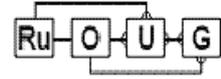
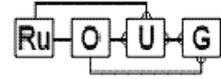


1. KEYNOTES	3
TOM KYTE	3
<i>What are we still doing wrong?</i>	3
TANEL PODER	3
<i>Back to Basics: Choosing the Starting Point of Performance Tuning and Troubleshooting Wisely</i>	3
MOGENS NØRGAARD	4
<i>Oracle's License Police, RAC for SQL Server, and counting correctly in 11g</i>	4
CANCELLED, UNFORTUNATELY: STEVEN FEUERSTEIN.....	4
<i>Coding Therapy for Software Developers aka "How does this code make you feel?"</i>	4
2. PRESENTATIONS.....	4
C.J. DATE	4
<i>The Closed World Assumption</i>	4
<i>Why Three-Valued Logic Doesn't Work</i>	5
<i>An Introduction to Logic</i>	5
MARK RITTMAN	6
<i>Oracle BI Enterprise Edition 11g New Features</i>	6
TOM KYTE	6
ALL ABOUT METADATA; WHY TELLING THE DATABASE ABOUT YOUR SCHEMA MATTERS – 1 HOUR	6
SUE HARPER.....	7
<i>Oracle SQLDeveloper Data Modeler</i>	7
JOŽE SENEGAČNIK	7
<i>Services – Why Should One Use Them?</i>	7
CANCELLED, UNFORTUNATELY: STEVEN FEUERSTEIN.....	7
<i>High Performance PL/SQL</i>	7
TANEL PODER	8
<i>Understanding LGWR, log file sync Waits and Commit Performance</i>	8
DANIEL MORGAN	8
<i>Edition-Based Redefinition: the Key to Online Application Upgrade</i>	8
MAREK LÄLL.....	9
<i>User Experience with Migrating Databases to Unicode</i>	9
PIET DE VISSER.....	9
<i>Good Indexing : Show CBO where to go</i>	9
RAIMONDS SIMANOVSKIS	10
<i>PL/SQL unit testing with Ruby</i>	10
FRANCISCO MUNOZ ALVAREZ	11
<i>Logging or Nologging: That is the question!</i>	11
DEBRA LILLEY.....	11
<i>Fusion Applications – what is behind it?</i>	11
ILYA DEEV	12
<i>SQL Tuning Advisor's work verification</i>	12
THOMAS ANTILA	12
<i>Variations (of the column order in a fat index) and how to decide</i>	12
PEKKA E LIUKKONEN.....	13
<i>Implementing archiving solution for Oracle EBS</i>	13
DMITRY VOLKOV, ORACLE	13
<i>Oracle Exadata Deep Dive</i>	13



RASMUS KNAPPE , ORACLE	13
<i>11g R2 New Features</i>	13
TAMMY BEDNAR, ORACLE	14
<i>A technical deep dive on both ASO network and at rest encryption.</i>	14
LAURI PIETARINEN	14
<i>Proactive Index Design using QUBE</i>	14
EVGENY GORBOKONENKO.....	14
<i>New strategy of Oracle RAC 11g R2 deploying.</i>	14
PETER ROBSON, JUSTSQL.....	14
<i>Reversing Committed Transactions in SQL *Plus.</i>	14
WOLFGANG SCHERRER.....	16
<i>How to design a Oracle Technology and Solution based IT Blueprint</i>	16
MARIS ELSINS.....	17
<i>Surviving the Crisis with the Help of Oracle Database Resource Manager</i>	17
3. SMALL ROOM.....	17
DANIEL MORGAN	17
<i>The Oracle ACE Program</i>	17
JOERGEN HOLMER-BRETLAU	17
<i>Development of Oracle Support interfaces</i>	17
SEBASTIAAN VINGERHOED, ORACLE	18
<i>Real Application Testing</i>	18
4. ROUNDTABLES/WORKSHOPS.....	18
ANTTI KOSKINEN	18
<i>(The Myth of a) Database-Independent Application - roundtable</i>	18
SUE HARPER, ORACLE	18
<i>Oracle SQLDeveloper Data Modeler</i>	18
TAMMY BEDNAR, ORACLE	18
<i>security hands-on lab</i>	18



