
Challenges in a World of Large Databases and Business Intelligence



Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

Syllabus

- Introduction
- Let's Define Our Terms
- Historic Road Trip
- What I See (far too often)
- Staffing
- Conclusion

Daniel A. Morgan



Oracle ACE Director



Consultant to Harvard University



University of Washington Oracle Instructor, ret.



The Morgan of Morgan's Library on the web



Board Member: Western Washington OUG

- Editor: Microsoft SQL Server 2008 DBA Training
- > 500 RAC clusters built
 - Largest RAC 24 nodes
 - Largest DB storage footprint 1.2PB



Let's Define Our Terms

The Next "Cool" Thing

Big Data Buzz

“Why big data is a big deal”

InfoWorld – 9/1/11

“The challenge—and opportunity—of big data”

McKinsey Quarterly—5/11

“Ten reasons why Big Data will change the travel industry”

Tnooz -8/15/11

“Keeping Afloat in a Sea of 'Big Data”

ITBusinessEdge – 9/6/11

“Getting a Handle on Big Data with Hadoop”

Businessweek-9/7/11

“The promise of Big Data”

Intelligent Utility-8/28/11

ORACLE

T | Copyright © 2011, Oracle and/or its affiliates. All rights reserved.

Jargon



Article [Talk](#)

[Read](#) [Edit](#) [View history](#)

Search



Very large database

From Wikipedia, the free encyclopedia

This article is about Large size databases. For International Conference on Very Large Databases, see [VLDB](#).

A **very large database**, or **VLDB**, is a database that contains an extremely high number of **tuples** (database rows), or occupies an extremely large physical **filesystem** storage space. The most common definition of VLDB is a database that occupies more than 1 **terabyte** or contains several billion rows, although naturally this definition changes over time.



Article [Talk](#)

[Read](#) [Edit](#) [View history](#)

Search



Business intelligence

From Wikipedia, the free encyclopedia

Business intelligence (BI) is the ability of an organization to collect, maintain, and organize knowledge. This produces large amounts of information that can help develop new opportunities. Identifying these opportunities, and implementing an effective strategy, can provide a competitive market advantage and long-term stability.^[1]

BI technologies provide historical, current and predictive views of business operations. Common functions of business intelligence technologies are reporting, [online analytical processing](#), [analytics](#), [data mining](#), [process mining](#), [complex event processing](#), [business performance management](#), [benchmarking](#), [text mining](#), [predictive analytics](#) and [prescriptive analytics](#).

More Jargon

- Big Data
 - "a collection of data sets so large and complex that it becomes difficult to process using on-hand database management tools." (wikipedia)
- Map Reduce
 - "... a programming model for processing large data sets ... typically used to do distributing computing on clusters of computers." (wikipedia)
- Hadoop
 - "... an open-source software framework that supports data-intensive distributed applications, licensed under the Apache v2 license." (wikipedia)
- Large Database
 - Any database whose size is larger than that which a team has comfortably worked with in the past. (morgan)

Define Large

- The number of tuples?
- The size of the storage footprint?
- The size of the queries?
- The size of the DML statements?
- The number of transactions per second?
- The number of simultaneously connected [active] users?

WORKLOAD REPOSITORY report for

| DB Name | DB Id | Instance | Inst num | Startup Time | Release | RAC |
|---------|-----------|----------|----------|-----------------|------------|-----|
| OPM01P | 782247420 | opm01p6 | 6 | 18-Aug-10 21:08 | 11.1.0.7.0 | YES |

| Host Name | Platform | CPUs | Cores | Sockets | Memory (GB) |
|-----------|-------------------------|------|-------|---------|-------------|
| usp9004b | Solaris[tm] OE (64-bit) | 128 | 64 | 16 | 503.16 |

| | Snap Id | Snap Time | Sessions | Cursors/Session |
|-------------|---------|--------------------|----------|-----------------|
| Begin Snap: | 7037 | 15-Sep-10 13:00:18 | 406 | 7.5 |
| End Snap: | 7038 | 15-Sep-10 14:01:28 | 318 | 8.5 |
| Elapsed: | | 61.17 (mins) | | |
| DB Time: | | 6,076.88 (mins) | | |

Big Queries



Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

This Is Page One Of A Single Query

Complete List of SQL Text

Page 2

```
(ontime_count-late_count), 2) < 0 then 0 else 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count,
2)percentage_of_total_by5am, percent(summarised_ontime_count+summarised_late_count, ontime_count+late_count, 2) current_percentage, ontime_count+late_count curr ent_raw_elements,
summarised_ontime_count+summarised_late_count current_summary_elements, summarised_ontime_count+summarised_late_count SUM_ELEMENTS_ATSAM, late_count late raw_elements, ontime_count+late_count available_at_cutoff,
percent(num_of_sumrows, num_of_rows, 2) accuracy from ( select count(distinct case when trunc(datetime_ins, 'mi') > trunc(sysdate)+3/24 then SGSNID end) late_count, count(distinct case when trunc(datetime_ins, 'mi') <
trunc(sysdate)+3/24 then SGSNID end) ontime_count, count (*) num_of_rows from ERICSSON_PSCORE_MM where datetime between trunc(sysdate)-1 and trunc(sysdate)-1/24/60 )rawt, ( select max(datetime)sunday, count(distinct case
when trunc(datetime_ins, 'mi') > trunc(sysdate)+7/24 then SGSNID end) summarised_ontime_count, sum(entries) num_of_sumrows from ERICSSON_PSCORE_MM where datetime = trunc(sysdate)-1 )sum union all SELECT 'ERICSSON_PSCORE_SGSNSTAT' AS TABLENAME, sumday, CASE WHEN percent(summarised_ontime_count,
(ontime_count-late_count), 2) > 100 THEN 100 ELSE percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff, case when 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0
else 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count, 2)percentage_of_total_by5am,
percent(summarised_ontime_count+summarised_late_count, ontime_count+late_count, 2) current_percentage, ontime_count+late_count current_raw_elements, summarised_ontime_count+summarised_late_count
current_summary_elements, summarised_ontime_count+summarised_late_count SUM_ELEMENTS_ATSAM, late_count late raw_elements, ontime_count+late_count available_at_cutoff, percent(num_of_sumrows, num_of_rows, 2)
accuracy from ( select count(distinct case when trunc(datetime_ins, 'mi') > trunc(sysdate)+3/24 then SGSNID end) late_count, count(distinct case when trunc(datetime_ins, 'mi') < trunc(sysdate)+3/24 then SGSNID end) ontime_count,
count (*) num_of_rows from ERICSSON_PSCORE_SGSNSTAT where datetime between trunc(sysdate)-1 and trunc(sysdate)-1/24/60 )rawt, ( select max(datetime)sunday, count(distinct case when trunc(datetime_ins, 'mi') <
trunc(sysdate)+7/24 then SGSNID end) summarised_ontime_count, sum(entries) num_of_sumrows from ERICSSON_PSCORE_SGSNSTAT_DY where datetime = trunc(sysdate)-1 )sum union all SELECT 'ERICSSON_UTRAN_RBS_CARRIER' AS TABLENAME, sumday, CASE WHEN percent(summarised_ontime_count, (ontime_count-late_count),
2) > 100 THEN 100 ELSE percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff, case when 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else
100-percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count, 2)percentage_of_total_by5am,
percent(summarised_ontime_count+summarised_late_count, ontime_count+late_count, 2) current_percentage, ontime_count+late_count current_raw_elements, summarised_ontime_count+summarised_late_count
current_summary_elements, summarised_ontime_count+summarised_late_count SUM_ELEMENTS_ATSAM, late_count late raw_elements, ontime_count+late_count available_at_cutoff, percent(num_of_sumrows, num_of_rows, 2)
accuracy from ( select count(distinct case when trunc(datetime_ins, 'mi') > trunc(sysdate)+3/24 then NODEB end) late_count, count(distinct case when trunc(datetime_ins, 'mi') < trunc(sysdate)+3/24 then NODEB end) ontime_count,
count (*) num_of_rows from ERICSSON_UTRAN_RBS_CARRIER where datetime between trunc(sysdate)-1 and trunc(sysdate)-1/24/60 )rawt, ( select max(datetime)sunday, count(distinct case when trunc(datetime_ins, 'mi') >
trunc(sysdate)+7/24 then NODEB end) summarised_ontime_count, count(distinct case when trunc(datetime_ins, 'mi') < trunc(sysdate)+7/24 then NODEB end) ontime_count, sum(entries) num_of_sumrows from ERICSSON_UTRAN_RBS_CARRIER_DY
where datetime = trunc(sysdate)-1 )sum union all SELECT 'ERICSSON_UTRAN_RBS_EDCHRES' AS TABLENAME, sumday, CASE WHEN percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else
100-percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff, case when 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else
100-percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count, 2)percentage_of_total_by5am,
percent(summarised_ontime_count+summarised_late_count, ontime_count+late_count, 2) current_percentage, ontime_count+late_count current_raw_elements, summarised_ontime_count+summarised_late_count
current_summary_elements, summarised_ontime_count+summarised_late_count SUM_ELEMENTS_ATSAM, late_count late raw_elements, ontime_count+late_count available_at_cutoff, percent(num_of_sumrows, num_of_rows, 2)
accuracy from ( select count(distinct case when trunc(datetime_ins, 'mi') > trunc(sysdate)+3/24 then NODEB end) late_count, count(distinct case when trunc(datetime_ins, 'mi') < trunc(sysdate)+3/24 then NODEB end) ontime_count,
count (*) num_of_rows from ERICSSON_UTRAN_RBS_EDCHRES where datetime between trunc(sysdate)-1 and trunc(sysdate)-1/24/60 )rawt, ( select max(datetime)sunday, count(distinct case when trunc(datetime_ins, 'mi') >
trunc(sysdate)+7/24 then NODEB end) summarised_ontime_count, count(distinct case when trunc(datetime_ins, 'mi') < trunc(sysdate)+7/24 then NODEB end) ontime_count, sum(entries) num_of_sumrows from ERICSSON_UTRAN_RBS_EDCHRES_DY
where datetime = trunc(sysdate)-1 )sum union all SELECT 'ERICSSON_UTRAN_RBS_HSDSCHRES' AS TABLENAME, sumday, CASE WHEN percent(summarised_ontime_count, (ontime_count-late_count), 2) > 100 THEN 100 ELSE
percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff, case when 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else
100-percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count, 2)percentage_of_total_by5am,
percent(summarised_ontime_count+summarised_late_count, ontime_count+late_count, 2) current_percentage, ontime_count+late_count current_raw_elements, summarised_ontime_count+summarised_late_count
current_summary_elements, summarised_ontime_count+summarised_late_count SUM_ELEMENTS_ATSAM, late_count late raw_elements, ontime_count+late_count available_at_cutoff, percent(num_of_sumrows, num_of_rows, 2)
accuracy from ( select count(distinct case when trunc(datetime_ins, 'mi') > trunc(sysdate)+3/24 then NODEB end) late_count, count(distinct case when trunc(datetime_ins, 'mi') < trunc(sysdate)+3/24 then NODEB end) ontime_count,
count (*) num_of_rows from ERICSSON_UTRAN_RBS_HSDSCHRES where datetime between trunc(sysdate)-1 and trunc(sysdate)-1/24/60 )rawt, ( select max(datetime)sunday, count(distinct case when trunc(datetime_ins, 'mi') >
trunc(sysdate)+7/24 then NODEB end) summarised_ontime_count, count(distinct case when trunc(datetime_ins, 'mi') < trunc(sysdate)+7/24 then NODEB end) ontime_count, sum(entries) num_of_sumrows from ERICSSON_UTRAN_RBS_HSDSCHRES_DY
where datetime = trunc(sysdate)-1 )sum union all SELECT 'ERICSSON_UTRAN_RNC_IURLINK' AS TABLENAME, sumday, CASE WHEN percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else
100-percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff, case when 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else
100-percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count, 2)percentage_of_total_by5am,
```


