IBM Software Group

# Filling In The IT Systems Management White Space Gap

*Ed Woods - IBM Corporation*

**Session #16331**
**Tuesday, March 3rd: 1:45 PM - 2:45 PM**

# Agenda

- Introduction
- Defining the ‘White Space Gap’
- Common Challenges
  - What drives and opens the gap
- How to address and bridge the gap
  - A strategy to address the unknown unknowns
- A look at Analytics

# Are There Gaps In You IT Systems Management Strategy?

- Many well documented methods and best practices
- The challenge may be identifying the shortfalls and filling in the gaps
  - It's what you miss, the "white space gap", that may cause the largest issues
- Every IT installation is unique
  - Each environment poses its own set of challenges
  - There are best practices that may be applied to most situations

# What Is The 'White Space Gap'?

- One example
  - The wasted space on a printout (a common Google search)
- A more relevant 'white space gap' definition
  - An informational or operational disconnect
    - A disconnect between various groups in an organization
      - A disconnect in determining requirements and defining processes
        - Assumed requirements versus actual requirements?
  - Gaps in knowledge, process, or procedure
    - Gaps may result in issues, outages, and longer MTTR

# The Challenge Of The Eye Of The Beholder

- May take many forms
  - Various SME groups, technical versus operations, technical versus line of business
    - Different requirements
    - Different priorities
- In most shops monitoring and management still tends to be ‘silo’ in nature
  - Focus is usually on the SME (Subject matter expert)
  - Different platforms, components and core technologies to manage
    - Often different tools and methodologies

# The Challenge Of Complex Applications

- Most new applications are composite in design and deployment
- Problem analysis will often be more complex
  - Problem resolution may require many groups and SMEs
- The management of composite applications may be challenged by the common issue of the white space gap
  - Problem analysis is often done in a ‘silo’ fashion
    - Tools are traditionally been used in a ‘silo’ fashion
    - Poses challenges in terms of identifying and resolving complex application issues
    - Islands of automation
  - Poses challenges when trying to become more proactive

# What You May Not Know
# What Is Happening Outside Of The z/OS Platform?

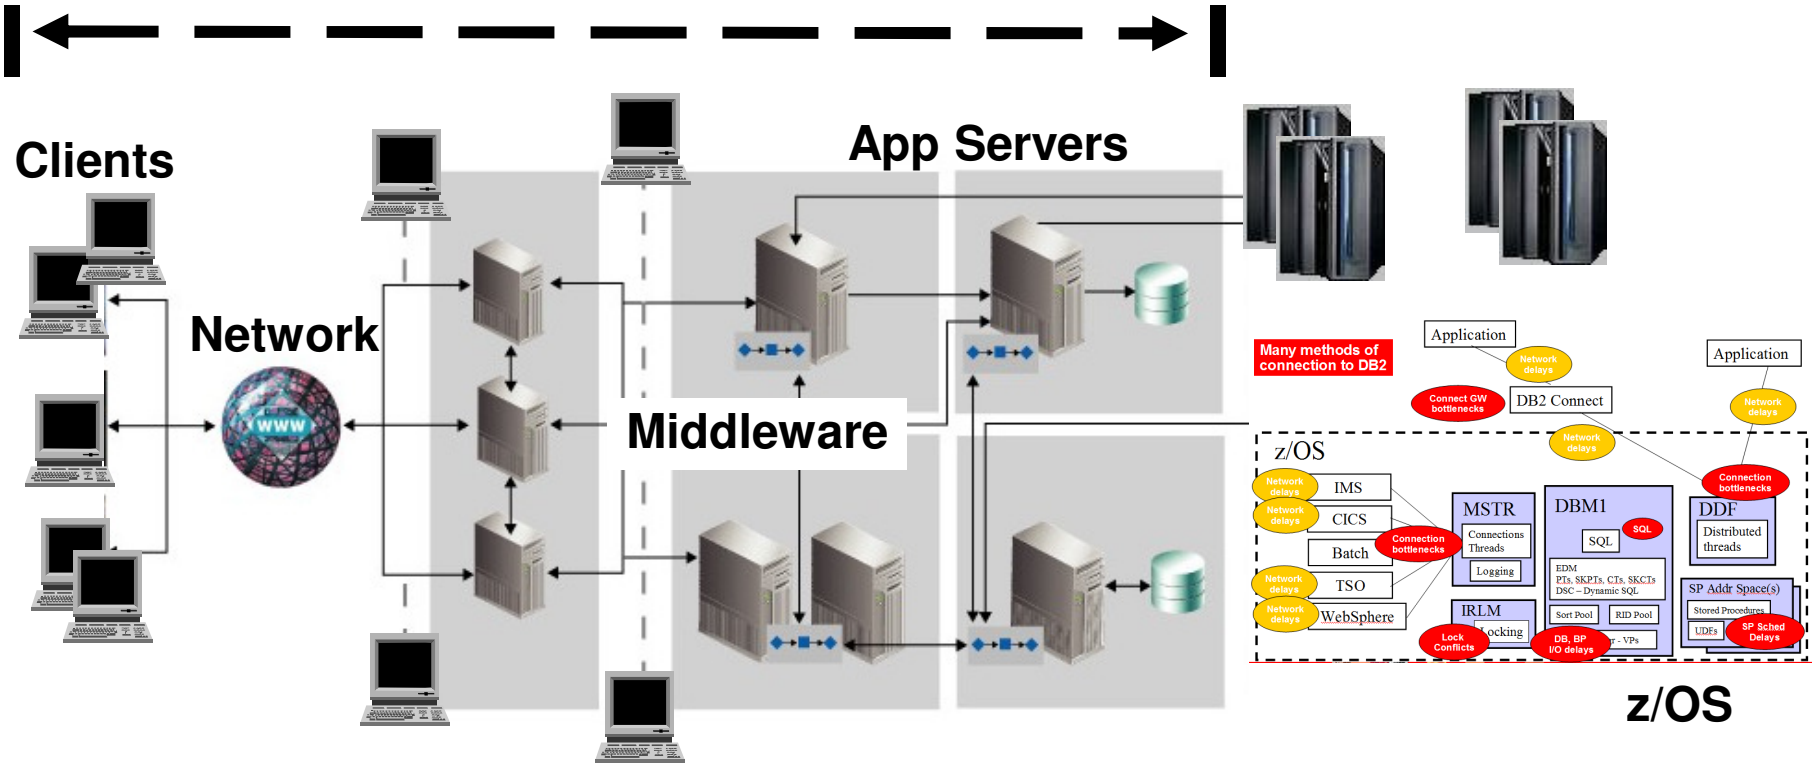


- Most applications are multi-layer and multi-component
  - Distributed servers, application servers, middleware, client layer, and network
- Where is the bottleneck? Is there queuing and at what level?

