White Paper

Netuitive Technical Overview

Behavior Learning Technology and Predictive Analytics for IT

NETUITIVE®

Introduction – A Tranformational Time for IT Operations

In the early 1980s, the first IT monitoring products came to market and there has hardly been any advancement since – even as infrastructures have grown more complex with the emergence of virtualization and cloud computing. It seems as if IT Operations teams have an intractable problem to solve, or as one analyst put it "the problem has grown beyond the capabilities of your best engineer".

But in a 2010 report on trends in IT Operations, Gartner described one technology – behavior learning – as being "tranformational"[1]. In another report, Gartner states: *"Behavior learning technologies have emerged as a new way to monitor the health of IT infrastructures. They move IT Operations to a more proactive state, where issues can be detected and addressed before affecting the business."*[2]

Netuitive software provides a predictive analytics software that forecasts performance and capacity problems before they impact your users. Netuitive does this by providing analytics that overlay your existing IT monitoring infrastructure, and the technology is powered by our own unique and patented Behavior Learning Engine. This paper will explain in plain language how this sophisticated technology works, and how it compares to other technologies that also promise to provide "predictive" analytics for IT.

The Evolution of Monitoring Tools

The inherent flaw with conventional monitoring tools and even newer analytics products is that they rely on human assumption. These tools often require operators to set static thresholds in order to generate monitoring alerts. Even newer tools that provide rudimentary "dynamic" thresholds still require administrators to manually program and maintain the logic in their performance models. With all of this labor-intensive guesswork it is no wonder why IT organizations are primarily reactive, always fire-fighting and remain at the mercy of their business systems – never knowing where things will break next.

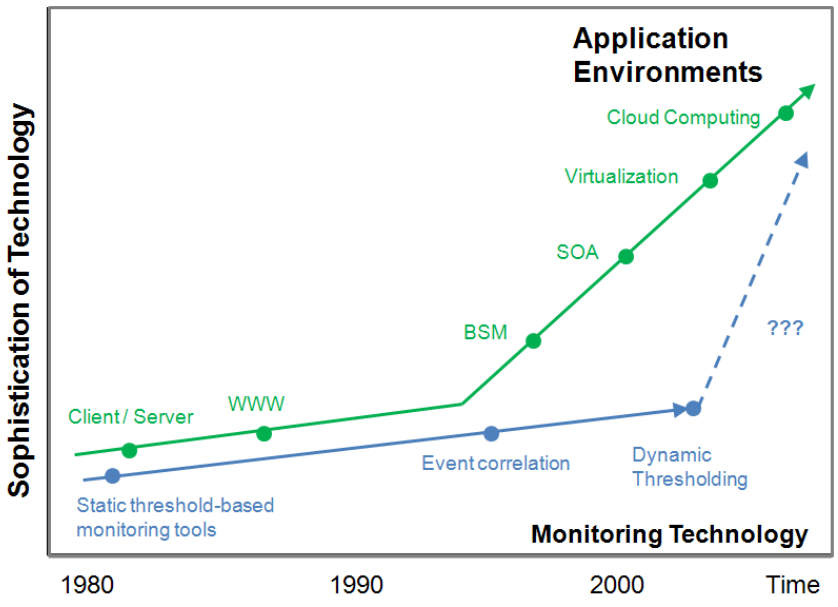

Figure 1. Evolution of Application Environments and Monitoring Technology

While this has mostly been accepted as the industry norm, we are now in the midst of several transformations that make it impossible to stay on the same path while remaining competitive.

- First, there is the move towards end-to-end service management – with a focus on overall application performance, not just availability.
- Second, there is the rapid emergence of virtualization, where infrastructure is dynamically shared and capacity issues quickly impact performance.
- Third, there is the emergence of cloud computing models, where distributed application models and dynamic infrastructure combine to create service performance behaviors that are humanly impossible to model or understand.

The increasing operational demands and complexity created by cross-silo IT service management in increasingly virtualized environments is quickly making the antiquated rules-based tools completely unusable. While some improved products have emerged to address this challenge such as "event correlation," "dynamic thresholding" and "pattern matching" tools, Netuitive has made a leap forward with it's self-learning analytics.

Rather than depending on human guesswork, Netuitive uses a statistical-based approach which automatically analyzes and correlates thousands of system metrics in real-time to learn the normal behavior patterns of a given environment, provide an end-to-end service health dashboard, isolate root-causes and forecast degradations.

Netuitive: Taking Guesswork out of the Equation

Based on nine patented technologies, Netuitive software is the industry's only self-learning analytics solution with the following key attributes:

- Self-learning – does not rely on manual rules, scripts or thresholds.
- Contextual – automatically learns performance dependencies between KPIs and related systems.
- Objective – analytics and alerts are based on objective mathematics, with no human assumption or bias.
- Adaptive – continuously learns and adapts to changes in the environment.
- Agentless – adds intelligence and predictive analytics to already-deployed monitoring tools.
- Cross-Platform – multi-vendor support that works across IT silos with any monitoring source.

