

AI/ML To The Rescue: How IT Operations Got Out Of A Monitoring Mess

Craig Shallahamer
craig@orapub.com
craig@viscosityna.com



This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

The Situation

Increase IT Operations Capacity

"Can you increase the quality of monitoring and alerting without increasing the associated operations cost?"



Why did the Operations team
believe this?

What is their plan?

Did it work?



My Machine Learning Journey



It all started in 1987




It all started in 1987, when I was intrigued with developing a program that could not only retain data but could answer my questions about that data.

At that time, I was using Borland's *Turbo Pascal* and noticed they released a Prolog product. It was amazing...

I was hooked!



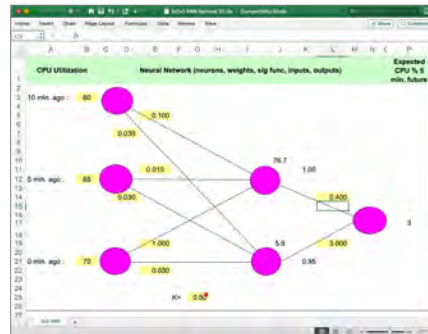
Then again in 1990



I created a program where I could play the game tic-tac-toe with a computer, but also enabled the computer to play itself... and learning through that process.



In 1997 Neural Networks...



I was obsessed is performance prediction... forecasting what was likely to happen or not happen in the future.



In 2001 at RMOUG

[illegible]

In 2001, after two years of rejections, I presented on ANNs at the RMOUG conference. My goal was to predict poor performance in the near future. While ANN is sound, I was never able to realize my "near future" dream. But I believe that dream could now be realized.




2007





This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>



Viscosity Core Pillars and Service Models

Proven Solutions

We're poised to handle all your database challenges flawlessly by drawing on proven, tested approaches and hands-on experience, such as zero downtime migrations.


Professional Services

No matter how complex your consulting needs, our team can make budgets work, reduce risk, eliminate surprises and make you a hero within your organization.

Managed Services

Our DBA's will ensure that your database environment receives the preventative care it requires for the future while promising to respond to your most critical needs with integrity 24/7.


DATA	APPS	CLOUD
Oracle & SQL Server Postgres	APEX	Azure Gold Partner
Performance Tuning	EBS	Cloud Migrations AVS
Data Replication	Web/Mobile Apps	Engineered Systems
Data Warehousing Analytics	.Net and C#	Oracle Cloud Platinum Partner
Data Integration	PeopleSoft	Google Partner
ERP Blue Prints	SAAS/PAAS	AWS Partner Hybrid Cloud

 **utoug**
utah oracle users group

Developer Resource Center /
Oracle ACE Program


Oracle ACE Program


The Oracle ACE Program recognizes and rewards community members for their technical contributions in the Oracle community.



Charles Kim,
CEO & Co-Founder


Twitter: @racdba


 Oracle ACE Director



Rich Niemiec,
Chief Innovation Officer


Twitter: @richniemiec


 Oracle ACE Director



Craig Shallahamer
Applied AI Scientist


Twitter: @orapub


 Oracle ACE Director



Sean Scott,
Consultant


Twitter: @oraclesean


 Oracle ACE Pro




Gary Gordhamer,
Consultant

Twitter: @ggordham

 Oracle ACE Pro

 **utoug**
utah oracle users group

 **VISCOSITY**

Big Business Problem

What drove the Operations to search for another monitoring and alerting strategy?

It goes beyond this specific project, because it business operations focused.



This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

The challenge begins here.



In the Oracle database systems **operations** world, significant effort goes into data collection and management, anomaly detection and alerting.



The challenge begins here.

In the Oracle database systems **operations** world,
significant effort goes into data collection and
management, anomaly detection and alerting.

And into the mix:
This organization has hundreds of Oracle systems.