Page 4

```
when trunc(datetime_ins, 'm') > trunc(sysdate)+7/24 then MSC end) summarised_late_count, count(distinct case when trunc(datetime_ins, 'm') < trunc(sysdate)+7/24 then MSC end) summarised_ontime_count, sum(entries) num_of_sumrows from NORTEL_PSCORE.VLR6_DY where datetime = trunc(sysdate)-1 )sumt union all SELECT 'NORTEL_PSCORE.GSCMM' AS TABLENAME, sumday, CASE WHEN percent(summarised_ontime_count, (ontime_count-late_count), 2) > 100 THEN 100 ELSE percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff, case when 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) end percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count, 2)percentage_of_total_by5am, percent(summarised_ontime_count-summarised_late_count, ontime_count+late_count, 2) current_percentage, ontime_count+late_count current_raw_elements, summarised_ontime_count+summarised_late_count current_summary_elements, summarised_ontime_count-summarised_late_count SUM_ELEMENTS_AT5AM, late_count late_raw_elements, ontime_count+late_count available_at_cutoff, percent(num_of_sumrows, num_of_rows, 2) accuracy from ( select count(distinct case when trunc(datetime_ins, 'm') > trunc(sysdate)+3/24 then SGSN end) late_count, count(distinct case when trunc(datetime_ins, 'm') < trunc(sysdate)+3/24 then SGSN end) ontime_count, count(*) num_of_rows from NORTEL_PSCORE.GSCMM where datetime between trunc(sysdate)-1 and trunc(sysdate)-1/24/60 )rawt, ( select max(datetime)sumday, count(distinct case when trunc(datetime_ins, 'm') > trunc(sysdate)+7/24 then SGSN end) summarised_late_count, sum(entries) num_of_sumrows from NORTEL_PSCORE.GSCMM_DY where datetime = trunc(sysdate)-1 )sumt union all SELECT 'NORTEL_PSCORE.GSCM_ACT' AS TABLENAME, sumday, CASE WHEN percent(summarised_ontime_count, (ontime_count-late_count), 2) > 100 THEN 100 ELSE percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff, case when 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) end percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count, 2)percentage_of_total_by5am, percent(summarised_ontime_count+summarised_late_count current_summary_elements, summarised_ontime_count-summarised_late_count SUM_ELEMENTS_AT5AM, late_count late_raw_elements, ontime_count+late_count available_at_cutoff, percent(num_of_sumrows, num_of_rows, 2) accuracy from ( select count(distinct case when trunc(datetime_ins, 'm') > trunc(sysdate)+3/24 then SGSN end) late_count, count(*) num_of_rows from NORTEL_PSCORE.GSCM_ACT where datetime between trunc(sysdate)-1 and trunc(sysdate)-1/24/60 )rawt, ( select max(datetime)sumday, count(distinct case when trunc(datetime_ins, 'm') > trunc(sysdate)+7/24 then SGSN end) summarised_late_count, count(distinct case when trunc(datetime_ins, 'm') < trunc(sysdate)+7/24 then SGSN end) summarised_ontime_count, sum(entries) num_of_sumrows from NORTEL_PSCORE.GSCM_ACT_DY where datetime = trunc(sysdate)-1 )sumt union all SELECT 'NORTEL_PSCORE.GSCSM' AS TABLENAME, sumday, CASE WHEN percent(summarised_ontime_count, (ontime_count-late_count), 2) > 100 THEN 100 ELSE percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) end percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count, 2)percentage_of_total_by5am, percent(summarised_ontime_count+summarised_late_count ontime_count+late_count, 2) current_percentage, ontime_count+late_count current_raw_elements, summarised_ontime_count+summarised_late_count current_summary_elements, summarised_ontime_count-summarised_late_count SUM_ELEMENTS_AT5AM, late_count late_raw_elements, ontime_count+late_count available_at_cutoff, percent(num_of_sumrows, num_of_rows, 2) accuracy from ( select count(distinct case when trunc(datetime_ins, 'm') > trunc(sysdate)+3/24 then SGSN end) ontime_count, count(*) num_of_rows from NORTEL_PSCORE.GSCSM where datetime between trunc(sysdate)-1 and trunc(sysdate)-1/24/60 )rawt, ( select max(datetime)sumday, count(distinct case when trunc(datetime_ins, 'm') > trunc(sysdate)+7/24 then SGSN end) summarised_late_count, count(distinct case when trunc(datetime_ins, 'm') < trunc(sysdate)+7/24 then SGSN end) summarised_ontime_count, sum(entries) num_of_sumrows from NORTEL_PSCORE.GSCSM_DY where datetime = trunc(sysdate)-1 )sumt union all SELECT 'NORTEL_PSCORE.GSDSTATS' AS TABLENAME, sumday, CASE WHEN percent(summarised_ontime_count, (ontime_count-late_count), 2) > 100 THEN 100 ELSE percent(summarised_ontime_count, (ontime_count-late_count), 2) END percent_of_cutoff, case when 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) < 0 then 0 else 100-percent(summarised_ontime_count, (ontime_count-late_count), 2) end percent_of_cutoff_not_summed, percent(summarised_ontime_count, ontime_count+late_count, 2)percentage_of_total_by5am, percent(summarised_ontime_count+summarised_late_count ontime_count+late_count, 2) current_percentage, ontime_count+late_count current_raw_elements, summarised_ontime_count+summarised_late_count current_summary_elements, summarised_ontime_count+summarised_late_count SUM_ELEMENTS_AT5AM, late_count late_raw_elements, ontime_count+late_count available_at_cutoff, percent(num_of_sumrows, num_of_rows, 2) accuracy from ( select count(distinct case when trunc(datetime_ins, 'm') > trunc(sysdate)+3/24 then SGSN end) late_count, count(*) num_of_rows from NORTEL_PSCORE.GSDSTATS where datetime between trunc(sysdate)-1 and trunc(sysdate)-1/24/60 )rawt, ( select max(datetime)sumday, count(distinct case when trunc(datetime_ins, 'm') > trunc(sysdate)+7/24 then SGSN end) summarised_late_count, count(distinct case when trunc(datetime_ins, 'm') < trunc(sysdate)+7/24 then SGSN end) summarised_ontime_count, sum(entries) num_of_sumrows from NORTEL_PSCORE.GSDSTATS_DY where datetime = trunc(sysdate)-1 )sumt
```

Historic Road Trip

Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

TO THE EDITOR

Right conclusion, wrong reason

I AGREE WITH SANDY REED's conclusion [that IT is undergoing a fundamental shift] but not for the reasons the Gartner Group cites. [See From the Editor in Chief, Nov. 16, page 81.]

My experience is that most IS/IT managers are incompetent. Not as managers of people and process, but rather as managers of technology. Many of the projects I have seen at major organizations, such as Boeing and King County (county seat of Seattle) and so on, are doomed due to the lack of technology background of key managers. They run at warp speed down dead ends. They choose the wrong technology. They don't understand the value of standards. They don't understand how to develop testing protocols. They focus wholly on the development side of a project and pay no attention to deployment and long-term maintainability issues.

Needless to say, the tens of millions of dollars being wasted are an incentive for change, and eventually upper management will try to stop the financial hemorrhaging. That, I think, will be the reason for change.

