Pulse

IBM SolutionsConnect 2013

Using Predictive Analytics to Prevent Business Disruption

Erin Burke

Program Director

Tivoli Product Management

IBM

Please note

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Key trends are fueling the need and urgency for analytics

1 The emergence of big data analytics

65%

65% of business are not using big data for business advantage

2 Increasing consumer expectations

84%

84% of consumers rely on social networks for purchase decisions

3 Accelerating pressure to do more with less

32%

Organizations using advanced analytics enjoy 32% higher return on invested capital

Managing Exploding Operations Data is a Huge Challenge

- **Too Little:** Limit Data Acquisition and risk missing important data
- **Too Much:** Flood IT Operations and risk missing important data
- **Just Right?** Traditionally IT Management evolved Tools, Best Practices, Process, and Experience to filter data for relatively static systems
- **Just Right: Automated Analytics** to examine all data, learn what is important, and escalate critical problems to Operations staff in a timely way

Why IBM for Analytics?

- IBM has invested $16B in analytics, more than any other company
- We have the most comprehensive portfolio from business to IT analytics to solve for your overall pain points, while most other vendors offer only point solutions
- Tivoli's suite of analytics products leverage best of breed capabilities from across IBM's portfolio, applied to the IT domain

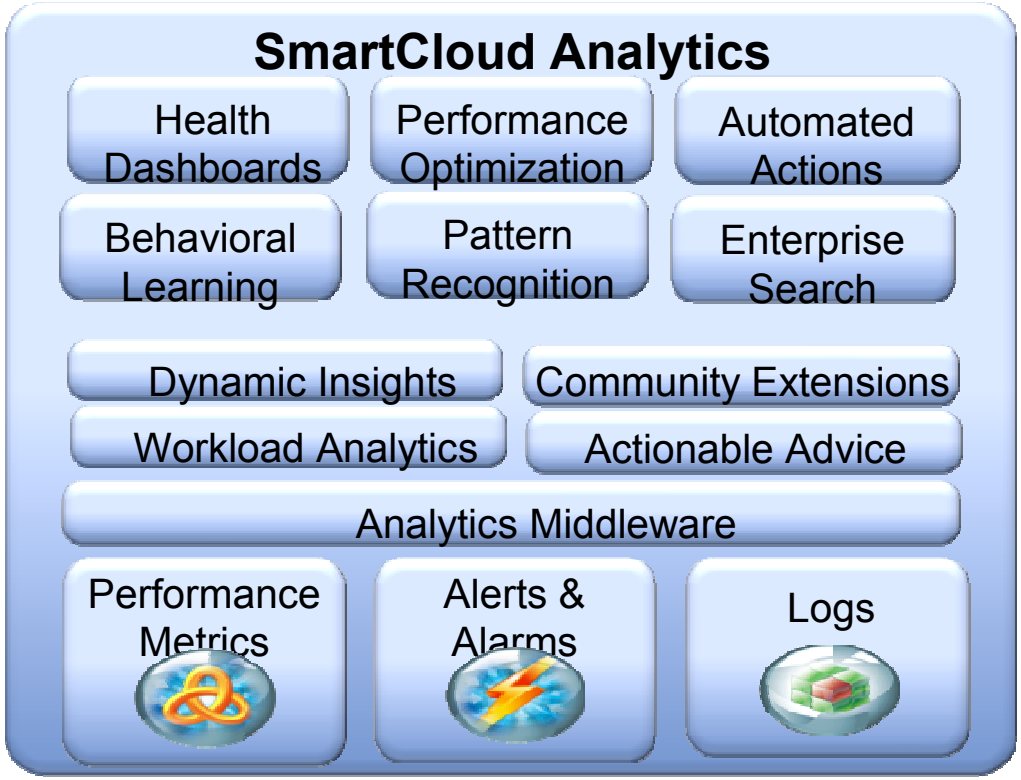

Business Value to Analytics Adoption

Optimized Performance	Predictive Outage Avoidance	Faster Problem Resolution	Improved Insight
Track, Optimize, and Predict capacity and performance needs over time	*Ensure availability of applications and services*	*Find & correct problems faster with tools that determine actions required to resolve issues*	*Enhance visibility into systems resource relationships while increasing customer satisfaction*
Perform • Track capacity and performance of applications and services in classic and cloud environments •Optimize resource deployment with what-if and best fit planning tools • Increase utilization of existing assets	**Predict** • Use learning tools to augment custom best practices • Leverage statistical methods to maximize predictive warning • Use past maintenance to predict part failures	**Resolve** • **Identify** problems quicker with insight to large unstructured repositories • **Isolate** problems quicker by bringing relevant unstructured data into problem investigations • **Repair** problems quicker with the right details quickly to hand.	**Know** • Determine what resources are interdependent to assess impact of failures • Gain insight into what is important to your customer • Decrease customer churn and acquisition costs while increasing customer retention and

Lower IT Administration Costs with Automated Analytics

- Escalate performance and capacity issues automatically, reducing manual analysis efforts
- Reduce manual customization using learning tools that automatically adjust to new normals
- Detect and present problems with a proposed resolution, to be able to do more with less

Proactive, Predictive and Preventative Management

- Few companies are genuinely proactive or preventative
- Most organization react to service outages in progress
- Diagnosis can be complicated by organizational silos, disparate tools, complexity and the sheer volume of data.
- Outages and degradation can cost millions of dollars, impact brand, customer churn & retention

Why aren't operations teams preventative today?

- Too much data to analyze manually
- Existing analytic techniques, such as standard thresholds, are not up to the task
- They cannot detect problems while they are emerging (before business impact)
- Set threshold too high, insufficient warning before total failure.
- Set threshold too low, too much noise, everything is ignored

