Predictive analytics for IT and service management

Daniel Roscigno,

IBM Tivoli Operations Architect

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.

How can we avoid business disruption? And decrease operational costs?

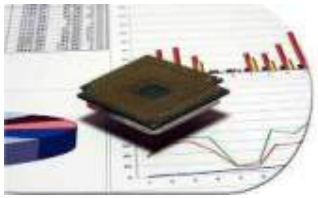

Operational costs of systems and networking are rising
Maintenance costs: $8 spent for every $1 spent on new infrastructure*

* Source: IDC 2007

There's an explosion in volume of data and information
Virtualization and cloud bring increased dynamicity and change
Unpredictable workload characteristics

Challenge: Improve the assurance of physical and virtual environments across applications, systems, networks and storage.

Analytics for Service Assurance approach: Add advanced analytics over operational data to detect problems before they become service affecting.

Move from "reactive" management to fixing things before they break.
Predictive analytics!

"I'd like 30 minutes warning to know when my user experience is going to deteriorate"

"After we fix the problem, then we set another threshold"

"What are realistic baselines? How can I reduce false alerts?"

Analytics…

why is this happening

what if these trends continue

what will happen next (that is, predict)

what is the best that can happen (that is, optimize)

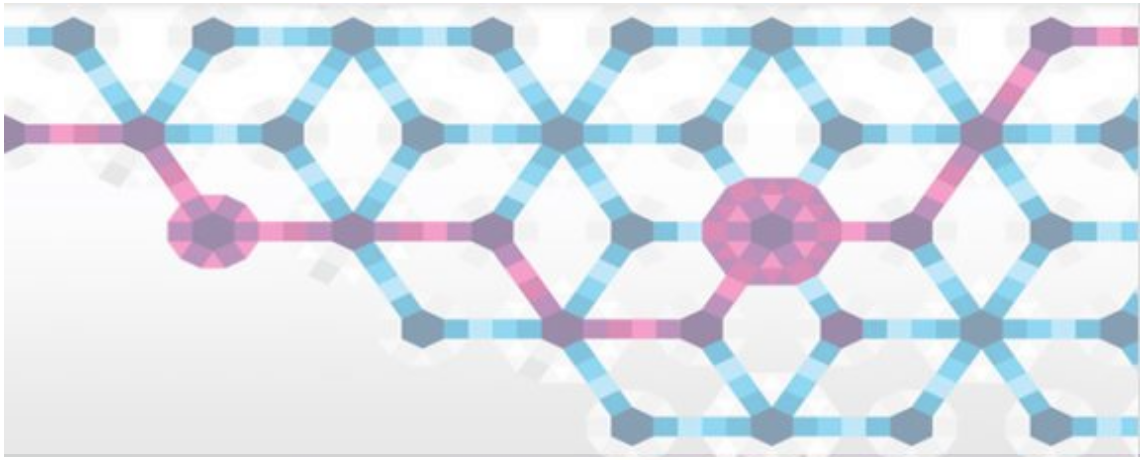

IBM is driving the future of integrated analytics through acquisitions and our strategic partnerships:

NETEZZA
SPSS
i2
EXEROS
unica
princeton softech
ILOG
Coremetrics
COGNOS
solid.
FileNet
DataMirror
dwl
Trigo
iphrase
SRD
Ascential
unicorn

IBM Research Business Analytics and Optimization

- Over 200 researchers with expertise in data analytics, operations research, mathematics, and industry applications of analytics
- Hold 300 patents and have an additional 450 pending on analytics and business applications
- Support IBM's "fact-based" management and processes in sales, supply chain, and services.
- Participate with IBM consultants in cutting-edge client projects
- Add differentiation through data analytics to outsourced accounts
- Provide new algorithms to IBM SW products
- Lead in the global scientific community
  - Over 250 publications in leading conferences and journals in recent years
  - Fellows at several leading professional societies
  - Successive wins at KDD Cup and INFORMS Data Mining Competitions (premier competitions)
  - Leaders in Optimization Open Source
  - Major INFORMS prizes and awards
  - Adjunct faculty at leading universities

Improved profitability through analysis of customer networks for a major telecom customer by providing better customer targeting

Deployed operational planning and scheduling to run steel plants at several leading Asian steel manufacturers – improved productivity

Design optimal maintenance plan for a set of interconnected offshore oil platforms - improve availability of oil platforms

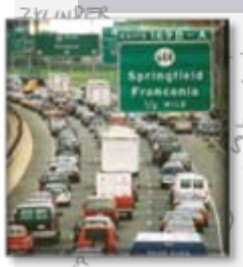

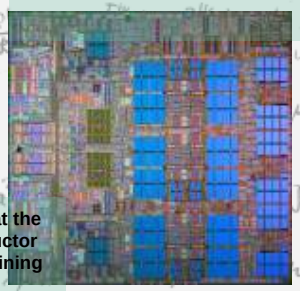

Analytics-Driven Solutions for Increased IBM sales force productivity – increased revenue and profitability

Improved wafer yield at the IBM 300mm semiconductor plant deploying data mining and machine learning

Creation of state-of-the art error correction code technology that is used in main memory systems of IBM's computers

Provide analysis of operational risk loss data for 36 leading banks from 13 countries - cross enterprise secure and anonymous data sharing

Centralized control and real-time visibility of the end-to-end supply chain for IBM supply chain – reducing inventory