Of course, the good news, as one contractor at Boeing said, is that those wasted dollars end up in our pockets. The way things are currently structured, contractors get paid the same whether a project succeeds or fails. And a badly organized project will last longer (thus pay more) than a well-thought-out one. Now there's something for one of your columnists to write about.

*Daniel A. Morgan
Mercer Island, Wash.*

InfoWorld: August 16, 1999

Low pay deserves low loyalty

TERRY STEYAERT in his letter to the editor had a lot to say. But your caption missed the most important point. [See To the Editor, July 26, page 54.]

Terry wrote, 'The company is still trying to hire senior-level people, but can't find them or won't pay them enough.' Separating the wheat from the chaff, the operative point is, 'won't pay them enough.'

Any company that thinks a senior-level IT person is going to study as much as any other professional, lawyer, or physician and not expect equal compensation didn't get on the elevator when the door opened. I have turned down dozens of job offers from Fortune 500 companies because they expected me to drop my salary requirement to match what they pay their employees.

Why? For what? The big question in my book is: Why would any competent employee stay with a company that treats them, and pays them, poorly?

*Daniel A. Morgan
Mercer Island, Wash.*

Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

Banging code responsibly

I STRONGLY AGREE WITH Don Glenn's comments. [See To The Editor, Jan. 11, page 68.] I have been in IS/IT since the early 1970s and find most people inadequately trained in Boolean logic and the fundamentals, not to mention their refusal to abide by standards or to document their work. Any idiot can learn to bang code — and many do.

We need to stop glorifying the barbarian image and start demanding responsibility. Corporations are not going to keep paying millions of dollars every year for inadequate systems. Sooner or later, top management will see past the smoke and mirrors.

*Daniel A. Morgan
DanM@aeimusic.com*

InfoWorld: February 22, 1999

I STRONGLY DISAGREE with the premise and conclusions in "Origins of a shortfall." There is no shortage of applicants. There is a shortage of companies willing to pay market value.

Take, for example, the company where I am

under contract as a senior Oracle DBA [database administrator] until July or August of this year. They have been looking for six months for a permanent employee to replace me. Many have applied, but none have accepted. Why? Because the company refuses to recognize that senior Oracle DBAs have the same amount of investment in their profession as physicians or attorneys. Anyone who is good expects commensurate compensation. If this company would pay market rates, there would be no shortage.

I turn away lots of job offers — some for lack of willingness to step up to the plate with my hourly rate, some because they are so poorly managed that I know the project is going to be a disaster. Bringing in [international] workers won't solve the problem. ... Unless "solving the problem" is reducing the income level of IS professionals.

Daniel A. Morgan
dmorgan@exesolutions.com

Self Image

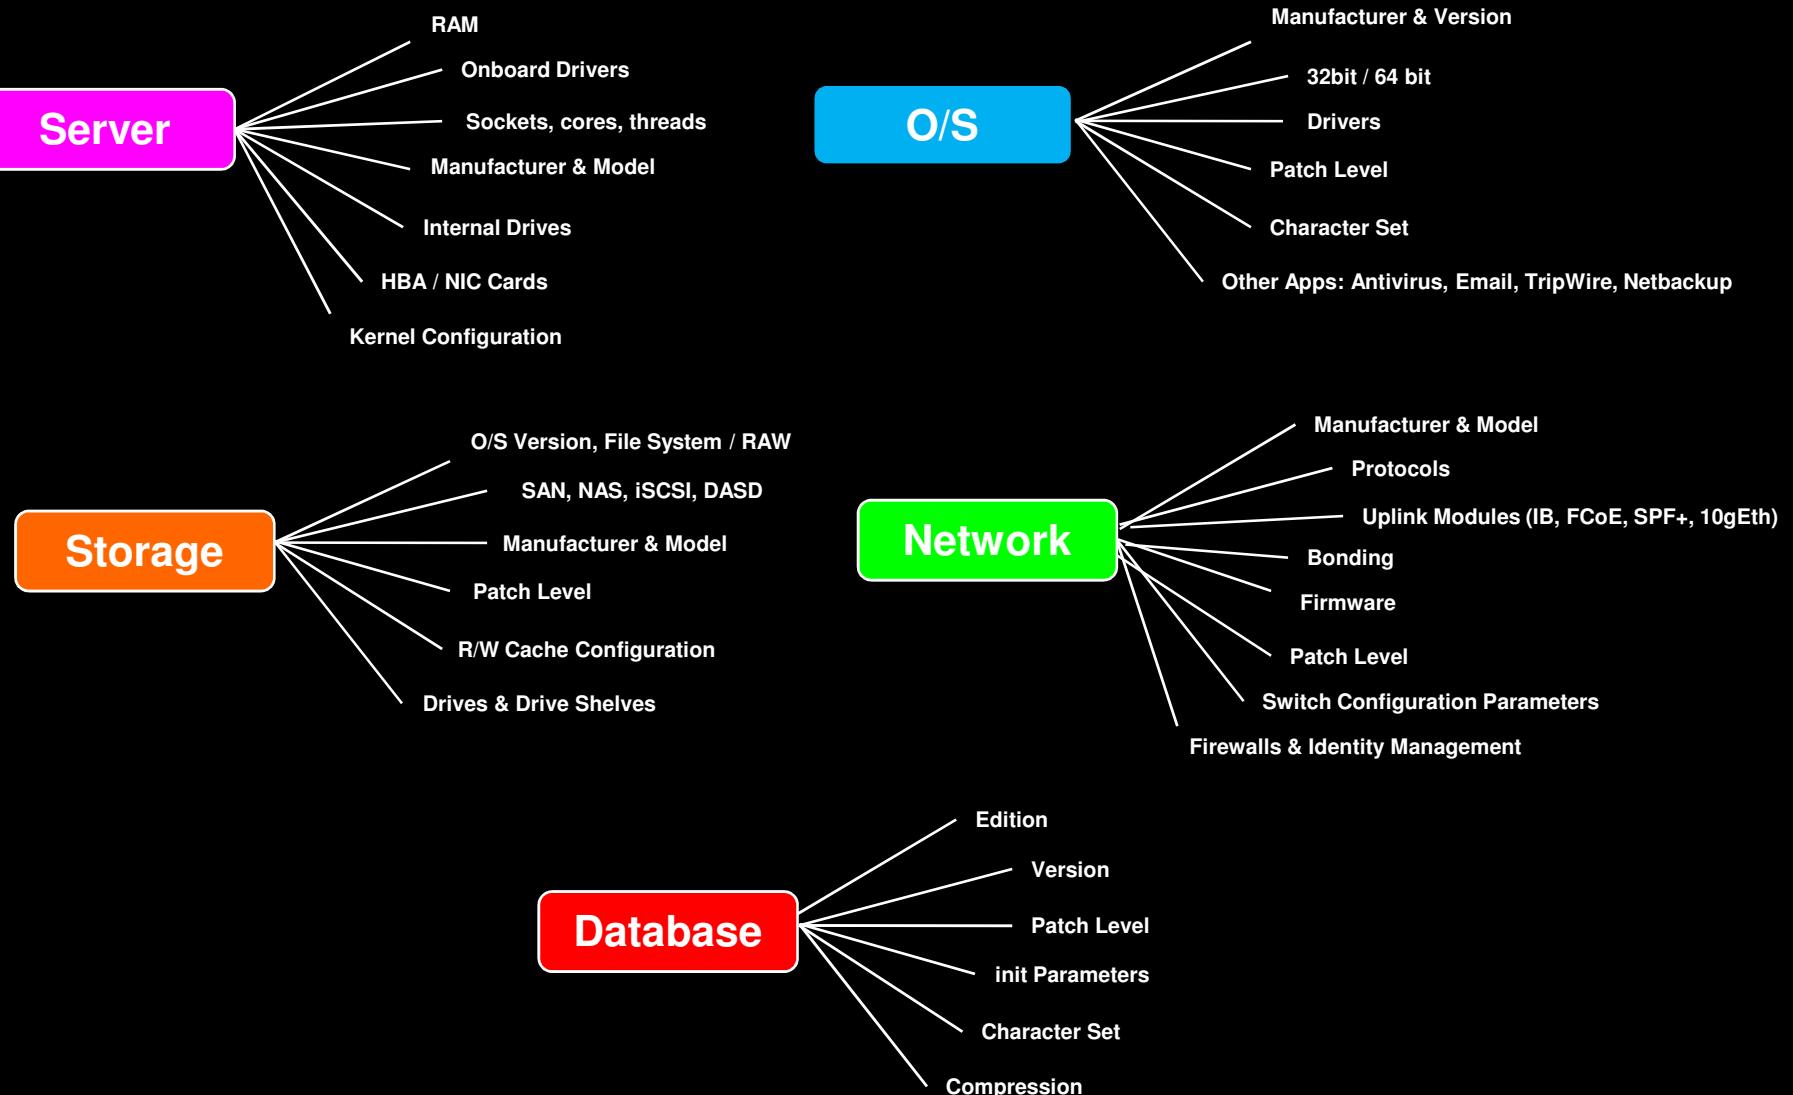
- Most administrators and developers are dinosaurs
- Their managers don't know what they do
- Worse their managers don't know what they should do



