

nocoug

JOURNAL

Official Publication of the Northern California Oracle Users Group

VOL. 16, No. 2 · MAY, 2002

\$15

Don't Miss Our Spring Conference!

The NoCOUG Board has planned a great conference that you won't want to miss. It's taking place on **Thursday, May 16 in Sunnyvale**. Sponsored by Lockheed Martin, the one-day conference will have **three parallel tracks:**

- Database Administration
- Application Development
- Data Warehousing



We're also happy to report that we will be holding two separate roundtable discussions after lunch focusing on development and DBA topics. This period will be moderated by a NoCOUG track coordinator and will be a time for attendees to gather to ask questions and exchange ideas. Bring your questions and latest challenges to discuss with the group!

Registration begins early. Please arrive early and allow extra time for check-in as Lockheed is a secure area. For more details, see our website at www.nocoug.org. ▲

Email Is the Way to Go!

The best way to stay up-to-date on NoCOUG happenings is through email notification. You'll hear about upcoming meetings and get updates of changes and additions to our schedule. If you do not receive email notifications of our meetings, please register on our website at www.nocoug.org.

TABLE OF CONTENTS

Don't Miss Our Spring Conference!	1
Email Is the Way to Go!	1
From the Editor	2
NoCOUG Board	2
User Groups.....	2
Publication and Submission Format.....	2
Advertising Rates.....	2
President's Message	3
Resource Corner	4, 18
Tech Tips	4, 14
Board Member Spotlight.....	5
OS Snapshots for Backup.....	7
Oracle Performance Triage	11
Encryption of Data at Rest	19
Treasurer's Report.....	23
Sponsorship Appreciation	23
Hackproofing Oracle Application Server	24
NoCOUG Spring Conference Session Descriptions.....	27
NoCOUG Spring Conference Schedule	28
—ADVERTISERS—	
XIOtech	3
Quest Software	6
BMC Software	10
Promatis	12
Informatica.....	12
dbDoctor	20
Promeria	25
Database Specialists	26
Embarcadero Technologies	26
LECCOTECH	27

From the Editor

A full year has gone by since I took over the responsibilities of editing the NoCOUG Journal. It's been a great experience, full of fun and challenges. I've enjoyed it so much that I decided to volunteer for the position for another year. Colleen Childers of Database Specialists has been a great help to me in putting together the NoCOUG Journal each quarter. Her creativity, writing, and editorial skills have been invaluable. I'm sure I could not have done it without her. I'd also like to thank the NoCOUG members who submit articles and tips for publication. NoCOUG exists for the education and representation of the users of Oracle Corporation's database and tools software, and those who contribute to the NoCOUG Journal help achieve that vision.

If you would like to get involved, now is a great time. There are a number of short- and long-term assignments that you can take on, depending on your time and interests. Just get in touch with me at journal@nocoug.org. I look forward to hearing from you.
Lisa Loper, *Journal Editor*

Other User Groups

Local

NorCalOAug – Northern California Oracle Applications Users Group

- **Contact:** Michael Capelle (650) 562-1167
- **Email:** capelle@tru-course.com
- **Website:** www.norcaloaug.org

Sacramento

SacOUG – The Sacramento Oracle User Group

- **Contact:** Ravi Verma (916) 705-3261
- **Email:** ravi.verma@telcommand.com
- **Website:** www.sacoug.org

International

IOUG-A – International Oracle Users Group of the Americas

- **Website:** www.ioug.org

U.S. Domestic

OAUG – Oracle Applications Users Group

- **Website:** www.oaug.org

ODTUG – Oracle Development Tools User Group

- **Website:** www.odtug.com

Canvassing calls by employment recruiters to local chapter contacts is strongly discouraged.

NOCOUG BOARD

President

Joel Rosingana, Independent Consultant
joelros@pacbell.net

Vice President

Roger Schrag, Database Specialists, Inc.
rschrag@dbspecialists.com

Treasurer/Secretary

Judy Lyman, Contra Cost County Public Works
gooma@california.com

Membership

Vacant Position

Webmaster

Vadim Barilko, Independent Consultant
vabarus@onebox.com

Journal Editor

Lisa Loper, Database Specialists, Inc.
lloper@dbspecialists.com

Vendor Relations

Ganesh Sankar, Providian Financial
bgs2k2@yahoo.com

IOUG-A Representative and Past President

Vilin Roufchaie, Cingular Wireless
vilin.roufchaie@cingular.com

Members, At Large

Hamid Minoui, Fritz Companies
hamid.minoui@fritz.com

Darrin Swan, LECCOTECH
darrin@leccotech.com

Publication and Submission Format

The NoCOUG Journal is published four times a year by the Northern California Oracle Users Group approximately two weeks prior to the quarterly regional meetings. Please send your questions, feedback, and submissions to: Lisa Loper, NoCOUG Journal Editor, at journal@nocoug.org.

The submission deadline for the upcoming August issue is July 1, 2002. Article submissions should be made in electronic format via email if possible. Word documents are preferred.

NoCOUG does not warrant the NoCOUG Journal to be error-free.

Copyright © 2002 by the Northern California Oracle Users Group. Permission to reproduce articles from this publication, in whole or in part, is given to other computer user groups for nonprofit use, with appropriate credit to the original author and the Northern California Oracle Users Group Journal. All other reproduction is strictly prohibited without written permission of the editor. Two copies of each reprint should be sent to the editor.

ADVERTISING RATES

Contact: Nora Rosingana

325 Camaritas Way
Danville, CA 94526
Ph: (925) 820-1589

The NoCOUG Journal is published quarterly.

The rates are:

Size	Per Issue	Per Year
Quarter Page	\$100	\$320
Half Page	\$200	\$640
Full Page	\$400	\$1,280

Personnel recruitment ads are not accepted.



A NoCOUG New Year

by Joel Rosingana, President, NoCOUG

The meeting at the Oracle Conference Center in February was a resounding success! We had over 200 attendees. It was heartening to see many new members in the mix with the “Old Timers” who attend most all of the meetings. It was a pleasure to meet many of you at the meeting and at the after-hours networking get-together at The Players Club.

This meeting could not have happened without the dedication and tireless work of the Oracle Staff. **Kate Kerner**, Senior Manager of Global User Group Programs and her staff were the prime movers of this event. They arranged for the keynote and a full track of Oracle speakers. The **Ken Jacobs** Keynote was a high point and a great success. I would like to thank Ken for taking time out of his busy schedule. This meeting was fully hosted by Oracle. This means that Oracle not only provided the facility, they also provided all the food. This really helps in keeping membership costs down. We all can thank Oracle, by showing our continued support. Thanks Kate, Brenda, and Jeannie.

Our next quarterly meeting will be at Lockheed Martin in Sunnyvale. Dave Erwin, of Lockheed and a long-time NoCOUG member, makes this event happen. Dave has always been there for us and sometimes on very short notice. Lockheed has since become the annual May meeting location. The meeting will be on Thursday, May 16. Check the website for the agenda updates. Security will be a little tighter this year, so, plan on a little extra registration time. The required documentation will be the same as last year. Everybody, including American citizens, will need a **Photo ID**. Foreign nationals will need a current **Passport**. Resident

aliens will need a **Green Card**. See you all there!

My message last quarter mentioned that we hoped to have online registration soon. Thanks to Vadim Barilko, our Webmaster, we are now online. The members frequently asked to register with a credit card. Now you can. Though we didn't have the system in production for the main registration rush, we have had quite a few use the system since. It seems to be working fine. In Phase 2 we plan to make the system more convenient for re-registration. We'll keep you posted.

I'll finish up with IOUG Live. Since the Journal will hit the street after IOUG Live in San Diego, I'll just say I hope you all attended and had a great time. I will be attending as your Alliance Partner Rep. Vilin Roufchaie, our board's IOUG Rep, has pressing job responsibilities and won't be able to attend. I shall use this space next time to update you on any IOUG/NoCOUG Alliance partner issues. Until then! ▲



Joel Rosingana

New Online Registration

We added a new option for membership registration on the NoCOUG website (www.nocoug.org). Now you can register online with secure payment through Paypal. All major credit cards are accepted. If you haven't renewed your membership for 2002, now is the time. Members are already using this option. Enjoy it!

—Vadim Barilko,
NoCOUG Webmaster



Oracle Racing Team Takes on Seattle and the World on the High Seas

This fall in Auckland, New Zealand, Larry Ellison and his Oracle Racing team will be battling nine sailing teams from around the world, including Paul Allen's One World Challenge from Seattle, for a chance to win back America's Cup from the Royal New Zealand Yacht Squadron. In addition to being an accomplished skipper in his own right, Mr. Ellison will have some local talent in Kiwi Chris Dickson, one of the best match racers in the world.

To enter the challenge for the cup, Larry Ellison and Oracle purchased the boats and assets of Paul Cayard's AmericaOne team, and declared through the Golden Gate Yacht Club of San Francisco. They have a budget of over \$80 million, the largest declared battle chest of any of the syndicates. Team Dennis Connor is also in the running from New York, with Computer Associates providing the major backing. In addition to the three American challengers, there will be teams from Great Britain, France, Sweden, Switzerland, and Italy.

Of course, technology will be playing a major part in the design of the boats. Oracle Racing is utilizing a



*Mike DeVito,
Oracle DBA and
avid sailing fan*

Compaq AlphaServer SC supercomputer and PTC's Pro/ENGINEER software to design two new 80-foot America's Cup Class (ACC) yachts. The action begins in October with the Louis Vuitton Cup, with America's Cup Matches starting in February 2003.

Official Sites for Information:

www.lvcup.org
www.oracleracing.com

'Scuttlebutt' (unofficial sites):

www.americas-cup.co.nz/
www.hauraki-news.com/

—Mike DeVito, Oracle DBA
mike@devito.com



TECH TIPS

Create Primary Key in Parallel

1. Create the primary key constraint (disabled)

```
alter table DEPT add (  
constraint PK_DEPT  
primary key (DEPTNO))  
disable primary key;
```

2. Create a unique index with the same name and columns as the primary key

```
create unique index PK_DEPT on DEPT (DEPTNO)  
tablespace DATA_1  
parallel (degree 5)  
unrecoverable;
```

3. Enable the primary key

```
alter table DEPT enable primary key;
```

— from Stewart McLaughlin, Oracle DBA
<http://www.stewartmc.com/oracle/tips.html>



A Butcher, a Blackbelt, a NoCOUG Board Member...

by Colleen Childers

Many of us have held interesting jobs in our past. Some are certainly more glamorous than others, and some, well, we'd probably rather forget about. Darrin Swan, member of the NoCOUG Board of Directors is a person who has had many intriguing jobs. His first was working the 4:00 a.m.-to-noon shift at Hewlett Packard where he ran the print servers of HP3000 and 9000 systems in their data center. The eclectic mix of jobs that followed was: martial arts instructor (in which he holds a black belt), butcher, and operations analyst, just to name a few. He then moved into his niche of solution selling, which he has been doing for the past eight years.

One job that Darrin takes very seriously is his position on the NoCOUG Board. Through his exposure as a vendor at our quarterly conferences, Darrin saw an opportunity to put his experience to good use. After all, as the regional account manager in charge of all sales-related efforts in the Bay Area and Pacific Northwest for LECCOTECH, he is quite used to helping groups develop solutions. So, in February 2001, Darrin joined the Board of Directors to see where he could help.

Darrin is a board member at large, and he wears a few different hats. Throughout the year, Darrin attends the board meetings to offer his insight. In preparation for the quarterly conferences, he helps to coordinate the speaker tracks. On the day of the conference, you can usually find him moving about, pitching in where needed.

Darrin enjoys working with the user community because he enjoys



Darrin Swan

learning firsthand what the problems and interests of the user community are and looks for ways he can help. He takes satisfaction in taking the "consultative" approach to problem solving

and works with others' best interests in mind.