This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
on January 26, 2023 for the Utah OUG (Utoug).
There is likely an updated version on the OraPub website at <https://www.orapub.com>
There is likely an updated version on the OraPub website at <https://www.orapub.com>

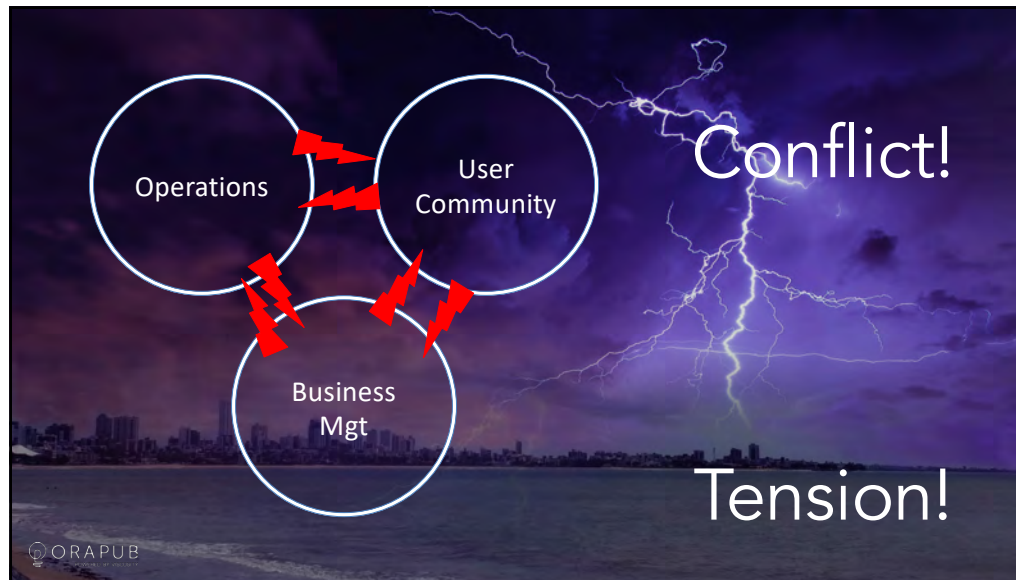
The challenge begins here.

In the Oracle database systems **operations** world,
significant effort goes into data collection and
management, anomaly detection and alerting.

And into the mix:
This organization has hundreds of Oracle systems.

*The Operations Director said, "No one person or
even my team can effectively monitor and respond to
all the systems alerts to the busines' satisfaction."*



This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

Big Business Objective

Increase IT Operations Capacity

"Can you increase the quality of monitoring and alerting without increasing the associated operations cost?"

 VISCOSITY

 **utoug**
utah oracle users group

"Can you increase the quality of monitoring and alerting without increasing the associated operations cost?"

Why does the Operations team believe this?

The Back Story...



Currently, if-then rules (heuristics) and statistical methods are used to detect anomalous for each specific system.

Why is this not working good enough?





This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

Rule Based Monitoring


Even with a single metric...

Rules start simple but as systems and workloads change so must the rules.

The natural result is a big pile of rules that become incredibly complicated and ultimately unmanageable and untrustworthy.

Imagine working with hundreds of metrics...



ORAPUB logo is in the bottom left corner. The UTOUG logo (Utah Oracle Users Group) is in the bottom right corner.



Statistical Based Monitoring

Statistical methods are used to detect high metric variations. The operations team uses over a thousand metrics.

It is difficult for a single metric to detect system anomalies. Combining metrics requires weighting and rules.



This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

What the Operations Team discovered

More systems results in more rules.

More simple rules results in more complexity.

