

ASSUMPTIONS

Killer of Tech Projects and First Dates Around the World

October 2014

hroug

hrvatska udruga oracle korisnika

TIM GORMAN AND KELLYN POT'VIN-GORMAN

- Married and live in Westminster, CO, near Denver
- Both active in Oracle users groups
- Kellyn's children - Samuel, Caitlyn and Joshua
- Tim's children - Peter and Marika
- Passionate about our work and our lives

HOW CAN IT PROJECTS AND FIRST DATES BE ALIKE?

- Stakeholders
- Requirements
- Expectations
- Planning
- Competing Goals
- Meetings

AGENDA

- First impressions
 - Assumptions and expectations first collide with reality
- Down The Rabbit Hole
 - The term comes from Lewis Carroll's "*Alice In Wonderland*"
 - Following a bad assumption with good reasoning
- Letting Evidence Lead
 - Quote from last year's Boston Marathon bombing investigator:
"As soon as you say to yourself, ***I know what happened*** or ***I know what is going on here***, that is when you are in deep trouble"
 - Certainly we've seen similar comments from investigators in the Malaysian flight 370 investigation
- Resolution
 - Stop, reset, and start again from the facts



THE PITFALL OF ALL OPTIMIZATION: THE “TINSEL MONKEY”



TINSEL MONKEY IS...

A term for...

- A monkey that haunts tourists and locals in India, China and Thailand countries that hunts down shiny objects.
 - The monkey has no concept of value.
 - Goes off just that the object is “shiny”
 - Simply assumes shiny object is gold or silver.
 - Doesn't realize difference in gold vs. tinfoil(tinsel).
- Only does hard work offer value.
- No analysis of data to ensure value of shiny object.



FIRST IMPRESSIONS – FIRST DATE

- Storyline... “John and Margo”
 - We met on a blind date, from a set up by friends...
 - ...OK, really we met from an online dating website...
 - Walking into a first-meet situation like that, one has expectations
 - In the absence of any data, one makes assumptions...

FIRST IMPRESSIONS – FIRST DATE



FIRST IMPRESSIONS – FIRST DATE



FIRST IMPRESSIONS – FIRST DATE



FIRST IMPRESSIONS – FIRST DATE


- Hmm, this is going to be a lot of fun!



FIRST ASSUMPTIONS- APPEARANCE



FIRST IMPRESSIONS- HIM



THE “EXADATA SOLUTION”

- Client has new ETL
- Will not complete on current hardware
- Purchase one, not two engineered solutions through Oracle.
- Told by development that specific Exadata solutions will “fix everything”.
- Wants to implement Auto-DOP, but not parallel..☺
- I’m brought into project after to manage Exadata
- The fun begins...



FIRST IMPRESSIONS - TECHNICAL

- Customer: “We have a data warehouse that is slow.”
 - “It uses a dimensional data model”
 - “Data loading is slow and is getting worse”
 - “Queries are very slow and are getting worse”
 - “We’re looking into obtaining an Exadata”



FIRST IMPRESSIONS - FORENSICS

- “ORA-03135: connection lost contact” errors occur sporadically
 - Google and MOS say to increase INBOUND_CONNECT_TIMEOUT value
 - Sure, fiddle with the only tunable parameter associated with the functionality
 - What assumptions have been made here?
- MOS also says to investigate using ASH
 - A more objective forensic approach

*

DOWN THE RABBIT HOLE - - FIRST DATE

- First sign of trouble...
 - Margo: Wow, I REALLY NEED a drink!

DOWN THE RABBIT HOLE – FIRST DATE



DOWN THE RABBIT HOLE – FIRST DATE

- More trouble...
 - John: If that waiter doesn't get here soon, I've got a permit for concealed carry of a weapon, so I just might have to take him hostage.

DOWN THE RABBIT HOLE – FIRST DATE

- “As far as interests are concerned, I like target shooting”





THE EXADATA PROBLEM

- The Challenge
 - AWR/ASH
 - High PGA/Temp Usage
 - 80% of overall waits
 - DML used for ETL load to FACT/DIM Tables
 - Traces show dynamic sampling is averaging 83% of total elapsed time for report queries.
 - No research was performed to find out what was causing actual performance issue.
 - ...but smart scans are working GREAT!! 😊

THE EXADATA CONUNDRUM

- PGA/Temp work are not “offloadable”.
- HCC isn’t feasible on tables using UPDATE and DELETE.
- 11.2.0.3 has the opportunity to decide DOP over parameter setting to perform dynamic sample.
- Due to base star schema without data presented in proper format, complex hash/sum/sort of data in each ETL and report query.
- None of the above are solved by Exadata features.

DOWN THE RABBIT HOLE – FORENSICS

- Troubleshooting using Google and MOS notes were tried
 - Didn't work
- Troubleshooting using ASH yielded fantastic results
 - Events observed in ASH initially showed library cache and sequence refresh issues
 - Subsequent analysis showed RAC issues related to sequence refresh
- So, ORA-03135/3136 is due to sequence refresh issues...

LETTING EVIDENCE LEAD – FIRST DATE

Just because I carry a concealed gun...

...doesn't mean that I like targeting other people...

...or that I'm not a responsible and careful gun owner...

...or that I see guns as a toy when I get drunk...

...or that I have a secret fantasy of myself as "Rambo"...

It just means that I like the technical and physical challenge of placing a bullet within a target.

LETTING EVIDENCE LEAD – FIRST DATE

Just because I like target shooting...

...doesn't mean that I don't believe in responsible gun ownership




LETTING EVIDENCE LEAD – FIRST DATE

“My nickname is “DD” ...

...no, not for “drunk and disorderly” ...

...it stands for “designated driver”, because when I’m out with friends, I’m the one who ensures that everyone gets home safely



LETTING EVIDENCE LEAD – FORENSICS

Boston Marathon bombing investigator:

“As soon as you say to yourself, “I think I know what happened” or “I think I’ve seen this before”, you have begun heading in the wrong direction

- Past success in troubleshooting does not apply to similar symptoms in future
 - An alarm in your head should start thinking “#FAIL”

RESOLUTION – FORENSICS

Even after resolving the ORA-03135 (“connect lost contact”) successfully using real data from ASH the first time...

...I fell right back into the trap of making assumptions...

...by making the assumption that this new situation was the same thing all over again...

- Understand the benefits and limitations of your tools
 - Active Session History (ASH)
 - trace files
 - extended trace files
 - Automatic Workload Repository (AWR)
- Apply forensic methods again and again...



THE EXADATA OLD-SCHOOL

- Implemented “dynamic sampling(0)” and “parallel(4)” hint at ODI universe level to stop dynamic sampling and keep parallel efficient for ETL processing.
- Located three, (currently) consistent joins that caused heavy temp reads/writes and created one simple mview and 1 rollup table, (more in future recommended)
- Reports went from 15-20 minutes down to under 12 seconds.

RESOLUTION – FIRST DATE

“My first impression was unfair”

“Can we start over again”



RESOLUTION

- We all come into situations, at work and in life, with assumptions
- The key is to recognize those assumptions and replace them with facts
- Sell the value of providing stakeholders with the right solution.

