



# Analyzing Data Warehouse Performance...

*...It's not just tuning SQL queries*



Presented by:  
Kevin Courtney, VP Solutions  
Teleran Technologies

# Agenda



- Data Warehouse Performance Process: Introducing Applications Usage Management Methodology
- Why it matters?
- Case Study: Fortune 50 Financial Services Firm
- Closing Thoughts

# Short Commercial & Disclaimer...



Hi, I'm Kevin...

Overall, 30 years in technology...

Last 15 years in BI / DW...as a practitioner

Currently, VP, Business Solutions at Teleran Technologies



# Some Thoughts on Information Management



“The computer is here to stay, therefore it must be kept in its proper place as a tool and a slave, or we will become sorcerer's apprentices, with data data everywhere and not a thought to think.”

-- Jesse Shera (father of modern library sciences)

“It is a capital mistake to theorize before one has data.”

-- Arthur Conan Doyle (creator of Sherlock Holmes)

“They have computers, and they may have other weapons of mass destruction.”

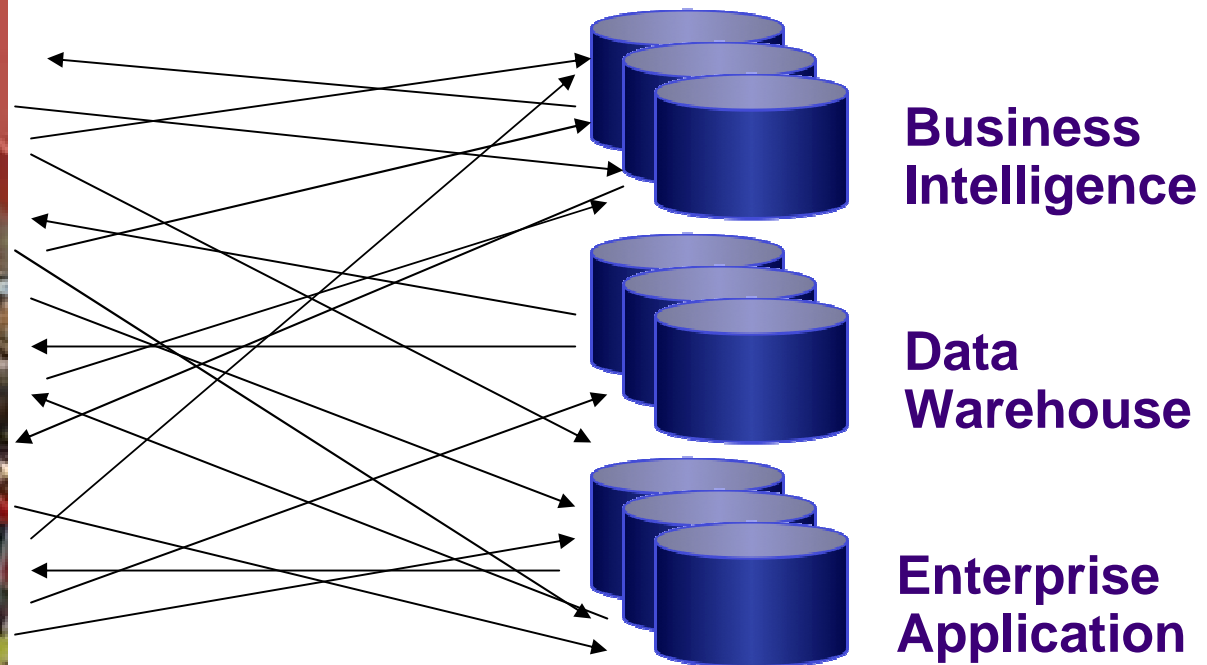
-- Janet Reno (former Attorney General)



# The Challenge: Growth and Complexity



More people, information, applications, environments threaten compliance, security, performance, and business value.



# Applications Usage Management



“Understanding and managing how users access and interact with applications and data to improve the performance, quality, compliance and TCO of business-critical applications”

*Claudia Imhoff*

# Some Simple Questions:



How many applications access the DW?

How many users access the DW?

How many resources are consumed, by whom and what?