A resident of North Beach in San Francisco, Darrin enjoys dining out and spending time with his family, which is very large. He's number five in the lineup of eight children. Growing up in the Bay Area, he learned a lot about the Hewlett Packard way, as his father has worked there for 37 years and one of his brothers has worked there for ten years.

One of Darrin's favorite pastimes is travelling with his wife. Some of his more memorable trips have been this year's trips to New York and New Orleans and last year's trip to Istanbul, Turkey. He's been to some fantastic places like Prague, Austria, and Italy. Darrin would be the first to tell you that the best way to get around Europe is to drive. He's an avid fan of driving trips, and swears that driving is much easier in Europe than in the Bay Area!


When asked what famous people he'd like to invite over for dinner, Darrin responds quickly with an impressive list: James Joyce, Bill Clinton, Bill Gates, and Madonna. The Irishman Joyce is Darrin's favorite author. He admires Bill Clinton's ability to "carry a group." Darrin appreciates the visionary qualities of Bill Gates

When asked what famous people he'd like to invite over for dinner, Darrin responds quickly with an impressive list: James Joyce, Bill Clinton, Bill Gates, and Madonna.

and his guerilla tactics that have allowed him to dominate the industry. And, to spice things up, Darrin would invite the talented Madonna, whom he considers an "amazing business woman."

So, at the next conference, please be sure to say hello to Darrin. Maybe you can swap stories about your interesting jobs. Maybe he'll tell you about his job selling golf balls to golfers as a kid! ▲

With the Right Tools, Anything is Possible



Quest Software Brings Higher Intelligence to Database Development

Smart people with good tools can mystify the masses and achieve the impossible. With Quest Software's market-leading tools for Oracle development, so can you. Quest Software offers the complete solution to design, develop, test and deploy code quickly and accurately, leaving you plenty of time to consider the possibilities.

The evidence is clear: more than 200,000 users worldwide believe Quest Software's development tools are the intelligent choice. Download your free trial versions from www.quest.com and see for yourself.



www.quest.com • 1.800.306.9329 (U.S.) • 1.949.754.8000 (outside U.S.)

OS Snapshots for Backup

Utilizing operating system snapshots for quick and painless Oracle database backup and restore.

By Kenny Gorman (kenny@kennygorman.com)

Toolbox: This example utilizes Oracle 8.1.7.2.0 on SuSE Linux 7.3 with the 2.4.16-64GB_smp kernel. It assumes basic knowledge of Oracle Hot Backup methods, and the Linux operating system. The Oracle database must be running in archive log mode. It is also necessary to have root access to perform the examples shown.

As the typical database size grows larger and larger, even the most ambitious backup plans begin to strain. With larger data volumes come longer backup times. In today's 24/7 world, allowing a database to be at risk and/or run at degraded performance due to a backup for long periods of time is just unacceptable. Many companies and institutions turn to incremental backups, or utilize expensive disk-mirroring technology to help keep the backup window time down. Some companies simply take fewer backups. Obviously, these techniques may require lots of expensive hardware or add complication and at worst, risk corporate data.

In addition, many database administrators take backups to disk in order to alleviate long backup windows while waiting for tape (and sometimes to reduce the Mean Time To Recover). Taking a backup to disk takes additional expensive disks, arrays, or NFS servers.

But there is another option that is under utilized by most database administrators. This technique is called OS file system snapshots. OS snapshots take a fraction of the time of file copies and are many times faster than backups directly to tape devices. OS snapshots are a simple and effective technique for reducing the backup window for large databases. OS snapshots can also reduce the amount of disk needed for backups to disk. In many cases it can be 10% of the total amount of disk.

Utilizing OS snapshots can be a useful technique to any database administrator and some system administrators running databases on UNIX and Linux.

Background / Overview

OS snapshots are a facility of the Linux LVM (Logical Volume Manager). A snapshot volume, once mounted, contains all the files that existed on the logical volume frozen at the point in time when the snapshot was taken. A snapshot volume can exist for one logical volume. You can have multiple snapshots of a single logical volume. In our case, the files on the snapshot volume are Oracle datafiles. Before we snapshot the logical volume where our Oracle datafiles reside, we will place Oracle in Hot Backup Mode. Once the snapshot is complete, we remove Oracle

from Hot Backup Mode. Depending on the activity on your file system, the time for a snapshot to complete can take mere minutes. Once the snapshot is taken, then the backup to tape can occur from the snapshot, and Oracle remains unaffected and is in service. Snapshots are a facility of the volume manager software. Increasingly, UNIX vendors and third parties are offering volume managers that have snapshot facilities. In this article we will explore Linux LVM snapshots.

Snapshots work for backups because Oracle has a facility called Hot Backup Mode. Hot Backup Mode allows for recovery of inconsistent (or "fuzzy") datafiles using the archive logs. When in Hot Backup Mode, Oracle writes not just the change vector for a block, but also the entire value for the change to the redo log. The Oracle datafiles are still written to, but that doesn't matter, because the archive logs hold all transactions that took place. When a snapshot takes a "picture" of the datafiles in Hot Backup Mode, they are inconsistent (or "fuzzy"). We don't care that the files are inconsistent because during recovery, Oracle will apply the archive logs and bring the database into a consistent state.

This article assumes your system is Linux 2.4.16 kernel with support for LVM and ReiserFS. The system used in this example is SuSE Linux 7.3, and comes with the needed software bundled right in. If you are unfamiliar with installation of software on Linux and rebuilding your kernel, I recommend you utilize the SuSE distributions. Of course because of the nature of Linux, pretty much any Linux distribution can be set up with the needed software. This article also assumes you are using Oracle 7.3+ and the database is running in archive log mode.

Step 1: Preliminary Setup

In order to create a snapshot, you must first have a logical volume to take a snapshot of. In order to create a logical volume, you must first have a disk to assign to the logical volume via a volume group. These steps are also outlined at www.sistina.com, the creator of the LVM software.

First the disk will need to be formatted using `fdisk`, and, in this case, a physical volume that spans the entire disk will



Kenny Gorman

be created. It is important to note that the disk type needs to be assigned the type 8e. This is shown below.

```
$>su -
$>fdisk /dev/sdb

The number of cylinders for this disk is set to 4492.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)

Command (m for help): p

Disk /dev/sdb: 255 heads, 63 sectors, 4492 cylinders
Units = cylinders of 16065 * 512 bytes

   Device Boot      Start         End      Blocks   Id  System
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-4492, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-4492, default 4492):
Using default value 4492

Command (m for help): t
Partition number (1-4): 1
Hex code (type L to list codes): 8e
Changed system type of partition 1 to 8e (Linux LVM)

Command (m for help): p

Disk /dev/sdb: 255 heads, 63 sectors, 4492 cylinders
Units = cylinders of 16065 * 512 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1          1           4492   36081958+   8e  Linux LVM

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: If you have created or modified any DOS 6.x
partitions, please see the fdisk manual page for additional
information.
Syncing disks.
$>
```

Now I have a physical volume to work with. To create the volume group and subsequent logical volume, we must first initialize the physical volume.

```
$>pvccreate /dev/sdb1
pvccreate - physical volume "/dev/sdb1" successfully created
```

This creates a volume group descriptor area (VGDA) at the start of the disks. Next, the disks need to then become part of a disk group. In this case, I define one disk group for the physical partition.

```
$>vgcreate datagr /dev/sdb1
vgcreate - INFO: using default physical extent size 4 MB
vgcreate - INFO: maximum logical volume size is 255.99 Gigabyte
vgcreate - doing automatic backup of volume group "datagr"
vgcreate - volume group "datagr" successfully created and
activated
```

This creates a disk group assigned to a physical partition. Now I need to assign the volume group to a logical volume. I will create the logical volume and leave some extra space for my snapshot. This extra space is arbitrary and depends on your transaction volume and length of snapshot for space consumption. In my case, I make it very large (relatively) so there is no problem with space. I

create the volume as 32GB, thus leaving approximately 4GB for the snapshot (36GB disk—32GB volume).

```
$>lvcreate -L32G -ndatavol datagr
lvcreate - doing automatic backup of "datagr"
lvcreate - logical volume "/dev/datagr/datavol" successfully
created
```

Once we have the logical volume then we need to create a filesystem on top of it. For this example, I will use ReiserFS. Details on ReiserFS can be found at www.reiserfs.com. You should carefully evaluate the requirements of your system to be sure ReiserFS is the correct choice.

```
$>mkreiserfs /dev/datagr/datavol
Initializing journal - 0%....20%....40%....60%....80%....100%
Syncing..ok
```

This may take some time to create depending on your system. Once this is done you can mount it. I have to create a directory to mount it to in normal UNIX fashion.

```
$>mkdir /oradata
$>chown oracle:dba /oradata
$>mount /dev/datagr/datavol /oradata
```

Now that the volume is mounted, then you can go ahead and create a database if you want. For this example, I am putting all the datafiles on one volume. In your implementation that probably won't make sense, but this should give you a good idea of how to use snapshots nonetheless.

Creating a snapshot volume

You have seen how to create a regular logical volume above. Now we need to create the snapshot. Snapshots are simply logical volumes created with special syntax. The snapshot volume syntax specifies the volume you are taking the snapshot of.

First we want to create a mount point for the snapshot.

```
$>mkdir /oradata_snap
$>chown oracle:dba /oradata_snap
```

Before you take the snapshot, you want to be sure to place the tablespace(s) that have datafiles on your logical volume in Hot Backup Mode. Otherwise, your backup will be corrupted. In our case, the entire database is on one volume. So it makes sense to put all the tablespaces in Hot Backup Mode all at once. You probably won't want to put all tablespaces in Hot Backup Mode on your production systems. Instead, you would cycle through your tablespaces, putting them in Hot Backup Mode one by one, taking snapshots, and then removing them from Hot Backup Mode.

Put all the tablespaces in Hot Backup Mode.

```
SQL> DECLARE
2   sqlstmt varchar2(100);
3   CURSOR tab_cur IS
4   SELECT tablespace_name FROM dba_tablespaces;
5   BEGIN
6   FOR tab_rec IN tab_cur
7   LOOP
8     sqlstmt := 'ALTER TABLESPACE '||tab_rec||' BEGIN BACKUP';
9     EXECUTE IMMEDIATE sqlstmt;
10  END LOOP;
11  END;
12 /
SQL> PL/SQL procedure successfully completed.
```


Once the tablespaces are in Hot Backup Mode, we can take the snapshot of the volume.

```
$>lvcreate -L2G -s -noradata_backup /dev/datagrpd/avol
lvcreate - WARNING: the snapshot will be automatically disabled
once it gets full
lvcreate - INFO: using default snapshot chunk size of 64 KB for
"/dev/datagrpd/oradata_backup"
lvcreate - doing automatic backup of "datagrpd"
lvcreate - logical volume "/dev/datagrpd/oradata_backup"
successfully created

$>mount /dev/datagrpd/oradata_backup /oradata_snap
mount: block device /dev/datagrpd/oradata_backup is write-
protected, mounting read-only
```

You have created the snapshot! You can now remove Oracle from Hot Backup Mode.

```
SQL> DECLARE
2  sqlstmt varchar2(100);
3  CURSOR tab_cur IS
4  SELECT tablespace_name FROM dba_tablespaces;
5  BEGIN
6  FOR tab_rec IN tab_cur
7  LOOP
8  sqlstmt := 'ALTER TABLESPACE '||tab_rec||' END BACKUP';
9  EXECUTE IMMEDIATE sqlstmt;
10 END LOOP;
11 END;
13 /
SQL> PL/SQL procedure successfully completed.
```