Example Scenario: Internet Banking Application

Goal: Automatically learn normal mathematical relationships between metrics

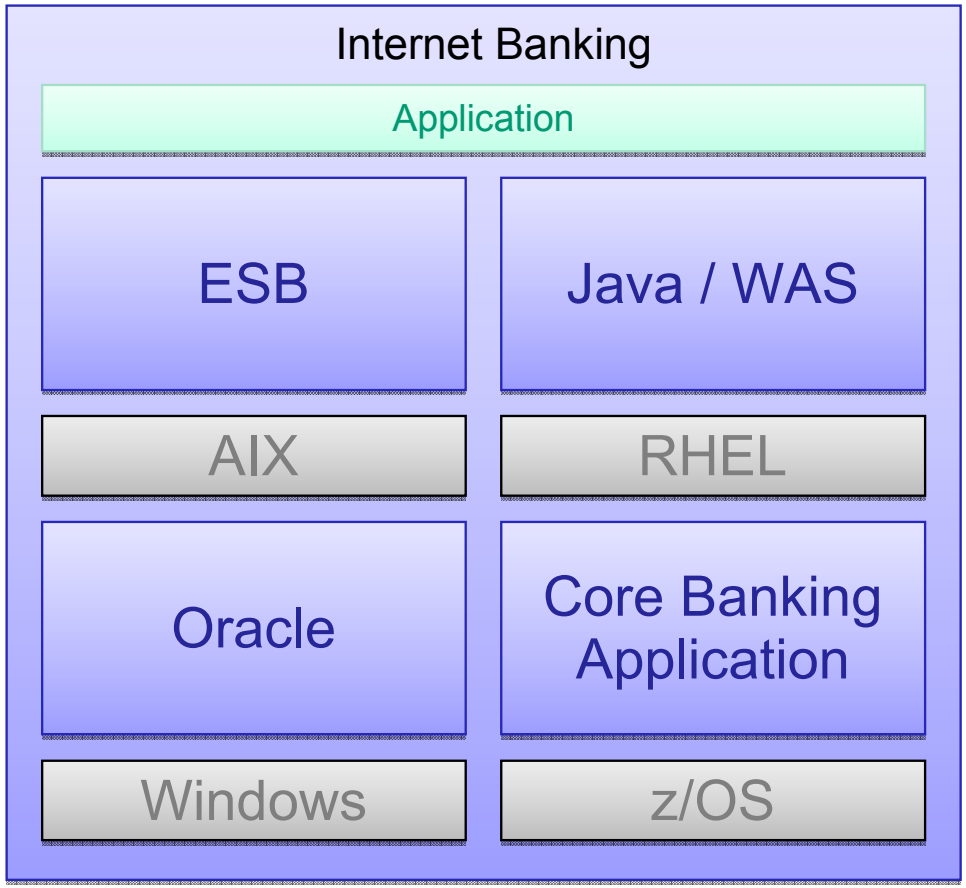

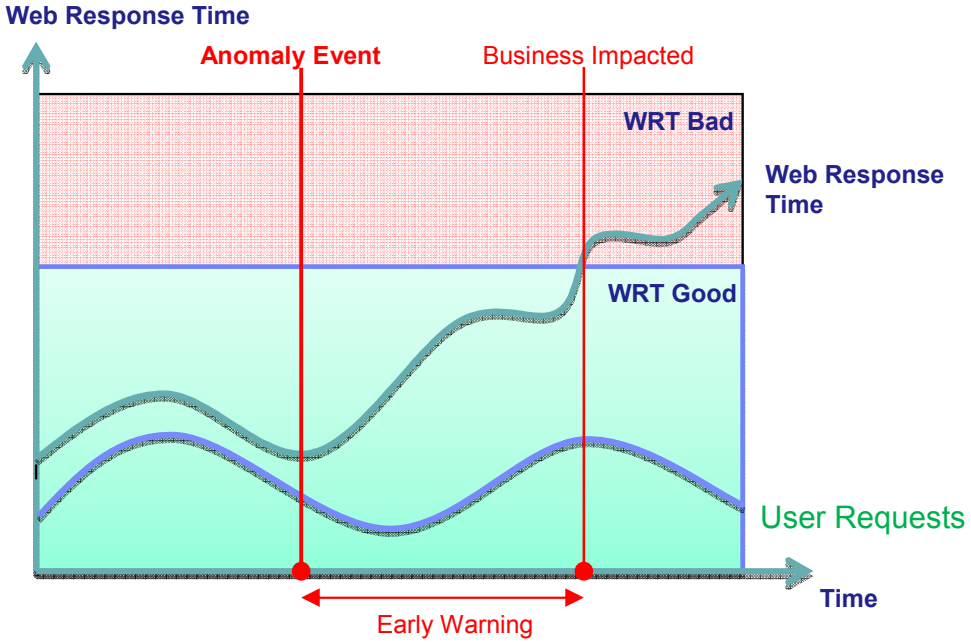

- Learns 'Web Response Time' has a normal causal relationship with 'User Requests' - WRT gets slower as user load gets higher.
- If this healthy historical relationship breaks down, say due to a memory leak, an anomaly is raised immediately
- The problem is detected even while WRT service is "good"

Emerging problems can be detected even while service levels are good in absolute terms