When are these resources consumed?

What is the weakest link (ie database, BI server, network)?

## Some Harder Questions:



Are my most successful *{brokers, sales reps, etc.}* using this stuff?

Which reports, analyzes and ah-hocs need to be tuned?

Who is causing bottlenecks, why?

Which LOBs use certain applications? Adoption?

What value does my DW / BI environment deliver?

# The Age Old Dilemma... Performance



IT says, “The DW is available 99.99% of the time...”

Business responds, “...But it is too slow.”

Management asks, “Are we getting our money’s worth from this stuff?”

# Case Study



## The Company:

- Fortune 50 Financial Services Firm

## The Environment:

- Result of several big mergers
- Many applications and tools
- Thousands of users, millions of queries, hundreds of TB

## The Challenge:

- Poor performance, no visibility



# Holistic AUM Approach



We need:

- to know what usage patterns to look for (methodology)
- the granular data to analyze
- tools, and 'know how' analyze and recommend
- the clout to address and remediate

**...as we prepare to enter the Performance Data Jungle...  
get ready for the "Aha moments" we will discover**

# Performance: Slow Reporting, A Vicious Cycle



**Scenario:** Report / Analysis takes 'too long'

**Behind the Scenes:** Error encountered in DB

**User Action:** Request it again and again

**Result:** Opportunity for user education

# Performance: Slow Reporting, A Vicious Cycle Part II



**Scenario:** Report / Analysis takes 'too long'

**Behind the Scenes:** Long running query

**User Action:** Request it again and again

**Result:** Opportunity for user education

**...sometimes, users shoot themselves in their own feet.**

# Performance: Exposing Big Queries



**Scenario:** Analyst launches complex query\* to load SAS ODS during 'prime time'

**Behind the Scenes:** Disproportionate amount of resources consumed impacting SLAs

**Result:** Response time suffers across the eco-system; IT installs shift balancing rules

\*- complex joins, wide and long, no constraints

**...sometimes, 'special users' shoot others in the foot.**

# Performance: Slow Data Movement



**Scenario:** Identified loading large downstream data mart using MS Access and Toad

**Behind the Scenes:** Series of long running queries, consuming resources, impacting SLAs

**Result:** IT installs shift balancing rules; BI COE to assist with best practices approach

# Performance: Uncovering Bad SQL



**Scenario:** Poorly written queries 'ping' the warehouse thousands of times per day

**Behind the Scenes:** Disproportionate amount of resources consumed impacting SLAs

**Result:** Conduct meetings with applications owners to remediate, based upon 'facts', not hunches

**...if they liked the results the first time, they really loved them by the 5,000<sup>th</sup> time that day...**

# Compliance: Finding Rogue Data Marts



**Scenario:** Identified several unauthorized, stealth, large downstream data marts

**Behind the Scenes:** Violates compliance and security policy; consuming resources, impacting SLAs

**Result:** Compliance and Security to investigate...

**Compliance Nightmare: Avoid being on Page 1 of WSJ**

# Operations: Locating Dormant Data



**Scenario:** Warehouse has 5,000 tables, 115,000 columns, over 100 TB of storage

**Behind the Scenes:** Many tables and columns unused

**Result:** Analyze based upon actual query patterns, develop recommendations for lower cost storage and off-line archive

# Design: Need an Index?



**Scenario:** Many 'highly used columns' might be index candidates.

**Behind the Scenes:** Balance based upon workload patterns

**Result:** Analyze based upon actual query patterns, develop recommendations for indices; improve overall performance

# Design: Which SQL to Tune?



**Scenario:** Millions of queries; thousands of models executed repeatedly; who are the candidates for tuning

**Behind the Scenes:** Identify the most used SQL and the longest running SQL in 'magic quadrant-style' approach

**Result:** Tune low hanging fruit first to improve overall resource consumption.


# The End Game...



Use a holistic AUM approach to:

- diagnose and collect evidence about the health of the DW / BI ecosystem...
- maintain and improve the health of the DW / BI ecosystem...
- manage expense and avoid cost...
- align IT resources to business imperatives...

# Closing thoughts... keep it simple



© UFS, Inc.