

April 10, 2008

Inside This Issue

- 1 Come Join Us at Collaborate 08!
- 1 TruTek TruTalk is Back!
- 1 Hot Presentations at Collaborate 08
- 3 More Hot Presentations and Papers
- 4 TruTek Training
- 7 Books
- 8 Tips and Tricks
- 11 Confio
- 12 AppsPerf
- 12 Collaborate Presentation Schedule

TruTek TruTalk is Back!

TruTek is excited to be able to bring back our TruTalk Newsletter. If you'd like to submit articles or suggest topics, send us a note at editor@trutek.com. If you have questions about the Oracle RDBMS, the E-Business Suite of Applications, or PeopleSoft, send us a note.

Come Join Us at Collaborate 08!

This year TruTek will be at **Booth 1515** in the vendor hall. Stop by and talk to [Mike Swing](#) about your Oracle RDBMS and E-Business Suite training and consulting needs. We specialize in Oracle Database consulting, Oracle Applications consulting, and Oracle Database and E-Business Suite training.

Mike is co-presenting [Migrating 11i to RAC Using DataPump](#) with Fred Linnenbrink, Director of Database Systems and Architecture at Metro St Louis on Tuesday, April 15th. If you've wondered about the most effective way to migrate Release 11i to a 64-bit RAC cluster on Linux from a single-node 32-bit 10gR2 database, this presentation is very informative. The paper includes an overview of RAC, ASM and Clusterware technologies, as well as the use of Oracle's DataPump tool as the fastest way to migrate data between operating systems.



Also, don't miss our drawing for an **iPod Touch MP3 Player!** Stop by Booth 1515 and sign up for your chance to win!



Hot Presentations at Collaborate 08

We're honored to collaborate with some of the top experts in the Oracle World. Don't miss these terrific presentations, as they're by folks we really respect, including clients as well as fellow consultants:



[Fueling the Workflow Engine – Writing Custom PL/SQL to be Called from your Workflow Processes](#)

Dan Stober, Intermountain Healthcare

When it comes time to customize your Workflow Processes, you'll need to know how to write code that interacts with Oracle Workflow. Among other topics, this session will feature an explanation of the Workflow data model, standard call signature, status_codes, and result_codes. We'll explore some uses of standard APIs as we examine how to write code to implement function activities, to populate the misunderstood document type attribute, and to create a selector function.



[Customization Survival Guide: How to Use E-Business Utilities to Migrate Your Custom Code](#) - Brad Simmons and Donna Campbell of Los Alamos National Laboratory

Have you ever had to migrate a form or a report and found yourself hand-typing the setups from your test environment to your production environment? Brad and Donna's paper thoroughly examines Oracle's FNDLOAD, XMLImporter, and XDOLoader utilities. They've included examples of commands and sample Unix scripts, and they've highlighted the potential 'gotchas' that you need to keep in mind.

[Scaling to Infinity: Partitioning Data Warehouses in Oracle](#) – Tim Gorman, Evergreen Database Technologies, Inc. - This seminar is designed to conclusively justify one of the most important technical decisions that can be made during the physical build of a data warehouse based on the Oracle RDBMS, regardless of size: whether or not to use partitioning. Using partitioning permits the full power of numerous advanced Oracle features to be brought to bear on the data warehouse, regardless of how large it grows or how busy it becomes in the future, which Tim characterizes as the "virtuous cycle". Not starting with partitioning, regardless of how humble the beginnings, can lead to a deadly embrace of unfortunate circumstances that can doom the project to eventual -- almost inevitable -- failure, which Tim characterizes as the "death spiral". This seminar is a full day Sunday University seminar that has a separate fee (\$459 in advance, \$549 onsite) from the conference fee.

[MAID for ILM Archival Storage for Data Warehouses - Information Lifecycle Management \(ILM\)](#) – Tim Gorman, Evergreen Database Technologies, Inc. – Information Lifecycle Management (ILM) is a set of techniques supported by the "ILM Assistant" plug-in to SQL Developer. ILM is the technique of utilizing the appropriate storage technology (i.e., fast and expensive, slower and less-expensive, etc.) according to the ways that the data is accessed. These techniques are particularly appropriate for data warehouses, but can be used for any application that generates enormous volumes of data with long or undefined retention policies.