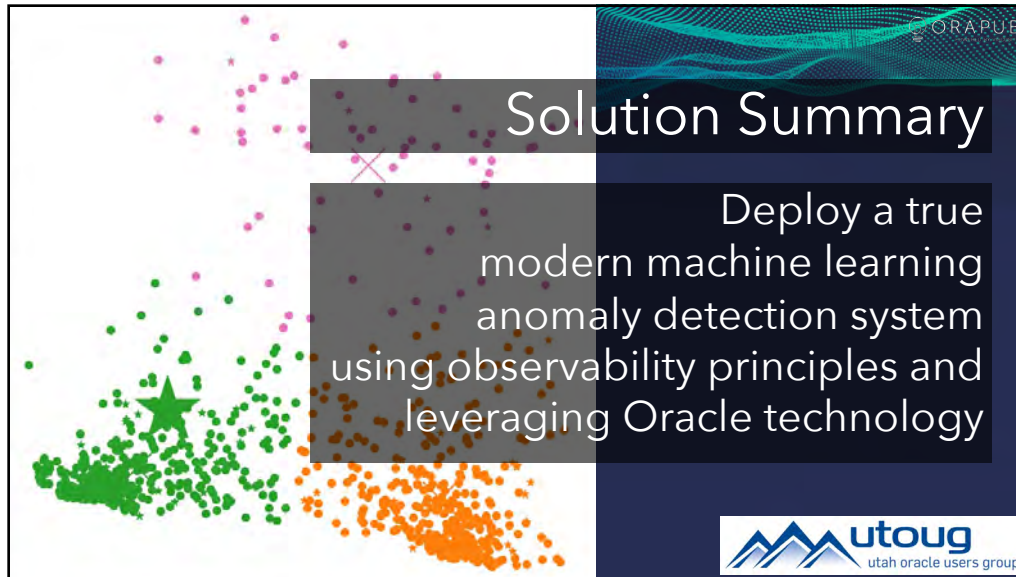
More statistical rules results in more complexity.

Real business workloads are incredibly complex.

Complexity kills scalability.







Solution Summary

Deploy a true modern machine learning anomaly detection system using observability principles and leveraging Oracle technology

ORAPUB

utoug
utah oracle users group

This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>



What is observability?

In IT and cloud computing, observability is the ability to measure a system's current state based on the data it generates, such as logs, metrics, and traces.

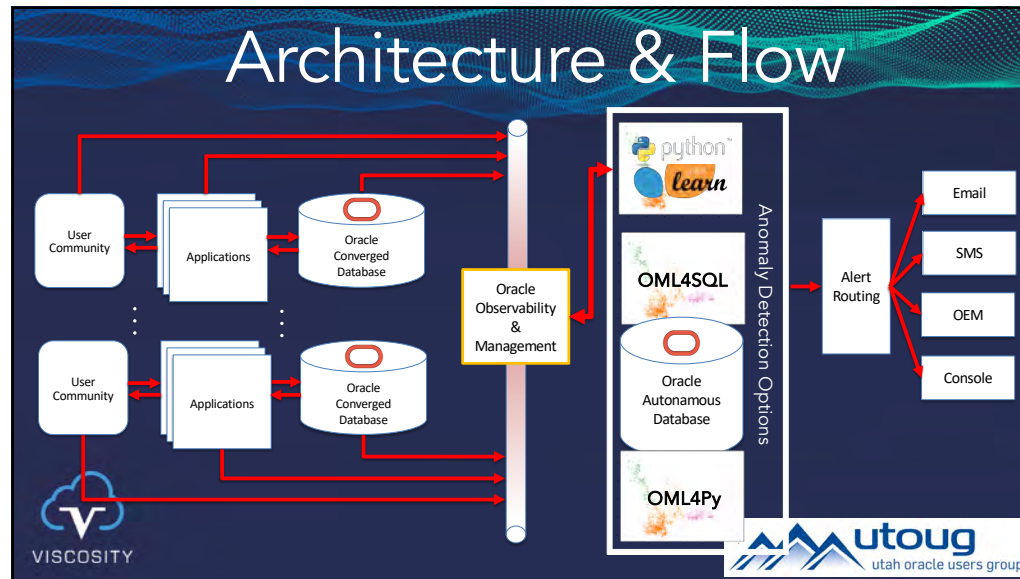
The goal of observability is to understand what's happening across all your environments and among all your technologies, so you can detect and resolve issues.