1. Keynotes

Tom Kyte

Bio

Tom Kyte is a Senior Technical Architect in Oracle's Server Technology Division. Before starting at Oracle, Kyte worked as a systems integrator building large-scale, heterogeneous databases and applications, mostly for military and government customers. Kyte spends a great deal of time working with the Oracle database and, more specifically, working with people who are working with the Oracle database. In addition, Kyte is the Tom behind the AskTom column in Oracle Magazine, answering people's questions about the Oracle database and its tools (<http://asktom.oracle.com/>). Kyte is also the author of Expert Oracle Database Architecture (Apress, 2005), Expert One on One Oracle (Wrox Press, 2001/Apress 2004), Beginning Oracle Programming (Wrox press, 2002/Apress 2004), and Effective Oracle by Design (Oracle Press, 2003). These are books about the general use of the database and how to develop successful Oracle applications.

Photos for adverts

<http://flickr.com/photos/tkyte/sets/72157607376761768/>

What are we still doing wrong?

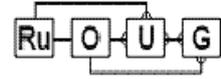
I've given many best practices presentations in the last 10 years. I've given many worst practices presentations in the last 10 years. I've seen some things change over the last ten years and many other things stay exactly the same. In this talk – we'll be taking a look at the good and the bad – what we do right and what we continue to do wrong over and over again. We'll look at why "Why" is probably the right initial answer to most any question. We'll look at how we get to "Know what we Know", and why that can be both a help and a hindrance. We'll peek at 'Best Practices' and tie them into what I term 'Worst Practices'. In short, a talk on the good and the bad.

On a somewhat more serious note, but still high level:

Tanel Poder

Back to Basics: Choosing the Starting Point of Performance Tuning and Troubleshooting Wisely

Choosing the right entry point to performance tuning and troubleshooting is the most important prerequisite for getting reliable and consistent results. It is the first decision in problem solving process, which will direct you in the right (or wrong) direction. In this session we will cover how to define and "normalize" the problem statement, quantify the problem and where to go on from there.



This is not a deep internals session. Tanel will instead speak about his 14-year journey from a "desperate switch flipping" DBA to a systematic problem troubleshooter, plus the most important lessons he has learnt during this time.

Mogens Nørgaard

[Oracle's License Police, RAC for SQL Server, and counting correctly in 11g](#)

The presentation will cover

- 1) RAC on SQL Server,
- 2) How to win - every time - over a multinational company,
- 3) The state of RAID-5, and
- 4) How to tell a good beer from a bad beer.

Cancelled, unfortunately: Steven Feuerstein

Coding Therapy for Software Developers aka "How does this code make you feel?"

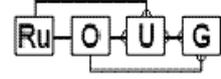
We can't write software without our brains, and our brains come with a full load of "issues." The way our brain remembers the past and projects into the future has a big impact on how we write code. Moving beyond physiology, human psychology also plays its role, making it difficult for us to acknowledge ignorance and ask for help. Steven will in this keynote address offer an intensive coding therapy session (including couples therapy, dream therapy and shock therapy) to help all attendees come to grips with their innate, unavoidable "issues", making it easier to write better code -- and help others on their team write better code.

2. Presentations

C.J. Date

[The Closed World Assumption](#)

The Closed World Assumption (CWA) is an extremely important concept in the database world, despite the fact that it isn't usually spelled out explicitly. Basically what it says is this: Everything stated by the database, either explicitly or implicitly, is true; everything else is false. This presentation explains the CWA in detail and shows why it's preferred over its rival, the Open World Assumption (OWA). In particular, it examines the claims that are sometimes heard to the effect that the database community operates under the CWA while the semantic web community operates under the OWA. It also explains how "missing information" can be handled without any need for nulls or three-valued logic.



Topics:

- Background review: what a database really is (?)
- Relation values vs. relation variables
- Relvar predicates and constraints
- Relations with no attributes
- The CWA and OWA defined
- Why the CWA is preferred
- Relvar predicates revisited
- Unknown and uncertain information
- Negation and disjunction

Prerequisites: Attendees will be expected to be familiar with basic relational database concepts. (However, the presentation will begin by briefly reviewing such matters.)