The real art of ILM is enabling the long-term "tier" of storage, usually termed "archival", which stores data the longest. In the past, offline or nearline solutions were common, but COPAN Systems offers a new technology called "MAID" (massive array of idle disk) that offers an extremely simple and cost-effective storage for massive multi-petabyte volumes of data to be retained for many years. MAID is intended to completely replace tape media while maintaining archived data fully online and available.



[PeopleSoft for the DBA: PeopleSoft Performance Optimisation](#) - David Kurtz, Go-Faster Consultancy Ltd. – This seminar is a subset of David's three day "PeopleSoft for the Oracle DBA" course. It examines how to obtain performance metrics from the PeopleSoft software, including the new performance Monitor, and then how to use that information to identify and resolve performance bottlenecks. It goes on to demonstrate how to identify source code, and how to use database tuning techniques in the context of a PeopleSoft system. This seminar is a full day Sunday University seminar that has a separate fee (\$459 in advance, \$549 onsite) from the conference fee.

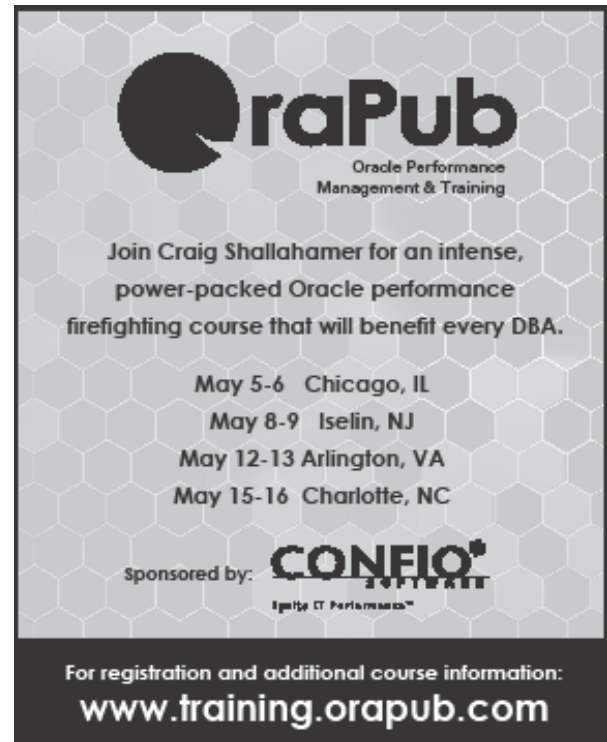
David is also presenting on three other topics at the conference: **[PeopleTools 8.48/8.49 – New Database Features](#)**; **[An Introduction to SQL*Trace, TKPROF and Execution Plans for Developers](#)**, and **[PeopleSoft: A Properly Instrumented Application](#)**.

More Hot Presentations and Papers

Getting Started Forecasting Oracle Performance – *Craig Shallahamer, OraPub* - Forecasting, capacity planning, and resource management try to answer questions like “When will my server run out of CPU power?” Craig’s presentation teaches DBAs that the math to answer this type of question can be simple, the required data has already been collected, DBAs can quickly do what Craig will demonstrate, and it is shockingly practical! DBAs of all levels will learn how to get a practical start with forecasting Oracle performance using freely available tools.

Getting Started Optimizing Oracle Performance – *Craig Shallahamer, OraPub* - Even the documentation states, “The Oracle server is a sophisticated and highly tunable software product.” Then doesn’t it follow that learning how to diagnose and optimize performance should be a complex, daunting, and overwhelming challenge? Absolutely not! Craig has spent years developing a number of methods and techniques that have allowed literally thousands to learn how to make complex Oracle server tuning simple and straightforward. This paper introduces new Oracle DBAs to the world of practical performance diagnosis and resolution. Craig begins by introducing the concepts and key points, surrounding them with proven methods, and then demonstrating how to successfully diagnose and set a path towards optimizing performance. Publicly available tools and methods are used throughout the paper.

