April 2006 • Nr 35, Årgang 7 • ISSN 1600-5147 • Pris: kr. 300,00 ex moms • www.OracleEkspert.dk

#35 ORACLE EKSPERT

Questioning Solutions Since 2000

Live 23

DesWeb SIG-møde om SQL Developer evaluering
Dato: 26. april 2006

Kursus: Objektorienteret Analyse
Dato: 1. maj 2006

Optimising Oracle—Performance by Design med Jonathan Lewis
Dato: 2. maj 2006

Kursus: Avanceret Systemanalyse
Dato: 25. maj 2006

Scaling Oracle: En tre dages Master Class med James Morle
Dato: 14. juni 2006

ODTUG-konferencen 2006
Dato: 17. juni 2006

OracleEkspert-konferencen 2006
Dato: 30. august 2006

Nyheder 17

Oracle lancerer Business Intelligence
Oracle vinder patentsag
SQL Developer frigivet
Staten Wisconsin dropper Oracle
Oracle laver erhvervs-Google
IT-budgetterne fastholdes i 2006
Ekstraordinær sikkerhedspatch
IT-folk skifter jobs i 2006
Oracle vil tage markedsandele fra SAP
Oracle anklager Microsoft

OracleEkspert-panelet 2
Af Marc de Oliveira

Udnyttelse af Kernens buffer-pools 4
Af Martin Jensen

Denne artikel søger at belyse hvordan disse buffer-cache områder anvendes, samt hvorledes man ser hvilke blokke, der ligger hvorhenne. Alle eksempler er afviklet på en Oracle kerne version 10.2.0.1.

There is No Object-oriented Analysis 6
By David C. Hay

The world's bookstores are now full of books concerned with "object-oriented analysis". Is it possible that the world of object-orientation has completely changed the nature of systems analysis?

SpørgeJørgen'et 12
Af Mogens Nørgaard

Oracle Development Tools User Group Bestyrelsesmøde 2006 14
Af Marc de Oliveira

Meningsmålingen 19
"Podcasting et godt medie til Oracle-stof"

Nørgård's News 20
Af Mogens Nørgaard

LEDER

Af Marc de Oliveira

OracleEkspert-panelet

Så er OracleEksperts seneste datter blevet født! Udover at være et fysisk tidsskrift er OracleEkspert desuden et elektronisk dokument, en årlig konference og siden 6. april 2006 også et podcast... Og et meget populært et af slagsen!

I skrivende stund ligger "OracleEkspert-panelet" nr 1 på podhead.dks Top 25 foran populære podcast som Boris Boll-Johansen og Bossy Bos vanvittige "Løgnehistorier", traileren til den nye Gasolin-film, Casper Rongsteds NLP-podcast og Biobums-en - et ugentligt podcast med anmeldelser af ugens præmierefilm.

Hvis du endnu ikke har lyttet til vores nye podcast, så skynd dig ind på www.podhead.dk og lyt til den første episode direkte på din pc eller via din mp3-afspiller.

I denne episode diskuterer Martin Jensen, Mogens Nørgaard, Sten Vesterli og jeg emnerne:

- Fusion
- Gartners kritik af Oracles sikkerhedshuller
- Larry Ellisons privatforbrug
- "Group by" bliver ikke længere sorteret
- Problemer med Collaboration Suite
- Oracle Secure Enterprise Search 10g
- Oracle Sql Developer
- Problemer med dual-core i Oracle 10g release 2

Når du har hørt episoden, er vi meget interesseret i at høre, hvad du synes om den og om podcastet i det hele taget. Er det noget, som burde være anderledes? Nye ideer? Ris? Ros? Alle henvendelser er velkomne til Marc@OracleEkspert.dk.

Og husk at fortælle andre om, at de kan høre OracleEkspert-panelet på www.podhead.dk. Det kræver mange lyttere at forblive på top 25! :-)

En anden vigtig begivenhed, som er ved at falde på plads er årets OracleEkspert-konference. I år afholdes konferencen onsdag den 30. august på Damhuskroen i København (mere præcist ved Damhussøen i Rødovre). Selv om Søpavillonen var et rigtig rart sted, så var det mindste foredragslokale alligevel for lille. På Damhuskroen kan man som på Trekroner og Søpavillonen opleve det klassiske københavn. Kroen er fra begyndelsen af 1600-tallet og har meget smukke lokaler og en hyggelig balsal, hvor morgenmad og frokost vil blive serveret.

På næste side kan du se de indlæg, som du vil kunne vælge imellem på konferencen. Og til dem, som når at tilmelde sig inden 1. juni er der 500 kr at spare på registreringsafgiften. Som OracleEkspert-abonnent sparer man desuden 300 kr ift andre deltagere. Husk dit brugernavn og din adgangskode, hvis du skal tilmeldes med abonnentrabat. Firmaer, der sender fem eller flere deltagere får den ene deltager gratis med.

Som det kan ses af programmet på næste side vil der blive optaget en episode af OracleEkspert-panelet med publikum ved OracleEkspert-konferencen. Her vil der være mulighed for deltagerne at stille spørgsmål til panelet og for panelet til at stille spørgsmål til deltagerne. Kom og vær med!

Endelig vil jeg gerne gøre opmærksom på "Meningsmålingen", som i denne måned handler om "open source". Oracle har i den seneste tid interesseret sig en del for open source database- og middleware-selskaber. Hvordan kan det være? Er der penge i open source? Er det fremtidens forretningsmodel? Gå ind på www.OracleEkspert.dk og giv din mening til kende. Jeg vil gerne opfordre alle til at deltage i meningsmålingerne, så vi dels kan få et så korrekt billede af opinionen i Danmark som muligt og dels kan få involveret så mange som muligt i debatten. Den mest interessante kommentar vinder et OracleEkspert-krus.

Oplag: . . . . . 150 kopier

Udgives af:
. . . . . PYTHIA Information
. . . . . Kongensvej 3
. . . . . 2000 Frederiksberg
. . . . . Danmark

Telefon: 2627 9991
Fax: 2632 9991
Email: . . . . . Info@OracleEkspert.dk
Web: . . . . . www.OracleEkspert.dk

Ansvarshavende redaktør:
. . . . . Marc de Oliveira
. . . . . Marc@OracleEkspert.dk

SpørgeJørgen'et
. . . . . Mogens Nørgaard
. . . . . Mogens@OracleEkspert.dk

Bentes Bamsers Bøger:
. . . . . Bente Rosenkrantz-Theil
. . . . . Bente@OracleEkspert.dk

Pris:

Enkeltnummer . . . . . DKK 300,00
1 års abonnement (6 numre):
1. abonnement . . . . . DKK 1600,00
Pr ekstra blad . . . . . DKK 600,00

Annoncer:

Annoncer til OracleEkspert skal være PYTHIA Information i hænde senest den 15. i måneden inden bladets udgivelse.

Annoncepriser kan findes på:
www.OracleEkspert.dk

Password: Henriette

ORACLE EKSPERT 2006

Damhus Kroen

30. AUGUST
kl 9:00 - 16:30
KØBENHAVN

PRIS

Inden 1. juni:
- **Alm pris: 2750.-**
- **Abonnenter: 2450.-**

Efter 1. juni:
- **Alm pris: 3250.-**
- **Abonnenter: 2950.-**

PROGRAM

Java-applikationer i Forms-fart? Prøv Oracle ADF
Sten Vesterli, scott/tiger

Oracle 10gR2 RAC på Windows
Martin Gamtofte, Miracle

Publicering af data via Web services udviklet i JDeveloper
Peter Lorenzen, WM-Data

Hvorledes man partitionerer en stor ikke-partitioneret tabel
Martin Jensen, Oracle Danmark

Rundt om OracleAS
Claus Jessing, WM-Data

.Net og Oracle
Dagur Dagsson og Johann Hermansson

3D Business Intelligence
Carsten, WM-Data

Erfaringer fra Miracle's Backup/Restore/Recovery-Center
Kaj Christensen & Henrik Højer, Miracle

Hvad laver mit PL/SQL-program egentlig?
Rune Mørk , NNE

Virtual Directory – hvad skal vi med det?
Søren Bach Andersen , WM-Data

OracleEkspert-panelet - Live podcast
Martin, Mogens, Sten & Marc

Oracle og øl - en overlegen kombination
Lasse Christensen, Miracle

Spørg Eksperterne - Om DBA
Panelleder: Mogens Nørgaard

Spørg Eksperterne - Om udvikling
Panelleder: Marc de Oliveira

Der kommer endnu et par sessioner, som ikke er blevet planlagt endnu. Der kan ske ændringer i programmet.

TILMELDING

www.OracleEkspert.dk
eller
telefon: 2627 9991

Af Martin Jonson

UDNYTTELSE AF KERNENS BUFFER-POOLS

Oracle Consulting. Martin har siden 1982 arbejdet med bl.a. Oracle's database-kerne, samt med forskellige aspekter af systemdesign og it-arkitektur.

INDLEDNING

I SGA'en af en Oracle-instance, ligger den såkaldte 'buffer-cache', hvor disk-blokke hentes ind, eventuelt ændres og skrives ud igen efter en 'least-recently-used' strategi. Det har i mange år været muligt at reservere særlige områder i denne buffer-cache for database-objekter med mere specielle anvendelses former. I gamle dage (Oracle8i) var det relativt enkelt at se hvilke blokke, der lå i hvilke områder.

Denne artikel søger at belyse hvordan disse buffer-cache områder anvendes, samt hvorledes man ser hvilke blokke, der ligger hvor. Alle eksempler er afviklet på en Oracle kerne version 10.2.0.1.