Netuitive provides an intelligent analytics layer on top of your existing monitoring solutions. The software leverages existing monitoring agents, collecting raw numeric data at the sub-system metric level for each key performance indicator (KPI), such as CPU and memory utilization, context switching, disk and I/O activity and hundreds of other system and application metrics. Netuitive learns the behavior patterns of each individual KPI for a given day of week, hour of day and minute of the hour. *Unlike any other solution*, it also learns how one KPI behaves in context of the others, which is essential for gaining a holistic picture of system and service health.

Contextual Analysis

Contextual understanding of KPI behaviors through an objective lens is fundamental for effective performance management. As an analogy, a medical doctor determines his patient's health based on a combination of vital signs (key performance indicators) such as blood pressure, heart rate, body temperature and others. Only by observing and analyzing (correlating) all of these conditions together can the doctor accurately diagnose the patient's health and even predict future illness.

Similarly, Netuitive assumes nothing about the "patient." Instead it analyzes and determines outcomes based on its own observation. Like an automated diagnostician, this technology continuously analyzes IT "vital signs," to determine the current and anticipated health of systems and the services being supported.

Netuitive self-learns the interdependencies between each "vital sign." It understands how each KPI influences the other, which can be seen in the software's "Correlation Assistant" interface (See diagram). All of these interdependencies, which are scored on a percentage basis -- from 0.0 to 1.0 -- are self-learned. None of these correlation coefficents are determined through manual means.

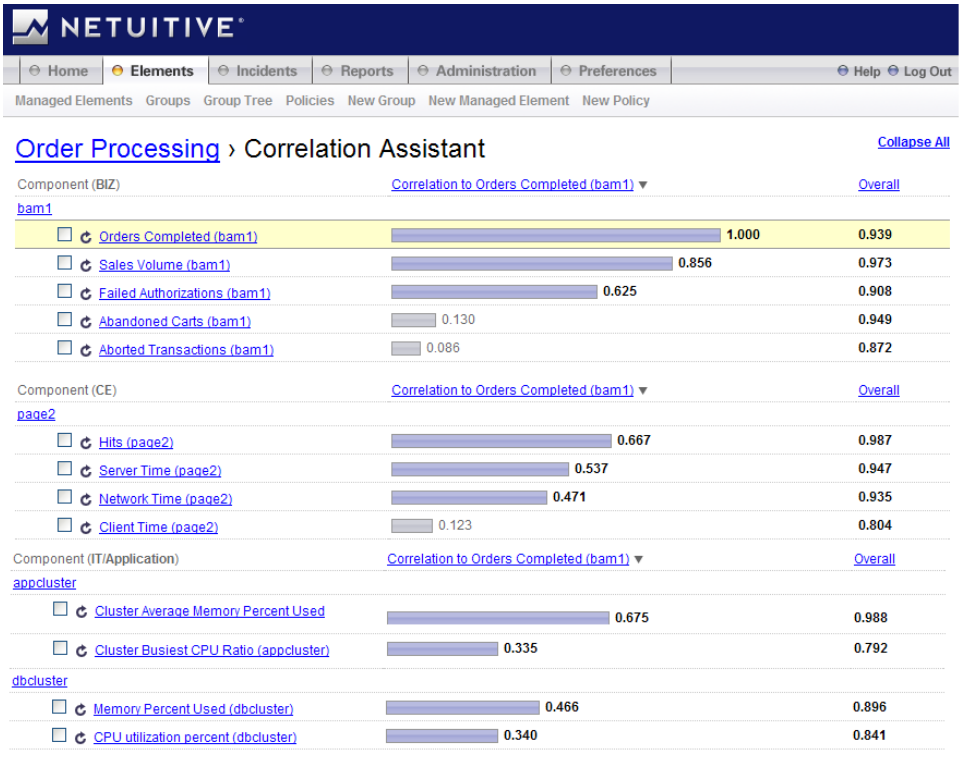

In this example alarm in an Order Processing service, "Orders Completed" – a business metric – has a value of 1.0 since it is of course 100% correlated 100% to itself.

Not only did Netuitive alert administrators of abnormal behavior in "Orders Completed", but it also isolated the problem by discovering a correlation with "Cluster Average Memory Percent Used" (.675) as well as "Server Response Time" (.537). Correlation values greater than 0.3 are statistically significant.

The absence of something like Netuitive's Correlation Assistant should cast doubt on whether a vendor's software does truly *automated* mathematical correlation.

Figure 2. Automated Performance Metric Correlation in Netuitive

To determine how each KPI behaves in context to the others, Netuitive uses multivariate regression analysis and other mathematical techniques to predict outcomes of a given KPI through the observation of other related KPIs. For each KPI, Netuitive calculates its "contextual" performance value in real-time -- using the correlated KPIs of the same time period and their previous history.