# The Challenge
# Islands Of Automation

- Many technical platforms, components and core technologies to manage
  - Often times each with it's own group of Subject Matter Experts (SMEs)
  - Each with it's own set of management tools
- The problems
  - Complex SME tools with different User Interfaces
  - SME tools that do not integrate or share information
    - More difficult to navigate
    - More difficult to do problem identification, isolation, and resolution
  - More challenging to automate corrective actions without clearly defined integration
    - More reliance on manual intervention

# The Challenge Of The Inherited Workload

- Acquisitions and mergers drive technical challenges
  - Potentially different workload types and formats
    - Different platforms, different core technologies
  - Potentially different management methodologies
    - May use different tools and management technologies
    - May use different management techniques and approaches
- Turnover in personnel and passage of time
  - Processes may stay in place that may need to be re-designed or eliminated
- Multiple core technologies make management more complex

# It's What You Don't Know That Can Hurt You

- There are known knowns. These are things we know that we know. There are known unknowns. That is to say, there are things that we know we don't know. But there are also unknown unknowns. There are things we don't know we don't know.
  - Donald Rumsfeld – former US Secretary of Defense

# What You Don't Know – An Example
# How An Unknown Unknown May Impact Availability

- Share Conference presentation, John Tobler & Nigel Slinger
  - "What Happened to My DB2? The Top Missteps in High Availability"
    https://share.confex.com/share/121/webprogram/Session13729.html

# What You Don't Know – An Example
# How What You Don't Know May Impact Availability

- Share Conference presentation, John Tobler & Nigel Slinger
  - "What Happened to My DB2? The Top Missteps in High Availability"
    https://share.confex.com/share/121/webprogram/Session13729.html
- ***Missing critical messages is among the most common source of DB2 outages***
  - Are you monitoring all the most critical messages?

# What You May Not Know
# Missing Critical DB2 Messages – Example Scenarios

- Example scenario - logging
  - DB2 is out or almost out of active log space
    - System may hang as logging comes to a halt
    - Messages DSNJ110E and DSNJ111E indicate the issue
    - May be due to insufficient log space or an offload failure (DSNJ115I)
- Example scenario – real storage utilization
  - DB2 is consuming large amounts of real storage
    - DB2 or entire LPAR may crash
  - DSNS003I message (more than 80% REALSTORAGE_MAX used)
    - REALSTORAGE_MAX is a DB2 zparm setting
  - Actions include reviewing DB2 storage usage with real time monitor
    - Presentation on DB2 memory management
      - http://www.slideshare.net/Caroldm/db2-10-memory-management

# The Challenge Of Message Management

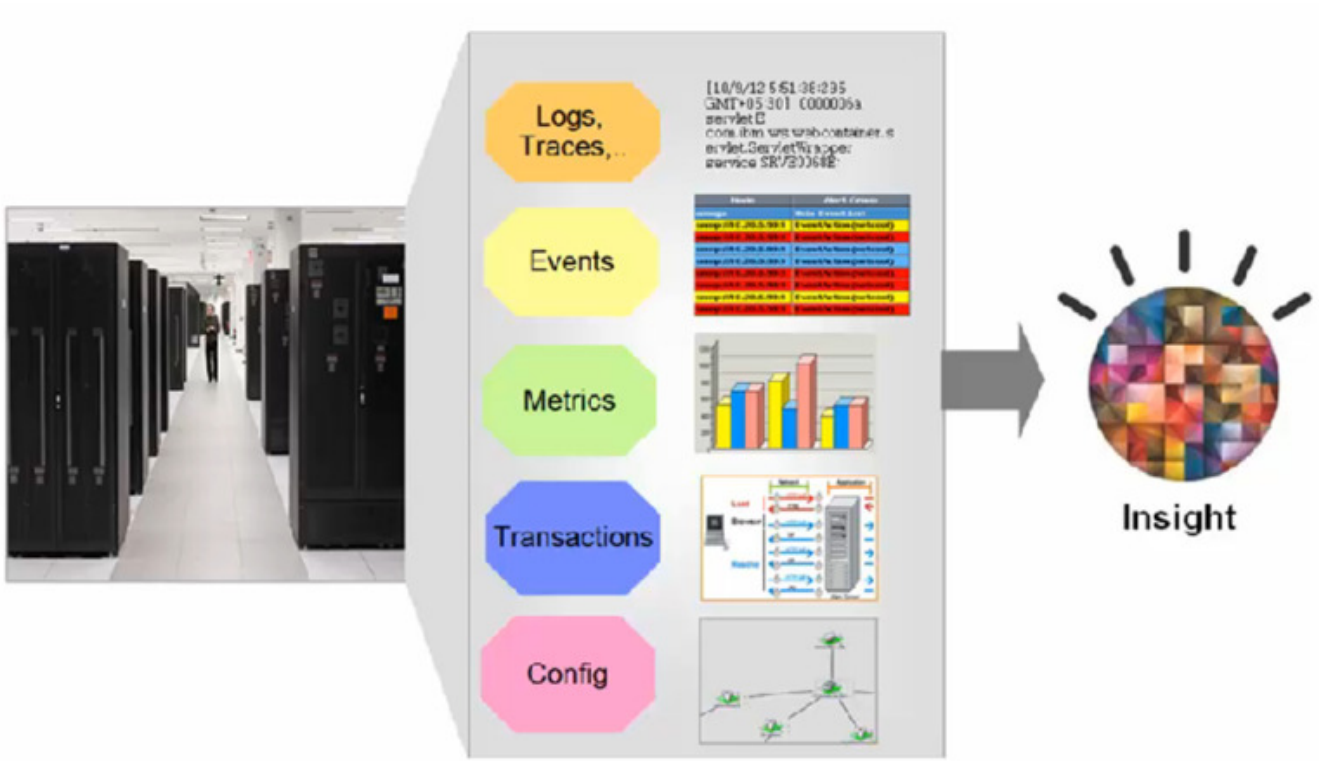


- Operators and SMEs are overwhelmed with volumes of log and message data
  - Manual process to determine the cause, location and scope of a problem
  - Example - an enterprise with 5000 servers may generate over 1 TB of log data daily