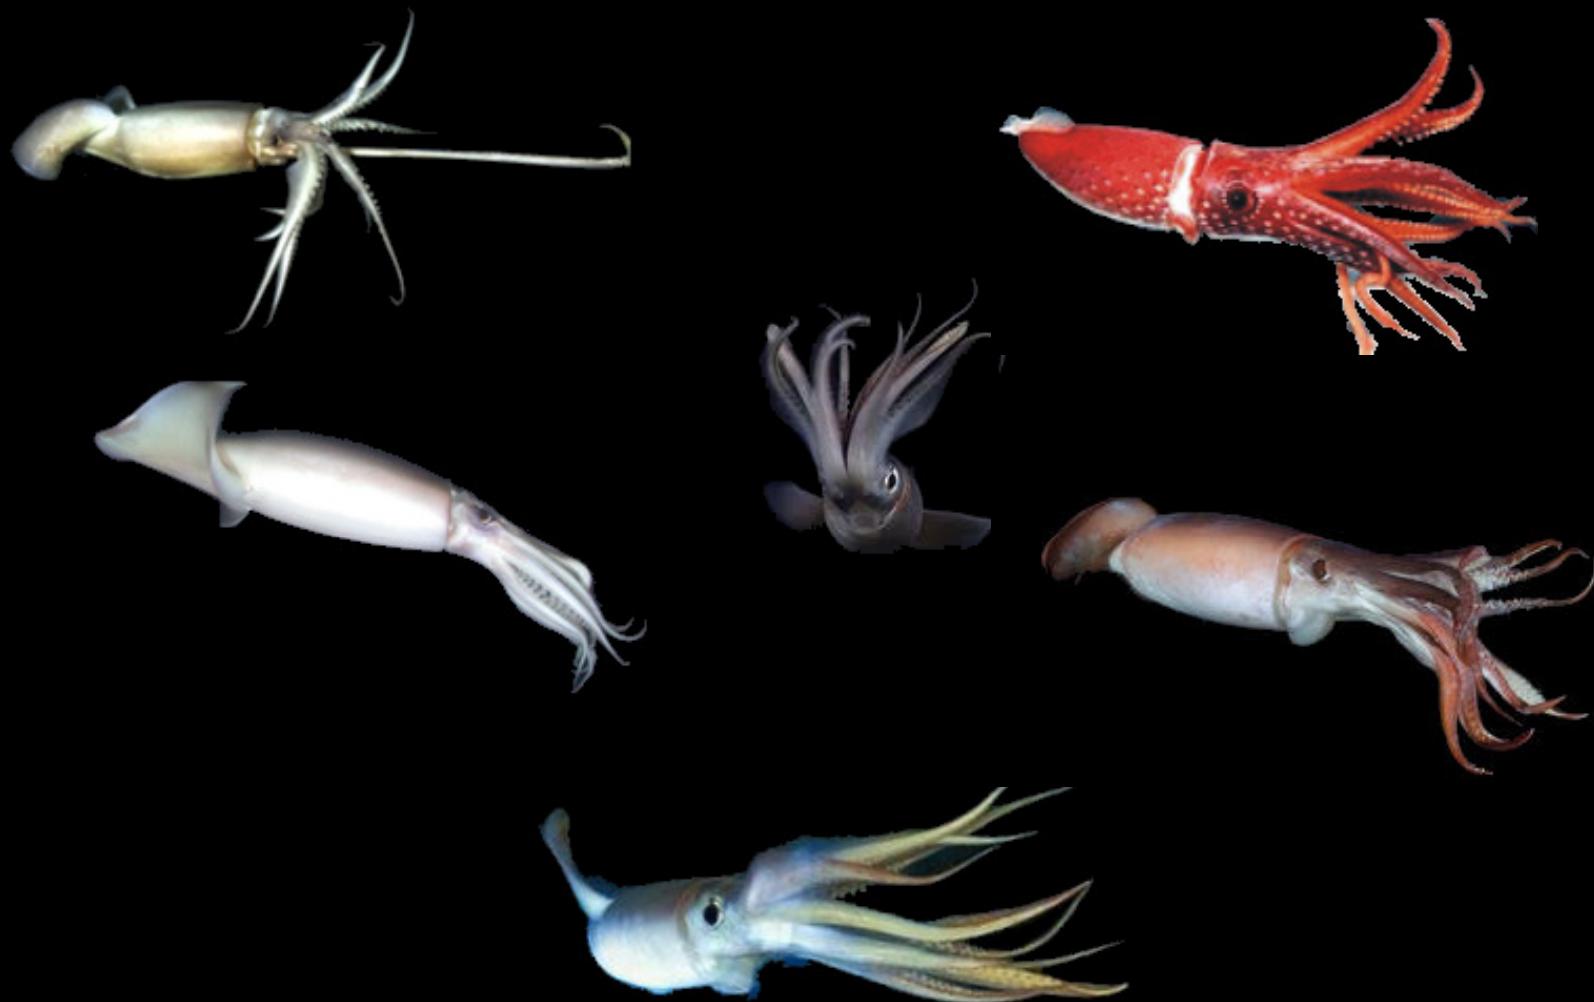
What I See

(far too often)