Why Three-Valued Logic Doesn't Work

This presentation is a kind of appendix to the presentation "An Introduction to Logic" (prior attendance at which is strongly advised). The logic on which the relational model is based is two-valued (2VL). In its attempt to deal with the so called "missing information" problem, however, SQL is based on a three-valued logic (3VL) instead. This presentation explains in detail why any such attempt is doomed to failure. More specifically, it shows why 3VL (a) doesn't solve the problem, (b) isn't useful, and (c) can actually be dangerous.

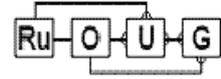
Topics:

- What is 3VL?
- Nulls and 3VL: the basic idea
- Effect on the connectives
- 3VL gives wrong answers!
- What about 4VL?
- Many-valued logics in general
- So what do we do?

Prerequisites: Attendees will be expected to be broadly familiar with the relational model (or at least with SQL).

An Introduction to Logic

Everyone knows the relational model is founded on logic, and moreover that it derives much of its strength, rigor, and robustness from that solid foundation. Few database professionals can claim to be familiar with that foundation, however, even though an elementary knowledge of logic is critical to successful use of a



relational DBMS. This presentation is offered as an attempt to rectify the situation: It explains certain key concepts from predicate logic and shows their direct relevance to a variety of database issues. A specific goal of the presentation is to enable attendees to use logic effectively in their day-to-day database activities.

Topics:

- The problem of ambiguity
- Propositions
- Connectives and truth tables
- Tautologies, contradictions, and inference
- Truth functional completeness
- Predicates and instantiation
- Quantification
- Free and bound variables

Prerequisites: Attendees will be expected to be broadly familiar with the relational model (or at least with SQL). No prior knowledge of logic will be assumed.

http://www.justsql.co.uk/chris_date/chris_date.htm

Mark Rittman

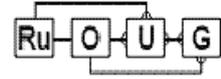
Oracle BI Enterprise Edition 11g New Features

"The 11g release of Oracle BI EE brings a number of new features including a redesigned user interface, an enhanced OLAP-style query tool, scorecarding and maps, the Action Framework for SOA integration and an enhanced metadata model. This presentation looks at the new features, demonstrates how they work and provides advice on which new features are the most valuable to apply in new business intelligence projects".

Tom Kyte

All about metadata; why telling the database about your schema matters – 1 hour

It is interesting to note how important metadata – data about data – is in the database. One important piece of metadata are database constraints. The database uses constraints for query optimization as well as data integrity. Most people understand that constraints are a 'data integrity feature', but are not as aware of the fact that the optimizer relies heavily on metadata to operate at its best. We'll explore the various types of metadata we as developers and DBA's can assert into the database and the positive impact this metadata will have on us.



Sue Harper

Oracle SQLDeveloper Data Modeler

Jože Senegačnik

Services – Why Should One Use Them? – 60 minutes

Services are most frequently used in conjunction with RAC. However, one can and should use services also in single-instance environments especially for managing and monitoring database workload and performance monitoring. This presentation will discuss all the required steps to create and manage services and utilize them for database monitoring and workload management.

Bio:

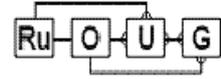
Jože Senegačnik has more than 20 years of experience in working with Oracle products. He began in 1988 with Oracle Database version 4 while working for the City of Ljubljana, where he had charge over the city's municipal and geographic information systems. From 1993 to 2003, he worked in developing GIS systems for the Surveying and Mapping Authority of the Republic of Slovenia, and in the development of applications for other governmental institutions, all based on the Oracle database. More recently, he has specialized in performance optimization, having developed his own toolset for monitoring performance and analyzing trace files.

Jože is an internationally recognized speaker, and a member of the highly respected OakTable Network (oaktable.net). He is a regular speaker at user-group conferences, especially those put on by the Slovenian Oracle Users Group (SIOUG), the Independent Oracle Users Group (IOUG), and the United Kingdom Oracle Users Group (UKOUG). He also speaks routinely at the Hotsos Symposium and Oracle Open World. In addition to sharing his knowledge through conference talks, Jože conducts technical seminars organized either by Oracle University or himself. He was awarded Oracle ACE membership for his long record of positive contributions to the Oracle community. Jože is also coauthor of the OakTable book titled "Expert Oracle Practices".