# Bridging The 'Whitespace Gap' A Methodology

- Concentrate
  - Focus on clarifying and defining the 'Knowns'
- Integrate
  - Eliminate the 'Known Unknowns'
    - Eliminate gaps through integration of information and management processes
- Automate
  - Identify and eliminate the 'Unkowns' through Analytics and Automated processes

# Concentrate
# Start With The ‘Knowns’

- Start with the knowns
  - Subject Matter Experts (SMEs) will always be essential
    - Knowledgable SMEs are key to knowing best practices
- Ensure you have the most appropriate information
  - You cannot always predict defects (Just assume there are some)
  - How well are you prepared?
    - Are you leveraging your technology?
      - Example - Do you know where all your JCL is?
        - Traces, Reports, Diagnostics
      - What history are you collecting?
      - What is your alert strategy?

# Concentrate
# Start With What You Know You Know
# Many Industry Standard 'Best Practices'

- Performance analysis best practices are well established for z/OS systems and its various subsystems
  - Many sources of well documented management and tuning best practices for DB2, z/OS, CICS, IMS
  - z/OS, DB2, IMS, CICS have decades established best practices
- Sources include
  - IBM Red Books - http://www.redbooks.ibm.com/
  - IDUG presentations - www.idug.org
  - Share conference - www.share.org
  - Computer Measurement Group - CMG – www.cmg.org

# Concentrate
# Define Your Numbers
# Is It A Good Number? Is It A Bad Number?

- Established 'best practices' document many recommended numbers and settings
- Many numbers reflect documented concerns
- Some numbers may not be so clear
  - Workloads are unique in each installation
    - Examples – In-DB2 times, CPU times, getpage counts may vary widely by shop, by workload and by workload type
- How do you know what is a bad number?
  - The first question should be – "What is a good number?"
  - How does the application appear when things are running well?
  - Define your baseline

# Integrate – Eliminate 'Unknowns' Bridge The Gap By Consolidating Information

- Eliminate 'Islands of Automation'
  - SMEs and SME tools are critical
  - Where feasible look at more integrated tool sets
- Integration expedites root cause analysis
  - Complex composite applications may require multiple tools
- Integration helps to address the key questions
  - Is the problem DB2? Is it in the SQL? Or the network? Or z/OS/ Or z/OS WLM? Is it an application issued? Is there an issue with middleware? Or CICS? Or IMS? Or distributed OS?
- Where feasible share information between tools
  - Integrated displays (dashboards), consistent monitoring and alerting

# Integrate
# The Value Of Dashboards


- We use dashboards every day
  - Integration of information is essential to effectively manage complex technology
- Dashboards provide important information
  - Resource status and resource utilization
  - Resource availability and projected availability
  - Problem identification and notification
- An effective dashboard provides the information in an easy to recognize and understandable format
  - Lights, graphs, indicators, gauges, status items

