "Replication is absolutely critical to scaling our technologies," says Singh. "Multiple

Database Development Choices

When it comes to tools for building database applications, developers have many choices—from Oracle and the open source community.

Oracle Application Express. Oracle Application Express is a no-cost rapid Web application development tool included with all editions and releases of Oracle Database 11*g*. For developers whose skills may be primarily SQL and PL/SQL, Oracle Application Express enables them to easily step into the world of Web-based applications by using a tool that's embedded in Oracle Database. Development and deployment are fast and easy, through a Web browser.

Oracle SQL Developer. Oracle SQL Developer is a free tool from Oracle that provides a graphical interface for database development, a migration path from other databases, and data modeling capabilities. "Oracle SQL Developer is a highly productive tool that helps you build the best SQL and PL/SQL possible," says Michael Hichwa, vice president of software development at Oracle.

Open source. Database developers can use many different open source dynamic scripting languages and tools with Oracle database products, including Ruby, Python, and PHP. PHP has long been associated with MySQL, and has been used with Oracle Database since Oracle Database 10*g*.

BOB MORRIS

Spatial Eye depends on Oracle Data Provider for .NET for integration with Oracle databases, says Patrick van Dijk, lead developer and architect.

replicas of the production database let us run sophisticated reports without affecting the performance of the core application." This approach also provides Lyris with higher availability, because one of the many replicas can take over processing automatically if necessary, allowing marketers rock-solid application reliability and access to up-to-the-minute reports that measure the impact of their digital marketing initiatives.

Singh also points out how much information and support is available to the MySQL developer and DBA communities, not only from Oracle but also from MySQL internet support channels, newsgroups, community forums, and the like. "MySQL is really easy to monitor, manage, and maintain. And, you don't necessarily need high-end hardware," says Singh.

Singh doesn't foresee any future problems scaling his MySQL solution to meet his company's evolving needs. "It's pretty simple. Our product wouldn't function without MySQL," says Singh. "No matter how our business evolves in the future, I'm certain MySQL will be a part of it, especially given how well it plays with emerging big data technologies."

DIGGING DEEPER WITH ORACLE DATA PROVIDER FOR .NET

Although Web-based delivery is the wave of the future for many applications, the approach is not a one-size-fits-all solution. For example, independent software vendor and Oracle partner Spatial Eye, located in the Netherlands, uses Microsoft tools and technologies such as Microsoft .NET to develop Microsoft Windows–based applications for spatial analysis.

For integration with Oracle databases, Spatial Eye's Windows-based solutions rely on Oracle Data Provider for .NET (ODP.NET) to move data. "From a developer perspective, ODP.NET is very easy to use," says Patrick van Dijk, lead developer and architect at Spatial Eye. "ODP.NET integrates very well with our environments."

Users of Spatial Eye software include utility companies, such as power, water, and telecom, and also municipalities and other agencies that may have underground physical assets that need to be identified and managed.

The Spatial Eye product portfolio includes a spatial data warehouse application and an application called Call Before You Dig. Spatial Eye Call Before You Dig is an off-the-shelf solution that can be configured and customized for various communication processes (phone, fax, e-mail) and datasources.

The company's flagship product, Spatial Workshop (available in Standard, Professional, and Ultimate editions), lets users explore, visualize, query, analyze, and integrate spatial data to varying degrees. Spatial Eye Call Before You Dig and Spatial Workshop run on Windows and work with spatial data contained in an Oracle database running the Oracle Spatial option.

Spatial Workshop can be used as a desktop application, or as a mobile field system. Geographic information from Oracle Spatial can be taken into the field using standard Oracle replication techniques, so users have access to up-to-date information before they start tearing up a roadway, for example. Spatial Eye's products and services also leverage the Oracle Workspaces API.

For example, a city planner designing a new electrical system for the city or a neighborhood can develop alternatives of the plan, and then use Oracle Workspaces to store and compare different versions. Spatial Workshop supports a plug-in, developed using the Oracle Workspaces API, to tap into that data and facilitate the comparison.

But ODP.NET is the primary Oracle development tool for Spatial Eye. "ODP.NET is very rich, and the implementation does a lot," says van Dijk. "It scales well, both in terms of analysis and in getting data from the Oracle server into the online application."

TOOLS FOR TOMORROW

Regardless of which tools developers use to create the next generation of applications, development will still revolve around the data held in relational databases, such as Oracle Database and MySQL. "SQL is still the backbone of the industry," says Oracle's Hichwa. It's still a very effective way to manipulate data, he adds, noting Oracle's goal is to make sure "developers have the best experience possible when creating applications that target Oracle databases." ◀

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

NEXT STEPS

READ **more about**

MySQL
oracle.com/us/products/mysql

Oracle Database 11*g* Release 2
oracle.com/us/products/database

Oracle Application Express
oracle.com/technetwork/developer-tools/apex/overview

Oracle SQL Developer
oracle.com/technetwork/developer-tools/sql-developer/overview

Oracle Data Provider for .NET
oracle.com/technetwork/topics/dotnet/whatsnew/index-101716.html

dynamic scripting languages
oracle.com/technetwork/articles/dsl

TON HENDRIKS

ORACLE APPLICATION DEVELOPMENT FRAMEWORK AND ORACLE JDEVELOPER

Service, Please!

Integrate Web services with Oracle ADF and Business Components applications.

Besides being used in SOA, Web services are a common access pattern in Web application development for remote data queries, application integration, and application departmentalization. Rather than enabling developers to directly access database tables owned by other lines of business, companies can expose the services interface to ensure data integrity and consistent application of business logic on the source data.

In this column, you will learn about the Web services integration options available for Oracle Application Development Framework (Oracle ADF) applications. Stepping through a sample application, you'll learn how to embed calls from Oracle ADF's Business Components models into Web services for create, read, update, and delete (CRUD) operations.

WEB SERVICES INTEGRATION IN ORACLE ADF

SOAP-based Web services and Oracle ADF can be integrated in several different ways. Oracle ADF's Web Service Data Control feature is easiest to configure at design time, but it does not provide much control of a service's runtime behavior, nor of the exchanged data. Thus, using Web Service Data Control is best for simple interactions only.

On the other hand, using the JavaBean Data Control feature or Business Components feature of Oracle ADF supports deeper, more complex integration scenarios. JavaBean Data Control can be used for integrations involving plain-old Java objects (POJOs). For applications that use Business Components as the business service layer, it makes sense to integrate the Web service application at the Business Components service integration layer and take full advantage of Business Components benefits such as participating in master-detail relationships and leveraging advanced framework features such as using model-driven lists of values.

In the JavaBean Data Control or Business Components approach, the Web services are accessed through the standard Java API for XML Web Services (JAX-WS) proxy client, which gives you programmatic control for manipulating data in transit. Rather than using the JAX-WS proxy client alone, I recommend using it in conjunction with a JavaBean wrapper. The wrapper provides a stable API for your Business Components application integration, so that if the Web service changes, you'll only need to regenerate your JAX-WS proxy client code and drop it into place.

The example provided with this column shows you how to integrate a Web service and a Business Components application by using this recommended approach. Before delving into the step-by-step development process, let's take a closer look at the integration architecture.

SAMPLE APPLICATION INTEGRATION ARCHITECTURE

The sample applications include a basic but complete Java Platform, Enterprise Edition Web service and a starter Business Components application that you'll extend to support integration with the Web service. The Web service obtains data from the Departments table in the Oracle HR sample database schema.

By default, Business Components queries data from database tables and views. For Business Components to work with other datasources, such as Web services, PL/SQL stored procedures, or Java Naming and Directory Interface (JNDI) services (LDAP, for example), developers create programmatic view objects and entities that override the Business Components framework default behavior.

Figure 1 shows the class hierarchy and the integration architecture used in the sample. The JAX-WS proxy client is accessed from a JavaBean wrapper in the Business Components project. The JavaBean wrapper is accessed from methods of the **WsDepartments** entity object and the **WsDepartmentsView** view object.

The **JAX-WS Proxy Client**, **JavaBean Wrapper**, **WsEntity**, and **WsViewObject** classes are all provided in the sample workspaces. **WsEntity** and **WsViewObject** are code templates that assist in the development of programmatic view objects and entities for Web services CRUD access. In general, I recommend building templates such as these for recurring tasks in your application development projects. (Both templates are located in the **oramag.sample.application.model.fmwExtensions** package of the Business Components project. Examine the code to learn more.)

In the sample integration this article covers, you'll first create the programmatic entity object **WsDepartments** (using a wizard in conjunction with the **WsEntity** class template) and then you'll use that object as the basis for the programmatic view object **WsDepartmentsView**, which you'll also create by using the wizard and the **WsViewObject** class template. You'll then complete and test the integration by adding the programmatic view object on the application module, for use from the Oracle ADF Data Controls panel.

GETTING STARTED

To begin, download the sample application at bit.ly/I5LIrR and unzip the file. The file contains two folders:

- **starter-app** contains two workspaces: a completed Web services application and a Business Components application workspace with the JAX-WS proxy client contained in its own project. This article shows you how to complete the business components in the project to integrate with

the Web services application.

- **completed-app** contains completed application workspaces for the Business Components application and the Web service. (As you'll see, the completed app also includes some advanced capabilities.)

To follow the steps in this column, you'll need the studio edition of Oracle JDeveloper 11*g* Release 2 (11.1.2.2) or later, available as a free download on Oracle Technology Network. You'll also need an Oracle Database instance with an unlocked HR schema.

Preliminary setup tasks include changing the database connection to point to your HR database schema from both the Business Components application project and the Web services project, as follows:

1. Launch Oracle JDeveloper 11*g* Release 2. Select **File** -> **Open**, and navigate to the directory containing the unpacked sample application.
2. Open the **starter-app** folder and then the **WebService** folder, and select the **WebService.jws** file, followed by **Open**. The workspace will now open in Oracle JDeveloper.
3. Select **View** -> **Database** -> **Database Navigator**, and expand the **WebService** node to find 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 setup. Test the changes, and then click **OK** to close the connection properties.

With a working connection to your database, you can now deploy and run the Web service using Oracle WebLogic Server, which is integrated into Oracle JDeveloper. In the Application Navigator

5. Expand the **hrDeptEmpService** project node and then the **oramag.sample.jaxws.hr.session** package structure.
6. Continue expanding the package structure until you locate the class file **hrDeptEmpServiceBean.java**.
7. Right-click **hrDeptEmpServiceBean.java** and choose **Run** from the context menu to deploy and run the service. (If a Create Default Domain dialog box appears, create a password for the default Oracle WebLogic Server domain associated with Oracle JDeveloper and click **OK** to continue.)

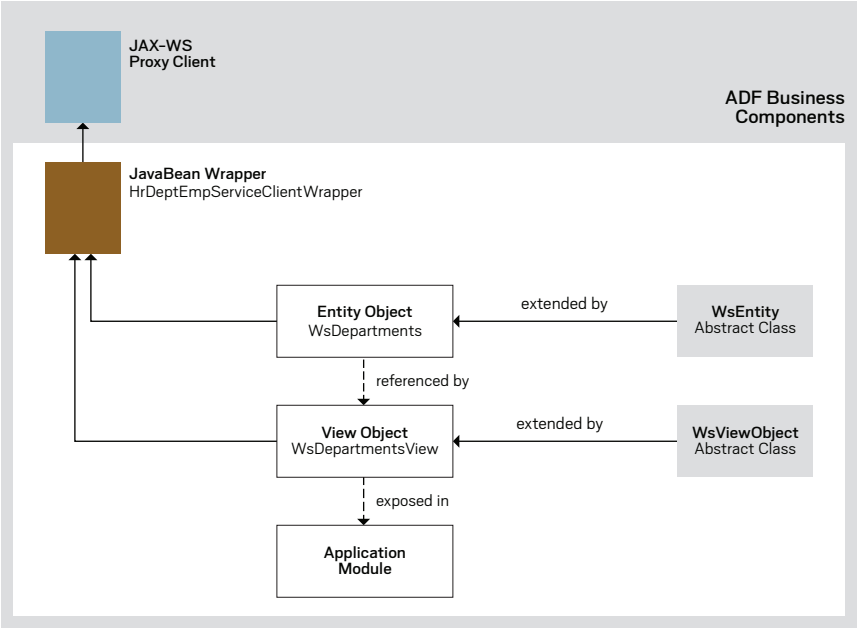

Figure 1: Sample Oracle ADF Web service integration architecture

At this point, the Web service has been deployed and is running on the integrated Oracle WebLogic Server. You must also provide the correct database connection information for the Business Components application:

8. Using **File** -> **Open**, navigate to the **OraMag-AdfBcWsApp** folder.
9. Select the **OraMag-AdfBcWsApp.jws** file.
10. Click **Open** to open the workspace.
11. Select **View** -> **Database** -> **Database Navigator**. Expand the **OraMag-AdfBcWsApp** node, and right-click the **hrconn** node.
12. Select **Properties** from the context menu to edit the database connection information to work with your database configuration. Test your changes, and then click **OK** to close the connection properties. You can now start integrating the Business Components application and the Web service.