Uanset om man nu har angivet hvor meget plads i SGA'en man vil afsætte til instansens buffer-cache eller man lader kernen selv afgøre dette, kan man med følgende select se hvor meget der pt. er afsat:

```
select value from v$sga
where name = 'Database Buffers';
```

På min 10g Release 2 kerne, hvor jeg kun har anvendt tablespaces med 8k blokke, er der afsat omkring 120 Mb. Lad os nu afsætte noget af denne plads til to særlige områder: KEEP og RECYCLE; den resterende plads kaldes for DEFAULT:

```
alter system set db_keep_cache_size=12M scope=memory;
alter system
 set db_recycle_cache_size=12M scope=memory;

select name, current_size, target_buffers from
V$BUFFER_POOL;

NAME     CURRENT_SIZE   TARGET_BUFFERS
-------  ------------   --------------
KEEP     12             1497
RECYCLE  12             1497
DEFAULT  104            12974
```

Vi opretter nu en almindelig tabel og et indeks med default parametre, således at blokke fra disse 2 database objekter vil komme ind i DEFAULT cachen, når de skal bearbejdes:

```
create table cache_test
 (id number not null, sv varchar2(20));

insert /*+APPEND */ into cache_test (id, sv)
select
 t.column_value id,
 'row '||to_char(t.column_value) sv
from table(system.counter(1,200000) ) t;
commit;

create unique index ct_pk on cache_test(id);
```

Ved at samle statistik om tabellen og kigge i user_tables ses det, at tabellen reelt fylder 564 blokke. Nu ændrer vi tabellen, således at kernen skal anvende KEEP området fremfor DEFAULT for dette objekts blokke, for på denne måde at holde denne tabels blokke i et særligt område, der ikke automatisk forældes af andre objekters vej ind til DEFAULT området. KEEP området bliver dermed en form for memory database, under forudsætning af at objekterne med anvendelse af dette KEEP område faktisk kan være i området:

```
alter table cache_test storage (buffer_pool keep);
```

Vi renser først buffer områderne, og så foretager vi et full-tablescan af tabellen, og ser hvordan det der ud i vores buffer statistik. Med flush kommandoen vil buffer-cache arealerne nulstilles (efter naturligvis at have skrevet eventuelle ændrede blokke til disk-systemet), mens selve tællerne i v$buffer_pool_statistics er uændret:

```
alter system flush buffer_cache;

select count(sv) from cache_test;

select name, buf_got, consistent_gets, physical_reads
from V$BUFFER_POOL_STATISTICS
where name = 'KEEP';

NAME     BUF_GOT        CONSISTENT_GETS PHYSICAL_READS
---      -------        ---------------- -------------
KEEP     565            568             565
```

Hvis vi foretog endnu et tablescan gennem tabellen CACHE_TEST, ville consistent_gets tælleren blive 568 større mens antallet af fysiske læsninger ville være uændret.

Hvis der allerede havde været blokke for cache_test objektet i en af de andre pools, ville det blive liggende og derfor ikke læst ind i KEEP arealet.

Lad os nu også markere at indexets blokke også skal ind i dette KEEP område:

```
alter index ct_pk storage (buffer_pool keep);
```

Ved nu at foretage en lille forespørgsel mod dette index, ses det at vores statistik på buffer_cachen tæller 2 op (en for rood-blokken og en for blad-blokken):

```
select count(id) from cache_test
where id between 1 and 10;

select name, buf_got, consistent_gets, physical_reads
from V$BUFFER_POOL_STATISTICS
where name = 'KEEP';

NAME     BUF_GOT        CONSISTENT_GETS PHYSICAL_READS
---      -------        ---------------- -------------
KEEP     567            570             567
```

I Oracle8i kunne man direkte ud af blokkens nummer i buffer-cachen direkte se i hvilken pool den lå (lo_bnum og hi_bnum i v$buffer_cache) - Metalink Note 135223.1 (Oracle Multiple Buffer Pools Feature). I Oracle9i og Oracle10g findes den information anderledes, da hele buffer cachen er delt i mindre enheder, for nemmere dynamisk at kunne skrue op og ned for de enkelte puljer. Her findes de enkelte blokke, der er placeret i KEEP arealet med følgende select:

```
SELECT
 bh.dbarfil, bh.dbablk, bh.ts#, bh.class,
 DECODE(bh.state,
        0,'FREE',1,'XCUR',2,'SCUR',3,'CR',4,'READ',
        5,'MREC',6,'IREC',7,'WRITE',8,'PI',9,'MEMORY')
FROM x$kcbwbpd bp, x$bh bh, x$kcbwds ds
 WHERE bp.bp_size > 0
   AND ds.set_id >= bp.bp_lo_sid
   AND ds.set_id <= bp.bp_hi_sid
   and bh.set_ds  = ds.addr
   AND bh.inst_id = USERENV('Instance')
   AND bp.inst_id = USERENV('Instance')
   AND ds.inst_id = USERENV('Instance')
   AND bp.bp_name = 'KEEP'
ORDER BY 1, 2, 3, 4;
```

For nu at finde ud af hvilke databaseobjekter de enkelte blokke i KEEP arealet de enkelte blokke kommer fra, kan vi joine med dba_extents. Desværre vil en sådan join tage mærkbar tid og ressourcer på et større system, hvorfor det anbefales at oprette en egentlig tabel med disse informationer før der joines:

```
create table mj_ext as
select * from dba_extents;

create index s_1 on mj_ext(file_id, block_id);

/* join condition: */
and
 bh.dbablk between e.block_id and e.block_id + e.blocks
```

Her vil buffer-cache arealerne nulstilles (efter naturligvis at have skrevet eventuelle ændrede blokke til disksystemet) – men selve tællerne i v$buffer_pool_statistics er uændret.

I grunden er der flere områder i kernen's buffer-cache, nemlig et område for hver blokstørrelse (bortset fra default). Så hvis blokstørrelsen er på 8k, har man mulighed for at konfigurere 4 andre arealer, til at håndtere blokke fra tablespaces med den pågældende blokstørrelse:

```
DB_2K_CACHE_SIZE
DB_4K_CACHE_SIZE
DB_16K_CACHE_SIZE
DB_32K_CACHE_SIZE
```

TECHNICAL ARTICLE

By David C. Hay

There is No Object-oriented Analysis

A veteran of the Information Industry since the days of punched cards, teletype machines, and paper tape, Dave Hay has been producing data models to support strategic information planning and requirements planning for nearly twenty years. He has worked in a variety of industries, including, among others, power generation, clinical pharmaceutical research, oil refining, forestry, and broadcast. He is President of Essential Strategies, Inc., a consulting firm dedicated to helping clients define corporate information architecture, identify requirements, and plan strategies for the implementation of new systems. He is the author of the book, Data Model Patterns: Conventions of Thought, published by Dorset House. He may be reached at dch@essentialstrategies.com, http://www.essentialstrategies.com, or (713) 464-8316.

The world's bookstores are now full of books concerned with "object-oriented analysis". Is it possible that the world of object-orientation has completely changed the nature of systems analysis? To listen to some of its aficionados, this must indeed be true. Information Engineering is passé, and has been completely replaced by object-oriented analysis.[1]

James Rumbaugh and his colleagues wrote one of the seminal books on analysis in the object-oriented world.[2] To their credit, they never used the expression "object-oriented analysis". They simply made analysis an important step in the object-oriented development process - as well they should have. The problem was that they then asserted that it was object-orientation that first made analysis data-centric[3], which is not the case. (Entity modeling has been around since 1977.) This attitude has had the unfortunate effect of permitting the object-oriented world to assume that any analysis that uses models of data structure is inherently "object-oriented". This then forms the basis for asserting that the old way of doing things is now completely invalid.

For an entire body of knowledge that has developed over twenty years or more to be summarily overturned by a new approach seems unlikely. This has never happened before, and it is difficult to believe that it has happened now.

More likely is the possibility that the object-oriented folks have added something to the body of knowledge. Perhaps, in addition to its contributions to design and programming, the object-oriented world has something to contribute to the field of analysis. Or perhaps not.

When data modelers refer to entities as the "things of significance to a business" (as they have done for some fifteen years, now), the argument could well be made that these are the "object classes" of the business. Occurrences of these entities are the business' "objects". When we address a business as object-oriented analysts, if we limit ourselves to business object classes and their associations, we are in fact being information engineers.

But creation of the expression "object-oriented analysis" itself suggests that more than the business objects we have been looking at are involved. The object-oriented analysis described by Mr. Rumbaugh and his colleagues is nothing other than that which we information engineers have been doing for a long time. For them to appropriate credit for these techniques is relatively harmless. More dangerous is the argument often made that object-oriented analysis is somehow *more* than regular analysis. The problem is that this "more" isn't really analysis.

Some of the books your author has encountered (including Rumbaugh's) do limit their view of the analysis of data to the objects of significance to the business.[4,5] Other authors, however, do not limit themselves to business objects, nor do they even seem to be oriented toward identifying business needs.

First of all, they are not truly oriented toward the business. This is evident in the language used. As a minor example, consider the names of entities and attributes. In UML and other "object modeling" notations, spaces are not permitted between words in entity and attribute names. If, however, this is to be a model of the business - validated by human beings - then conventional English must be used. The expression most business people will recognize is "purchase order", not "PurchaseOrder" or "Purchase_order". The statement being made by this convention is that the model is being prepared, not in service to the people of the business, but in service to the computer.

(It is true that some practitioners do use spaces between words,[6] and the original specification for UML does not specify a rule either way[7], but the prevalence of those who don't allow spaces says something profound about the state of the industry.)

The problem in fact goes beyond spaces between words in names: in at least one of the books your author has read, object classes (definitions of things of significance to the business, remember) are called things like "CustInfo" instead of "Customer".[8]

Another author at least spelled out "Student-Information", but she then defined this class as "Information needed to register and bill students." As an afterthought she acknowledged that a "student is someone who is registered to take classes at the University",[9] but this wasn't the definition of the class.

The business object classes are not "student information" and "customer information" (even spelled out with spaces). The business object classes are "students" and "customers".

Grady Booch, in his seminal work on *Object-oriented Analysis and Design*, was unable to define his principle concepts without using program code. This made it a little tough for those of us who do not know C++ to understand what he was talking about. For example, he defined the sub-type/super-type relationship with

the following bit of pithy text:

```
// Yield type
typedef unsigned int Yield;
class FruitGrowingPlan : public GrowingPlan {
public:
  FruitGrowingPlan(char* name);
  virtual "FuitGrowingPlan();
  virtual void establish(Day, Hour, Condition&)
  void scheduleHarvest(Day, Hour);
  Boolean isHarvested() const;
  unsigned daysUntilHarvest() const;
  Yield estimatedYield() const;
protected
  Boolean repHarvested;
  Yield repYield;
};[10]
```

Apparently this identifies a "FruitGrowingPlan to be a sub-type of "GrowingPlan". Please forgive your author if he does not use this definition when next he tries to introduce the business community to data modeling.

Second, these books add object classes that are not oriented toward identifying business needs. There are two categories of such classes.

First, there are object classes that constitute what are simply bad models:

In one case, for example, the additional object classes "CustInfoMulti" and "CustInfo1" were defined as separate sub-types of the unfortunately named CustInfo entity referred to above. These sub-types were defined to distinguish customers with just one address from those with multiple addresses.[11] This is nonsense, since the underlying object class is still "customer". The number of addressees it has should be shown by the association between Customer and Address. The association "Each Customer may be at one or more Addresses," asserts that if a Customer may have one or more addresses, it certainly may have one. Or it may have more. There is absolutely no reason for there to be two subtypes.