Cancelled, unfortunately: Steven Feuerstein

High Performance PL/SQL

Users hate to wait - for anything. For our applications to be successful, they not only must be correct (meet user requirements) and maintainable. They must also execute efficiently enough to avoid user frustration. This session reviews the most important techniques for improving PL/SQL performance, including data caching, FORALL and BULK COLLECT. After attending this session, you will be able to proactively identify



opportunities for applying techniques that will most dramatically (generally, an order or magnitude or more) improve the performance of your PL/SQL code.

Tanel Poder

[Understanding LGWR, log file sync Waits and Commit Performance](#)

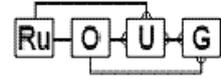
The commit's *log file sync* wait event is often misunderstood to be entirely an IO related wait. In reality there are more factors which affect your commit performance, such as the process scheduling latency and application commit rate.

This session walks you through some redo generation and writing internals, how commits happen and how the LGWR process fits into all this. After this session you will be able to easily determine whether your log writing wait events have anything to do with bad IO performance, CPU starvation or just bad application coding practices.

Daniel Morgan

[Edition-Based Redefinition: the Key to Online Application Upgrade](#)

Large, mission-critical applications built on Oracle Database are often unavailable for tens of hours while the application's database objects are patched or upgraded. Oracle Database 11g Release 2 introduces revolutionary new capabilities that allow online application upgrade with uninterrupted availability of the application. Existing sessions can continue to use the pre-upgrade application until their users decide to finish; and, at the same time, new sessions can use the post-upgrade application. When no sessions are any longer using the pre-upgrade application, it can be retired. The application as a whole therefore enjoys hot rollover from the pre-upgrade version to the post-upgrade version. The capability depends on these new kinds of object: the edition, the editioning view, and the crossedition trigger. Code changes are installed in the privacy of a new edition. Data changes are made safely by writing only to new columns or new tables not seen by the old edition. An editioning view exposes a different projection of a table into each edition to allow each to see just its own columns. A crossedition trigger propagates data changes made by the old edition into the new edition's columns, or (in hot-rollover) vice-versa. The capability as a whole is called edition-based redefinition – EBR for short. This session explains how it all works.



Marek Läll

User Experience with Migrating Databases to Unicode

Overview of previous configuration and motivation that forced us to move towards Unicode. Overview of Estonian, Latvian and Lithuanian specific characters. Typical mistakes in encoding related configuration (and in thinking). Init.ora parameters that change Unicode specific behavior. References to data dictionary views that help to understand differences between single-byte and Unicode database. "exp" and "imp" tool restrictions/issues related to Unicode. List of issues raised during migrations of different database versions. Tips and tricks used to solve those issues. Issues and restrictions that Unicode database have today. Conclusions, suggestions and Q&A.

SPEAKER BIOGRAPHICAL INFORMATION

Marek Läll (36) - Database Area Manager, IT Operations division, Swedbank Baltics (SBB). He has been working in SBB since 2000. He has experience with Oracle databases from version 7.0.16+, since 1994. He has passed Oracle8 OCP and upgrade exam to Oracle10. He is responsible for maintaining and developing infrastructure of 7x24 business critical Oracle databases in SBB. The last challenges have been physical movement of 7x24 business critical databases to centralized datacenters; performing upgrades from 8.x to 10.x and from 9.x to 11.1; implementing stretch RAC cluster 11.1 on intel / linux platform; implementing Grid Control 10.2.

DETAILS OF PREVIOUS SPEAKING EXPERIENCE

Oracle User Group Finland 2009 autumn seminar: "User experience with implementing 11g (11.1) RAC in Stretch cluster" He has instructed ~50 trainings related to Oracle Database technologies and administration. Oracle East Central Europe has granted to him "qualified to teach" certificates for 12 different courses related to SQL, PL/SQL and Oracle Database administration.

Piet de Visser

Good Indexing : Show CBO where to go.

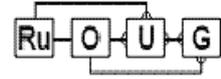
So how can we make this database thing work for us instead of the other way round? This presentation will generally look for the simplest, and the laziest road to to Efficient systems.