Static Puzzle Pieces



Animated Puzzle Pieces



Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

As easy to embrace as a barrel of squid



The Wrong Hardware

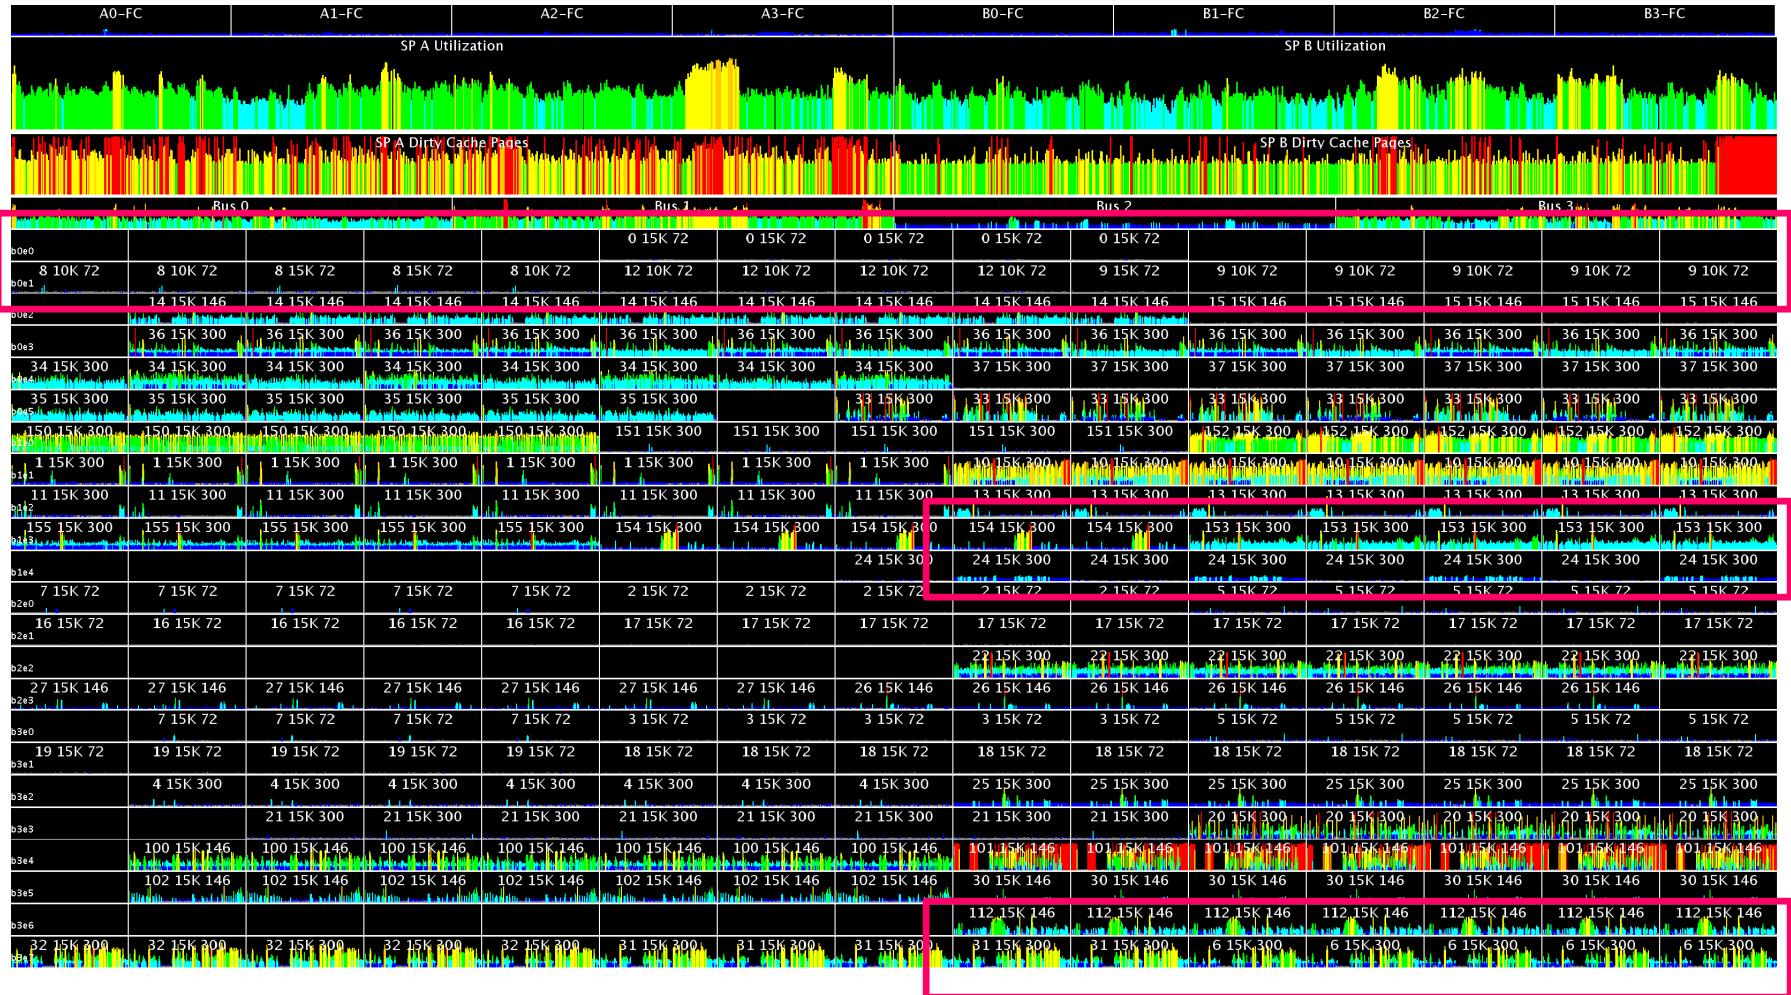


3 Sun M9000s



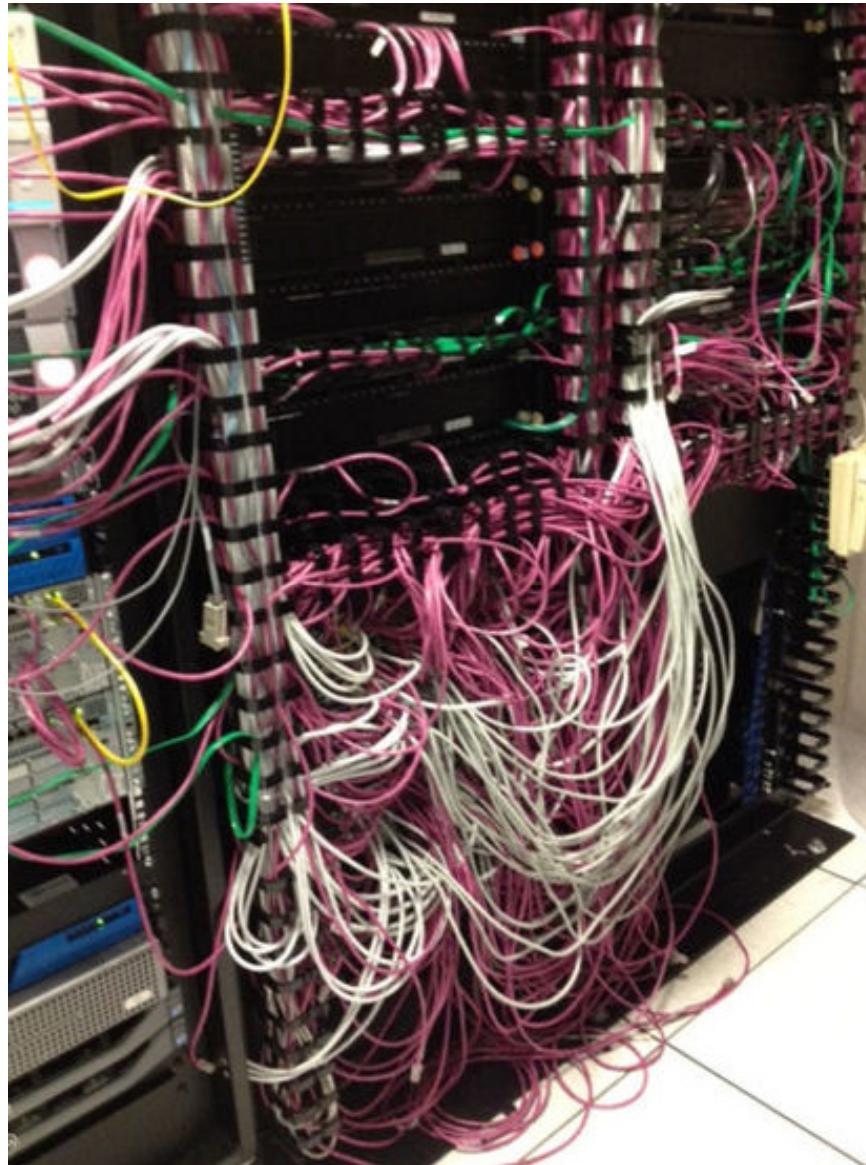
3 EMC VMax SANS

Storage Management?



1. Usage over 50% guarantees any component failure will overwhelm remaining resources producing an outage
2. Metrics coded in yellow are warnings
3. Metrics code in red indicate the database has been forced into a wait state while resources are recovered

Network Management?



Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

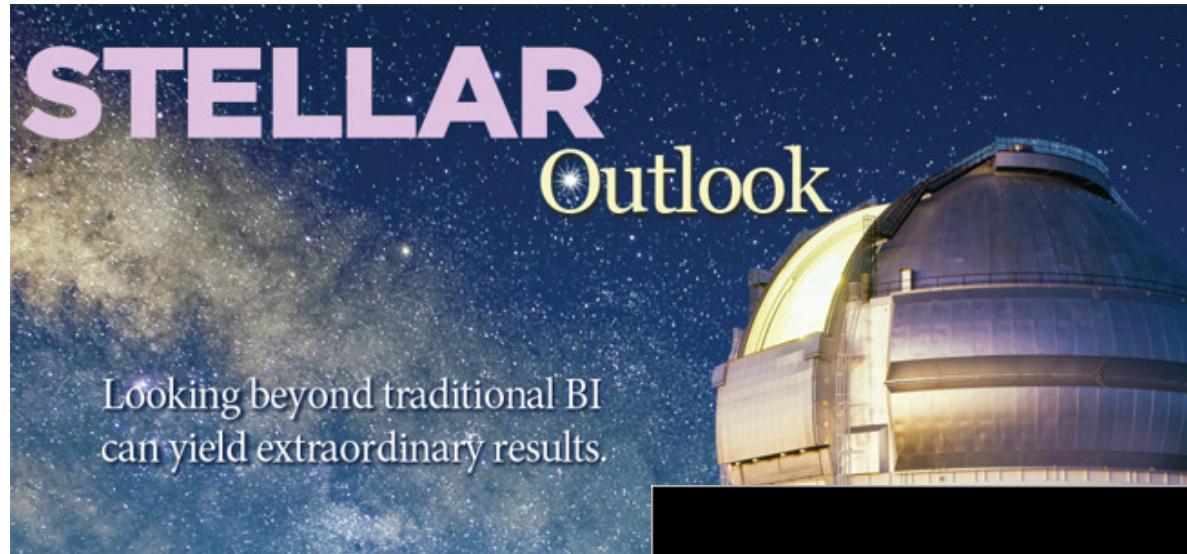
The Wrong Software

- Software vendor in Surrey, UK
- Vendor's developers in
 - India
 - South Africa
 - UK
 - [none initially in the] US
- The vendor has never come with an order of magnitude of scaling a system this size
- MongoDB and similar new products
 - No LDAP
 - No Active Directory
 - No audit trail
 - No tool vendors
 - No long-term financial track record

The Wrong Skill Set

- Client DBAs do not have development experience
- Client DBAs have little tuning experience ... in fact they really don't know what tuning is
- Vendor DBAs are 8 time zones away and support multiple customers
- Software upgrades that utilize none of the new features
- Audit trails no one looks at
- Far too much projector-ware floating around

Another Dead End



TERADATA
THE BEST DECISION POSSIBLE



Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

What Would Have Worked



Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

Exadata and ZFS



Why?

- Simplified management
- Lowered cost
- Addressed loading issues
- Could run all existing code
- 40GB Infiniband a substantial improvement over 10MB Ethernet
- One neck to choke ... eliminated two major finger-pointing vendors

Warning: Appliances are not childproof



Can handle 85% of most database loads



Why?

- Simplified deployment
- Simplified management
- Lowered cost
- Removes multi-vendor engineering requirement
- One neck to choke ... eliminates multi-vendor finger-pointing

Other Issues I See Regularly

- Lack of documented procedures
- Lack of documented processes
- Lack of tech writers

Staffing

Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

Considerations

- Be careful what you wish for ... you might just get it
- Do you have anyone that can interview people that have the skills you need ... rather than the skills you already have?