Now bad models should not be held against object-orientation. After all, information engineers have produced their share of bad models as well. Your author has found only a few object-oriented books,[12,13,14] however, where the models were not perverse in the manner of the above example. The issue is not what could be done, but what object modelers are taught to do. The fact of the matter is that rarely are they taught to respect the underlying structure of the information.

A second category of object-oriented "extensions" to the data model are classes that in fact belong to the domain of function modeling and physical design. These come from the object-oriented developer's emphasis on behavior. To describe this, though, it is worth a minute to discuss entities and behavior:

OBJECTS AND BEHAVIOR

The most profound difference between the object-oriented approach to modeling and the information engineering approach is the inclusion of behavior in the class definitions. As an approach to software *design* this has been a significant addition to the body of knowledge in our industry. The extension of the idea to *analysis* models, however, is more problematic.

In the case of basic classes, it is simple enough to add documentation of behavior, without affecting the underlying structure of the model. The issues determining what classes should be there in the first place would not be affected. Indeed, the idea of adding behavior to an entity definition is not a bad one, and could result in useful insights. As Bob Brown has said, "entities are just objects that don't know how to behave."[15] To do so during analysis, however, the behavior involved would have to be true business behavior, and not the behavior anticipated to be designed into a system.

The examples found in most texts don't describe business behavior. They are written in a kind of structured English that looks suspiciously like program code.

Indeed, Grady Booch says, "We must emphasize that analysis focuses upon behavior, not form." While your author, a data bigot, wouldn't be happy about it, this is not an unreasonable position to take. The problem is that he then goes on to say, "It is inappropriate to pursue issues of class design, representation, or other tactical decisions during this phase. Rather *analysis* must yield a statement of what *the system* does, not how it does it."[16] [Emphasis added.]

In other words, his version of object-oriented analysis is not concerned with object data models at all. (So much for Rumbaugh's assertion that object-oriented invented datacentric analysis.) In Mr. Booch's view, object-oriented analysis isn't about the articulation of the business requirements for a system. He as much as said that object-oriented analysis is about system design. Indeed, even to the extent that analysis is concerned with functions, they are the functions of a yet-to-be-designed system, not of the business.

Things get stranger when a logical conclusion of the object model approach is that you can have "objects" with behavior but no data. In particular, there are two categories of object classes like this: "Control classes" are used to control processing in a system, and "boundary" or "interface classes", that "interact with things outside the application".[17] Neither of these describe business data structure, and indeed, following Mr. Booch's definition above, they don't describe business function, either. They essentially programming constructs, describing a hypothetical system.

For example, one text specified "ProfessorCourseManager" and "AddACourseOffering" as classes.[18] The first is neither data nor function, but merely the fact that a professor can control a course. If this were presented as an associative class between Professor and Course (and named something like "CourseManagement"), where an occurrence of the class is the role of a Professor's managing a Course, it would be a legitimate class. But as presented, it is only related to Course, and represents control of the course, not a fact about it.

The second is simply a function in disguise. In fact even an object-oriented analyst would recognize it as a method that should be captured *inside* the CourseOffering class.

Your author once sat in on a presentation of an object model which consisted of an object class that performed some function, then handed off data to another object that did another function, which in turn handed other data off to yet a third object that did something else.

Excuse me, but how is this different from information engineering's data flow diagram? (Except that it is being presented as a data structure diagram.)

Yes, it is true that information engineering produces process models. At this point the argument could be made that it is an extension of the object-oriented philosophy that function boxes and control boxes should appear along with data boxes on object models - but

adding behavior to an object class is not the same thing as adding function boxes to an entity/class diagram. They really are different things, and should be the subject of different models.

As mentioned above, this doesn't mean that it wouldn't be useful, in a business context, to describe the life cycle of an entity/object class. What functions create occurrences of it? What update them? What delete occurrences? These are legitimate extensions of the entity/relationship modeling technique. A technique called "entity life histories" (which predates UML and other object-oriented analysis techniques by many years), addresses exactly this issue.[19] It was originally described by Michael Jackson (no, not the popular singer) in his 1983 book, *System Development*,[20] and it has been incorporated into the SSADM methodology widely used in Europe.[21]

While there is a good argument for adding the behavior of an entity to the entity's description, cobbling together object and function *models* simply confuses the issue.

In short, there is no such thing as "object-oriented analysis". There is only "analysis", to which the object-oriented folks have added a dollop of object-oriented design. While object-oriented analysis has been offered as a fundamentally new way of doing analysis, what is really new is that this version of analysis contains elements specifically oriented toward the *implementation* of systems using object-oriented technology.

To understand why this is misguided, it is necessary to understand John Zachman's "Framework for Information System Architecture".

THE ZACHMAN FRAMEWORK

By 1989, even though the industry was awash with methodologies, modeling notations, and different ways to "communicate" with each other, we in fact rarely communicated very well at all - with each other or with our clients. In response, John Zachman came up with some important insights, and from them developed his "Framework for Information Systems Architecture".[22] Among other things, Mr. Zachman realized that a source of our communications problems was that each of us views the problem of information system development from a quite different point of view. Mr. Zachman's recognition of these points of view was important, because it finally clarified why it is that we have such difficulty talking to each other. In particular, he identified six different perspectives in any system development project:

- ***Scope*** - The understanding of why the organization exists, how it is like other organizations in the same industry, and what makes it distinctive.
- ***Business owner's view*** - This refers to the perspective not of the stockholders, but of the people who operate the business. This encompasses all the jargon of the business, as well as an understanding of how everything actually works.
- ***Architect's view*** - Recognition that there are fundamental, technologicallyindependent structures present, and representation of the business in terms of these structures.
- ***Designer's view*** - The application of technology to address any information system requirements discovered in the views above. The perspective here is technological: relational data base systems, object orientation, network protocols, and so forth.
- ***Builder's view*** - This is the view of the inside of the programs and technology. The builder knows the finer points of the programming language, or the communications technology, or whatever.
- ***Production view*** - The view of the completed system.

If two people are coming at the problem from different perspectives, their language and terms of reference will be different, and they will be working at cross purposes, unless they recognize the differences in their points of view - and translate. With the translation it becomes clear to both why they believe what they do - and it becomes clear that, as long as they recognize and respect each other's positions, this doesn't matter. Each has a different job to do, and each's perspective is appropriate to the carrying out of those jobs.

Note that this is not an argument for the "waterfall approach" to developing systems. The waterfall addresses the perspectives in sequence: first, strategy addresses scope; then, analysis deals with the business owner's and architect's views; later, design captures the designer's view; and so forth. But this is not necessary. If you choose to start with design, you may do so. The Framework simply tells you what it is you don't know when you do that. That is, you don't know the business owner's perspective, so you will have to guess at what the system is to do. You don't know the architect's perspective, so you will have to cobble the system together as best you can, without benefit of any understanding of fundamental structures.

Are you willing to risk the consequences of not knowing those things? If so, go right ahead.

If you decide to do a little bit of analysis, followed by a little bit of design, and then return to analysis, that's fine. Again, it's your choice. You can now understand the shortcuts you are taking and the potential costs of those shortcuts. If you decide to do analysis and design at the same time, the Zachman view simply requires that you recognize you are dealing with two perspectives at once, and you must be willing to accept the implications of this.

WHAT OO MISSES

The Zachman Framework now makes it possible to describe what is wrong with the object-oriented world's approach to analysis. Specifically, object-oriented analysis begins with a cross between the business owner's view and the designer's view, moves directly to the designer's view, and never addresses the architect's view at all.

Yes, Grady Booch specifically describes the importance of architecture, saying that "A system with sound architecture is one that has conceptual integrity and, as Brooks firmly states 'conceptual integrity is *the* most important consideration in system design.'"[23,24] And yet we saw above just how much architecture Mr. Booch puts in analysis. It is true that Mr. Rumbaugh and his colleagues devote a large part of their book to analysis, but if you look at the details of how object-orientated analysis is too often carried out, you can see that architecture - based on the architecture of the enterprise to be served - gets pretty short shrift, on both the function and data side of the effort.

Functions

A popular technique in the object-oriented world is "use cases". These are represented as diagrams of specific anecdotes offered by people in the business being examined. They are supposed to describe the specific activities carried out by an individual. As such they should be effective at capturing the business owner's view.

As it happens, however, they don't actually do that.

"A use case diagram provides a functional description of a *system* and its major processes and places a boundary on the problem to be solved. It also provides a graphic description of who *will* use the system and what kinds of interactions they *can expect to* have with the *system*."[25] [Emphasis added.]

Note that a use case describes a future "system", not the existing enterprise. Moreover, it is about a future world when a new system will make all the business person's problems magically disappear. A use case does not actually describe the physical world presently experienced by the user. ("There are five copies of a purchase order and the gold one goes to accounting".) Instead, it describes a job in terms of the technology anticipated for the system (complete - in 1998 - with lap-top computers and web access[26]). Rather than describing the current business, it describes an imagined system - one that hasn't even been designed yet.

It is the beginning of a user-interface design. That's all.

As it happens, many of the business processes observed by the analyst are not really intended to carry out the functions of the business at all. They are only there to accommodate inadequacies in current systems, and it is a waste of time for us to try to automate them. Instead of trying to do that, we should be asking, what is the business trying to accomplish with these processes the analyst sees? What are the underlying functions to be carried out? What are the "essential" processes (without regard for the technology that might be used) required to carry out the functions the business wants to do? How can technology bring us closer to providing the true functions of the business?

These questions are never asked.

Even if the use cases could be used to infer something about the current nature of business systems, developers then typically go from this supposedly business owner's view directly to the designer's view. There is often no attempt to examine the set of use cases either to determine which should be automated or to identify general patterns of behavior or underlying functional structures. Nor is there provision for analyzing a set of use cases to make sure that they neither overlap nor leave out important activities.

Data

If the object-oriented approach doesn't come to grips with the underlying *functions* of an enterprise, what about its *data*? We have already seen Mr. Booch's view of the importance of understanding data structure at all during analysis. Indeed, to the extent that it is done, it typically only addresses the objects seen by the users, without trying to understand their inherent structure.

First of all, since object-oriented technology does not require normalization, this step is often skipped, precluding development of a sound understanding of the underlying structure of the data involved.

When Dr. Codd invented the relational model, and with it "normalization", the idea was that the normalization techniques were required to make relational databases work. With the appearance of object-oriented database technology, presumably normalization is no longer required.

This would be true, except for the fact that the advent of normalization had a much more profound effect on the industry than the mere facilitation of relational databases: it has helped us systematically to eliminate data redundancy, and, more importantly, to better *understand* the *nature* of data. Each datum in an enterprise is *about* something - and only one thing. Our job as data analysts is to determine what that thing is. Moreover, there is only one datum that is that fact about that thing, and we must determine where to put that fact so it can be found again.