For example, in the diagram below, Netuitive learns the contextual real-time behavior "A8-*cxt*" through an algebraic regression calculation using the other real-time KPIs (B8… X8) along with the historically observed KPIs (A1… X7).

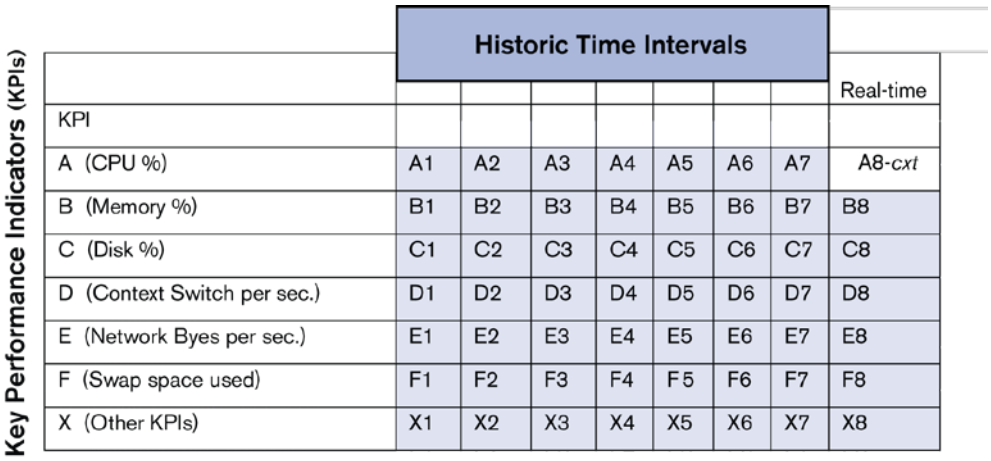

Figure 3. Contextual Analysis and Correlation of Multiple Metrics in Netuitive

Netuitive then compares the *actual* performance of A8 with its *expected* contextual performance, A8-*cxt*. Netuitive applies this same calculation and comparison for each and every other real-time KPI. By comparing the expected contextual behavior with the actual, Netuitive can detect meaningful deviations while disregarding benign anomalies.

Forecasting

In addition to determining real-time contextual performance, multivariate regression enables Netuitive to calculate the forecasted performance of KPIs many hours in advance. All of the historic KPIs (A1… X8), real-time KPIs (B8… X8), along with trending calculations are used to determine *each* forecasted KPI (e.g. A*f*).

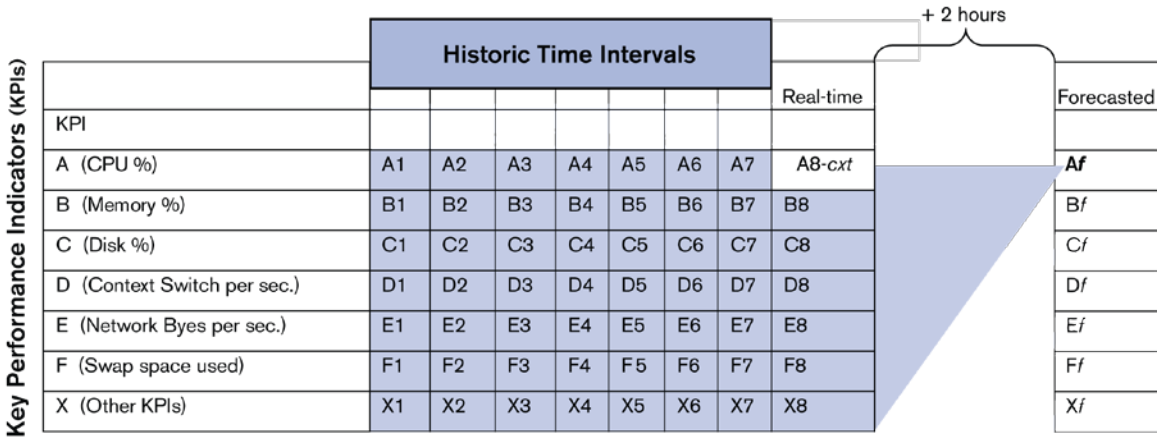

Figure 4. Concurrent Analysis and Forecasting of Multiple Metrics in Netuitive

Adaptive Behavior Profiles

Netuitive's Behavior Learning Engine performs three different types of analysis for each KPI: Real-time, Contextual, Forecast. For each, Netuitive automatically determines a range of acceptable performance for a given point in time. As newer data is updated into the model, the influence of the older data is gradually reduced.

The *real-time* tolerance band compares the actual performance of a given KPI against its historic behavior. This baseline is derived by averaging the performance data for its proceeding time intervals – hourly, daily and weekly time intervals are estimated into the calculation. The diagram at the right represents weekly calculation.