IBM Business Analytics

Software that Addresses Key Customer Needs

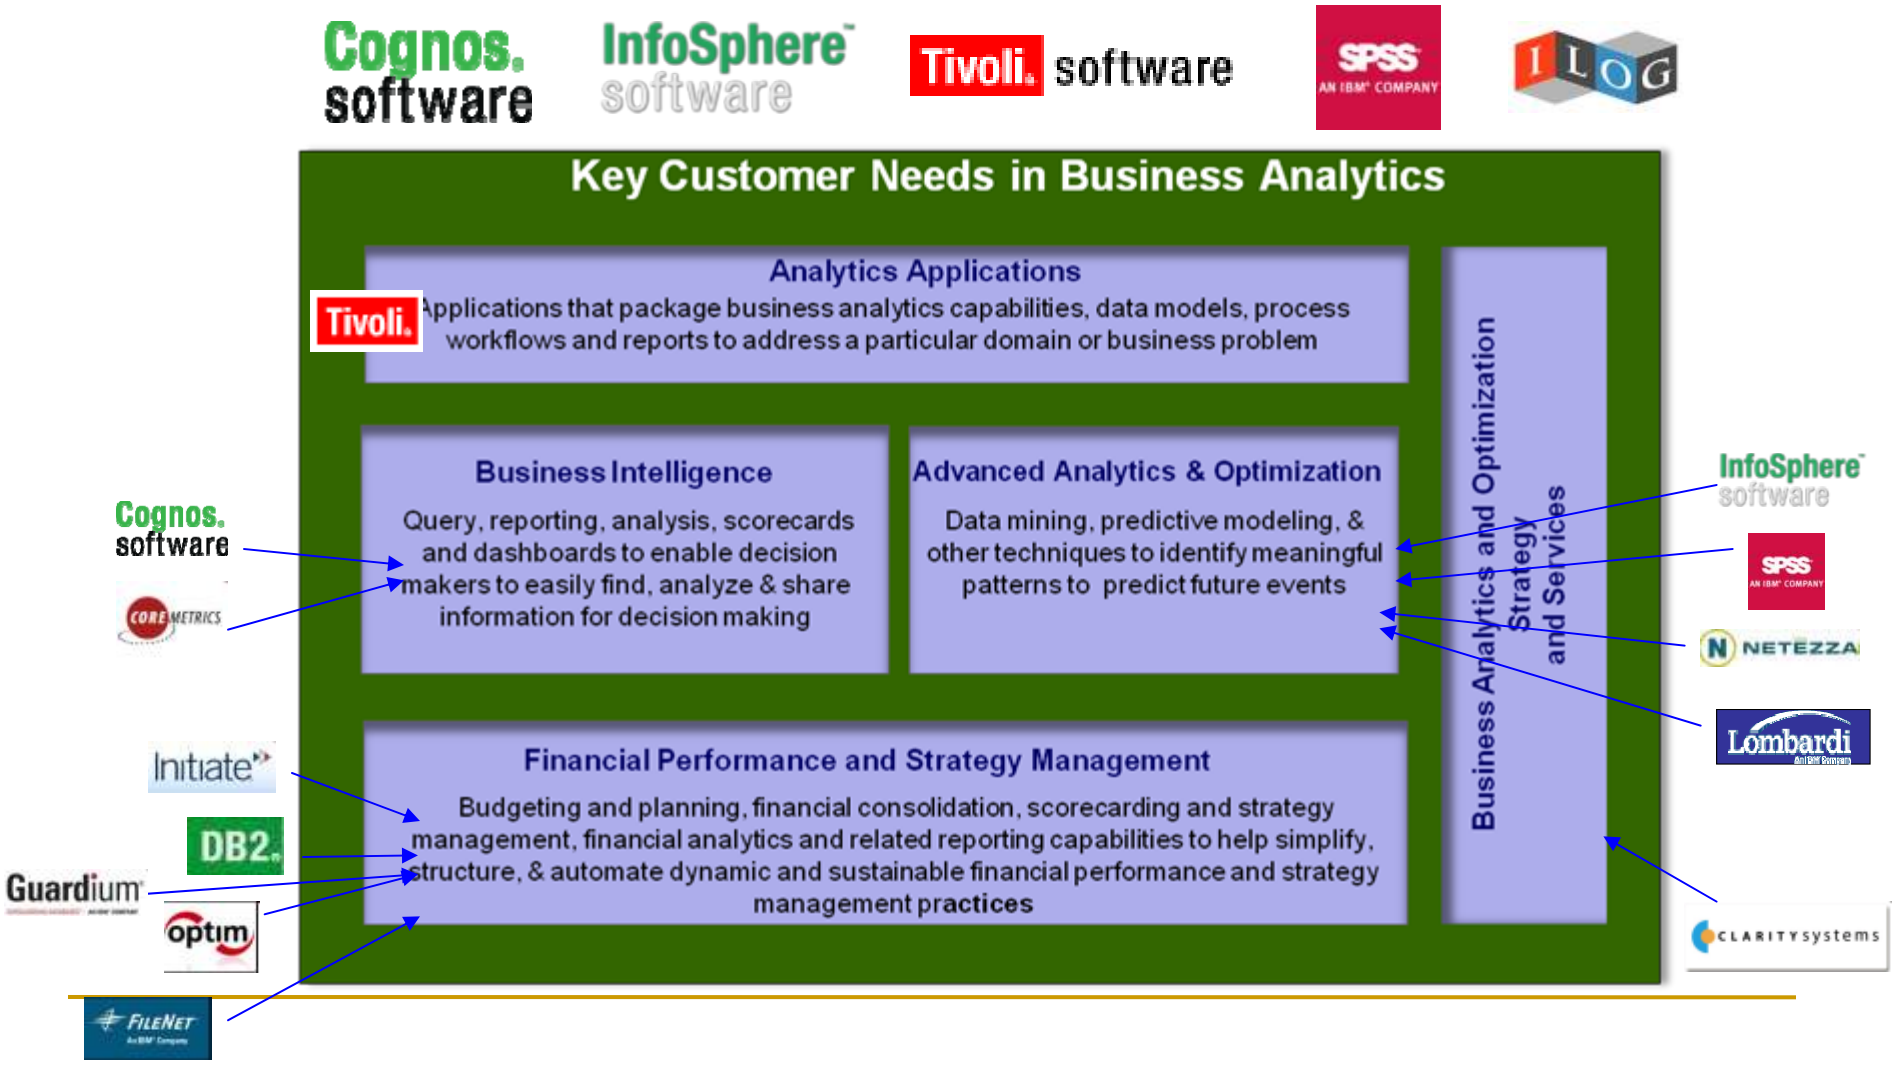

Breadth of operational analytics for ISM

Tivoli. software

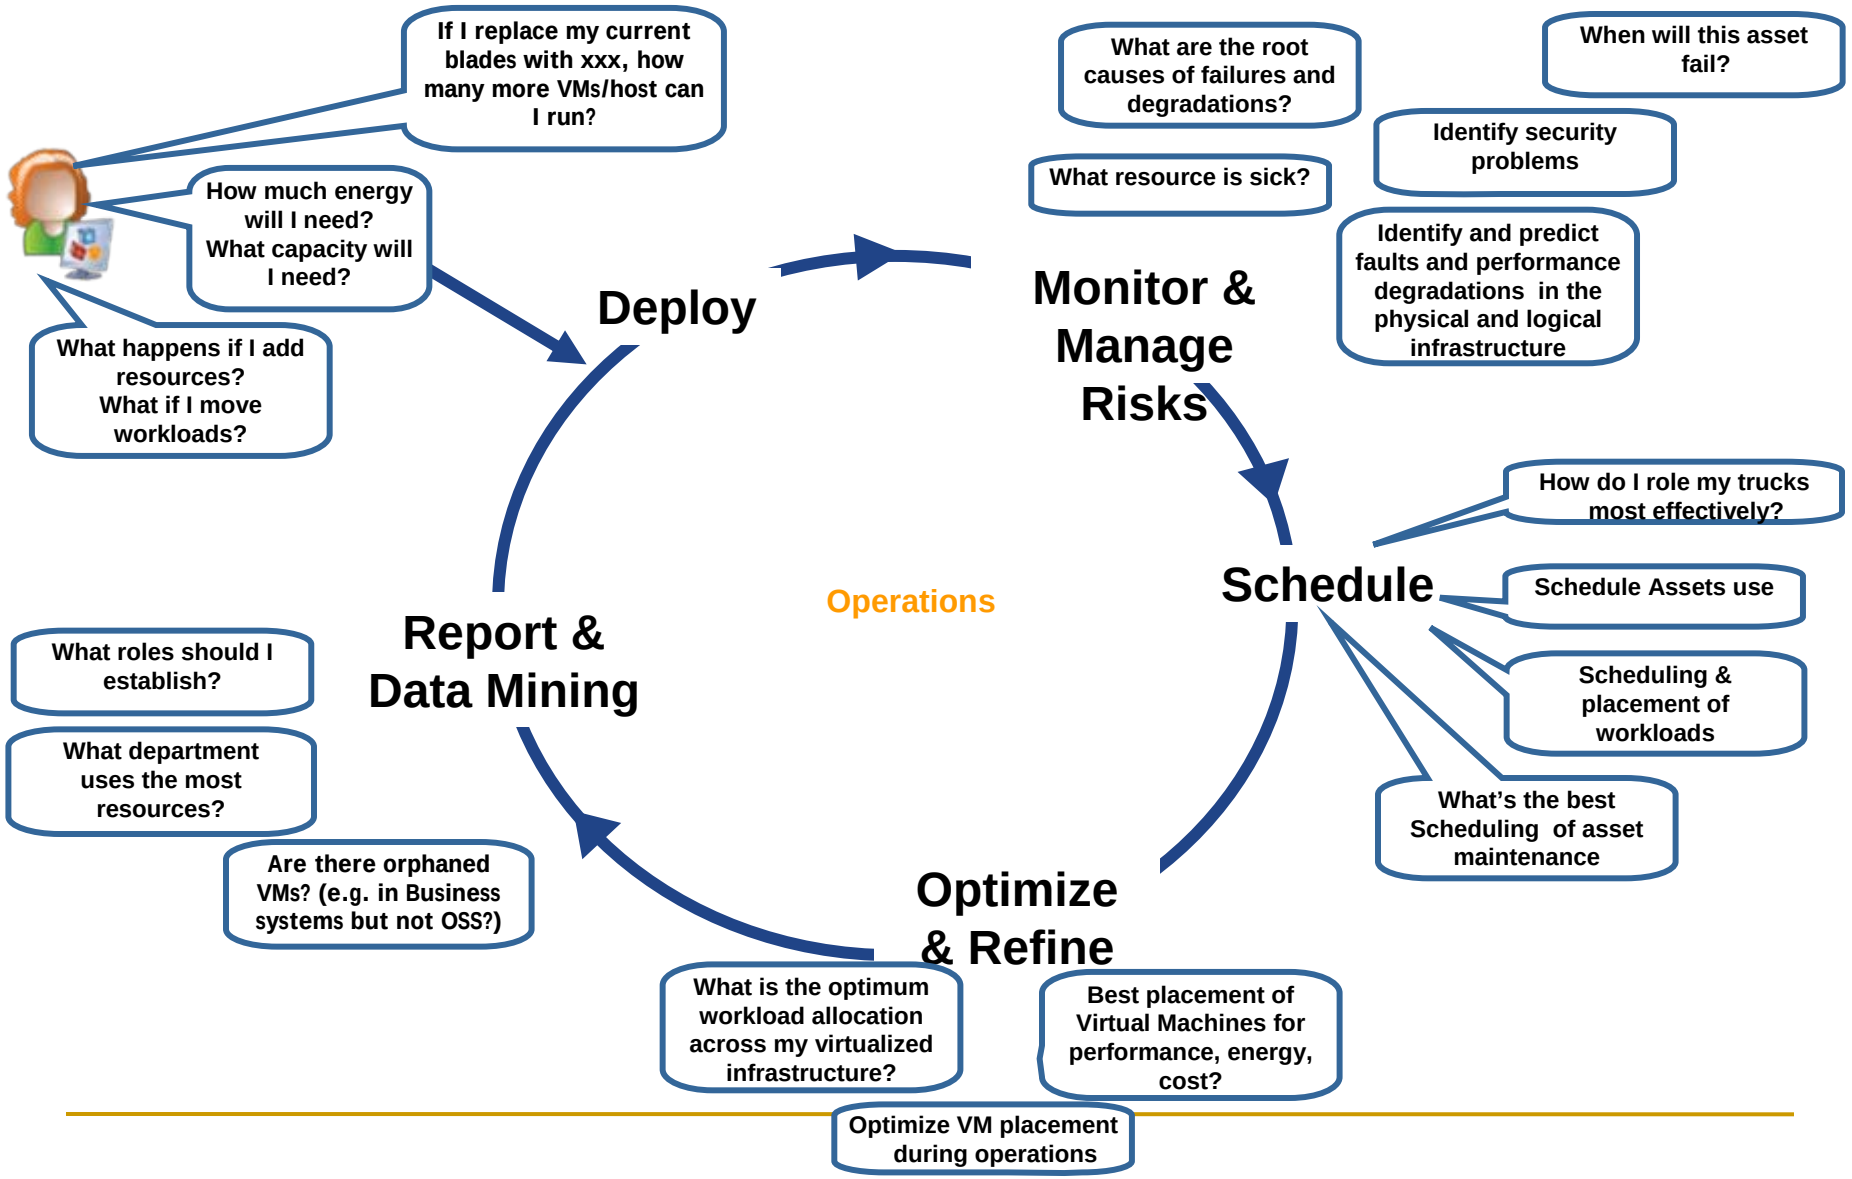

Analytics in IBM Tivoli Service Management Today

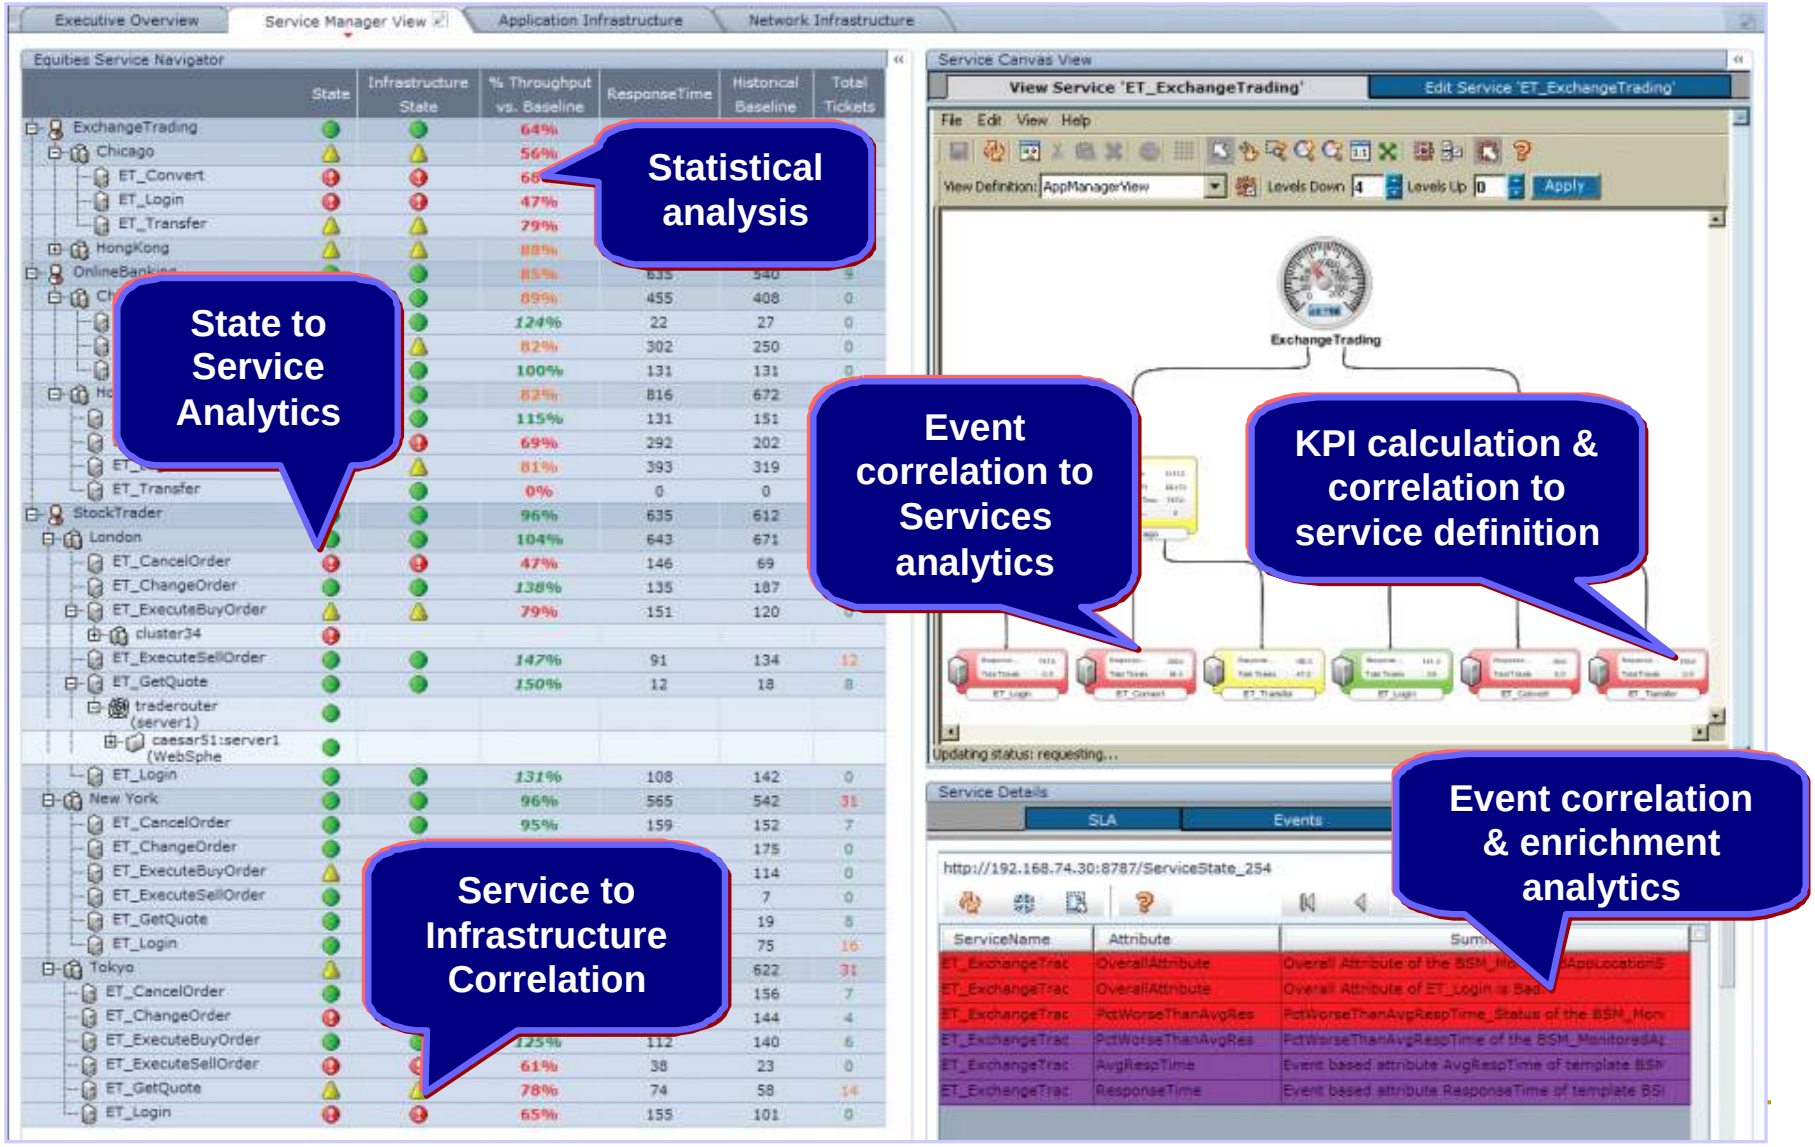

What is Predictive Analytics?

Adaptive Monitoring

Topology Correlation

Predictive Modelling

Self-learning

Predictive Analytics enable IT organizations to move from reactive to proactive management of services, reducing outages and improving business performance.

"Analytics leverage data in a particular functional process (or application) to enable context-specific insight that is actionable." - *Gartner*

Move the sensing and alerting, and eventually actions to earlier and earlier

Forecasting

Anomaly Detection

Dynamic Thresholding

Forward Trending

- Advance warning of service impact, deterioration or outage
- Realistic service baselines
- Avoidance of expensive and time-consuming false alerts
- Detection of service impacts that are not identified by fixed thresholds alone
- Swifter diagnosis of certain events and patterns
- Identification of the underlying root cause to implement fixes

IDC study: Predictive analytics initiatives show an average ROI of 145%, in comparison to 89% for non-predictive analytics*

* Source: "Predictive Analytics and ROI: Lessons from IDC's Financial Impact Study" paper, Henry D. Morris

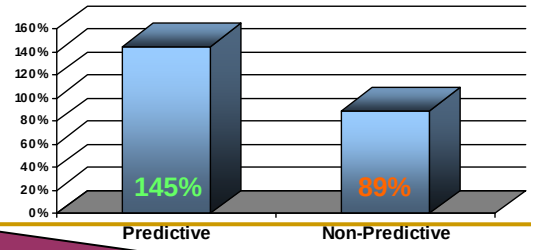

"After we fix the problem, then we set another threshold"

"I'd like 30 minutes warning to know when my user experience is going to deteriorate"

Operations challenge: Balancing the need to manage more with improved service levels & lower costs

"..multi-dimensional relationships between dynamic infrastructure and changing business services are too complex for IT staff to continue reacting to event storms and constantly tweaking static monitoring thresholds." *

* Source: PNA, 2008

"40% of unplanned downtime due to operator error" **

