IBM Software Group

Predictive Analytics And IT Service Management

Session – 11479
Wednesday, August 8th
1:30 – 2:30 PM

Tivoli software

Ed Woods
Consulting IT Specialist
IBM Corporation

Agenda

- What is Predictive Analytics?
- Examples
- How is predictive analytics relevant to IT Service Management?
- Typical monitoring and management paradigms
- Real time information versus historical data collection
- Univariate versus multivariate analysis
- Examples of relevant metrics
- Where to begin

What Is Predictive Analysis?

- 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)

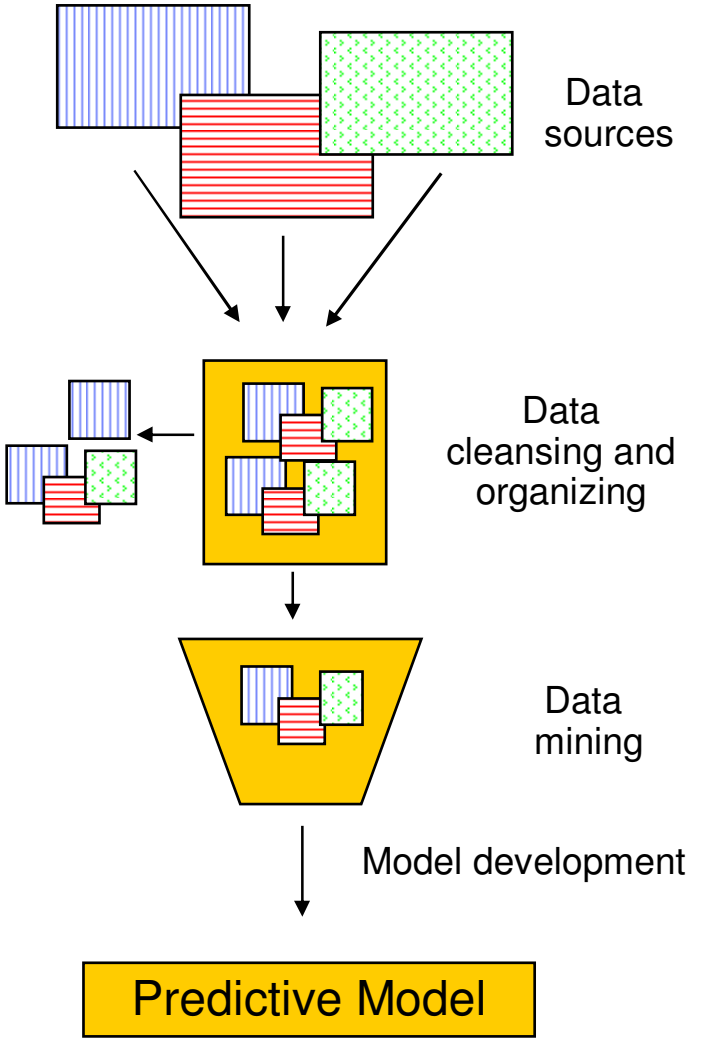

Predictive Analytics
About Regression Models And Types Of Models

- Regression models are the core of predictive analytics
- A wide variety of models can be applied
  - Linear regression model
    - Analyzes the relationship between the response or dependent variable and a set of independent or predictor variables
  - Partial or Stepwise regression
    - Modeler does not specify all the explanatory variables
    - Variables are added iteratively
  - Logit or Probit regressions
    - Allow one to predict a discrete outcome (yes/no) from a set of variables
  - Time series models
    - Used for predicting or forecasting the future behavior of variables
    - Data points taken over time may have an inherent time relation
    - Developed to decompose the trend, seasonal and cyclical component of the data
  - Many more models……

Examples Of Predictive Analytics Commonly Applied to IT

- Performance modeling
  - z/OS workload right sizing and load balancing
    - Model workload placement using SMF data as input
- Trending and forecasting of workload/resource utilization
  - Workload performance trends
    - Discern patterns in resource utilization
    - Capacity planning
  - The common question >> When will a critical resource reach breaking point?
- 'What If' Analysis examples
  - DB2 buffer pool analysis
    - DB2 performance trace data to determine optimal pool sizing and object placement
  - DB2 SQL and object tuning
    - DB2 Explain analysis based on DB2 Catalog statistics and SQL call changes

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?