BUILDING A WEB SERVICE–BASED ENTITY OBJECT

Web services are integrated into Business Components via programmatic view objects. To implement CRUD behavior, the view objects must be based on an entity object that overrides the built-in framework methods, **doDML** and **doSelect**, of the Business Components entity object. You provide information to Oracle JDeveloper about the properties and datatypes provided by the Web service. For the example integration, the Web service **Departments** entity has the following properties defined:

- BigDecimal departmentId;
- String departmentName;
- BigDecimal locationId;
- BigDecimal managerId;

You can obtain this information from the Departments.java class file or by examining the structure in the **oramag.sample.proxy.client.types** package, located in the JaxWsProxyClient project.

You'll use this information as you step through the wizard. With the **OraMag-AdfBcWsApp** application still selected, open the Application Navigator.

1. Expand the **AdfBcModel** -> **Application Sources** node and subsequent nodes until you can select the **oramag.sample.application.model.entities** package node.
2. Right-click the **oramag.sample.application.model.entities** package node, and select **New Entity Object** from the context menu. The Create Entity Object wizard launches. In step 1 of 6 of the wizard
   - In the **Name** field, type `WsDepartments`.
   - Uncheck the **Existing Object** checkbox.
3. Click **Next** and then **Yes** to accept

WSDEPARTMENTS as the schema object name.

4. In step 2 of 6, click the **New** button (repeatedly) to open a New Entity Attribute properties page and create the four entity attributes detailed in Table 1. When you have finished creating all four attributes, click **Next**.
5. Click **Next** to skip the wizard step 3 of 6. In the wizard step 4 of 6
   - Check the **Generate Entity Object Class** checkbox.
   - Click the **Classes Extend** button to launch the Override Base Classes dialog box.
6. In the Override Base Classes dialog box, click the **Browse** button next to the **Row** field. The Find Superclass search dialog box appears.
7. Type `WsEntity` in the **Match Class or Package Name** field. When the **WsEntity** (**oramag.sample.application.model.fmwExtensions**) package appears in the dialog box, click it and then click **OK** to save the change. Click **OK** again to dismiss the dialog box and redisplay the wizard step 4 of 6.
8. Click **Finish** to conclude the wizard.

Oracle JDeveloper generates a custom entity class, **WsDepartmentsImpl**, based on the template, and places its metadata file on the XML edit pane. Now you override the abstract methods of the **WsEntity** template class in **WsDepartmentsImpl** to make Business Components write data updates to the Web service.

9. In the WsDepartments.xml editor, select the **Java** menu option and the link next to the **Entity Object Class** label to open the entity implementation class, **WsDepartmentsImpl**, in the Java code editor.
10. In the Java code editor, right-click anywhere in the file and choose **Source -> Override Methods** from the context menu.
11. In the Override Methods dialog box, check the **Group By Class** checkbox and then locate **WsEntity** (and its methods) in the list.

    Select the **doDML** and **doSelect** methods under the **WsEntity** class. Also check the **Copy Javadoc** checkbox (at the bottom of the dialog box), so that the override instructions from the template populate the file.
12. Click **OK**.

    Oracle JDeveloper generates a Java class containing method signatures and comments instructing you how to access the JavaBean wrapper class in the AdfBcModel project from this class. To simplify your tasks for this example integration, I've provided the complete code for you.
13. Open the **entity-code-complete.txt** file, located in the **entities** package, and copy and paste the file's contents over **WsDepartmentsImpl.java**, replacing the entire generated content.
14. Close all tabs in Oracle JDeveloper, and save all your work.

    You now have a programmatic entity object that can be used as the basis for the view object that will access the Web service. The next step is to create the programmatic view object.

BUILDING THE PROGRAMMATIC VIEW OBJECT

When using programmatic view objects, you must override the Business Components framework methods that query the database for data. The view object methods you will typically need to override include create, executeQueryForCollection, createRowFromResultSet, getQueryHitCount, and hasNextForCollection. (See "Key Framework Methods to Override for Programmatic View Objects" in *Oracle Application Development Framework Developer's Guide For Forms/4GL Developers* to learn more.) For this sample integration, the template provides much of the code, so you won't be selecting all these methods.

Name	Type	Other property to set
DepartmentId	BigDecimal	Check **Primary Key** checkbox
DepartmentName	String	Set **Database Column Type** to **VARCHAR2(30)**
ManagerId	BigDecimal	None
LocationId	BigDecimal	None

Table 1: Attribute definitions for the custom entity object

As with the programmatic entity object creation, a wizard guides you through the process. These steps begin in the OraMag-AdfBcWsApp application in the Application Navigator:

1. Expand the **AdfBcModel -> Application Sources** node until the **oramag.sample.application.model.views** node appears.
2. Right-click the **oramag.sample.application.model.views** package node and select **New View Object** from the context menu. The Create View Object wizard launches. On the Name page (Create View Object – Step 1 of 9)
   - Enter `WsDepartmentsView` in the entity **Name** field.
   - Keep the **Data Source** option set to **Entity object**.
3. Click **Next** to continue. On the Entity Objects page (step 2 of 9)
   - Expand the **oramag.sample.application.model.entities** package, and select the **WsDepartments** entity object.
   - Click the right arrow to move **WsDepartments** to the **Selected** pane. Click **Next**.
4. Click the double-arrow to move all **Available** attributes to the **Selected** pane.
5. Click **Next** until the wizard step 7 of 9 appears. Set the Java properties for the view object as follows: check the **Generate View Object Class** checkbox (causes the wizard to create a custom view object implementation class whose default framework methods you can override). Under this checkbox
   - Uncheck the **Include bind variable accessor** checkbox (the sample application does not use bind variables).
   - Uncheck the **Include custom data source methods** checkbox. (In other applications, you might use this feature to generate signatures for the methods to be overridden, but this sample application provides you with a template class for this task).
   - Click the **Classes Extend** button. In the Override Base Classes dialog box, click the **Browse** button (next to the **Object** field) to open the Search dialog box. In the Search dialog box, type `WsViewObject` in the **Match Class or Package Name** field. When the package

name appears in the **Match Class or Package Name** field, click the name to select it, and then click **OK** to exit the Search dialog box. Click **OK** again to redisplay step 7 of 9 of the wizard.

6. Click **Finish** to close the view object creation wizard. In the WsDepartmentsView.xml editor, click the **Java** menu.
7. Select the link next to the **View Object Class** label to open the class in the Java code editor.
8. Move the cursor anywhere in the code within the Java editor, right-click, and choose **Source -> Override Methods** from the context menu.
9. In the Override Methods dialog box, check the **Group By Class** checkbox to sort the methods by class. Locate **WsViewObject** in the list, and select its **createRowFromResultSet**, **executeQueryForCollection**, and **getQueryHitCount** methods. Be sure to check the **Copy Javadoc** checkbox.
10. Click **OK**.

The generated file is now populated with methods for accessing the JavaBean wrapper class; the file includes comments (from the template's javadoc) that include override instructions.

For your own integration projects, this is the approach you'll take, overriding the methods as needed. For the purposes of this article, I've provided that code for you.

11. Open the **viewObject-code-complete.txt** file, located in the **views** package, and copy and paste its entire contents over the entire **WsDepartmentsViewImpl** class, replacing all the generated content.
12. Save your work, and close all the open tabs.

FINAL STEPS

At this point, you've created a programmatic entity object and a custom view object that uses the entity. To make these components available to integrate the Web service with the Business Components application, you must expose the custom view object by adding it to the application module. You can also then test the integration. Starting from the Application Navigator

1. Expand the **AdfBcModel -> Application Sources** node and the **oramag.sample.application.model.services** package.

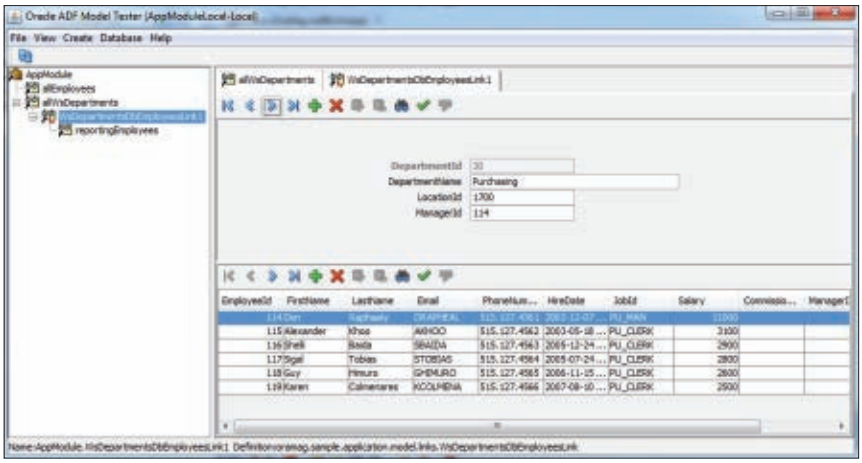

Figure 2: Web services integration application with advanced functionality

2. Right-click the **AppModule** entry, and select **Open AppModule** from the context menu.
3. In the AppModule.xml editor window, click the **Data Model** menu.
4. In the **Available View Objects** field, expand the **oramag.sample.application.model.views** package and select **WsDepartmentsView**.
5. Click the right arrow to move the selected view object to the Data Model list.
6. Right-click the **WsDepartmentsView1** instance, and choose **Rename** from the context menu.
7. Rename the view object instance `allWsDepartments` and click **OK** before saving your work.

Assuming that the Web service you deployed during initial setup is still running, you can test the application by using the Oracle JDeveloper component tester.

Right-click the **AppModule** entry in the Application Navigator and choose **Run** from the context menu to launch the Business Component tester.

In the opened Oracle ADF Model Tester, double-click the **allWsDepartments** instance to read data from the Web service.

Click the arrows to browse and modify the data, and see how the entity and view objects you created populate the UI and make changes to data on the Web service. The Business Components project also contains a pregenerated **Employees** entity and view object that query employee data directly from the database. As an advanced self-study exercise, create an entity association and view link between the **WsDepartments** entity and the **Employees** entity, using the **DepartmentId** attribute to define the parent-child relationship.

The completed-app workspaces include this extra functionality as well as a model-driven list of values on the **DepartmentId** attribute of the **EmployeesView** object (see Figure 2). I encourage you to explore the code and learn about these powerful features.

SUMMARY

Oracle ADF enables application developers to integrate Web services into their Business Components applications at the service layer, using a standard JAX-WS proxy client in conjunction with programmatic entity and view objects. ◀

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

NEXT STEPS

READ more about Oracle ADF
bit.ly/KsyrKB

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

the sample application for this article
bit.ly/I5LIrR

ORACLE DATABASE

Working with Collections

Part 8 in a series of articles on understanding and using PL/SQL

I-HUA CHEN

In the previous article in this series, I showed you how to work with a PL/SQL record, which is a *composite datatype* composed of one or more fields. In this article, I will explore another composite datatype, the collection. An Oracle PL/SQL collection is a single-dimensional array; it consists of one or more elements accessible through an index value.

Collections are used in some of the most important performance optimization features of PL/SQL, such as

- **BULK COLLECT.** SELECT statements that retrieve multiple rows with a single fetch, increasing the speed of data retrieval.
- **FORALL.** Inserts, updates, and deletes that use collections to change multiple rows of data very quickly
- **Table functions.** PL/SQL functions that return collections and can be called in the FROM clause of a SELECT statement.

You can also use collections to work with lists of data in your program that are not stored in database tables.

This article introduces you to collections and gives you a solid foundation in both collection syntax and features.

COLLECTION CONCEPTS AND TERMINOLOGY

Before exploring collections, it is helpful to have a common collections vocabulary that includes the following terms.

Index value. The location of the data in a collection. Index values are usually integers but for one type of collection can also be strings.

Element. The data stored at a specific index value in a collection. Elements in a collection are always of the same type (all of them are strings, dates, or records). PL/SQL collections are *homogeneous*.

Sparse. A collection is sparse if there is at least one index value between the lowest and highest defined index values that is not defined. For example, a sparse collection has an element assigned to index value 1 and another to index value 10 but nothing in between. The opposite of a sparse collection is a *dense* one.

Method. A collection method is a procedure or function that either provides information about the collection or changes the contents of the collection. Methods are attached to the collection variable with dot notation (object-oriented syntax), as in *my_collection*.FIRST.

TYPES OF COLLECTIONS

Collections were first introduced in Oracle7 Server and have been enhanced in several ways through the years and across Oracle Database versions. There are now three types of collections to choose from, each with its own set of characteristics and each best suited to a different circumstance.