** Source: Gartner, March 2009

Challenge: Can you create and maintain increasing numbers of thresholds and situations in a constantly changing IT environment? How can you minimize the number of alerts that operators must handle?

Service Assurance Analytics approach: Intelligent or 'Predictive Events' that result from speeding the ability to detect abnormal trends before end users and mission critical applications are impacted

"I'd like my operators to be able to more quickly diagnose certain events and patterns to implement fixes."

"We have too many thresholds to maintain"

"I'd like to adjust for repeating traffic situations"

Maturation of Monitoring and Analytics

Receive trouble tickets. Fault!
React to user report

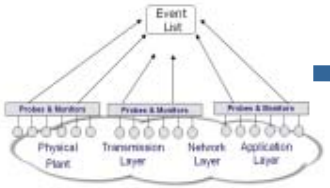

Gather alerts
React to events, before user report

Gather &Correlate alerts
React to events, w/ improved RCA

Add performance monitor
React to performance changes

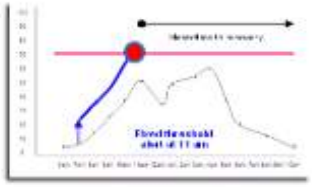

Static Thresholds
Set reasonable thresholds thresholds, create alerts when violated

Linear Forecasting
Linear prediction
One metric

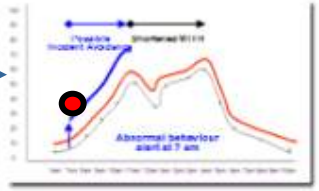

Dynamic Thresholds
Earlier Detection based on history

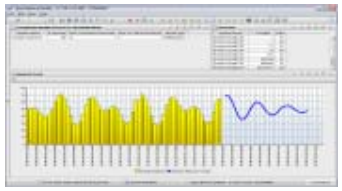

Nonlinear Forecasting
More accurate prediction of single metric

Monitoring and Analytics Approaches

Where we're going

•Low latency analysis of performance and wellness data in motion across physical and virtual infrastructures, from across the service delivery stack (from servers to hypervisors to networks to applications)

- Advanced univariate, and multivariate predictive analysis
  - With behavioral learning algorithms that can learn normal behavior during operation, and react to changes

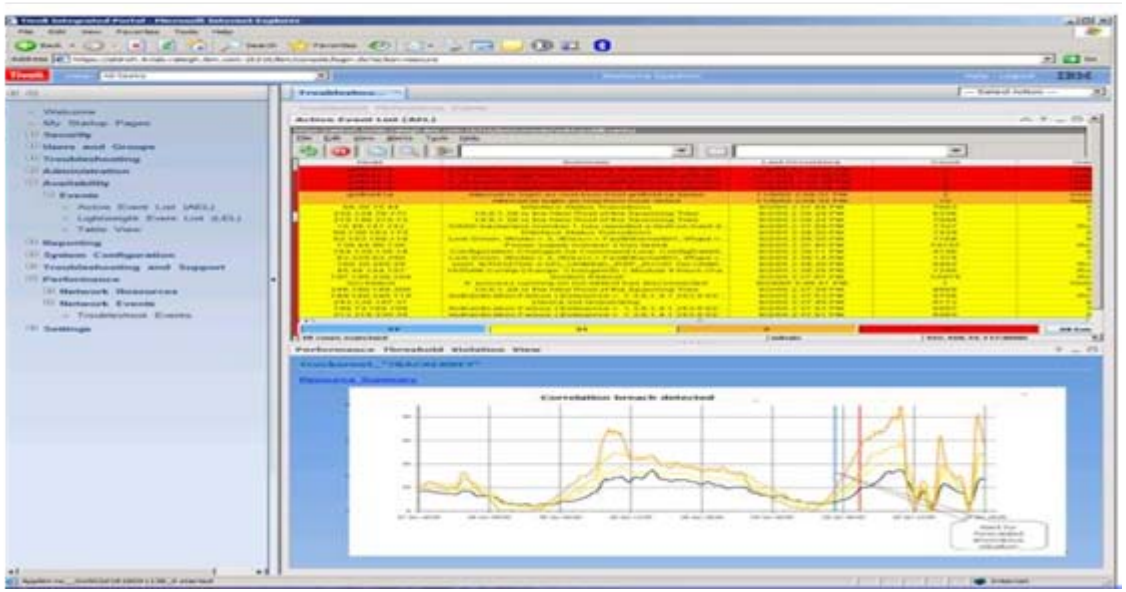

What are the use cases? ….

Learning algorithms build a model of relationships between KPIs

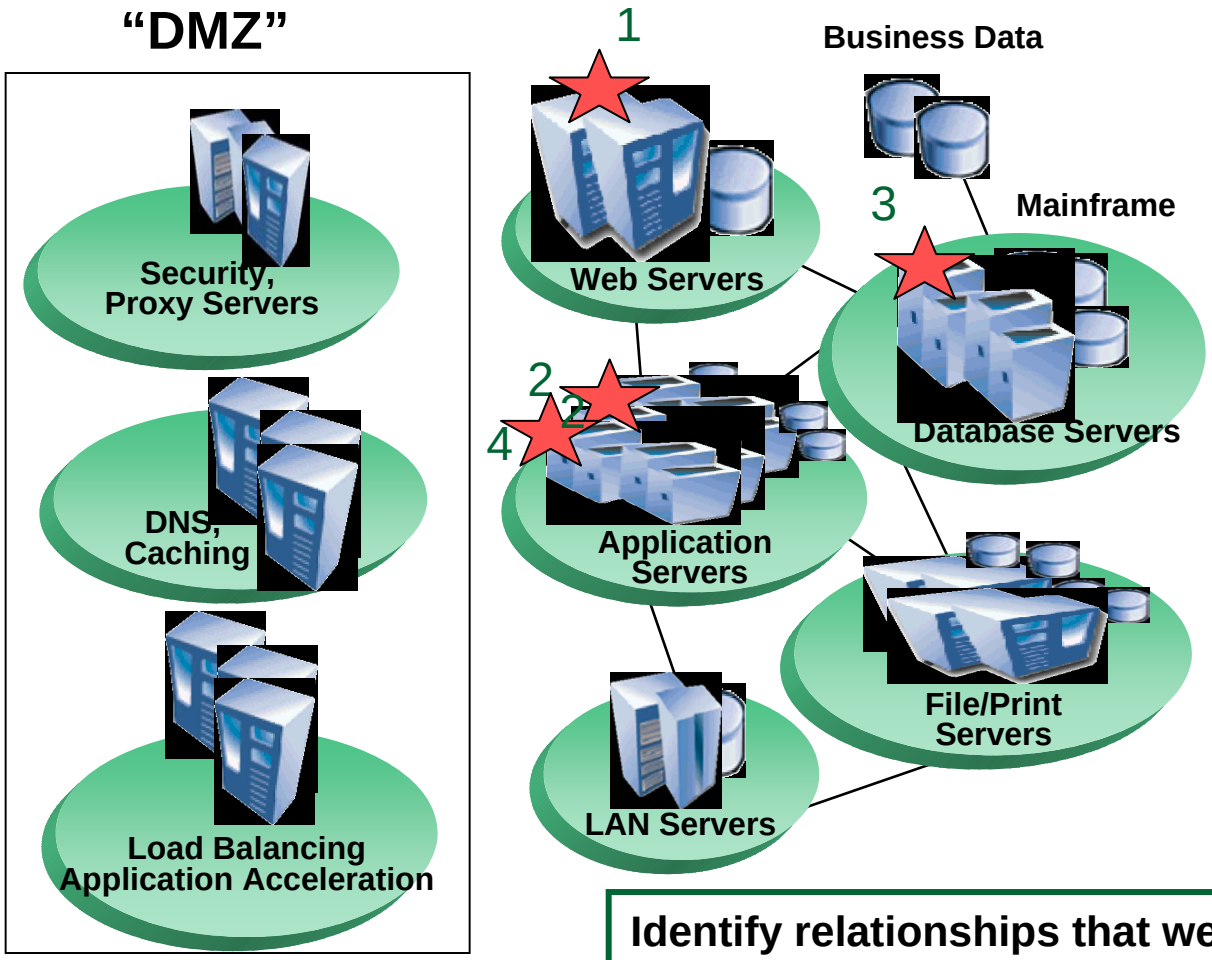

1. Web server traffic
2. Application memory
3. I/O accesses
4. Application performance

Web server traffic, application memory, and I/O accesses trend up and down together. Application performance trends opposite.