Oracle technologies are known for their deep instrumentation. But this also generates a tremendous amount of data.

Oracle's Observability and Management platform brings all this instrumentation together into a single source. An analysts gold mine.

ORAPUB

utoug
utah oracle users group

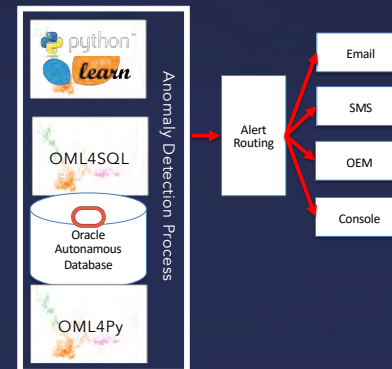


This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
 This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
 on January 26, 2023 for the Utah OUG (UTOUG).
 There is likely an updated version on the OraPub website at <https://www.orapub.com>
 There is likely an updated version on the OraPub website at <https://www.orapub.com>





Business Requirements #1

- Fast! Less than 2 minutes
- Alerts logged, emailed, SMS, OEM and console
- Alerts to contain diagnostic info... explainability



Business Requirements #2

- Understand real **business workloads**
- Use multiple **self adjusting** alert thresholds
- **View system** as multiple metrics



ORAPUB
POWERED BY VISCOSITY

Business Requirements #3

- ML Tools:
 - Python, Scikit Learn, Jupyter
 - OML4PY, OML4SQL
- Deployment:
 - Linux, Conda, Oracle DB
 - Autonomous DB w/OML4PY and OML4SQL
- Maintenance: self adjusting threshold, live tuning and easy restart
- Data: Oracle AWR Sysmetric from Oracle's Observability & Management platform.



This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

ML Particulars

- General ML process
- General anomaly detection strategy
- Self-Adjusting threshold determination
- Providing diagnostic clues; explainability



ML Particulars

General ML Process



This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

General Machine Learning Process

- **Load Data**
 - Pull new Sysmetric data and/or combine with existing
- **Preprocessing**
 - Denormalize Sysmetric data
 - Scale using standardization
 - Engineer new features
 - Date: hour of day, day of week, day of month
 - Moving average: 2 and 3
 - Power function: X^y where y is 0.5 and 2.0
 - Stepping/Lagging: 2 previous samples
 - Binning: percentile based (0.0, 0.25, 0.50, 0.75, 1.00)
 - Stationarity: remove trends
 - Feature selection using a High Correlation Filter, which resulted in around 7 to 15 features.
- **Build n clusters** using K-Means algorithm, usually 3 to 5 clusters
- **Thresholds** determination, usually above 98%-tile
- **Check for anomaly**, if true, send alerts