Real Time Monitoring Provides A Starting Point For Analysis

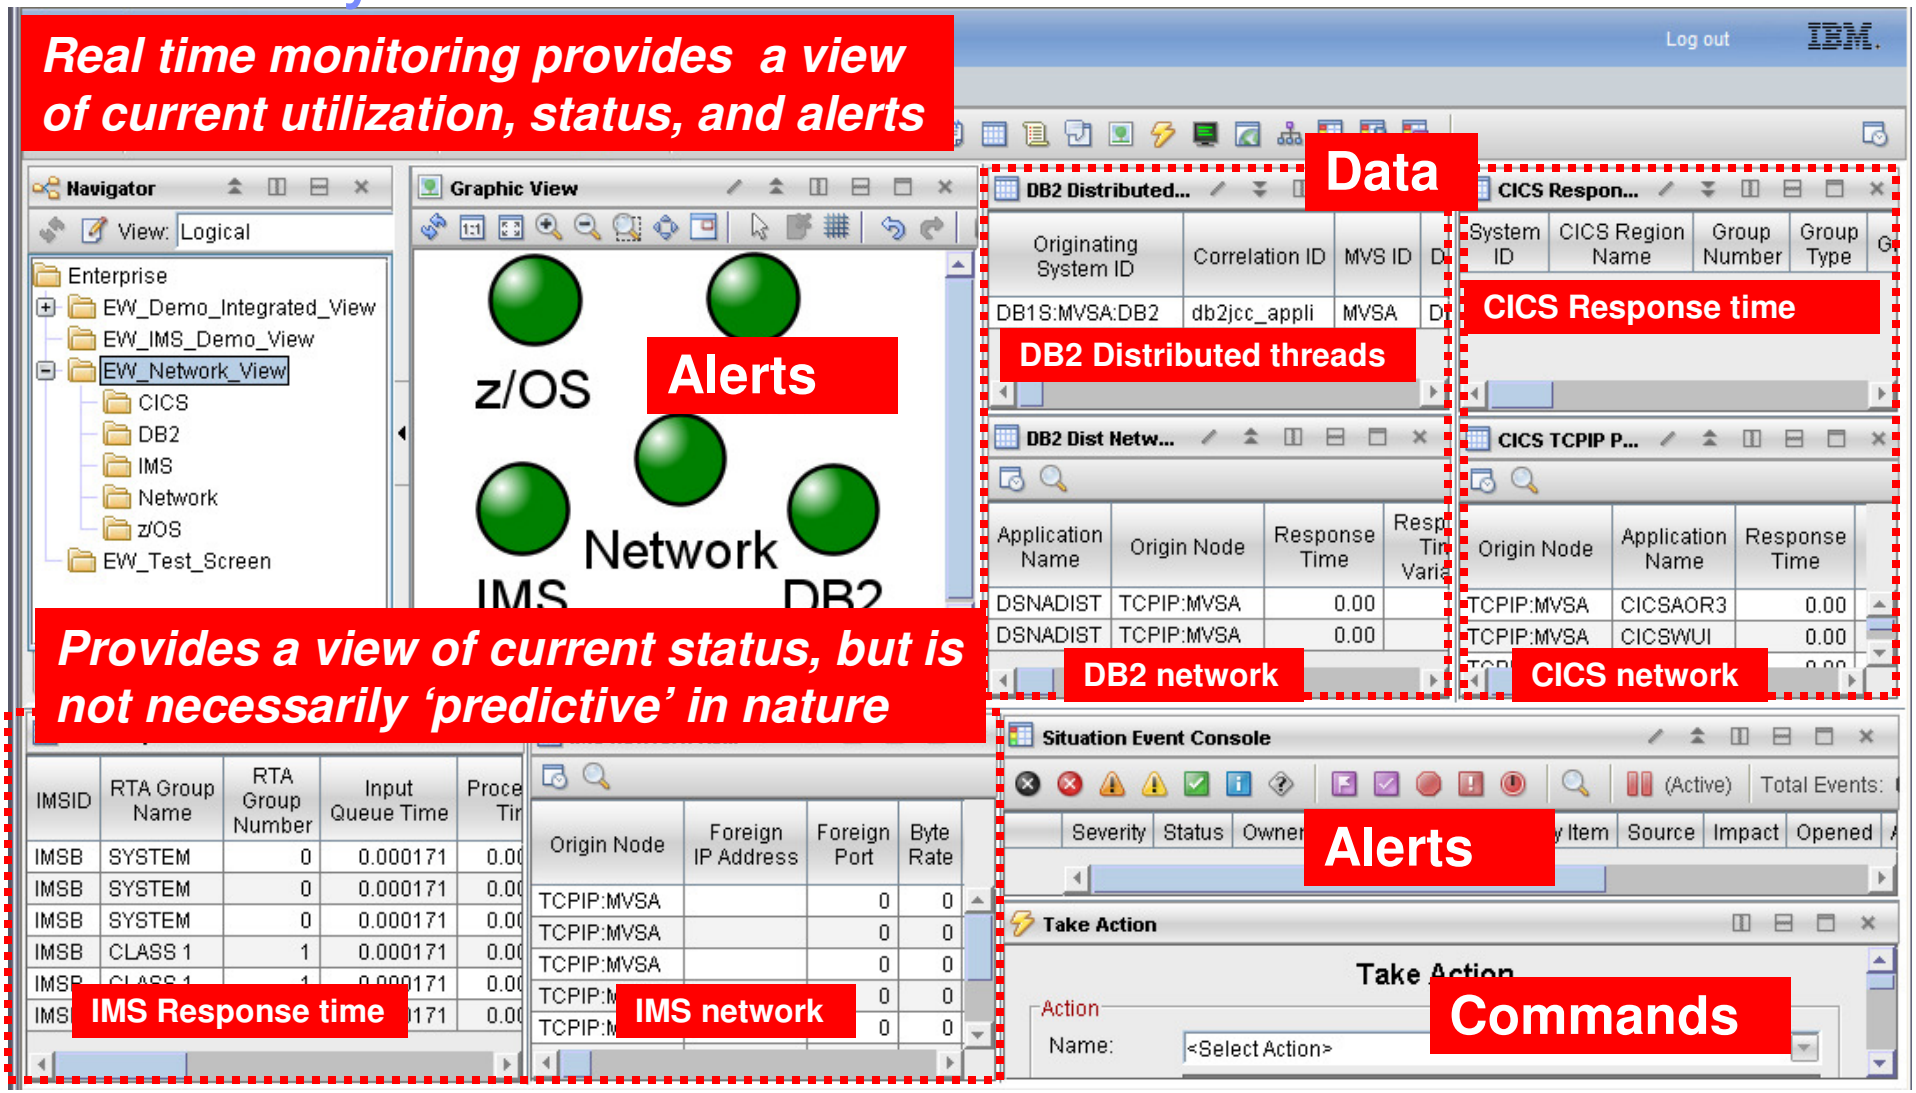

Historical Data Analysis

Helps Identify Critical Metrics, Trends, Usage Patterns And Potential Issues

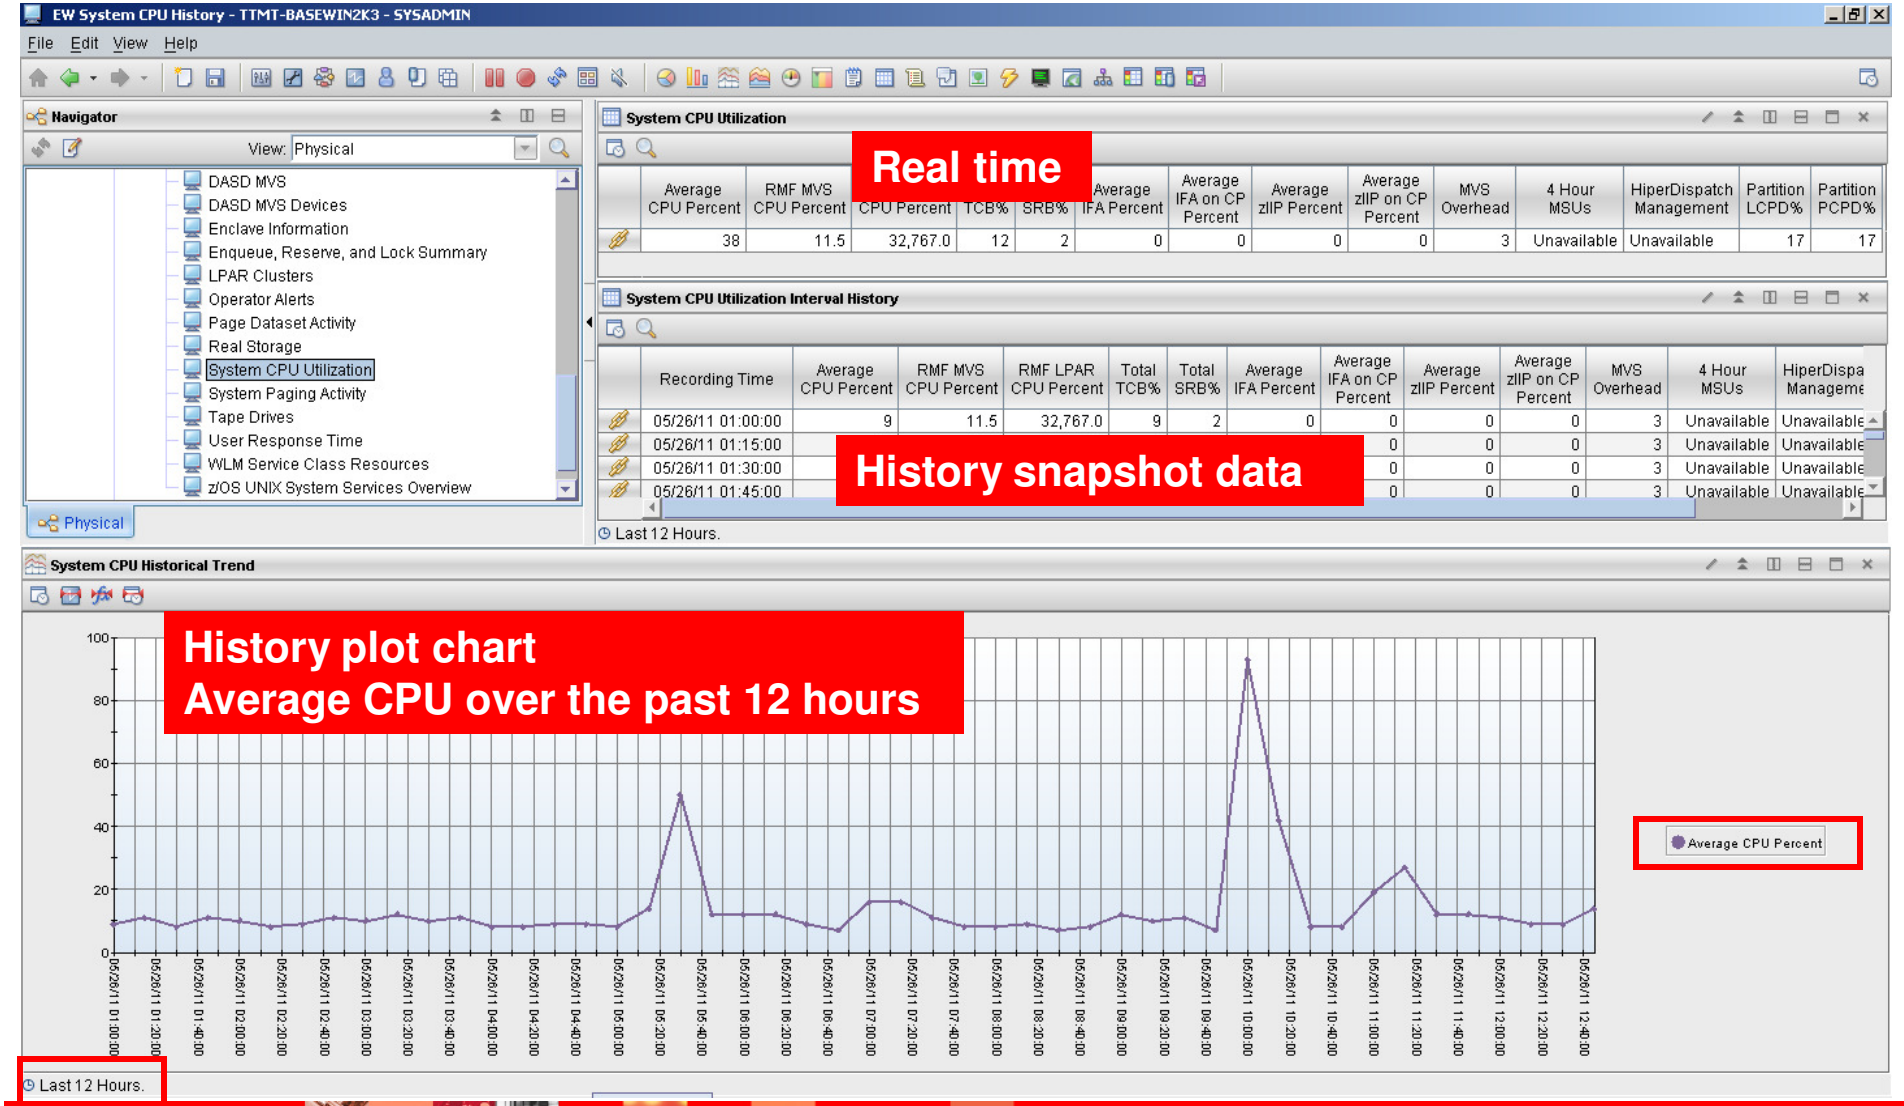