If we don't correctly identify the relationship between data and the things they describe, we haven't done our job.

A second problem with the object-oriented approach to data is that it tends to identify object classes simply from the use cases, as the objects seen by the potential system users. This approach fails to ask an important question: What are the *fundamental* things of significance to the business, of which the things most people see are but examples?

The business objects described in use cases are those that the individual sees in his own area of activity. Different individuals may be perceiving what are really the same objects, but with different names and characteristics. One department sees "customers"; another sees "vendors"; another sees "clients"; yet another deals only with relationships between departments in the same company.

Are not all these things simply instances of "parties" (people and organizations) in different roles? Indeed will not many parties play several of these roles?

Going beyond the boundaries of the enterprise itself, are there not similarities between the fundamental structure of this enterprise and those of other enterprises in the same or different industries?

Object-orientation is supposed to promote re-use: the place to start should be in the identification of standard entity/object models for standard business situations.[27]

In fairness, many object-oriented people have acknowledged the importance of understanding common structures while doing requirements analysis,[28] and at least one book has now taken up the issue of doing a proper conceptual architecture[29,30], but these views are not reflected in a large part of the literature.

To summarize, in ignoring the architectural view, many object-oriented practitioners fail to ensure that they are addressing the entire enterprise. They fail to examine the organization as a whole and to understand fundamental structures and processes - those which cut across the views of the individuals who run the business. There is no attempt to understand those structures which will exist regardless of any technology employed to deal with them.

Yes, it seems to be true that object-oriented programming makes it easy to change a design when something turns out to be wrong. But this approach to analysis will guarantee that there are a lot more wrong things to fix than would be the case if some architectural thought were given to the system before it was designed.

CONCLUSIONS

It has been difficult to write this article. The criticisms of object-oriented analysis presuppose that we know what object-oriented analysis is. As with any technique it is in fact practiced very differently by different practitioners. Many who claim they are object-oriented are in fact following the best of information engineering techniques.

Indeed, many information engineers are guilty of modeling sins far more egregious than those presented here. Your author is the first to acknowledge both of those points.

The problem is language. Our industry is highly susceptible to fads, especially in terminology. The fact of the matter is that the subject at hand is - as it has always been - nothing other than, how do we do a better job of building systems? Each year people come up with new techniques to further the craft. Unfortunately, new sets of techniques invariably get packaged with a new name and are presented as if they have completely overturned everything that went before. This is never true. The new techniques are always incremental additions to the body of knowledge and only that - fancy names notwithstanding.

This is the problem with the way object-oriented analysis has been presented. Of course object-orientation has contributed greatly to the programmer's art.

The assignment for requirements analysis, however, has not changed. It remains that of understanding an enterprise in technologically neutral terms. The objective is to determine the data, function, and other requirements for information processing - expressed in terms that are both clear to the business customer of that processing, and which enable any technology to be applied in addressing them. Nothing in the objectoriented arsenal of tools has changed that. Requirements analysis is fundamentally a dialogue between the system architect and the potential system user.

Many of its practitioners in fact do recognize that if object-oriented development is to succeed, its practitioners must acknowledge and adopt the valuable techniques which preceded the advent of object-oriented programming. A proper business analysis is still required. Disciplined understanding of the nature of the data that describe objects remains important. We have known how to do these things for a long time. We can't afford to throw this knowledge out.

In short, there is no such thing as "object-oriented" analysis.

1 Oracle Education Services, "Oracle Education Mini-lesson: Object Database Designer for Oracle8 Objects", Oracle Development Tools User Group, (Palm Springs, CA:1998), p. 3.

2 James Rumbaugh, Michael Blaha, William Premerlani, Frederick Eddy, William Lorensen, Object-oriented Modeling and Design, Prentice Hall (Englewood Cliffs, NJ:1991), PP. 153-156.

3 Rumbaugh, et al., ibid, p. 146.

4 Rumbaugh, et al., ibid.

5 Craig Larman, Applying UML and Patterns: An Introduction to Object-Oriented Analysis and Design, Prentice-Hall PTR (Upper Saddle River, NJ:1998).

6 Martin Fowler, UML Distilled, Addison-Wesley (Reading, MA:1997).

7 Rational Software Corporation, Unified Modeling Language Notation Guide, Rational Software Corporation, (Santa Clara, CA:1997), pp. 42-44.

8 Paul Harmon and Mark Watson, Understanding UML: The Developer's Guide, Morgan Kaufmann Publishers, Inc., (San Francisco: 1998).

9 Terry Quatrani, Visual Modeling with Rational Rose and UML, Addison-Wesley, (Reading, Massachusetts:1998).

10 Grady Booch, Object-oriented Analysis and Design, Benjamin/Cummings Publishing Company, Inc., (Redwood City, CA:1994), p. 61.

11 Paul Harmon and Mark Watson, ibid. p. 150-152

12 James Rumbaugh, op.cit.

13 Martin Fowler, op. cit.

14 Craig Larman, op. cit.

15 Bob Brown, "Extended Modeling Language", GUIDE Proceedings, November 1990.

16 Grady Booch, op. cit., p. 252.

17 Paul Harmon & Mark Watson, op. cit. pp. 123-125.

18 Terry Quatrani, op. cit. p. 54.

19 David C. Hay, "Object Oriented Data Modeling: Entity Life Histories", Oracle CASE Special Interest Group, 1993. Also available at http://www.essentialstrategies.com/publications/objects/elh.htm.

20 Michael Jackson, System Development, Prentice Hall, (Englewood Cliffs, NJ: 1983).

21 Ed Downs, Peter Clare, and Ian Coe Structured Systems Analysis and Design Method: Application and Context, Prentice Hall INternational (UK) Ltd, (Hemel Hempstead, Hertfordshire:1988).

22 Zachman, John, "A Framework for Information Systems Architecture", IBM Systems Journal, Vol 26, No 3, 1987.

23 F. Brooks, The Mythical Man Month, Addison-Wesley, (Reading, MA:1975), p. 42.

24 Grady Booch, op. cit., p. 230.

25 Paul Harmon and Mark Watson, op. cit. p. 112.

26 Paul Harmon and Mark Watson, op. cit. p. 116.

27 See David Hay, Data Model Patterns: Conventions of Thought, Dorset House, (New York: 1997)

28 Fowler, M, Analysis Patterns: Reusable Object Models, Addison-Wesley (Reading, MA:1997).

29 James Rumbaugh, et al., op. cit.

30 Craig Larman, op. cit.

Du Bi Du Bi Du Bu

Træt af at høre
den samme musik
igen og igen?

Tank din MP3-afspiller op
med danske podcasts, der
underholder, udfordrer,
og bevæger dig

LYT TIL
ORACLEEKSPERT-PANELET
NR 1 PÅ
WWW.PODHEAD.DK

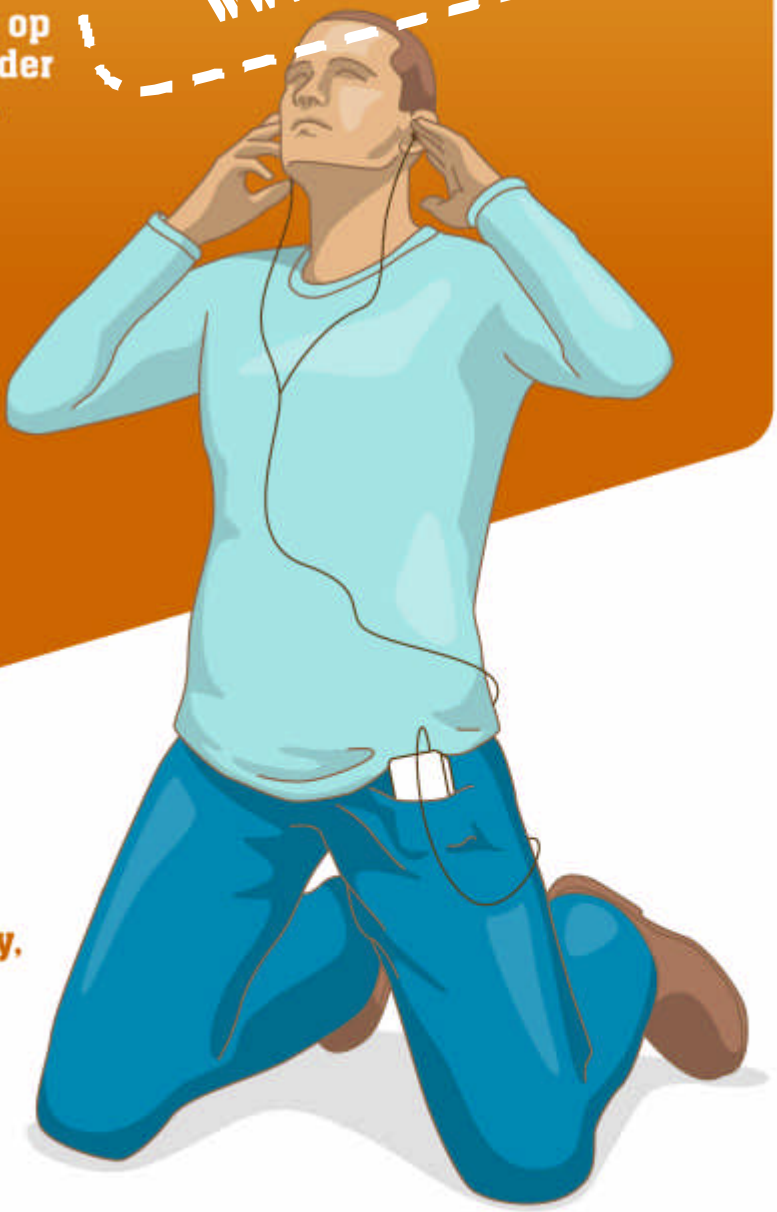

BRÆNDSTOF
TIL DIN POD

På podhead.dk finder du comedy,
talkshows, debat og litteratur,
lige til at tage med på farten

podhead.dk

iSQL*Plus er død

En læser skrev noget i retning af:

Metalink Note:359859.1 fortæller, at iSQL*Plus for Windows de-supporteres – man bør bruge andre ting, såsom SQL Developer, Application Express, eller lignende. Det synes en del folk er træls, og de vil gerne have det tilbage.

Det sker ikke. Oracle har – så vidt jeg ved – aldrig reageret på brugernes reaktioner omkring et produkts desupport ved at tage produktet tilbage. Det sker nok heller ikke her.