You can now copy (tar, cpio, Netbackup, etc.) your datafiles to tape from the snapshot at your leisure without needing Oracle to be in Hot Backup Mode, or worse, taking Oracle down. You will also need to be sure to back up the archive logs for the period that the database was in Hot Backup Mode.

```
$>tar -cvf /dev/rmt0 /oradata_snap/* /oraarch/*
```

Once the backup is complete, unmount and remove the volume. Once you do this, the snapshot is lost forever. Snapshots don't keep the data between reboots, so you will not want to rely on the snapshot itself for a backup, you will want to copy the data off to a tape drive or some other type of permanent media.

```
$>umount /oradata_snap
$>lvremove -f /dev/snappgrp/oradata_backup
lvremove - doing automatic backup of volume group "datagrpd"
lvremove - logical volume "/dev/datagrpd/oradata_backup"
```

Restoration

Of course, none of this is any good if you don't test the restore. In order to restore, the tape media or other storage would need to be utilized, and the copy of the data redirected back to the original source directory. For instance, using the above example, I backed up my snapshot from /oradata_snap and archive logs from /oraarch. I would redirect the restore to put the files from /oradata_snap on /oradata, and the files from /oraarch back to the /oraarch volume. At this point, the database could be recovered and all would be well.

```
$>tar -xvf /dev/rmt0 /oradata
```

```
SQL>...
SQL>Recover database until cancel.
SQL>...
```

Weaving snapshots into your backup scheme

Obviously, there is more to using snapshots when you consider most production environments. Mainly you will want to script both the snapshots and the Hot Backup Mode. You will also need to consider archive logs, and make sure you have the correct set of archive logs to recover the database. In my environments, I use the following order of events:

1. switch archive logs
2. start backup for tablespace 1
3. snapshot filesystem for tablespace 1
4. end backup for tablespace 1
5. repeat for all tablespaces
6. create control file to trace
7. switch archive logs
8. snapshot the archive log filesystem
9. run tape backup software to 'grab' all of the snapshot volumes

Monitoring the completion and ensuring a successful snapshot are, of course, essential. I recommend reading the documentation at www.sistina.com to understand all of the different parts of the LVM. There are lots of powerful things you can do with it. You may over time need to extend the sizes of the data and/or snapshot volumes. Management of the sizes and interdependencies of the data and its snapshot are important.

As of this writing, LVM was still in beta. So care should be taken in how you or if you deploy it in a production environment. I recommend keeping up to date on the latest version and subscribing to the LVM mailing list at: <mailto:linux-lvm@sistina.com>.

Conclusion

Snapshots can be really cool, and can reduce your backup window substantially. They can also save on disk space by alleviating the need to have extra disks just for backups. Additionally, the components shown in this article are free. They come with some major Linux distributions, but can also be downloaded and compiled separately.

Kenny Gorman (kenny@kennygorman.com) is an independent Oracle consultant. He provides Oracle DBA services to clients around the San Francisco Bay Area. He specializes in High Availability, Monitoring and Management, and Veritas. He is an OCP and Oracle Master and has over six years of experience as an Oracle DBA at various companies and startups in and around the Bay Area.



SPACE EXPERT™ GIVES YOU A SMARTER WAY TO
visualize YOUR DATABASE'S SPACE PROBLEMS.

It's not always easy to see where pesky space problems are lurking, ready to degrade the performance of the database that your business depends on. And it's not always easy to keep your database operating at peak levels without lots of highly trained DBAs.

Until now, that is.

Introducing Space Expert™ for Oracle from BMC Software. It intelligently and automatically visualizes, isolates, analyzes and corrects space-related problems. Enabling even a novice DBA to perform the work of many. Which means the IT department's SLAs are good as gold. IT managers can concentrate on managing. Your database's performance is enhanced. And your company is more competitive than ever.

See for yourself. **Register for a 30-day trial copy of Space Expert for Oracle today at www.bmc.com/distdata/spaceexpert.** Then you'll discover just how easy we make it for you to help your database, and your business, really perform. Because it takes intelligence, not hocus pocus.



Oracle Performance Triage: Stop the Bleeding!

Craig A. Shallahamer, OraPub, Inc.

Abstract

Triage: A system designed to produce the greatest benefit from limited treatment facilities for battlefield casualties . . . If you manage Oracle performance, you live in a battlefield. While people's lives usually are not at risk, you, your colleagues, and your family's peace of mind are definitely at risk. This paper will help you set up and execute Oracle performance triage. We start by introducing a powerful Oracle triage method followed by an introduction to a number of fundamental performance concepts. Once this foundation has been established, we move into an actual triage case study. The case study is very real, follows the presented triage method, and uses common UNIX performance tools and free OraPub tools.

Introduction

Welcome to the world of Oracle triage. It's not pretty, but it's not boring either. If you are taking the time to read this paper, I suspect you have had the opportunity to experience, to some degree, the joy of intense work pressure surrounded by panic combined with the reality of a significant portion of your company's business being in serious jeopardy.

I wrote this paper to help DBAs enthusiastically triumph while performing Oracle triage and to help better prepare students for taking my Advanced Performance Management for Oracle Based Systems class.¹ The type of training required to do Oracle triage is very unique, requiring a broad range of abilities, such as, effective communication, Oracle internals, Oracle bottleneck detection, and operating system bottleneck detection . . . just to name a few. Oracle triage is not your typical DBA day (hopefully) and so the required training is not typical either.

Overview: How to Perform Oracle Triage

In your mind's eye, live this for a few minutes . . . It's two weeks before Christmas and the company's massive mission-critical production Oracle application just finished a significant upgrade that would be extremely painful to back out of. Within a few minutes after the application was placed back online, not only was the system as a whole obviously very slow, but specific OLTP and batch processes were taking an unacceptable period of time to complete. The pressure quickly escalated and you are just as quickly surrounded by your peers, your boss (who is now perspiring profusely), and your boss's boss. The CFO is so upset he is actually grabbing people by their ties and yelling at them! People begin to panic and act irrational. You are their savior and they are watching your every move!

As if this situation couldn't get any worse, it is complicated by the fact that it is completely reactionary. No one expected this to happen and everyone is getting ready for a nice peaceful Christmas celebration. You know you have to act quickly in your problem determination, analysis, recommendations, and implementation. But it all begins to happen so fast and with such intensity, you soon begin to question every keystroke you feebly attempt to make.

1. Don't panic. No matter how bad you mess this up, Oracle DBAs are in such high worldwide demand you could easily get another DBA job . . . and probably make more money. So you got that going for you.

2. Scope out the situation. Before you begin digging into the details it is important to technically detach yourself from the situation and think like a manager. The first question you ask is, "Can I do this myself?" If you're not sure, you will need to ask for reinforcements. Do not play games here. If the problem is very serious, you have every right to ask for an operating system expert, an Oracle internals expert, an Oracle SQL tuning expert, an application specialist, and someone to help shield you from all the incoming bullets that will be shot directly at you. Hopefully your manager can be your shield. In a triage situation, if you try to be a hero and do it yourself, make sure to save some time to update your resume.

3. Document the current performance. It's as simple as this: You cannot prove you significantly improved performance if you don't write down what performance was before you got involved. To set yourself up for success, make sure to carefully and quantitatively document the key performance areas.

4. Install your tools. Triage requires a different type of tool kit than daily database administration. You need a simple tool kit that will allow you investigate the fundamental computing systems. This needs to be done both quickly and easily from both a detailed and a wide breadth perspective. When I speak of systems, I am referring to the Oracle subsystem, the operating subsystem, and the application subsystem. When I triage, I use my OraPub System Monitor (OSM) tool kit.²

Since you will need to collect data for later analysis as well as look at the systems interactively, get your tools installed and gathering data as soon as possible. You will need good data to build a strong case to support your recommendations.

5. Develop a simple communication strategy. DBAs are not trained in public relations, but you will need to demonstrate some skills in this area. Unfortunately, when the situation heats up, everyone from your boss, your colleagues, your operating system vendor, your application vendor, and

your Oracle sales team will be pointing their grimy little fingers directly at YOU. So be prepared. Come up with a simple, timely, and graphical way to outline your approach, your progress, and your gleaming successes.

6. Isolate the problem. Once your tools have been installed immediately begin to interactively isolate the problem. Once your historical tool kit has had some time to gather data begin to add breadth and depth to your analysis. The fastest and most complete way to isolate the problem is to perform a Holistic Problem Isolation Method (HPIM) analysis [3]. While this is discussed in more detail below, using the HPIM is the proven way to isolate the problem with nearly a zero percent risk of making a mistake. Even on massive and very complex Oracle based systems, you should be able to isolate the problem in only a few hours.

7. Quickly perform your rock-solid analysis. I cannot emphasize this enough; your analysis must be bullet proof. People will be looking for a fall guy and you do not want it to be you. Make sure you use the HPIM and the session wait contention identification approach. Both of these are explained in more detail below. Your recommendations should be very concise, backed up with solid and obvious data, and make sense to nearly everyone. Remember, a weak but well understood analysis is usually better received than a solid but misunderstood analysis.

8. As soon as possible, implement your recommendations. You don't want to make mistakes, so carefully plan out exactly what you are going to do and review it. I like to have someone with me so I don't make any typographi-

cal errors. The more tired you are, the more likely you are to make a mistake, so it's important to have someone check your work.

9. Monitor the situation. After your changes have been implemented, monitor and document their effect. This provides a way to prove you have really improved performance, but it also sets you up for the next round of changes. It is very difficult to effectively optimize a system when changes are continually made without proper monitoring and documentation.

10. Repeat until the heat is off. It is perfectly normal to go through multiple rounds of changes and monitoring periods. Multiple carefully planned and executed cycles are far better than making a bunch of changes in a desperate attempt to improve performance. Mature DBAs show their wisdom by repeatedly making calculated changes that improve performance as opposed to a desperate attempt to do everything at once.

While this all seems relatively simple on paper, it is amazing what a carefully executed Oracle triage can do in just a few short days or even hours. The trick is to not be pulled into the chaos.

Key Triage Concepts

There are a handful of key performance concepts that most Oracle DBAs don't quite understand. Yet if they did, their effectiveness would radically improve. While these concepts are more fully explained in my Total Performance Management paper, I feel they are so important to successful Oracle triage that I included them below.

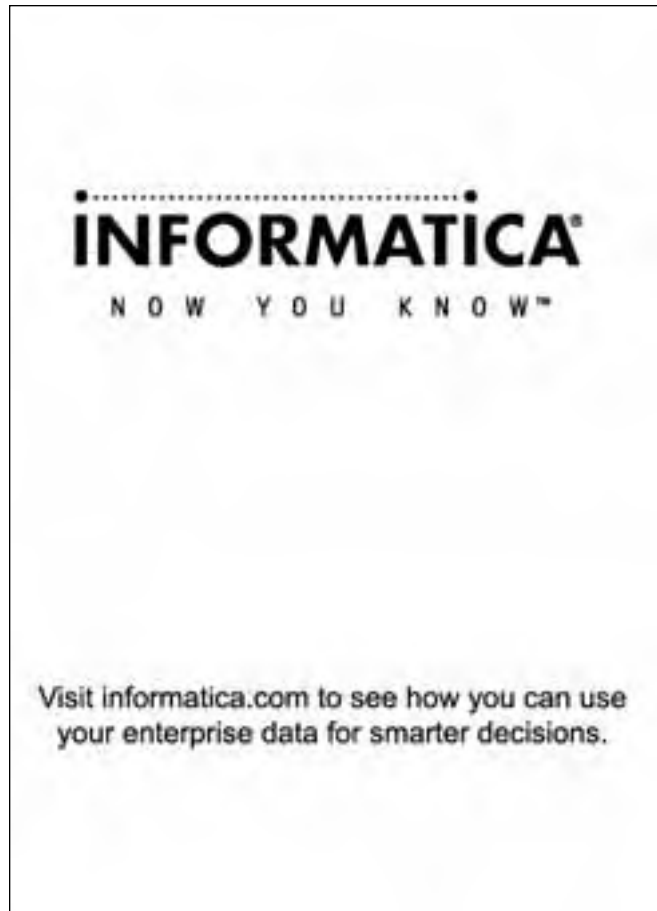


Be the HERO

ROI. Revenue generation. Extending usage. Issues every IT department worries about; especially now.