Another Example - Historical Baseline Data To Compare Past Trends To Current Trends

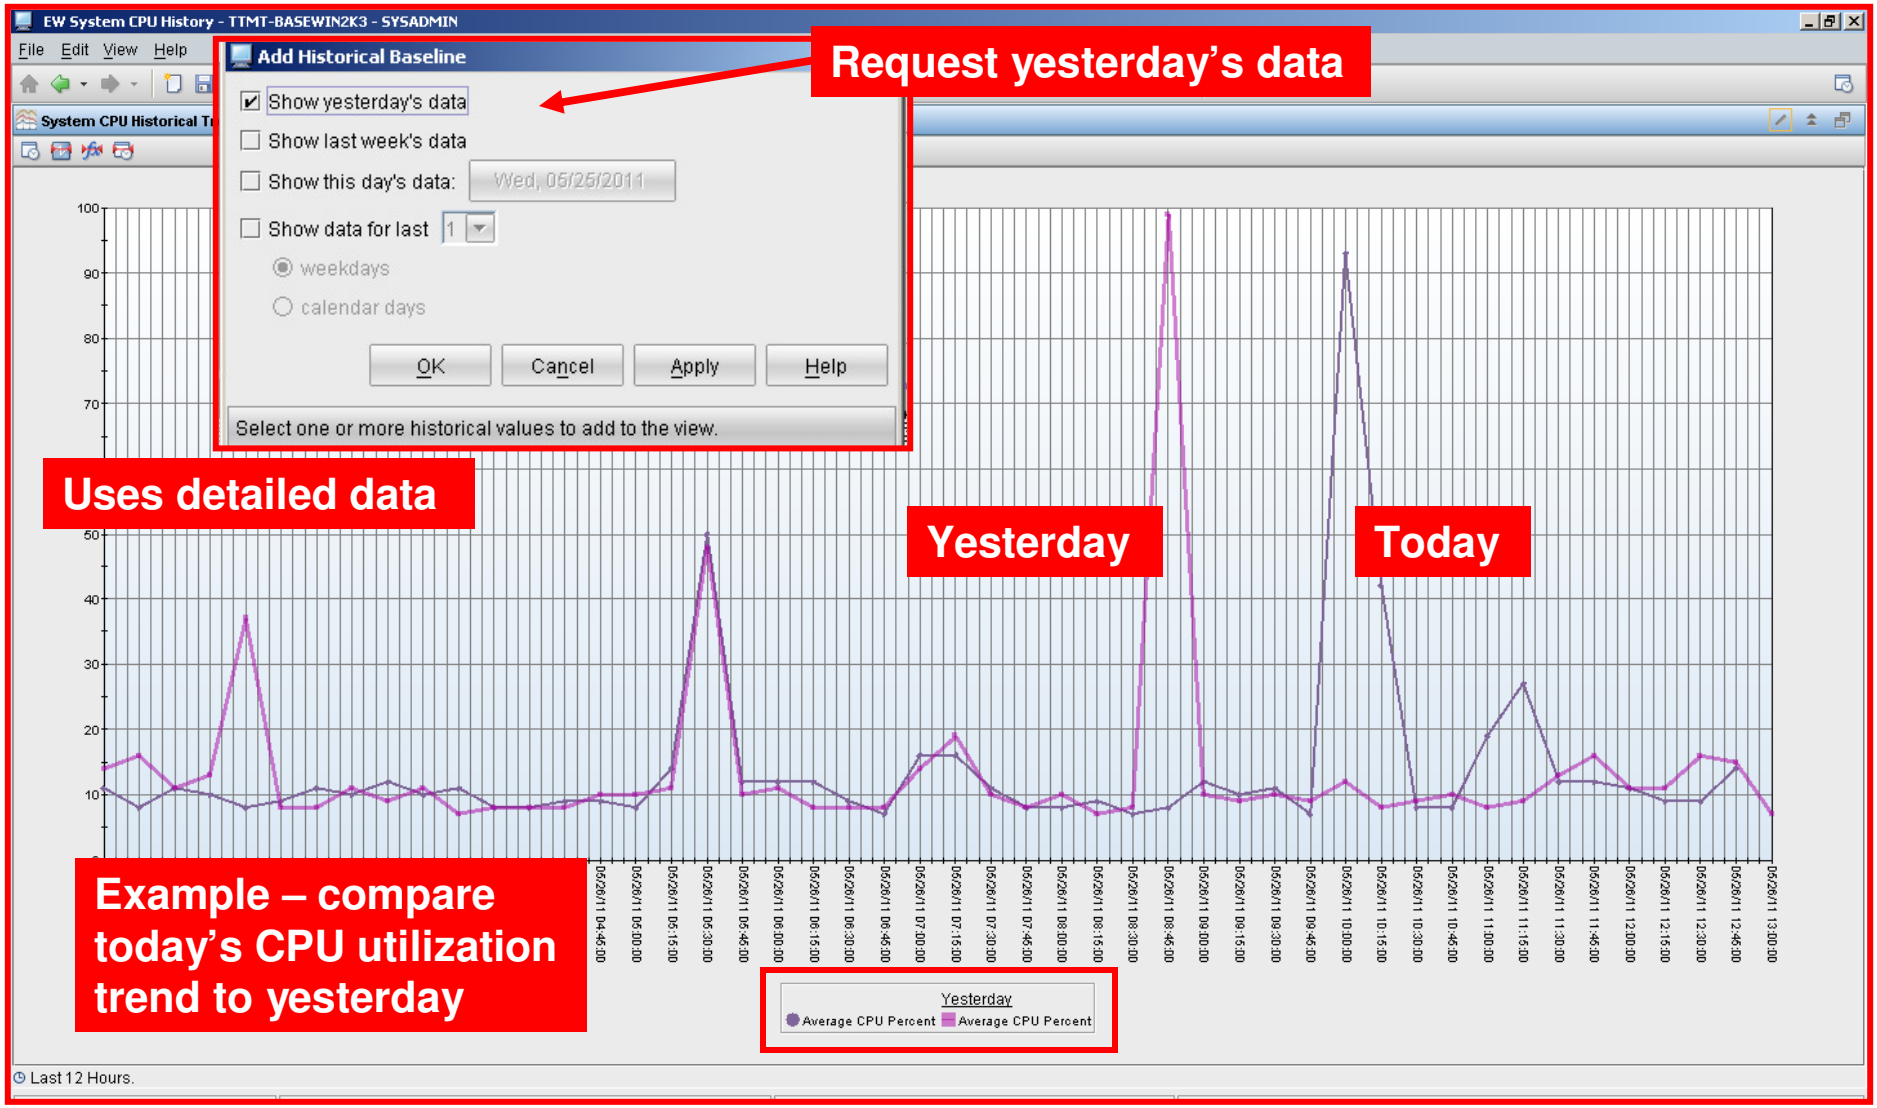

The Problem: Traditional Monitoring Approaches Have Limitations

- 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

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

Why Multivariate Analysis?