All About Oracle’s In-Memory Undo – *Craig Shallahamer, OraPub* – This paper is so new it didn’t make the Collaborate deadline. The Oracle database is extremely complex software. Part of the solution to performance optimization is Oracle Corporation itself optimizing the core database kernel code. One such optimization is the introduction of Oracle’s patented *In-memory undo* (IMU). Essentially, instead of maintaining undo in Oracle segments, the undo is managed, as much as possible, in memory using structures optimized for in-memory operations. But as you will read, how Oracle does this is fascinating and foreshadows even greater things to come. But with any piece of code, there is always the possibility of a bottleneck, so this paper covers how to detect IMU performance issues and multiple ways to potentially solve the problem (in addition to just turning it off).



The poster features the OraPub logo at the top, which includes a stylized 'O' icon and the text 'OraPub Oracle Performance Management & Training'. Below the logo, the text reads: 'Join Craig Shallahamer for an intense, power-packed Oracle performance firefighting course that will benefit every DBA.' The event dates and locations are listed: 'May 5-6 Chicago, IL', 'May 8-9 Iselin, NJ', 'May 12-13 Arlington, VA', and 'May 15-16 Charlotte, NC'. At the bottom, it says 'Sponsored by: CONFIO! Equalize IT Performance™' and 'For registration and additional course information: www.training.orapub.com'.

Sign up now for
TruTek Training!

Sign up now for
TruTek Training!

TruTek offers premium quality training at a reasonable price. Our classes include E-Business Suite technical and functional topics. We plan to offer PeopleSoft classes soon. Look for additional classes at <http://www.trutek.com>.

Date	Class	Instructor	Location
April 21-25	<u>OBI EE</u> – Students will learn to build, configure and maintain Interactive Dashboards. Students will also learn how to build queries in Oracle BI Answers, customize the Oracle BI user interface, use Oracle BI Content in Microsoft Office documents and use Oracle BI Delivers to configure, schedule and deliver personalized alerts.		SLC, UT
April 22-24	<u>Workflow Builder, Business Events and Administration</u> – This class covers Workflow Builder, Business Events and Administration. Labs are intended to help attendees learn to create and modify a workflow process and to administer workflow.		SLC, UT
May 5-9	<u>Release 11i Advanced Pricing</u> – Learn about the following Advanced Pricing features and concepts: generic setup steps, pricing security, qualifiers, modifiers, price lists, formulas, attribute management and troubleshooting. Learn how to set up and use price lists with pricing lines to define item and/or item category prices. Students will also learn how to use modifiers to set up price adjustments, benefits, and freight and special charges.		SLC, UT
May 13-16	<u>Release 12 Upgrade</u> – A hands-on workshop oriented review of the entire Release 12 upgrade process, including: RDBMS 10gR2 upgrade, Release 12 Architecture, Concepts and Administration, pre-upgrade steps, Release 12.0.4 upgrade and post-upgrade steps.	Mike Swing	SLC, UT
May 20-22	<u>BI Publisher</u> – Build your expertise with Oracle BI Publisher 10g by gaining a foundation in its basic concepts, architecture, and underlying standards. By the end of the course, you will be able to create Oracle BI Publisher templates in RTF and PDF formats. You will also learn how to deploy those templates to integrate with the the E-Business Suite of Applications. Hands-on exercises reinforce concepts taught in class.		SLC, UT
June 3-5	<u>Release 11i System Administration</u> - Students will learn about function and data security, using Oracle Applications Manager to manage the E-Business Suite, creating and managing concurrent managers, managing	Mike Swing	SLC, UT