Jeg tror såmænd ikke der ligger andet bag denne desupport end at den person, der virkelig vidste noget om det har forladt skuden og man derfor ikke gider investere i, at en anden kan blive lige så god til at vedligeholde og udbygge produktet. Det er et typisk one-man produkt, og dem er der faktisk mange af i Oracle's historie: En gut får en ide, laver en prototype, som er god og nyttig, og så er det et produkt.

Raptor er et eksempel. Nu håber og beder vi bare til, at idémanden får lov til at være i fred for de store, fælles integrationer og andre dræbere.

Jeg ved også, at Miracle's egen Michael Möller er rigtigt godt irriteret over, at iSQL de-supporteres, fordi det faktisk var nyttigt og nemt at bruge.

Jo, men I VED jo også ganske udmærket, at når noget som helst IT-relateret når et stadie af stabilitet, forudsigelighed og god funktionalitet vil det blive afløst af noget, som ikke er dét. Men som til gengæld kan så meget andet...

Michael Möller addendum/bemærkninger:

Lidt præcisering af de berørte produkter:

Der er "iSQL*Plus" (kig nøje efter det foranstillede "i") som er en "webserver" så man fra en almindelig browser får en HTML side hvor man kan skrive SQL sætninger og få svar i HTML Tabeller. iSQL*Plus kan køre på alle platforme, og kan glimrende bruges som test/demo af 3-tier setup (browser – isqlplus – database) Desupport: 359855.1

Der er "SQL*Plus for Windows" som er det mærkelige enkel-vindue med Notepad lignende skrift. Kører kun på Windows. I dets vindue bruger det non-standard cut-n-paste. Tastaturgenvejen i mange versioner var shift-INSERT eller noget i den stil (der kan I se, hvor lidt hvor jeg bruger det) og mærker man tekst med musen er det en firkant, ikke begynd-slut af hele teksten. Sidst men ikke mindst, hvis man brugte EDIT eller HOST så var vinduet frosset, og indholdet ikke gentegnet hvis man flyttede et overliggende vindue væk. Desupport: 359859.1

"SQL*Plus" er kommandolinje værktøjet og det går ikke væk. Tænk på alle de scriptfiler Oracle selv bruger. Ingen Desupport (kun noget om versionen 8.1.0)

Jeg vil ikke savne "SQL*Plus for Windows". Jeg vil savne iSQL*Plus. Jeg kunne godt lide dens formatering af brede tabeller (V$DATABASE-48 kolunner, DBA_TABLES-49 kolunner og … V$SESSION-77 kolunner!), output i en browser med scrollbars (jeg blev dog vanvittig af default efter version 92 var at lave en "Next page" knap hver 24 linjer – fandt aldrig hvor man kunne omkonfigurere det default permanent) og at man havde fuld editering af alle sql-linjer i sin input-boks. Jeg har jo lavet en del undervisning, og isqlplus kan bruges for at demonstrere setup af SQLNet, 3-tier og også NLS problemer.

Men ingen panik – det sker først om 3½ år.

MIRACLE

Kære Hr. eller Fru. DBA

Igen i år tilbyder Miracle kurser, hvor kun de bedste folk underviser. Programmet byder i år på følgende:

Planlagte events for 2006:

14. marts	Besøg af Amir Kahana fra Symantec om I3 (tidl. Veritas). Information på den fede måde! Anjo Kolk er med i vores bestyrelse, så vi tror på, at vi kan få svar på spørgsmål og evt. tekniske problemer. (Gratis seminar)
2. - 4. maj	Master Class med Jonathan Lewis - Optimizing Oracle Performance by Design Manden selv er notorisk i Oracle verdenen og i 2005 udkom endnu en fantastisk bog (om Oracle's optimizer), som bogstaveligt talt blev revet væk på UKOUG's bruger-konference i Birmingham-UK.
14. - 16. juni	Master Class med James Morle - Scaling Oracle James er i gang med at opgradere sin bog "Scaling Oracle8i" til "Scaling Oracle10g", og det kommer der en Master Class ud af.
28. - 30. September	DatabaseForum 2006 (det allersidste, nogensinde, nogensteder, for altid)

Læs mere på www.miracleas.dk under Events & Uddannelse.

Andre nyheder fra Miracle:

ViBrugerIkkeCarlsberg.dk er gået i luften - opleves bedst med et glas koldt MiracleØl eller andet godt øl.

Aktietegningen er i fuld gang for Miracle Breweries, som i første omgang har udstedt 250 stk. folkeaktier, hvor ingen nogen sinde vil se deres penge igen. Til gengæld følger fordele som rabat på MiracleØl og lov til at bære hvide handsker til aktionærmøderne.

Nytårsfortsæt 2006: Test din backup....også den du har outsourcet!!

Venlig hilsen
Mogens Nørgaard

Oracle Development Tools User Group Bestyrelsesmøde 2006

Marc har været indlægsholder på ODTUGs konferencer siden 1999, og har siden 2003 desuden været medlem af ODTUGs bestyrelse.

ODTUG (www.odtug.com) er en uafhængig Oracle-brugergruppe lige som den danske Oracle-brugergruppe OUGDK (www.oug.dk).

De to væsentligste forskelle på ODTUG og OUGDK er 1) at ODTUG primært beskæftiger sig med systemudvikling, mens OUGDK dækker alle emner, der er relateret til Oracle, og 2) at det er lykkedes ODTUG at holde sig aktiv, mens OUGDK kun har meget få aktiviteter.

Som optakt til ODTUGs kommende konference i juni i Washington DC har jeg her samlet nogle billeder fra det seneste face-to-face bestyrelsesmøde i januar, hvor ODTUGs bestyrelse planlagde konferencen og brugergruppens andre fremtidige aktiviteter.

Øverste billede viser lokalet, hvor vi tilbragte det meste af vores tid i to dage i januar 2006. En vigtig del af mødet handler om at få defineret ODTUGs identitet. Hvad skal ODTUG stå for? Hvem skal vi arbejde sammen med? Hvem skal vi IKKE samarbejde med? Hvilke medlemmer skal vi tiltrække? Osv osv.

Herunder ses jeg i diskussion med ODTUGs bogholder Jeff Canclini, som forøvrigt er testpilot for det amerikanske luftvåben. I baggrunden hænger nogle plakater med visioner, som vi definerede ved bestyrelsesmødet New Orleans i marts 2005. Hakkerne angiver de visioner, som blev nået, mens tidligere bestyrelsesmedlem Paul Dorseys vision om at ODTUG skulle merge med IOUG er streget ud, da han stod meget alene med den holdning. Umiddelbart efter en afstemning om emnet trak han sig ud af bestyrelsen under meget dramatiske omstændigheder.

I 2006 vil ODTUG lancere en helt ny hjemmeside. Som det ses på det tredie billede, blev det nye design gennemgået ved bestyrelsesmødet. Her mødte selve Scott Spendolini (se billede 4) op, og demonstrerede mulighederne. Det er ham, som står bag implementeringen af den nye hjemmeside, der udvikles i HTML DB. Scott er tidligere Oracle product manager for HTML DB. ODTUGs hjemmeside er hans første opgave som selvstændig.

På modstående side har jeg valgt nogle billeder af

konferenceområderne. For lige at slå tonen an skal jeg minde om at Washington DC er USAs hovedstad, hvor det hvide hus ligger overfor den enorme figur af præsident Lincoln, der nedkæmpede sydstaterne og ophævede slaveriet.

Selve hotellet hedder Wardman Park Mariott og rummer såvel store som små foredragssale. Billede to herunder viser salen, hvor keynote-talerne kommer til at foregå.

rencedage. Sidst jeg var i Washington DC i juni var der 40 grader varmt...

ODTUG har valgt at fortsætte med at holde sin egen konference med ca 600 deltagere, selv om vi i de seneste par år er blevet presset voldsomt af Oracle til at slå vores konference sammen med IOUGs og AOUGs. Den første "fælles" konference bliver afholdt den 23. - 27. april i Nashville under titlen "Collaborate 06". Jeg vil tro, at der vil være mindst 10 gange så mange deltagere på Collaborate 06.

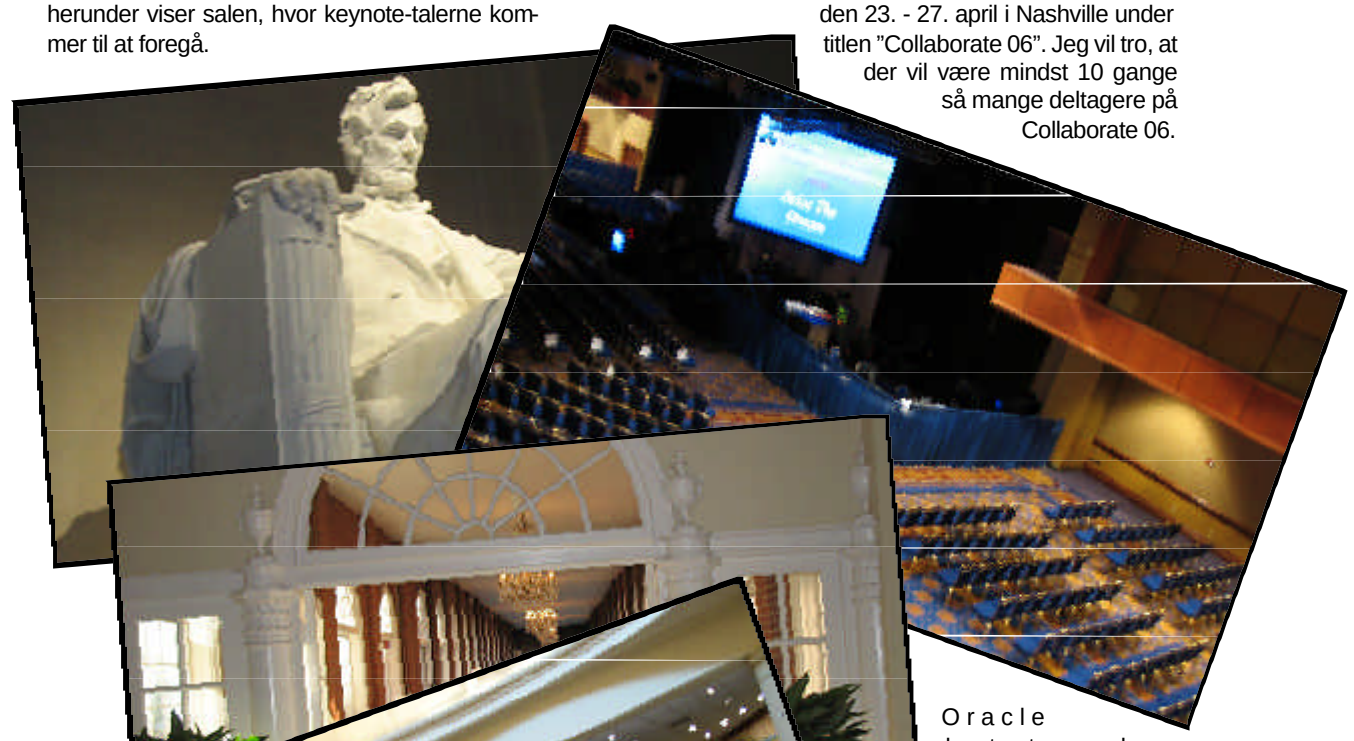

Oracle har truet os med ikke at ville sende product managers til ODTUG-konferencen, hvis vi fortsatte med at afholde vores egen konference, men vi ved, at det er der ikke enighed om internt hos Oracle, så lad os se, hvad det ender med. I år har alle de sædvanlige Oracle-folk i hvert fald meldt deres ankomst.