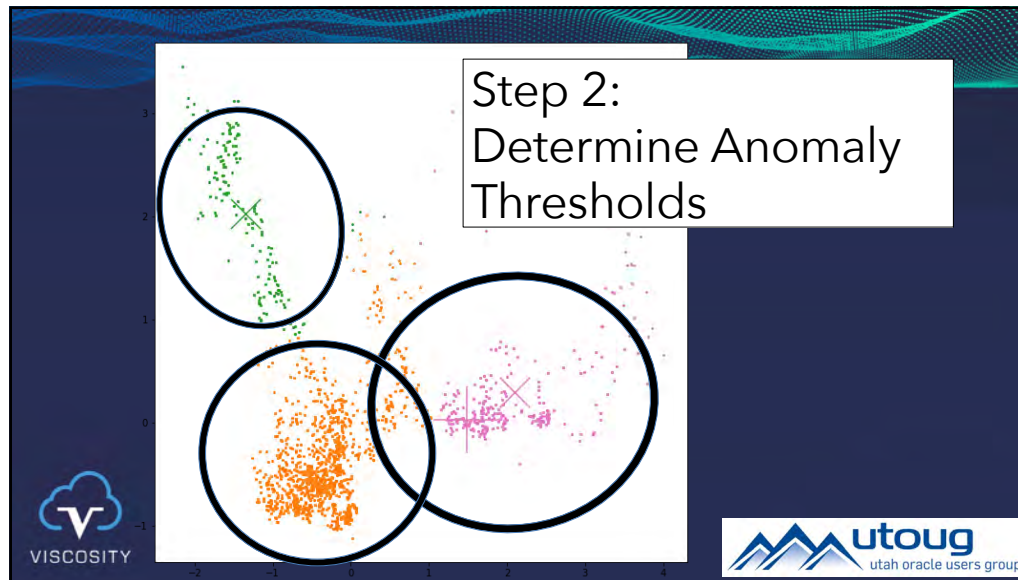
ML Particulars

General Anomaly Detection Strategy



Step 1:
Create Cluster System
K-Means algo with 3 clusters





This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>



ML Particulars


Self-Adjusting Threshold Determination



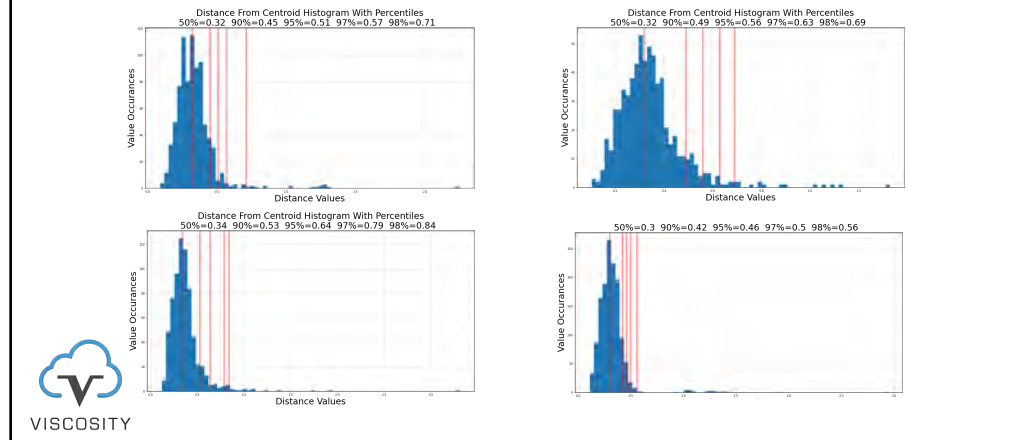
This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

Alert Threshold Determination is based on cluster specific distances



Percentile thresholds adapt to cluster specific distance values



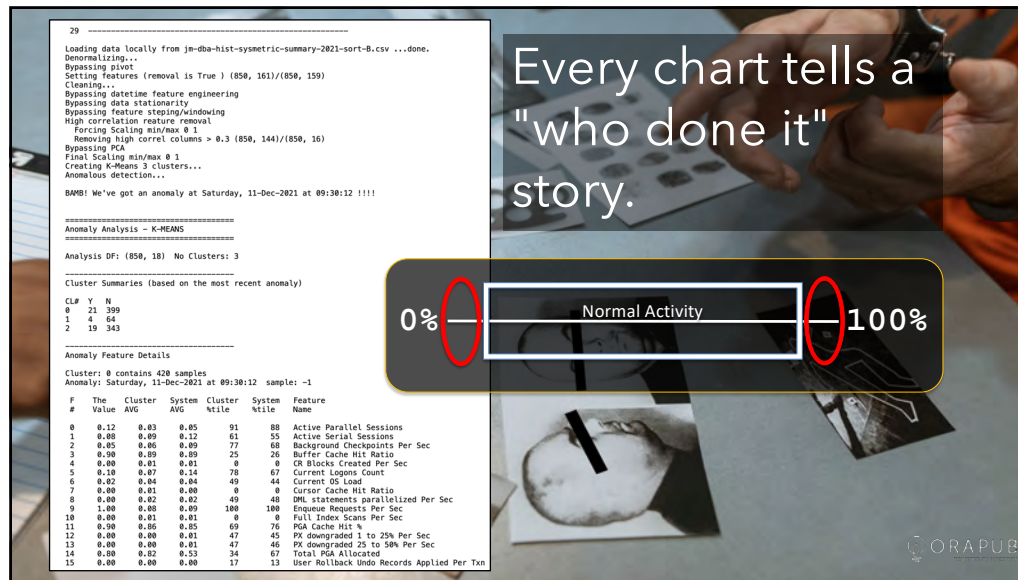
This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