Common Joke: The Best way to fix inefficient queries is to not have them.

The next best way is to present them the most efficient path into the data, which is why I will cover some indexing basics items and a few lesser used techniques like IOT and Clusters. The presentation includes some demos which will only be done if time permits.

The 2nd half of the presentation covers basic CBO behavior, some hard-learned lessons and some additional tricks that should only be used in "emergencies". There are no silver bullets to CBO problems. But it wont hurt to look and learn. As long as we realize that we are using "tricks", and realize our risks and limitations. We might get lucky.

Level: Beginner and Intermediate.



Raimonds Simanovskis

PL/SQL unit testing with Ruby

Unit testing and TDD (test driven development) practices are nowadays one of the key software development practices. In some languages (like Java, Ruby, Python, C# etc.) there is quite good tools and frameworks support for unit testing and as a result there is quite high testing culture among top developers in these communities. But unfortunately in PL/SQL community so far automated unit testing is not used very often.

Why is it so and what are are current options for doing automated PL/SQL unit testing? The first framework was utPLSQL which was based on ideas of other xUnit frameworks but nowadays it is not maintained anymore. Then there are some visual GUI tools like Quest Code Tester (commercial tool) and new SQL Developer 2.1 includes unit testing support. But still majority of PL/SQL developers have not found them very attractive and easy to use as a result they are not used very frequently.

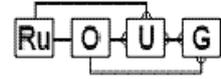
If in comparison we look at Ruby community then it has very high testing culture and has many good tools for testing support (e.g. RSpec testing framework). Therefore presenter suggests to use the approaches and tools that are working well in Ruby community and use them for PL/SQL unit testing as well. Presenter has created ruby-plsql library which provides very easy API for calling PL/SQL procedures from Ruby and presenter has created ruby-plsql-spec framework which combines ruby-plsql and RSpec libraries to provide framework for PL/SQL unit testing. ruby-plsql spec provides much shorter and readable syntax for unit tests compared to utPLSQL and it provides much more flexibility and better readability compared to GUI testing tools.

More information about PL/SQL unit testing with Ruby can be found at <http://blog.rayapps.com/2009/11/27/oracle-plsql-unit-testing-with-ruby/> Screencasts with ruby-plsql-spec demonstration can be found at <http://blog.rayapps.com/2010/01/06/screencasts-of-oracle-plsql-unit-testing-with-ruby/> Presentation will include overview of existing PL/SQL unit testing frameworks, motivation and description of ruby-plsql-spec framework as well as live demo of ruby-plsql-spec usage.

About speaker:

Raimonds Simanovskis is Technical Director at TietoEnator Alise (a Tieto corporation company). His primary role is technological development and strategy as well as participation and supervision of new technology projects. Raimonds has participated in many Oracle E-Business Suite implementation projects as well as Oracle based software development projects. Nowadays he is very interested in Ruby language and Rails framework and is using Ruby on Rails together with Oracle database and applications. He also uses and promotes Agile software development practices.

He is also Oracle ACE and blogs about Oracle, Ruby and Mac at <http://blog.rayapps.com>. Raimonds received Oracle Magazine's Developer of the Year award in 2009 (<http://www.oracle.com/technology/oramag/oracle/09-nov/o69awards.html#simanovskis>) for his open source contributions and community support in Ruby and Oracle interoperability area.



Francisco Munoz Alvarez

[Logging or Nologging: That is the question!](#)

The main question about NOLOGGING I hear all the time is: does creating a table with the NOLOGGING option means there is “no generation of redo ever”, or just that the initial creation operation has no redo generation, but that DML down the road generates redo? How and when can the NOLOGGING option be employed?

Redo generation is a vital part of the Oracle recovery mechanism. Without it, an instance will not recover when it crashes and will not start in a consistent state. Excessive redo generation is the result of excessive work on the database. This presentation covers the subject of reducing redo generation using LOGGING and NOLOGGING options, the differences between them, how it happens, how to reduce it and when to use. Also, you will find examples and tips regarding each one of them.

