GLOC

Great Lakes Oracle Conference

MAY 16 & 17, 2018

CLEVELAND PUBLIC AUDITORIUM, CLEVELAND, OHIO

WWW.NEOOUG.ORG/GLOC

Everything You Know Is Wrong: DBA Intuition, Meet Hard-Core Metrics

Jim Czuprynski
ViON Corporation
May 17, 2018

My Credentials

- 35+ years of database-centric IT experience
- Oracle DBA since 2001
- Oracle 9i, 10g, 11g, 12c OCP
- Oracle ACE Director
- 100+ articles on databasejournal.com and IOUG SELECT
- Co-author of 4 Oracle books
- Oracle-centric blog (**Generally, It Depends**)
- Regular speaker at Oracle OpenWorld, IOUG COLLABORATE, Hotsos Symposium, and Regional OUGs

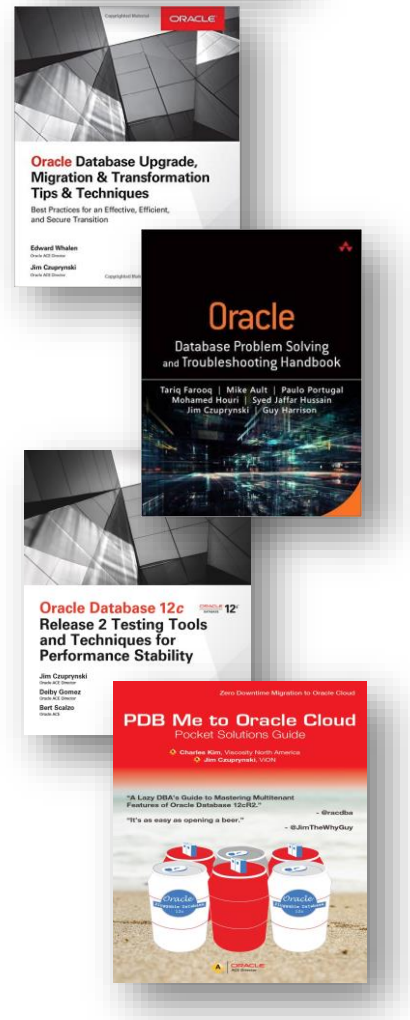

Experience. Expertise. Solutions.

ViON

Veteran-owned, privately-held, with financial strength and flexibility

Experience

Scaling to serve government and commercial customers

Solutions

Address today's and tomorrow's IT challenges

Expertise

Data management **Expertise**

Customer Intimacy is our Mission

Our Agenda

- Evaluating Your Career: Applying Metrics to Your Future
- Horror Story #1: “What’s Wrong With the Database?”
- When Intuition Fails, There’s Always Metrics
- Understanding the Underlying Genius of Quadruple-A
- Horror Story #2: “But VMSTAT Says …”
- Horror Story #3: “Your SAN’s Performance S*cks!”
- Conclusions

Evaluating Your Career: Applying Metrics to Your Future

- Data gathered from **600,000**+ Oracle Developers and Oracle DBAs over the past **five years**
- **Advanced analytics** were applied via **multiple** iterations of **robust** modeling
- Predictive **accuracy** has been verified at **92.8%** (**+/- 3.2**%)
- The result? An **accurate profile** of your current skillset … and what it bodes for your future job prospects!

Next … a demonstration of this **highly accurate** and **valuable** career counseling tool!


And now, a brief demonstration
of the incredible accuracy
of this new predictive analytics tool!

Relax. You've Just Been Forer'd.

Psychologist Bertram R. Forer established the ***Forer Effect*** in 1948, which states that people tend to rate sets of statements about themselves as ***highly accurate …*** even though those same statements could apply to ***a multitude of people***.