Associative array. The first type of collection available in PL/SQL, this was originally called a "PL/SQL table" and can be used only in PL/SQL blocks. Associative arrays can be sparse or dense and can be indexed by integer or string.

Nested table. Added in Oracle8 Database, the nested table can be used in PL/SQL blocks, in SQL statements, and as the datatype of columns in tables. Nested tables can be sparse but are almost always dense. They can be indexed only by integer. You can use the MULTISET operator to perform set operations and to perform equality comparisons on nested tables.

Varray. Added in Oracle8 Database, the varray (variable-size array) can be used in PL/SQL blocks, in SQL statements, and as the datatype of columns in tables. Varrays are always dense and indexed by integer. When a varray type is defined, you must specify the maximum number of elements allowed in a collection declared with that type.

You will rarely encounter a need for a varray (How many times do you know in advance the *maximum* number of elements you will define in your collection?). The associative array is the most commonly used collection type, but nested tables have some powerful, unique features (such as MULTISET operators) that can simplify the code you need to write to use your collection.

NESTED TABLE EXAMPLE

Let's take a look at the simple example in Listing 1, which introduces the many aspects of collections explored later in the article.

When I run the block in Listing 1, I see the following output:

```
Veva
Steven
```

Listing 1 also includes references to the lines in the code block and descriptions of how those lines contribute to the nested table example.

DECLARE COLLECTION TYPES AND VARIABLES

Before you can declare and use a collection variable, you need to define the *type* on which it is based. Oracle Database predefines several collection types in supplied packages such as DBMS_SQL and DBMS_UTILITY. So if you need, for example, to declare an associative array of strings whose maximum length is 32767, you could write the following:

```
l_names DBMS_UTILITY.maxname_array;
```

In most cases, however, you will declare your own application-specific collection types. Here are examples of declaring each of the different types of collections:

1. Declare an associative array of numbers, indexed by integer:

```
TYPE numbers_aat IS TABLE OF NUMBER
   INDEX BY PLS_INTEGER;
```

2. Declare an associative array of numbers, indexed by string:

```
TYPE numbers_aat IS TABLE OF NUMBER
   INDEX BY VARCHAR2(100);
```

3. Declare a nested table of numbers:

```
TYPE numbers_nt IS TABLE OF NUMBER;
```

4. Declare a varray of numbers:

```
TYPE numbers_vat IS VARRAY(10)
OF NUMBER;
```

Note: I use the suffixes _aat, _nt, and _vat, for *associative array type*, *nested table type*, and *varray type*, respectively.

You might be wondering why the syntax for defining a collection type does not use the word *collection*. The answer is that the IS TABLE OF syntax was first introduced in Oracle7 Server, when there was just one type of collection, the PL/SQL table.

From these examples, you can draw the following conclusions about collection types:

- If the TYPE statement contains an INDEX BY clause, the collection type is an associative array.
- If the TYPE statement contains the VARRAY keyword, the collection type is a varray.
- If the TYPE statement does not contain an INDEX BY clause or a VARRAY keyword, the collection type is a nested table.
- Only the associative array offers a choice of indexing datatypes. Nested tables as well as varrays are always indexed by integer.
- When you define a varray type, you specify the maximum number of elements that can be defined in a collection of that type.

Once you've declared a collection type, you can use it to declare a collection variable as you would declare any other kind of variable:

```
DECLARE
   TYPE numbers_nt IS TABLE OF NUMBER;
   l_numbers numbers_nt;
```

Code Listing 1: Nested table example

```
 1  DECLARE
 2     TYPE list_of_names_t IS TABLE OF VARCHAR2 (100);
 3
 4     happyfamily   list_of_names_t := list_of_names_t ();
 5     children      list_of_names_t := list_of_names_t ();
 6     parents       list_of_names_t := list_of_names_t ();
 7  BEGIN
 8     happyfamily.EXTEND (4);
 9     happyfamily (1) := 'Veva';
10     happyfamily (2) := 'Chris';
11     happyfamily (3) := 'Eli';
12     happyfamily (4) := 'Steven';
13
14     children.EXTEND;
15     children (children.LAST) := 'Chris';
16     children.EXTEND;
17     children (children.LAST) := 'Eli';
18
19     parents := happyfamily MULTISET EXCEPT children;
20
21     FOR l_row IN 1 .. parents.COUNT
22     LOOP
23        DBMS_OUTPUT.put_line (parents (l_row));
24     END LOOP;
25  END;
```

Lines	Explanation
2	I declare a new nested table type. Each element in a collection declared with this type is a string whose maximum length is 100.
4–6	I declare three nested tables—happyfamily, children, and parents—based on my new collection type. Note that I also assign a default value to each variable by calling a *constructor function* that has the same name as the type.
8	I "make room" in my happyfamily collection for four elements by calling the EXTEND method.
9–12	I assign the names of the members of my immediate family (my wife, Veva; my two sons, Chris and Eli; and myself). Note the use of typical single-dimension array syntax to identify an element in the array: *array_name (index_value)*.
14–17	Now I populate the children nested table with just the names of my sons. Rather than do a "bulk" extend as on line 8, I extend one index value at a time. Then I assign the name to the just-added index value by calling the LAST method, which returns the highest defined index value in the collection. Unless you know how many elements you need in advance, this approach of extending one row and then assigning a value to the new highest index value is the way to go.
19	Both of my children are adults and have moved out of the ancestral home. So who's left in this place with too many bedrooms? Start with the happyfamily and subtract (with the MULTISET EXCEPT operator) the children. Assign the result of this set operation to the parents collection. It should be just Veva and Steven.
21–24	The result of a MULTISET operation is always either empty or densely filled and starts with index value 1. So I will iterate through all the elements in the collection, from 1 to the COUNT (the number of elements defined in the collection) and display the element found at each index value.

INITIALIZING COLLECTIONS

When you work with nested tables and varrays, you must *initialize* the collection variable before you can use it. You do this by calling the constructor function for that type. This function is created automatically by Oracle Database when you declare the type. The constructor function constructs an instance of the type associated with the function. You can call this function with no arguments, or you can pass it one or more expressions of the same type as the elements of the collection, and they will be inserted into your collection.

Here is an example of initializing a nested table of numbers with three elements (1, 2, and 3):

```
DECLARE
   TYPE numbers_nt IS TABLE OF NUMBER;
   l_numbers numbers_nt;
BEGIN
   l_numbers := numbers_nt (1, 2, 3);
END;
```

If you neglect to initialize your collection, Oracle Database will raise an error when you try to use that collection:

```
SQL> DECLARE
  2    TYPE numbers_nt IS TABLE OF
NUMBER;
  3    l_numbers numbers_nt;
  4  BEGIN
  5    l_numbers.EXTEND;
  6    l_numbers(1) := 1;
  7  END;
  8  /
DECLARE
*
ERROR at line 1:
ORA-06531: Reference to uninitialized
collection
ORA-06512: at line 5
```

You do *not* need to initialize an associative array before assigning values to it.

POPULATING COLLECTIONS

You can assign values to elements in a collection in a variety of ways:

- Call a constructor function (for nested tables and varrays).
- Use the assignment operator, for single elements as well as entire collections.
- Pass the collection to a subprogram as an OUT or IN OUT parameter, and then assign the value inside the subprogram.
- Use a BULK COLLECT query.

The previous section included an example that used a constructor function. Following are examples of the other approaches:

1. Assign a number to a single index value. Note that with an associative array, it is not necessary to use EXTEND or start with index value 1.

```
DECLARE
   TYPE numbers_aat IS TABLE OF
NUMBER
      INDEX BY PLS_INTEGER;
   l_numbers numbers_aat;

BEGIN
   l_numbers (100) := 12345;
END;
```

2. Assign one collection to another. As long as both collections are declared with the same type, you can perform collection-level assignments.

```
DECLARE
   TYPE numbers_aat IS TABLE OF
NUMBER
      INDEX BY PLS_INTEGER;
   l_numbers1 numbers_aat;
   l_numbers2 numbers_aat;
BEGIN
   l_numbers1 (100) := 12345;
   l_numbers2       := l_numbers1;
END;
```

3. Pass a collection as an IN OUT argument, and remove all the elements from that collection:

```
DECLARE
   TYPE numbers_aat IS TABLE OF
NUMBER
      INDEX BY PLS_INTEGER;
   l_numbers   numbers_aat;

   PROCEDURE empty_collection (
      numbers_io IN OUT numbers_aat)
   IS
   BEGIN
      numbers_io.delete;
   END;
BEGIN
   l_numbers (100) := 123;
   empty_collection (l_numbers);
END;
```

4. Fill a collection directly from a query with BULK COLLECT (covered in more detail in the next article in this series):

```
DECLARE
   TYPE numbers_aat IS TABLE OF
NUMBER
      INDEX BY PLS_INTEGER;

   l_numbers   numbers_aat;
BEGIN
     SELECT employee_id
       BULK COLLECT INTO l_numbers
       FROM employees
   ORDER BY last_name;
END;
```

ITERATING THROUGH COLLECTIONS

A very common collection operation is to iterate through all of a collection's elements. Reasons to perform a "full collection scan" include displaying information in the collection, executing a data manipulation language (DML) statement with data in the element, and searching for specific data.

The kind of code you write to iterate through a collection is determined by the type of collection with which you are working and how it was populated. Generally, you will choose between a numeric FOR loop and a WHILE loop.

Use a numeric FOR loop when

- Your collection is densely filled (every index value between the lowest and the highest is defined)
- You want to scan the entire collection, not terminating your scan if some condition is met

Conversely, use a WHILE loop when

- Your collection may be sparse
- You might terminate the loop before you have iterated through all the elements in the collection

You should use a numeric FOR loop with dense collections to avoid a NO_DATA_FOUND exception. Oracle Database will also raise this exception, however, if you try to "read" an element in a collection at an undefined index value.

The following block, for example, raises a NO_DATA_FOUND exception:

```
DECLARE
```

Answer to the Challenge

The PL/SQL Challenge question in last issue's "Working with Records in PL/SQL" article tested your knowledge of how to declare a record variable based on a table or a cursor. The question asked which of the following could be used in the question's code block so that a value ("Keyboard") from the question's table would be displayed. All the choices are listed below; only (c) and (d) are correct.

a.

```
l_part   plch_parts%TYPE;
```

b.

```
l_part   plch_parts;
```

c.

```
l_part   plch_parts%ROWTYPE;
```

d.

```
CURSOR parts_cur
  IS
     SELECT * FROM plch_parts;

  l_part   parts_cur%ROWTYPE;
```

```
  TYPE numbers_aat IS TABLE OF NUMBER
      INDEX BY PLS_INTEGER;
  l_numbers numbers_aat;
BEGIN
  DBMS_OUTPUT.PUT_LINE (l_numbers (100));
END;
```

When, however, you know for certain that your collection is—and will always be—densely filled, the FOR loop offers the simplest code for getting the job done. The procedure in Listing 2, for example, displays all the strings found in a collection whose type is defined in the DBMS_UTILITY package.

This procedure calls two *methods*: FIRST and LAST. FIRST returns the lowest defined index value in the collection, and LAST returns the highest defined index value in the collection.

The following block will display three artists' names; note that the index values do *not* need to start at 1.

```
DECLARE
   l_names   DBMS_UTILITY.maxname_array;
BEGIN
   l_names (100) := 'Picasso';
   l_names (101) := 'O''Keefe';
   l_names (102) := 'Dali';=
   show_contents (l_names);
END;
/
```

If your collection may be sparse or you want to terminate the loop conditionally, a WHILE loop will be the best fit. The procedure in Listing 3 shows this approach.

In this procedure, my iterator (l_index) is initially set to the lowest defined index value. If the collection is empty, both FIRST and LAST will return NULL. The WHILE loop terminates when l_index is NULL. I then display the name at the current index value and call the NEXT method to get the next defined index value higher than l_index. This function returns NULL when there is no higher index value.

I call this procedure in the following block, with a collection that is not sequentially filled. It will display the three names without raising NO_DATA_FOUND:

```
DECLARE
   l_names   DBMS_UTILITY.maxname_array;
BEGIN
   l_names (-150) := 'Picasso';
   l_names (0) := 'O''Keefe';
```

Take the Challenge!

Each PL/SQL 101 article offers a quiz to test your knowledge of the information provided in it. The quiz appears below and also at PL/SQL Challenge (plsqlchallenge.com), a Website that offers online quizzes on the PL/SQL language as well as SQL and Oracle Application Express.