- Multivariate analysis expands the relevance of the predictive analytic approach
  - 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

An Example Of Multivariate Analytics

- Goal -> Predict/Identify Issues (Early Warning)
- Analyze performance data and combine Univariate with Multivariate Analytics
- Identify Metric Inter-relationships
- Detect anomalies rapidly, as metrics deviate from normal behaviour individually and from the correlated group

Predict: Eruption forecasting using seismic energy..

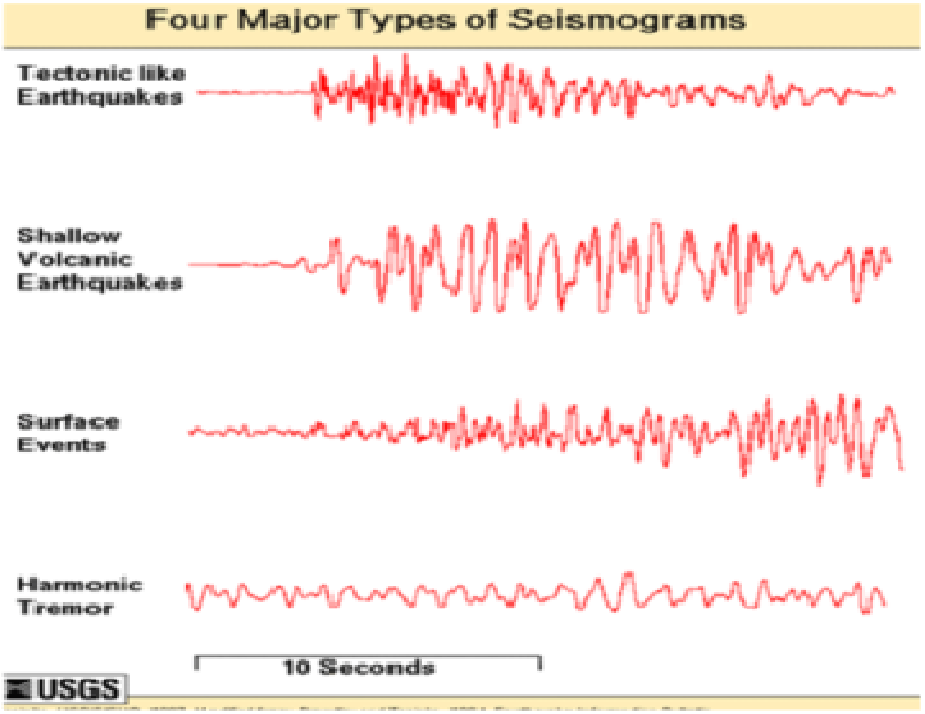

Using Predictive Analytics To Expand The Typical Analysis Paradigm

- Univariate metrics may be useful for some forms of analysis
- Many typical IT challenges are multivariate in nature
  - Most applications are multi-component and multi-layered
  - Most applications cross software platforms and may involve multiple ‘hops’
  - Complete analysis may require metrics from multiple sources
- Exploit the knowledge and expertise of the IT staff to begin to form a multivariate approach
  - Many may use predictive analytics in an informal manner
  - Proven knowledge of business applications (‘school of hard knocks’)
  - Knowledge as documented in established industry ‘best practices’

How Does Multivariate Analytics Detect Problems Sooner?