Date	Class	Instructor	Location
	concurrent programs, request sets, profile options, and managing customizations.		
June 9	<u>Troubleshooting and Tuning</u> – Renowned database expert Jonathan Lewis describes the different approaches to troubleshooting and tuning Oracle databases, including details about understanding and interpreting the content of Statspack and AWR reports.	Jonathan Lewis	SLC, UT
June 10	<u>Writing Optimal SQL</u> – Don't miss this opportunity to learn from one of the best instructors in the industry. In this one day seminar, renowned database expert Jonathan Lewis covers indexing, joins, hints, and coding strategies and how to pick an efficient execution path..	Jonathan Lewis	SLC, UT
June 11	<u>Troubleshooting and Tuning</u> – Renowned database expert Jonathan Lewis describes the different approaches to troubleshooting and tuning Oracle databases, including details about understanding and interpreting the content of Statspack and AWR reports.	Jonathan Lewis	Denver, CO
June 12	<u>Writing Optimal SQL</u> – Don't miss this opportunity to learn from one of the best instructors in the industry. In this one day seminar, renowned database expert Jonathan Lewis covers indexing, joins, hints, and coding strategies and how to pick an efficient execution path.	Jonathan Lewis	Denver, CO
June 13	<u>Explain Plan</u> – In Jonathan Lewis's third one day class, students will learn how to find execution paths, how to interpret execution plans, they types of omissions and errors that can lead you astray with explain plans, and advanced features.	Jonathan Lewis	Denver, CO
June 17- June 18	<u>Scaling to Infinity: Partitioning Data Warehouses in Oracle</u> - This two-day seminar is designed to conclusively justify one of most important technical decisions that must be made during the physical build of an Oracle data warehouse, regardless of size: whether or not to partition. Partitioning brings to bear the full power of numerous advanced Oracle features to the data warehouse, regardless of how large it grows or how busy it becomes. Tim characterizes this as the "virtuous cycle". Not starting with partitioning can lead to a "death spiral" of unfortunate circumstances that can doom the project to eventual -- almost inevitable -- failure. This seminar provides all levels of DBAs, technical architects, and technical managers with a clear and detailed understanding of how partitioning is used, and how it unlocks the features and power of Oracle in a	Tim Gorman ♠ (Tim is an <u>Oracle ACE</u>)	Englewood, CO

Date	Class	Instructor	Location
	decision-support application environment.		
June 23- June 26	<u>11i Financial Business Process Overview</u> – This class discusses the integration and shared entities of the Financial Suite Applications. This includes General Ledger, Purchasing, Payables, Fixed Assets, Receivables and Cash Management. The class will also discuss how Inventory, Order Management, Projects, iExpense, iSupplier, iProcurement, iReceivables, TCA and Human Resources tie in with these applications. This class will cover key setup areas from each of the applications. Students will have the opportunity to perform key transactions in each application and follow them through the process cycle of the Financial Applications.		Denver, CO
June 23- June 27	<u>11i/R12 Apps DBA Concepts and Administration</u> – This class is a very hands-on, unique class that approaches Applications DBA challenges from the Application, Application Server, Database and Operating System perspective. Students will learn how to install and configure Oracle E-Business Suite Release 11i and Release 12, and how to consistently and quickly understand complex Oracle Applications problems.	Mike Swing	Denver, CO
June 30- July 3	<u>Release 11i System Administration</u> – Students will learn about function and data security, using Oracle Applications Manager to manage the E-Business Suite, creating and managing concurrent managers, managing concurrent programs, request sets, profile options, and managing customization.		Denver, CO
July 7-8	<u>Discoverer for End Users</u> – This course introduces students to the querying and analytical capabilities of Oracle Discoverer. Students will learn to create, modify, run, and refine ad hoc queries as well as to view, chart, and analyze multi-dimensional data.	Dan Hotka	SLC, UT
July 9-10	<u>Discoverer for Administrators</u> – This class highlights the strengths and weaknesses of standard BIS EUL and how to modify it to get what you need in Discoverer. Topics include installation prerequisites, setup and configuration, implementation steps for BIS, and Discoverer EUL in an 11.5.10 environment.	Dan Hotka	SLC, UT
July 8-11	<u>Procure-to-Pay</u> – This course provides an understanding of the Procure-to-Pay life cycle and its implementation in the Oracle Purchasing and Payables applications. It also covers information about the suppliers, sites, and quotations used in the Procure-to-Pay process.		SLC, UT


Our Instructors Wrote the Book(s) on...

Check our Training Schedule. Several of our favorite authors/instructors are offering training classes this quarter!