ODTUG er igen blevet inviteret med til Collaboration 07, som vil bliver afholdt i Las Vegas, men vi regner stadig med at holde vores egen konference

Hotellet er bygget med en del fløje og niveauforskelle, som dels giver charme, men nok også kan give visse udfordringer, når man skal finde rundt blandt de ca ti foredragsale, som hele tiden vil kæmpe om deltagernes opmærksomhed.

Selv sidst i januar var vejret meget pænt, så man skal nok forvente nogle ret varme konfe-

DEVELOPERS,
LOOK NO FURTHER

ODTUG 2006

KALEIDOSCOPE

CLARIFYING THE DEVELOPER'S CHANGING UNIVERSE

Highlights:

- Web Architecture Symposium Saturday, June 17
- Hands-on Training
- BI Summit
- More Than 100 Technical Presentations
- Built-in Time to Meet and Talk with the Experts

June 17-21, 2006
Wardman Park Marriott Hotel
Washington, D.C.

To obtain the special conference rate of $175/night at the Wardman Park Marriott Hotel, call 1-800-Marriott by May 12, 2006.
Attention government employees: Government rates are available, but conference room rates are lower.

TOPICS:

Application Development • Traditional Tools • Business Intelligence • Methodology
Development DBA • Professional Development • Non-Oracle Development Tools

www.ODTUG.com

2006.03.23

Oracle lancerer Business Intelligence

Charles Phillips annoncerede at Business Intelligence er et nyt "ben" for Oracle, som de regner med at kunne komme til at tjene over USD 1 mia på.

Med dette produkt er Oracle ude at konkurrere med firmaer som Microsoft, Cognos og Business Objects.

Produktet sælges i tre moduler: Standard Edition One til USD 25.000 pr server, Standard Edition til USD 400 pr bruger og Enterprise Edition tilUSD 1.500 pr bruger.

2006.03.17

Oracle vinder patentsag

MangoSoft lagde i 2002 sag an mod Oracle med anklager om at Oracle havde krænket MangoSofts patent om memorydeling i forbindelse med Oracles database-cluster software. MangoSoft kræver over USD 500 mill i erstatning.

En foreløbig dom afgivet ved US District Cout i New Hampshire giver Oracle ret i at MangoSofts patent ikke er blevet krænket.

2006.03.13

SQL Developer frigivet

Oracles "Toad-dræber", som tidligere gik under navnet "Raptor", er blevet frigivet officielt under navnet "SQL Developer".

2006.03.09

Staten Wisconsin dropper Oracle

En investering på USD 2,1 mill i Oracle Collaboration Suite, der skulle have nedbragt antallet af servere fra 222 til 19, er blevet opgivet efter tre måneders forsøg på at flytte administrationens 1200 mail-konti.

Erfaringerne inkluderede at mails forsvandt fra "Sendt"-folderen, mødeindkaldelser dukkede ikke op i kalenderen og distributionslister distribuerede ikke.

Projektet har været planlagt igennem de sidste tre år.

2006.03.02

Oracle laver erhvervs-Google

Larry Ellison præsenterede Oracle Secure Enterprise Search 10g under sin key note tale ved Oracle OpenWorld Tokyo. Han kaldte Oracle Secure Enterprise Search 10g Oracles vigtigste produktannoncering i mange år.

Oracle Secure Enterprise Search 10g ligner Google til forveksling. Den væsentligste forskel er at SES 10g har et login-billede, som tillader brugere at søge i kalendere, mailbokse, lokale filer osv.

2006.03.01

IT-budgetterne fastholdes i 2006

Ifølge 2006 The Controller's Report vil IT-budgetterne i 2006 svare nogenlunde til 2005.

Fordelingen vil ligeledes være den samme som i 2005 med 30% til løn, 22% til teknologi og hardware, 18% til applikationer, 10% til konsulenter og outsourcing.

46% af IT-budgetterne kommer til at stige med i gennemsnit 19%, mens 15% af budgetterne vil falde med 17% i gennemsnit.

2006.02.28

Ekstraordinær sikkerhedspatch

Oracle har netop annonceret en ny patch til E-Business Suite, som inkluderer et antal sikkerhedsforbedringer. Det specielle ved denne patch er at den kommer næsten to måneder før Oracles næste planlagte patch.

Oracle har i et års tid udsendt kvartalsvise patch frem for at udsende patches løbende.

2006.02.24

IT-folk skifter jobs i 2006

Ifølge en CareerBuilder.com-undersøgelse er 20% af alle IT-folk utilfredse med deres nuværende stilling og 1/3 planlægger at skifte job i 2006.

De primære kritikpunkter er 1) For højt arbejdspres, 2) Dårlig løn og 3) Dårlig ledelse.

61% mener at arbejdspresset er steget i løbet af de seneste seks måneder.

Kun 20% har i 2005 fået mere end 5% i lønforhøjelse. 2/3 fik ikke nogen bonus.

2006.02.24

Oracle vil tage markedsandele fra SAP

Senior Vice President John Wookey udtalte til Financial Times Deutchland at Oracle vil genoptage kampen om SAPs markedsandele.

Oracle har kunne se det internet i de seneste to kvartaler. I løbet af dette og næste kvartal vil det blive synligt for eksterne analytikere.

2006.02.23

Oracle anklager Microsoft

IBM, Oracle, Nokia, Sun Microsystems m fl har via European Committee for Interoperable Systems (ECIS) lagt sag an imod Microsoft angående deres nye version af Windows-operativsystemet, som kommer til at hedde Vista, og office-pakken, der kommer til at hedde Office 12.

Kritikken falder på følgende punkter:

- Office 12 er ikke kompatibelt med andre office-produkter som OpenOffice og StarOffice.
- Vista udskifter HTML med XAML, som hævdes at give Microsoft en urimelig fordel ift konkurrenter.
- Vista kommer indbygget med Microsofts DRM-teknologi. Det svarer til den sag om bundling af Windows og Media Player, som Microsoft blev dømt for i 2004.

DBVision

Nyt kursuskatalog på

www.dbvision.dk

eller ring på 35 36 21 44

Er du enig?

Af Marc de Oliveira

Jeg bestyrer denne lille sektion, hvor jeg i hvert nummer af OracleEkspert diskuterer et emne, som I alle gerne må give mig jeres mening om ved at afgive jeres stemme på www.OracleEkspert.dk (under menupunktet Meningsmålingen). Der uddeles en præmie for en af de mest interessante kommentarer. Ideer til spørgsmål, som det kunne være spændende at høre andres mening om er velkomne til Marc@OracleEkspert.dk.

Påstanden som blev fremlagt i OracleEkspert nr 34 lød:

"Podcasting et godt medie til Oracle-stof"

Opinion: 3.8

(på en skala fra 1-5, hvor 1 betyder "helt uenig" og 5 betyder "helt enig")

OE35 Vinder: Mette Stephansen, PFA

Selv om Mette Stephansen fra PFA vandt Meningsmålingen i OracleEkspert nr 31, så var jeg nødt til at kåre hende igen i dette nummer. Hendes kommentar om Oracle og podcasting er til inspiration helt ind på Larry Ellisons kontor:

"Jeg synes da at podcasts selvfølgelig skal være en del af den nye Oracle 11 - altså INDBYGGET i databasen, så den selv kan sende!!! Så kan jeg sidde derhjemme og nyde solen. Og så afbryder de lige musikken med en vigtig meddelelse (jeg går da ud fra at vi kan lade den gå ind over både autoradio og det hele, ikk?) - at SKRIIIG, min SQL er gået ned med en 'snapshot too small'". Ideen er hermed givet videre...

Også Henrik Verup fra Sigma Business Intelligence kigger lidt ud i fremtiden. Han siger blandt andet at: "En af de store udfordringer med "Always On" er jo netop tilgængelighed, og det løser podcasting på glimrende vis. Nu mangler vi bare at bilfabrikanterne for alvor tager sig sammen og indfører USB-stik, der både kan levere strøm til opladning og modtage audio-signal til bilradioen."

Det er nok ikke så overraskende at jeg er helt enig i denne måneds påstand. Podcasting er særligt velegnet til smalle emner, da det er en forholdsvis billig kommunikationsform, som stort set alle både kan forbruge og producere. Podcasting er derfor ikke blot en modedille, som vil blive erstattet af den næste smarte ting om noget tid. Podcasting er en unik kommunikationsform på linie med blogging og videocast.

Mange mener at videocasting vil erstattet podcasting, men det tror jeg bestemt ikke. Videocasting er et meget spændende område, men det har to helt centrale problemer, som gør at det aldrig vil blive et lige så stort medie som traditionel podcasting: 1) Det er væsentligt dyrere at producere video end audio. Det kræver flere folk og flere ressourcer at lave video end audio, og 2) Det er meget mere begrænset hvornår man kan forbruge video i forhold til audio. For det første er der meget færre transportable video-afspillere end audio-afspillere (men det vil muligvis udjævne sig med tiden), og for det andet (og det er den vigtigste grund) så kan man ikke se video mens man kører bil eller på cykel, mens man maler plankeværk eller vasker gulv, mens man programmerer eller overvåger servere osv osv. Video-mediet kræver meget mere opmærksomhed fra forbrugeren end audio-mediet, så muligheden for at forbruge video vil altid være væsentlig mere begrænset end muligheden for at forbruge audio.

Påstanden til meningsmålingen i OracleEkspert nummer 36 lyder:

"Det er vigtigt for Oracle at dominere open source markedet"

Oracle har i den seneste tid interesseret sig en del for open source database- og middleware-selskaber. Hvordan kan det være? Er der penge i open source? Er det fremtidens forretningsmodel?

Skynd dig ind at stemme på www.OracleEkspert.dk!