Identify relationships that were previously unknown.
Identify patterns in KPI performance between resources.

Case 1: Identify anomalies that could not be found with single KPI thresholds

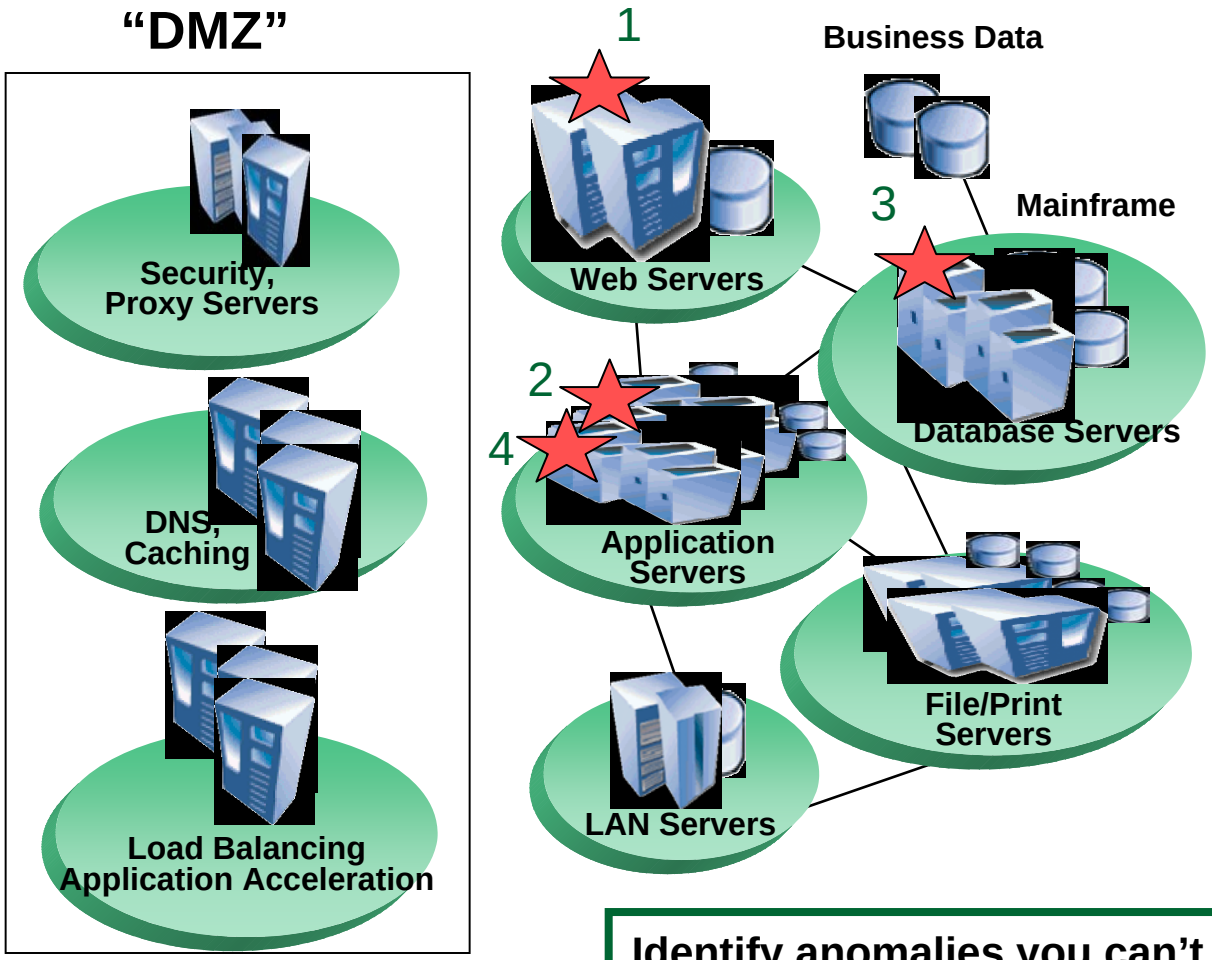

1. Web server traffic stable

2. Application memory rising

3. I/O accesses stable

4. Application performance within normal range

Identify emerging problem, even if all still "green"

Memory leak!

Identify anomalies you can't find with single KPI thresholds

Case 2: Identify problems in resources not in your management domain

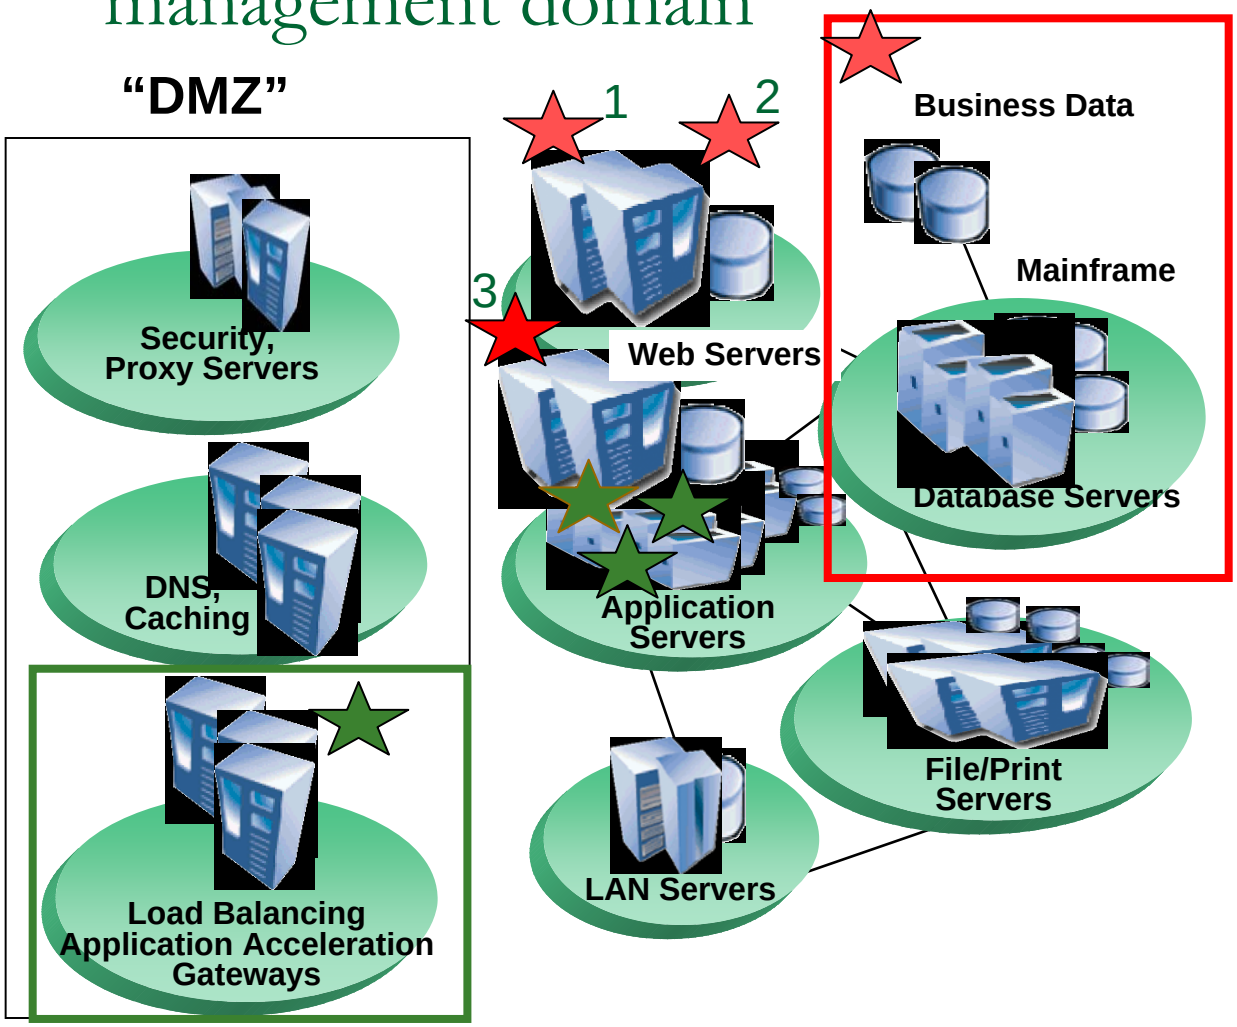

1,2,3 – Web servers response times

- Previously uncorrelated Web servers response times become correlated with each other.
- Problem in a downstream resource in a different management domain- SAN errors

Green stars – Applications performance degradation

- Gateway problems

Identify "outside" problems that will affect service

Case 3: Identify problems in use of shared resources

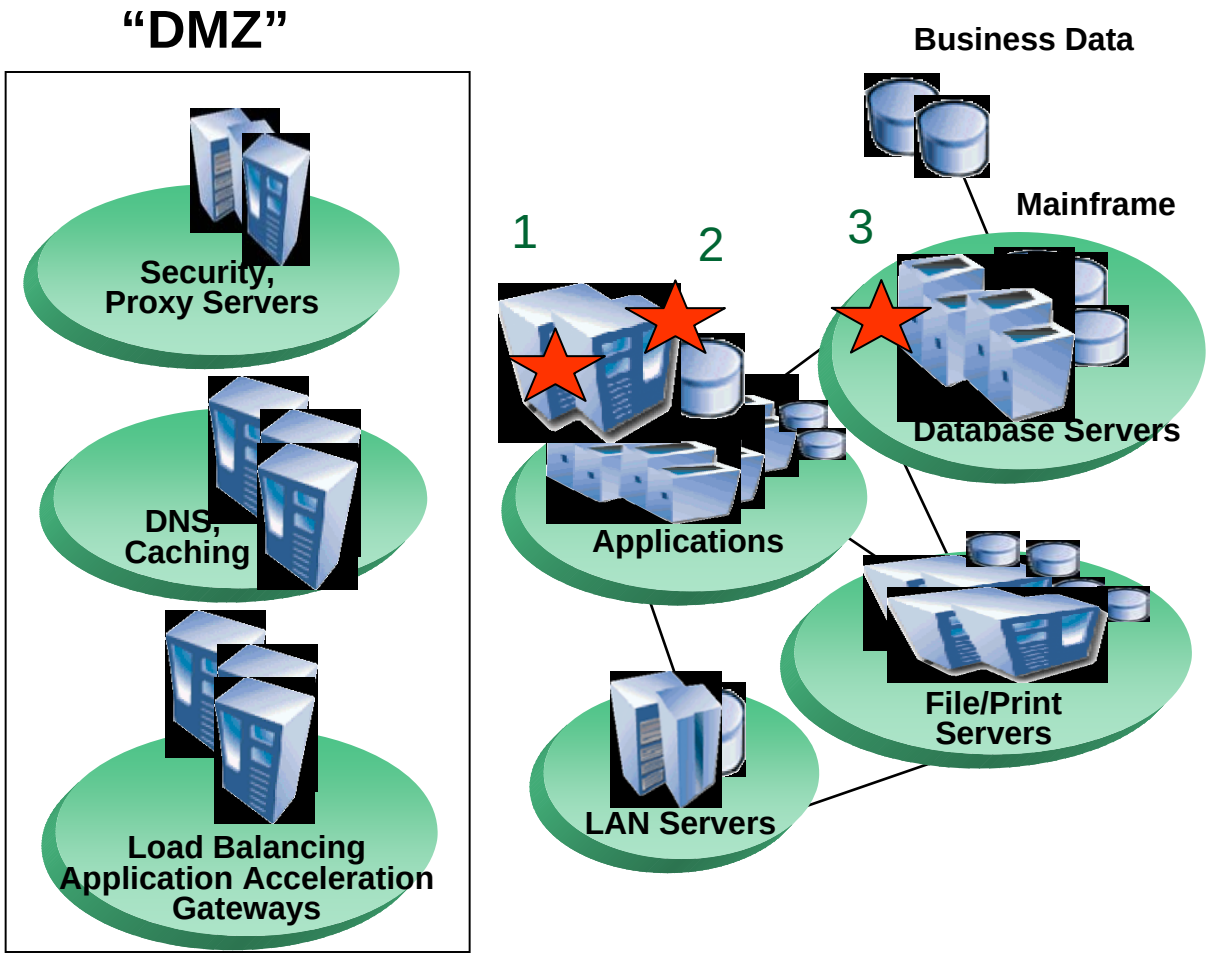

- Application performance (1), VNIC utilization (2) and DB traffic (3) are related. More traffic, lower performance.
- Analysis shows: Application performance degrades, but VNIC utilization and DB do not show corresponding increases.
- Root cause: New application deployed on same physical server. Server I/O being "hogged" by other VMs. Issue identified before retraining complete, before service levels breached