Correlation of Multiple Metrics

Statistical models can discover mathematical relationships between metrics

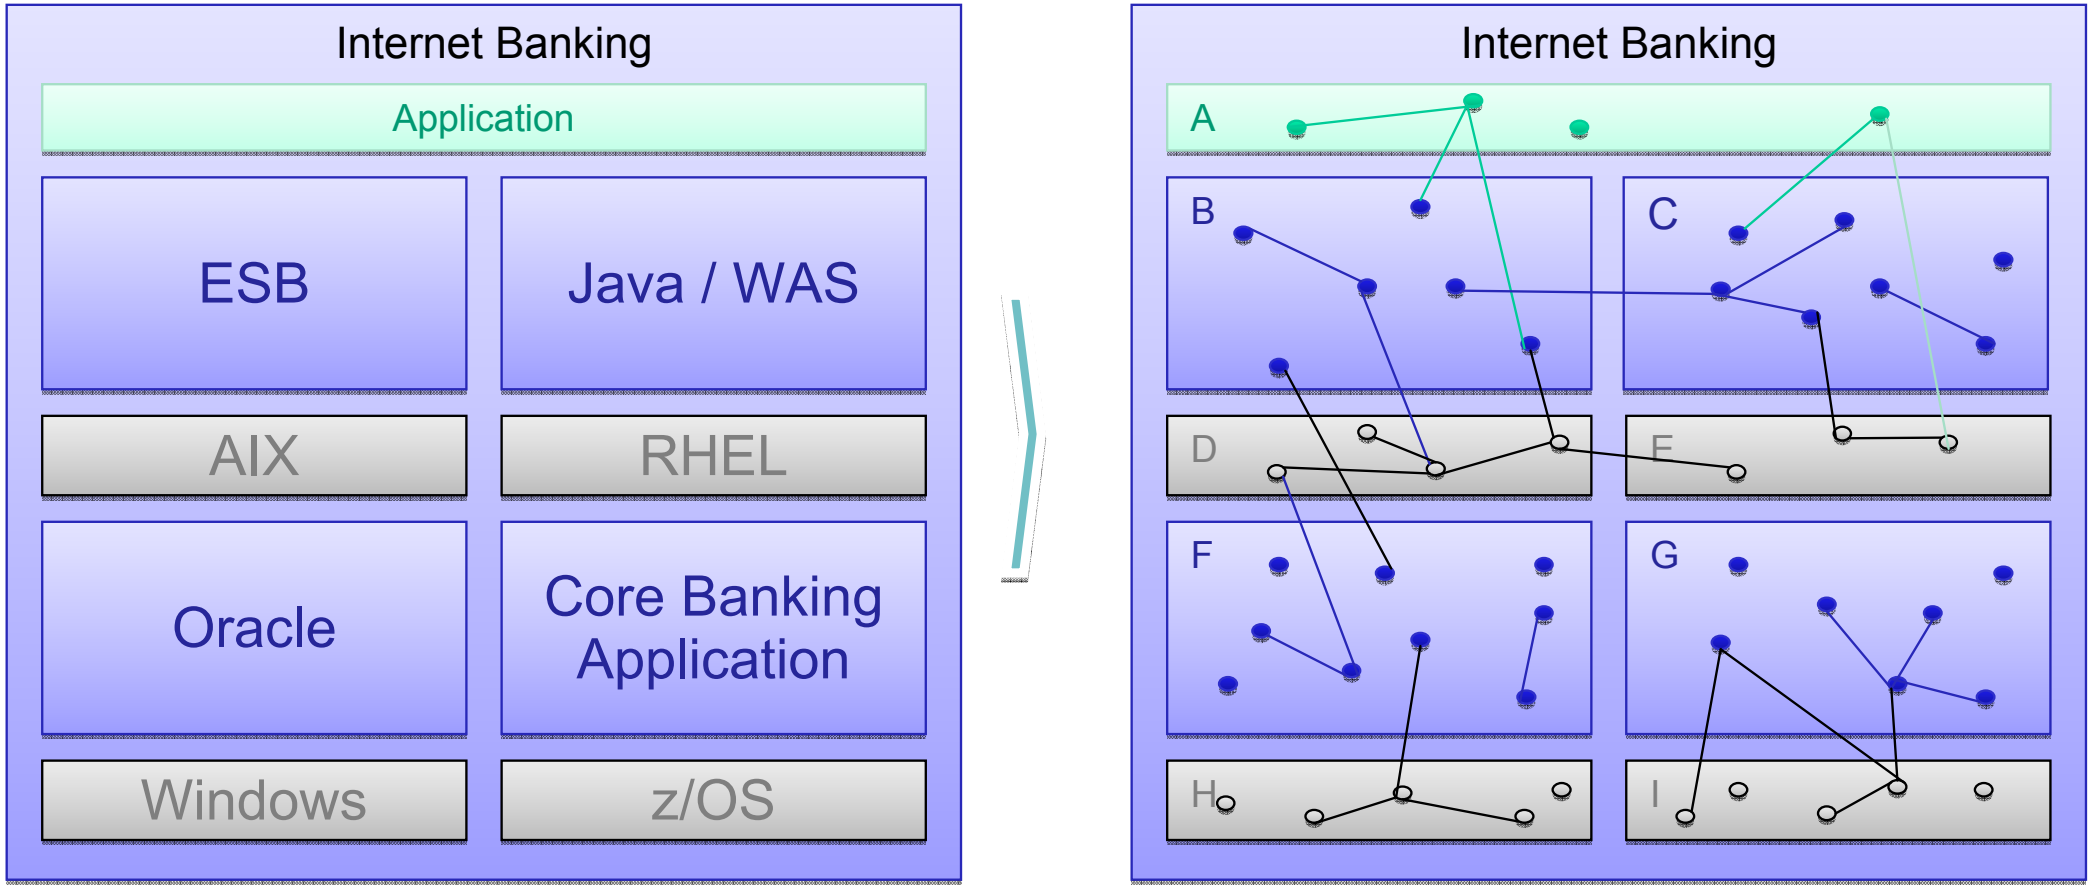

The extent this can be achieved depends on a number of factors, such as: range and type of data, availability of data, and stability of environment. Analytics falls back to a single metric if metrics are unrelated.

Multiple Metrics Analysis - Value of this approach

- Learns normal operational behaviour across the infrastructure, including how metrics behave together.
- Maximize Advance Warning: Identifies metric relationship changes that signal a problem long before traditional thresholds
- Identifies problems before you know to look for them
- Detects service impacts that are not identifiable by fixed thresholds alone.
- Assists with root cause analysis by indicating the most offending metrics.
- Reduces expensive and time consuming false alerts.

Provides a more intelligent real-time assessment of data, able to detect problems as they are emerging

Available Now: Tivoli Analytics for Service Performance

Proactive and self-learning Performance and BSM intelligence from Tivoli

- Real-time analytics for detecting and avoiding service disruption
- Uses advanced Watson research algorithms
- Correlates metrics across multiple domains and heterogeneous data

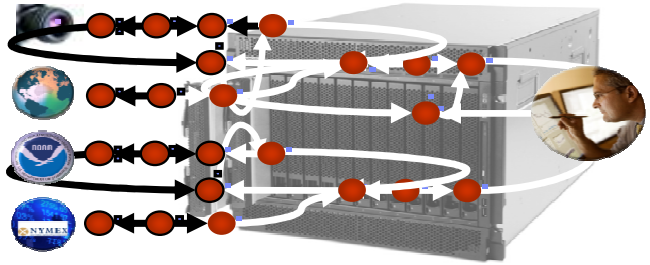

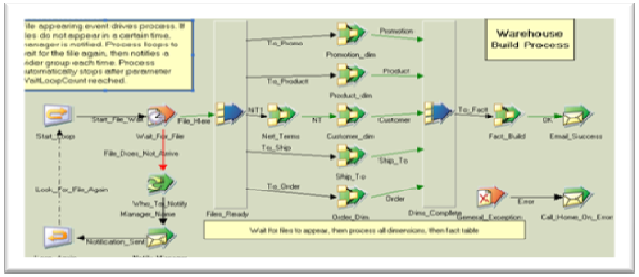

- Leverages IBM Big Data technology
- Embeds InfoSphere Streams, IBM's unique streaming analytic engine
- Enables ultra-high scalability commodity server computing clusters and large algorithm sizes to maximize machine intelligence value
- Embeds InfoSphere Datastage, IBM's market leading mediation solution
- Quickly integrate to any monitoring source using a large library of out-of-the-box connectors
- Leverages your Tivoli and non-Tivoli environments

Solution Architecture - Mediation

TASP

- **User Interface & Management**
  Tivoli Integrated Portal
- **Post-Processing Rules**
  Uses OMNIbus Rule Engine
- **Anomaly Consolidation**
- **Analytic Application**
- **Analytic Engine**
  IBM InfoSphere Streams
- **Mediation**
  IBM InfoSphere Datastage

- Market leading mediation - provided as component.
- Proven rapid integration to new data sources.
- Productivity tooling & collaboration included
- High performance and scalability.
- Large framework of connectors.
- Fast integration to common monitoring data formats.

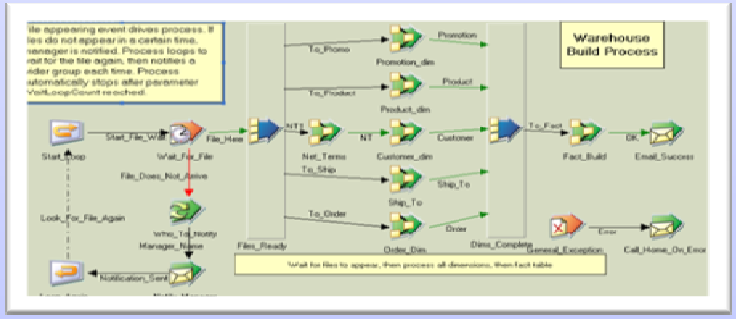

Mediation Rapid Common Extraction

TASP provides a quick setup ‘Common Extractor’ feature that allows fast extraction from the most common interface types such as:

- CSV
- Databases and database connectors, e.g. JDBC

Monitoring Suites	Interface	Implemented in trials
HP Sitescope	JDBC	Yes
Quest Foglight	Script dump to CSV	Yes
CA Wily Introscope	JDBC	Yes
IBM ITM TDW	DB2	Yes
IBM TDW Proxy Agent (low lat)	CSV	Yes
IBM ITCAM TDW	DB2	Yes
VAM	Script dump to CSV	Yes
HP Mercury BAC	JDBC	Yes
IBM Performance Manager	CSV	Yes
Brix	CSV	Yes
IBM Service Quality Manager	CSV	Yes

Other extractions can be quickly built from a large library of Datastage connectors

Mediation Connector Library – InfoSphere DataStage

RDBMS
DB2 (on Z, I, P or X series)
Oracle
Informix (IDS and XPS)
Ingres
Netezza
Progress
RDB
RedBrick
SQL/DS
SQL Server
Sybase (ASE & IQ)
Teradata
Universe
UniData
NonStop SQL
InfoSphere Federation Server
InfoSphere Classic Federation
And more…..

General Access
Sequential File
Complex Flat File
File Set
Data Set
Named Pipe
iWay
FTP
SFTP
Compressed / Encoded Data
External Command Call
Parallel/wrapped 3rd party apps

Enterprise Applications
JDE/PeopleSoft OneWorld
Oracle Applications
PeopleSoft
SAS
SAP BW
SAP R/3
Siebel
Ariba
Manugistics
I2

Standards & Real Time
WebSphere MQ
Java Messaging Services (JMS)
Java
XML & XSL-T
EBXML
Web Services (SOAP)
Enterprise Java Beans (EJB)
EDI

CDC
DB2 (on Z, I, P, X series)
Oracle
SQL Server
Sybase
Informix
IMS
VSAM
ADABAS
IDMS
Datacom

Legacy
Allbase/SQL
C-ISAM
D-ISAM
Datacom/DB
DS Mumps
Enscribe
Essbase
FOCUS
IDMS/SQL
ImageSQL
Infoman
KSAM
M204
MS Analysis
Nomad
Nucleus
RMS S2000
Supra
TOTAL
TurboImage
Unify
And many more….