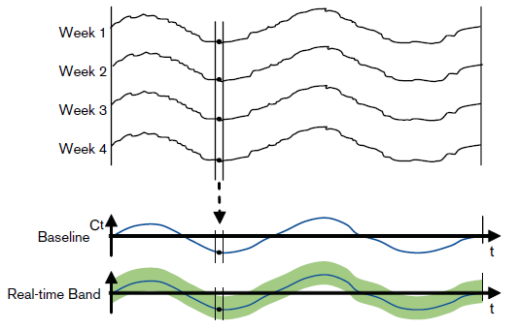

Figure 5. Weekly time-based calculations

The *contextual* and *forecast* tolerance bands are similarly derived, but using the multivariate regression techniques described in the previous section.

To accelerate deployment, Netuitive uses a rapid learning algorithm so that the initial baselines are established within a matter of hours when first deployed. As the sample shows, the contextual bands can provide a much more refined baseline (Figure 6).

Netuitive's statistic-based tolerance bands are dynamic, time-based models that capture rhythms of performance and automatically adapt as environmental changes are detected:

1. *Real-time*: Compares the actual value of a given KPI against its historic behavior.
2. *Contextual*: Compares the actual value of a given KPI against its real-time *expected* behavior, which is calculated from its other interdependent metrics (see page 4).
3. *Forecast:* Anticipates performance two-hours out and is computed using Netuitive's trends-based analysis algorithms (see page 5).

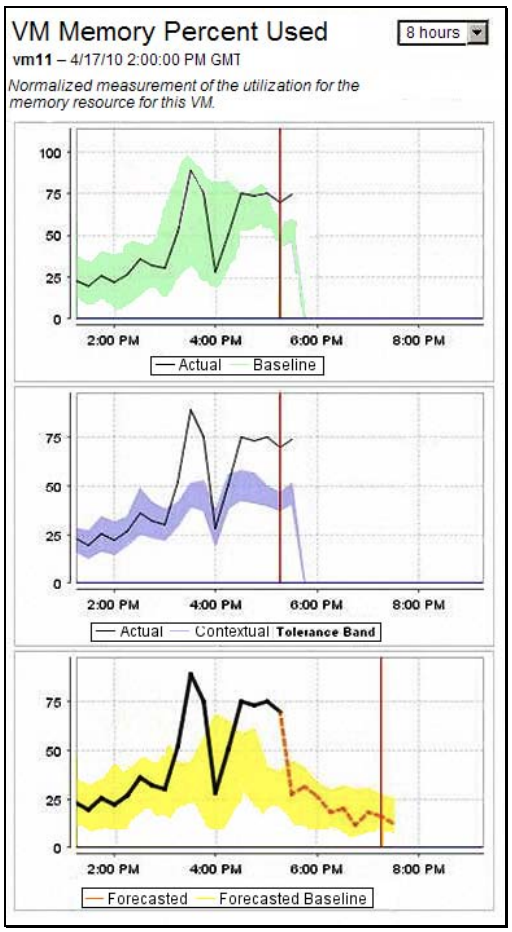

Figure 6. Netuitive's tri-graph display

Each of the described profiles are made up of a tolerance band that incorporates both an upper and lower dynamic threshold based on an adjustable number of standard deviations. These statistical deviations can tuned by the user to adjust for sensitivity.

All three methods – actual, contextual, and forecasted – are calculated simulateneously. By using multiple, simultaneous methods of anayzing each individual KPI, Netuitive delivers unrivaled accuracy for IT systems and service management.

One of a Kind Features: Netuitive Predictive Analytics in Action

Trusted Alarms

Trusted Alarms are one of Netuitive's most valued features and only made possible through Netuitive's automated contextual analysis approach. Delivered on both a real-time and forecasted basis, these alerts provide an easy-to-understand composite view of impending issues, and can be integrated into existing monitoring consoles and trouble ticketing systems. They are generated using an accumulated score of of real-time, contextual and forecasted deviations – taking into account the number, frequency and severity of each deviation. Each system-verified Trusted Alarm results from analyzing dozens of real-time and forecasted calculations. A Trusted Alarm can be generated for an individual system issue or an end-to-end business service. Third-party analyses have shown that, when compared to alerts generated from manual static-thresholds or even rudimentary "dynamic" thresholds, Netuitive Trusted Alarms are not only proactive, but reduce false alerts by 90 to 99%.

Figure 7. Trusted Alarms in Netuitive

It should be noted that while Netuitive's forecasting projects metric baselines two-hours ahead, real customer case studies have shown that Netuitive's Trusted Alarms and predictive analytics are often capable of finding "anomalies" that give operators dozens of hours or even days of advanced notice to impending problems.

Health Index

System and service "health" indicators are only made possible through the use of contextual analysis. Just as a person's overall health can be determined through observation and analysis of multiple vital signs, Netuitive determines health by understanding related KPIs that make up a service delivery ecosystem – business activity, applications, servers, databases, network, etc..