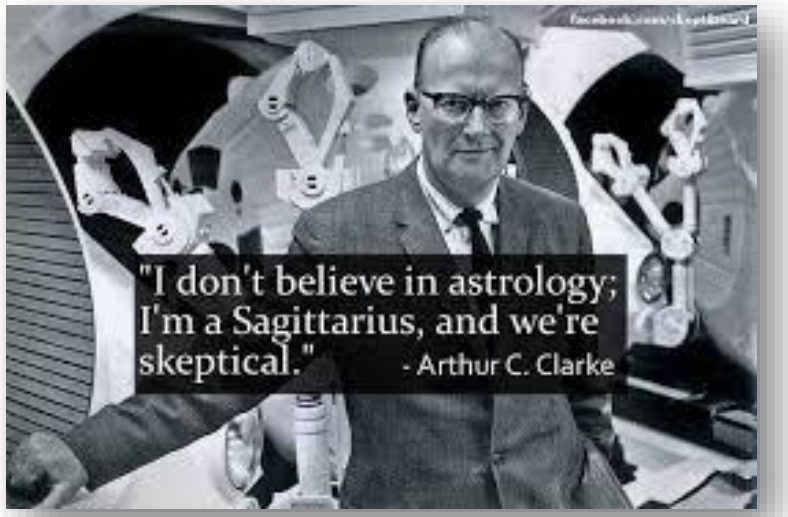

So, Three Scientists Are Having Coffee One Morning …

After 10 r
the man
house, nc

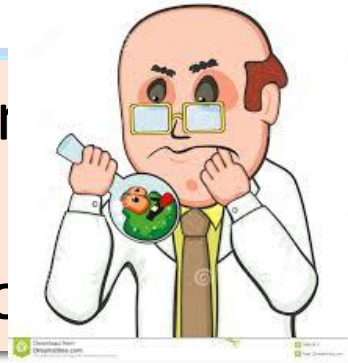

ss, they see
n leave the
hild in tow

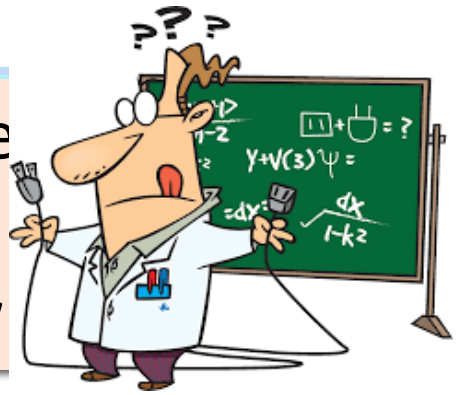

The mathematician thinks:
If just one more person enters the house …
then it will be empty.

Horror Story #1: “What’s Wrong With the Database?”

Oracle licensing audit revealed need to reduce number of licensed CPUs by **50%** mission-critical production database

- Solution: Move database to new server with 4 vs. 8 CPUs
- Before the move to new server:
  - **13%** CPU utilization at steady state (i.e. no workloads running)
  - Monthly batch processing spikes CPU utilization to **~ 40%**
- During workload simulation testing on new server:
  - **25%** CPU utilization at steady state
  - Batch processing spikes CPU into **75%** range, but application performance is still acceptable
- Conclusion: **This will work!**

Horror Story #1: "What's Wrong With the Database?"

- Everything works fine for **3 days** … and then:
  - Off-peak CPU utilization spikes to **55%**
  - Next batch processing cycle will **deplete** server's CPU capacity
- **Boss:** ***Yeah … I thought you tested this?***
- **Me:** [gulp] Something must've changed. I'll figure it out …
- **Solution:** Installed STATSPACK
  - One statement is executing 1M per hour:
    SELECT SYSDATE FROM DUAL;
  - Hidden in base Java code in **two** applications
- **Resolution:**
  - Fix the code
  - Terminate the developer
  - Keep STATSPack in my arsenal

When Intuition Fails, There's Always Metrics

Snapshots
Active
from my
Session
Mexico
History
vacation

(**ASH**)