Who Wants These Skills?

| Job Title | Company | Location | Date Posted |
|--|-------------|--------------|-------------|
| Java Developer - Hadoop/Map Reduce experience needed | Intersource | Seattle, WA | Oct-17-2012 |
| Software Development Engineer (Java, Big Data) | TRIAD Group | Bellevue, WA | Nov-09-2012 |
| Software Developer - Ruby on Rails Engineer - Web Publisher | CyberCoders | Seattle, WA | Nov-14-2012 |
| Software Engineer: Big Data and Data Warehouse | Amazon | Seattle, WA | Nov-14-2012 |
| Sr. Software Developer - Big Data | Amazon | Seattle, WA | Nov-14-2012 |
| Software Development Engineer-Distributed systems(Amazon EMR and DynamoDB) | Amazon | Seattle, WA | Nov-14-2012 |
| Software Engineer: Big Data and Data Warehouse | Amazon | Seattle, WA | Nov-14-2012 |

Who Wants These Skills?

| Job Title | Company | Location | Date Posted |
|--|--------------------|-------------|-------------|
| Full Time Position For Hadoop Developer Job in Seattle, WA | Resource Logistics | Seattle, WA | Nov-14-2012 |
| Hadoop / HBase / ElasticSearch and Scala | Kshunya Inc | Seattle, WA | Nov-14-2012 |
| Hadoop Developer | HCL America Inc. | Seattle, WA | Nov-14-2012 |
| Java Developer - Hadoop/Map Reduce experience needed | Intersource | Seattle, WA | Oct-17-2012 |
| Hadoop/No SQL, Java, ETL) Engineer | People Tech Group | SEATTLE, WA | Nov-13-2012 |
| Sr. Software Engineer - Big Data, Hadoop, Amazon Elastic MapReduce | Amazon | Seattle, WA | Nov-14-2012 |
| Software Engineer - Big Data, Hadoop, Amazon Elastic MapReduce | Amazon | Seattle, WA | Nov-14-2012 |

What You Don't Really Want

| Job Summary | |
|---|-------------------------------------|
| Location: | Bellingham, WA 98227 |
| Job Category: | IT/Software Development |
| Occupations: | Database Development/Administration |
| Industry: | Insurance |
| Job Status/Type: | Full Time Employee |
| Career Level: | Experienced (Non-Manager) |
| Contact Information | |
| Reference Code: | 1289 |
| Database Administrator | |
| Position Description: | |
| As one of the nation's most, progressive and dynamic insurance organizations, Windsor Health Group is a Health Specialty Management Company comprised of Windsor Health Plan, Sterling Insurance, and Olympian Health Partners. Collectively we offer a full range of products and services aimed to improve health, peace of mind and financial security for our members and policy holders when navigating health care coverage options. We aim to make working at Windsor an energizing and fulfilling experience. We believe that if we expect people's continued energy and commitment at work, we must provide the right environment for that to happen. Our organization is looking for individuals with a strong commitment to quality service, a clear sense of purpose and pride in their individual contribution to the company's success. Summary: The Database Administrator works with senior database administration staff to design, implement, administer, troubleshoot and develop database resources to meet company and departmental goals and standards for availability, capacity and security. | |
| Duties and Responsibilities: | |
| <ul style="list-style-type: none">• Install, configure, troubleshoot, and performance tune Oracle and MS SQL databases on Sun UNIX, Linux, and Windows based operating systems.• Evaluate and install new software releases, perform system upgrades, evaluate and install patches and resolve database related problems.• Work jointly with networking engineers to install, configure, troubleshoot and performance tune database connectivity.• Ensure uninterrupted database availability by managing physical and logical space. Measure growth and project for continued optimal performance and server upgrades.• Maintain database security and where appropriate host security in accordance with company standards and industry best practices. This includes but is not limited to deployment of appropriate security patches, hot fixes, vulnerability scans and software and firmware upgrades.• Work with other IT staff members to automate management tasks, streamline processes and perform standard administration functions as needed.• Assist in IT disaster recovery planning.• Participates in the development of policies, procedures, and resources that standardize the functioning of the systems and software and in accordance with company standards and industry best practices.• Perform and adhere to established policies, and procedures.• Assist more experienced staff in resolution of system and software issues related to the company's database systems.• Coordinate with fellow IT staff to provide training and technical support to the user community on use of IT services.• On-call responsibilities as required.• Obtain and/or maintain technical certifications necessary for continued company growth and success.• Perform additional duties as assigned by management. | |

What You Don't Really Want

Knowledge, Skills and Abilities:

- Ability to evaluate critical systems, prioritize work, and determine solutions
- Excellent verbal and written communication skills
- A working knowledge of these technologies:
 - Oracle RDMBS; 9i, 10g, and 11g
 - Oracle RAC and Grid
 - Microsoft SQL Server 2005
 - SQL, Oracle PL/SQL, and Microsoft Transact-SQL
 - Sun Unix, Linux, and Microsoft Windows operating systems
 - IIS or Apache web server architecture.
 - Network storage technologies.
 - Backup software installation, configuration, and operation
 - TCP/IP
 - UNIX shell scripting (Korn) and Perl
 - Data center technologies including enterprise monitoring and event management.
 - Strong project management skills
 - Strong security knowledge and ability for securing and maintaining data for safety and to meet or exceed standards required by HIPAA, Sarbanes/Oxley and other regulations.
 - Technical expertise of personal computing systems, peripheral equipment, and applications
 - Strong troubleshooting skills
 - Strong organizational skills
 - Ability to work well under various urgent and time sensitive situations.
 - Excellent verbal and written communication skills
 - Strong attention to detail

Are my customer's happy?

Instance Efficiency Percentages (Target 100%)

| | | | |
|------------------------------|-------|-------------------|--------|
| Buffer Nowait %: | 99.80 | Redo NoWait %: | 100.00 |
| Buffer Hit %: | 97.34 | In-memory Sort %: | 99.99 |
| Library Hit %: | 99.97 | Soft Parse %: | 98.79 |
| Execute to Parse %: | 99.29 | Latch Hit %: | 99.57 |
| Parse CPU to Parse Elapsd %: | 0.00 | % Non-Parse CPU: | 96.60 |

No!

Foreground Wait Events

- s - second, ms - millisecond - 1000th of a second
- Only events with Total Wait Time (s) $\geq .001$ are shown
- ordered by wait time desc, waits desc (idle events last)
- %Timeouts: value of 0 indicates value was < .5%. Value of null is truly 0

| Event | Waits | %Time -outs | Total Wait Time (s) | Avg wait (ms) | Waits /bn | % DB time |
|-----------------------------|------------|-------------|---------------------|---------------|-----------|-----------|
| cursor: pin S wait on X | 559,981 | 100 | 11,250 | 20 | 15.41 | 23.77 |
| db file sequential read | 1,824,756 | 0 | 4,231 | 2 | 50.23 | 8.94 |
| unspecified wait event | 462,648 | 0 | 1,996 | 4 | 12.73 | 4.22 |
| gc buffer busy acquire | 322,335 | 0 | 1,342 | 4 | 8.87 | 2.84 |
| external table misc IO | 57,038 | 0 | 1,131 | 20 | 1.57 | 2.39 |
| db file scattered read | 69,467 | 0 | 989 | 14 | 1.91 | 2.09 |
| PX Deq: reap credit | 68,532,223 | 100 | 630 | 0 | 1,886.33 | 1.33 |
| IPC send completion sync | 3,694,966 | 98 | 389 | 0 | 101.70 | 0.82 |
| PX Deq: Slave Session Stats | 132,529 | 12 | 350 | 3 | 3.65 | 0.74 |
| external table read | 96,704 | 0 | 339 | 4 | 2.66 | 0.72 |
| PX Nsq: PQ load info query | 1,668 | 97 | 328 | 197 | 0.05 | 0.69 |
| read by other session | 91,147 | 0 | 300 | 3 | 2.51 | 0.63 |
| latch free | 4,480 | 0 | 297 | 66 | 0.12 | 0.63 |
| gc cr grant 2-way | 417,285 | 0 | 272 | 1 | 11.49 | 0.57 |
| DFS lock handle | 13,528 | 11 | 261 | 19 | 0.37 | 0.55 |
| gc cr multi block request | 547,105 | 0 | 217 | 0 | 15.06 | 0.46 |
| enq: PS - contention | 186,170 | 56 | 201 | 1 | 5.12 | 0.42 |
| kksfbc child completion | 2,097 | 100 | 138 | 66 | 0.06 | 0.29 |
| cursor: pin S wait on X | 559,981 | 100 | 11,250 | 20 | 15.41 | 23.77 |

No!

Dictionary Cache Stats

- "Pct Misses" should be very low (< 2% in most cases)
- "Final Usage" is the number of cache entries being used

| Cache | Get Requests | Pct Miss | Scan Reqs | Pct Miss | Mod Reqs | Final Usage |
|----------------------|--------------|----------|-----------|----------|----------|-------------|
| dc_awr_control | 94 | 1.06 | 0 | | 1 | 1 |
| dc_constraints | 75 | 81.33 | 0 | | 75 | 268 |
| dc_database_links | 12,557 | 0.00 | 0 | | 0 | 3 |
| dc_files | 18,320 | 0.00 | 0 | | 0 | 482 |
| dc_global_oids | 414,904 | 0.00 | 0 | | 0 | 355 |
| dc_histogram_data | 303,148 | 0.46 | 0 | | 307 | 146,041 |
| dc_histogram_defs | 4,462,445 | 0.10 | 0 | | 714 | 496,635 |
| dc_object_grants | 23,089 | 0.00 | 0 | | 0 | 1,094 |
| dc_objects | 2,748,959 | 0.45 | 0 | | 2,384 | 281,443 |
| dc_partition_scns | 3 | 0.00 | 0 | | 0 | 5 |
| dc_profiles | 12,918 | 0.00 | 0 | | 0 | 4 |
| dc_rollback_segments | 6,824,069 | 0.00 | 0 | | 0 | 1,662 |
| dc_segments | 242,656 | 4.68 | 0 | | 9,231 | 877,062 |
| dc_sequences | 396 | 3.28 | 0 | | 396 | 33 |
| dc_table_scns | 43 | 37.21 | 0 | | 0 | 8 |
| dc_tablespace_quotas | 8,495 | 21.27 | 0 | | 2,365 | 66 |
| dc tablespaces | 11,893,072 | 0.00 | 0 | | 0 | 38 |
| dc_users | 10,119,305 | 0.00 | 0 | | 15 | 6,827 |
| global database name | 8,789 | 0.00 | 0 | | 0 | 2 |
| outstanding_alerts | 182 | 75.82 | 0 | | 1 | 26 |