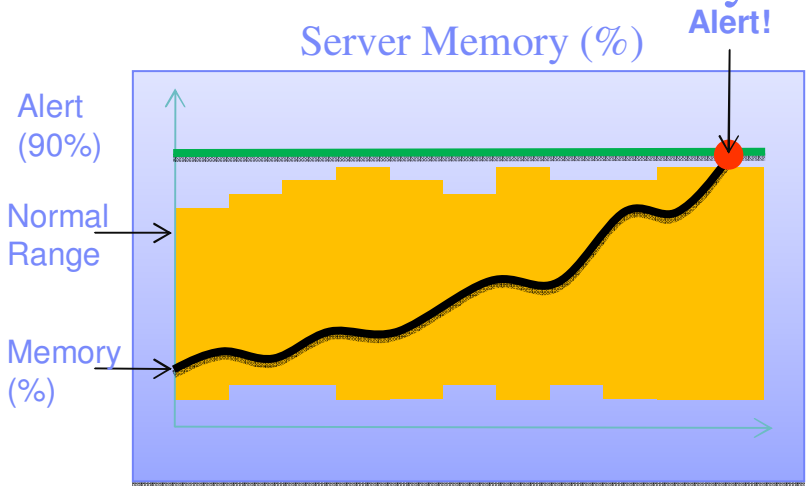

Static Threshold = Short Warning

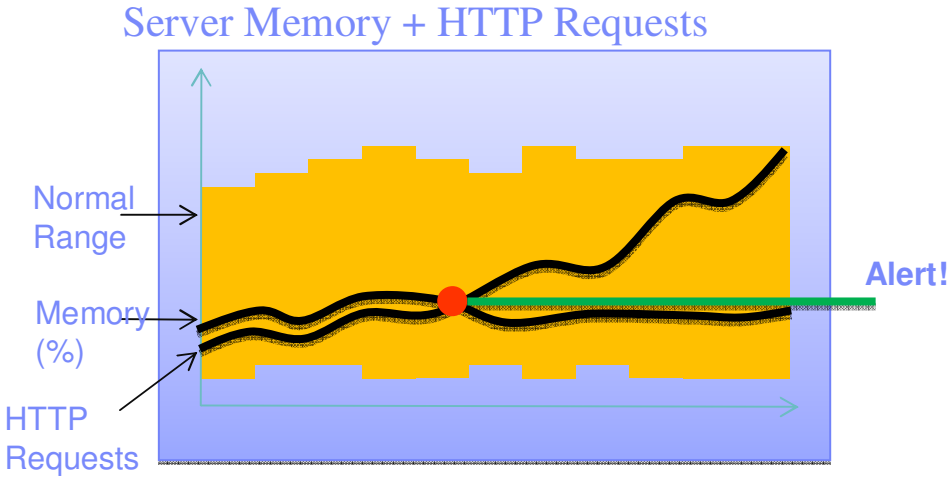

Multivariate = Alerts earlier on Deviation

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

An Example Of Multivariate Analysis For IMS Performance

Monitor and trend multiple IMS performance metrics over time

Plot chart analysis of key IMS performance metrics

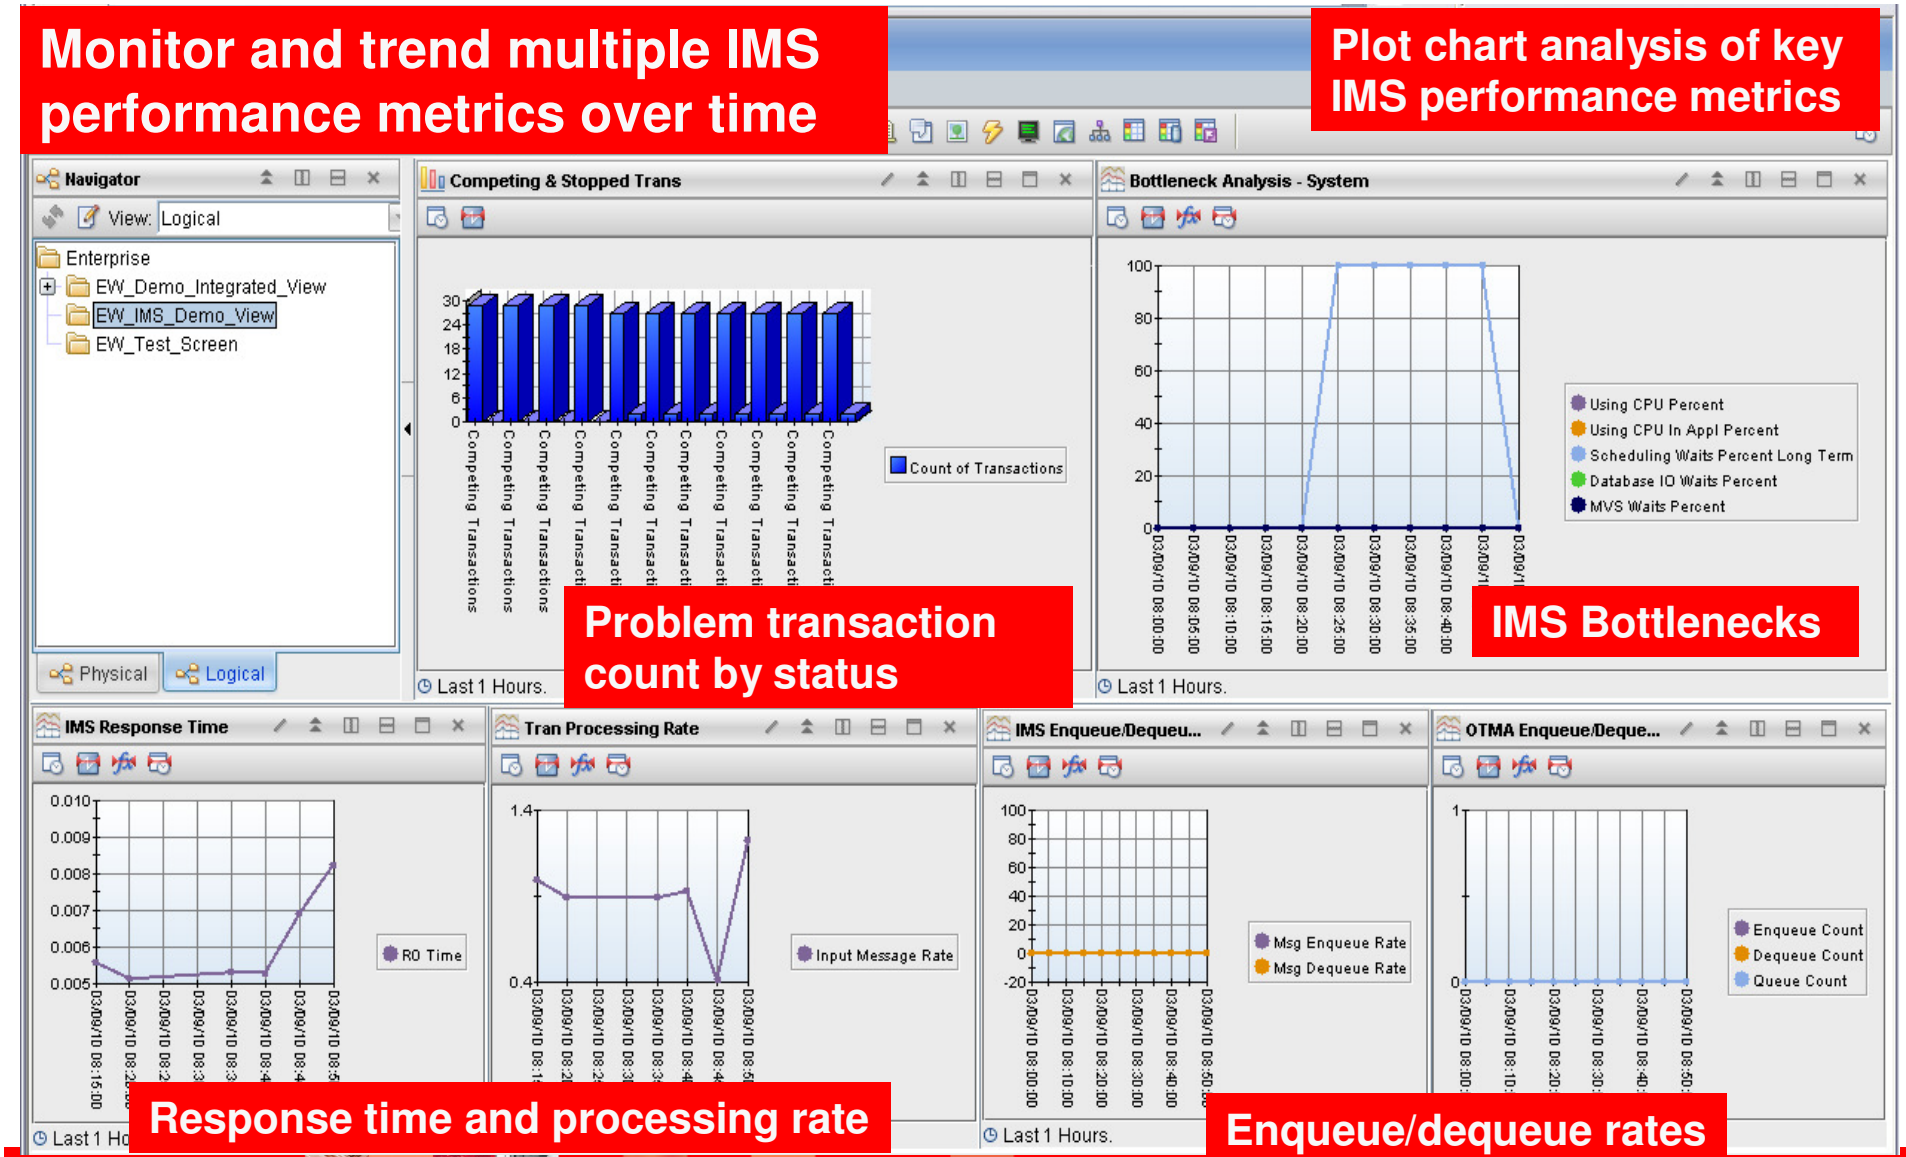

Problem transaction count by status

IMS Bottlenecks

Response time and processing rate

Enqueue/dequeue rates

Identifying The Critical Metrics

- Knowledge of business applications
  - Internal operational processes
  - Known issues based upon prior operational experience
  - Maintaining a history of common alerts/events
- Identify critical performance metrics as established by ‘best practices’ documented in commonly available sources
  - IBM documentation and IBM Red Books
  - Share, CMG, IDUG, Pulse, IOD and other user group presentations
- Define a list of the most critical metrics to track
  - Consider each component/platform for the application(s)
  - Consider various data sources
    - Monitoring, automation, console logs, application data sources