# Integrate
## An Effective Dashboard Helps Bridge The 'Whitespace Gap'

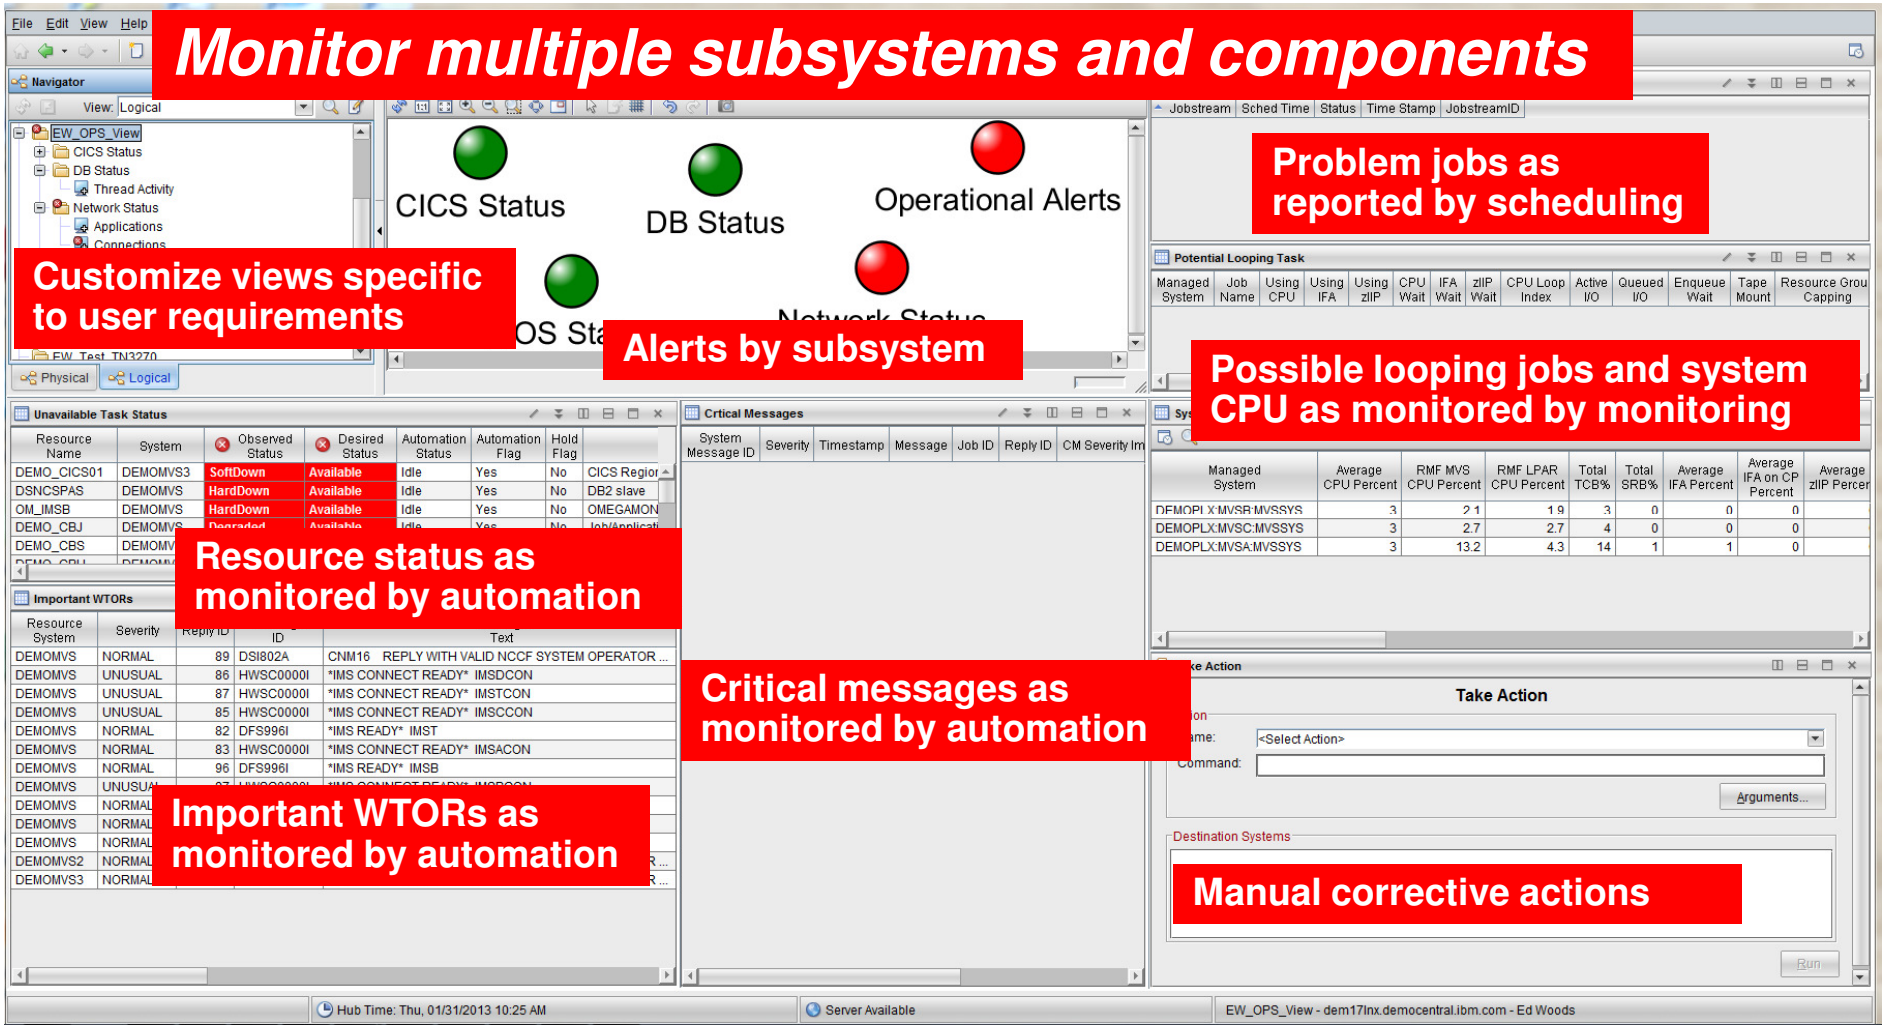

# Integrate
# Important Characteristics Of A Dashboard

- Integration
  - Pull critical information together to the “single pane of glass”
    - Gather essential information from a variety of sources
  - End to end views for complex composite business applications
- Flexibility
  - Different views for different audiences
    - Management and Line of business/end users
    - Operations
    - Help desk
    - Technical Subject Matter Expert (SME) views
  - Optimize the views as the environment or requirements change
- Ease of Use
  - Eliminate the clutter and tune out the “noise”
  - Focus on critical metrics

# Automate
## Finding The 'Unknown Unknowns" With Predictive Analytics

- An area of analysis that deals with extracting information from data and using it to predict future trends and behavior patterns
- Relies on capturing relationships between explanatory variables and the predicted variables from past occurrences
  - Exploit the information to predict future outcomes
- Accuracy and usability of results will depend greatly on the quality of data analysis and the quality of assumptions
- Predictive analysis is used in many facets of business
  - Common example would be credit score
    - Function of many data items
    - Income, payment history, amount of outstanding debt, etc…

# Steps In The Predictive Analytics Process

- Data organization and cleansing
  - Identify data sources
- Data Mining
  - Analysis of data to identify underlying trends, patterns, or relationships
  - Identify data to be used to develop the predictive model
- Model Development - Regression models
  - Regression modeling describes the relationship between dependent variable (the variable to be predicted) and independent explanatory variables
  - Regression models imply some level of causation (versus correlation)