Be the hero with total and cost-effective enterprise solutions from PROMATIS INCOME Suite. Build business intelligence portals quickly with pre-configured Knowledge Bases. Control and organize knowledge in multiple languages. Integrate all components through the web. Know everything's working with wireless monitoring. Others have pieces; only PROMATIS has it all.


Make Your Business Fly
get-INCOME.com



INFORMATICA[®]
NOW YOU KNOW™

Visit informatica.com to see how you can use your enterprise data for smarter decisions.

The Holistic Problem Isolation Method (HPIM)

The HPIM was first presented in my original 1994 Total Performance Management³ paper. The concept is very simple; yet, following this method has prevented me from making rash contention identification conclusions. This method is so fundamental to performance management, I offer a free Internet Video Seminar specifically on this topic.⁴

Most performance specialists tend to focus on either the Oracle system, the operating system, or the application system. This results in an ill-defined problem definition that will translate into a lopsided solution. In many cases, the solution, while appearing to solve the problem from one viewpoint, results in overall system performance degradation.

The HPIM identifies the bottleneck in each system and then looks for their overlap or where they correlate. Within this overlap lies the first bottleneck. Identifying the bottleneck from three different perspectives and having them support each other builds an extremely powerful base for you to perform your analysis. And the risk of identifying the wrong bottleneck is substantially reduced because it has been correlated from the other two systems.

Methods of Contention Identification

Quickly identifying Oracle system contention will provide you with more time to perform your analysis. There are two basic ways to identify contention within Oracle. The first method, known as the ratio method, basically creates ratios by placing one statistic in the numerator and another in the denominator. The data block buffer cache hit ratio is an example of a ratio. When enough ratios have been calculated over a period of time, coupled with system knowledge and ratio contention identification experience, they will direct one towards where the Oracle contention resides.

A superior method of Oracle contention identification is performed by querying the various session wait performance views. Identifying Oracle contention using the session wait views is so fundamental to quickly identifying contention, I wrote a paper specifically about using the session wait views [5] and offer an Internet Video Seminar on the topic [6]. The paper is freely available on OraPub's website. The session wait views explain specific Oracle contention for the entire system and for specific processes. For example, the session wait views could show session 645 is waiting for a specific database block because of a full table scan. With session wait information you can quickly identify Oracle contention.

The Decreasing Relevance of Increasing Data

Remember when you were a freshman in college and it was easy to swing your grade point average either up or down? Or, when you were a senior and it was nearly impossible to bring your grade point average up very much? It is because historical activities locked you into the past and did not reflect the current reality. When looking at Oracle performance data, one can get caught into this same trap.

Because nearly all Oracle performance views are reset when the Oracle instance re-starts, after a few days of counter incrementation¹, what would radically change a ratio becomes increasingly non-reflective of current activity.

Consider the Oracle data block cache hit ratio (DBCHR). Suppose the instance has been running for 30 minutes with a 70% DBCHR. The 70% is the average since the instance has started, not the average of the past 15 minutes. The average over the last five minutes could have been 95%, but the data from the first five minutes has masked the more recent data. So when looking at many "v\$" related statistics, make sure you know whether you are looking at data from the last x minutes or since the instance has started.

Understanding "Deltas"

A startling example of deltas can be seen when one runs a "top SQL" report at 10 a.m. each morning. Will the "top SQL" reflect intensive nightly batch processing or will it reflect business morning activity? The answer is, there is a good chance the "top SQL" will be skewed towards reflecting late night activity. The way to get around this is to periodically gather "v\$" data and display the change from one time period to the next. I call this change a delta. Deltas highlight more recent activity and provide a much better performance indicator. Good tools, such as the OraPub System Monitor (OSM) tool kit, will always use deltas whenever possible.

Quickly determine which series is increasing and which is decreasing.

```
154833232, 154833357, 154833437, 154834889  
154833149, 154833271, 154933200, 154933240
```

Actually they both are increasing, but it takes time to make this determination. Part of performance triage is quickly doing things. We'll try anything to increase our ability to spot trends or notice significant changes in data. One way to do this is to display the differences in two data points—that is the deltas, instead of displaying the raw data. When looking at six or twelve digit numbers, it is very difficult for our eyes and minds to pick up trends and find when the trend peaks. However, by looking at the deltas, our eyes will quickly notice these trends. Usually the deltas are more important to us than the raw data. Good performance triage tools will always show display deltas.

Two Basic Approaches to Gather Performance Data

An Oracle performance tool will either look at historical data or at current activity. Each approach has significant differences that, when used properly, will improve Oracle triage.

First, let me further define the difference. Most DBAs have a set of SQL scripts they frequently use. These tools usually will directly reference the system as it is currently running. For example, doing a select count(*) from v\$session to determine the current number of Oracle sessions is referencing current system activity. I typically call this the "interactive approach." Contrast this with tools that periodically gather and store data for later retrieval and analysis. I

typically call this the “historical approach.” Both approaches offer distinct advantages compared to the other when used properly.

The interactive approach allows one to quickly dive down into excruciating detail. Detail that most historical based tools don’t capture, such as `v$session_wait` data. However, in a triage environment, things can happen so fast, you will miss important activity while involved in something else. With historical data saved, you can look at your data in a different way or follow different performance analysis paths. The historical approach also offers the ability to highlight trends and overall system activity and contention. The historical approach also allows historical analysis in a way that was not thought of during crunch-time triage. It is important that your tools support both the interactive and the historical based approach. There are many tools on the market today which meet this requirement. OraPub’s OSM tool kit supports both interactive and historical approaches.



TECH TIPS

Handy New Script for Comparing Data in Two Oracle Schemas

Just published on the Database Specialists website is a valuable script for comparing data in two Oracle schemas. This script reads all the rows in a specified list of tables in one Oracle schema and compares the data row for row against tables with the same names in another schema. You’ll get a report with summary counts of differing rows, and SQL queries you can run to view all of the data discrepancies. You can specify columns to exclude from the comparison, such as creation timestamps or IDs generated by a sequence.

This script can be quite helpful when data in multiple schemas is supposed to be the same and you want to make sure that it is. This includes situations such as comparing seed data in separate databases, validating regression test data in a QA environment, or testing data load scripts.

Download this script and several others at <http://www.dbspecialists.com/scripts.html> ▲

Choosing the Appropriate Data Gathering Frequency

Now that you decided to gather historical data, the next question is, “How often should the data be gathered?” I have trained myself to look at data gathered with a frequency between 30 to 60 minutes. Any longer than 60 minutes I can miss out on an important event, but any less than 30 minutes, I will have pages and pages of data to sift through making it more difficult to spot a trend. Or I’ll end up staring at my computer screen for hours watching amazing animated graphics and mumbling words like, “Wow! Did you see that line jump?” or “I wonder how this would look if it had a green background?” Anyway, an important consideration is that gathering, storing, and retrieving tons of data puts a load on the various computing systems involved. It’s also not real exciting when one of your performance tool SQL statements shows up as the most resource intensive statements.

Total Performance Management (TPM)

I first published the Total Performance Management (TPM) [3] method in 1994 at an Oracle Applications User Group (OAUG) conference. It has since been one of the most frequently downloaded technical papers on OraPub’s website. It’s maintained its usefulness because it is a method about how to quickly transform an Oracle environment characterized by explosive surprises and poor performance, into an environment where users do not even think about performance . . . it is just there.

To quickly summarize the TPM method, it comprises three phases: the audit, tuning, and the proactive maintenance phase. The audit phase focuses on scoping out the situation and determining the overall approach to solving the problem. The tuning phase is an iterative process that over the course of many cycles can quickly eliminate performance problems allowing the DBA to break out of the performance-is-never-good-enough cycle. Once the DBA does break out of the tuning phase, and into the proactive maintenance phase, time is spent putting tools, processes, and communication pieces in place to thwart future performance threats.

Case Study: Performing Oracle Triage

Below is a relatively simple triage situation. However, the same methods and tools can be used in any size situation. I have written this section like a DBA might write an email or log as he or she makes notes during triage. I also followed the triage method outlined earlier in this paper. So here we go!

Don’t Panic

OK. I’m not panicking but I’m not dancing either . . . at least not dancing for joy. I realize that as a DBA my skills are in great demand so I’ll try to enjoy this escapade.

Scope Out the Situation

After walking around a bit, talking with a few people, and poking around on the system I’ve determined:

- The main system having problems is the company public website.

- This is an e-commerce company with no standard retail outlet.
- The marketing department is complaining the most.
- During peak hours, the website receives over 2,000 hits each minute.
- Each hit is recorded in a log file (actually an Oracle table).
- The system is so slow, ten Marketing Analysts were sent home yesterday.
- The Marketing Director is livid because he is putting together a new targeted marketing campaign and must run a number of reports to properly prepare and execute the campaign.
- The Marketing Director and the Information Systems director are good friends.
- The UNIX System Administrators feel the Oracle application is the cause of this problem.
- The Oracle DBAs feel the operating system is poorly configured thereby constraining Oracle.
- The initial bottleneck appears to be I/O.

After taking this all in, I've decided to form a Triage Team consisting of a SQL tuning, operating system, Oracle internals, and Marketing Department user expert. With these folks, I'm hoping to quickly isolate and confirm the bottleneck(s). After my analysis, I'm hoping this team will quickly reach consensus on our first-round tuning recommendations.

I'm also setting expectations that this could take a week to resolve, and even longer if the problem turns out to be extremely complex requiring a team of people to resolve, or an additional system capacity purchase.

Document the Current Performance

My Marketing Department user told me there is one key online marketing query which all the analysts run multiple times a day. This query provides an online summary of hourly website activity broken down by each website area's page type. The online report can be easily executed by a simple click. Usually the query returns in around ten seconds, but this past week it is taking around three minutes.

I had the Marketing Department user actually run the report for me on my computer three times. We timed the runs on our watches and they took 193, 177, and 210 seconds. That's an average of 197 seconds or 3.28 minutes.

I also heard that website activity has substantially increased. I checked with one of our UNIX System Administrators and he said that during peak we received around 2,000 hits each minute and around 720,000 hits per day. I'm going to verify this with the Oracle DBA because each web hit should be stored in the database.

Install Your Tools

I installed OraPub's OSM tools²—both the interactive and the historical tool kits. I modified the historical tool kit's driver script, rock, to enable the gathering of the Oracle "v\$" views and operating system cpu, memory, disk I/O, and network activity. I didn't have to write any scripts to gather website hit activity because it's already recorded in the web log. I did modify the key marketing report so it inserts a line into a log file every time the report is run. The line includes

After taking this all in, I've decided to form a Triage Team consisting of a SQL tuning, operating system, Oracle internals, and Marketing Department user expert.

both the start and stop time, so I can calculate response time. I'll use this to document triage success. I set the OSM tools to gather data every 30 minutes.

Develop a Simple Communication Strategy

To minimize triage distraction, I created one simple yet key communication graphic. By hour, it shows the number of website hits, the number of marketing queries run, and the average marketing query response time. This report will show if query response time correlates with website and marketing query activity . . . which I suspect it does. I am also creating a brief status email twice a day that I'll send to anyone who asks. The email consists of a two-sentence summary paragraph followed by more details if appropriate. I'm also going to host an open meeting every morning at 10 a.m. to anyone who is interested. I'm hoping with my frequent but brief communications, combined with an open attitude about the situation, trust will develop between the triage team and the suffering users.