Example – Critical IMS Performance Metrics

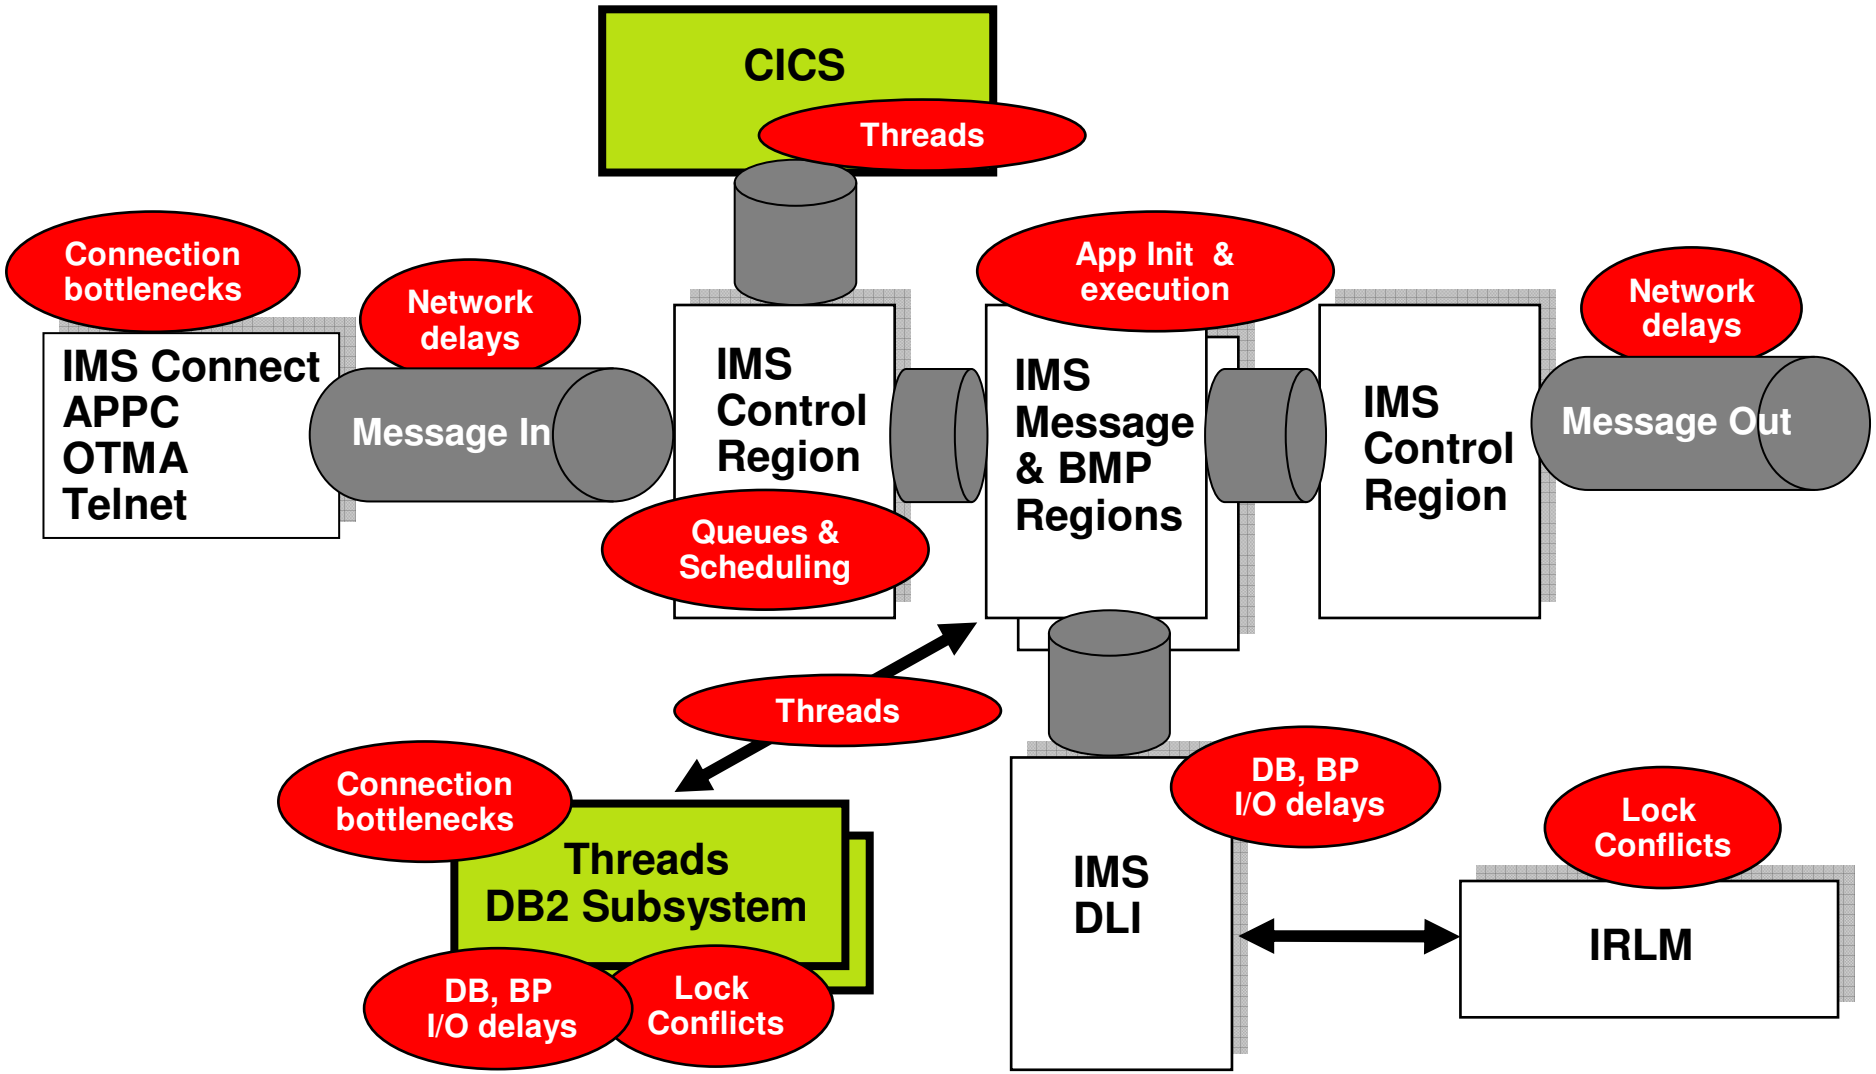

Another Example – Typical DB2 Performance Metrics

DB2 Performance Metrics
- DB2 thread connection issues
- Long running threads
- DB2 SQL delays and long running SQL
- DB2 database I/O delays and BP perf.
- Network and/or OS delays
- DB2 lock conflicts