Solution Architecture – Analytic Engine

TASP

- **User Interface & Management**
  Tivoli Integrated Portal
- **Post-Processing Rules**
  Uses OMNIbus Rule Engine
- **Anomaly Consolidation**
- **Analytic Application**
- **Analytic Engine**
  IBM InfoSphere Streams
- **Mediation**
  IBM InfoSphere Datastage

- Real-time streaming analytic engine, provided as a component
- High volume and low latency.
- Supports server clustering and redundancy (next rel)

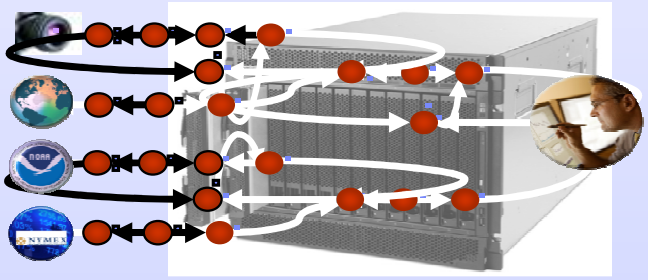

- Enables large algorithm capacity – 80,000 metrics in a single algorithm instance (a typical banking application produces ~30,000 - 60,000 metrics)
- Allows multiple algorithm instances spread across commodity server computing clusters, making maximum advantage of multi-core parallelism (next rel)

Solution Architecture – Analytics

TASP

- **User Interface & Management**
  Tivoli Integrated Portal
- **Post-Processing Rules**
  Uses OMNIbus Rule Engine
- **Anomaly Consolidation**
- **Analytic Application**
- **Analytic Engine**
  IBM InfoSphere Streams
- **Mediation**
  IBM InfoSphere Datastage

• Automated anomaly detection and prediction on time-series performance metrics

• Behavioural learning to model not only one metric at a time, but the relationships between them for anomaly detection...

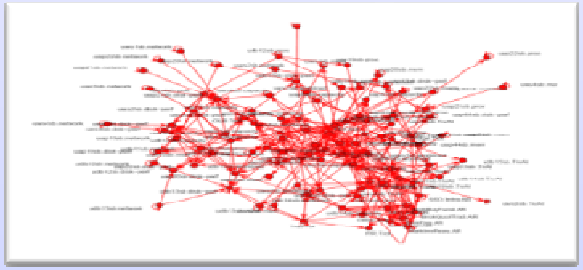

• Single metric evaluation replacing many manual thresholds for any time series data

• Multiple metric correlation enabling earlier detection than traditional thresholds with higher confidence

Solution Architecture – Anomaly Consolidation

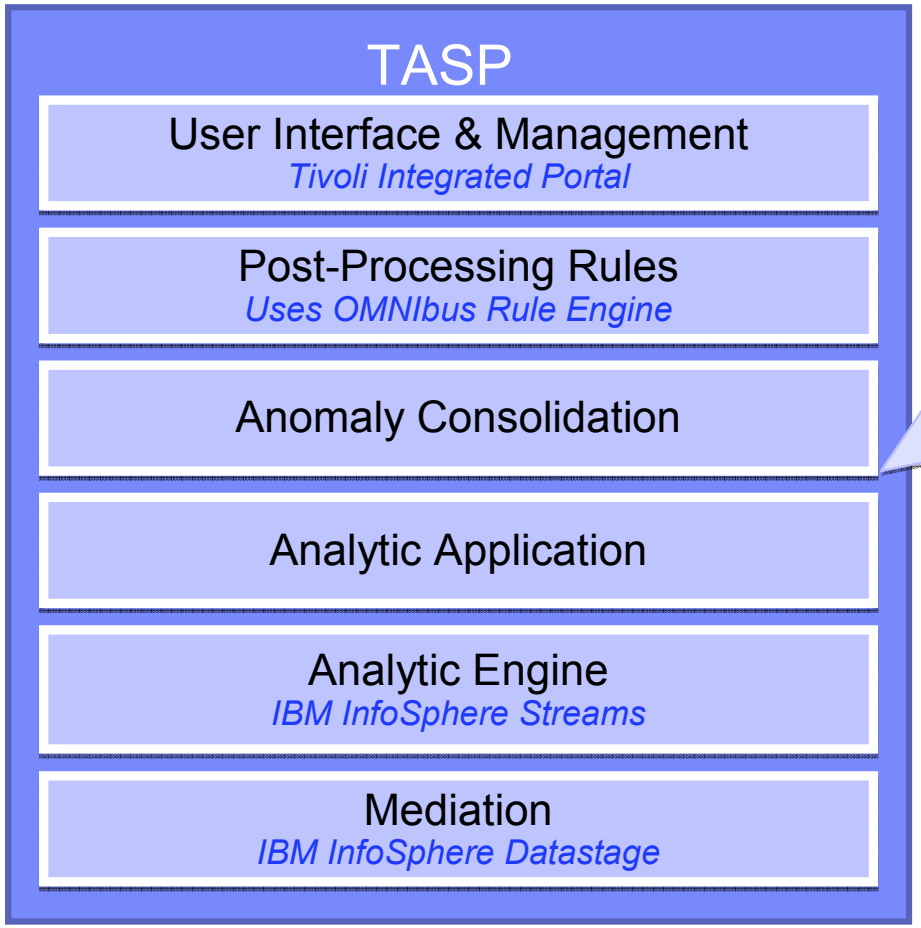

- Targeted for next release, the alarm consolidation framework reduces the events that are presented externally allowing for efficient processing and accurate alerts.

- Different techniques will be selectable depending on the richness of the data processed.

Internal Events	External Events
UV: Node A: Metric 1	EXT: Node A, B, C, Metric 1, 2, 3
UV: Node B: Metric 2	EXT: Node M, Metric 47
U V: Node C: Metric 3	
MV: Node B, C: Metric 2, 3	
MV Node A, B, C: Metric 1, 3, 2	
UV: Node M: Metric 47	

- It reduces the volume of external alarms forwarded to event consoles or application/domain administrators, without removing any information that could be useful in prediction, detection or RCA.

Clustering Analytics – Consolidation of Events

A problem in the application will produce an event for each multiple relationship and an event for each unrelated metric – results in *more events than desired...*