Isolate the Problem

Since I just installed the OSM tool kit, I don't have enough historical data to review. So I ran some of the interactive reports to quickly isolate the bottleneck—either cpu, memory, I/O, or network. The triage chart in figure 1 outlines the various tools I used, organized by their areas of investigation, followed by my comments. After there was a reasonable amount of historical data to review, I went back and updated the triage chart. Ruminating on the data I gathered and my comments (as shown in figure 1), there is an obvious I/O bottleneck caused by both the marketing reports and the increased website activity.

Quickly Perform Your Rock-Solid Analysis

Summary and Recommendations. The computing system is seriously I/O bound as the result of increased website activity, intensive and repetitive marketing queries, and potentially un-optimized SQL. However, there appear to be a number of items, when combined with overall system performance, as well as marketing query response time, that should dramatically improve. While a more detailed analysis is below, here are my recommendations.

1. Move the HITS table into its own Oracle tablespace.
2. Move the HITS tablespace to a non-busy RAID 0+1 array.
3. Move TEMP tablespace database files to a non-busy

RAID 0+1 array.

4. Tune the marketing query SQL.
5. Reduce marketing query executions by having the query automatically run every 15 minutes and allowing the results to be instantly viewed via a standard web browser.
6. Increase Oracle's data block buffer cache to better absorb activity bursts.

Any one of these recommendations should substantially improve performance, but combined I anticipate a dramatic performance improvement.

Operating System Performance Analysis. There is a clear I/O bottleneck supported by CPUs waiting for I/O before processes can be run, no memory paging or swapping, no network collisions and latency problems, and four extremely active disks. The hot disks are part of the same RAID array and contain the web log (an Oracle table). There are many other disks which are less than 10% busy, so there is an opportunity to spread out the I/O across more disks or simply place the heavy I/O files on a different RAID array. I'll be talking with my O/S expert and administrator about this.

Oracle System Performance Analysis. Oracle sessions are waiting predominately for I/O because of full-table scans on the HITS tables. In fact, the MKTG tablespace and related database files, where the HITS table resides, are significantly more active than the other database files. The MKTG tablespace is also heavily full-table scanned, supporting my wait event analysis. Because each web page is logged, that is, written to the HITS table, when website activity peaks there is so much insert activity, the server processes have to wait to get a free block in the buffer cache. This is exacerbated by the MKTG reports, which are also looking for free blocks to fill with HITS blocks used for the marketing queries.

My initial approach to improve performance in this area is to increase the database block buffer cache and to isolate the HITS table into its own tablespace which will reside on a very non-active RAID array. Since there is plenty of memory, I will increase the database block buffer cache to better absorb bursty buffer activity thereby increasing the possibility of more free buffers. This is a short-term solution.

Other key contention possibilities like latching contention and enqueue contention are not an issue at this point. Once we deal with the immediate problems, they could raise their ugly heads.

Another interesting point is because of the MKTG query related full-table scans, there is a tremendous amount of temporary tablespace activity. However, I suspect that this will be substantially reduced once the MKTG query is tuned. But before that happens, I'm going to move the TEMP tablespace data files to a very non-busy RAID array.

Application System Performance Analysis. The MKTG and the web server processes are by far the most resource consuming processes. I looked at the SQL for both processes. The SQL is fairly straightforward and I suspect our SQL tuning expert can significantly reduce the number of I/Os the MKTG query requires (I hope so anyway!).

I also can't figure out why so many marketing queries are being run. It seems silly and a waste of resources. I'm going to talk with the Marketing Director about setting up some process where the key marketing query will be run every 15 minutes and write to a text file where any web browser can instantly look at it.

Doing a little math: During the peak bottleneck time (around 2 p.m.), there are 12 marketing queries being run simultaneously. They each touch around 26,000 Oracle blocks and take an average of 197 seconds to run. So on average, within 197 seconds there are 312,000 Oracle blocks touched. Our database consists of 8KB blocks, so this means 2,496,000 KB of data is touched every 197 seconds. The data block buffer cache hit ratio during peak is around 90% and the UNIX buffer cache ratio is also around 90% resulting in only 24,960 KB of data being physically read every 197 seconds. Doing a little more math, this means that only 126.7 KB/sec is physically read from oxide. Even with the associated insert activity, any descent RAID 0+1 array (which we have) should be able to handle the load. Moving the HITS table to its own non-busy RAID array should dramatically improve I/O response times allowing some time for the SQL to be tuned!

Once the "15-minute marketing query" process is in place, the HITS table is moved to a non-busy RAID 0+1 array, and the query is tuned, I anticipate a dramatic response time improvement.

Quickly Implement Your Recommendations

I met with both the Marketing Director and the lead UNIX System Administrator. The meetings were very profitable. The Marketing Director agreed to my "15-minute marketing query" idea and the UNIX Administrator agreed to provide virtually idle RAID 0+1 arrays for the HITS and TEMP tablespaces. The SQL tuning expert is already working on the marketing query. The lowest system activity occurs between 2 a.m. and 6 a.m., so that's my window of opportunity.

Here's the plan:

The SQL tuning expert is already working on the marketing query. As soon as it is optimized, it can be placed easily into production. We should experience absolutely no associated downtime.

Before 2 a.m.:

Because of advanced volume management tools, the UNIX Administrator will move existing files around resulting in the creation of two idle RAID 0+1 arrays. There should be absolutely no associated downtime.

Once the UNIX Administrator has one of the RAID arrays ready, I'll create a new temporary tablespace, alter all Oracle users to point to the new temporary tablespace, then remove the old temporary tablespace. There should be absolutely no associated downtime.

I created a script to quickly move the HITS tablespace and its associated database files to the second non-busy RAID array. There should be absolutely no associated downtime.

At 2 a.m.:

I will run the script to move the HITS tablespace and all

Figure 1. Oracle Triage Chart. This is an example chart that can be used as a checklist to triage any Oracle-based system. While additional areas can and will probably need to be investigated, the chart will bring out those additional areas. This specific chart is being used for the discussed case study.

Sub System	Approach	Component	Tool(s)	Comments	
UNIX O/S	Interactive	CPU	\$sar-u, \$mpstat	User=30 system=10 wio=60 idle =0	
		Memory	\$sar-g, \$sar-w	No significant paging or swapping	
		I/O	\$sar-u, \$sar-d	Only four disks very busy (80%) w/avg service time of 50ms. Devices c18t[1,3,4,6]s6	
		Network	\$netstat-i, \$ping	Bandwidth: <2% collisions, Latency: pings <5ms from DB server to Web server and Web server to users	
		Top processes	\$top	Oracle server processes (MKTG report!)	
	Historical	CPU	@cpu (OSM-H)	During peak wio=50-65 with idle=0-5	
		Memory	@mem (OSM-H)	No significant paging or swapping	
		I/O	@diskl (OSM-H), @iosum (OSM-H)	Only four disks very busy (80%) w/avg service time of 50ms. Devices c18t[1,3,4,6]s6.	
		Network	@netl (OSM-H) @pingl (OSM-H)	Same as interactive analysis	
		Oracle	Interactive	System wide waits	@swpct (OSM-I), swpctx (OSM-I)
		Session level waits	@swwc (OSM-I), @swwp (OSM-I), @swenq (OSM-I)	Most sessions were waiting for I/O, especially for disks with marketing data. MKTG users had heavy waits for FTS reads.	
		Busiest TBS	@dfio (OSM-I)	TBSs are very active, but the MKTG and the TEMP TBS are the most active.	
		Busiest DB files	@dfio (OSM-I)	The /u67/oradata/prod6/mktg08.dbf and the /u18/oradata/prod6/temp01.dbf are the really hot files. The mktg dbf gets 80% FTS'd.	
		Latching	@latch (OSM-I)	There most sig latching is library cache, but it's not a big deal at this point.	
	Historical	System wide waits	@swpctl (OSM-H), @swsys (OSM-H)	Just like interactive. Scattered reads and db file parallel write waits.	
		Session level waits	@swl (OSM-H)	Caught a few scattered read waits.	
		Busiest TBS	@tbsl (OSM-H)	TEMP and MKTG are by far the busiest. TEMP is 30% of all reads and MKTG is 60% of all reads.	
		Busiest DB files	@filel (OSM-H)	Same as the interactive.	
		Latching	@latchl (OSM-H)	Top latches are shared pool, library cache, and cache buffer chains, but sessions are not waiting for these, so not a big deal now.	
Application	Interactive	Top SQL	@sqls1 (OSM-I), @sqls2 (OSM-I)	Most resource consuming SQL address is 21889D24. It FTS's the MKTG HITS table! The next one is 4A44529C. It inserts into the HITS table.	
		Top Oracle users	@tp (OSM-I)	The top users fluctuates a lot. The top are the web user server processes following by the marketing user's server processes.	
		Number of sessions	@cpu (OSM-H)	There are 210 sessions. Oracle=9, Mktg=8, other=193	
			Number of concurrent MKTG rpts	\$mktgconcur (custom)	When I looked (many times) there were consistently 4 reports being run.
	Historical	Top SQL	@topsql (OSM-H), @sqlrepo (OSM-H)	Just like interactive.	
		Number of sessions	@sessl (OSM-H)	The number fluctuates quite a bit. From a min of 25 to a max of 250. During peak activity, there are around 200 sessions.	
		Max number of MKTG reports/hour	@mkstats (custom)	12. This happened around 2 p.m.	
		Avg number of MKTG reports/hour	@mkstats (custom)	Only 2, but this is because during the night, no reports are run.	
		Response time of MKTG reports	@mkstats (custom)	Avg: 197 sec., stddev 20 sec.	
		Max number of web hits/hour	@ws (custom)	Because this is an international site, the max and average are not too different. The max is 523.	
		Avg number web hits/hour	@ws (custom)	The average is 502.	

its associated database files. I anticipate five minutes of downtime. If the move fails, the system will be back up within ten minutes and a worst case scenario of sixty minutes of downtime (database file level point-in-time recovery).

Next day:

I will work with my application specialists to begin designing the “15-minute marketing query” system. Because of the concept and integration simplicity, I suspect this project will take three days to complete.

Monitor the Situation

All my “2 a.m.” recommendations have gone into effect and I’m working with the applications specialists to complete the “15-minute marketing query” solution. This is going very well and should be implemented in a couple of days. While the SQL tuning expert has improved query response time (50% logical I/O reduction), it is not what I’m expecting. I’m hoping we get, at a minimum, a 90% logical I/O reduction.

I performed an interactive and historical performance analysis just as I did in the Isolate The Problem triage phase. I won’t show all the details here, but the situation has substantially improved and the bottleneck has shifted (as I expected it would). The bottleneck is still I/O, but it is now

associated with the e-commerce application itself, not the activity logging or the associated marketing component.

I had the Marketing Director come over to my computer and we re-ran the “3-minute” query again and it took an average of 57 seconds. He was pleased but expecting sub-second response. I gently explained the SQL is still being tuned and the “15-minute marketing query” solution has not yet been implemented. And when either one of these items is implemented, I am expecting a less than 3-second response time. He was pleased with that.

So at this point the heat is off! I’m still working with the SQL tuning expert and the applications expert, but the pressure has substantially subsided, business is booming, and people don’t think about performance so much anymore. Time to get some sleep . . . zzzz.

Concluding Thoughts

Oracle performance triage is an exciting place to be, but no one can live there all the time. I’m hoping through my research, experience, and conversations, that I have been able to accurately convey how to appropriately apply triage. How successful I am will be determined when you are faced with a triage situation and have successfully applied the concepts, methods, and tools described in this paper. If this paper has been useful to you, please let me know. Feel free to email me and let me know how this paper has helped you or any other comments you may have to improve its usefulness.