Du kan vinde et OracleEkspert-krus

NØRGAARD'S NEWS

Af Mogens Nørgaard

Dear reader,

Welcome to the Nørgaard's News column. I'll talk about a new discipline called Anti-tuning, spend some time on the topics of Windows vs. Linux and fencing. Finally, a word to our sponsors: Oracle Support is not getting better.

ANTI-TUNING

I'd like to make some additional comments on two of the topics I talked about in the last column: Revolutionary IT and RAC nodes booting when CPU-saturated.

'Fencing' is the term used in clustered environments for situations where a node cannot reach another node and therefore must assume that it is dying, dead, or at least unreachable.

This is a problem in all environments, but because Oracle has implemented the important background processes for RAC in user-mode on Windows and Linux, and not in kernel-mode as on UNIX and VMS, it is more likely to happen on these two platforms, since there's no guaranteed way of ensuring enough spare CPU cycles for the necessary RAC overhead (for synchronisation and coherency purposes).

So on Windows and Linux the non-reachable node is basically executed (it can be described in lyrical terms as 'shot in the head' - SITH).

Unfortunately, this means that no trace/log is left behind. The node dies ungracefully. This probably explains most of the cases where a RAC node has died spontaneously on you without leaving any hints as to why.

This doesn't have to happen just because the CPU's of one node reaches 100% for a short duration, but this saturation will certainly slow down the operation of RAC itself, and eventually it will lead to one or more nodes dying ungracefully.

So a few OakTable members (Jeff Needham, James Morle, Kevin Closson and myself) have been asking ourselves why this doesn't happen all the time?

We think we know the reason: People have successfully implemented RAID-5!

Well, to be more specific: Since a processor like the Opteron is basically a watershed technology, it takes SO many disks to match the IO throughput that few people have it. So most systems become IO-bound.

Even lesser processors than the Opteron can have this effect.

The fast, nice 15K disks today can deliver an average of only 128 IO's per second. That's not impressive at all, is it? Now, if a single Opteron dual-core can kick off 16,000 IOPS (real IO's, sustained), then few people have the number of disks to match it.

So we have always wanted our systems to be CPU-bound, right? Now that we really can have it for a cheap price (a fully-loaded DL-585 with four dual-core Opterons and 32 GB of RAM has a list price of only $25,000), we have naturally introduced IO as a bottleneck.

With RAC systems on Windows and Linux, because of the user-mode madness, we really, really want to avoid using our CPU's to their fullest.

That's IT irony for you.

Because Linus Torvalds (for whatever reason – it's really strange and, uhm, stupid) doesn't personally like the idea of introducing a real scheduler in Linux (the debate has been raging for years) there is no way of guaranteeing spare CPU cycles for, say, RAC.

(I believe there's some special edition or whatever of Windows that allows you to reserve a part of the CPU cycles, so there might be hope there. I haven't seen it verified.)

So the only logical thing to do appears to be this: ensure that the system is IO-bound!

IO-bound is not good. CPU-bound is good. But we must now strive to make our systems IO-bound.

Fortunately that seems to have been the case in most places where they have implemented new IO-sub-systems.

Particularly RAID-5 SAN's have therefore actually been saving cluster nodes everywhere from dying ungracefully, because they're never allowed to work at full capacity thanks to inadequate number and setup of disks.

I ALWAYS knew there was something good to be said about RAID-F implementations, apart from giving us lots of performance and availability problems to solve, of course.

Here it is, ladies and gentlemen: It will do the necessary anti-tuning that will keep your Linux-cluster more alive.

This Anti-tuning will also help you in non-RAC environments where you can now suddenly run many batch jobs per CPU without maxing out the CPU's. That was meant as a terrible joke, by the way…

Of course there are other means of anti-tuning: Very small Oracle caches (log buffer, buffer cache), etc., but nothing beats a lousy IO sub-system.

Anti-tuning will become big business, mark my words. There will be positive words invented to describe it (can't have the A-word – sounds negative), and nonsense will be written into papers and even books on the subject.

Here's a line from the Dire Straits song 'Industrial Disease':

Sociologists invent words that mean 'Industrial Disease'

In some ways it seems to me that certain authors and consultants have actually been recommending this in their books and articles for years. I hope it was planned.

ONE FURTHER NOTE ON FENCING

Three years ago I wrote a paper titled You Probably Don't Need RAC (YPDNR). Three years is a long time in our industry, so I'm currently revising it thoroughly and will probably publish YPDNR II in a month or so.

From a Dutch guy who has recently implemented a number of RAC or High Availability systems (notice the or), I received this note after he had read my alpha draft of the paper ('STONITH' is my acronym for Shoot The Other Node In The Head):

The STONITH action is a less wanted action. If you use a proper clustering file system (read: Polyserve) this system can reach the fibre channel switches and disable ports. This is a better way to fence out a system because you don't have to start the node again (with the undesired behaviour you might get). To my knowledge CRS is not capable of this means of io fencing. So much for Oracle CRS on Linux...

In other words: It would appear that the Polyserve Cluster File System is implemented in kernel mode, and therefore it can "catch" the kill node commands and either bring the node down gracefully (leaving a log/trace) or even make sure the node doesn't die at all.

I don't know if it's only Polyserve CFS that can do this (haven't asked Veritas yet), and it's not a cure for the larger user-mode madness problem in Oracle, but it could help a lot. It's certainly worth checking out and worth testing before deciding.

CRS is Oracle's Cluster Ready Services, which has been renamed Clusterware in 10gR2. It is really good in 10gR2, by the way, and since it's both supported and documented by Oracle, I strongly advise you to use 10gR2 clusterware on all 10gR1 RAC installations.

It's not possible to use 10gR2 Clusterware on 9i RAC. It could be argued that 9i was really version 1.0 of RAC, so get off it anyway.

Windows or Linux?

Some of you may know or remember Morten Egan, the young guy from Miracle whom I interviewed about iAS a year or two ago. A nerd if ever there was one.

So he and I were recently conducting some seminars around a benchmark made by Morten and another Miracle guy involving a 14-node RAC setup on 64-bit Windows using 10gR2. Opteron, of course.

The benchmark went alright in itself. No problems there.

But when Morten and I were asked by the audiences whether to choose Windows or Linux, we surprised each other by answering "Windows".

This is not easily said, of course. I, for one, hate monopolies, and I really don't like certain things done by Mr Gates and his henchmen. Morten is a true nerd, who has had his second life among the open source guys for many years (and partly as a result of this he worked for Oracle's Wim Coekaerts for a couple of years, too, on Project Raw Iron).

But from a business perspective we would now both recommend Windows. It's just as stable as Linux. Things run fast enough, and probably most things run as fast as on Linux.

Windows is probably more expensive. They are likely to take full advantage of their near-monopoly, and I will truly hate them for that. Yes, I know that Larry has done the same. I just hate monopolies that abuse their position, and they all end up doing that – private or public, it doesn't matter.

The real reason for our recommendation is that smaller companies will find it harder and harder to find qualified personnel to run Linux-installations. Windows is known by more young people, and it doesn't seem so difficult to use.

Linux is confusing and some of the code is sub-optimal (lousy, really) – and there's no guarantee that they will ever have, say, a scheduler class – who should donate the code or the coding effort? Customers have no power in the Linux world.

If Microsoft customers really want a good scheduling class they will get it. Customers have some power in that respect.

Oracle Support?

Yes, I know, it's easy to shoot at Oracle Support. No, I don't like doing it. I spent 10 good years there. A lot of Miracle's employees worked there. Big pieces of our hearts are still there.

But it is really bad these days, I'm sorry to say. No amount of extra analysts in countries that don't speak good English, no number of emergency procedures, no new search facilities on Metalink, and, please God, no more presentations from Support people about how to work efficiently with and get the full benefit of Support, can hide it.

It's really, really bad. It's not getting better.

When Oracle Support did the thing that all its competitors had already done in the late 90's it wasn't for lack of a good business. Margins where often in the 80's. It was greed for even higher margins.

So the MBA boys did the usual "plagiarise with pride" thing, called it something unique, and ended up with the usual four elements:

1. Global (i.e. no ownership)
2. Internet-based (i.e. non-personal)
3. Low-cost countries (i.e. communications problems)
4. Self-service (i.e. no service)

Yet, to everybody's great, utter surprise, customers are less satisfied. Really. After a while, where the Excel Video Game showed all the right results and projections, customers started to complain.

Look in a typical TAR these days. The answers from Support are often so dumb that the author of the TAR (the customer) asks again and again whether the analyst has actually READ the TAR. To that there is never a direct reply, since saying no in certain (most Asian) cultures is seen as losing head, and therefore possible job.

Instead customers with decades of experience are asked by people with days of experience to do things like "Could you move the iAS from Windows to Sun, because the OID works there?" and "Have you tried to set spin_count" and things in general that leave people speechless.

A couple of our customers asked us for help recently with two TAR's that basically was about the same problem, namely a massive latching issue in Oracle on 10gR2 on dual-core processors – be they Itanic or Opteron.

When one of the clients' 700 engineers had not been able to schedule the building of a couple of factories for clients for two or three business days (no progress on the TAR, of course) they fell back to their old system, which was using single-core Itanic, and things worked very slowly, but they worked.

Guess what the Support analyst did? Lowered the pri-

ority to 2, of course, since the production system wasn't down anymore! That kind of thinking could be turned into a Monty Python sketch.

In both cases we pointed out the known issue with dual-core and Oracle doing excessive latching in the TAR's. No acknowledgement. Just more of the same old same old: Have you tried… Could you provide a trace…

Fortunately we knew the necessary bug numbers and could request the relevant one-offs, backports, etc.

It's bad. It's getting worse.

And Support is now in the hamster wheel like all its competitors because they removed the one thing that makes price irrelevant: Superior, personal service.

The next thing is to replace Indians with Chinese or Indonesians, West Europeans with Romanians. There are some low-cost areas in the US, too, like Oklahoma.

I really like that we're creating real jobs out there. I really do. We should have done that as a society many decades ago. We should also open our markets fully to them to give them a chance (especially the Africans, I'd say).

But Support analysts should be local, specialised, highly experienced, and more. Why local? Because they can then communicate efficiently, spot local trends, know specifics about their customers, be aware of cultural issues… in other words make a difference.

It's tempting to ask if the organisation should be renamed Support Oracle, since that seems to be what most just moderately complex TAR's end up with: The customer giving the analyst(s) hints and tips and finally working out what to do either to fix or work around the problem. That's On-The-Job training for you right there.

Oracle un-bundled Support and Update Rights in late 1999. They recently bundled it again, which means you have no choice but to sign up for both. Another sign of a monopoly in trouble.