Francisco Munoz Alvarez (Oracle ACE Director) is the founder and CEO of DBIS - Database Integrated Solutions and has worked with Oracle since 1989. He is also the President of CLOUG (Chilean Oracle Users Group), LAOUC (Latin American Oracle Users Groups Council), and NZOUG (New Zealand Oracle Users Group). Francisco worked as associate technologist at Oracle Brazil and Chile, was an Oracle instructor for New Horizons Center, Chile and for Oracle education (Brazil and Chile), and also worked on the first team to introduce Oracle to South America (Oracle 6 and Beta version of Oracle 7). He was also the first Master Oracle 7 Database Administrator in South America. Oracle Certification Levels: OCA (SQL & PL/SQL, 10g AS), OCP (10g/11g DBA, E-Business 11i) and 10g RAC OCE.

Debra Lilley

[Fusion Applications – what is behind it?](#)

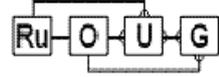
Introduction

I have talked for a while about Fusion Application concepts. What blew me away was the vision, all of the technology was there and I could identify with it all, so if it was available could I build my own Fusion Applications? And yes if I had the resources, the investment and above all the vision I could.

In the past 12 months Oracle has moved closer to the release of the suite with a promised release date of 2010. At Oracle Open World Larry Ellison showcased the product and immediately afterwards my NDA was lifted. So I can now talk more about the product but actually it is the technology that this paper concentrates on.

Objective

My objectives remain the same. I want is for the audience to think about what Oracle is doing, how it fits together and what they need to be planning and building skills around Fusion, not just if they are intending to use Oracle applications



Initially the term 'Fusion' was ignored by the technical community it was just about applications, to get the message out that the applications would be about the technology Oracle renamed their middleware Fusion MiddleWare. The tools are available today and many of you are using them. My concern is that as the business community start to see the Fusion Applications they will not appreciate that behind the great 'user interface' there is a great deal of technology that needs to be maintained and exploited for business benefit.

If you understand today, how Fusion Applications will work, how the tools are here now and think about how you can improve the investment of your organization you can build the business case to invest in those skills today and be ahead of the game.

Ilya Deev

SQL Tuning Advisor's work verification

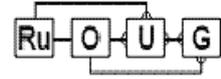
I think SQL Tuning Advisor is an amazing tool, but unfortunately it does not always work properly. How could we check its work and how could we automate our verification work, especially with Oracle 10g? The answer is quite simple. You do not need SPA, which comes with Oracle 11g. (And the RAT option is quite expensive!) We can build a simple test framework instead, which can perform the same work and even more! :)

I have published an article on this theme in Oracle Magazine Russian Edition (http://www.oracle.com/global/ru/oramag/july2009/russia_sta.html). Please, see the attached Power Point file with a small XML-styled presentation. The renewed article in English will be later.

Thomas Antila

Variations (of the column order in a fat index) and how to decide.

- How to calculate the optimal column order in a fat index.
- Can we trust testing with virtual indexes?
- Event 10053
- How to choose indexes in the database design phase.



Pekka E Liukkonen

Implementing archiving solution for Oracle EBS

Tieto has implemented Informatica Archive product for Tieto's ERP system to move inactive data to an archive database on inexpensive high capacity disks but still seamlessly accessible by authorised users via normal EBS user interface. The presentation can be either 30 minutes short introduction of the solution or 60 minutes to include also experiences of the implementation project and operations. This doesn't fit to the subjects of the planned agenda, but if there is room and interested attendees, I am ready to give this presentation.

Presenter: Pekka Liukkonen, Tieto Oyj

Manager, TERP Operational Services projects

Managing Tieto ERP (EBS) operational services since 2006. Developing processes of Tieto Outsourcing services and working in acquisition and sales projects since 2000. Working as a production manager of ICT operations services for many Finnish industries. Building first online applications in 1980's for Finnish metal industry.

Dmitry Volkov, Oracle

Oracle Exadata Deep Dive

This presentation will cover technical features of product and some lessons learned after real testing in Oracle. I want also include small demonstration of Exadata Software.