The SQL Challenge

find matching values present in two different tables

```
SELECT srvr_id
  FROM servers
INTERSECT
SELECT srvr_id
  FROM serv_inst;
```

```
SELECT srvr_id
  FROM servers
 WHERE srvr_id IN (
    SELECT srvr_id
  FROM serv_inst);
```

```
SELECT srvr_id
  FROM servers s
 WHERE EXISTS (
    SELECT srvr_id
  FROM serv_inst i
 WHERE s.srvr_id = i.srvr_id);
```

```
SELECT DISTINCT s.srvr_id
  FROM servers s, serv_inst i
 WHERE s.srvr_id = i.srvr_id;
```

Optimizer Plans (the way it was)

```
SELECT DISTINCT E1_2.OBJECT_ID
  FROM PMCM.ELEMENT_DETAIL E1_1, PMCM.ELEMENT_DETAIL E1_2, PMCM.MARK_NETW_HIERARCHY H1,
        PMCM.ELEMENT_DETAIL E2_1, PMCM.ELEMENT_DETAIL E2_2, PMCM.MARK_NETW_HIERARCHY H2
 WHERE E1_1.OBJECT_ID = H1.PARENT_ID
   AND E1_2.OBJECT_ID = H1.OBJECT_ID
   AND E2_1.OBJECT_ID = H2.PARENT_ID
   AND E2_2.OBJECT_ID = H2.OBJECT_ID
   AND E1_1.CURRENT_IND = 'Y' AND E2_1.CURRENT_IND = 'Y'
   AND E2_1.CURRENT_IND = 'Y' AND E2_2.CURRENT_IND = 'Y'
   AND H1.CURRENT_IND = 'Y' AND H2.CURRENT_IND = 'Y'
   AND H1.HIERARCHY_TYPE = 'NETWORK' AND H2.HIERARCHY_TYPE = 'NETWORK'
   AND H1.PARENT_TYPE IN ('BSC', 'RNC') AND H2.PARENT_TYPE IN ('BSC', 'RNC')
   AND E2_2.ELEMENT_TYPE = 'CELL' AND E1_2.ELEMENT_TYPE = 'CELL'
   AND H1.PARENT_TYPE IN ('BSC', 'RNC')
   AND E1_1.ELEMENT_NAME = E2_1.ELEMENT_NAME
   AND E1_1.ELEMENT_ID = E2_1.ELEMENT_ID
   AND E1_2.ELEMENT_NAME = E2_2.ELEMENT_NAME
   AND E1_2.ELEMENT_ID = E2_2.ELEMENT_ID
   AND E1_2.USEID LIKE '%%' AND E2_2.USEID NOT LIKE '%%';
```

| Id | Operation | Name | Rows | Bytes | TempSpc | Cost (%CPU) | Time | Pstart | Pstop |
|------|---------------------------|-----------------------------|------|-------|---------|-------------|----------|--------|-------|
| 0 | SELECT STATEMENT | | 1 | 78 | | 74M (40) | 50:54:42 | | |
| 1 | TEMP TABLE TRANSFORMATION | | | | | | | | |
| 2 | LOAD AS SELECT | | | | | | | | |
| 3 | PARTITION RANGE ALL | | 22M | 1111M | | 38153 (11) | 00:01:34 | 1 | 29 |
| * 4 | TABLE ACCESS FULL | ELEMENT_DETAIL | 22M | 1111M | | 38153 (11) | 00:01:34 | | |
| 5 | LOAD AS SELECT | | | | | | | | |
| 6 | PARTITION HASH ALL | | 337K | 9231K | | 3514 (15) | 00:00:09 | 1 | 16 |
| * 7 | TABLE ACCESS FULL | MARK_NETW_HIERARCHY | 337K | 9231K | | 3514 (15) | 00:00:09 | | |
| 8 | SORT AGGREGATE | | 1 | 78 | | | | | |
| * 9 | HASH JOIN | | 927G | 65T | 534M | 74M (40) | 50:53:00 | | |
| 10 | VIEW | | 22M | 277M | | 16808 (12) | 00:00:42 | | |
| 11 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485F_6A66C42E | 22M | 1111M | | 16808 (12) | 00 | | |
| * 12 | HASH JOIN | | 21G | 1272G | 534M | 1616K (43) | 01:06:04 | | |
| 13 | VIEW | | 22M | 277M | | 16808 (12) | 00:00:42 | | |
| 14 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485F_6A66C42E | 22M | 1111M | | 16808 (12) | 0 | | |
| * 15 | HASH JOIN | | 476M | 23G | 524M | 97327 (22) | 00:03:59 | | |
| * 16 | HASH JOIN | | 10M | 401M | 8704K | 34520 (10) | 00:01:25 | | |
| * 17 | HASH JOIN | | 234K | 5948K | 8256K | 783 (10) | 00:00:02 | | |
| 18 | VIEW | | 337K | 4286K | | 142 (14) | 00:00:01 | | |
| 19 | TABLE ACCESS FULL | SYS_TEMP_0FDA74860_6A66C42E | 337K | 3956K | | 142 (14) | 00:00:01 | | |
| 20 | VIEW | | 337K | 4286K | | 142 (14) | 00:00:01 | | |
| 21 | TABLE ACCESS FULL | SYS_TEMP_0FDA74860_6A66C42E | 337K | 3956K | | 142 (14) | 00:00:01 | | |
| 22 | VIEW | | 22M | 277M | | 16808 (12) | 00:00:42 | | |
| 23 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485F_6A66C42E | 22M | 1111M | | 16808 (12) | 00:00:42 | | |
| 24 | VIEW | | 22M | 277M | | 16808 (12) | 00:00:42 | | |
| 25 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485F_6A66C42E | 22M | 1111M | | 16808 (12) | 0 | | |

Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Optimizer Plans

(tuning gone terribly wrong)

| Id | Operation | Name | Rows | Bytes | TempSpc | Cost (%CPU) | Time | Pstart | Pstop |
|------|---------------------------|-----------------------------|------|-------|---------|-------------|-----------|--------|-------|
| 0 | SELECT STATEMENT | | | | | | | | |
| 1 | TEMP TABLE TRANSFORMATION | | | | | | | | |
| 2 | LOAD AS SELECT | | | | | | | | |
| 3 | PARTITION RANGE ALL | | 22M | 1111M | | 38153 (11) | 00:01:34 | | |
| * 4 | TABLE ACCESS FULL | ELEMENT_DETAIL | 22M | 1111M | | 38153 (11) | 00:01:34 | 1 | 29 |
| 5 | LOAD AS SELECT | | | | | | | | |
| 6 | PARTITION HASH ALL | | 337K | 9231K | | 3514 (15) | 00:00:09 | 1 | 16 |
| * 7 | TABLE ACCESS FULL | MARK_NETW_HIERARCHY | 337K | 9231K | | 3514 (15) | 00:00:09 | | |
| 8 | SORT AGGREGATE | | 1 | 78 | | | | | |
| 9 | MERGE JOIN | | 471P | 15E | | 14T (100) | 999:59:59 | | |
| 10 | MERGE JOIN | | 10P | 616P | | 694G (81) | 999:59:59 | | |
| 11 | MERGE JOIN | | 231T | 10P | | 377G (64) | 999:59:59 | | |
| 12 | SORT JOIN | | 334T | 11P | 28P | 377G (64) | 999:59:59 | | |
| 13 | MERGE JOIN CARTESIAN | | 334T | 11P | | 140G (14) | 999:59:59 | | |
| * 14 | HASH JOIN | | 989M | 23G | 534M | 96010 (38) | 00:03:56 | | |
| 15 | VIEW | | 22M | 277M | | 16808 (12) | 00:00:42 | | |
| 16 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485B_6A66C42E | 22M | 1111M | | 16808 (12) | 00:00:42 | | |
| 17 | VIEW | | 22M | 277M | | 16808 (12) | 00:00:42 | | |
| 18 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485B_6A66C42E | 22M | 1111M | | 16808 (12) | 00:00:42 | | |
| 19 | BUFFER SORT | | 337K | 4286K | | 140G (14) | 999:59:59 | | |
| 20 | VIEW | | 337K | 4286K | | 142 (14) | 00:00:01 | | |
| 21 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485C_6A66C42E | 337K | 3956K | | 142 (14) | 00:00:01 | | |
| * 22 | SORT JOIN | | 337K | 4286K | 12M | 844 (14) | 00:00:03 | | |
| 23 | VIEW | | 337K | 4286K | | 142 (14) | 00:00:01 | | |
| 24 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485C_6A66C42E | 337K | 3956K | | 142 (14) | 00:00:01 | | |
| * 25 | SORT JOIN | | 22M | 277M | 855M | 65084 (16) | 00:02:40 | | |
| 26 | VIEW | | 22M | 277M | | 16808 (12) | 00:00:42 | | |
| 27 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485B_6A66C42E | 22M | 1111M | | 16808 (12) | 0 | | |
| * 28 | SORT JOIN | | 22M | 277M | 855M | 65084 (16) | 00:02:40 | | |
| 29 | VIEW | | 22M | 277M | | 16808 (12) | 00:00:42 | | |
| 30 | TABLE ACCESS FULL | SYS_TEMP_0FDA7485B_6A66C42E | 22M | 1111M | | 16808 (12) | 0 | | |