ML Particulars

Providing Diagnostic Clues
Explainability... big deal





This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

CL# Y N

0	21	399
1	4	64
2	19	343

Anomaly Feature Details

Cluster: 0 contains 420 samples
Anomaly: Saturday, 11-Dec-2021 at 09:30:12 sample: -1

F #	The Value	Cluster AVG	System AVG	Cluster %tile	System %tile	Feature Name
0	0.12	0.03	0.05	91	88	Active Parallel Sessions
1	0.08	0.09	0.12	61	55	Active Serial Sessions
2	0.05	0.06	0.09	77	68	Background Checkpoints Per Sec
3	0.90	0.89	0.89	25	26	Buffer Cache Hit Ratio
4	0.00	0.01	0.01	0	0	CR Blocks Created Per Sec
5	0.10	0.07	0.14	78	67	Current Logons Count
6	0.02	0.04	0.04	49	44	Current OS Load
7	0.00	0.01	0.00	0	0	Cursor Cache Hit Ratio
8	0.00	0.02	0.02	49	48	DML statements parallelized Per Sec
9	1.00	0.00	0.00	100	100	Enqueue Requests Per Sec
10	0.00	0.01	0.01	0	0	Full Index Scans Per Sec
11	0.90	0.86	0.85	69	76	PGA Cache Hit %
12	0.00	0.00	0.01	47	45	PX downgraded 1 to 25% Per Sec
13	0.00	0.00	0.01	47	46	PX downgraded 25 to 50% Per Sec
14	0.80	0.82	0.53	34	67	Total PGA Allocated
15	0.00	0.00	0.00	17	13	User Rollback Undo Records Applied Per Txn

Important Points & Lessons Learned



This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com)
on January 26, 2023 for the Utah OUG (utoug).
There is likely an updated version on the OraPub website at <https://www.orapub.com>
There is likely an updated version on the OraPub website at <https://www.orapub.com>

Demanded Very Well Defined Scope

- Once people get a glimpse of ML, they go crazy...
"Let's predict when world peace will occur."
- Scope was documented and agreed upon.
- Scope was very tight.
 - Anomalous Oracle Database activity
 - Sysmetric data
 - Alerting: email, SMS, console
 - 15 minute AWR snap
 - 2 minutes to check for anomaly



Domain Expertise Was Essential

- Our team had both Oracle performance and applied ML expertise.
- Real-world forecasting and modeling experience was very helpful.
- True Data Scientist was not needed because we used common and existing algorithms, for both core cluster algo and preprocessing.



This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

Big Team Not Required

- Our team had only three people part-time over a 2 month period.
- The second month was more focused on "tuning" the system and alert details.
- We leveraged the hardware, security and network teams.
- Our IT requirements were tiny compared to other systems.



AWR Sysmetric data is gold!