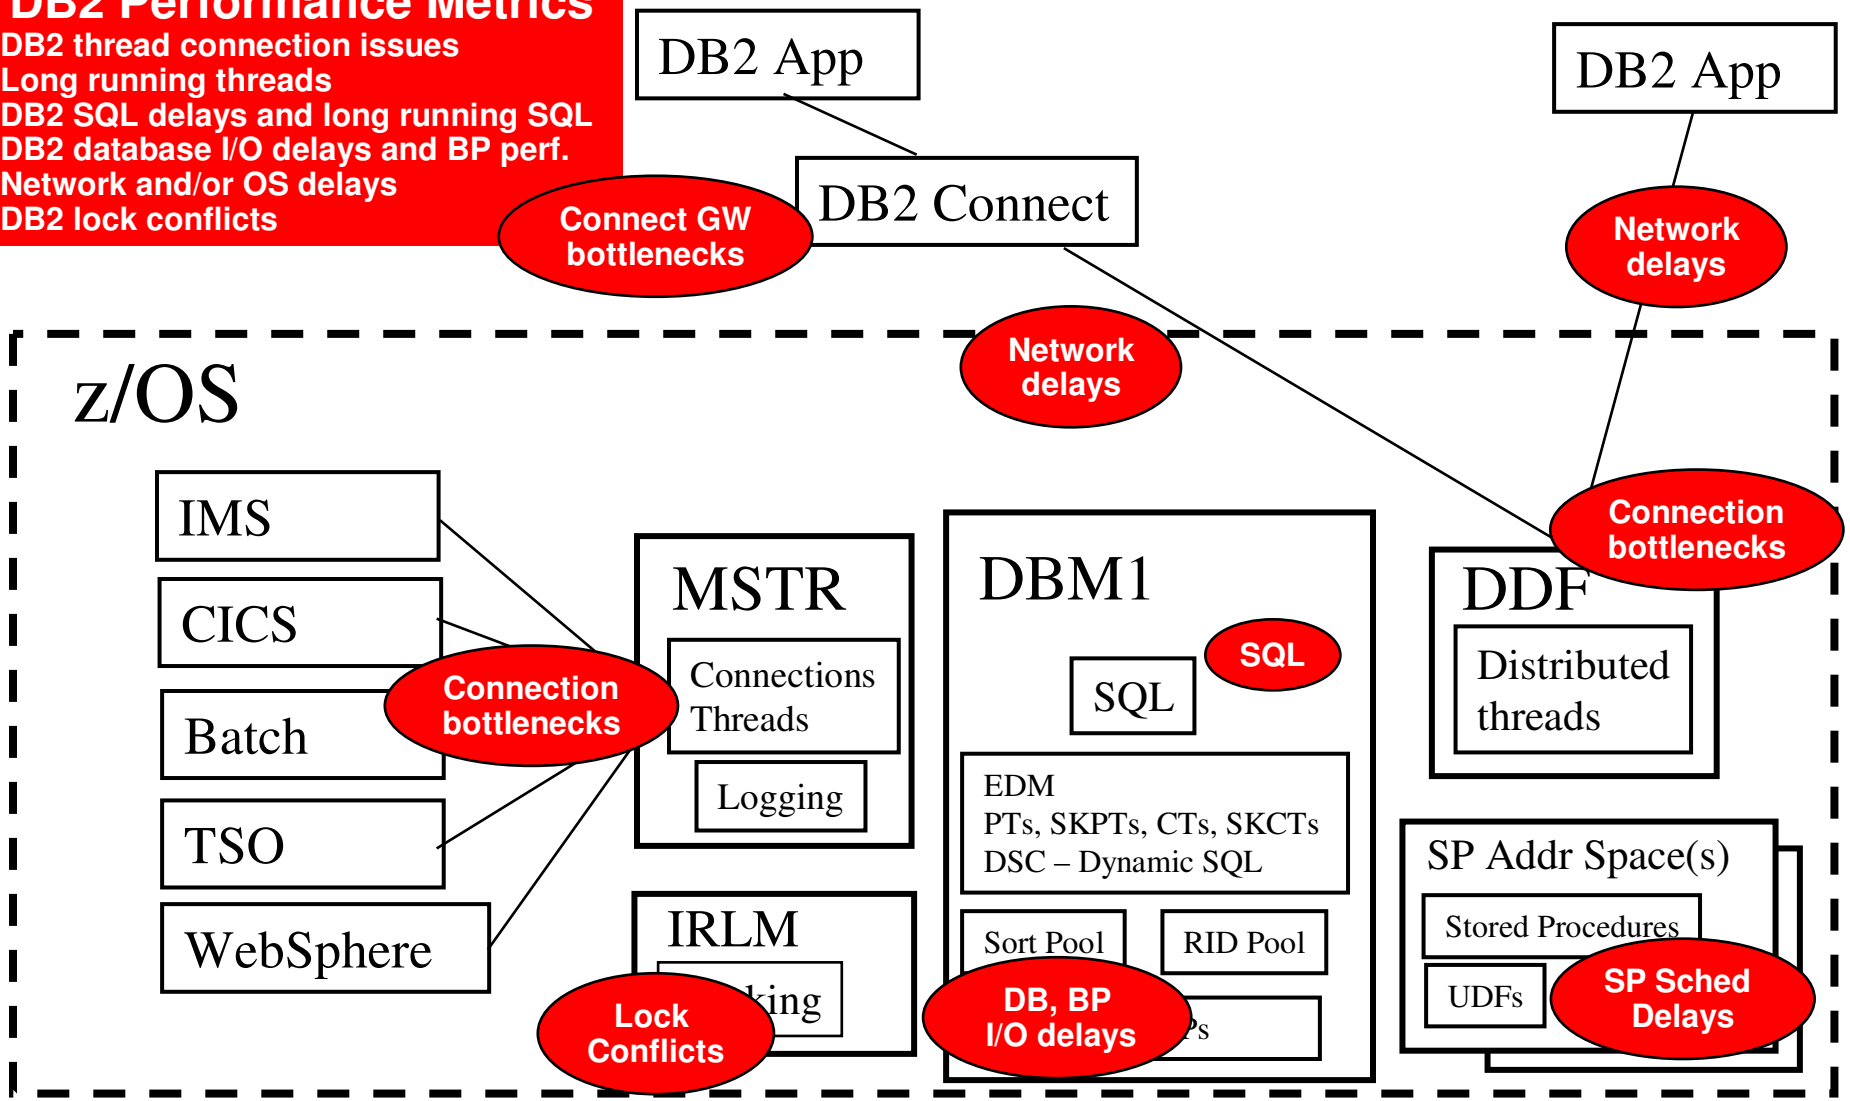

Other Examples Of Common z/OS Critical Performance Metrics

WebSphere MQ
Queue depth
Message send/receive rate
DLQ depth
Channel status and performance

CICS
Transaction response time
Transaction rate
Region CPU rate
File I/O count
String waits
Abend messages

z/OS
System CPU rate
Paging rate
WLM Performance Index
DASD I/O MSR time and rate
Critical console messages

WebSphere
Method call count and elapsed time
Heap size
Garbage collection
Connection pool utilization

Network
Network Connection status and performance
Network interface utilization

Don't Overlook Alerts
Alerts Can Provide Valuable Metrics

Alerts may be a useful source of metrics for analysis
Number of alerts and frequency of alerts may be useful
Correlate alerts to identify critical metrics