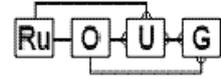
Part One: What's inside the non black box

Part Two: What we can see from application side

Part Three: How to prepare for testing, what we need to mind before, and some Real World experiences after testing in Oracle Reading Center

Rasmus Knappe, Oracle

11g R2 New Features



Tammy Bednar, Oracle

A technical deep dive on both ASO network and at rest encryption.

Lauri Pietarinen

Proactive Index Design using QUBE

Surprisingly, after years of using database technology, indexes are in many sites a dark area and one often finds databases with loads of one-column indexes. But what is a good index and how do we estimate it's effectiveness before the fact? One solution is to use QUBE, the Quick Upper Bound Estimate, a simple method developed by Tapio Lahdenmäki. In this presentation we also talk about the benefits and trade offs of fat (covering) and semi fat indexes and what makes an index (or a set of indexes) "perfect" for a specific SQL.

Evgeny Gorbokonenko

New strategy of Oracle RAC 11g R2 deploying.

Since the version 11g R2, Oracle supports integration with network services DHCP, DNS and NTP at the Grid Infrastructure level. The given integration, on the one hand, allows to use Oracle Clusterware new features effectively, such as SCAN-addresses, pools of servers, Grid plug'n'play. On the other hand, DBA must be more careful with planning of the required configuration and work in cooperation with network and system administrators.

In the presentation I wish to show on working examples how to prepare an environment for RAC deploying and how to integrate Oracle Clusterware with network services. The very short overview of new features is given in presentation and pros and cons of new environment are considered.

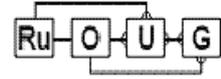
Peter Robson, JustSQL

Reversing Committed Transactions in SQL*Plus.

CTRL-Z in SQL*Plus - Infinite Rollback!

Objectives:

1 – Present a simple method of reversing committed transactions – even from previous sessions run days, weeks or months earlier.



2 – Understand transaction logic within a referentially constrained environment

3 – Appreciate the power of SQL*Plus when used with nested scripts

Abstract

Long experience of building and running audit systems resulted in the realisation that for very little extra work, a totally comprehensive roll back system can be built. This makes it possible to reverse the state of any audited tables back to any point in time - irrespective of the number of intervening commits, so long as the audit tables are complete.

A suite of dynamic interdependent SQL scripts have been constructed to enable any user (permission dependent, of course) to roll back through as many previous commits as required.

The technique is based on a table and transaction auditing system. These recoveries respect all referential integrity constraints, and do not require access to DBA system privileges. They can be implemented on individual schemas – or across the enterprise.

The process has similarities with Log Miner, but is easier to use, does not require the intervention of a dba, can be run by an individual user, and with care may be directed at specific tables of interest. It is powered exclusively within SQL*Plus, and is therefore fully functional against all versions of Oracle.

The presentation will describe the system design, and will demonstrate in a live session how it can be used.

Outline

I Objectives of the Presentation

- 1 Define the Problem
- 2 Describe the Solution
- 3 Demonstrate the specialist techniques involved

II The Problem Defined:

- 1 Requirement to rollback 'n' committed transactions
- 2 Requirement to do this without use of conventional backup / recovery systems

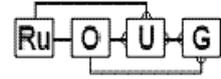
III Criteria for Success

- 1 Easy to Understand
- 2 Easy to Implement
- 3 Easy to Use
- 4 Can be used by any user

IV The Solution Architecture

- 1 Audit logs on each table
- 2 Single Transaction audit log for all tables
- 3 Triggers drive the auditing

V Design Requirements



- 1 Must be accessible by any user
- 2 Must be independent of DBA processes
- 3 Must respect full referential integrity
- 4 Use must be as simple as CTRL-Z in Windows

VI Implementation –

- 1 Architecture Details, Overview
- 2 Audit Tables, structure
- 3 Data –to- Audit Table Relationships
- 4 Describe SQL*Plus Nested Scripts engines

VII Building the Infrastructure

- 1 The Audit Tables
- 2 The Transaction Audit Tables

VIII Building the Triggers

- 1 Automatic trigger generation, using SQL*Plus
- 2 Detailed analysis of the scripts
- 3 Examples

IX Audit Data Flow

- 1 How data moves within the system
- 2 Rollback achieved by Reversing this flow

X Rollback Data Flow

- 1 Referential Integrity – Crucial
- 2 Transactions are Reversed
- 3 Example Schematic

XI Live Demonstration to Show:

- 1 – Building Tables
- 2 – Building Triggers
- 3 – Typical Transaction, Insert, Update, Delete
- 4 – Commit!
- 5 – Repeat 'n' times
- 6 – Interactive Rollback as required.