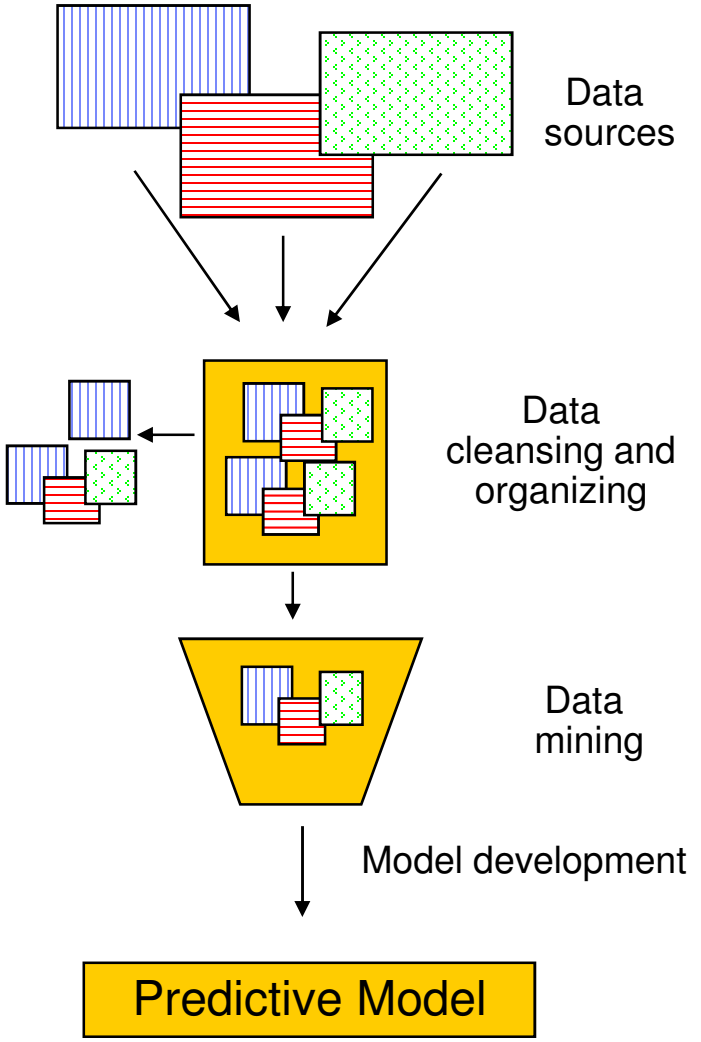

# A Goal For Many Shops
# Make Systems Management More 'Proactive'

- In many shops systems management tends to be done 'ad hoc'
  - Some alert generation – varies by shop
    - Some shops very alert driven – many are not
  - Often notification consists of 'call the help desk'
- Many customers want to be more 'proactive'
  - Definition of proactive may vary
    - Proactive for some installations may mean more rapid alert and notification of technical and/or business application issues
    - Proactive for some installations may mean notification ***prior*** to the problem
      - Alert when utilization indicates a potential issue in the future
      - Alert when I'm within 90% of the wall

# The Typical Monitoring Paradigm

- Traditional monitoring strategy
  - Monitor key resources based upon established 'best practices'
    - Resource utilization and resource bottlenecks
  - Monitor performance and availability
    - Key Performance Indicators (KPIs)
      - Examples – Response time, transaction rate, technical component, software subsystem, or business application availability
    - Monitor based on established SLA's
  - Alert notification about performance bottlenecks and outages
    - Notification via monitoring UIs, paging, emails
- Real time monitoring versus historical
  - Real time monitoring for current utilization and status
  - Historical data collection for trending and after the fact analysis

*Most shops monitor – but how predictive is it?*

# Problem Analysis And Resolution In Many IT Environments

- Problem identification and notification may be ad hoc
  - Alert notification via phone calls, emails, or paging
- Problem analysis is often after the fact
- Problem analysis and resolution often involves rounding up the usual suspects (and getting them to confess)
- Issue resolution relies heavily on the knowledge and intuition of the technical staff
  - Knowledge of the systems and business applications
  - Understanding ***complex problems*** will be ***multivariate*** in nature

# The Problem: Traditional Monitoring Approaches Have Limitations

- Many tools, data sources and metrics available
  - Many are Resource/Single Metric Focused (Univariate)
- Often many missed, or misinterpreted events
- In many shops not enough time, and/or resources to correlate completely
  - May require many people and groups to collaborate effectively
  - Many resources and no obvious resource inter-relationships

Univariate - refers to an expression, equation, function or polynomial of only one variable

Multivariate - encompasses the simultaneous observation and analysis of more than one statistical variable

# Why Multivariate Analysis?

- Multivariate analysis expands the relevance of predictive analytics
  - Provides context through correlation
- Example – credit rating metrics
  - Payment history – how relevant if I do not consider other metrics?
  - Income – again how relevant if I do not consider other metrics?
- Multivariate is important for IT Service Management
  - Many business applications are composite in nature
    - Many components, platforms, core technologies
  - Many critical resources are shared and inter-related
    - Mainframes support many applications
    - Networks may support a wide array of workloads

# Multivariate Analysis In An IT Context – An Example

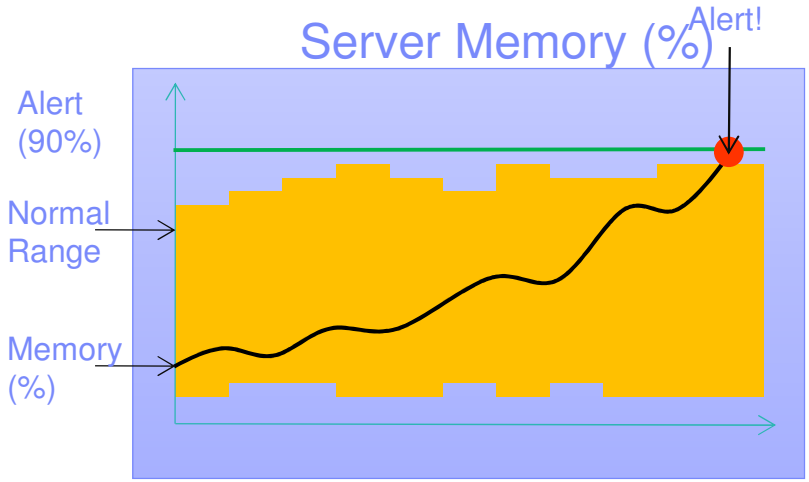


Static Threshold = Short Warning

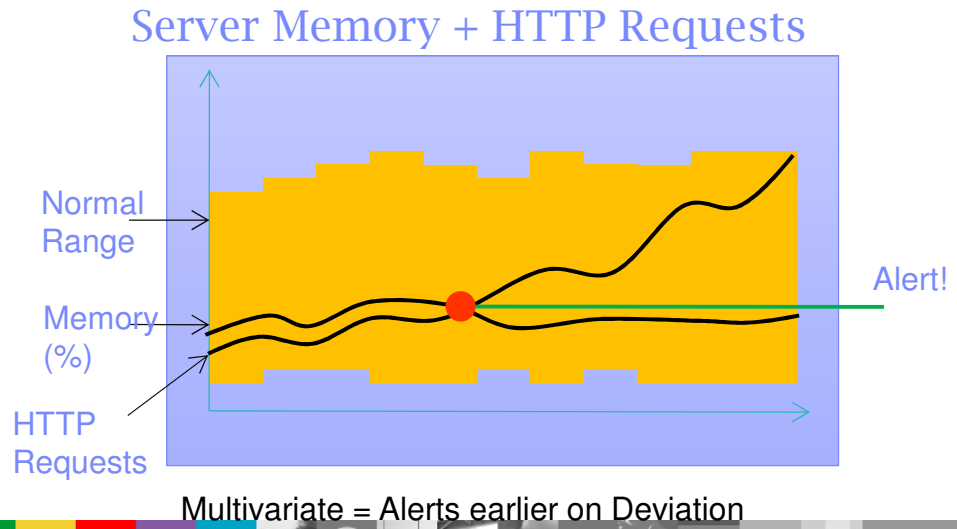


Multivariate = Alerts earlier on Deviation

Multivariate analytics detects problems sooner by detecting the deviation of metrics that normally move together.