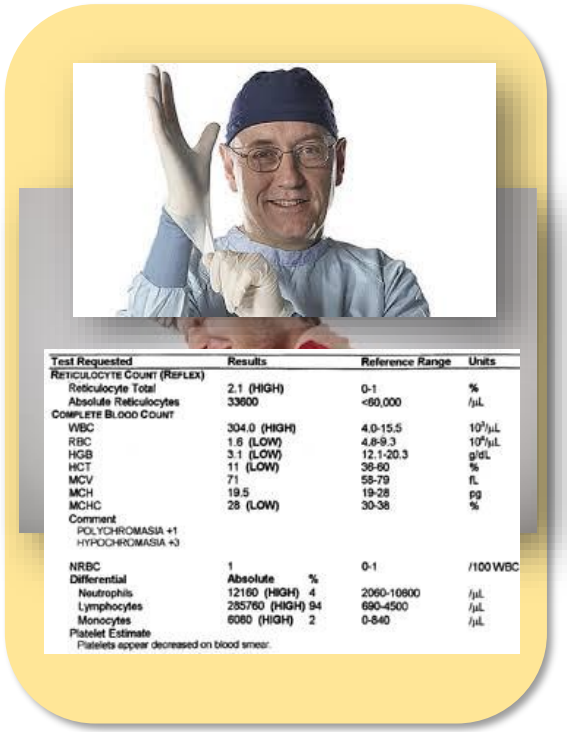

Test Requested	Results	Reference Range	Units
RETICULOCYTE COUNT (REFLEX)			
Reticulocyte Total	2.1 (HIGH)	0-1	%
Absolute Reticulocytes	33600	<60,000	/µL
COMPLETE BLOOD COUNT			
WBC	304.0 (HIGH)	4.0-15.5	10^3/µL
RBC	1.6 (LOW)	4.8-9.3	10^6/µL
HGB	3.1 (LOW)	12.1-20.3	g/dL
HCT	11 (LOW)	36-60	%
MCV	71	58-79	fL
MCH	19.5	19-28	pg
MCHC	28 (LOW)	30-38	%
Comment			
POLYCHROMASIA +1			
HYPOCHROMASIA +3			
NRBC	1	0-1	/100 WBC
Differential	Absolute %		
Neutrophils	12160 (HIGH) 4	2060-10800	/µL
Lymphocytes	285760 (HIGH) 94	690-4500	/µL
Monocytes	6080 (HIGH) 2	0-840	/µL
Platelet Estimate			
Platelets appear decreased on blood smear.			

My Automatic
Tir Workload
nR Repository

(**AWR**)

Solutions:
Automatic Database
- Orthopedic surgeon
Diagnostic Monitor
- Physical therapy

(**ADDM**) and Advisors
- No more free weights!

All Hail Dr. Deming!

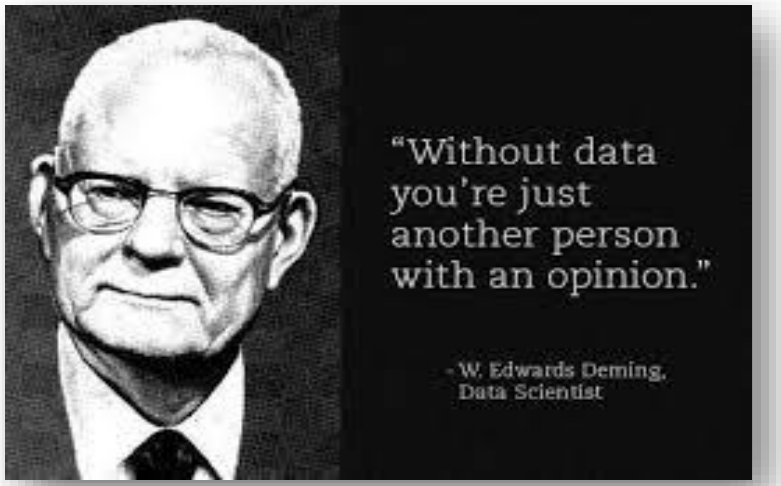

Oracle's metrics are based on **Dr. W. Edwards Deming's** principles of ***Statistical Process Control*** (**SPC**):

- A certain amount of acceptable **variation** in performance exists within a closed system.
- Events **within** a range of acceptable performance aren't worth worrying about.
- But when an **outlier** event occurs, either:
  - (a) **Re-evaluate** the performance model, or
  - (b) **Attack** the outlier with extreme prejudice

Andy Rivenes is the senior PM for all things Oracle Database In-Memory. He's codified several rules of thumb for understanding Oracle Database performance.
His **most famous rule**:

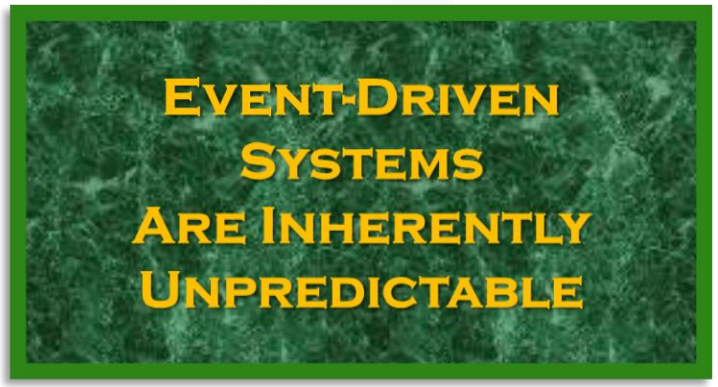

Don’t Panic. Just Call the Quadruple A!