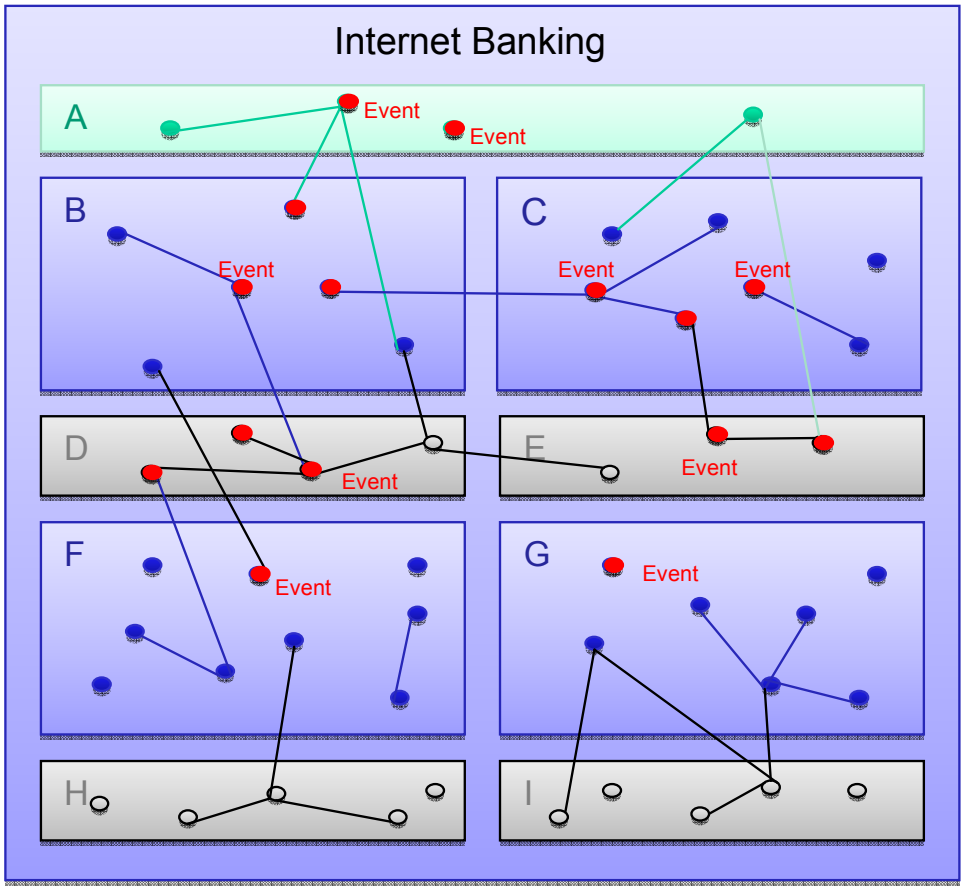

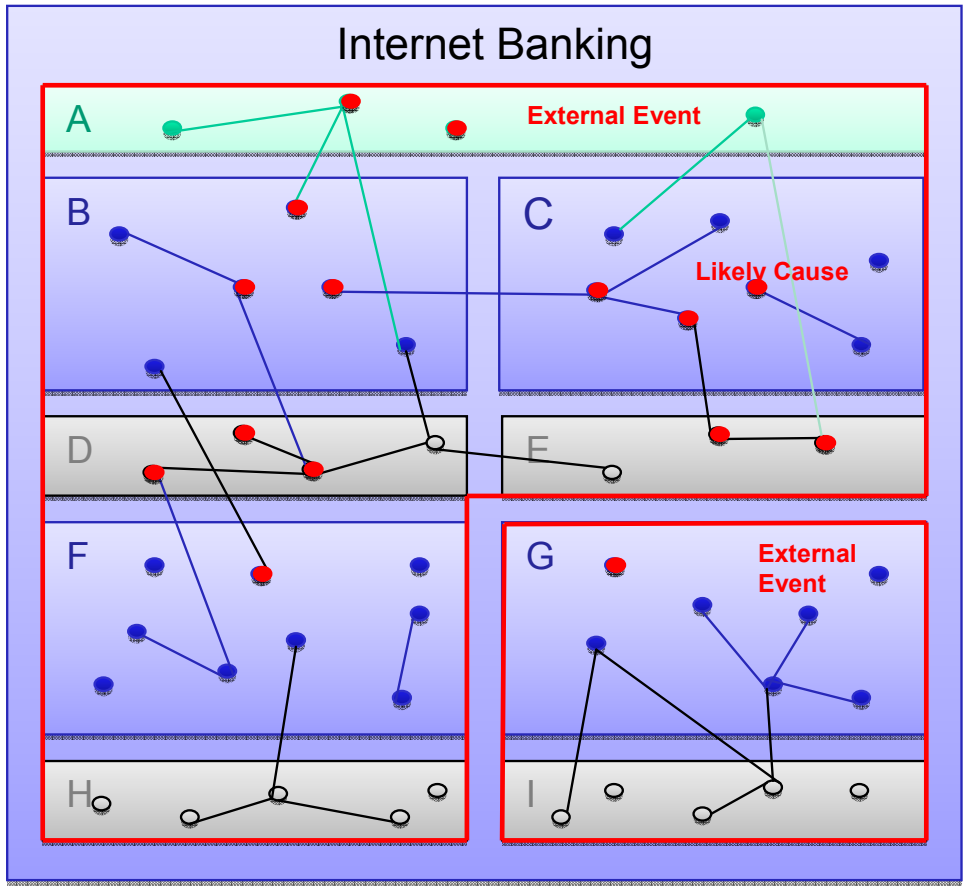

Clustering analytics is required to consolidate events **and** identify metrics likely to be close to the problem location (first symptom)

Solution Architecture – Anomaly Post Processing

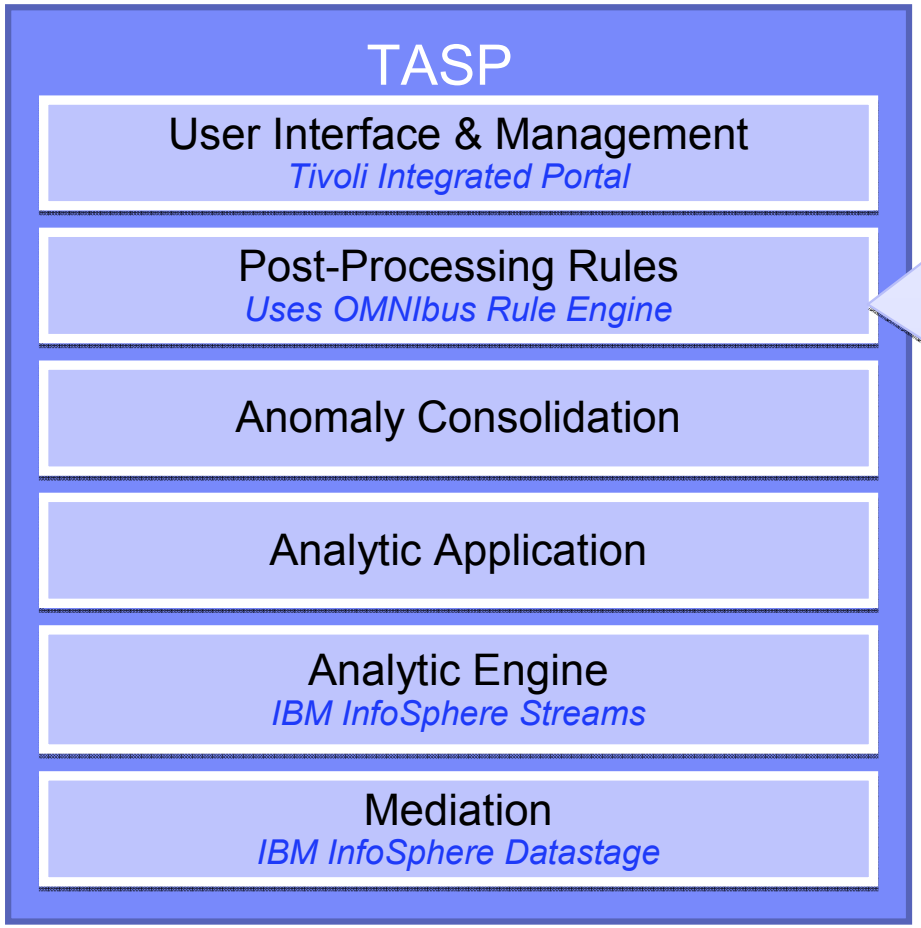

- The post-processing engine allows anomaly events to modified, customized, or enriched
- It can an be optionally used to put some 'business/domain' context around the 'domain agnostic' analytic anomaly events.
- Typically this will be used to re-prioritize anomaly severity to 'major' if it is service impacting.