Question
Which of the following blocks will display these three lines after execution:

```
Strawberry
Raspberry
Blackberry
```

a.

```
DECLARE
   l_names   DBMS_UTILITY.maxname_array;
BEGIN
   l_names (1) := 'Strawberry';
   l_names (10) := 'Blackberry';
   l_names (2) := 'Raspberry';

   FOR indx IN 1 .. l_names.COUNT
   LOOP
      DBMS_OUTPUT.put_line (l_names (indx));
   END LOOP;
END;
/
```

b.

```
DECLARE
   l_names   DBMS_UTILITY.maxname_array;
BEGIN
   l_names (1) := 'Strawberry';
   l_names (10) := 'Blackberry';
   l_names (2) := 'Raspberry';

   indx := l_names.FIRST;

   WHILE (indx IS NOT NULL)
   LOOP
      DBMS_OUTPUT.put_line (l_names (indx));
      indx := l_names.NEXT (indx);
   END LOOP;
END;
/
```

c.

```
DECLARE
   l_names   DBMS_UTILITY.maxname_array;
BEGIN
   l_names (1) := 'Strawberry';
   l_names (10) := 'Blackberry';
   l_names (2) := 'Raspberry';

   DECLARE
      indx   PLS_INTEGER := l_names.FIRST;
   BEGIN

      WHILE (indx IS NOT NULL)
      LOOP
         DBMS_OUTPUT.put_line (l_names (indx));
         indx := l_names.NEXT (indx);
      END LOOP;
   END;
END;
/
```

d.

```
DECLARE
   l_names   DBMS_UTILITY.maxname_array;
BEGIN
   l_names (1) := 'Strawberry';
   l_names (10) := 'Blackberry';
   l_names (2) := 'Raspberry';

   FOR indx IN l_names.FIRST .. l_names.LAST
   LOOP
      DBMS_OUTPUT.put_line (l_names (indx));
   END LOOP;
END;
/
```

```
   l_names (307) := 'Dali';
   show_contents (l_names);
END;
/
```

I can also scan the contents of a collection in reverse, starting with LAST and using the PRIOR method, as shown in Listing 4.

DELETING COLLECTION ELEMENTS

PL/SQL offers a DELETE method, which you can use to remove all, one, or some elements from a collection. Here are some examples:

1. Remove all elements from a collection; use the DELETE method without any arguments. This form of DELETE works with all three kinds of collections.

```
l_names.DELETE;
```

2. Remove the first element in a collection; to remove one element, pass the index value to DELETE. This form of DELETE can be used only with an associative array or a nested table.

```
l_names.DELETE (l_names.FIRST);
```

3. Remove all the elements between the specified low and high index values. This form of DELETE can be used only with an associative array or a nested table.

```
l_names.DELETE (100, 200);
```

If you specify an undefined index value, Oracle Database will *not* raise an error.

You can also use the TRIM method with varrays and nested tables to remove elements from the end of the collection. You can trim one or many elements:

```
l_names.TRIM;
l_names.TRIM (3);
```

GET COMFY WITH COLLECTIONS

It is impossible to take full advantage of PL/SQL, including some of its powerful features, if you do not use collections. This article has provided a solid foundation for working with collections, but there are still several advanced features to explore, including string-indexed and nested collections, which will be covered in a future article.

The next article in this PL/SQL 101 series will explore how to *use* collections with PL/SQL's most important performance-related PL/SQL features: FORALL and BULK COLLECT. ◀

Code Listing 2: Display all strings in a collection

```
CREATE OR REPLACE PROCEDURE show_contents (
   names_in IN DBMS_UTILITY.maxname_array)
IS
BEGIN
   FOR indx IN names_in.FIRST .. names_in.LAST
   LOOP
      DBMS_OUTPUT.put_line (names_in (indx));
   END LOOP;
END;
/
```

Code Listing 3: Use WHILE to iterate through a collection

```
CREATE OR REPLACE PROCEDURE show_contents (
   names_in IN DBMS_UTILITY.maxname_array)
IS
   l_index   PLS_INTEGER := names_in.FIRST;
BEGIN
   WHILE (l_index IS NOT NULL)
   LOOP
      DBMS_OUTPUT.put_line (names_in (l_index));
      l_index := names_in.NEXT (l_index);
   END LOOP;
END;
/
```

Code Listing 4: Scan a collection in reverse

```
CREATE OR REPLACE PROCEDURE show_contents (
   names_in IN DBMS_UTILITY.maxname_array)
IS
   l_index   PLS_INTEGER := names_in.LAST;
BEGIN
   WHILE (l_index IS NOT NULL)
   LOOP
      DBMS_OUTPUT.put_line (names_in (l_index));
      l_index := names_in.PRIOR (l_index);
   END LOOP;
END;
/
```

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 1–7
bit.ly/fcOuoJ

READ more Feuerstein
stevenfeuerstein.com
toadworld.com/sf

Oracle PL/SQL Programming
amzn.to/IDhqef

Oracle PL/SQL Language Pocket Reference
amzn.to/KuEMl6

Oracle PL/SQL Best Practices
amzn.to/JTXAcg

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

Beginning Performance Tuning

Resolve session performance issues in Oracle Database.

It's the middle of the night, and you get a frantic call from someone complaining that the database is slow. The caller demands to know why—and what you're going to do about it. Sound familiar? If it does, you are not alone. High performance is a common expectation of database system users: they get very unhappy when they don't get it, and they are usually not shy about letting you know. What should you do next? In this article, you will learn some techniques for troubleshooting Oracle Database performance issues.

To use the scripts in this article, you need to create some tables in a test schema and access some dynamic performance views. The database user SYS has all privileges to access the views, so you need the password for the SYS user. The script for setting up the example tables is available in the online version of this article, at bit.ly/IvwqLE.

SESSION STATE

Before you start troubleshooting why a database is slow, you have to first understand that the database itself is never slow or fast—it has a constant speed. The *sessions* connected to the database, however, slow down when they hit a bump in the road. To resolve a session performance issue, you need to identify the bump and remove it. Fortunately, it's very easy to do this in most cases.

The first step to resolving a session performance issue is to ascertain what that database session is doing now. An Oracle Database session is always in one of three states:

1. **Idle**. Not doing anything—just waiting to be given some work.
2. **Processing**. Doing something useful—running on the CPU.
3. **Waiting**. Waiting for something, such as a block to come from disk or a lock to be released.

If a session is waiting to be given work (idle), it's really not slow at all—it just has nothing to do. If a session is waiting for some resource, such as a block or a lock, it has stopped processing. Until it gets that resource, the session will continue to wait. When it gets that resource, it does some processing and then moves on to the next resource it needs, waits for that to be available, and starts processing . . . and the cycle goes on until the session has nothing else to do. If it waits for resources often, the session will appear slow. But it's not really slow—it's just following a pattern of go, stop, go again, stop again, and so on. Your mission is to find and eliminate the "stop" issues in the session.

How difficult is it to get information about what's causing the session to stop? It's actually *very* easy: Oracle Database is instrumented to talk about what the database sessions are doing. All you need to do is to listen attentively or, more precisely, look for that information in the right place, and that place is a view called V$SESSION. Everything you need for your analysis is in this view.

To explain how to use the V$SESSION view, I will use a very common scenario—row locking—as an example. To follow along, first set up the previously mentioned tables as described in the online version of this article. Then connect as user ARUP from two different sessions. From the first session, issue the following SQL statement:

```
update t1
set col2 = 'x' where col1 = 1;
```

The output will show "1 row updated," indicating that the row was updated. Do not issue a COMMIT after the statement. By not committing, you will force the session to get and hold a lock on the first row of the T1 table. Now go to the second session and issue the following SQL statement:

```
update t1
set col2 = 'y'
where col1 = 1;
```

This statement will hang. Why? The answer is simple: the first session holds a lock on the row, which causes the second session to hang and the user to complain that the session is slow. To know what the second session is doing, the first thing you need to check is the STATE column in V$SESSION:

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

SID   STATE
————  ————————————————
3346  WAITING
2832  WAITED KNOWN TIME
```

Study the output carefully. Session 3346 (in the SID column) indicates that it is waiting for something—and therefore not working. That should be your first clue that the session is experiencing one of those performance bumps in the road. But before you can determine what the session is waiting for, let's study the state of session 2832 in the output, which shows that it waited for some known amount of time earlier. The important point is that session 2832 is *not waiting* for anything right now, meaning that it's working productively.

Next, let's see what the second session (3346) is waiting for. That information is readily available in the EVENT column in the same V$SESSION view. The EVENT column not only shows an event a session is waiting for currently, but also shows an event a session has waited for earlier. The query against V$SESSION in Listing 1 displays information from the EVENT column for both sessions.

The output in Listing 1 shows that session 3346 is waiting right now for an event: "enq: TX – row lock contention"—short for "enqueue for transaction-level lock on row" or, in plain English, a row-level lock. The session is waiting because it wants to lock one or more rows, but another session has already placed locks on the row or rows. Unless that other session commits or rolls back its transaction, session 3346 will not get the lock it needs and will have no choice but to wait. On the other hand, the state of session 2832, "WAITED KNOWN TIME," means that it is *working*—not waiting—right now. It was, however, waiting earlier for an event called "SQL*Net message from client" (I will discuss this specific event later.) There is one very important lesson in these results: you cannot look at the EVENT column alone to find out what the session is waiting for. You must look at the STATE column first to determine whether the session is waiting or working and then inspect the EVENT column.

After you determine that a session is waiting for something, the next thing you need to find out is how long the session has been waiting. A very long wait usually indicates some sort of bottleneck. Where can you get information on the length of the waiting period? The answer is right there in the V$SESSION view, in the SECONDS_IN_WAIT column.

Getting the amount of time a session has been waiting makes sense for sessions that are waiting *right now*, but what about the sessions that are working now? Recall that the EVENT column shows not only the event a session is experiencing now but also the last wait event the session has experienced. Another column—WAIT_TIME—in the same V$SESSION view shows how long that wait lasted. (Note that WAIT_TIME is shown in centiseconds [hundredths of a second].)

Now that you know how to get information on the sessions waiting and working, let's put all the information together in a single query, shown in Listing 2. It clearly shows the state of the sessions: whether they are working or waiting; if they are working, what they were waiting for earlier and for how long; and if they are waiting, what for and for how long.

Code Listing 1: Query for displaying sessions, session state, and events

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

 SID  STATE               EVENT
----  ------------------  ----------------------------------
2832  WAITED KNOWN TIME   SQL*Net message from client
3346  WAITING             enq: TX - row lock contention
```

Code Listing 2: Query for displaying sessions, session state, and wait details

```
col "Description" format a50
select sid,
       decode(state, 'WAITING','Waiting',
              'Working') state,
       decode(state,
              'WAITING',
              'So far '||seconds_in_wait,
              'Last waited '||
              wait_time/100)||
       ' secs for '||event
       "Description"
from v$session
where username = 'ARUP';
```

Output:

```
 SID  STATE     Description
----  --------  --------------------------------------------------
2832  Working   Last waited 2029 secs for SQL*Net message from client
3346  Waiting   So far 743 secs for enq: TX - row lock contention
4208  Waiting   So far 5498 secs for SQL*Net message from client
```

IDLE EVENT

Note the details of session 4208 in Listing 2; it's currently waiting 5,498 seconds for a "SQL*Net message from client" event. Recall from the previous section that an Oracle Database session can be in one of the three states: working, waiting for a resource, or waiting for work. But how can a session determine whether it is idle? It will expect to be given work by clients via SQL*Net, but there is no way for it to know in advance if any work is coming from the clients. All it can do is wait for some instruction coming through SQL*Net. Until then, it will have nothing else to do but eagerly stare at the SQL*Net interface, and this condition is reported as "SQL*Net message from client" in the V$SESSION view's EVENT column, which is practically the same thing as just being idle.

You can disregard another EVENT column value, "rdbms ipc message," because it is an event status for sessions that are idle. Note that an idle session does not show IDLE as the STATE column value; it still shows "Waiting." You have to check the EVENT column to determine whether the session is truly idle.

You may be tempted to modify the query in Listing 2 to filter sessions that include the "SQL*Net message from client" and "rdbms ipc message" idle events. Although you can do that, I highly discourage doing that, for multiple reasons. First, not all instances of the "SQL*Net message from client" event indicate that the session is idle. Consider the possibility that the network might be truly slow, in which case the session will also wait for these events. Remember, the session doesn't have the ability to determine whether the client is truly idle or is sending instructions that are slow or stuck in the network. All it can do is wait, and it will wait with the "SQL*Net message from client" event. Second, idle events may provide some clues to Oracle Support about what else is going on inside a session. So I recommend displaying these "idle" EVENT values.

DIAGNOSIS OF LOCKING

The output of Listing 2 provides enough information to enable you to make a diagnosis about the performance of these three sessions. Session 4208 is idle, so any complaints that session 4208 is slow just aren't related to the database. Any performance issues related to this session could be related to a bug in the code that's going through an infinite loop or high CPU consumption on the application server. You can redirect the performance troubleshooting focus toward the application client.