For example, if a person is hyperventilating it might be because they're having a heart attack, not a breathing problem. Or maybe this is a normal response to a vigorous jog. By correlating heart rate to breathing patterns along with other indicators, and analyzed with normal behavior patterns for the given time period, the underlying problem can be accurately diagnosed. Likewise, in IT environments, CPU spikes can be caused by a failing hard drive, a badly behaving application, a sudden flood of users or a virus attack. It could also just be a harmless anomaly. But without contextual analysis there is no easy way to isolate the root-causes.

Netuitive's Behavior Learning Engine self-learns regression weights and correlation coefficients between all the customer experience (latency measurements) and infrastructure metrics automatically and in real-time. In addition, Netuitive generates its health index by also considering the frequency and severity of a given set of anomalies. Health is represented in the software through the Netuitive dashboard.

Features from other vendors (sometimes called "super metrics") are really just a way of saying that users must go through the laborious manual process – for each service – of determining how to develop and maintain models for a single composite service health index.

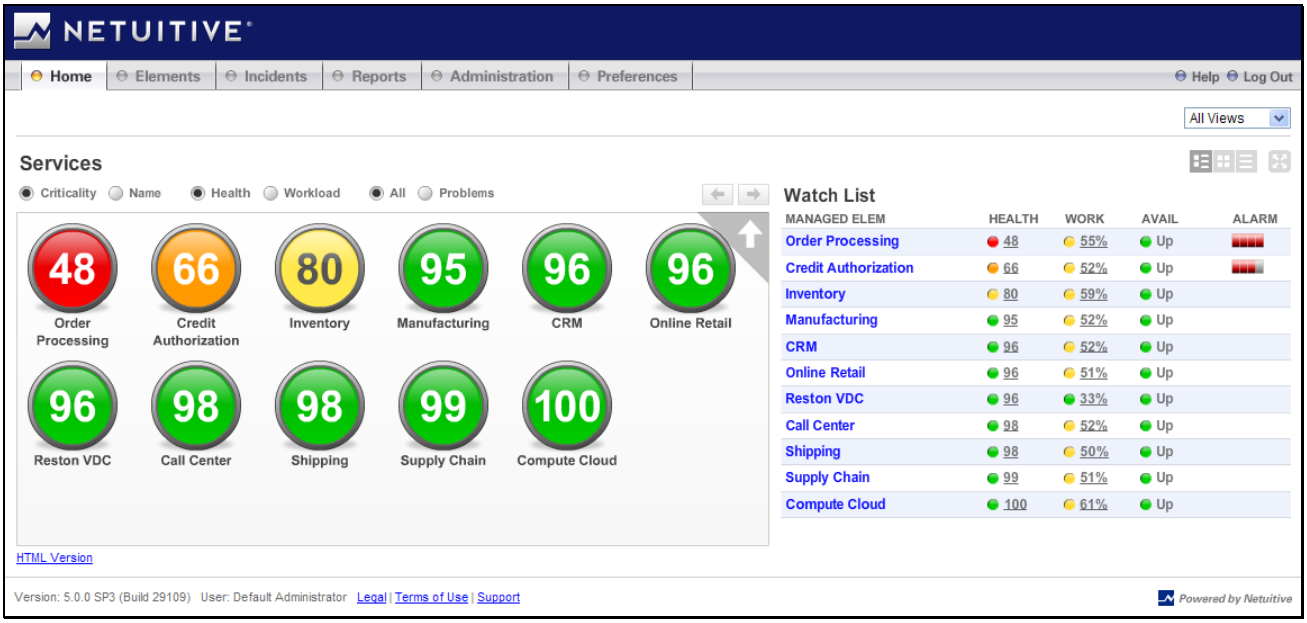

Figure 8. The Netuitive Service Health and Workload Dashboard

Real-Time Workload Index Capacity

The Workload Index is also a powerful derivative from the Netuitive Behavior Learning Engine. It represents a composite index of resource utilization or how “hard” a component or application is working to do its job.

The Workload Index is automatically calculated in real-time, without the need for any manual configuration, using KPIs collected from standard monitoring tools. The Netuitive Workload Index is represented by a value between 0 and 100, and factors in multiple KPIs by resource type. As an example, for an operating system -- resources consist of CPU, memory, network and disk -- where each resource is represented by multiple KPIs. In addition to OS-related workloads, Netuitive uniquely builds workload indexes for any hardware or software component such as the application, storage area network, middleware or server clusters.