References

1. “Advanced Performance Management For Oracle Based Systems” Class Notes (2001). OraPub, Inc., <http://www.orapub.com>
2. “OraPub System Monitor (OSM)” tool kit (2001). OraPub, Inc., <http://www.orapub.com>
3. Shallahamer, Craig A. (1995). Total Performance Management. Published and presented at various Oracle-related conferences worldwide. <http://www.orapub.com>
4. Shallahamer, Craig A. (2000). Holistic Problem Isolation Method. *OraPub Internet Video Seminar*. <http://www.orapub.com>
5. Shallahamer, Craig A. (1999). Direct Contention Identification Using Oracle’s Session Wait Views. Published and presented at various Oracle-related conferences worldwide. <http://www.orapub.com>
6. Shallahamer, Craig A. (1999). Direct Contention Identification Using Oracle’s Session Wait Views. *OraPub Internet Video Seminar*. <http://www.orapub.com>

Mr. Shallahamer’s seventeen-plus years of experience in the IT marketplace brings a unique balance of controlled creativity to any person, team, or classroom. As the President of OraPub, Inc., his objective is to empower Oracle performance specialists and capacity planners. Mr. Shallahamer can be contacted by email at craig@orapub.com.

This article is excerpted and reprinted with permission from the IOUG-A Live! Conference, April 2001.

¹ I don’t think incrementation is a real word, but it seems to fit nicely here.



RESOURCE CORNER

Oracle Resources

If you’ve got questions about Oracle, you might want to check out a handy resource from Oracle Corporation. In addition to MetaLink and OTN, there is Ask Tom at <http://asktom.oracle.com>. The “Tom” at Ask Tom is Thomas Kyte, author of *Expert One on One: Oracle*. He’s been answering questions about Oracle for a long time. You can ask a question yourself, or search through a long list of archives on various topics.

On his website, Tom says that he takes between 20 and 40 questions on a typical day. He answers most questions during the day but says, “If I have time and no one will play pool with me, I’ll do another batch at night.”

In searching through frequently asked questions, you’ll find the following categories:

- 9i specific
- DBA
- Developer
- General Database
- Java/JDBC
- PL/SQL
- Pro*C
- SQL

Thanks to NoCOUG member Ali Pankey for providing this tip for the Resource Corner. ▲

Encryption of Data at Rest

By Aaron C. Newman, CTO/Founder, Application Security, Inc.

Introduction

The digital age is upon us. No longer is instant access to information a request—it is now a requirement. Serving as a backbone for instant access is the relational database management system. Databases serve as the warehouses of digital information and hold our most critical assets. As such, to properly maintain the integrity and confidentiality of this data, the need for securing databases is growing. One of the requirements for securing databases is to encrypt the information stored within them.

Unfortunately, there are many misconceptions surrounding what database encryption is and how it should be performed. Encryption is a complex subject and properly implementing it requires a grasp of not only the theories behind encryption, but also the practical applications in the real world. All too often, the line between access control and encryption is blurred, and encryption solutions simply supplement the access controls already in place. What we hope to outline in this paper is an appropriate use of encryption as well as its proper implementation.

“Defense In-Depth”

No single security solution can properly protect a system. What is most important is “defense in-depth.” This means that more than a single layer of security is required in order to adequately protect a system. A good example of “defense in-depth” is a castle. The castle contains multiple defense systems—a moat, castle walls, archers on the walls, etc. Individually, each defense system would not be able to deter an attacker, but combined, the castle becomes very difficult to penetrate.

Encryption is one of the layers of security needed to secure your database. Of course, without implementing other security measures first, encryption is an inefficient and ineffective solution. Attempting to encrypt data that is not “locked down” utilizing the proper access controls leads to poor performance and poor security. For example, on your laptop you may encrypt your PGP private key using a strong password. If you did so, would it feel appropriate to email the encrypted key or provide public access to the key? Certainly not!

Access Controls

Now before even considering the implementation of encryption, you need to ensure that proper access controls are in place. Setting up access controls require the configuration of users and the actions they should be able to perform within the database. Within a database, access control consists of creating users, and granting them the privilege to act on objects together with performing certain commands and tasks. The built-in controls and mechanisms within the database are your best means of providing access controls.

Encryption

Once you have access controls in place, encryption should then be implemented. Encryption provides an additional restriction if access controls are circumvented. In other words, encryption should stop someone who has already broken through the “first line of defense.” Even when the hacker has broken through the first door—ACCESS CONTROL—encryption forms the next barrier of entry.

To demonstrate an appropriate method of using encryption on an operating system flat file, let us take a look at Microsoft’s Encrypted File System (EFS) within a Windows environment. In our example, we will walk through securing a document that contains secrets critical to your organization’s success stored on a file server.

The first thing you should do is set NTFS permissions on the file to prevent unauthorized users from reading the file. This is access control and should always be the first line of defense. However, there are several weaknesses with access control:

- NTFS permissions do not prevent system administrators from accessing the files
- If an attacker gains control of the operating system, they can use system administrator privileges to read the data
- An attacker can bypass the permission checking by booting the server to a different operating system to get around any access controls implemented by the original operating system

For example, let us say an attacker is aware of a buffer overflow in Microsoft Windows 2000 that allows him/her to run shell commands on the server. The attacker can then reset the Administrator password on the file server. Even though the appropriate permissions are established, they can be evaded by taking control of the operating system using the buffer overflow. Now the attacker has the ability to read the files on your file server using the Administrator account.

Is this type of attack avoidable? First off, your system should be patched well enough to withstand most buffer overflows. However, we live in a world where buffer overflows and other attacks are discovered on a daily basis. There is little chance you can “guarantee” invulnerability to this type of attack.



Aaron C. Newman

Encryption offers a reasonable, although far from perfect, solution. One way to protect your data even if an attacker gains full control over the data is by encrypting it. EFS can encrypt the file based on the user's password. By establishing a password as a key to the encryption, we have prevented an attacker with control of the operating system from reading the file. Two things to note about encryption are the following:

1. Encryption does not protect data from being deleted
2. Encryption does not protect data from being modified, although it does provide a way to tell if an unauthorized change has been made

Keep in mind the capabilities of encryption and its purpose. It is important that you maintain the proper backups so that if someone deletes or changes your encrypted data, you can restore the data.

Encryption of Data-in-Motion

Encryption can be categorized into two types: encryption of "data-at-rest" and encryption of "data-in-motion." The problems and approach of each type of encryption is very different. This paper addresses the issue of "data-at-rest." However, we should touch briefly on encryption of "data-in-motion."

Encryption of "data-in-motion" hides information as it moves across the network from the database to the client or from the client to the database. Data-in-motion includes traffic moving over your local network, the Internet, or

even over a wireless network. The various standards for this type of encryption include SSL (Secure Sockets Layer), TLS (Transport Layer Security), and IPSEC (Secure Internet Protocol). Most database vendors have adopted the SSL standard, and include the ability to send traffic between the client and database vendor over an SSL tunnel using some combination of RSA, RC4, DES, or Diffie-Hellman algorithm.

Encryption of "data-in-motion" is necessary to prevent someone from intercepting traffic as it goes back and forth from the client and the database. Encryption of data-in-motion is also effective at preventing such attacks as session hijacking and replay attacks.

Encryption of data-in-motion is typically implemented at a session level—the network layer above the protocol being encrypted. Network communications are encrypted as they are being transmitted over the wire and decrypted as they are received at the other end. Each command sent by the client is encrypted as it is sent and decrypted as it is received by the database. Each result is returned from the database, and is encrypted as it is sent and decrypted as the client receives it.

Encryption of "Data-At-Rest"

Encryption of "data-at-rest" is the encryption of information stored in the database. Encrypting "data-in-motion" does nothing to protect data that is attacked at the end points. Consider that most attacks do not occur on data-in-motion. Most attacks occur against the end points of data, where data sits for long periods of time. This leaves us in an interesting situation because encryption of data-in-motion is already widely adopted. Even the most "security-conscious" database administrators have not adopted encryption of data-at-rest.

How Do We Encrypt Data-At-Rest?

There are several possible strategies to encrypt "data-at-rest," and each strategy has certain advantages and disadvantages. The rest of this paper will outline the strategies and propose the best solution.

Encryption of data-at-rest can be performed in several ways. One way is to encrypt the actual database files at the operating system level. An example of using this strategy is to encrypt an entire database file using Microsoft's EFS within a Windows environment.

There are many weaknesses to using this strategy. Four of them are the following:

- 1) You cannot selectively encrypt individual pieces of data. This approach results in encrypting the entire file, which means all the data is encrypted. This causes serious performance problems for reading from the database. Every time data is read from the database, it is encrypted regardless of whether or not the data really needs to be secured. This adds significant overhead to any action performed against the database.

- 2) Encrypting the entire file not only adds the overhead of reading all data, but also leads to other additional overhead when recording pointers, indexes, and other internal data structures that must be encrypted and decrypted for

db Doctor
Emergency Database & OS Support

Health Watch for Oracle

**AUTOMATED, PROACTIVE, SYSTEMIC
DATABASE REPORTING**

**Save time and
avoid pain with**

- **80+ Automatic Reports**
- **90 Day Performance History Archive**

visit us on the web, or
call for a free trial

<http://www.dbDoctor.net> (866) 323-6286

any operation against the database. Ideally, when an insert is made to a database, the only encryption required should be the encryption of the data being inserted. Using file-based encryption, the information to determine where in the file to store the new record, the index, and many other internal file structures must be decrypted in addition to the data being inserted.

3) Another weakness is that different pieces of data cannot be encrypted with different keys. Imagine if you had a database that contained both sales and personnel information. The human resources department should have access to the personnel data but not the sales data. The sales department should have access to the sales data but not the personnel data. Using file-level encryption, this cannot be achieved because operating system file encryption encrypts the entire file, not sections of the file.

4) File-based encryption only protects the data from operating system-level attacks. If an operating system user copies the physical database file, the information is secured from that user. It does not protect the data from a user who breaches the database. When someone breaks into the database, it will gladly decrypt the information for the database user. This is because it appears to be a properly authenticated user, thereby circumventing the encryption.

A more efficient and effective way to encrypt information in a database is to perform the encryption on a column and row basis. To further explain this concept, think of a table containing a list of customers. Within this customer table, there is the following information:

- Customer ID
- Customer name
- Customer address
- Customer credit card number

In this table there is little reason to encrypt the customer ID. It is most likely that you would only want the credit card information encrypted. You gain several advantages by only encrypting your most sensitive data, which in this case is credit card data. One advantage is that you can minimize the performance hit incurred by only encrypting sensitive information. For instance, when a user attempts to search the table for a specific user, they incur very minimal overhead because only the data that must be decrypted is the data found row—a small subset of the data. Even better is the fact that when you select from other tables, which do not require encryption, there is absolutely no additional overhead added.

One of the serious problems encryption solves is protecting data from being read by administrators. This is accomplished by encrypting data utilizing a secret not known by the database administrator. Of course, the most important part of this statement is that the encryption must be dependent on restricting the administrator from discovering this secret, and utilizing it to decrypt the information within the database. For instance, if the administrator can simply reset the password of an account, logon to the account, and access the data, encryption has failed to protect the data from administrators. Encryption should be based on a secret such as a password. Therefore, if encryp-

tion were implemented utilizing a password as a secret, the database administrator could not just simply reset the password of an account to decrypt the data.

Of course, this means that when the user needs to change his or her password, this also involves resetting the decryption keys. Investigating this statement a little closer, the encryption system that we are referring to here would utilize a single key to encrypt and decrypt data in each column. A copy of this key, called the column key, is then stored encrypted with the user's password. To further illustrate this concept:

When I decide to change my password, I must first login with my current password, decrypt each column key, and then re-encrypt each column key with my new password replacing the old copy of the key. Next time I login with my new password, I can decrypt the keys with the new password. If an administrator simply changes the database password and logs in, he/she will decrypt the key to the wrong values because they are using the wrong password, and will be unable to decrypt the values in the table.