AAAA

ASH captures extremely detailed snapshots of database performance in real time – once every *second*! – across several performance *dimensions*

AWR retains only the most significant events within snapshot periods for current and/or later analysis

WORKLOAD REPOSITORY report for

DB Name	DB Id	Unique Name	Role	Edition	Release	RAC	CDB
DB12CR2	1416159483	DB12CR2	PRIMARY	EE	12.2.0.1.0	NO	NO

Instance	Inst Num	Startup Time
DB12CR2	1	02-May-17 16:05

Host Name	Platform	CPUs	Cores	Sockets	Memory (GB)
nuvola2	Linux x86 64-bit	1			2.62

	Snap Id	Snap Time	Sessions	Cursors/Session
Begin Snap:	229	02-May-17 18:58:21	66	1.4
End Snap:	234	02-May-17 19:29:20	65	1.6
Elapsed:		31.00 (mins)		
DB Time:		120.84 (mins)		

Report Summary

Top ADDM Findings by Average Active Sessions

Finding Name	Avg active sessions of the task	Percent active sessions of finding	Task Name	Begin Snap Time	End Snap Time
Top SQL Statements	4.90	89.93	ADDM:1416159483_1_232	02-May-17 18:58	02-May-17 19:15
Top SQL Statements	2.76	94.98	ADDM:1416159483_1_233	02-May-17 19:15	02-May-17 19:29
Undersized PGA	4.90	30.48	ADDM:1416159483_1_232	02-May-17 18:58	02-May-17 19:15
Undersized PGA	2.76	46.67	ADDM:1416159483_1_233	02-May-17 19:15	02-May-17 19:29
Top Segments by "User I/O" and "Cluster"	2.76	35.93	ADDM:1416159483_1_233	02-May-17 19:15	02-May-17 19:29

ADDM analyzes performance every time an AWR snapshot is taken and recommends actions to be taken, or recommends a referral to a trusted **Advisor** (e.g. SQL Advisor)

```
ADDM Task ADDM:1416159483_1_232

          ADDM Report for Task 'ADDM:1416159483_1_232'
          --------------------------------------------

Analysis Period
---------
AWR snaps
Time peri
Time peri

Analysis
---------
Database
Database
ADDM perf
nuvola2.

Activity
---------
Total dat
The avera

Summary o
---------
   Descri

   ---------
1  Top SQL
2  Undersiz
3  Top Segm
4  Hard Par
```

```
          Findings and Recommendations
          ----------------------------

Finding 1: Top SQL Statements
Impact is 4.41 active sessions, 89.93% of total activity.
```

```
Finding 2: Undersized PGA
Impact is 1.49 active sessions, 30.48% of total activity.
-------------------------------------------------------------
The PGA was inadequately sized, causing additional I/O to temporary
tablespaces to consume significant database time.
The value of parameter "pga_aggregate_target" was "416 M" during the analysis
period.

   Recommendation 1: Database Configuration
   Estimated benefit is 1.05 active sessions, 21.34% of total activity.
   ---------------------------------------------------------------------
   Action
      Increase the size of the PGA by setting the value of parameter
      "pga_aggregate_target" to 745 M.

   Symptoms That Led to the Finding:
   ---------------------------------
      Wait class "User I/O" was consuming significant database time.
      Impact is 1.63 active sessions, 33.23% of total activity.


Finding 3: Top Segments by "User I/O" and "Cluster"
Impact is .43 active sessions, 8.76% of total activity.
```

Horror Story #2: “But VMSTAT Says …”

Complaint

Client reports poor ***database server*** performance

- CPU count, storage, and memory sizes are identical for two different database hosts
- SysAdmin: *I should be seeing significantly better CPU utilization on my bare-metal system than on my VMWare Linux system!*

Diagnosis

- **Me:** *Can I see your most recent AWR report for each system?*
- **SysAdmin:** *What’s an AWR report?*

Solution: Use the ***right*** metrics!

- SysAdmin had no visibility as to how the ***databases*** were performing
- DBA had not been contacted for performance appraisal
- The QA/Test Linux host’s database implemented AWRs …
- … but the production database running on bare metal host only had STATSPACK installed (to save on licensing costs)

Horror Story #3: “Your SAN’s Performance S*cks!!!”