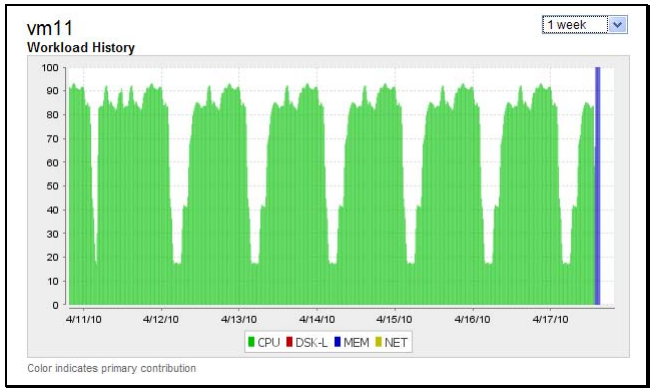

Figure 9.Historical workload profile of a virtual server.

Capacity Management and Forecasting

Having a reliable model for workload and utilization is also extremely important for capacity management. By analyzing historical health and workload, Netuitive’s behavior learning engine is also able to trend and forecast near-term growth in workload and resource consumption as well as the impact on performance (health) for a given server cluster. This provides unique joint insight into the intertwined relationship between capacity and performance that is extremely important in virtualized / cloud environments. One of the key technologies powering Netuitive’s capacity management reporting and trending is the Netuitive Performance Management Database (PMDB).

Figure 10. VM Cluster Forecast

Netuitive Performance Management Database (PMDB) for Search, Reporting and Data Export

One of the key challenges to implementing a solution for capacity management is getting access to the right data in timely fashion. Complex and dynamic infrastructures (e.g. virtualization) means that planning horizons are shorter and data needs to come from multiple sources in real time.

This essentially defines the Netuitive Performance Management Database (PMDB) – which is the foundation for Netuitive's self-learning analytics, search, and reporting capabilities. The Netuitive PMDB integrates real-time performance data you are already collecting from tools such as BMC, CA, IBM, HP, Microsoft, Oracle, NetApp, and VMware into the richest single source of IT performance data available anywhere.

While designing and building a PMDB may seem straightforward, Netuitive users can save themselves a lot of work because the Netuitive PMDB already supports the following capabilities:

- Search millions of data points to quickly evaluate and optimize system performance
- Collect raw performance data of 10-1000 KPIs per system (physical or virtual)
- Integrate and normalize data from any monitoring source
- Very low storage requirements using Netuitive's unique retention profiles
- Support metadata to describe system attributes
- Integrate with CMDBs
- Export of data to 3rd party tools for planning, analysis or reporting; available data includes
  - Metric baselines
  - Correlation coefficients
  - Statistical regression weights
  - Trends
  - Events
  - Attributes
  - Relationships, etc.

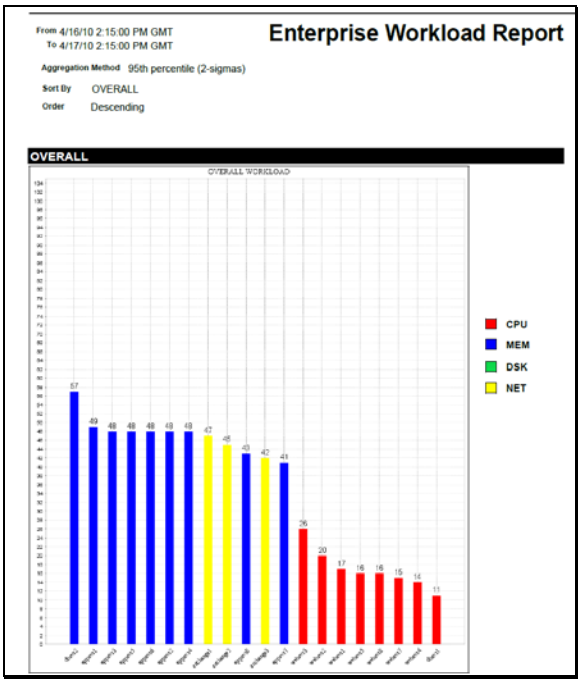

Figure 11. Enterprise Workload Report

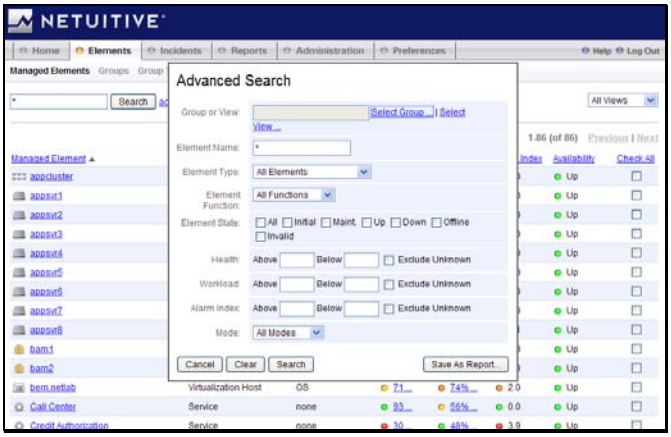

Figure 12. Advanced Search Interface to the PMDB

Intellectual Property

Just as important to the technology's unique self-learning attributes, Netuitive has developed a breakthrough by correlating tens-of-thousands of KPIs simultaneously using off-the-shelf hardware – massive calculations that generally require cost-probibitive mainframe computing. Resulting from nearly 20 years of academic and commercial research, Netuitive has nine issued patents for its intellectual property, including:

- *Concurrent Learning and Performance Information Processing System*
- *Multi-Kernel Neural Network Concurrent Learning, Monitoring & Forecasting System*
- *Automatic Data Extraction, Error Correction and Forecasting System*
- *Enhanced Computer Performance Forecasting System*
- *Engine Analyzer and Configuration Utility used in a Computer Performance Forecasting System / Automated Analyzers for Estimation Systems*
- *Computer Performance Estimation System Configured to Take Expected Events into Consideration*
- *Method and system for analyzing and predicting the performance of computer network using time series measurements*
- *Methods and System for Self-Learning Performance Monitoring and Decision Support*

Comparing Netuitive Self-Learning Analytics to Other Approaches

Beyond the well-known threshold-based monitoring tools – such as BMC Patrol, HP OpenView, Tivoli, etc. -- a newer generation of analysis products have come to market in recent years which do provide some improvement over the conventional approach, though still fall far short of the ultimate requirement. These tools use such analytics techniques as event correlation, single-metric baselining (dynamic thresholding) and pattern matching (fingerprinting). Some also claim to do automated metric correlation and "intelligent" alarms similar to Netuitive. So let's take a closer look under the covers.

Event Correlation: Many users have investigated or tried event correlation products.. The inherent flaw with event correlation techniques is that they rely on the threshold-based events coming from conventional monitoring tools. Because these thresholds were set according to the guesswork of operators, the events being correlated are inaccurate to begin with. It is not uncommon for upwards of 90% of events coming from such tools to be false-positives. Therefore, the accuracy of event correlation analysis is severely compromised because of the data it relies on. Furthermore, event correlation products require their own manual rules, which can be tedious to program and maintain.

Single-Metric Baselining or Dynamic Thresholding: Single-metric baselining or dynamic thresholding is becoming more common and is often included for free in tools from companies like HP, IBM, BMC, CA and Microsoft to improve alarming and event correlation.

This type of dynamic thresholding involves the baselining of one KPI at a time – in isolation – to generate a tolerance band based on a standard deviation from the observed historical behavior. And while this approach provides some value by automating thresholding for each KPI, is has some limitations. First, when a metric has a significant amount of historical variation, the "dynamic" thresholds tend to "flatten" and end up being no better than fixed thresholds (see Figure 13). In fact, this can generate even *more* alerting noise than conservatively set manual thresholds.

Metric Correlation:. Second, when single metric behavior is analyzed in isolation, there is no contextual analysis or "metric correlation" (as other vendors may call it). Any correlation between KPIs requires the operator to determine the associations manually. The reality is that when other solutions claim to automatically "correlate" metrics, what they are doing is manually grouping metrics (e.g. OS) or creating a composite index or "super metric". This approach can hardly be called "automated" when compared to Netuitive's contextual analysis, since manual maintenance of correlation rules and logic is required to accurately depict system or service health or accelerate root cause isolation.

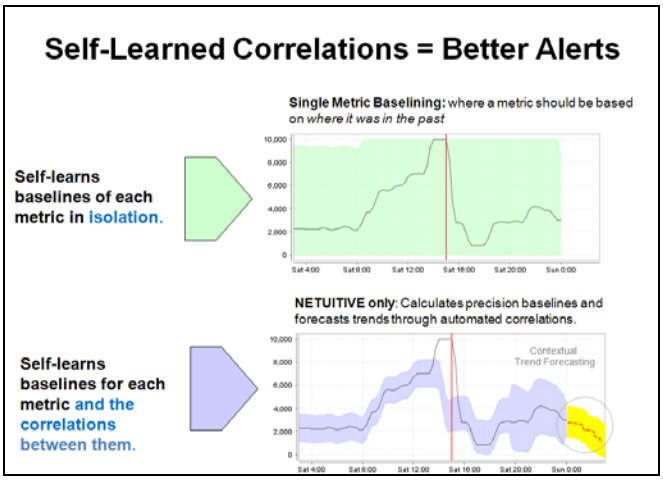

Figure 13. Single Metric Baseline vs. Contextual Baseline

"Smart" or "Predictive" Alerts: Netuitive's sophisticated baselining approach (real-time, contextual, forecasted) is behind Netuitive's accurate and predictive Trusted Alarms. Other vendors also claim similar alerting accuracy, but their approach is much more simplistic. In one solution, for example, the number of deviations is counted for grouped metrics or the "super metric". When the number of deviations reaches a set threshold in a fixed amount of time, a "smart" alert is issued. You want "predictive" alerts? Tell the software to alert you to the first deviation of any in the grouping.