The story of session 3346 is different. This session is truly a bottleneck to the application. Now that you know why this session appears slow—it is waiting for a row lock—the next logical question is which session holds that lock. The answer is also found in—I hope you guessed it—the V$SESSION view, in, more specifically, the BLOCKING_SESSION column. (Note that in an Oracle Real Application Clusters [Oracle RAC] environment, the blocking session may exist in a different instance. In such a case, the blocking instance is displayed in the V$SESSION view's BLOCKING_INSTANCE column.)

You can find out the blocking session and instance by issuing the following SQL statement:

```
select
  blocking_session B_SID,
  blocking_instance B_Inst
from v$session
where sid = 3346;

B_SID  B_INST
_____  ______
 2832       1
```

The output shows clearly that SID 2832 is holding the lock that SID 3346 is waiting for. Now you can follow a cause/effect relationship between the session in which an update to a row is being blocked and the session that holds the lock on that row.

You can find the specific row that is locked by first finding the table containing that row. To find that table, use the same V$SESSION view; in this case, the information is in the ROW_WAIT_OBJ# column, which shows the object number of the table whose row is being locked. You can then get the name of the table from the DBA_OBJECTS view, using this object number, as shown in Listing 3.

Code Listing 3: Getting row lock information

```
select row_wait_obj#,
       row_wait_file#,
       row_wait_block#,
       row_wait_row#
from v$session
where sid = 3346;

ROW_WAIT_OBJ#  ROW_WAIT_FILE#  ROW_WAIT_BLOCK#  ROW_WAIT_ROW#
_____________  _____________  _______________  _____________
       241876           1024          2307623              0
```

To get the object information:

```
select owner, object_type, object_name
from dba_objects
where data_object_id = 241876;

OWNER  OBJECT_TYPE  OBJECT_NAME
_____  ___________  ___________
ARUP   TABLE        T1
```

Code Listing 4: Finding the row information

```
REM Filename: rowinfo.sql
REM This shows the row from the table when the
REM components of ROWID are passed. Pass the
REM following in this exact order
REM  1. owner
REM  2. table name
REM  3. object_id
REM  4. relative file ID
REM  5. block ID
REM  6. row Number
REM
select *
from &1..&2
where rowid =
        dbms_rowid.rowid_create (
                rowid_type      =>  1,
                object_number   => &3,
                relative_fno    => &4,
                block_number    => &5,
                row_number      => &6
        )
/

SQL> @rowinfo ARUP T1 241876 1024 2307623 0

COL1  C
____  _
   1  x
```

The output shows that some row in the T1 table is the point of the row lock contention. But which specific row is locked? That data is available in three V$SESSION view columns—ROW_WAIT_FILE#, ROW_WAIT_BLOCK#, and ROW_WAIT_ROW#—which show the relative file ID, the block ID in that file, and the row's slot number inside that block, respectively, for that specific row. Using this information, you can identify the ROWID of the row. The ROWID, the physical address of every row in an Oracle Database instance, can be used to uniquely identify a row.

Listing 4 shows a SQL script that enables you to select the specific blocking row from the table with the information gathered so far. Save this script in a file named rowinfo.sql. The script expects the input in the following order: owner, table name, object#, file#, block#, and row#. You can call this script and pass all the requested parameters by copying and pasting the corresponding output from Listing 3.

The output in Listing 4 shows the specific row on which a lock is being requested but that is locked by another session. So far you have identified not only the source session of the locking but the specific row being locked as well.

Is it possible that the session holding the lock (SID 2832) is somehow disconnected from the client? That can occur in connection pools or when users access the database with thick-client tools such as Oracle SQL Developer. After you identify the session holding the lock, you may want to wait until it commits or rolls back the transaction. Either action releases the lock.

In the case of a dead connection, you may alternatively decide to kill the session, which will force a rollback releasing the locks held by the blocking session and enabling the waiting sessions to continue. Occasionally the problem can be pretty simple: for instance, someone issued an UPDATE statement from a thick-client tool but forgot to commit and thus caused every session to wait for those updated rows. Identifying that blocking session enables you to send a gentle reminder to rectify that situation immediately.

Code Listing 5: Sessions from a specific user

```
select SID, osuser, machine, terminal, service_name,
       logon_time, last_call_et
from v$session
where username = 'ARUP';

SID   OSUSER  MACHINE   TERMINAL  SERVICE_NAME  LOGON_TIME  LAST_CALL_ET
----  ------  --------  --------  ------------  ----------  ------------
3346  oradb   prodb1    pts/5     SYS$USERS     05-FEB-12          6848
2832  oradb   prodb1    pts/6     SERV1         05-FEB-12          7616
4408  ANANDA  ANLAP     ANLAP     ADHOC         05-FEB-12             0
```

OSUSER. The operating system user as which the client is connected. The output indicates that session 4408 is connected from the ANLAP machine, where a Windows user, ANANDA, has logged in.

MACHINE. The name of the machine where the client is running. This could be the database server itself. For two of the sessions, the machine name shows up as "prodb1." Session 4408 runs on a different machine—ANLAP—presumably a laptop.

TERMINAL. If the session is connected from a UNIX server, this is the terminal where it runs.

LOGON_TIME. This shows when the session was first connected to the Oracle Database instance.

LAST_CALL_ET. This shows when the session last issued some SQL. The output indicates that session 3346 made its last SQL call 6,848 seconds ago.

Code Listing 6: Session waits for a specific machine

```
col username format a5
col program format a10
col state format a10
col last_call_et head 'Called|secs ago' format 999999
col seconds_in_wait head 'Waiting|for secs' format 999999
col event format a50
select sid, username, program,
       decode(state, 'WAITING', 'Waiting',
              'Working') state,
last_call_et, seconds_in_wait, event
from v$session
where machine = 'appsvr1'
/
                                          Called    Waiting
SID   USERNAME  PROGRAM      STATE    secs ago  for secs  EVENT
----  --------  -----------  -------  --------  --------  ----------------------
2832  ARUP      sqlplus.exe  Waiting       152       151  SQL*Net message
                                                         from client
3089  ARUP      sqlplus.exe  Waiting       146       146  enq: TX - row lock
                                                         contention
3346  ARUP      sqlplus.exe  Working        18        49  SQL*Net message
                                                         from client
```

MORE ON THE SESSION

In many troubleshooting situations, just knowing the SID of each session is not enough. You may need to know other details, such as the client machine the session is connecting from, the user (of both the database and the operating system), and the service name. All of this information is also readily available in the same V$SESSION view you have been using. Let's briefly examine the columns that provide that information, by running the script shown in Listing 5.

Using the columns shown in Listing 5, you can get very detailed information on a user's sessions.

Suppose you receive a complaint that the applications running on the application server named appsvr1 are experiencing performance issues. Listing 6 shows a query against the V$SESSION view—including columns you've used in previous queries in this article—for the sessions connected from that machine and the output.

From the output, you can easily see that three sessions are connected from the appsvr1 application server. All of them are running SQL*Plus (as shown in the PROGRAM column). SID 3346 is the only one that is working (indicated by "Working" in the STATE column). Because it's working, the EVENT column shows the last time the session waited. The wait time in this case is meaningless, because the session is not waiting but actually working. The "Called secs ago" column (representing the "last_call_et" column in V$SESSION) displays 18,

which means that the session made a SQL call 18 seconds ago.

The other sessions are waiting. SID 3089 is waiting for a row lock. From the output, you can see that the session has been waiting for 146 seconds and that it also made its last SQL call 146 seconds ago. This indicates that the session has been waiting for that particular lock ever since it made that SQL call.

Finally, session 2832 is also waiting; in this case, it is waiting with a "SQL*Net message from client" event, which means it is idle, waiting to be given some work. The session issued its last SQL statement 152 seconds ago and has been idle for 151 seconds.

Armed with this information, you can diagnose performance issues very accurately. You can tell the complaining user that of the three sessions connected from the appsvr1 application server, one session is idle, one is working, and one is waiting for a lock. The user is probably referring to the slowness of this last session. Now you know the reason and how you can rectify it.

GETTING THE SQL

Another key piece of performance tuning information is the SQL statement a session is executing, which will provide more insights into the workings of the session. The same V$SESSION view also shows the SQL statement information. The SQL_ID column in the V$SESSION view shows the ID of the last SQL statement executed. You can get the text of that SQL statement from the V$SQL view, using the SQL_ID value. Here is an example of how I have identified the SQL statement executed by the session that appears slow to the user.

```
select sql_id
from v$session
where sid = 3089;

SQL_ID
-------------
g0uubmuvk4uax

set long 99999
select sql_fulltext
from v$sql
where sql_id = 'g0uubmuvk4uax';

SQL_FULLTEXT
--------------------------------------------------
update t1 set col2 = 'y' where col1 = 1
```

DATA ACCESS ISSUES

I have used row-level locking as the cause of a slowdown in this article. Although locking-related contention is a very common cause, it is not the only cause of performance problems. Another major cause of contention is disk I/O. When a session retrieves data from the database data files on disk to the buffer cache, it has to wait until the disk sends the data. This wait shows up for that session as "db file sequential read" (for index scans) or "db file scattered read" (for full-table scans) in the EVENT column, as shown below:

```
select event
from v$session
where sid = 3011;

EVENT
----------------------------
db file sequential read
```

When you see this event, you know that the session is waiting for I/O from the disk to complete. To make the session go faster, you have to reduce that waiting period. There are several ways to reduce the wait:

1. Reduce the number of blocks retrieved by the SQL statement. Examine the SQL statement to see if it is doing a full-table scan when it should be using an index, if it is using a wrong index, or if it can be rewritten to reduce the amount of data it retrieves.
2. Place the tables used in the SQL statement on a faster part of the disk.
3. Consider increasing the buffer cache to see if the expanded size will accommodate the additional blocks, therefore reducing the I/O and the wait.
4. Tune the I/O subsystem to return data faster.

The online article at bit.ly/IvwqLE includes additional information on how to address performance issues related to disk I/O.

CONCLUSION

In summary, this article presented the following steps for starting a successful performance tuning session:

1. Check whether the session is working or waiting. If the latter, determine what it is waiting for and how long it has been waiting.
2. Compare the waiting period of the session with how long ago it issued a SQL call.
3. If the cause of the wait is a lock contention, find the session holding the lock and get the details of the session. (If the session holding the lock is an orphan session, you may want to kill it to release the lock.)
4. Find the SQL statement the session is executing.
5. If the session is waiting for I/O, find out which segment (table, materialized view, index, and so on) the I/O is waiting for.

The techniques presented in this article will help you resolve about 20 percent of the performance issues you encounter as a DBA. Oracle Database is instrumented to provide information on its inner workings so that you can zero in on the exact cause of an issue—all you have to do is listen.

I sincerely hope that this article has helped you realize how simple it is to diagnose some common but seemingly thorny performance issues in Oracle Database by identifying the right information sources. Happy tuning! ◀

Arup Nanda (arup@proligence.com) has been an Oracle DBA for more than 16 years, 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.

NEXT STEPS

READ online-only article content
bit.ly/IvwqLE

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

Oracle Database Performance Tuning Guide 11g Release 2 (11.2)
bit.ly/IBPsBW

ORACLE DATABASE

A Function of Character

Part 6 in a series on the basics of the relational database and SQL

Part 5 in this series, "An Order of Sorts" (*Oracle Magazine*, May/June 2012), introduced the ORDER BY clause of a SQL SELECT statement (or query) and how it behaves in conjunction with certain options and keywords to order (or *sort*) the data in query results. Now you are ready to start learning how to use SQL *functions* in your queries to transform result set data so that it displays differently from how it is stored in the database. This article focuses on the SQL *character functions* (also known as *string functions* or *text functions*), which enable you to manipulate how character data is displayed. Subsequent articles in this series will introduce the *date* and *number* functions.

To try out the examples in this and subsequent articles in the series, you need access to an Oracle Database instance. If necessary, download and install an Oracle Database edition for your operating system from bit.ly/fherki. I recommend installing Oracle Database, Express Edition.

If you install the Oracle Database software, choose the installation option that enables you to create and configure a database. A new database, including sample user accounts and their associated schemas, will be created for you. (Note that SQL_101 is the user account you'll use for the examples in this series; it's also the schema in which you'll create database tables and other objects.) When the installation process prompts you to specify schema passwords, enter and confirm passwords for SYS and SYSTEM and make a note of them.

Finally—whether you installed the database software from scratch or have access to an existing Oracle Database instance—download, unzip, and execute the SQL script available at bit.ly/JsDOX2 to create the tables for the SQL_101 schema that are required for this article's examples. (View the script in a text editor for execution instructions.)