Internal Events	External Events
UV: Node A: Metric 1	EXT: Node A, B, C, Metric 1, 2, 3
UV: Node B: Metric 2	EXT: Node M, Metric 47
U V: Node C: Metric 3	
MV: Node B, C: Metric 2, 3	
MV Node A, B, C: Metric 1, 3, 2	
UV: Node M: Metric 47	

ng

Web Response Time', then the anomaly severity can be changed to 'Major'.

- This reuses OMNIbus Probe rules libraries, but is dependent on having a northbound OMNIbus object server to receive the anomaly events.

Solution Architecture – User Interface & Management

TASP

- User Interface & Management
  Tivoli Integrated Portal
- Post-Processing Rules
  Uses OMNIbus Rule Engine
- Anomaly Consolidation
- Analytic Application
- Analytic Engine
  IBM InfoSphere Streams
- Mediation
  IBM InfoSphere Datastage

• TIP based anomaly visualization

• Allow all anomalous metric to be visualized together

• Normalizes metric scales, and allows, pan/zoon etc, so that anomalous conditions are more readily apparent.

• In-context linking between OMNIbus, TBSM, ITMM AEL and anomaly charts

Example: Field Trial at Large Retail Bank

Retail Bank experiencing severe problems with their online banking application

Trial Scope:

- Online banking service with back end application
- ITM AIX, Linux, Windows, ITCAM for WAS, ITCAM for WRT
- ~80 servers
- ~40k metrics

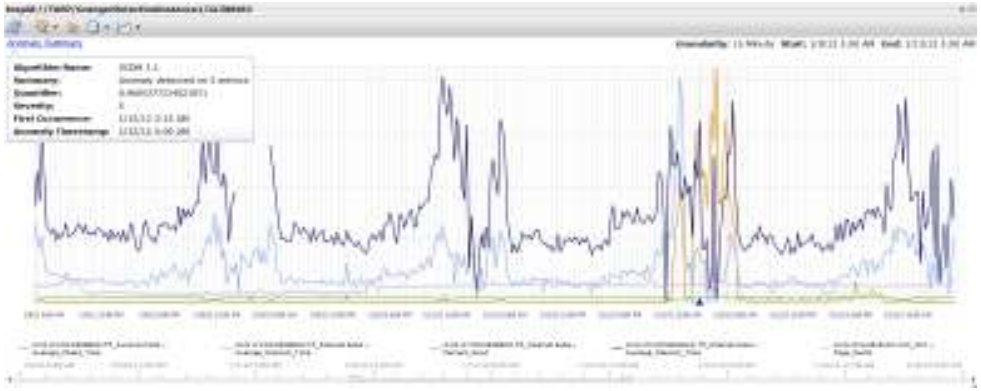

Results:

- 15 Major Incidents reported during the 4 week trial period
- 10 major incident were detected or predicted by TASP
- 5 missed incidents were application code problems and not manifest in health metrics
- 100% of “detectable problems” detected

Prediction & Detection Intervals:

Report included a ‘Problem Start Time’, a ‘Problem Detection Time’ and ‘Problem Resolution Time’

- 6 out of the 10 detected incidents were predicted before the customer’s ‘Problem Start Time’
- All 10 out of 10 detected problem were detected before or around the customer’s ‘Problem Detection Time’ interval

Results for this Customer

- Using industry average outage costs, potential outage avoidance savings for 4 weeks: ***$600k***
- Event reduction savings for 4 weeks***: $53k***

Download Open Beta Today

- Open Beta and Demo available for download today

- http://www.ibm.com/developerworks/servicemanagement/bsm/tasp/index.htm

Operations / Performance Data is Exploding

A typical enterprise with 5000 servers, running 125 applications across 2 to 3 data centers generates in excess of 1.3 TB of data per day

- Only 3% of the data generated is operations oriented metric data.
- 97% is made up of unstructured/semi structured data
- Workloads are running on heterogeneous platforms.

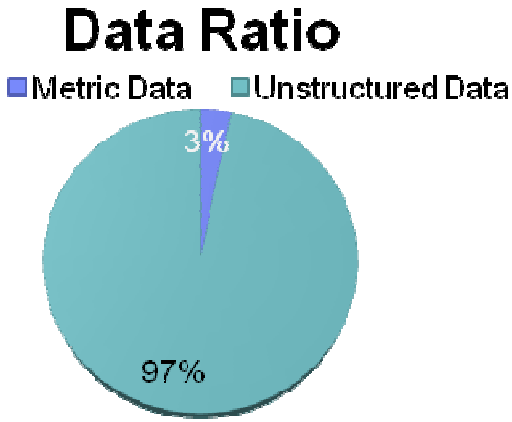

SmartCloud Analytics – Log Analysis Delivers Faster Problem Resolution

- Search, and Index unstructured data to provide consolidated view

Faster Problem Resolution
Find & correct problems faster with tools that determine actions required to resolve issues

Resolve

- **Identify** problems quicker with insight to large unstructured repositories
- **Isolate** problems quicker by bringing relevant unstructured data into problem investigations

- Built on IBM's Big Data platform
- Integrate structured and unstructured data for better problem identification and resolution
- Extensible, with IBM and partner expertise built-in
- Get the last critical piece of data for identifying, isolating, and correcting problems faster

IBM SmartCloud Analytics – Log Analysis in Brief

Collects large volumes of obscure unstructured data and transforms it through analytics into actionable intelligence.