XII Summary

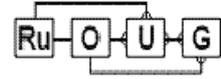
- 1 Other Areas of Application

Wolfgang Scherrer

[How to design a Oracle Technology and Solution based IT Blueprint](#)

The session will cover all aspects of designing a complete IT Blueprint Architecture.

The architecture includes Oracle Fusion Middleware (SOA, Content Management) together with Oracle E-Business Suite R12.1 . It will show the way from a first basic concept to a final solution with all problems and open questions in the middle.



Maris Elsins

Surviving the Crisis with the Help of Oracle Database Resource Manager

This is the time of cost savings and the options of extensive performance tuning of applications or upgrading the HW might be limited. Despite the fact that the Resource Manager is included in Oracle Database Enterprise Edition for no extra cost, this feature seems to be underestimated and rarely used.

This presentation will provide information on how Database Resource Manager can be used to prioritize the execution of most important processes when the limit of server resources is reached. The presentation will cover the basic theory about the Resource Manager and will provide live demo on how it can help distributing server resources between processes optimally.

Speaker: Maris is an Oracle Database and Oracle e-Business Suite Administrator currently employed by Tieto Alise located in Riga, Latvia. He has led or taken part in Oracle eBS 11i and R12 implementation, maintenance, migration and upgrade projects. Maris has an extensive experience in troubleshooting and performance tuning of Oracle Databases and e-Business Suite systems, his knowledge has been proven by acquired certificates - 9i/10g/11g DBA OCP and 11i eBS DBA OCP. Maris has also been a speaker on several UKOUG and LVOUG events.

3. Small room

Daniel Morgan

The Oracle ACE Program

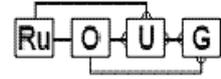
The Oracle ACE Program is designed to recognize and reward Oracle experts for advocating Oracle Technology and Applications. Oracle ACE recipients are chosen based on their significant contributions to, and activity in, their respective community. The program currently has two levels: Oracle ACE and Oracle ACE Director. This presentation will highlight the Oracle ACEs at OUGN and explain the ACE program

Joergen Holmer-Bretlau

Development of Oracle Support interfaces

Short Description: Story about the development of Support Interfaces. What was good what was bad. Ending up with hints & tips for how to filter away information that might be irrelevant for the user in My Oracle Support.

Duration: 45-60 minutes.



Sebastiaan Vingerhoed, Oracle

Real Application Testing

4. Roundtables/workshops

Antti Koskinen

(The Myth of a) Database-Independent Application - roundtable

The purpose of this roundtable is to illustrate why there still aren't, and probably never will be, truly database-independent applications, and also to provide ideas on how applications could be designed to allow support for additional DBMSes to be added later. The discussion will begin from the data model level, progressing through differences in physical database implementation, connection management, query execution, transaction isolation and so on. Or, to be more exact, these are among the topics covered in discussion slides; the discussion will by no means be constrained to follow a premeditated track, but will rather be allowed to sidetrail and flow as freely as possible. The support slides will cover Oracle, MySQL, PostgreSQL and MSSQL, but attendees with experience on other DBMSes are most encouraged to join the discussion as well.

Sue Harper, Oracle

Oracle SQLDeveloper Data Modeler
Workshop

The attendees should bring laptops, and maybe even download material before arrival. The Data Modeler tool can be installed in the workshop.

Learn more about the Data Modeler

<http://www.oracle.com/technology/products/database/datamodeler/index.html>

and SQL Developer on OTN

http://www.oracle.com/technology/products/database/sql_developer/index.html

Watch online demos, read supporting collateral and try sample code and tutorials

Blog: <http://sueharper.blogspot.com/>

Book: Oracle SQL Developer 2.1 <http://www.packtpub.com/oracle-sql-developer-2-1/book/>

Tammy Bednar, Oracle

security hands-on lab