Identify resource issues caused by injudicious deployment of applications on a shared infrastructure

Streaming analytics engine

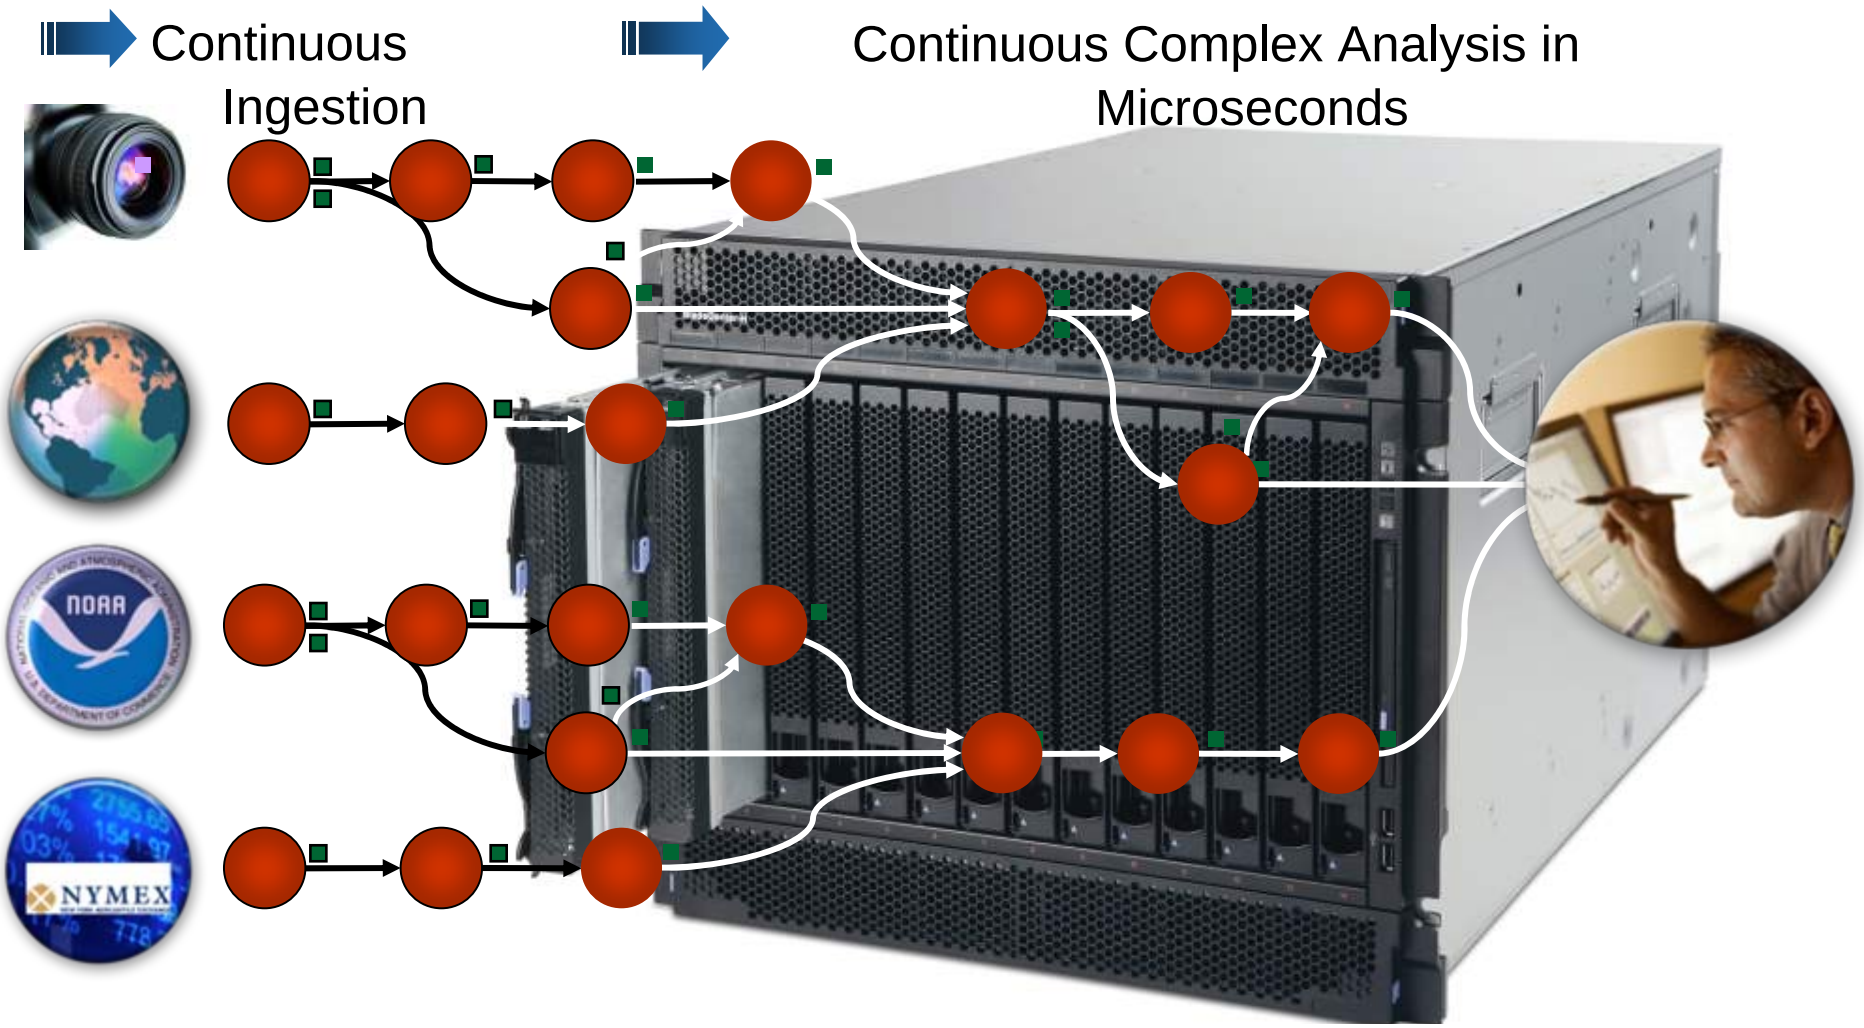

- Processes millions of events per second
- Used in finance, manufacturing, law enforcement

TASP:

- Ryhu#wkh#wrs#dqdo|wlfv/#zklfk#bhyhudjh#h{lvwlqj#p rqlwrulqj#v|vwhp v/#dqdo|}lqj#p hwulfv# iurp #sk|vlfdo#dqg#yluxdo#hqylurqp hqwv
- Xvhv#dqdo|wlfv#wr#bhduq#qrup dc#rshudwlrqdo#ehkdylrxu#dfurvv#wkh#forxg#lqiudvwuxfwxuh
- Ohduqv#p dwkhp dwlfdo#uhodwlrqvkisv#ehwzhhq#p hwulfv#dfurvv#k|shuylvruv/#YP v dqg#rwkhu# hdhp hqwv#lq#wkh#vhuylfh#ghoyhu|#vwdfn1
- Ghwhfwv#surehp v#ehiruh#wkh|#ehfrp h#exvlqhvv#p sdfwlqj1
- Vhqgv#dqrp do|#nyhqwv#wr#p dqdjhp hqw#frqvrdnv1