I-HUA CHEN

PRETTY PRINTING

The most-basic character functions enable you to change the way alphanumeric data is formatted in a result set. For example, the simple query in Listing 1 obtains all unique last name values from the EMPLOYEE table and displays them in all capital letters. It does this by applying the UPPER character func-

Code Listing 1: Query that lists every unique last name value in uppercase letters

```
SQL> set feedback on
SQL> select distinct UPPER(last_name) "Uppercase Employee Last Name"
  2    from employee
  3  order by UPPER(last_name);

Uppercase Employee Last Name
------------------------------------
ECKHARDT
FRIEDLI
JAMES
LEBLANC
MICHAELS
NEWTON
PETERSON

7 rows selected.
```

Code Listing 2: Query that displays all department locations in lowercase letters

```
SQL> select name, LOWER(location) "Lowercase Department Location"
  2    from department
  3  order by location;

NAME                           Lowercase Department Location
------------------------------ ------------------------------------
Accounting                     los angeles
Payroll                        new york

2 rows selected.
```

Code Listing 3: Query that shows certain names converted and with the initial letter capitalized

```
SQL> set lines 32000
SQL> select first_name, last_name, INITCAP(first_name) "First Name",
INITCAP(last_name) "Last Name"
  2    from employee
  3   where employee_id in (6569, 6570);

FIRST_NAME           LAST_NAME            First Name           Last Name
-------------------- -------------------- -------------------- --------------------
michael              peterson             Michael              Peterson
mark                 leblanc              Mark                 Leblanc

2 rows selected.
```

tion to the LAST_NAME column. Similarly, the query in Listing 2 uses the LOWER character function to display all department location names from the DEPARTMENT table in lowercase letters. All functions take some kind of input parameter(s). Character functions require input parameters that are alphanumeric—either a *character* (or *string*) *literal* or a column with a VARCHAR2, CHAR, or CLOB datatype. The data in the EMPLOYEE table's LAST_NAME column and the DEPARTMENT table's LOCATION column is stored with a datatype of VARCHAR2. Recall that a literal character value is any list of alphanumeric characters enclosed in single quotation marks, such as 'Smith', '73abc', or '15-MAR-1965'.

Listings 3 and 4 demonstrate the INITCAP function. The query in Listing 3 uses INITCAP to convert certain first and last names from being stored in all lowercase in the EMPLOYEE table to being displayed with initial capital letters. The INITCAP function capitalizes the first letter of a string *and* lowercases the remainder of the string, as demonstrated by the query in Listing 4. That query also shows that the input parameter for an INITCAP function can consist of a character function's application to a string or a database column that stores alphanumeric data. Specifically, the query applies the UPPER function to the LAST_NAME column of the EMPLOYEE table for certain employees. The UPPER function is said to be *nested* inside the INITCAP function. The Oracle Database server applies nested functions in order, from innermost function to outermost function. In Listing 4, the UPPER function converts the values peterson and leblanc to PETERSON and LEBLANC. Then the INITCAP function converts those uppercase values to Peterson and Leblanc.

Code Listing 4: Query that demonstrates the INITCAP function

```
SQL> select INITCAP('eMPLOYEE lAST nAMES') "INITCAP Literal",
INITCAP(UPPER(last_name)) "Converted Employee Last Name"
  2    from employee
  3   where employee_id in (6569, 6570);

INITCAP Literal      Converted Employee Last Name
-------------------  ---------------------------------
Employee Last Names  Peterson
Employee Last Names  Leblanc

2 rows selected.
```

Code Listing 5: Query that applies the RPAD and the LPAD functions

```
SQL> select RPAD(name, 15, '.') department, LPAD(location, 15, '.') location
  2    from department;

DEPARTMENT      LOCATION
--------------  ---------------
Accounting..... ....LOS ANGELES
Payroll........ .......NEW YORK

2 rows selected.
```

Code Listing 6: Query that demonstrates the CONCAT function

```
SQL> select CONCAT(first_name, last_name) employee_name
  2    from employee
  3  order by employee_name;

EMPLOYEE_NAME
--------------------------------------------------------------------------------
BetsyJames
DonaldNewton
EmilyEckhardt
FrancesNewton
MatthewMichaels
RogerFriedli
markleblanc
michaelpeterson

8 rows selected.
```

PADDING YOUR RESULTS

To *pad* something is to add to it. The LPAD and RPAD functions enable you to pad your character-data results by repeating a character, space, or symbol to the left or right of any string. LPAD pads to the left of a string; RPAD pads to the right.

Listing 5 demonstrates the power of the RPAD and LPAD functions. Note that each takes three input parameters: the column name or string literal you want to pad; the length to which the string should be padded; and the character, space, or symbol (the *filler*) to pad with. For example, the query in Listing 5 specifies that the department name should be right-padded to a total length of 15 with the "." filler character. If any department name is exactly 15 characters or longer, no filler character will be added. Because Accounting is 10 characters long, the RPAD function adds five filler characters to its right. The query also specifies that the location should be left-padded to a total length of 15. Because LOS ANGELES is 11 characters long, counting the space, the LPAD function adds four filler characters to its left.

THE HELPFUL DUAL

Oracle Database provides a single-row, single-column table called DUAL that is useful for many purposes, not the least of which is learning about Oracle functions. DUAL is an Oracle system table owned by the SYS user, not the SQL_101 schema. Many Oracle system tables are made available to all users via *public synonyms*. Synonyms will be discussed in subsequent articles in this series.

The DUAL table contains no data that's useful in and of itself. (It has one row with

one column—called the DUMMY column—that contains the value X.) You can use DUAL to try out functions that work on string literals and, as you'll see in subsequent articles in this series, on number literals and even on today's date.

The following demonstrates the single-row, single-column output of a SELECT statement executed against the DUAL table:

```
SQL> select *
  2    from dual;

D
-
X

1 row selected.
```

To display the current date, you can query the DUAL table as follows:

```
SQL> select sysdate
  2    from dual;

SYSDATE
---------
18-APR-12

1 row selected.
```

And finally, the following example shows how you can practice any function in the SELECT clause of a SQL statement, using the DUAL table:

```
SQL> select rpad('Melanie', 10, '*')
Melanie, lpad('Caffrey', 10, '.')
Caffrey
  2    from dual;

MELANIE    CAFFREY
---------- ----------
Melanie***  ...Caffrey

1 row selected.
```

Note that functions work even though there is no usable data in DUAL. In the preceding examples, the SYSDATE function displays the current date and time of the operating system hosting the database, and the RPAD and LPAD functions add padding to my name.

Code Listing 7: Query that demonstrates the concatenation operator, ||

```
SQL> select first_name||' '||last_name employee_name
  2    from employee
  3  order by employee_name;

EMPLOYEE_NAME
--------------------------------------------------------------------
Betsy James
Donald Newton
Emily Eckhardt
Frances Newton
Matthew Michaels
Roger Friedli
mark leblanc
michael peterson

8 rows selected.
```

Code Listing 8: Query that demonstrates nested CONCAT calls

```
SQL> select CONCAT(first_name, CONCAT(' ', last_name)) employee_name
  2    from employee
  3  order by employee_name;

EMPLOYEE_NAME
--------------------------------------------------------------------
Betsy James
Donald Newton
Emily Eckhardt
Frances Newton
Matthew Michaels
Roger Friedli
mark leblanc
michael peterson

8 rows selected.
```

Code Listing 9: Query that trims extra spaces

```
SQL> select '''' ||TRIM(TRAILING ' ' FROM 'Ashton   ') || '''' first_name,
'''' || TRIM(LEADING ' ' FROM '  Cinder   ') || '''' last_name
  2    from dual;

FIRST_NA  LAST_NAME
--------  ------------
'Ashton'  'Cinder   '

1 row selected.
```

Code Listing 10: Query that trims extra spaces, including rightmost extra spaces

```
SQL> select '''' || TRIM(TRAILING ' ' FROM 'Ashton   ') || '''' first_name,
'''' || TRIM('   Cinder   ') || '''' last_name
  2    from dual;

FIRST_NA  LAST_NAM
--------  ----------
'Ashton'  'Cinder'

1 row selected.
```

STRINGING STRINGS TOGETHER

Sometimes it makes sense to combine certain strings, such as the FIRST_NAME and LAST_NAME values from the EMPLOYEE table, in the result set display. You can use *concatenation* to accomplish this task—with either the CONCAT function, illustrated in Listing 6, or the (more commonly used) concatenation operator || (two pipe characters), illustrated in Listing 7.

The CONCAT function takes two parameters and concatenates them. You can also nest multiple CONCAT function calls, as shown in Listing 8. The queries in Listings 7 and 8 concatenate literal strings with column data values. (I prefer the concatenation operator, because it has unlimited input parameters and makes the concatenated output more readable.)

GIVING YOUR DATA A TRIM

Sometimes you want to remove unwanted spaces or characters from data when you display it. For example, data inserted into a table column via a form application might include extraneous characters or spaces—preceding and/or following the actual data value—that the form input field doesn't trim.

Listing 9 shows a query that trims extra spaces from string values. The TRIM function in Listing 9 takes two parameters. The first parameter is the character, symbol, or space (delimited by single quotes) to be removed. The second parameter specifies the string literal or column value to be trimmed. The TRIM function supports three keywords: LEADING, TRAILING, and BOTH. The example in Listing 9 uses the TRAILING keyword to right-trim the FIRST_NAME value. The TRIM function applied to the LAST_NAME value specifies the LEADING keyword to left-trim the spaces from that value. And, as you can see, the spaces to the right of the LAST_NAME value remain and are included in the output.

Compare the output in Listing 9 with that in Listing 10, which trims the rightmost extra spaces from the LAST_NAME value. When no keyword is specified, the default behavior for the TRIM function is to trim leading as well as trailing characters. The older RTRIM and LTRIM functions are available for backward compatibility.

Code Listing 11: Query that demonstrates the INSTR character function

```
SQL> select last_name, INSTR(last_name, 'ton') ton_starting_point
  2    from employee
  3  order by last_name;

LAST_NAME                      TON_STARTING_POINT
------------------------------ ------------------
Eckhardt                                        0
Friedli                                         0
James                                           0
Michaels                                        0
Newton                                          4
Newton                                          4
leblanc                                         0
peterson                                        0

8 rows selected.
```

Code Listing 12: Query that demonstrates the SUBSTR character function

```
SQL> select last_name, SUBSTR(last_name, 1, 3)
  2    from employee
  3  order by last_name;

LAST_NAME                      SUB
------------------------------ ---
Eckhardt                       Eck
Friedli                        Fri
James                          Jam
Michaels                       Mic
Newton                         New
Newton                         New
leblanc                        leb
peterson                       pet

8 rows selected.
```

Code Listing 13: Query that demonstrates the INSTR and SUBSTR character functions

```
SQL> select last_name, INSTR(last_name, 'ton') ton_position, SUBSTR(last_name,
INSTR(last_name, 'ton')) substring_ton
  2    from employee
  3  order by last_name;

LAST_NAME                      TON_POSITION SUBSTRING_TON
------------------------------ ------------ --------------------
Eckhardt                                  0 Eckhardt
Friedli                                   0 Friedli
James                                     0 James
Michaels                                  0 Michaels
Newton                                    4 ton
Newton                                    4 ton
leblanc                                   0 leblanc
peterson                                  0 peterson

8 rows selected.
```

SEARCHING FOR STRINGS WITHIN STRINGS

When you need to search column values for similar string pattern values, you can do so with the INSTR character function. INSTR—which stands for *in string*—returns the position of a substring within a string value. Listing 11 demonstrates the INSTR function applied to the LAST_NAME column of the EMPLOYEE table to locate all occurrences of the "ton" substring. As you can see, the INSTR

Code Listing 14: Query that demonstrates the LENGTH function

```
SQL> select first_name, LENGTH(first_name) length
  2    from employee
  3  order by length desc, first_name;

FIRST_NAME                          LENGTH
----------------------------------  ------
Frances                                  7
Matthew                                  7
michael                                  7
Donald                                   6
Betsy                                    5
Emily                                    5
Roger                                    5
mark                                     4

8 rows selected.
```

function takes as input the literal or column value you want to search, followed by the substring pattern to search for. In Listing 11, the INSTR function finds the "ton" pattern in only two column data values—both of them Newton—and returns 4 as their position. Because it did not find the search string in any other values, the output for those values is 0.

Two additional parameters—*starting position* and *occurrence*—are optional. The starting position specifies the character in the string from which to begin your search. The default behavior is for the search to begin at the first character—otherwise known as character position 1. The occurrence parameter lets you specify which occurrence of the substring you'd like to find. For example, the word *Mississippi* includes two occurrences of the "issi" substring. To search for the starting-position location of the second occurrence of this pattern, you must provide the INSTR function with an occurrence parameter of 2:

```
SQL> select INSTR('Mississippi', 'issi',
1, 2)
  2    from dual;

INSTR('MISSISSIPPI','ISSI',1,2)
-------------------------------
                              5

1 row selected.
```

EXTRACTING STRINGS FROM STRINGS

Sometimes you need to extract a portion of a string for your desired output. The SUBSTR (for *substring*) character function can assist you with this task. Listing 12 shows a query that uses the SUBSTR function to extract the first three characters of every LAST_NAME value from the EMPLOYEE table. The SUBSTR function takes two required parameters and one optional input parameter. The first parameter is the literal or column value on which you want the SUBSTR function to operate. The second parameter is the position of the starting character for the substring, and the optional third parameter is the number of characters to be included in the substring. If the third parameter is not specified, the SUBSTR function will return the remainder of the string.