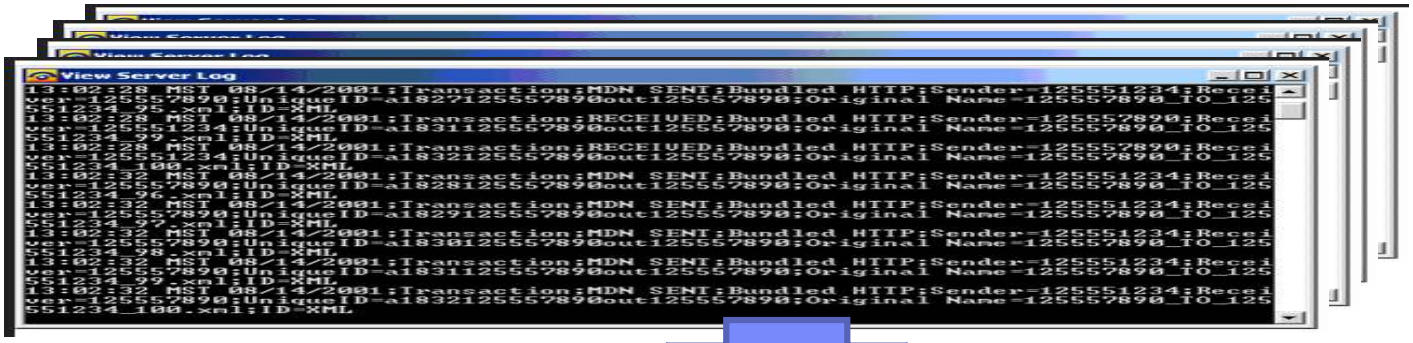

GBs of Obscure Log Files

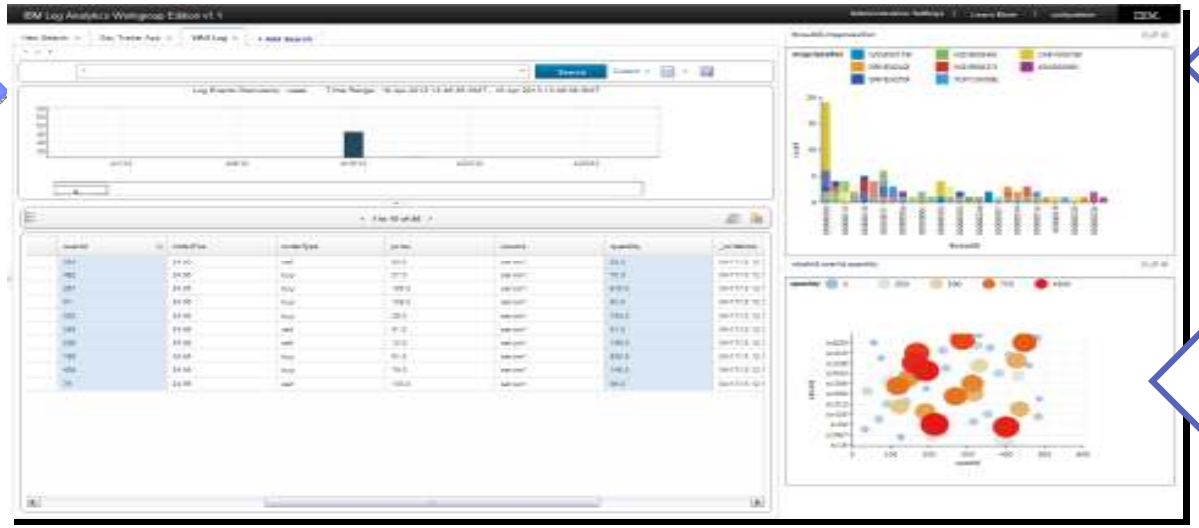

Single Actionable Dashboard

Intelligent Support Docs Integration through Advanced Text Analytics

Insight

IBM Log Analytics Client Value

IBM Log Analytics helps IT Generalists and Application Specialists accelerate problem resolution through rapid analysis of unstructured data

Value	Highlights
• **Faster Problem Identification and Isolation** • Quickly search structured and unstructured data. • Perform cross domain analysis on this data.	• Collection and Annotation of data • Generic Logs Support • Federation of Data
• **Faster Problem Repair** - By linking expert knowledge to log error/warning messages	• Advanced Text Analytics • Downloadable insight packs on the ISM Library starting with WebSphere and DB2
• **Improved Service Availability and Maintainability of Custom Apps** - Provide users with advanced insights into custom applications quickly	• Tools to create custom insight packs for your own applications