Dear Oracle: Start up slowly. Find the few good analysts that are left, and get them used to the whole concept of working closely with a local customer again.

Then make sure some MBA boys discover the huge efficiency gains to be had from closing TAR's in real time instead of having them linger in no-ownership queues for days or weeks. They will know how to come up with a snappy title for the program.

Good consultants will always suggest that the customer does exactly the opposite of what he's currently doing. I think that's what I'm suggesting that Oracle Support's own MBA boys (they have a load of them, both in EMEA land and in the US) advise their mother company to do.

Damhuskroens historie går helt tilbage til første fjerdedel af 1600-tallet, hvor Christian IV lod et beskedent vagthus opføre ved dæmningen over Damhussøen.

Vagthuset var bolig for en opsynsmand, der skulle føre tilsyn med slusen ved dæmningen.

Første gang vi hører at der bliver drevet kro ved vagthuset er i 1682, hvor der i en beskrivelse af Rødovre står:" Ved bemeldte by ligger et Krohus "Langvejs Damhus".

Det var åbenbart så som så med bevillingen og det lykkedes først så sent som i 1757 at få kongeligt privilegium på at holde kro efter en hård kamp med myndighederne og ejerne af de priviligerede kroer længere ude ad Roskildevej.

Ved folketællingen i 1834 boede der foruden kroejeren og hans kone 10 tjenestefolk i Damhuskroen. Hertil kom Bomhuset, der lå lige øst for Damhuskroen. Her boede der en bomhusforpagter, dennes kone og deres to børn samt et "tjenestetyende". Det var et helt lille samfund, der lå ved dette hjørne af Damhussøen.

Efter anlæggelsen af jernbanen fra København til Roskilde i 1847 gik landevejstrafikken tilbage og det gjorde det også for Damhuskroen. Værelserne blev stort set ikke udlejet, og kroen blev mest besøgt af prangere, studedrivere og torvebønder, der fik sig en tår øl på vejen.

Mod slutningen af 1800-tallet fik kroen en renæssance. Det blev moderne for københavnere at rejse derud og nyde landlivets glæder ved bredden af Damhussøen, de begyndte at ligge på landet på Damhuskroen og lejede sig ind på kroens værelser på førstesalen. Forfatterne Holger Drachmann og Viggo Stuckenberg hørte blandt andre til kroens gæster.

Senere gik det igen tilbage for kroen. Omkring 1930 blev den beskrevet som en temmelig lurvet knejpe med rejselade og slyngelstue og et tvivlsomt publikum. Imidlertid fik nye ejere rettet op på stedets renommé, bl.a. familien Stephansen, der overtog kroen i 1936. Fra den tid kom kroen til at lægge lokaler til mange foreningsarrangementer, baller, dilettantforestillinger, juletræsfester m.m.

I de senere år blev Damhuskroen kendt for sine enkeballer om tirsdagen. De er nu flyttet til torsdagen og hedder Ladies Night. Kroen er også kendt for sine julefrokostarrangementer og for skiftende optræden af skuespillere og entertainere som f.eks. Monrad og Rislund.

Eventkalenderen

Følg med i de nyeste databaserelaterede arrangementer og opret helt nye arrangementer direkte på vores hjemmeside www.OracleEkspert.dk under punktet Eventkalender.

Læs yderligere oplysninger om de enkelte arrangementer på Eventkalenderens hjemmeside.

Det er gratis for vores abonnenter at oprette arrangementer i Eventkalenderen.

DesWeb SIG-møde om SQL Developer evaluering

Dato: 26. april 2006

Tidspunkt: Kl 13:00 - 16:00

Oracle SQL Developer ("Raptor") evaluering. Oracle SQL Developer er Oracles bud på en TOAD-lignende værktøj. Indlægget holdes af Kim Alex Olsen

Kontakt: Tilmelding skal ske til .

Sted: Oracle Danmark, Ballerup

Pris: Gratis

Kursus: Objektorienteret Analyse

Dato: 1. maj 2006

Tidspunkt: 1.-3. maj 2006

Dette kursus giver deltagerne en indføring i den Objektorienterede udviklingsmetode og de tilhørende analyseteknikker gennem teori og øvelser

Kontakt: +45 2627 9991

Sted: Fensmarkgade 3, 2200 København N

Pris: DKK 9.450 (rabat på DKK 500 til OracleEkspert-abonnenter)

Optimising Oracle—Performance by Design med Jonathan Lewis

Dato: 2. maj 2006

Tidspunkt: 2.-4. maj 2006

Oracleeksperten Jonathan Lewis, kommer til Danmark til foråret, hvor han vil afholde en Master Class, et tre-dages seminar om Oracle. Sæt allerede nu kryds i starten af maj og gå ikke glip af muligheden for at møde Jonathan og få lejlighed til at stille ham en masse spørgsmål.

Kontakt: www.MiracleAS.dk

Sted: Gentofte Hotel, Gentoftegade 29, 2820 Gentofte.

Pris: 11.200 / 9.000 kr, med/uden overnatning

Kursus: Avanceret Systemanalyse

Dato: 25. maj 2006

Tidspunkt: 25.-27. maj 2006

Lær teknikkerne til at komme fra en model af hvordan virksomheden fungerer i dag, til en model af hvordan arbejdsgangene bedst optimeres i fremtiden.

Uanset om du er en een-mands udviklingsafdeling, der understøtter en lille virksomhed eller sidder med mange udviklere og laver systemer til en milliard-virksomhed, så har din arbejdsgiver brug for, at du løser virksomhedens problemer på den mest effektive måde.

Kurset er baseret på The Zachman Framework og giver en forståelse for de mekanismer, der bringer udviklingsprojekter frem til det ønskede mål til den planlagte tid, pris og kvalitet.

Kontakt: 2627 9991

Sted: Fensmarkgade 3, 2200 Kbh N

Pris: DKK 10.900 (rabat på DKK 1000 til OracleEkspert-abonnenter)

Scaling Oracle: En tre dages Master Class med James Morle

Dato: 14. juni 2006

Tidspunkt: 14.-16. juni 2006

James Morle, forfatteren til "Scaling Oracle8i", grundlæggeren af Scaleabilities og en af stifterne af The OakTable kommer til Danmark for at afholde sin første tre dages Master Class i København.

Kontakt: www.MiracleAS.dk

Sted: IDA (Ingeniørforeningen i Danmark), Kalvebod Brygge 31-33, 1780 København V

Pris: 9.000 kr

ODTUG-konferencen 2006

Dato: 17. juni 2006

Tidspunkt: 17.-21. juni 2006

Fem dage med over 120 indlæg om udvikling mod Oracle vha Designer, Deveoper, JDeveloper, Warehouse Builder etc.

Kontakt: www.odtug.com

Sted: Marriott Wardman Park Hotel, Washington DC

Pris: ca USD 1.000

OracleEkspert-konferencen 2006

Dato: 30. august 2006

Tidspunkt: kl 9-16

Beskrivelse: 15 Oracle-indlæg i tre parallelle spor.

Kontakt: www.OracleEkspert.dk

Sted: Damhuskroen i København

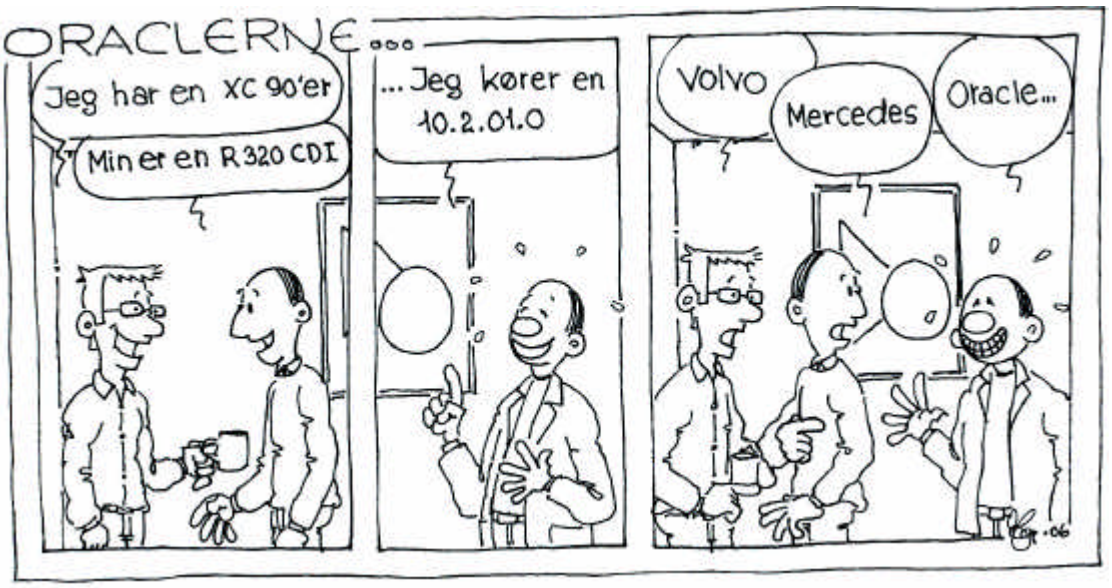

Af Jesper Dyjak

HP offers Oracle Patch Parity

As problems are discovered in released products, Oracle is very diligent about generating patch fixes for the reporting customer and the platform that is involved. The patch fixes are not ported and deployed to other platforms and other customers, until a periodic patch release or unless the customer explicitly requests it. As a result, a patch fix may not be ported and deployed to other platforms for some time, perhaps as much as 4-6 months. This has been a serious issue for customers migrating to the Itanium platform. They are aware of a bug and it's patch fix on their PA platform. However, the bug also exists on their Itanium platform, sometimes for a long time, before the patch fix is ported and deployed to Itanium. To relieve the burden on the customer to request a patch-fix, HP has stepped in to assure patch parity across all HP platforms.

HP has assigned an engineer onsite at Oracle to proactively port patches to HP platforms, with Integrity being a primary target. Now, a patch fix initially introduced on any platform is immediately ported, tested, and deployed to all HP platforms. The latest fixes from Oracle are applied with parity across all HP platforms. So Oracle products will run with highest available quality and equality in near real time on all HP platforms. Customers will not have to wait for the next big patch release to get fixes. This is especially important and of great benefit for customers migrating to Itanium while they are still running Oracle software on other platforms.

For more information

For more information about this program, contact Keith Carlson at keith.carlson@hp.com, or to learn more about the HP and Oracle alliance, visit www.hp.com/go/oracle.

With HP solving the problem of Patch Parity, Oracle products will run with highest available quality and equality in near real time on all HP platforms.

invent