- Already collected and available
- 150 metrics
- Most metrics in per-second
- Wide range of Oracle activity and performance metrics
- Covers Oracle, OS and some application centric metrics



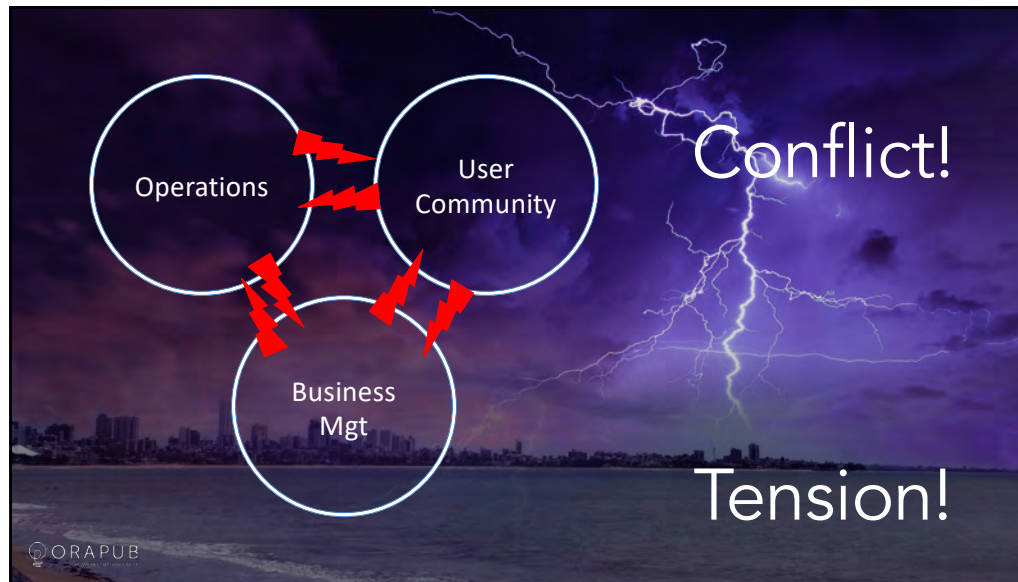
This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

Keep metric names

- From a pure ML perspective, it is tempting to reduce dimensionality that removes the Sysmetric feature names.
- Keeping metric names makes *next step diagnosis* easier and faster.
- Explainability is virtually lost without metric names.








How to get **all the materials** for this presentation and updated versions.
Plus **all** Viscosity's webinars and conference presentations.

After the conference:
www.orapub.com/members > Tools & Presentations > Presentations



Get Trained!



MACHINE LEARNING 1
Anomaly Detection

ORAPUB



MACHINE LEARNING 2
Performance Prediction

ORAPUB




This presentation was given by Craig Shallahamer (craig@orapub.com / craig@viscosityna.com) on January 26, 2023 for the Utah OUG (UTOUG).

There is likely an updated version on the OraPub website at <https://www.orapub.com>

<https://www.orapub.com/ml-ecourse>



Machine Learning
For
Oracle Professionals
E-Course

ORAPUB











FREE from OraPub.com


OraPub is committed to cross-training
Oracle Professionals in Machine Learning


Developer Resource Center /
Oracle ACE Program

Oracle ACE Program

The Oracle ACE Program recognizes and rewards community members for their technical contributions in the Oracle community.

 <p>Charles Kim, CEO & Co-Founder</p> <p>Twitter: @racdba</p> <p> Oracle ACE Director</p>	 <p>Rich Niemiec, Chief Innovation Officer</p> <p>Twitter: @richniemiec</p> <p> Oracle ACE Director</p>	 <p>Craig Shallahamer Applied AI Scientist</p> <p>Twitter: @orapub</p> <p> Oracle ACE Director</p>
 <p>Sean Scott, Consultant</p> <p>Twitter: @oraclesean</p> <p> Oracle ACE Pro</p>	 <p>Gary Gordhamer, Consultant</p> <p>Twitter: @ggordham</p> <p> Oracle ACE Pro</p>	

 VISCOSITY

 utoug
utah oracle users group

Thank You!

 VISCOSITY

 utoug
utah oracle users group

AI/ML To The Rescue: How IT Operations Got Out Of A Monitoring Mess

Craig Shallahamer
craig@orapub.com
craig@viscosityna.com