For example:

- Memory consumption is normally correlated to HTTP requests

- But when memory deviates from HTTP Requests, as would happen with a memory leak, this indicates a problem and an alert is generated.

- The alert is generated much sooner than waiting for a static threshold violation.

This advanced warning time helps you become proactive and mitigate damage before customer service is impacted.

It also help reduce threshold alerts due to normal threshold violation correlated with HTTP Requests.

# Examples Of IT- related Multivariate Metrics

- DB2 example
  - DB2 object lock conflict >>
    - long running SQL call >> high In-DB2 time >> longer thread elapsed time >> longer DB2 query time
- IMS example
  - High IMS message region occupancy time >>
    - IMS transactions queued >> longer IMS transaction scheduling time >> longer IMS response time >> lower IMS transaction processing rate
- MQ example
  - Lower MQ message input rate >>
    - Higher MQ message queue depth >> lower transaction processing rate >> longer CICS/IMS transaction response time

# Understanding Critical z/OS Messages About zAware

- IBM zAware – IBM System z Advanced Workload Analysis Reporter
- Monitors z/OS OPERLOG including all messages written to z/OS console, including ISV and application generated messages
  - Early detection and focused diagnosis can help improve time to recovery
- Technology based on machine learning developed by IBM Research
  - Pattern recognition techniques look at the health of a system to pinpoint deviations from the ‘norm’
  - High speed analytics facilitates the ability to consume large quantities of message logs
- Allow establishment of procedures to prevent reoccurrence
- IBM Red Book - http://www.redbooks.ibm.com/redbooks/pdfs/sg248070.pdf