In [Oracle Database 11g New Features](#), Robert Freeman clearly explains Oracle Database 11g new concepts and includes real life examples to illustrate the value of those features. Every new and updated feature is covered and presented with screenshots, code samples, tables, and charts. Find out how to take full advantage of all the new and improved capabilities of Oracle Database 11g, including: new Oracle Automatic Storage Management features; improved Automated Memory Management capabilities; new database design features; new Database Advisors for data recovery and repair, partitioning, and streams performance; new tuning features; enhancements to Oracle RMAN; improved auditing and security capabilities; advanced BI, data warehousing, and partitioning features and new SQL and PL/SQL enhancements.


If you're thinking about Oracle Database 11g, then this is the must-have reference.



Have you ever wondered just how much activity your database server can handle before performance degrades? Or if it makes sense to upgrade your server? If you are responsible for ensuring there is no unplanned downtime and that your database servers can handle future loads, then you need to know how to forecast performance. Craig Shallahamer's [Forecasting Oracle Performance](#) focuses on teaching you, step-by-step, how to forecast the future risk of over utilized database server resources and then developing ways to mitigate the risk...before the fires begin!



In [Oracle Insights: Tales of the Oak Table](#), Tim Gorman and Jonathan Lewis are two of eleven contributing authors (along with *Mogens Nørgaard, Dave Ensor, Kyle Hailey, Anjo Kolk, Connor McDonald, Cary Millsap, James Morle, David Ruthven, and Gaja Krishna Vaidyanatha*) who provide insights about their successful (and less-than-successful!) experiences in the world of database technology. This anthology is educational, entertaining, and alarming. The members of the [Oak Table Network](#) dispense wisdom and humor reflecting years of experience in Oracle-based technology. Courtesy of Apress, Tim's chapter from the book is posted on [Simple-Talk.com](#).



In [Cost-Based Oracle Fundamentals](#), the first book in a series of three, Jonathan Lewis - one of the foremost authorities in this field - describes the most commonly used parts of the Cost Based Optimization model, what the optimizer does with your statistics, and why things can go wrong. With this information, you'll be in a position to fix entire problem areas, not just single SQL statements, by adjusting the model or creating more truthful statistics.

Tips and Tricks: Approvals Management Engine Example

1. Log into the Applications as the operations user and choose the Approvals Management Business Analyst responsibility.
2. Choose the “Business Analyst Dashboard” function.
3. In the “Approval Process Setup” box on the right side of the screen, enter “Oracle Self Service Human Resources” and click on the flashlight icon to select “Oracle Self Service Human Resources” off the list of transaction types.
4. Click on the “Rules” link in the “Approval Process Setup” box.
5. In the “Rules” screen, click on the “Create” button.
6. In the “Create New Rule: Enter rule details” screen, enter “XX Custom rule with two approvers” above the requestor in the “name” field. Leave the rule type set to “List Creation” and click on the “Next” button.
7. In the “Create New Rule: Add conditions” screen, click on the “Add Condition” button.
8. In the “Add Conditions” screen, click on the “Create” button.
9. In the “Create New condition” screen, leave the “Condition type” set to “Ordinary” and choose “WORKFLOW_PROCESS_NAME” from the list of values in the “attribute” field. You should see a string value field.
10. In the string value box, enter “XXHR_PERSONAL_INFO_JSP_PRC” (the name of your personal information workflow process) as the value for your condition. Click on the “Apply” button. You should be returned to the Oracle Approvals Management main screen, and should see a message that the condition has been created successfully.
11. In the “Add Condition/Confirmation” screen, select your condition using the checkbox and then click on the “Continue” button.
12. In the “Create New Rule: Add Conditions” screen, check that you see your condition and click on the “Next” button.
13. In the “Create New Rule: Add Actions” screen, choose “Require approvals up to the first two superiors, at most” from the list of values in the “Action” field. Click on the “Next” button.
14. In the “Create New Rule: Review” screen, review your rule and click on the “Finish” button.

To Test Your Rule:

1. Click on the “Home” link in the upper right corner of the screen.
2. Choose the Business Analyst dashboard function.
3. In the Approvals Process Setup box at the right side of the screen, enter “Oracle Self Service Human Resources” and click on the flashlight icon to pick “Oracle Self Service Human Resources” off the list of transaction types.
4. In the “Approvals Process Setup” box at the right side of the screen, click on the “Test Workbench” link.