Since there is human bias in how the metrics are grouped and in how many deviations constitute a significant incident, this approach seems neither "smart" nor "predictive". In fact, part of the reason some vendors do this is simply to suppress the large number of alarms generated by their simplistic or rules-based thresholding approaches.

Pattern Matching (fingerprinting): Pattern matching or "fingerprinting" tools use a snapshot of a performance problem that were the warning indicators for a given issue. This performance pattern -- or "fingerprint" -- is then captured and stored to identify the same problem in the future, if it happens to recur. Pragmatically speaking, this techonology is really an alternative approach to traditional event correlation. These tools can work well for a standalone packaged application, but they have severe limitations.

- Vendors sometimes pre-populate known patterns into a "code book." If you look at the problem from a crime-fighting analogy: how can you catch a crook if his fingerprint is not in the database? (i.e. how will code books recognize yet-to-be-seen problems?)
- Fingerprinting can also prove impractical for monitoring across system silos for virtualized environments or cloud-based applications. In these dynamic environments, there is simply no way to effectively capture the thousands of unique fingerprints that represent a complex web of applications, servers and network devices. The number of permutations are limitless and incapable of scaling for such an environment.

Because of these technology limitations, vendors who use "fingerprinting" or other approaches have had a tough time proving their value in large enterprise environments. They may even "bury" the technology under the general disguise of metric correlation and behavior learning. But as outlined in this paper, predictive IT analytics software based on true "behavior learning" technology requires automated contextual analysis and multivariate regression techniques. These solutions should not rely on human assumption, but on unbiased observation. By this standard, and the impressive list of large enterprises that have successfully deployed the software, Netuitive stands on its own.

Conclusion: The Value of Predictive Analyticts and Behavior-Learning Technology

While today's businesses rely on complex 21st Century applications, the tools to manage them still use technology almost unchanged from two decades ago. Even as these tools have provided incremental improvements and updated user interfaces, the reality is that the complexity of the IT environment has grown beyond the capability of the manual approaches they still rely on.

It is no wonder that every year, analyst polls still find that most IT shops first learn about incidents when users call their help desks. Essentially, they have no reliable end-to-end visibility into their application infrastructure, let alone the ability to forecast problems before users notice them. Now with the growing adoption of virtualization, and the migration of applications to cloud environments, the complexity is accelerating beyond the breaking point. And the volume of data is greater than users can possibly analyze and act on.

"The problem has grown beyond the ability of your best engineer."

In an attempt to address this challenge, a second generation of "predictive" analytics tools have come along. And while some of these claim "behavior-learning" capabilities, their reliance on manual event correlation, single-metric dynamic thresholding and pattern matching simply does not live up to the assertion.

Predictive analytics powered by behavior learning *is the new essential technology* to meet the challenges in the next decade for IT performance management in data centers. In choosing such a solution, consider that Netuitive is the pioneer and industry leader in behavior learning technology – a game-changing approach to performance management that Gartner has called "transformational".

IT executives at 7 of the top ten banks, the two largest telcos and hundreds of enterprises around the world recognize and trust Netuitive as the only self-learning, predictive analytics platform for managing the performance of their physical, virtual, and cloud environments. Only Netuitive works across multi-vendor infrastructures to deliver end-to-end visibility, forecast performance issues, and speed problem resolution. Unlike policy or rules-based management solutions that rely on human guesswork, Netuitive behavior learning technology applies patented advanced mathematics and statistical analysis to automatically self-learn an environment's normal operating rhythms and continuously adapt to changing business conditions.

References

1. Govekar, Milind; "Hype Cycle for IT Operations Management, 2010"; 12 July 2010; Gartner Research ID Number: G00201203
2. Williams, David; "An Introduction to IT Operations Behavior Learning Tools"; 16 December 2009; Gartner Research ID Number: G00172843

Netuitive

Predictive Analytics for IT

Netuitive, Inc.
12700 Sunrise Valley Drive
Reston, VA USA 20191-5804
Toll Free: 877.492.9672
Main: +1.703.464.1500

Netuitive Europe Ltd
Chineham Business Park
Crockford Lane
Chineham, Basingstroke
RG24 8AL United Kingdom
Main: +44 (0)844.546.5001

Web: www.netuitive.com

Trust your performance to Netuitive.

Netuitive provides predictive analytics software for IT. Netuitive replaces human guesswork with automated mathematics and analysis to forecast, identify and resolve IT performance issues before they impact quality of service. Hundreds of customers, including seven of the 10 largest banks, rely on Netuitive to proactively manage the performance and capacity of their IT infrastructures -- physical, virtual and cloud. Industry recognition includes the 2010 CODiE Award for "Best Systems Management Solution," the 2010 EMA Award for "Best Analytics," and the 2009 "Best of VMworld" Award for Virtualization Management.

Proactive. Predictive. Preventative.

Contact Netuitive to learn more about self-learning performance management solutions for IT infrastructure and services, or to schedule an on-site demonstration.