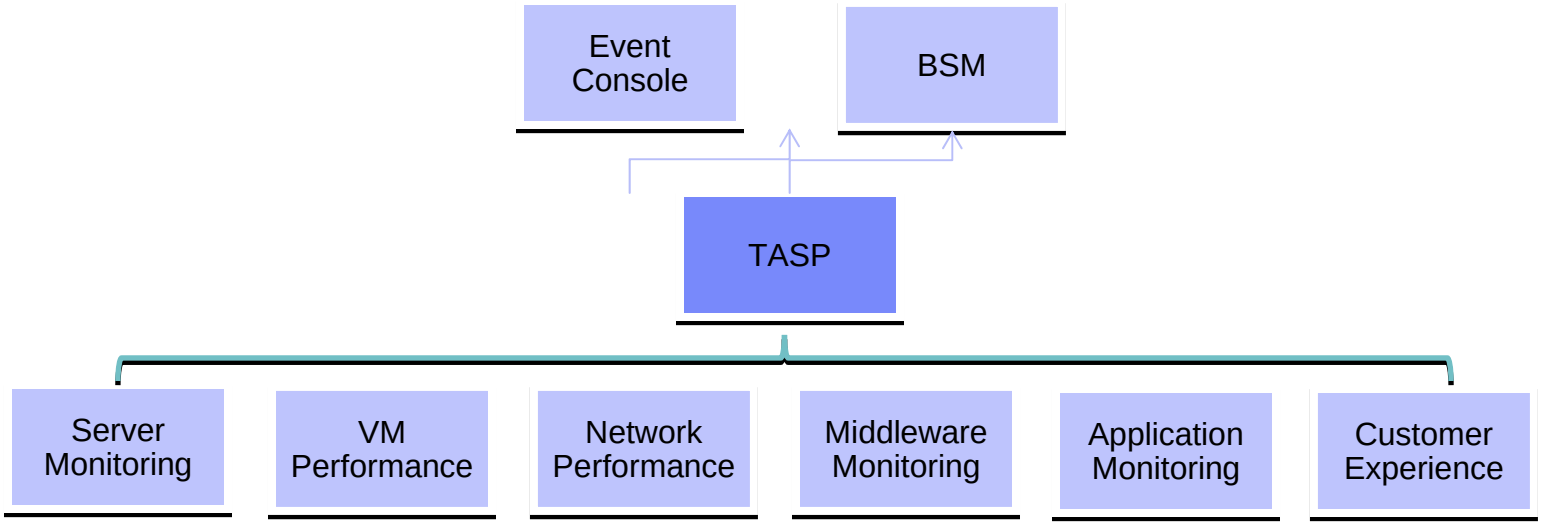

Tivoli's Predictive Analytics Lifecycle

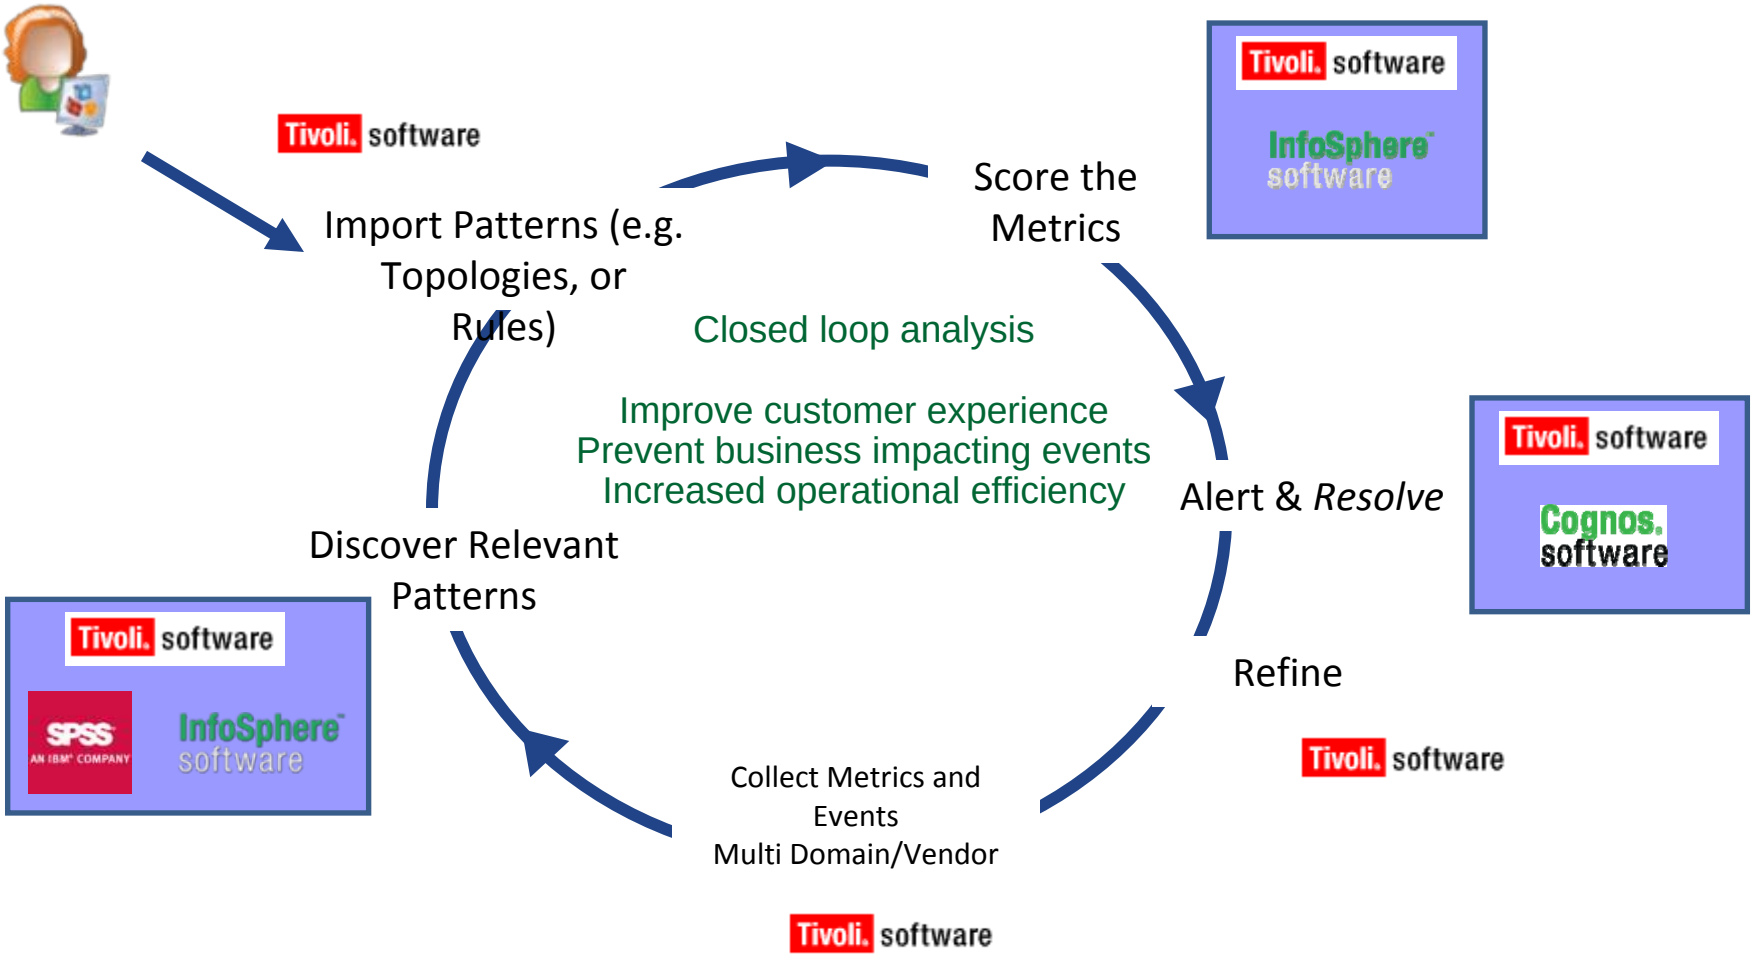

Analytics Landscape

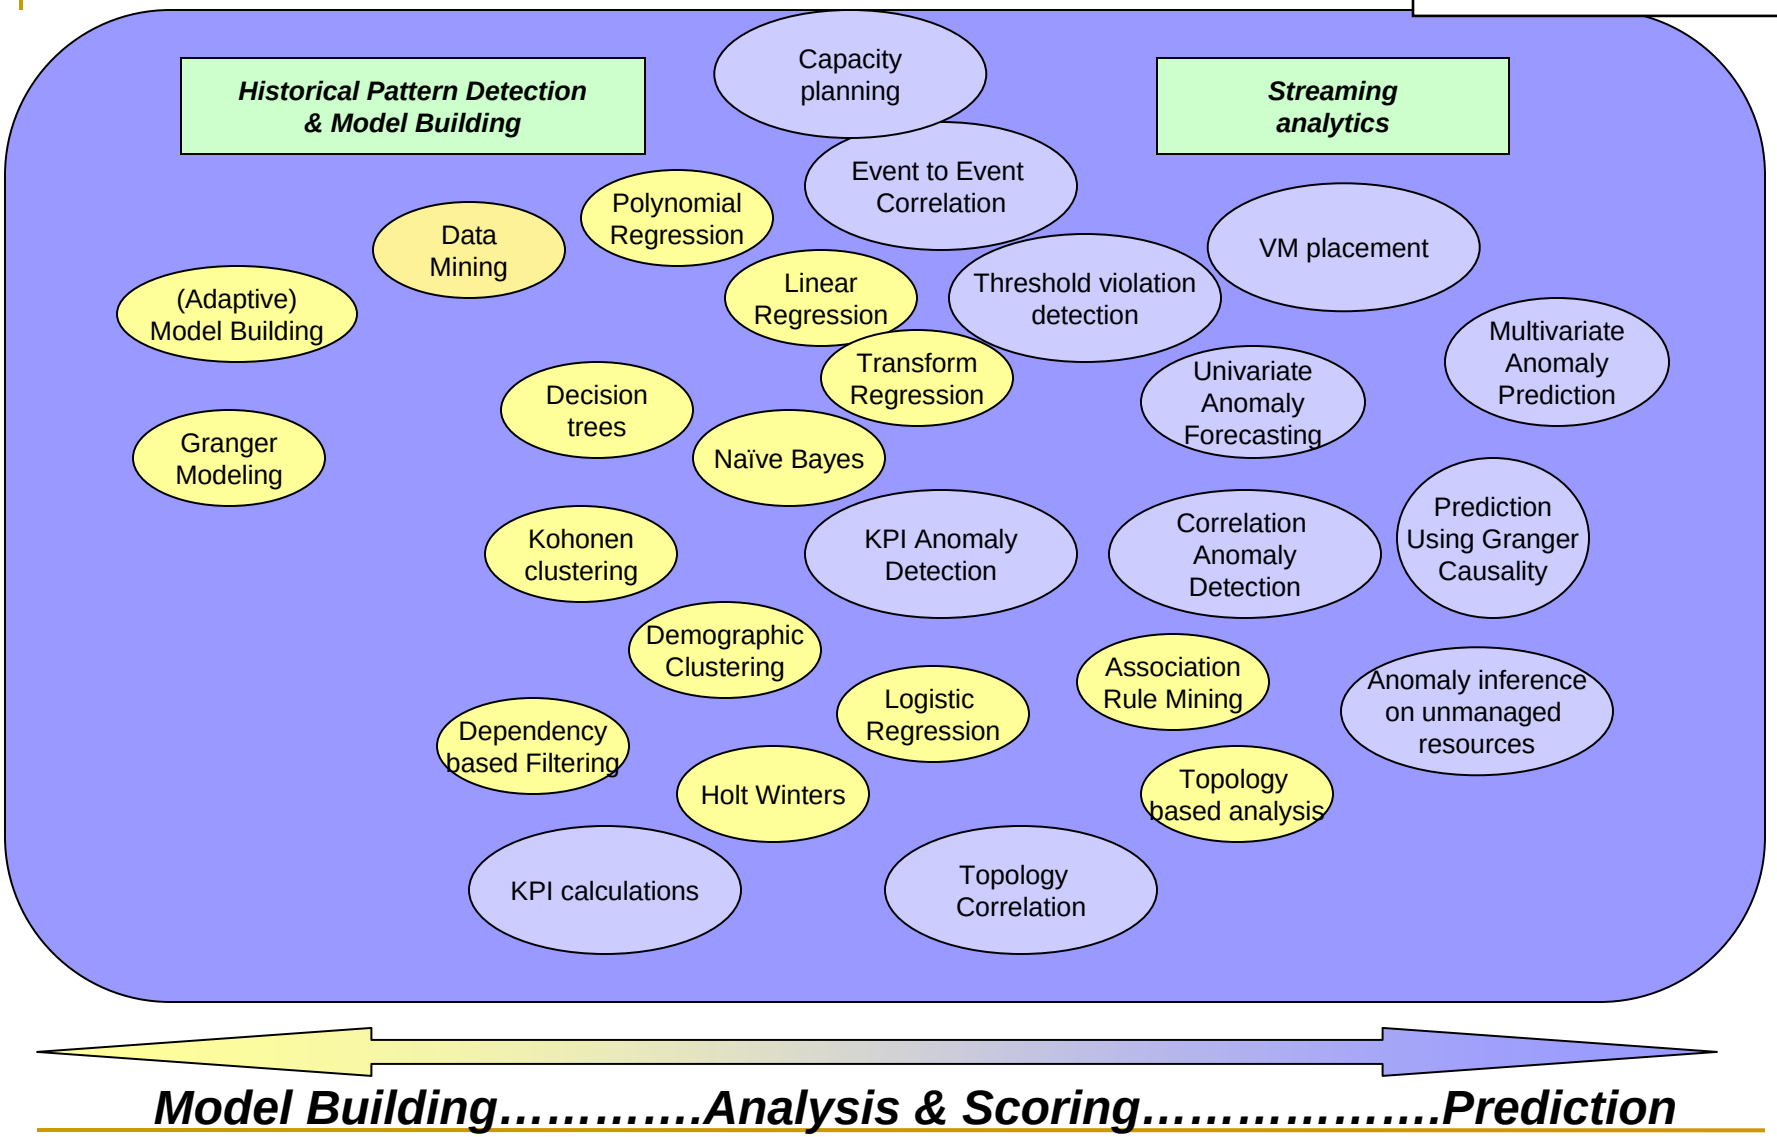

Methodology: Temporal Causal Modeling by Graphical Granger Modeling

- Granger causality
  - First introduced by the Nobel prize winning economist, Clive Granger
- Definition: a time series x is said to “Granger cause” another time series y, if and only if regressing for y in terms of both past values of y and x (1) is statically significantly better than that of regressing in terms of past values of y only (2)

$$y_t \approx A \cdot \vec{y_{t-1}} + B \cdot \vec{x_{t-1}} \quad (1)$$

$$y_t \approx A \cdot \vec{y_{t-1}} \quad (2)$$

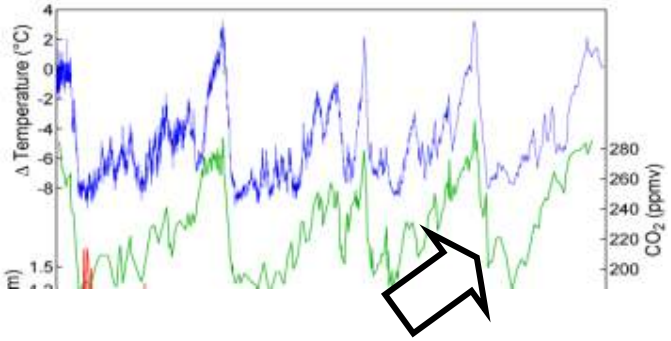

- Combination of Granger Causality and cutting-edge modeling techniques provides efficient and effective methodology for Granger causal modeling of a large number of time series variables