5. In the "Test workbench" screen, click on the "Create" button to create a test case.
6. In the "Create Test Case" screen, enter "XX TESTCASE 1" in the name field.
7. In the "Create Test Case" screen, choose "Stock, Ms. Pat" from the list of values in the "TRANSACTION_REQUESTOR_PERSON_ID" field.
8. In the "Create Test Case" screen, enter "XXHR_PERSONAL_INFO_JSP_PRC" in the "WORKFLOW_PROCESS_NAME" field.
9. Click on the "Run Test Case (2)" button.
10. In the "Run Adhoc test case" screen, you should see "Casey Brown (Pat Stock's supervisor)" and "William Tucker (Casey Brown's supervisor)" in the "Final Processed Approver" list.

Tips and Tricks: How to Set Up HR Workflows

Log in as the Application Developer responsibility. Choose "Function" and then set up a new function with a type of "SSWA jsp function" (you'll see that on the "Properties" tab).

The linking of the workflow to the function (menu item) happens on the "Form" tab. Inside the field for "Parameters" you have to include a string of parameter name/value pairs separated by ampersands. This builds a standard URL call with HTTP parameters. Those parameters should match up to item attributes that are defined in the workflow item. It might sound difficult, but the truth is that it sounds more confusing than it really is. When you are done, your string might look something like this:

```
pProcessName=HR_PERSONAL_INFO_JSP_PRC&pItemType=HRSSA&pCalledFrom=IHR_BASIC_SS&pPersonID=&pFromMenu=Y
```

Then, there's another piece on a different tab within the same form, "Web HTML". In this field, put "OA.jsp", which is the name of the the JSP that drives all of this, followed by a question mark, which is the character that precedes paramaters in a URL, and then a couple more parameter pairs:

```
OA.jsp?akRegionCode=HR_CREATE_PROCESS_TOP_SS&akRegionApplicationId=800&OAFunc=HR_ADDRESS_SS
```

Chapter 3 and Chapter 7 in *Implementing Oracle Self-Service Human Resources 4.0, Release 11i, February 2002* cover this set up in more detail.

Tips and Tricks: Measuring Disk Activity

When was the last time you took a look at your disk activity? If you are like most DBAs, you work in a shop that has storage administrators, which means you probably have not had to take a look. Recently, we were working with a client that wanted us to review their backup strategy and make some recommendations. We started to dig into the environment and found that the internal folks were managing the environment well, but could use a little help with some broader perspectives. For example, what would they do if their SAN failed? Of course, the DBAs had a decent plan in which they would restore to a NAS device that they had onsite. They had tested this feasibility earlier and found that they could run their E-Business Suite on the system, albeit at a significant loss of performance. The performance degradation was acceptable to the end users because this was a failure type of scenario.

As we reviewed the environment, we drew one interesting conclusion related to multiplexing control files, redo logs and archive logs. We found that the customer was doing an excellent job of getting the archive logs synched to the NAS for protection using the "Rsync" utility. However, they were not doing anything to protect their control files or redo logs. The client told us about testing they had performed

earlier where it was determined the NAS was ultimately too slow to multiplex to. The storage administrators argued that there was no need to worry about losing the SAN because everything was mirrored, redundant and protected.

Being a DBA though, I could not accept that answer on its merit because there is always a chance – though perhaps only a very slim one - that the SAN could have an issue. The customer could handle protecting the control files much the same way they were handling the archive logs, by entering the database and using the backup control file commands to dump both binary and text-based copies. These copies would then be synched down to the NAS. With that problem solved, we went on to considering the redo logs.

There are many ways to handle the redo scenario, but these are not the real focus of this newsletter's tip or trick. In fact, the information above was only useful to this tip as a back story. So, why did we talk about the disk in the opening paragraph? Well, one of the exercises we went through was to test how fast the disks were actually able to write. While there are several good monitoring programs that will help with this, we decided to use the old fashioned Unix `dd` command to write a bunch of zeros to the file. We wrote a script that writes a 1GB file and uses the `/usr/bin/time` command to determine the "real" execution time. This script writes the file, records the time, erases the file, and records the time. We placed it in a loop with about 10 minutes or so between executions. This gave us multiple data points and also spread the times out so they were more indicative of the processing environment. We'll run this script several times throughout the day to make sure we have data during busy times and not busy times (if those actually exist). We'll take the findings and chart them afterwards, which is often a telling story.