Listing 13 demonstrates the SUBSTR and INSTR functions working together to display the portion of every LAST_NAME value from the EMPLOYEE table that contains the "ton" substring. In this example, the output from the INSTR function provides the value for the input parameter that specifies the position for the SUBSTR function's starting character. In the LAST_NAME values in which the substring "ton" is not found, the entire LAST_NAME value is returned, for two reasons: SUBSTR treats a starting position of 0 the same as a starting position of 1 (that is, as the first position in the string), and because the query omits the optional length parameter, the full remainder of the string is returned.

WHEN SIZE MATTERS

Occasionally you need to determine a string's length—for example, to determine the maximum number of characters a form entry field should permit. Listing 14 shows a query that uses the LENGTH function to display the length of all FIRST_NAME values from the EMPLOYEE table.

The online version of this article at bit.ly/JAQPk3 includes examples of LENGTH and other character functions in WHERE and ORDER BY clauses.

CONCLUSION

This article has shown you how character functions can be used in SELECT statements to manipulate the ways data is displayed. You've seen how to convert data values to uppercase, lowercase, and mixed cases and how to search for strings within strings. You've also seen how to pad and trim data and how to specify a string's total length. By no means does this article provide an exhaustive list of the Oracle character functions. Review the documentation for more details: bit.ly/HZUBC5.

The next installment of SQL 101 will discuss number functions and other miscellaneous 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
online-only article content
bit.ly/JAQPk3

SQL 101, Parts 1–5
bit.ly/JKZzDt

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/l9w52z

DOWNLOAD the sample script for this article
bit.ly/JsDOX2

ORACLE DATABASE

On Characters, Pivots, and Endings

Our technologist substitutes characters, looks at fishy results, and avoids redo.

I want to insert a value into a table as follows:

```
create table test (name varchar2(35));
insert into test values ('&Vivek');
```

But the system asks for a value for the substitution variable. How can I insert an ampersand (&)?

This is probably one of the most frequently asked questions out there, not only on AskTom (asktom.oracle.com) but on any Oracle Database forum.

Before I give the answers, I'll explain the problem fully. By default, SQL*Plus will scan each line of input and look for an & character. Upon finding it, SQL*Plus will scan the characters after the ampersand and use those as a variable name (the variable name in the above example is Vivek). SQL*Plus will then prompt the user for a value for Vivek—like this:

```
SQL> insert into test (name)
values ('&Vivek');
Enter value for vivek: Hello World
old   1: insert into test (name)
values ('&Vivek')
new   1: insert into test (name)
values ('Hello World')

1 row created.
```

Here you can see how SQL*Plus turned &Vivek into Hello World. Now the question is how to stop it from doing that. The easiest method is simply to issue the SQL*Plus `set define off` command:

```
SQL> set define off
SQL> insert into test (name)
values ('&Vivek');

1 row created.

SQL> select * from test;

NAME
-----------------------------------
Hello World
&Vivek
```

That prevents SQL*Plus from scanning the input to try to find the substitution character. Another approach is to use a different substitution character:

```
SQL> set define @

SQL> insert into test (name)
values ( '&Vivek @X' );
Enter value for x: this was x
old   1: insert into test (name)
values ( '&Vivek @X' )
new   1: insert into test (name)
values ( '&Vivek this was x' )

1 row created.
```

In this case, the @ character is doing what the & used to do.

There are other approaches, such as avoiding the & character in your SQL:

```
SQL> insert into test
values (chr(38)||'Vivek  xxx');

1 row created.

SQL> select * from test
where name like '% xxx';

NAME
-----------------------------------
&Vivek  xxx
```

Although that approach works, I'm not a fan of it, because you have to change your SQL statement.

Yet another approach is to use a zero-length substitution variable name, which will make SQL*Plus just leave that & character alone:

```
SQL> insert into test
values ('&'||'Vivek yyy');

1 row created.

SQL> select * from test
where name like '% yyy';

NAME
-----------------------------------
&Vivek yyy
```

This is perhaps better than using chr(38) to avoid the & character in your SQL, but it still is not something I recommend. By the way, even though I've been using SQL*Plus for almost 25 years, I had no idea that the last solution—with '&' ||'Vivek'—would work that way. That was something new I learned from a reader of AskTom!

DYNAMIC PIVOT

I have a table like this:

```
create table fish (
fish_id number,
fish_type varchar2(3),
fish_weight number);

insert into fish values (1,'COD',20);
insert into fish values(1,'HAD',30);
insert into fish values(2,'COD',45);
insert into fish values(2,'HKE',10);
insert into fish values(2,'LIN',55);
```

insert into fish values(3,'CTY',90);
insert into fish values (3,'HAD',60);
insert into fish values (3,'COD',52);

I would like it to be displayed as

COD HAD HKE LIN CTY .....
1 20 30 X X
2 45 X 10 55
3 52 60 X X
....

The columns aren't fixed in number or name, because there can be multiple species.
How can I create this display?

In SQL you need to know the number, name, and datatype of every single column at parse time, so you'll have to use a bit of dynamic SQL. Before I show the dynamic SQL, I will first develop a static SQL statement that works against the existing data. Listing 1 contains a query that works in all releases of Oracle Database. (Note that in Oracle Database 11*g* and later releases, I could have used the built-in PIVOT syntax, but it too requires dynamic SQL.)

Now, to make the SQL in Listing 1 dynamic, I will create a stored procedure that executes a query to determine what the column names are and use that information to dynamically construct the pivot query, as shown in Listing 2.

Note: In Listing 2, the fish_type column is obviously a foreign key to another table in which fish_type is the primary key. The table in which fish_type is the primary key is the lookup table for valid fish types. Replace my SELECT DISTINCT... with a simple SELECT against that table.

What I did in the stored procedure in Listing 2 was to generate a list of distinct fish types and add a column to the query for each one. I used the string

```
q'|, sum(decode(fish_type,'$X$',
fish_weight)) $X$|'
```

as a template for the original static SQL:

```
sum(decode(fish_type,'COD',
fish_weight)) cod,
```

The only thing I had to do with the template was replace X (a string I chose at random to represent COD, HAD, and so on) with the value x.fish_type. Note that I did not just replace X with the value x.fish_type blindly. I used the DBMS_ASSERT package to validate that the data I was concatenating into the SQL statement was "safe"—that it was a simple SQL name and not some SQL that would change the meaning of my SQL statement. In short, that DBMS_ASSERT call is protecting against SQL injection. If SQL injection is a

Code Listing 1: Static SQL for existing data

```
SQL> select fish_id,
  2         sum(decode(fish_type,'COD',fish_weight)) cod,
  3         sum(decode(fish_type,'HAD',fish_weight)) had,
  4         sum(decode(fish_type,'HKE',fish_weight)) hke,
  5         sum(decode(fish_type,'LIN',fish_weight)) lin,
  6         sum(decode(fish_type,'CTY',fish_weight)) cty
  7    from fish
  8   group by fish_id
  9   order by fish_id
 10  /

   FISH_ID        COD        HAD        HKE        LIN        CTY
__________ __________ __________ __________ __________ __________
         1         20         30
         2         45                    10         55
         3         52         60                               90
```

Code Listing 2: Stored procedure that determines column names

```
SQL> create or replace procedure go_fishing( p_cursor in out sys_refcursor )
  2  as
  3      l_query long := 'select fish_id';
  4  begin
  5      for x in (select distinct fish_type from fish order by 1 )
  6      loop
  7          l_query := l_query ||
  8             replace( q'|, sum(decode(fish_type,'$X$',fish_weight)) $X$|',
  9                      '$X$',
 10                      dbms_assert.simple_sql_name(x.fish_type) );
 11      end loop;
 12
 13      l_query := l_query || ' from fish group by fish_id order by fish_id';
 14
 15      open p_cursor for l_query;
 16  end;
 17  /

Procedure created.
```

Code Listing 3: Ref cursor, cursor, and fetch

```
SQL> variable x refcursor
SQL> exec go_fishing( :x )

PL/SQL procedure successfully completed.

SQL> print x

   FISH_ID        COD        CTY        HAD        HKE        LIN
__________ __________ __________ __________ __________ __________
         1         20                    30
         2         45                                10         55
         3         52         90         60
```

topic you are not up to speed on, or if you just want to see some really interesting SQL injection techniques, I encourage you to check out bit.ly/IgU3YQ and bit.ly/K7aAKW.

After I build the query in a string, I use a ref cursor to open a cursor that can be sent back to a client application, as shown in Listing 3.

FINDING THE LAST ROW

I have created a view based on multiple tables in Oracle Database, and I am able to fetch the records within a view via SELECT. My question is: If there are newly added records and I want to fetch only the newly added records inside the view, what is the SQL query for fetching those? I tried using the rowid, but it didn't return the results as expected.

You would have to tell me how you could identify "new rows" if you printed them out on a piece of paper. If you cannot tell which rows are "new" or what the "last row" was, neither can I. Rowids are an address of a row on a block in a file—they are not monotonically increasing values. Even if you only insert into a table, you will find that the rows might not be in "insert order"—sometimes they might be, but other times they won't. In short, you can never count on their being in any sort of order. For example, I used an automatic segment space managed (ASSM) tablespace with an 8 K block size and inserted some data, as shown in Listing 4.

Now, arguably row "4" is the "last" and newest row in the T table, but if I start looking at the rows—and what database blocks they are on—I'll see a different story, as shown in Listing 5.

If I were to try to use ROWID to find the "last" row or the newest row, I'd be very much let down at this point. What I did was put a small row 1 on the first block and place a larger row 2 on the same first block, but then when I tried to insert row 3, it was too big to fit on the first block with those rows, so it went to the second block in the table. However, when row 4 came along—it was small again—there was room for it on the first block.

Rows go on a block that has free space on it—enough free space to hold those rows. They do not go to the "last block" in the table; if they did, tables would only grow. We would never be able to reuse space left behind after a delete.

If you need to find the "newest rows" in a table, you'll have to have some bit of data—a time stamp, for example—associated with each row that would enable you to identify a new row. Some people might mention ORA_ROWSCN as a solution, but the overhead of using it to find "new rows" would be overwhelmingly huge—you'd have to inspect every single row every single time.

REDO AND GLOBAL TEMPORARY TABLES

Do temporary tables generate redo for standard data manipulation language] operations?

I think that because every DML operation generates undo, every DML operation on global temporary tables will generate redo.

In short, there is no redo generated on temporary table blocks. However, any undo generated for those blocks will have redo generated. Thus, many operations against global temporary tables will generate redo as

Code Listing 4: Creating table for "last row" test

```
SQL> create table t ( x int, y varchar2(4000), z varchar2(4000) );
Table created.

SQL> insert into t values ( 1, rpad('*',1,'*'), rpad('*',1,'*') );
1 row created.

SQL> connect /
Connected.

SQL> insert into t values ( 2, rpad('*',3000,'*'), rpad('*',3000,'*') );
1 row created.

SQL> connect /
Connected.

SQL> insert into t values ( 3, rpad('*',3000,'*'), rpad('*',3000,'*') );
1 row created.

SQL> connect /
Connected.

SQL> insert into t values ( 4, rpad('*',1,'*'), rpad('*',1,'*') );
1 row created.
```

Code Listing 5: Query to see order of rowids

```
SQL> connect /
Connected.

SQL> select dbms_rowid.rowid_block_number(rowid), x from t;

DBMS_ROWID.ROWID_BLOCK_NUMBER(ROWID)          X
_____________________________________  __________
                               18948           1
                               18948           2
                               18948           4
                               18949           3

SQL> select rowid, x from t order by rowid desc;

ROWID                    X
__________________  __________
AAAaxGAAEAAAEoFAAA           3
AAAaxGAAEAAAEoEAAC           4
AAAaxGAAEAAAEoEAAB           2
AAAaxGAAEAAAEoEAAA           1
```

a side effect of generating undo.

The example in Listing 6 demonstrates operations that generate and do not generate redo.

In Listing 6 there was 0 redo size with a direct path INSERT (but you have to commit the INSERT before you can read it) because the direct path INSERT bypassed undo generation. The conventional path load generated 412,112 bytes of redo, but that was to protect the UNDO information only—not the data loaded. That is apparent if you use a conventional path INSERT into a "normal" table:

```
SQL> create table t
  2  as
  3  select * from all_objects where 1=0;
Table created.

SQL> set autotrace traceonly statistics;
SQL> insert into t select * from
all_objects;

72862 rows created.

Statistics
----------------------------------

…
    8546004  redo size
…
      72259  rows processed
```

A global temporary table typically significantly reduces the amount of redo generated, but it will not typically eliminate it.