# Inside IBM zAware

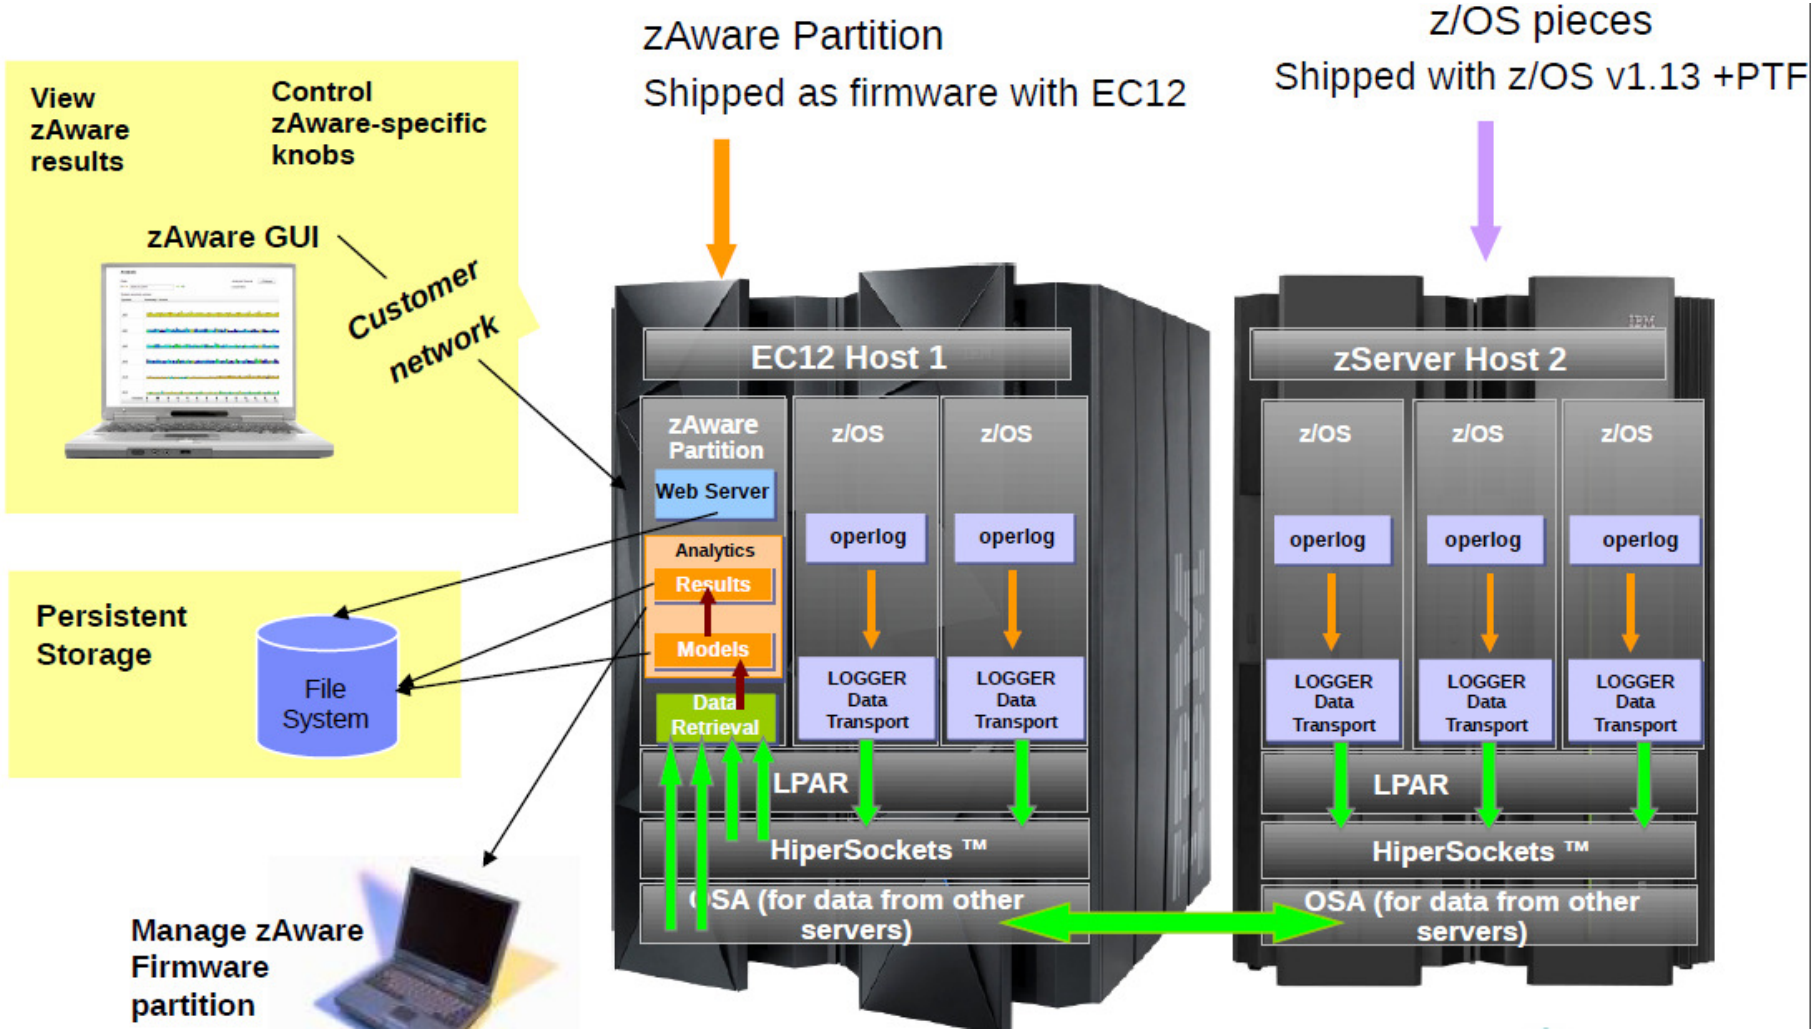

# How zAware Operates

- OPERLOG is processed per-system
  - zAware recognizes any well-formed message Ids, including IBM and non-IBM products and customer applications
- zAware ***builds a model*** of normal behavior based on the last 90 days
  - Called “Training”, automatically trains every 30 days, customizable
  - Unusual days can be excluded from future models
- Real-time OPERLOG data is compared to the model
- ***Assigns a message anomaly score*** to indicate deviation from the model
  - Rare messages, Out of context from normal patterns, high counts
- Uses z/OS-specific knowledge to influence the scores
  - Generates an interval anomaly score
- Provides a GUI for analysis
  - GUI shows number of unique message IDs and interval anomaly score
- Provides API for real time monitoring

# Messages Provide Important Input To Analytics Process

View, track, trend, and analyze critical messages

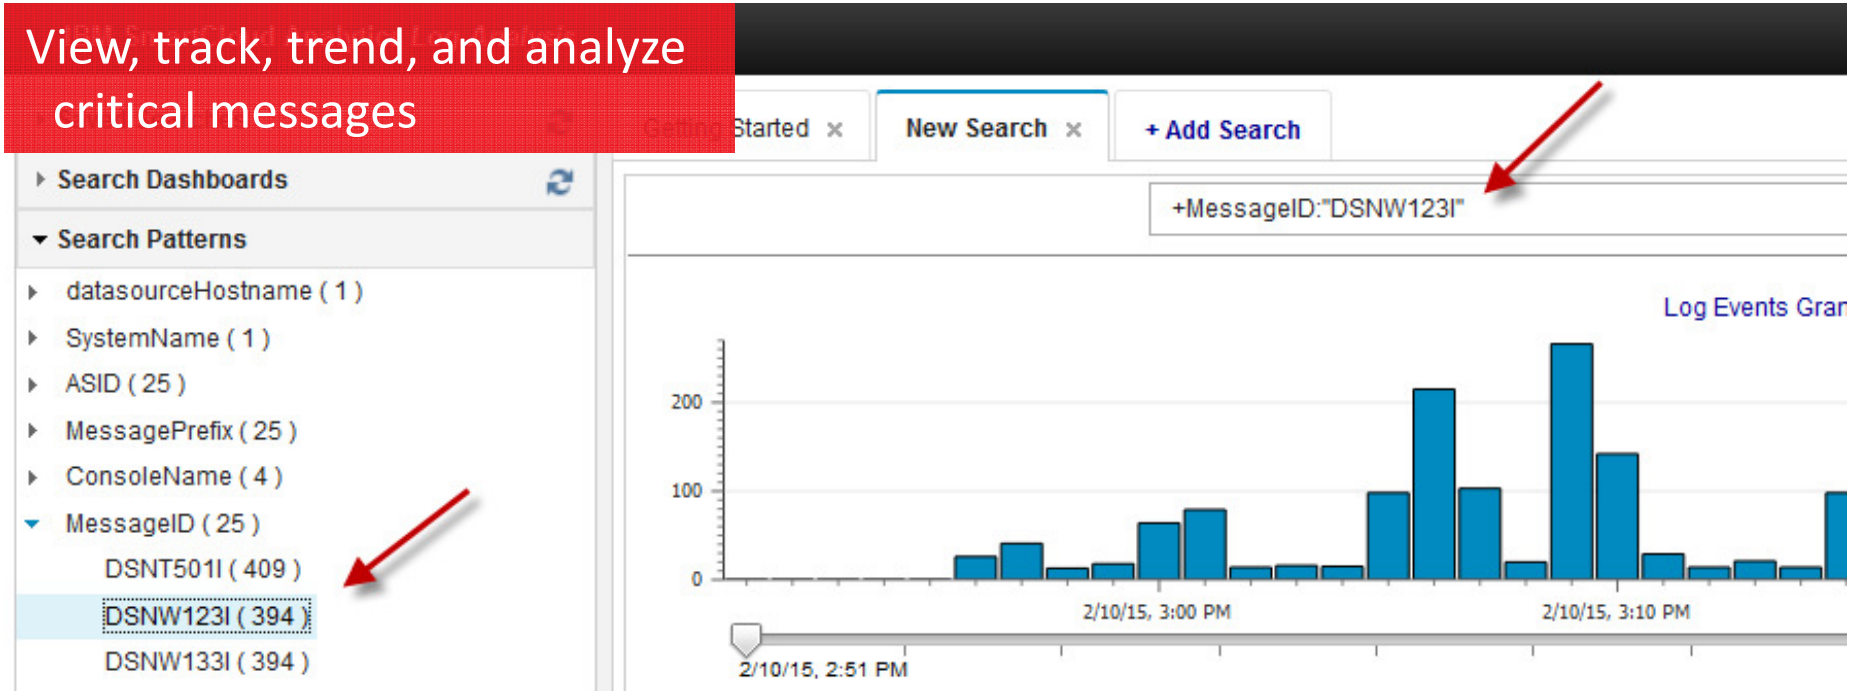


- Messages highlight issues and events IT platforms
  - z/OS subsystem messages, application errors, abends, notifications, alerts