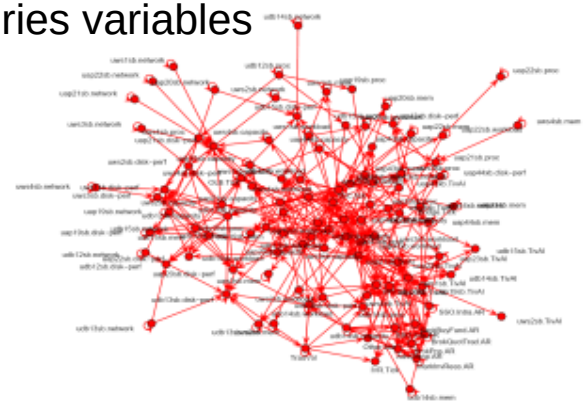

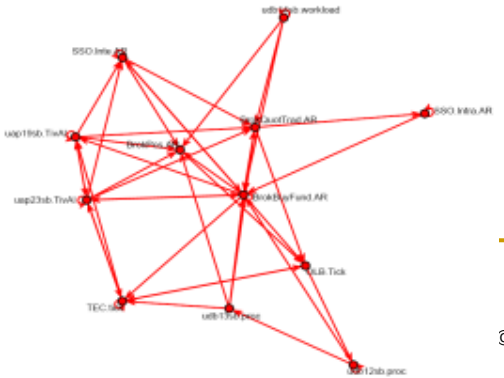

Service Assurance Analytics

3rd Party Event Consoles

SNMP

TIP Visualization

Analytic Algorithm Application

Streaming Analytic Engine

Mediation

InfoVista

ca

bmcsoftware

NetApp

Compuware

hp

Tivoli.

- Leverages IBM Information Management assets to field a state of the art solution
- Highly scalable and resilient streaming analytics engine
- Powerful analytics algorithms, combining multiple approaches, designed to leverage the analytics engine for extensive scalability
- Highly flexible and scalable data mediation layer providing turn key integrations and easily extendable capabilities
- SNMP and Netcool/Omnibus native predictive alerts