TRICKY UNIQUE CONSTRAINT

I have a data rule for a table that says a certain pair of columns must be unique if some other field is a certain value. Specifically, if a column t_resource_type value is in the set of values 100000, 1000001, and 1000002, the t_resource_address1 and the t_resource_hst_id values must be unique. How can I construct this?

This is easy to accomplish via a function-based index or virtual columns (in Oracle Database 11*g* and later only). I'll demonstrate both. First, I'll create your table:

```
SQL> create table test_data
  2  (
  3  t_resource_type number(8)
  4  ,t_resource_address1 varchar2(50)
  5  ,t_resource_hst_id number(11)
  6  );
Table created.
```

And then I'll add a unique index—more specifically, a unique function-based index:

Code Listing 6: No redo created in direct path insert

```
SQL> create global temporary table gtt
  2  on commit preserve rows
  3  as
  4  select * from all_objects where 1=0;
Table created.

SQL> set autotrace on statistics
SQL> insert into gtt select * from all_objects;
72259 rows created.

Statistics
----------------------------------------------------------------------------

…
     412112  redo size
…
      72259  rows processed

SQL> insert /*+ append */ into gtt select * from all_objects;
72259 rows created.

Statistics
----------------------------------------------------------------------------

…
          0  redo size
…
      72259  rows processed

SQL> set autotrace off
```

Code Listing 7: A real constraint with virtual columns

```
SQL> alter table test_data
  2  add (
  3  t_resource_address1_unq varchar2(50)
  4  generated always as
  5  (case when t_resource_type in (100000, 1000001, 1000002) then
      t_resource_address1 end)
  6  )
  7  /

Table altered.

SQL> alter table test_data
  2  add (
  3  t_resource_hst_id_unq number(11)
  4  generated always as
  5  (case when t_resource_type in (100000, 1000001, 1000002) then
      t_resource_hst_id end)
  6  )
  7  /

Table altered.

SQL> alter table test_data
  2  add constraint address_hst_id_unique
  3  unique (t_resource_address1_unq,t_resource_hst_id_unq);

Table altered.
```

```
SQL> create unique index t_idx
     on test_data (
  2  case when t_resource_type in
     (100000, 1000001, 1000002)
     then t_resource_address1
     end,
  3  case when t_resource_type in
     (100000, 1000001, 1000002) then
     t_resource_hst_id
     end
  4  );

Index created.
```

The two case statements will return either

Null, Null – when the t_resource_type column value is not in the set of values 100000, 1000001, and 1000002

or

t_resource_address1, t_resource_hst_id – when the t_resource_type column value is in that set of values.

Because the null, null values are always considered unique, they will not appear in the index—the index will index rows in the table only such that the resource type is in the specified set of values. You'll have a unique index on only some of the rows in the table.

With Oracle Database 11*g* and later, I can use a real constraint by using virtual columns, as shown in Listing 7.

This approach adds two columns to the table, and these columns appear to have values only when the resource type is in the specified set of values—otherwise they are null. Because they are "real" columns in the table, you can apply a traditional constraint to them. They consume no storage, because they are virtual columns, but you can do pretty much anything to them that you can do to a "regular" column. In the end, the two approaches are nearly equivalent, in that they both create a function-based index. ◀

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

DOWNLOAD Oracle Database 11*g* Release 2
bit.ly/epBiUG

LEARN more about SQL injection
bit.ly/IgU3YQ
bit.ly/K7aAKW

Testing Makes Perfect

Successful application upgrade testing combines leadership, preparation, and the best people.

People say "practice makes perfect," but in the enterprise application and enterprise resource planning (ERP) world, a more appropriate statement might be "testing makes perfect."

A recent *OAUG ResearchLine* report[1] about enterprise application/ERP suite upgrade strategies indicates that proper testing is of great importance. The report, based on the results of a research study conducted with members of the Oracle Applications Users Group (OAUG), measured how respondents perceive the business opportunity and the process for moving to the latest release of an enterprise application suite. The testing process was identified by a majority of respondents (63 percent) as a key risk area, leading the list in terms of ERP upgrade challenges.

While there are a number of diverse software development methodologies, testing is one of the most critical. Here are three key ways to make your organization's application upgrade testing experience as smooth and as "perfect" as possible.

LEADERSHIP MUST EDUCATE

The first strategy for ensuring a smooth-to-perfect testing experience starts with leadership. The leadership team should educate users on the overall testing process, the various types of testing, when each type occurs, and how users are involved. Users are in the best position to determine if a system is going to meet their requirements when it is turned over to them, and it is the leadership team's responsibility to make sure users know what is expected of them during the testing process.

For example, users involved in testing should not just run a process or report to see that it has been executed; they should instead review the results and make sure the results are exactly what was expected.

Some people naturally make better testers than others.

In addition to using correct testing processes, users should be educated and prepared to do *negative testing*. Negative testing helps to ensure that the system prevents or traps errors or properly unwinds transactions that are being reversed or canceled.

With their big-picture vantage point, project leaders are in the best position to engage users, establish expectations and objectives, and direct testing efforts.

BE PREPARED

In addition to leadership, preparation is key to a successful testing experience. Proper preparation includes training and educating the users on the appropriate way to go about testing and how to log, report, and track errors. Users should be instructed on how to communicate perceived issues by including screen shots or other documentation with the issue being logged. Those tasked with issue resolution must understand *why* the test case failed so they can work on a proper fix. The testing team should understand who is responsible for tracking the issue and ensuring that the issue is fixed.

Additionally, testing teams should be provided with proper test cases, test scripts, and test data. A systematic effort that includes a repeatable approach allows testing teams to approach the job in the same way each time patching and upgrades occur.

GETTING THE BEST

After focused leadership is in place and necessary preparation is done, a third strategy for successful testing is to make sure the best people are involved in the hands-on testing process. Some people naturally make better testers than others. These are the people who can anticipate how the system may break or how someone may try to use it, and they have an innate sense of the overall processes and functions of the software. These team members will understand if a system as built or configured is actually meeting business requirements. Sometimes these people may be wearing multiple hats within the organization, or they may have gained their special insight from working in a wide variety of roles within a business. People with an audit background frequently make great additions to a testing team.

Unfortunately, due to their critical roles in an organization, the best testers do not always have the time to dedicate to the testing effort. In this case, they should be tapped to help design test cases and test scripts that capture their expertise so that others may apply this knowledge.

More organizations should view testing as an opportunity for feedback and refinement as well as user education. Too many teams view testing as a necessary evil or something that must be overcome prior to go-live, and they miss out on the benefits that a proper testing effort can provide. ◀

Mark C. Clark is president of OAUG and has been an active member since 1992. As a senior partner at O2Works, Clark is currently engaged with customers worldwide in their Oracle E-Business Suite Release 12 and 12.1 efforts.

NEXT STEPS

EXPLORE the survey
bit.ly/IBZr3o

JOIN OAUG
oaug.org

1 "ERP Upgrades: What's Your Philosophy? 2012 OAUG Survey on Enterprise Application/ERP Suite Upgrade Strategies," Unisphere Research, February 2012

The Best Defense

Today's database security tools protect your data at multiple levels.

Oracle Magazine spoke with Martin Kuppinger, founder and principal analyst at KuppingerCole, about database security as the cornerstone of an end-to-end security architecture.

Oracle Magazine: Why is a multilevel approach to security important?

Kuppinger: When you look at the well-publicized breaches of IT security, in many cases the attacker is an internal person who had access to a database. A layered security approach protects each part of your technology stack, from the network to the application, including the database. While identity management technology authenticates people at the application level, if the data is still readable and in plain text, then there are plenty of ways that a malicious intruder can access it.

Oracle Magazine: What's the difference between securing data in the cloud and securing data on premises?

Kuppinger: Your security approach should differ depending on the type of clouds you're using. Running a private cloud in a well-defined data center at a specific location is different from simply renting a virtual machine in a public cloud. One of the issues is that you often don't know where your data resides. You rely on service-level agreements for security, which comes down to trusting the vendor. In those instances, it's important to protect your data—generally using encryption.

Oracle Magazine: What responsibilities do enterprises have to secure personally identifiable information [PII], and what are the primary risks that they must address?

Kuppinger: There are two types of risk here: monetary penalties and breach notification. PII regulations differ from country to country. International organizations must meet the highest levels of security to ensure that they are fully compliant. As to the risks, if you lose data then you face breach notification penalties and you might end up making the headlines the next day. Three or four years ago, you could lose a lot of data and it would be noted in some computer magazines. Now you might find yourself on the front page of the business news, which can have huge ramifications on the enterprise, on shareholder value, and on your reputation with customers.

A layered security approach protects each part of your technology stack.

Oracle Magazine: How do database security technologies help enterprises mitigate these risks?

Kuppinger: You need a multifaceted database security portfolio to fulfill regulatory compliance criteria. Auditability and traceability are very important, as are labeling data and segmenting it into different domains. Encryption and strong authentication are also essential.

Organizations must look at the requirements for their industry, region, and country. They must identify risks and select a variety of technologies to make sure that they have covered everything that pertains to them.

Oracle Magazine: DBAs, system administrators, and other technical personnel need access to database resources. How do organizations secure database information from their own administrators?

Kuppinger: Limit the actions of privileged users. For starters, you can segment data into domains and limit administrative access to financial data and PII. HR data is often confidential as well. Privileged users are important, but they are also a big risk. Don't give them access that they don't need, and encrypt sensitive data so that they can work on the database without seeing things that they don't need to see.

Oracle Magazine: How does a database firewall differ from a network firewall?

Kuppinger: Both of them are called firewalls, but they do different things. A network firewall guards the perimeter of the network, while a database firewall works from within to detect SQL injections and rogue transactions that shouldn't be allowed. Place a database firewall in front of the database within your data center to analyze the SQL statements and prevent the execution of malicious programs or loss of data.

Oracle Magazine: What are the pros and cons of database-level encryption—such as transparent data encryption—and full-disk encryption?

Kuppinger: Important data should be encrypted, partly to protect it from privileged users who have broad access to information. Transparent data encryption is applied to the specific needs of a database environment, whereas full-disk encryption protects data at rest on the disk but in no other situation. Of course, even transparent data encryption doesn't protect data while somebody is using it. But it does protect some part of the communication when the data is in motion. ◀

David Baum (david@dbaumcomm.com) is a freelance business writer based in Santa Barbara, California.

Founded in 2004, KuppingerCole (kuppingercole.com) is a leading analyst company for identity-focused information security, in classical and cloud environments.

NEXT STEPS

READ about Oracle Database security
oracle.com/us/products/database/security

Unbeatable power protection now beats energy costs, too.

Only APC Smart-UPS delivers unsurpassed power protection and real energy savings.

Today's more sophisticated server and networking technologies require higher availability. That means you need more sophisticated power protection to keep your business up and running at all times. But that's not all. In today's economy, your UPS must safeguard both your uptime *and* your bottom line. Only APC™ by Schneider Electric™ helps you meet both of these pressing needs. Specifically, the APC Smart-UPS™ family now boasts models with advanced management capabilities, including the ability to manage your energy in server rooms, retail stores, branch offices, network closets, and other distributed environments.

Intelligent UPS management software.

PowerChute™ Business Edition, which comes standard with Smart-UPS 5 kVA and below, enables energy usage and energy cost reporting so you can save energy and money by tracking energy usage and costs over time; CO_2 emissions monitoring to reduce environmental impact through increased understanding; and risk assessment reporting so you can identify and proactively manage threats to availability (e.g., aging batteries).

Best-in-class UPS.

Our intelligent, interactive, energy-saving APC Smart-UPS represents the combination of more than 25 years of Legendary Reliability™ with the latest in UPS technology including an easy-to-read, interactive, alphanumeric LCD display to keep you informed of important status, configuration, and diagnostic information, a unique battery life expectancy predictor, and energy-saving design features, like a patent-pending "green" mode.

Now, more than ever, every cost matters and performance is critical. That's why you should insist on the more intelligent, more intuitive APC Smart-UPS.

Why Smart-UPS is a **smarter solution**

Intuitive alphanumeric display

Get detailed UPS and power quality information at a glance – including status, about, and diagnostic log menus in up to five languages.

Configurable interface

Set up and control key UPS parameters and functions using the intuitive navigation keys. On rack/tower convertible models, the display rotates 90 degrees for easy viewing.

Energy savings

A patent-pending "green" mode achieves online efficiencies greater than 97 percent, reducing heat loss and utility costs.

Download our FREE Server Room Efficiency Kit and enter to win one of five Smart-UPS units (SMX1000)!

Visit www.apc.com/promo Key Code p675v
Call 888-289-APCC x6410

APC™
by Schneider Electric