In the case of this client we expanded the normal scope to cover their SAN mountpoints, NAS mountpoints and their local disk. When we worked through the results, we found something that was interesting. Many DBAs read whitepaper after whitepaper, sometimes confirming what we already know, and sometimes learning something new. This will undoubtedly be the same for many of you. What we saw was that the SAN, on average, was significantly faster than either the NAS or the local disk. The NAS device tested better than the local disk, which we would not have expected (given what we were testing), but they were generally close in times.

The "ah-ha" we found was that there were periods where both the NAS device and the local disk outperformed the SAN. What we had found was several cases of "cache cram". Cache cram is a situation where the cache in the SAN device is filled faster than it can write to disk. When this happens the SAN stops accepting input until it can catch up by writing a predetermined amount of the cached information to the disk. These events cause performance slow downs and may be responsible for intermittent or transient slowdowns in a system.

Below is the command we used to write out the 1 GB files:

```
${TIME} -p dd if=/dev/zero of=${OUTFILE} bs=8192 count=131072  
2>>${REPORT}
```

The variable `${TIME}` equals `/usr/bin/time` and `${REPORT}` is set to the output filename.

So, when was the last time you checked your disk activity?



Confio builds performance management software that improves the effectiveness of IT systems and the people who run them. The Confio Igniter Suite currently includes products that help Oracle DBAs, managers, and developers continuously monitor the performance of databases and the applications that depend on them, isolate specific problems, and identify solutions in production systems.

Ignite for Oracle is an Oracle database performance tool that identifies the toughest problems affecting Oracle databases and any application that depends on them.

Ignite for RAC is an extension of Ignite for Oracle that identifies the overhead and imbalance associated with Oracle RAC™ database configurations. Ignite for RAC identifies the Wait-Time caused by RAC overhead, and isolates performance problems at the cluster as well as the individual node level.

Ignite for SQL Server extends Wait-Time analysis for the first time into the Microsoft environment. For sites with mixed SQL Server and Oracle systems, it provides a single monitoring system to identify and diagnose problems on both platforms. Wait-Time analysis provides detailed, actionable data on database performance issues down to the SQL statement layer.

Ignite for Java works together with Ignite for Oracle in environments using a 3 tier Web/Java/Oracle architecture. Ignite for Java goes deep into the application running on J2EE app servers like Weblogic, Websphere, JBoss and Oracle AS to expose the bottlenecks, whether they are at the J2EE layer or database layer.

Come check out **Confio** at the Convention Center!



APPSPERF 2008 DENVER COLORADO APRIL 17 OUR INAUGURAL WORKSHOP

APPSPERF 2008 is a workshop at *COLLABORATE 08* that will be held at the Hyatt Regency Hotel on Thursday April 17.

APPSPERF 2008 will provide attendees with detailed tuning expertise and tips based on real-life deployments and extensive experience. Attendees will also learn how to effectively tune and manage the entire E-Business Suite environment using a holistic approach and well-proven techniques. Attendees will have a unique opportunity to network with the world's top Applications performance experts.

This workshop is a full day Conference Workshop that has a separate fee (\$595) from the conference fee. Register onsite at the Convention Center. The workshop is being held at the Hyatt Regency Hotel.

Agenda Overview:

8:30 AM Keynote - Cliff Godwin – Senior Vice President of Applications Development, Oracle Corporation.

9:00 AM - 10:30 AM Sessions

Room 1: Tuning the Application Server and JVM - Ahmed Alomari, Cybernoor Corporation

Room 2: Managing Your E-Business Suite Systems Workload - Larry Klein, Sleuth Systems

10:30 AM - 11:00 AM Break

11:00 AM - 12:30 PM Sessions

Room 1: 11g New Optimizer Features - Dr. Mohamed Zait, Oracle Corporation

Room 2: Tuning the Concurrent Manager - Jerry Ireland and Mark Farnham, Rightsizing, Inc.