Optimizer Plans (making it better)

```
WITH ed AS (SELECT object_id, element_id, element_name, element_type, useid
            FROM pmcm.element_detail
           WHERE element_type = 'CELL'
             AND current_ind = 'Y'),
      mn AS (SELECT parent_id, object_id
            FROM pmcm.mark_netw_hierarchy
           WHERE current_ind = 'Y'
             AND hierarchy_type = 'NETWORK'
             AND parent_type IN ('BSC', 'RNC'))
SELECT COUNT(*)
  FROM ed e1_1, ed e1_2, ed e2_1, ed e2_2, mn h1, mn h2
 WHERE e1_1.object_id = h1.parent_id AND e1_2.object_id = h1.object_id
   AND e2_1.object_id = h2.parent_id AND e2_2.object_id = h2.object_id
   AND e1_1.element_name = e2_1.element_name
   AND e1_1.element_id = e2_1.element_id
   AND e1_2.element_name = e2_2.element_name
   AND e1_2.element_id = e2_2.element_id
   AND e1_2.useid LIKE '*%'
   AND e2_2.useid NOT LIKE '*%';
```

| Id | Operation | Name | Rows | Bytes | TempSpc | Cost | (%CPU) | Time | |
|------|-----------------------------------|---------------------|------|-------|---------|-------|--------|----------|--|
| 0 | SELECT STATEMENT | | 1 | 214 | | 100K | (6) | 00:04:08 | |
| 1 | HASH UNIQUE | | 1 | 214 | | 100K | (6) | 00:04:08 | |
| * 2 | HASH JOIN | | 1 | 214 | 12M | 100K | (6) | 00:04:08 | |
| 3 | PARTITION HASH ALL | | 337K | 9231K | | 3514 | (15) | 00:00:09 | |
| * 4 | TABLE ACCESS FULL | MARK_NETW_HIERARCHY | 337K | 9231K | | 3514 | (15) | 00:00:00 | |
| * 5 | HASH JOIN | | 207K | 36M | 22M | 95860 | (6) | 00:03:56 | |
| 6 | PARTITION RANGE ALL | | 586K | 15M | | 16233 | (2) | 00:00:40 | |
| 7 | TABLE ACCESS BY LOCAL INDEX ROWID | ELEMENT_DETAIL | 586K | 15M | | 16233 | | ???:?? | |
| * 8 | INDEX SKIP SCAN | ED_ET_TECH_CI | 586K | | | 12791 | (1) | 00:00:3? | |
| * 9 | HASH JOIN | | 207K | 31M | 22M | 77982 | (7) | 00:03:12 | |
| 10 | PARTITION RANGE ALL | | 586K | 15M | | 16233 | (2) | 00:00:40 | |
| 11 | TABLE ACCESS BY LOCAL INDEX ROWID | ELEMENT_DETAIL | 586K | 15M | | 16233 | | ???:?? | |
| * 12 | INDEX SKIP SCAN | ED_ET_TECH_CI | 586K | | | 12791 | (1) | 00:00:?? | |
| * 13 | HASH JOIN | | 179K | 22M | 12M | 60372 | (8) | 00:02:29 | |
| 14 | PARTITION HASH ALL | | 337K | 9231K | | 3514 | (15) | 00:00:09 | |
| * 15 | TABLE ACCESS FULL | MARK_NETW_HIERARCHY | 337K | 9231K | | 3514 | (15) | 00:00:?? | |
| * 16 | HASH JOIN | | 184K | 17M | 10M | 55886 | (8) | 00:02:18 | |
| 17 | PARTITION RANGE ALL | | 184K | 9008K | | 37137 | (8) | 00:01:32 | |
| * 18 | TABLE ACCESS FULL | ELEMENT_DETAIL | 184K | 9008K | | 37137 | (8) | 00:01:32 | |
| 19 | PARTITION RANGE ALL | | 576K | 28M | | 17383 | (8) | 00:00:43 | |
| * 20 | TABLE ACCESS BY LOCAL INDEX ROWID | ELEMENT_DETAIL | 576K | 28M | | 17383 | (8) | ???:?? | |
| * 21 | INDEX SKIP SCAN | ED_ET_TECH_CI | 583K | | | 13939 | (9) | 00:00:35 | |

Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

Result Cache

SQL ordered by Executions

- Total Executions: 29,717,627
- Captured SQL account for 77.4% of Total

| Executions | Rows Processed | Rows per Exec | CPU per Exec (s) | Elap per Exec (s) | SQL Id | SQL Module | SQL Text |
|------------|----------------|---------------|------------------|-------------------|---------------|---------------------------------------|-------------------------------------|
| 10,128,178 | 2,506,529 | 0.25 | 0.00 | 0.00 | 932srzg1krc33 | ASN_07B_DIP(004110016) | SELECT NE_TIMEZONE FROM CMPM.E... |
| 7,576,759 | 7,579,197 | 1.00 | 0.00 | 0.00 | 1h698sb62un99 | asci_56_RANAPProtocolStats(016110006) | SELECT DISTINCT NE_TIMEZONE FR... |
| 3,914,621 | 3,848,268 | 0.98 | 0.00 | 0.00 | 5tbzddgguu8cc | asci_56_RANAPProtocolStats(016110006) | SELECT SYS_VERSION FROM CMPM.T... |
| 311,645 | 311,604 | 1.00 | 0.00 | 0.00 | 7gtztzv329wg0 | | select c.name, u.name from co... |
| 301,428 | 301,325 | 1.00 | 0.00 | 0.00 | 36s446f9cnwhw | | SELECT C.NAME FROM COL\$ C WHERE... |
| 200,692 | 200,669 | 1.00 | 0.00 | 0.00 | 4vs91dev7utp6 | OMS | insert into sys.aud\$ (sessioni... |
| 65,044 | 65,035 | 1.00 | 0.00 | 0.00 | fz9xwpt2cvt0k | | SELECT par_type, param_clob, ... |
| 64,949 | 3,945,482 | 60.75 | 0.00 | 0.00 | f5ra7dru5fk5n | XML_P7R_RNC_RCS(003110008) | SELECT NAME, PATH, READ, WR... |
| 64,801 | 64,807 | 1.00 | 0.00 | 0.00 | fhzj09a7fmrnb | XML_V7I_IN_LP_DC(00811000V) | SELECT DBTIMEZONE, LENGTH(DBT... |
| 64,632 | 64,542 | 1.00 | 0.00 | 0.00 | 15jnrrb6016nd | XML_V7I_IN_LP_DC(00811000V) | SELECT SESSIONTIMEZONE, LENGTH... |

```
CREATE OR REPLACE FUNCTION rcache(p_srvr_id IN servers.srvr_id%TYPE) RETURN BOOLEAN
RESULT_CACHE RELIES_ON(servers) IS
  srvrow servers%ROWTYPE;
BEGIN
  SELECT *
  INTO srvrow
  FROM servers
  WHERE srvr_id = p_srvr_id;
  RETURN TRUE;
EXCEPTION
  WHEN OTHERS THEN
    RETURN FALSE;
END rcache;
/
```

```
SELECT /*+ RESULT_CACHE */ srvr_id
FROM (
  SELECT srvr_id, SUM(cnt) SUMCNT
  FROM (
    SELECT DISTINCT srvr_id, 1 AS CNT
    FROM servers
    UNION ALL
    SELECT DISTINCT srvr_id, 1
    FROM serv_inst)
  GROUP BY srvr_id)
WHERE sumcnt = 2;
```

Conclusions

Is Anyone Ready For The Next Generation?

- No one is ready for the next generation
- No one has an appropriately sized training budget?
- There really is no such thing as "Big Data"



LONELINESS

IF YOU FIND YOURSELF STRUGGLING WITH LONELINESS, YOU'RE NOT ALONE.
AND YET YOU ARE ALONE. SO VERY ALONE.

Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

Questions

**ERROR at line 1:
ORA-00028: your session has been killed**



Thank you

InfoWorld: November 17, 1997

Ask the right questions

MICHAEL VIZARD ASKS in a recent column, "So the question is, do we just have to hope that the courts will serve us well as the government gets more involved in technology?" [See From the News desk, Oct. 27, page 3.] No, I don't think that is the question, not at all! Part of the social contract we have all entered into, as citizens of the United States, is one in which the government represents the will of the governed for the mutual benefit of the governed. If you don't subscribe to this concept, as imperfectly as it may at times be applied, you should not be a citizen. A better question would have been: "Does technology deserve to be exempt from the social contract governing every other business enterprise in the United States?" I would say emphatically no.

We are no better and we are no worse and we are no more deserving.

Daniel A. Morgan

EXE Solutions

danmorgan@exesolutions.com

Daniel A. Morgan | damorgan12c@gmail.com | www.morganslibrary.org

Challenges in a World of Large Databases and Business Intelligence

Presented: Association for Women in Computing - 14 November, 2012

InfoWorld: July 19, 1999

Linux is stable, period

I'VE READ Nicholas Petreley's response to Bob Metcalfe [about Linux] and I think neither of them gets it. [See Down to the Wire, June 28, page 78.]

Economics is important. No question about that. But the reason why Linux is making inroads is only minimally due to pricing and economic models. The reason, very simply, is that no one in IS takes the money for a server operating system out of his or her own pocket. The cost is just figured into the cost of a project, which includes hardware, software (purchase or development), deployment, and maintenance. An operating system that is free or an operating system that costs \$1,000 per server is just pennies in the total scheme of things.

Out here in the cubicles where the work is

done, we are moving toward Linux for one primary reason: stability.

If Windows NT was free, heck, if Microsoft paid us money to deploy NT Servers, the OS still isn't stable enough for mission-critical applications. Do I want to take support calls during the business day because the server crashed again? Do I want to drive to the office in the middle of the night to reboot the server because of memory leakage? Do I want to have a closer relationship with my cell phone and pager than I do with my girlfriend? I don't think so.

When Microsoft wants to be taken seriously, it will stop development of new versions of its OS and fix the one it already has. Given the choice between Windows 2000, Version 1.0, and a stable, secure version of NT 3.51, I'd go back to 3.51 in a second.

*Daniel A. Morgan
Mercer Island, Wash.*