Using this technique, encryption is truly dependent on a secret, providing you a way to store data within your database that even an administrator cannot view.

Minimizing Performance Problems

One of the most important decisions you will make when utilizing encryption within your database is when you answer the following question: "What data should I encrypt?" Database lookups are designed to be very efficient. Unlike typical file systems, databases are expected to look through millions of rows searching for specific items in seconds. This need for fast access and retrieval places additional hardships on encrypting databases. A database cannot afford to encrypt and decrypt each piece of data it must search. Therefore, it is critical to properly plan encryption based on how an application will use the database. For example, let us imagine that we have a table with five columns and one million rows. The table contains customer information with the following columns: CustomerID, CustomerName, CustomerAddress, SalesRegion, and CreditCardNumber. Here are some alternatives with the respective performance pros and cons for each implementation:

Encrypting All Columns in a Table

What if you encrypt all five columns? If you select from the table for a specific CustomerID, you will be forced to decrypt the CustomerID of all 1 million rows. This will result in a huge amount of overhead. When you insert into the table, the overhead is not substantial, however, if you update the column based on the CustomerID, you will again be forced to decrypt the CustomerID for every single row.

Encrypting All But One Column in a Table

What if instead, you encrypt all columns except the CustomerID? When you select from the table for a specific

customer ID, you are not required to perform any decryption until you find the actual row that matches the selection criteria. This is because the engine in the database will only look at the column you “selected from” for all rows. Only when it finds a row that meets the criteria you indicated would decryption need to occur.

The additional time required to select from the table in this query has been reduced to almost nothing. Decryption of the rows found will cause some overhead, however, this overhead is minimal since the row set is relatively smaller.

In the same case, if I selected from the table looking for a specific CustomerName, I would again be forced to decrypt all 1 million rows, resulting in a very slow query. This is because the database engine will need to decrypt the customer name for every row in order to find the rows that match the search criteria. The same would apply to Updates and Deletes. Now if I selected from the table based on the CustomerID and the CustomerName, the search would be substantially faster because decryption would only occur for rows that matched the CustomerID. Of course, if the database’s query processor decided to search based on CustomerName before CustomerID, the results would be entirely different. When the query processor decides to perform the lookup on CustomerName first, the database is forced to decrypt CustomerName for every row.

Encrypting Only the Credit Card Number

If I encrypted only the credit card numbers, I would substantially reduce the chance that any query might significantly slow down the database. The only significant performance hit would be if the query processor decided that searching on the CreditCardNumber column would be the most efficient search. This would only occur if you were to search for a specific credit card number without providing any other search criteria.

Before actually deciding which columns to encrypt, you should first gather a list of the most common statements executed against the database. Most large applications are highly dependent on a handful of queries. Analyzing the use and frequency of SQL statements allows you to make an informed decision on how encrypting a column will affect performance.

Key Management for Database Encryption

One of the main problems with encryption is key management. Very often we see encryption implemented by hard-coding a password into a procedure or script. Even if you use the most robust encryption algorithms and the strongest keys, this implementation of database encryption is inherently flawed.

The problem is reduced to the fact that you are relying on access control to implement security. The idea behind “defense in-depth” should not rely on one layer of security to protect another layer. In this case, once access control is circumvented, encryption is also immediately circumvented. This does little to improve the security posture of your database.

Encryption should be layered on top of access control.

Encryption should protect the data when access controls are circumvented. When an operating system user is able to gain full control of the database, they can then circumvent the access controls. Encryption, if implemented properly, uses mathematics to prevent someone with full control over the system from reading data.

Solutions

Deciding to build your own security system is not a task most people should endeavor to undertake. There are a myriad of details to deal with, such as padding or dealing with NULL values that result in complications, and leaves you open to attack if implemented incorrectly. Your best bet is to find a system that provides all the features you need. While encryption will indubitably require some additional administrative work, you should find a solution that minimizes this impact.

When evaluating the possible solution, you should carefully consider each system. It is most important to understand how the system works and judge for yourself whether the solution provides “true encryption,” or does it break down when access controls break down. Talk to the vendor to understand the underlying architecture. If the vendor is not open about the underlying architecture, chances are they do not want you to know that something is wrong. An encryption system should not attempt to hide the implementation details by concealing how the code works. If it does, be wary.

About Application Security, Inc.

Application Security, Inc. (ASI) is an organization dedicated to the security, defense, and protection of one of the commonly overlooked areas of security—the application layer. Application Security, Inc. provides solutions to proactively secure (penetration testing/vulnerability assessment), actively defend/monitor (intrusion detection), and protect (encryption) your most critical database, groupware, and ERP applications. Application security-specific links, white papers, presentations, and FREE EVALUATION COPIES of DbEncrypt™ (database encryption) and AppDetective™ (penetration testing/vulnerability assessment) are available for download from the Application Security, Inc. website at: www.appsecinc.com

Application Security, Inc.

Web: www.AppSecInc.com

E-Mail: info@appsecinc.com

Tel: 1-866-9APPSEC • 1-212-490-6022



Many Thanks to Our Sponsors

NoCOUG would like to acknowledge and thank our generous sponsors for their contributions. Without this sponsorship, it would not be possible to present regular events while offering low-cost membership dues. If your company is able to offer sponsorship at any level, please contact NoCOUG President Joel Rosingana at joelros@pacbell.net.

Long-term full event sponsorship:

LOCKHEED MARTIN

CHEVRON

Long-term supplemental event sponsorship:

ORACLE

GENENTECH

Thank you! Year 2002 Gold Level Support Vendors:

- BMC Software
- Cast Software, Inc.
- Database Specialists, Inc.
- dbDoctor
- Embarcadero Technologies
- Informatica
- LECCOTECH
- Promatis
- Quest Software, Inc.
- XIOTech Corporation

For information about our Gold Level Vendor Program, contact Ganesh Sankar, Vendor Relations, at: bgs2k2@yahoo.com.



TREASURER'S REPORT

Judy Lyman, Treasurer

Beginning Balance		
December 31, 2001		\$ 45,486.30
Revenue		
<hr/>		
Advertising	960.00	
Sponsorships	-	
Membership Dues	21,774.35	
Meeting Fees	920.00	
Vendor Receipts	6,500.00	
Interest	44.04	
Vendor Mailing	-	
Total Revenue		\$ 30,198.39
Expenses		
<hr/>		
Regional Meeting	836.78	
Journal	5,541.15	
Membership	152.85	
Administration	940.00	
Website	477.13	
Board Meeting	182.00	
Miscellaneous	55.00	
IRS	1,718.00	
FTB Tax	739.00	
Total Expenses		\$ 10,641.91
Ending Balance		
March 31, 2002		\$ 65,4042.78

Hackproofing Oracle Application Server

The following focuses on the PL/SQL module of Oracle 9iAS. It is an excerpt from the paper: "Hackproofing Oracle Application Server: A Guide to Securing Oracle 9" by David Litchfield of Next Generation Security Software Ltd. in the United Kingdom. The full white paper can be downloaded from <http://www.nextgenss.com/papers/hpoas.pdf>.

PL/SQL

PL/SQL is Oracle's Procedural Language extension to Structured Query Language. PL/SQL packages are essentially stored procedures in the database. The package exposes procedures that can be called directly, but also has functions that are called internally from within another package. The PL/SQL module for Apache extends the functionality of a web server, enabling the web server to execute these stored PL/SQL packages in the database. The best way to imagine the PL/SQL module is like a gateway into an Oracle database server over the Web using stored procedures. By default all requests to the web server leading with /pls are sent to the PL/SQL module to be dispatched. The client request URI will contain the name of a Database Access Descriptor or DAD, the name of the PL/SQL package in the database server and the procedure being accessed. Any parameters that are to be passed to the procedure will be in the query string.

<http://oracleserver/pls/bookstore/books.search?cname=War+and+Peace>

The URL above has a DAD of "bookstore", a PL/SQL package called "books", a procedure called "search" which takes a parameter "cname", the name of the book to search for. The DAD describes a section in the wbsvr.app file that describes how Apache is to connect to the database server and contains details such as the UserID and password to authenticate with. If no credentials are supplied, the request would result in the web client being prompted for credentials. On connecting to the database server the database will load the books package and execute the search procedure and the search results would be passed back to the web server, which would then pass them on to the requesting client.

PL/SQL Buffer Overflows

The PL/SQL module contains several buffer overflow vulnerabilities. These can be exploited to run arbitrary code on the vulnerable web server. On Windows NT/2000 the Apache process is running in the security context of the local SYSTEM account so any code that is executed will run with full privileges. The first vulnerability occurs when a request is made for an administration help page. Even if the admin pages have been protected and require a user ID or password, this is not true of the help pages. To test if your site is vulnerable request

http://oracleserver/pls/dadname/admin_/help/AAAAA.....

Where AAAAAA . . . is an overly long string of around 1000 bytes. If the Apache process access violates or core

dumps, then the server is vulnerable and the patch should be applied from the Metalink site. If the patch can't be applied then, as a measure to help mitigate the risk of an attack, the default adminPath of /admin_/ should be changed to something difficult to guess or brute force. To do this, edit the wbsvr.app file found in the \$ORACLE_HOME\$Apache\modplsql\cfg directory.

Another overrun occurs when a similar request is made but this time without the DADname.

http://oracleserver/pls/admin_/help/AAAAA.....

This causes the Apache server to redirect the request to /pls/dadname/admin_/help/AAAAA . . . where "dadname" is the name of the default DAD. It is here the buffer overflow occurs. Again, the patch should be installed to address this problem and the default adminPath should be changed.

Another buffer overflow occurs when a request is made, by a client, presenting credentials to the web server using the "Authorization" HTTP header. An overly long password will cause the overflow. Any exploit code would be base 64 encoded so isn't easily recognizable.

For all of these buffer overrun vulnerabilities the patches should be downloaded and installed from the Metalink site.

PL/SQL Directory Traversal

The PL/SQL module can be abused to break out of the web root and access arbitrary files readable by the operating system account Apache is running under. To check if your site is vulnerable open

http://oracleserver/pls/dadname/admin_/help/..%255Cplspl.conf

This problem is due to the fact that the PL/SQL module has a double URL decoding problem and on the first pass converts %255C to %5C and on the second pass converts %5C to "\ " and the directory traversal becomes possible. To protect against this install the patch from the Metalink site.

PL/SQL Administration

By default it is possible to administer PL/SQL DADs remotely without needing to authenticate. This is obviously not a good thing. Whilst this doesn't allow an attacker an opportunity to run commands, they could attempt to change the user ID and password used to connect to the database server trying to boost privileges by using a default user login and password such as SYS, SYSTEM or CTXSYS. At the "best" they could deny service. Requesting

http://oracleserver/pls/dadname/admin_/

will show whether the site is vulnerable. If the administration page is returned then of course it is. To secure against this several steps are required. Firstly the wdsrvr.app file located in the \$ORACLE_HOME\$\Apache\modplsql\cfg directory should be edited. The adminPath entry should be modified to something difficult to guess or brute force. A password should be added.

PL/SQL Authorization Denial of Service

There exists a denial of service issue with the PL/SQL module. When a request is received by the module with a malformed Authorization HTTP client header with no authorization type set such as Basic Apache will access violate or core dump. The resolution to this is to install the patch provided by Oracle. This is available from the Metalink web site.

The OWA_UTIL PL/SQL Package

The OWA_UTIL package exists to provide web related services, along with packages like HTP, used for creating HTML content, and HTF which has functions that produce HTML tags. These and others are all installed as part of the PL/SQL Toolkit.