Demo Screenshots: Event received

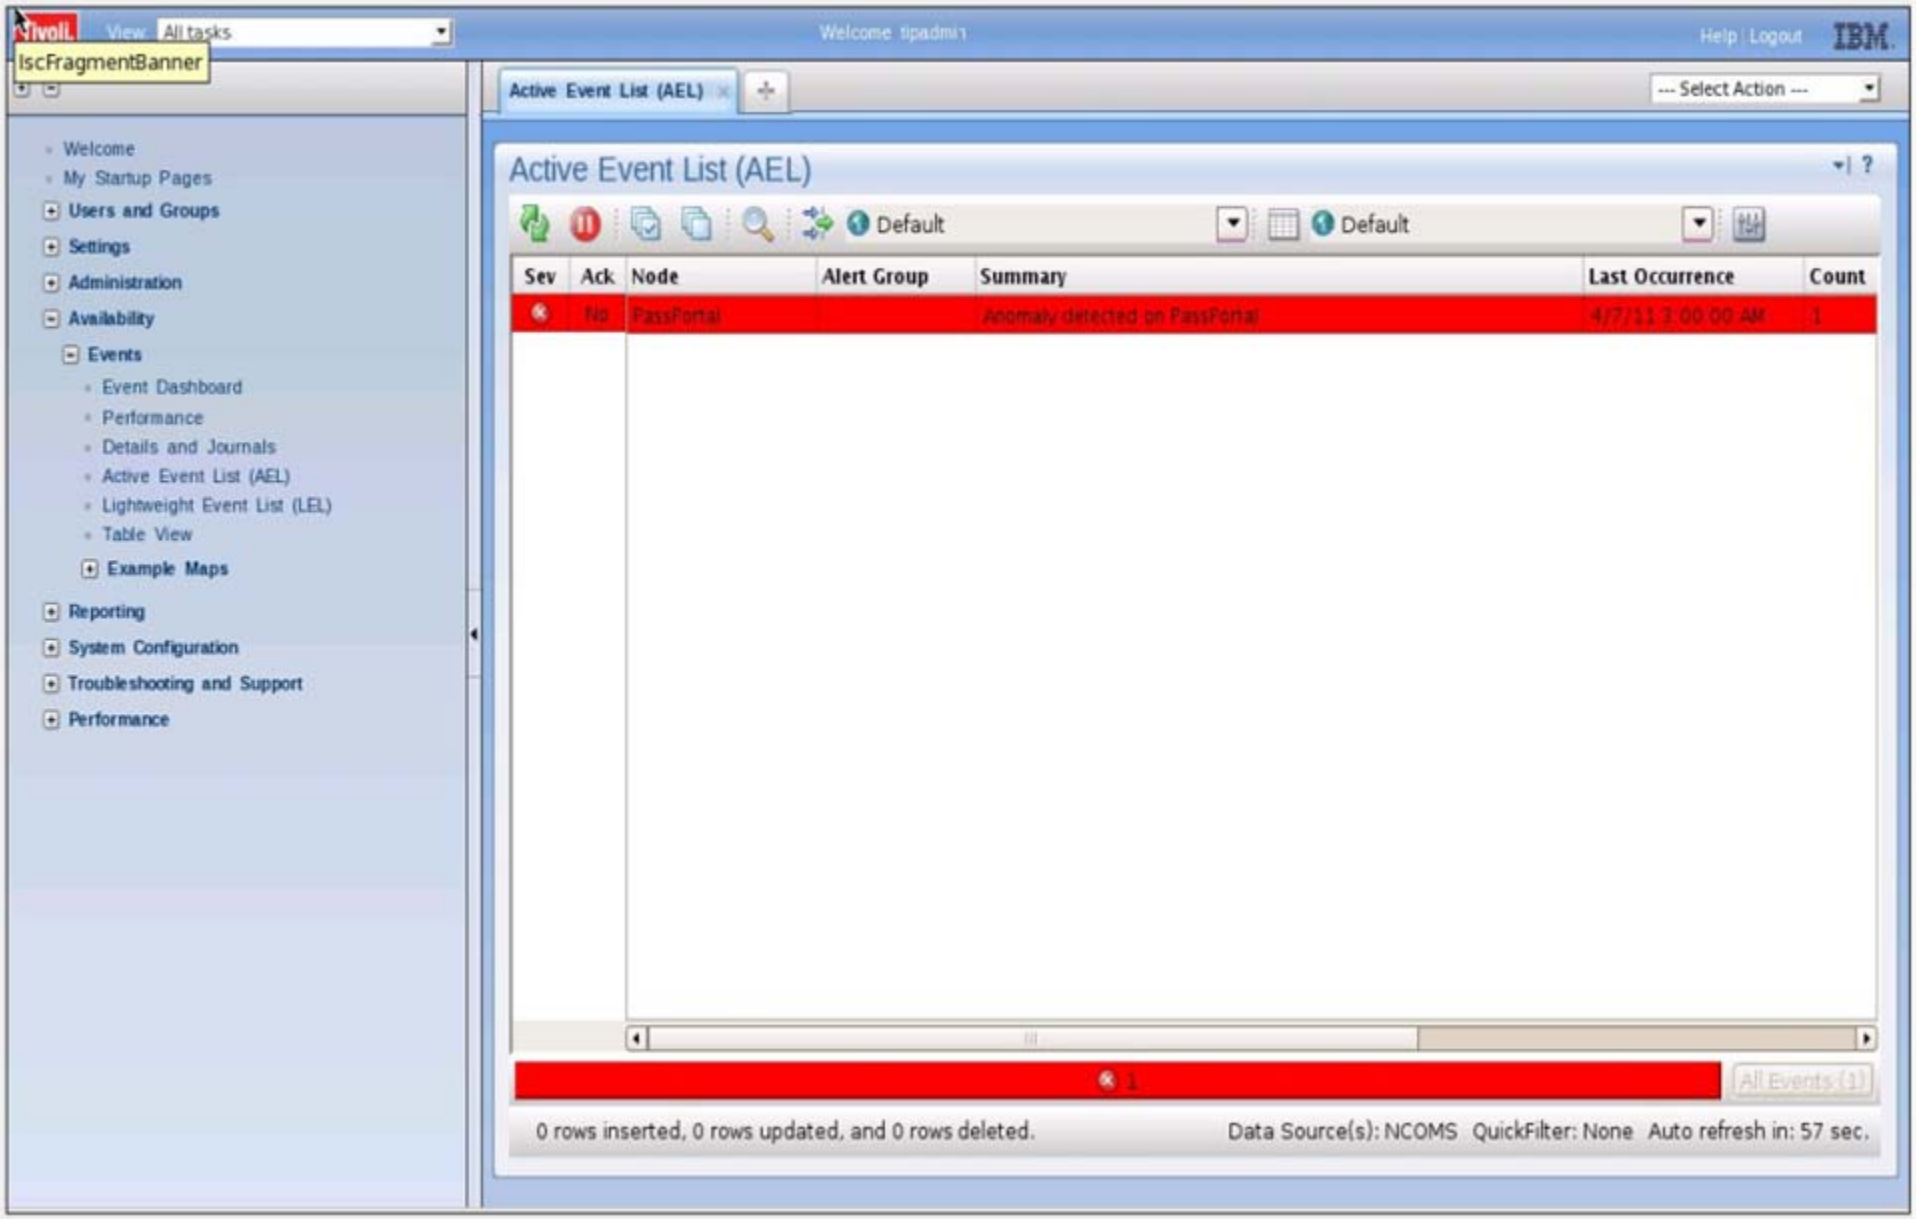

Investigation

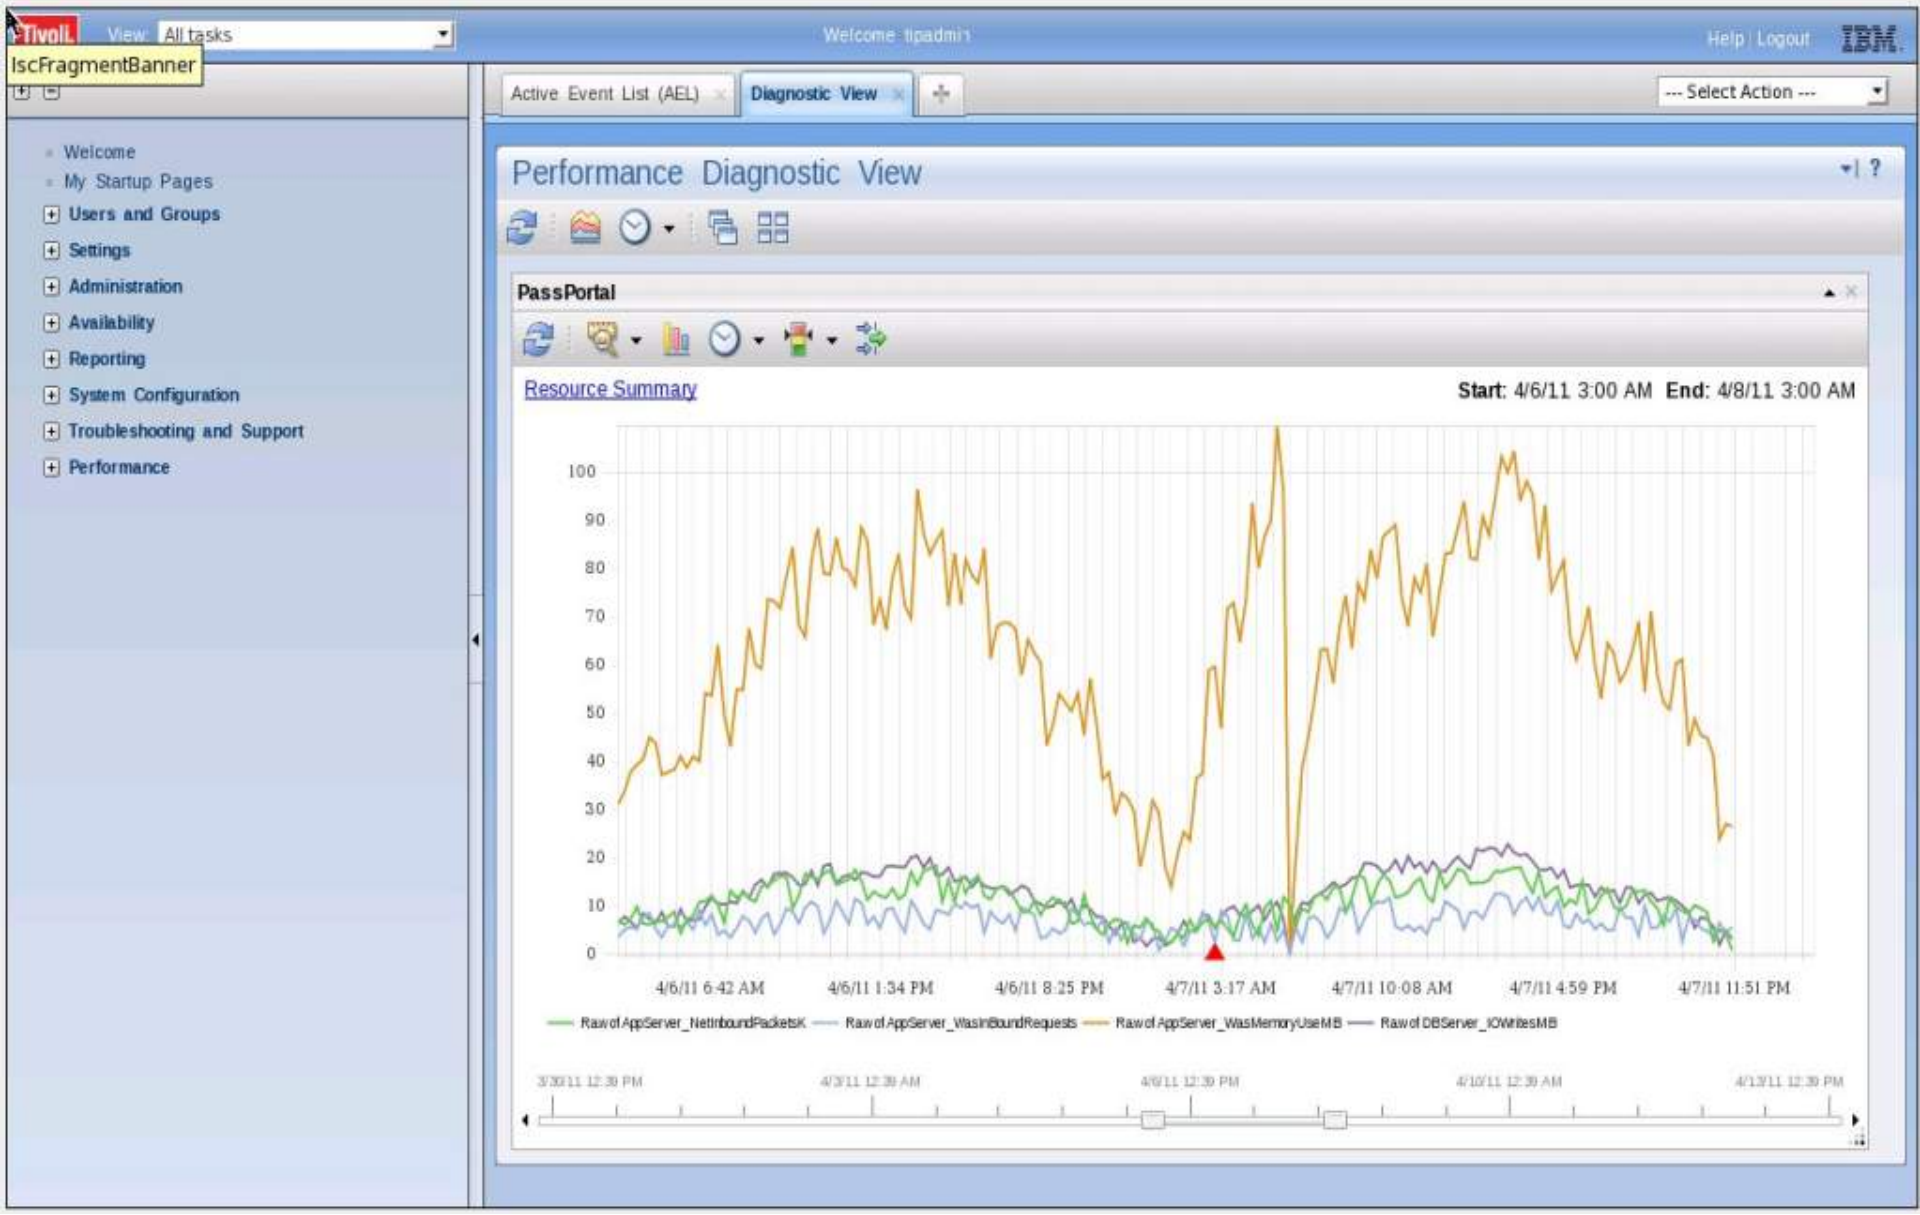

Adding in Pertinent Events

Visualize Predictive Analytics in Business Context

BSM helping further drive IT Operations by aligning Correlated Actual and Predictive States into Business Reprioritized actions

IT Operations now managing from the Predictive and Actual states of Services. Proactive Management

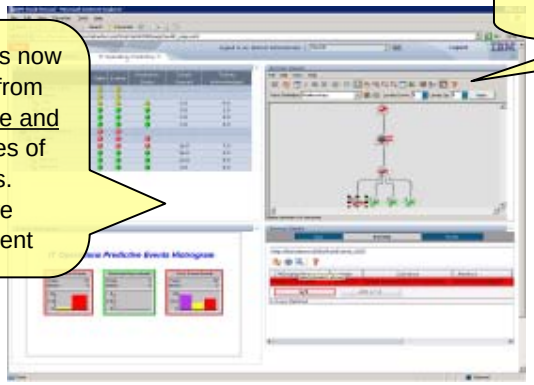

Integrated Analytics Dashboard

Analytics User Interface

Visualize Results (TBSM)

Add Service and Operations Context (Impact)

Insert anomaly events and their associated KPI relationships

Learn normal behaviour and alert on anomalies (TASP)

Historical Data

Real-time Data

Topology, Configuration, Fault, Maintenance

Data

Tivoli's solutions allows you see anomalous conditions priortized for business impact associated with other environmental data, such as faults, configurations changes, maintenance activities, etc...

Service Assurance Analytics

- Identify anomalous KPI behavior without any thresholds.
  - Leverage existing managed data.
- Leverage near real time streaming analytics to identify complex, multi-domain interactions and subtle emerging problems across domains
- Warn users in advance of service impact, deterioration or outage.
  - Learning algorithms which learn normal behavior, with ability to adapt to changes
- Focus on usefulness of results, not on individual algorithms.

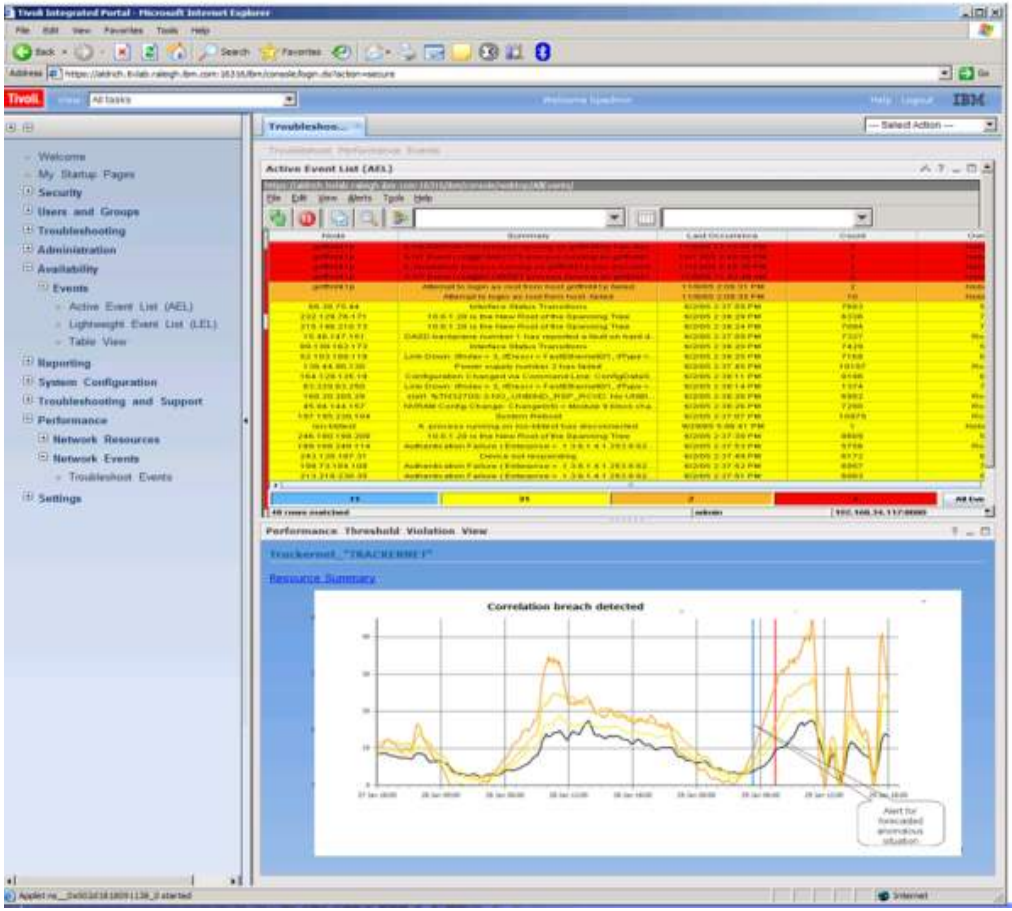

•Add new algorithms over time, without requiring users to become analytics experts.

Acknowledgements and Disclaimers:

Availability. References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS-IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Communities

- **On-line communities, User Groups, Technical Forums, Blogs, Social networks, and more**
  - Find the community that interests you …
    - **Information Management** **ibm.com**/software/data/community
    - **Business Analytics** **ibm.com**/software/analytics/community
    - **Enterprise Content Management** **ibm.com**/software/data/content-management/usernet.html

- **IBM Champions**
  - Recognizing individuals who have made the most outstanding contributions to Information Management, Business Analytics, and Enterprise Content Management communities
    - **ibm.com**/champion