# Predictive Analytics
# Categories And Sources Of Information For Analysis

- Messages and events
  - System console messages
    - z/OS console messages, CICS, IMS, MQ messages
    - System message logs from open systems sources
  - Application message logs (including error messages)
  - Various abend and error messages
- Alerts
  - Alerts from various monitoring sources
- Monitored metrics
  - Real time monitoring – critical system and resource metrics
  - Historical monitoring and collection
    - Critical system and resource metrics collected for historical analysis
    - Detail and summary historical data

# Operational Analytics Of Transaction Workload

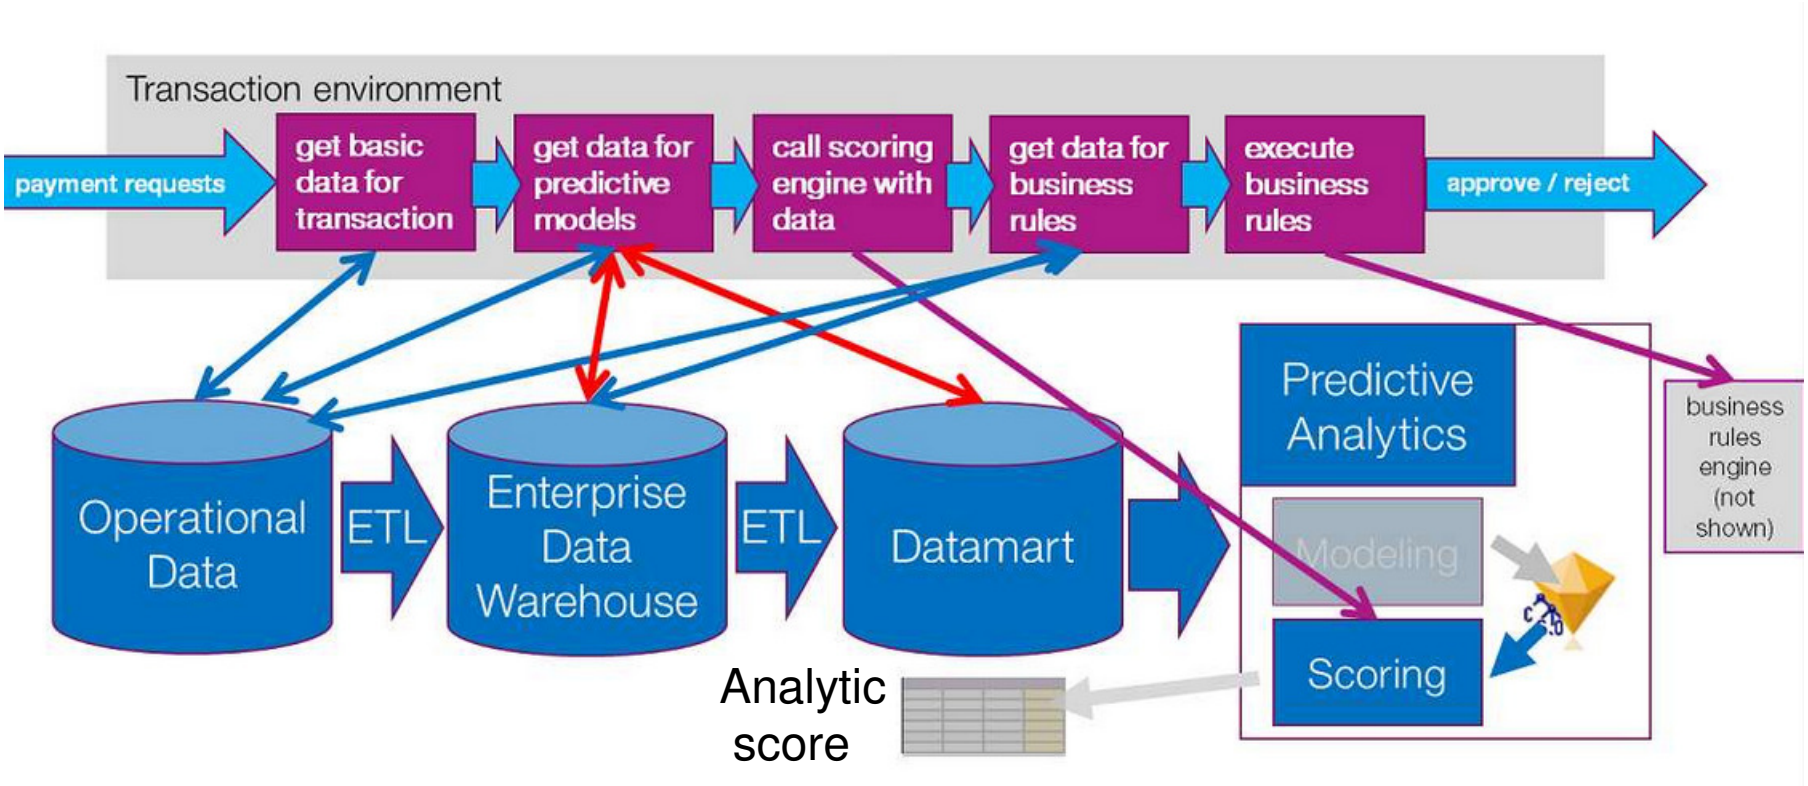


- Analytics may be applied to transactional workloads
- Use analytics to analyze and score workload processing

# Summary
# Filling In The White Space Gap

- Every shop has some level of ‘White Space Gap’
  - Organizational disconnect and human nature
- The gap can be bridged
  - Requires effort and planning
- It all begins with a process
  - Define and understand your ‘knowns’
    - Define your best practices, have clearly defined sources of information
  - Try to determine and define your ‘known unknowns’ and try to come up with a way to unearth your ‘unknown unknowns’
    - Concentrate, Integrate, Automate
    - Incorporate Analytics into your management methodology

# Thank You!

# Check Out My Blog

## http://tivoliwithaz.blogspot.com