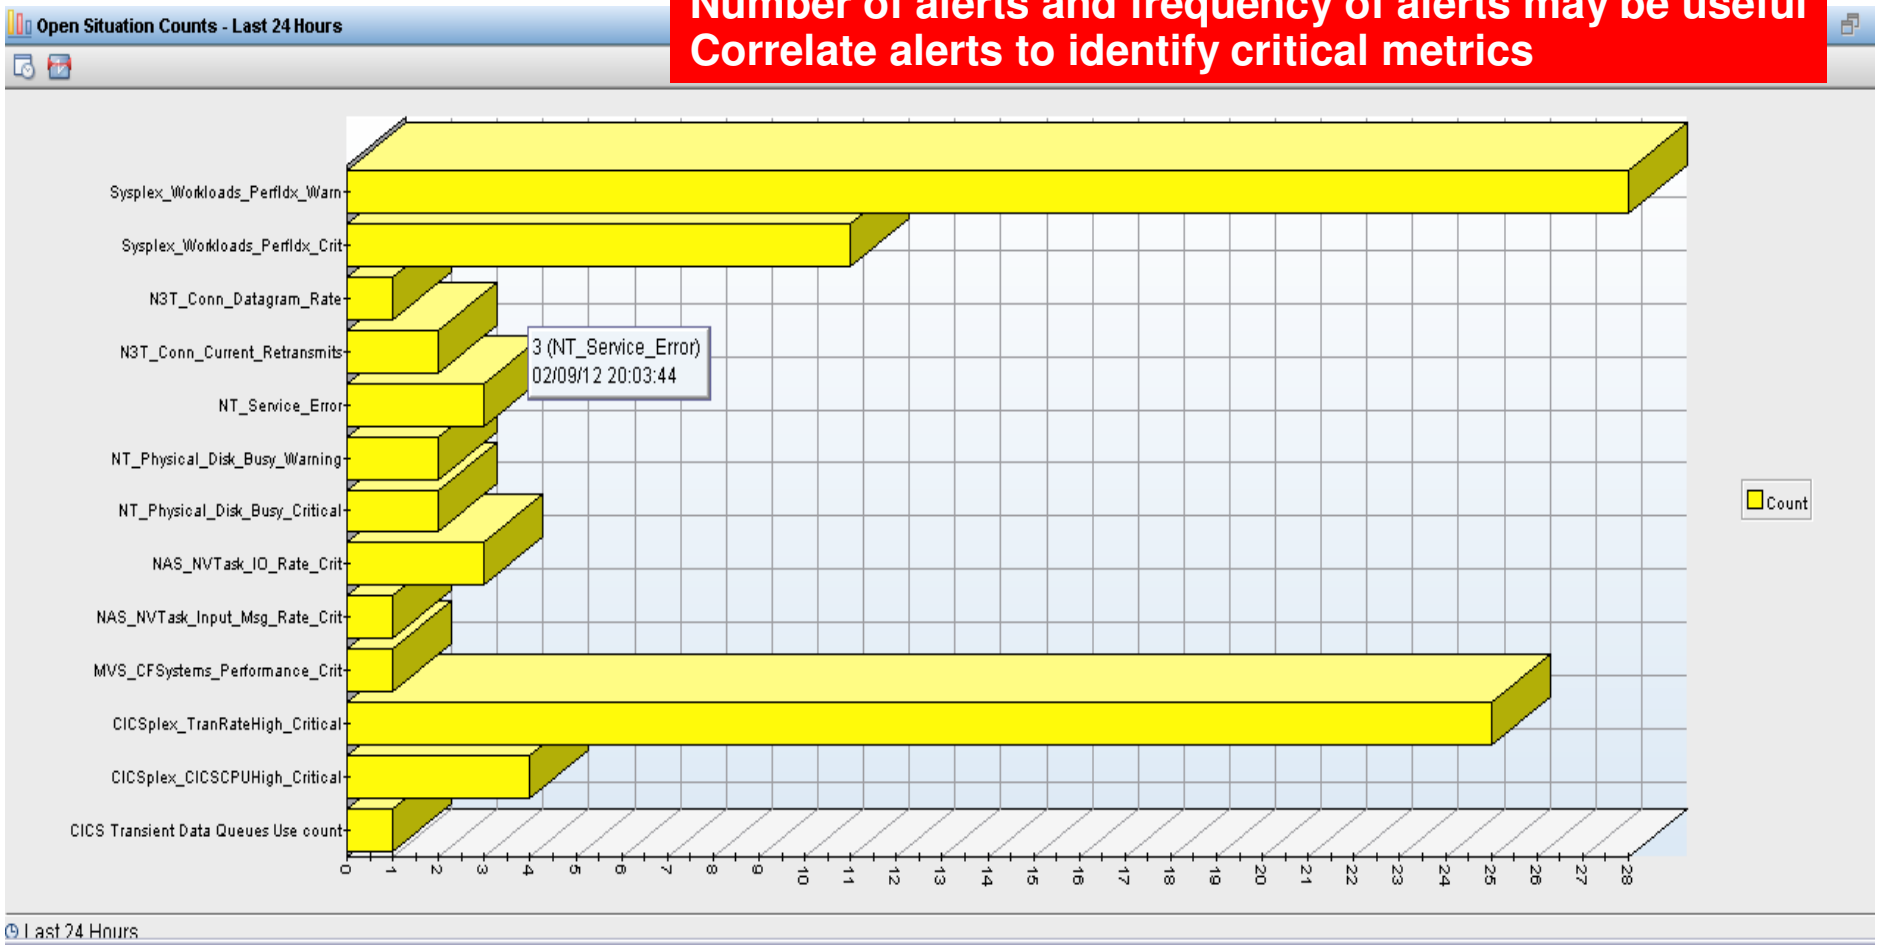

Predictive Analytics
Defining A Strategy – What's Required

- A predictive analytics approach starts with a comprehensive historical data collection strategy
  - Gather history to a common point, if technically feasible
- Consider historical collection challenges carefully
  - Different source platforms (example - LUW versus z/OS or z/VM)
  - Different instrumentation, trace options, data formats
    - z/OS SMF, DB2 logs, DB2 trace IFCIDs, CICS SMF records, IMS logs, WebSphere PMI trace interface, Windows performance counters, and much more…….
  - Collection frequency and quantity
    - Avoid flooding with useless information
    - Focus on ***quality versus quantity***
  - What metrics to collect
  - What quantity of data
  - Where to house the data
  - What level of detail and/or summarization

Predictive Analytics Begins With History
Historical Data Collection Considerations

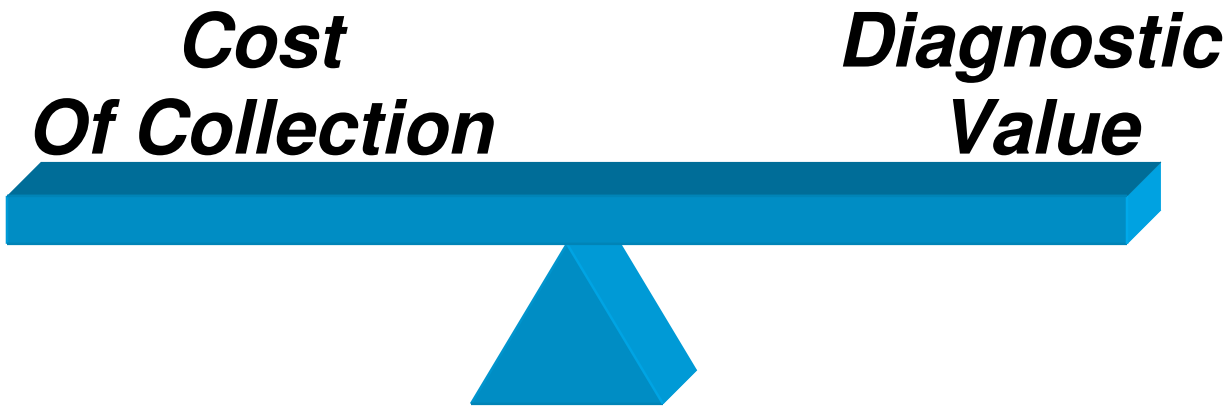

Monitoring And History Collection Trade-off

- Historical data collection varies in cost and quantity
  - CPU, memory, and software process cost of collection
  - Cost of data storage and retention
  - Cost of retrieval and post processing
  - Ease of review and analysis
- Some historical data will be more relevant and useful than other data
  - Consider the context, nature, and meaningfulness of the data

Types Of Historical Monitoring Data

- Know the nature and characteristics of the history data being collected

- Detail data
  - Data that documents/measures detail of a specific event
  - Often high quantity data and the most detailed for analysis
  - May pose the greatest challenge in terms of cost, retention, post processing
  - Examples – DB2 Accounting records, CICS SMF 110 records, IMS log records

- Summary data
  - Data that summarizes underlying detail data
  - Either an aggregation or an averaging of underlying detail records
  - May be useful for longer term trending and analysis
  - Reduces quantity of data and reduces cost of retention, post processing
  - Less detail may mean less diagnostic value

Types Of Historical Monitoring Data - continued

- Interval data
  - History data that includes an encapsulation of one or multiple events to a specified time interval
  - The data will include all activity within that given time interval
  - Useful for problem analysis and trending analysis
  - Examples – DB2 statistics records

- Snapshot data
  - Typically a point in time snapshot of activity
  - Snapshots are usually based on a specified time interval
  - Snapshots may be taken of types of history (detail, summary, or interval)
  - Snapshots will show activity at time of the snapshot, but may/may not reflect activity between snapshots
  - Useful for problem analysis and trending analysis
  - Useful as an aid in setting alert thresholds
  - Examples –snapshot history captured by performance monitoring,

Predictive Analytics
Defining A Strategy – What's Required - continued

- Pursue a multivariate approach where feasible
  - Univariate is useful for certain trending and modeling uses
  - Multiple variables provide a more meaningful indication of potential issues
  - Multivariate is more relevant to today's composite applications
- Multivariate poses challenges
  - How to best identify the key metrics to collect
  - How much data to collect
  - How to analyze and correlate the information
  - How to display the data correlation result
  - How to feed the result to other systems
    - Monitoring, automation, business application views, help desk views

Predictive Analytics – Summary
What's Needed

- Identify information metric sources
  - Monitoring tools, platforms, alerts, console and application logs
- Define metrics
  - What are the most critical metrics to track?
- History
  - Define a collection strategy that allows for the aggregation of data
- Correlation methodology
  - How to correlate metrics in real time and in history
  - Is there a way to automate the correlation process?
- Display and analysis methodology
  - How to analyze and display critical metrics – data and alerts
- Prediction
  - How predictive are the chosen metrics?

Thank You!!

Check Out My Blog

http://tivoliwithaz.blogspot.com