Sample App dashboard

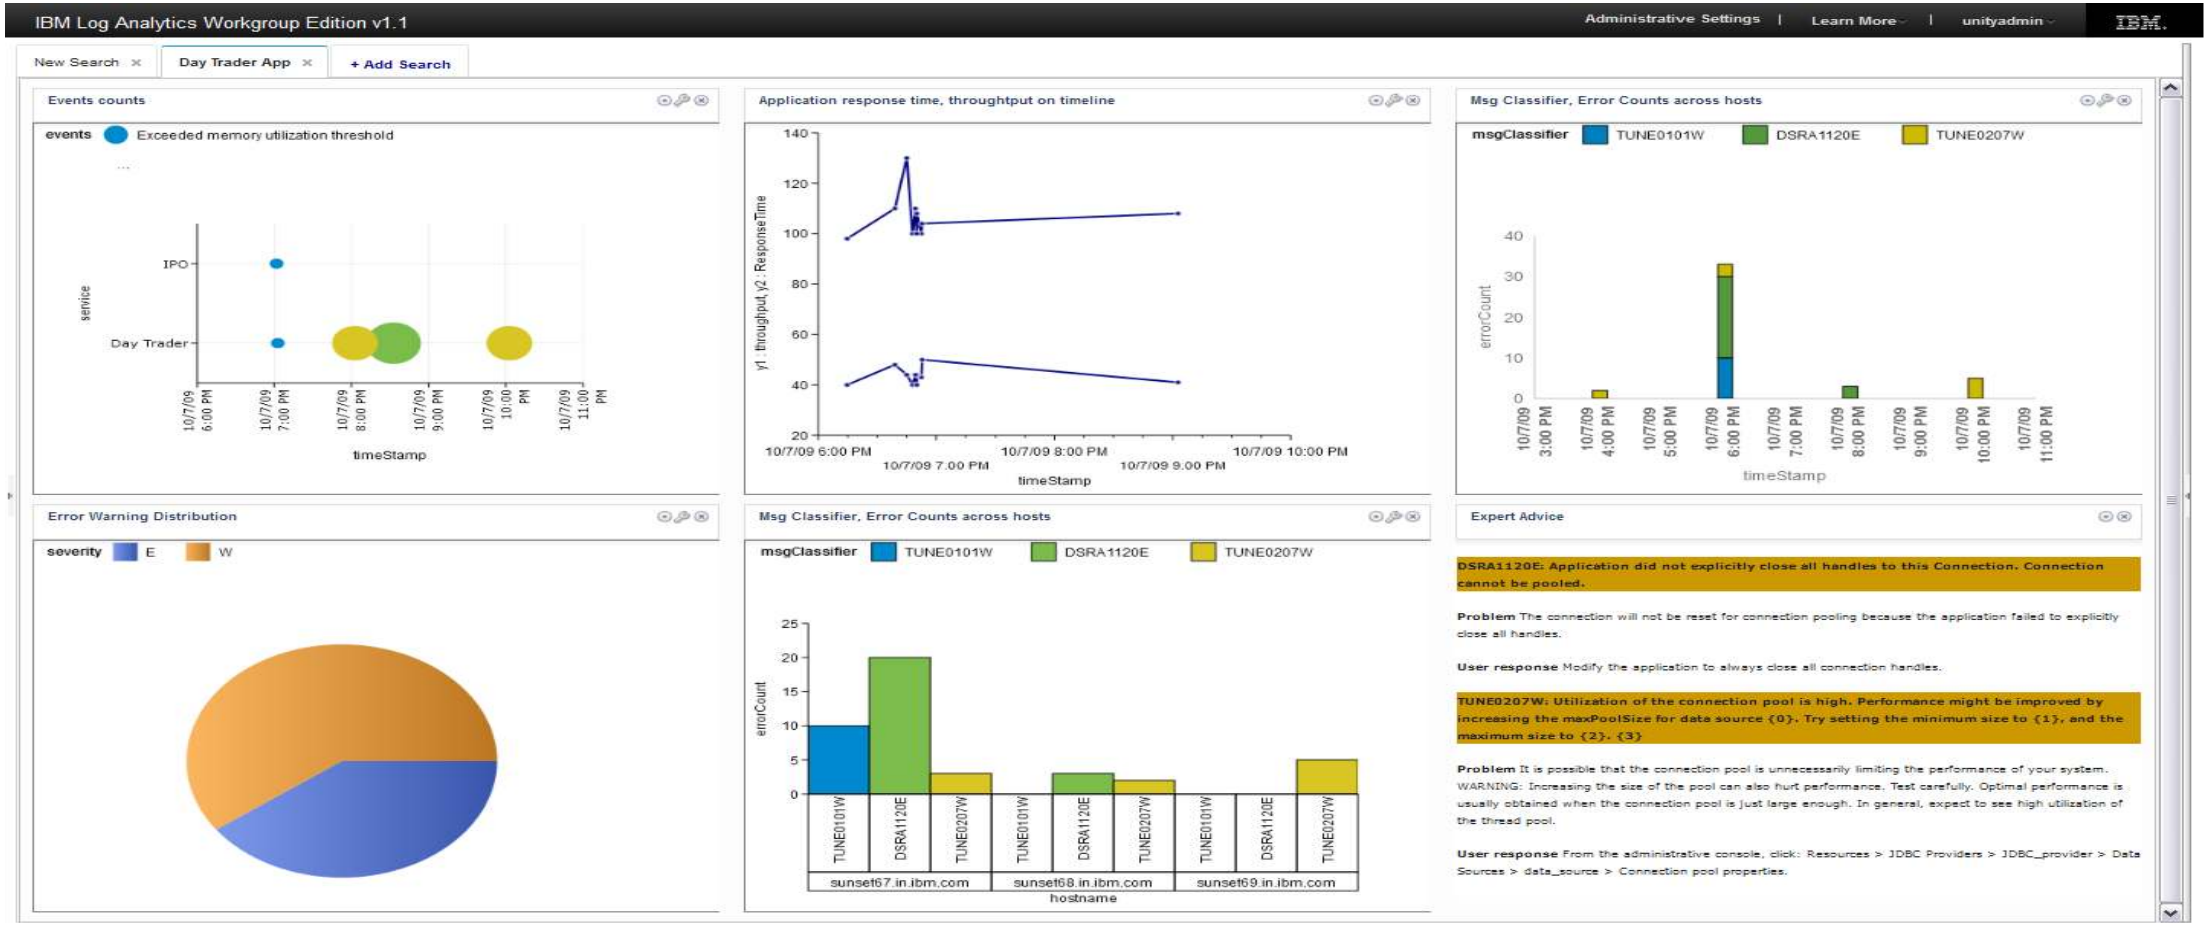

Log Analytics Free Trial – Available TODAY!

- Download Log Analytics today via the service management connect website

http://www.ibm.com/developerworks/servicemanagement/bsm/log/index.html

A Healthcare Provider reduces time to diagnose system problems by providing a holistic view of all relevant data

Need

- Have too many tools across structured and unstructured datasets making problem resolution difficult and time consuming
- Desired a solution to time-correlate a view into many sources of data to perform problem detection, isolation and repair

Benefits

- Reduced time to determine root cause of problems by leveraging performance, event and log data
- Skills required to diagnose problems were easily saved and repeated to reduce overall costs

Next Release - Improved UI Diagnostics & Visualization

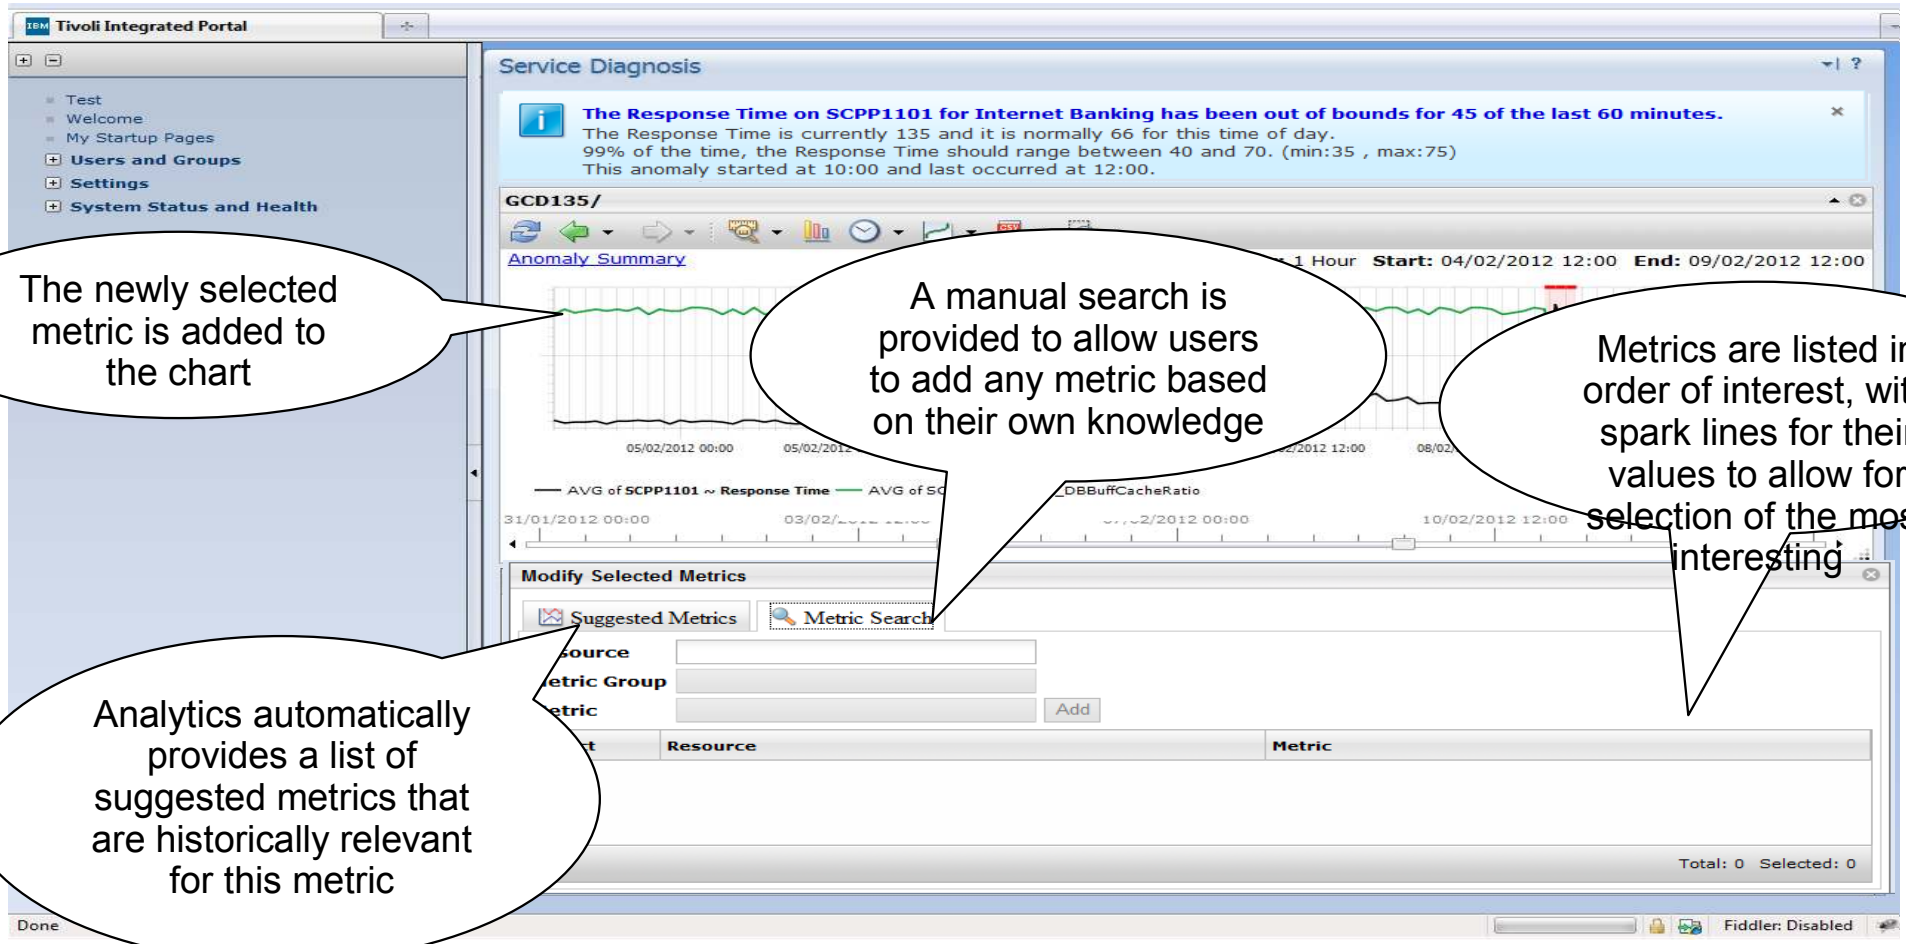

Next Release - Improved UI Diagnostics & Visualization