12:30 PM - 1:30 PM Break for Lunch

1:30 PM - 3:00 PM Sessions

Room 1: SQL Performance Management in Oracle 11g with Automatic SQL Tuning and Real-Time SQL Monitoring; Real Application Testing in Oracle 11g with SQL Performance Analyzer and Database - Dr. Benoit Dageville, Oracle Corporation

Room 2: Customer Success Stories: Optimizing the E-Business Suite with Real Application Clusters (RAC) - Ahmed Alomari, Cybernoor Corporation

3:00 PM - 3:15 PM Break

3:15 PM - 4:45 PM Sessions

Room 1: Understanding and Tuning the Shared Pool - Kumar Rajamani, Oracle Corporation

Room 2: Tuning Oracle Workflow - Karen Brownfield, Solution Beacon

4:45 PM - 5:00 PM Break

5:00 PM - 6:00 PM Meet the Applications Performance Experts: Q&A

PRESENTATION	DATE	LOCATION
Scaling to Infinity: Partitioning Data Warehouses in Oracle – <i>Tim Gorman, Evergreen Database Technologies, Inc.</i>	Seminar 9:00am-5:00pm Sunday, April 13th	Korbel Ballroom 1F
PeopleSoft for the DBA: PeopleSoft Performance Optimisation – <i>David Kurtz, Go-Faster Consultancy Ltd.</i>	Seminar 9:00am-5:00pm Sunday, April 13 th	Korbel Ballroom 1D
PeopleTools 8.48/8.49 – New Database Features – <i>David Kurtz, Go-Faster Consultancy Ltd.</i>	3:30pm-4:30pm Monday, April 14 th	
Getting Started with Forecasting Oracle Performance – <i>Craig Shallahamer, OraPub, Inc.</i>	3:30-4:30pm Monday, April 14th	Korbel Ballroom 3A
Fueling the Workflow Engine – Writing Custom PL/SQL to be Called from your Workflow Processes – <i>Daniel Stober, Intermountain Healthcare</i>	9:45-10:45am Tuesday, April 15th	Room 404
Migrating 11i to RAC Using DataPump – <i>Mike Swing, TruTek and Fred Linnenbrink, Director of Database Systems and Architecture at Metro St Louis</i>	3:30-4:30pm Tuesday, April 15th	Korbel Ballroom 1C
PeopleTools 8.48/8.49 – New Database Features – <i>David Kurtz, Go-Faster Consultancy Ltd.</i>	4:45-5:45pm Tuesday, April 15 th	Room 502
Getting Started Optimizing Oracle Performance – <i>Craig Shallahamer, OraPub, Inc.</i>	8:30-9:30am Wednesday, April 16th	Korbel Ballroom 2A
An Introduction to SQL*Trace, TKPROF and Execution Plans for Developers – <i>David Kurtz, Go-Faster Consultancy Ltd.</i>	1:30-2:30pm Wednesday, April 16 th	Korbel Ballroom 3A
MAID for ILM Archival Storage for Data Warehouses – <i>Tim Gorman, Evergreen Database Technologies, Inc.</i>	3:15-4:15pm Wednesday, April 16th	Korbel Ballroom 2A
PeopleSoft: A Properly Instrumented Application – <i>David Kurtz, Go-Faster Consultancy Ltd.</i>	3:15-4:15pm Wednesday, April 16 th	Room 505
Customization Survival Guide: How to Use E-Business Utilities to Migrate Your Custom Code – <i>Brad Simmons and Donna Campbell, Los Alamos National Laboratory</i>	11:00am-12:00pm Wednesday, April 16th	Room 207
PeopleSoft: A Properly Instrumented Application – <i>David Kurtz, Go-Faster Consultancy Ltd.</i>	8:30-9:30am Thursday, April 17th	
AppsPerf 2008 – <i>Ahmed Alomari, Jerry Ireland, Mark Farnham, Michael Brown</i>	Seminar 8:30am-6:00pm Thursday, April 17th	Hyatt Regency Hotel