OWA_UTIL exposes many procedures that can be called directly from the web - this document will look at signature, showsource, cellsprint, listprint and show_query_columns.

`owa_util.signature`

Signature does nothing but simply returns a message - it can be used to verify access can be gained to owa_util. It doesn't require any parameters (though it can take some - but not needed)

`http://oracleserver/pls/dadname/owa_util.signature`

If a message is returned along the lines of

This page was produced by the PL/SQL Cartridge on December 21, 2001 04:50 AM

then access can be gained to owa_util.

If it doesn't return this and the web server returns a 500 or 403 response then it may be the package is protected. More often than not, depending upon how it has been protected, this protection can be bypassed by inserting a space, tab or new line character before:

`http://oracleserver/pls/dadname/%20owa_util.signature`
`http://oracleserver/pls/dadname/%0Aowa_util.signature`
`http://oracleserver/pls/dadname/%08owa_util.signature`

Regardless of how access is gained once it has been the other procedures can be called.

owa_util.showsource

Showsource will give the source code back of a package. It takes one parameter, "cname" - the name of the package to be viewed.

`http://oracleserver/pls/dadname/owa_util.showsource?cname=owa_util`

This will give the source code of the owa_util package.

`http://oracleserver/pls/dadname/owa_util.showsoucre?cname=books`

owa_util.cellsprint

Cellsprint allows the running of arbitrary SELECT SQL queries. It requires one parameter, "p_theQuery" but can also take a second "p_max_rows"—which specifies how many rows to return. If p_max_rows is not specified 100 rows are returned.

`http://oracleserver/pls/dadname/owa_util.cellsprint?p_theQuery=select * from sys.dba_users`

would return the first 100 rows from the dba_users table in the sys schema.

`http://oracleserver/pls/dadname/owa_util.cellsprint?p_theQuery=select * from sys.dba_users&p_max_rows=1000`

would return 1000 rows from the same table.

Of particular interest is the sys.link\$ table. This table contains a list of other database servers that connect to the one being queried. There are also clear text userids and passwords stored here - used when connections are made. If there are any connections then it is possible to "proxy" off of the first database server.

`http://oracleserver/pls/dadname/owa_util.cellsprint?p_theQuery=select * from sys.dba_users@other.world`

By specifying Other.world here, from the name column of sys.link\$ would cause the database server to look up the name of the other db server from the host column and then connect with the user id and password defined and retrieve sys.dba_users.

owa_util.listprint

Listprint is like cells print - arbitrary SQL queries can be run - with one difference - if select * is executed only one of the columns - the first - will be returned. Rather than selecting *, select a column name instead

`http://oracleserver/pls/dadname/owa_util.listprint?p_theQuery=select %20userna`
`me%20from%20sys.dba_users&p_cname=&p_nsize=`

This begs the question - how does one find the name of the columns in a given table?

`owa_util.show_query_columns`

To do this use the show_query_columns procedure. This takes one parameter - "ctable" - the name of the table.



http://oracleserver/pls/dadname/owa_util.show_query_columns?ctable=sys.dba_users

The HTML page returned will have a list of the column names.

As can be seen if an attacker can access the OWA_UTIL package they can almost peruse the database at will. It is therefore imperative to protect against access. OWA_UTIL is an example of one pack you'll want to prevent access to. Further to this all of the dbms_* packages should be protected, the htp packages, utl* packages and anything else you deem to be dangerous. To do this edit the wdbsvr.app file and add an entry in the exclusion list.

PL/SQL Authentication By-pass

In certain circumstances it may be possible to by-pass the authentication process when attempting to access a PL/SQL package. Imagine an on-line banking system that allows customers to register on-line. There will be a PL/SQL application to do this and will be given its own database access descriptor configured with a user ID and password thus allowing "anonymous" access to any who wish to register. Once registered the user would then access the actual banking PL/SQL application, which has its own DAD, too. This DAD has not been configured with a user ID and password though and as such when a user attempts to access the banking application they are prompted for credentials. The banking site URLs would be similar to the following

<http://oracleserver/pls/register/reg.signup>
<http://oracleserver/pls/banking/account.welcome>

As DADs are simply descriptions of how to access the database server the authentication of the banking app can be by-passed by simply substituting the DADs.

<http://oracleserver/pls/register/account.welcome>

By requesting this access is given because the user ID and password configured in the "register" DAD have authenticated the user to the database server. To protect against this there are two ways. Firstly the banking application should be created in a different database schema than the register application. The user ID used to gain access to the register application should not be able to access anything in the schema used for the banking application. The other way to prevent this kind of problem is to edit the wdbsvr.app file and add an entry to the "exclusion_list" entry.

`exclusion_list= account*, sys.*, dbms_*, owa*`

PL/SQL Cross-site scripting vulnerabilities.

By default access to the htp PL/SQL package is allowed. This package exports procedures for outputting HTML and HTML tags. Many of the procedures can be used in cross-site scripting attacks.

[http://oracleserver/pls/dadname/htp.print?cbuf=<script>alert\('Doh!'\)</script>](http://oracleserver/pls/dadname/htp.print?cbuf=<script>alert('Doh!')</script>)

Cross-site scripting attacks have been widely discussed before and pose a potential threat and, as such, access to the htp package should be disallowed by adding it as an exclusion entry to the wdbsvr.app file.



Need Part-Time DBA Support?

When you need no-nonsense, straightforward Oracle expertise, call on Database Specialists. We'll roll up our sleeves and help you get things done. With our DBA Pro service, we can configure a part-time or remote DBA program that best suits your needs. You receive:

- A cost-effective and flexible extension of your IT team
- Proactive database maintenance and quick resolution of problems
- Increased database performance and minimized database downtime
- Constant database monitoring with Database Rx™
- Onsite and offsite flexibility
- Reliable support from a stable team of DBAs familiar with your databases

Call Us Today!

415-344-0500 • 888-648-0500
www.dbspecialists.com

ORACLE | CERTIFIED SOLUTION PARTNER


Database Specialists

Dive into any database administration challenge. Without the risk.

DBArtisan. The fastest, easiest—and safest—way to manage multiple databases.

www.embarcadero.com/download/artisan.htm

 **Embarcadero Technologies**

do more now

©2000 Embarcadero Technologies, Inc. All rights reserved. Embarcadero Technologies and DBArtisan are registered trademarks of Embarcadero Technologies, Inc. All other products are the trademarks of their respective companies.

NoCOUG Spring Conference

Thursday, May 16, 2002
Session Descriptions

For up-to-date information, see www.nocoug.org

DBA TRACK

10:45-11:45

Oracle9i Database Release 2 Overview

Mark Townsend, *Director, Product Management*
Oracle Corporation

This session provides an overview of the new enhancements in Oracle9i Database Release 2. It discusses enhancements that increase and improve availability, data integration, application development, content management and manageability. It also covers new features in the areas of business intelligence, OLAP, and data mining. *(Intermediate Level)*

1:30-2:30

Online Object Redefinition in Oracle 9i

Chris Lawson, *Independent Consultant, Performance Solutions*, and Roger Schrag, *Senior Consultant, Database Specialists, Inc.*

With Oracle 9i's online object redefinition feature, applications can update data in the database while DBAs are relocating tables to new tablespaces or even adding new columns to tables. In this presentation, two experts on high availability databases explain how to use this new feature to maximum benefit. First they will discuss the `dbms_redefinition` package provided in Oracle 9i, and how you use this package to reorganize or redefine tables without downtime. They will cover the finer points like interim tables, preserving integrity constraints and triggers, and rolling back if an error occurs midway. Finally, we will walk through examples of how `dbms_redefinition` can be used to deploy a schema change without data unavailability. We will look at how well the feature performs and various gotchas to look out for. *(Intermediate Level)*

3:15-4:15

Storage Management and Oracle 9i

Dean Richards, *Senior System Engineer*,
Insession Technologies

The presentation will be a discussion of Oracle 9i's new features for storage management. Topics covered include external tables, unused index identification, multiple block sizes, dynamic memory configuration, auto segment space management, automatic undo management, resumable space allocation, online schema changes, and others. This presentation will discuss some issues in Oracle 9i storage management and present ways AutoDBA, from Insession Technologies, can assist. *(Beginner to Intermediate Level)*

DEVELOPER TRACK

10:45-11:45

Performance Tune the SQL

Scott Sumner, *Systems Consultant*
Quest Software

This presentation goes into considerable detail on the tuning of actual SQL statements. The attendee will understand the Oracle optimizer choices, the Explain Plan, and how to interpret the output of an Explain Plan. This presentation will also include a discussion on the differences of the various Explain Plan steps such as Merge-Join and Nested-Loop, when is it best to use each, etc. The presenter also discusses how indexes are accessed with the Cost-Based optimizer as well as some important index-oriented INIT.ORA parameters. This presentation is a MUST for anyone doing SQL statement level tuning. *(Intermediate Level)*

1:30-2:30

Develop for Optimum Application Performance

Claudia Fernandez, *Technical Services Manager*,
LECCOTECH

In this session, attendees will learn how to develop the highest performing applications prior to deploying in production. The specific methodology for developing for performance not just functionality will be presented. Developers will also learn SQL benchmarking methodology and techniques for delivering the highest performing applications the first time every time. *(All Levels)*

LECCO
SQL Expert
SQL with a Higher IQ

LECCO SQL Expert automates the task of performance tuning and optimization of database applications by providing developers and DBAs with expert knowledge of SQL. Through its proprietary Feedback Searching engine, it provides optimum SQL statements—guaranteed to provide the correct results. No guesswork or hopeful suggestions, statements are actually rewritten.

Join the ranks of the experts and download a FREE evaluation copy today!

- Oracle
- Sybase
- MS SQL Server

Where the experts go for performance.

LECCOTECH
www.leccotech.com

NoCOUG Spring Conference

Thursday, May 16, 2002

Location: Lockheed Martin in Sunnyvale. Check the NoCOUG website for directions. Be sure to arrive early to allow time for security. Please note that Lockheed Martin has special security requirements: All attendees must bring government-issued photo identification in order to enter the facility. Resident Aliens must also bring their green card, and Foreign Nationals on work visas must bring their passport.

- 8:00-9:30 **Registration and welcome. Refreshments served.**
- 9:30-10:15 **Opening Remarks and Announcements/Vendor Introduction**
- 10:15-10:45 **Morning Break**
- 10:45-11:45 **Parallel Sessions**
DBA Track: Oracle9i Database Release 2 Overview
Developer Track: Performance Tune the SQL
Data Warehouse: Taming the Customer in Oracle Dimensional Modeling - an Example from the BioTech World
- 11:45-12:45 **Lunch Break**
- 12:45-1:15 **Roundtable Discussions for DBA and Developer Tracks**
- 1:15-1:30 **Break**
- 1:30-2:30 **Parallel Sessions**
DBA Track: Online Object Redefinition in Oracle 9i
Developer Track: Develop for Optimum Application Performance
Data Warehouse: TBA
- 2:30-3:15 **Afternoon Break**
- 3:15-4:15 **Parallel Sessions**
DBA Track: Storage Management and Oracle 9i
Developer Track: TBA
Data Warehouse: TBA
- 4:30-?? **Networking and Happy Hour: Faz Restaurant and Bar, 1108 North Mathilda, Sunnyvale**
- Cost: **\$40 admission fee for non-members. Members free. Includes lunch voucher.**



View session descriptions at www.nocoug.org.

Thank you to Lockheed Martin, our meeting sponsor.

RSVP online at <http://www.nocoug.org/rsvp.htm>

NoCOUG
P.O. Box 3282
Danville, CA 94526

FIRST-CLASS MAIL
U.S. POSTAGE
PAID
SAN FRANCISCO, CA
PERMIT NO. 11882