New high-powered (and high-priced!) SAN installed at client

- Intense load testing begins on two-node RAC system
- One node is performing perfectly; the other node is crawling along
- CIO storms in and, umm, **expresses his displeasure** at new SAN

Diagnosis

- Ran AWR against database, then instances separately
- Workload evenly split (based on **session count**) across both nodes
  - Node #1:Tearing it up!!
  - Node #2: High I/O waits for queries and DML

Solution

- Consulted with Network and Storage Admins
  - Network admin says, “*Network is healthy! Must be the SAN.*”
  - Storage admin says, “*Hmmm … let me check everything again*.”

The click … was deafening.

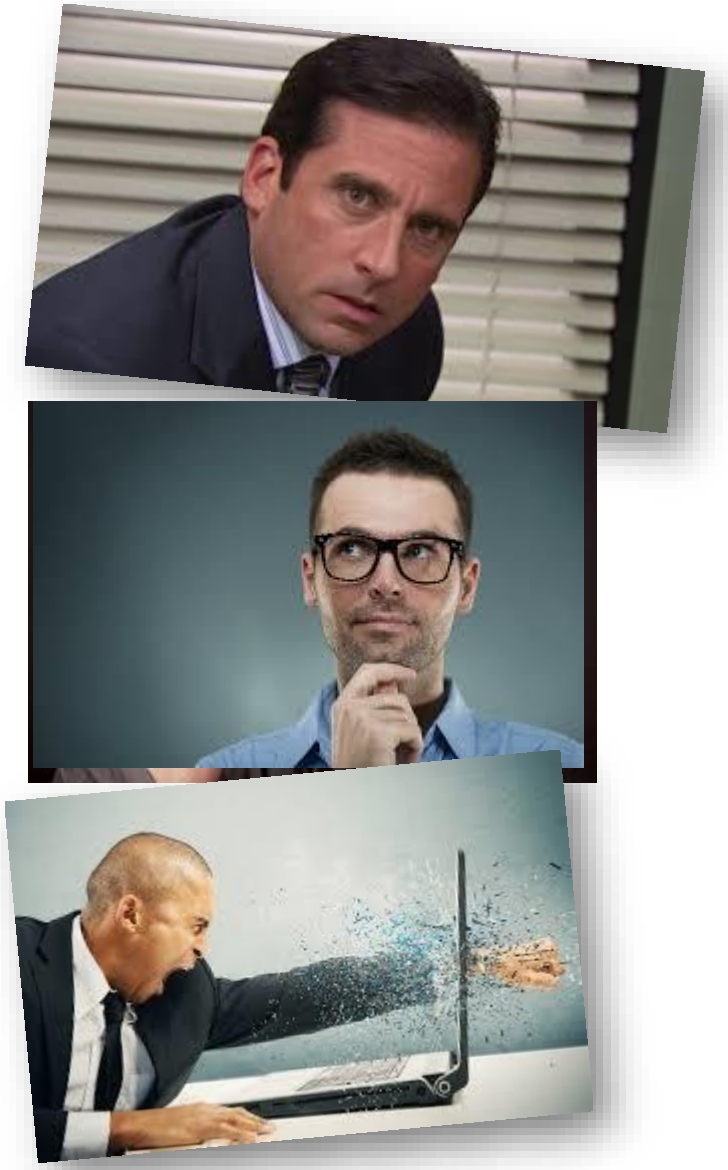

Conclusions

Metrics can tell you what's ***right*** *with your system*, as well as what ***isn't wrong***

- Going with your ***gut*** can cause ***indigestion*** … and often ***does***

Metrics are the key to understanding what part of your system is:

- Performing as expected
- Performing beneath expectation
- Not at fault at all

Be sure to use the right tools for the job!

- **ASH** for real-time performance analytics
- **AWR** for identifying the most "offensive" SQL statements and pernicious problems
- **ADDM** for getting the right advice on how to fix issues
- **Advisors** to delve deeply into problems that ADDM reports

Thank You For Your Kind Attention!

If you have any questions or comments, feel free to:

- E-mail me at Jim.Czuprynski@vion.com
- Follow my blog (***Generally, It Depends***): http://jimczuprynski.wordpress.com
- Follow me on Twitter (@**JimTheWhyGuy**)
- Connect with me on LinkedIn (**Jim Czuprynski**)