White Paper

Self-Learning Performance Management Comes of Age

Technology offers relief for IT Operations staff drowning in useless data

October 2008

NETUITIVE®

Corporate IT Struggles to Serve Business Owners While Managing Chaos

In recent years, enterprise IT organizations have been under pressure to be more aligned with corporate and business objectives. They have responded by implementing new business-oriented operational philosophies–with names like Business Service Management (BSM), IT Service Management (ITSM), and ITIL (IT Infrastructure Library) V3. To support this shift, analyst firms like Forrester have called for IT operations staff to evolve from monitoring not just service component availability, but business service performance–described as "understanding the health and availability of a particular service (application) to meet the expectations of the business."[1]

With the transition from availability management to business service management you would expect the major legacy IT management vendors to be well positioned to meet the evolving needs of users. For three years running, however, survey results suggests otherwise. At the Gartner IT Operations Summit in 2008, a poll of attendees revealed that over 70% gave their IT system management vendors a "C" or worse[2] indicating a high degree of dissatisfaction. The results showed no change from similar Gartner survey results for the previous two years.[3]

Even as IT organizations are trying to become more business-aligned, the IT landscape is changing at an incredible pace. Application architectures are growing in complexity to incorporate service-oriented, highly distributed n-tier architectures. Adding fuel to the fire are the latest disruptive "hot" technology trends such as Web 2.0 and x86 server and desktop virtualization.

This torrid pace of change has caused the legacy systems management vendors to constantly adapt to what infrastructure to manage, with little room for true innovation on how to better manage infrastructure. As a result, many enterprises today are trying to manage their 2008 IT operations challenges using tools with decades-old development roots[2].

Where Traditional Performance Management Fails

The high level of dissatisfaction with traditional monitoring tools shows up as "ease of use" issues in Gartner surveys.[4] A deeper analysis reveals a lack of usable functionality to address user needs in today's chaotic data center environment[2]

- **Ineffective monitoring thresholds** - fixed or simplistic approaches result in inaccurate alerting and require labor-intensive maintenance
- **Event storms (false alerts) and lack of relevant diagnostics** - make root cause isolation nearly impossible.
- **Reactive monitoring vs. proactive analytics** - focus on Mean Time to Repair (MTTR) and not proactive response when service performance begins to degrade.
- **No end-to-end service monitoring** - makes problem diagnosis and communications difficult across traditional IT silos and provides no visibility to business application owners.

In the sections below, we will explore these challenges in more detail and hear what leading analysts are saying about an emerging technology that can help.

Ineffective Monitoring Thresholds

Legacy monitoring tools are dependent on fixed monitoring thresholds for performance metrics to alert users of potential performance problems. This approach is not only labor intensive - it can be extremely inaccurate.

Matthew Mayer, a veteran systems administrator at financial services giant Wells Fargo, considers the IT performance management process a shocking productivity sinkhole. "You're never ever done doing threshold settings, which makes engineers miserable", he laments, referring to the tasks of monitoring Wells Fargo's nearly 20,000 servers.

What Mayer is referring to is the traditional paradigm that has existed for the past 10 to 15 years for IT performance management simply cannot adapt to today's dynamic IT environment. This paradigm involves manually crafted performance models for servers and other components that include the setting of thresholds and the creation of rules to attempt to detect when IT components deviate from "normal" behavior.

"How could we set good thresholds when 'normal' baselines determined yesterday are not going to be good for tomorrow?" states Carl Truitt, lead IT monitoring engineer at the National Institutes of Health (NIH). "We're constantly guessing."

Some vendors of the legacy tools have responded by adding limited automated thresholding capabilities to their products. However, users are finding that these rudimentary single-variable approaches fall woefully short of being usable. The legacy tools still require user configuration; they do not offer a full view of performance by correlating multiple metrics; and they can cause performance problems in the monitoring tools themselves due to high computational overhead.

Nevertheless, an effective solution to threshold monitoring is crucial. Without it, operations staff can be faced with "event storms" - or a flood of alerts and irrelevant data.

Event Storms (False Alerts) and Lack of Relevant Diagnostics

False positive alerts occur when a threshold is met or exceeded, but there is no real problem. A typical operations group can receive thousands of false positives daily–representing more than 90 percent of all alerts received. Often administrators and engineers unknowingly end up working on the same event, or get contacted after hours to check events that turn out to be normal deviations.

"When we didn't have good performance data to go on, we just acted on IT action requests submitted by end-users," relates Glen Huey, server administrator from Teledyne. "We didn't have anything to compare these events against… we had no way of knowing if the event was really abnormal, or if it was an indication of something harmless that happened periodically."

When alerts do appear they generally lack context. There is no way to determine if the alert is an indication of the symptom or the underlying problem. To use an analogy, if a person is hyperventilating it might be because they're having a heart attack, not a breathing problem. If one could correlate their heart rate to their breathing pattern, there would be a much higher likelihood for diagnosing the patient. In IT environments, for example, 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 context–or correlation–there is no good way to figure this out.

Virtualization adds an even greater degree of complexity and difficulty in problem diagnosis. "Root-cause analysis is important because of the growing number of interdependencies–many of which are fluid or mobile because of the underlying virtualization technology making the construction of potential fault paths more problematic," explains Cameron Haight from Gartner. "Companies that are doing some interesting things here include Netuitive."[5]

Reactive Monitoring vs. Proactive Analytics

To avoid a flood of false positives, overwhelmed staff will overcompensate by setting thresholds too high, which results in real alerts being missed. Without solid performance data, IT organizations remain in a pure reactionary mode. They depend on user complaints, rather than their performance monitoring system.

"The vast majority of my day was being spent going through each alert to figure out which ones needed to be addressed and which should be ignored," recalls Truitt from NIH. False positives can set off a stampede of operators, administrators, engineers and help-desk personnel, all chasing after a problem that doesn't exist.

The inability to provide proactive service can have serious business impact. Jim Hirschauer, Technology Architecture Manager for Wachovia's Corporate and Investment Banking Group, recalls one poignant example of this. "It was a foreign trading application, and we were having intermittent performance issues, which was causing some significant problems for some of our customers," he says. "They started complaining about delays, and with the speed of that market, delays aren't good." The problem, Hirschauer says, was that Wachovia's battery of systems-management and monitoring tools reacted too slowly[6].

The capability to visualize and be alerted to negative forecasted trends in performance metrics can help users be proactive and address service problems before users are impacted. "If you have a performance impact on an important system, you're doing a disservice to the business if you're not keeping up with it and keeping it from happening," Hirschauer says. "Customers will just pick up and move elsewhere."[6]

Incident Timeline: Netuitive Proactively Forecasts 2+ Hours Before Outage Occurs

The timeline and screen shot below shows an actual real-life example of how statistical analytics helped alert systems administrators to a performance problem 2 hours before service was impacted. Administrators initiated failover processes which prevented any services from being impacted despite the eventual server crash.

2:15 am *Netuitive alerts NOC on disk issue. NOC which investigates alert details: self-learned thresholds had been exceeded for Physical Disk Utilization and IO Wait Percent. Legacy monitoring tool is showing no alarms at this point. NOC continues to investigate.*

2:45 am *NOC ran an error report and saw many i/o errors. NOC opens support ticket, problem escalated to Sys Admin.*

3:30 am *Sys Admin recommends contacting hardware vendor.*

4:20 am *Application support reports they cannot connect to the applications on this system. Attempts application failover.*

4:35 am *Server crashes and legacy monitoring tool finally begins generating filesystem alarms.*

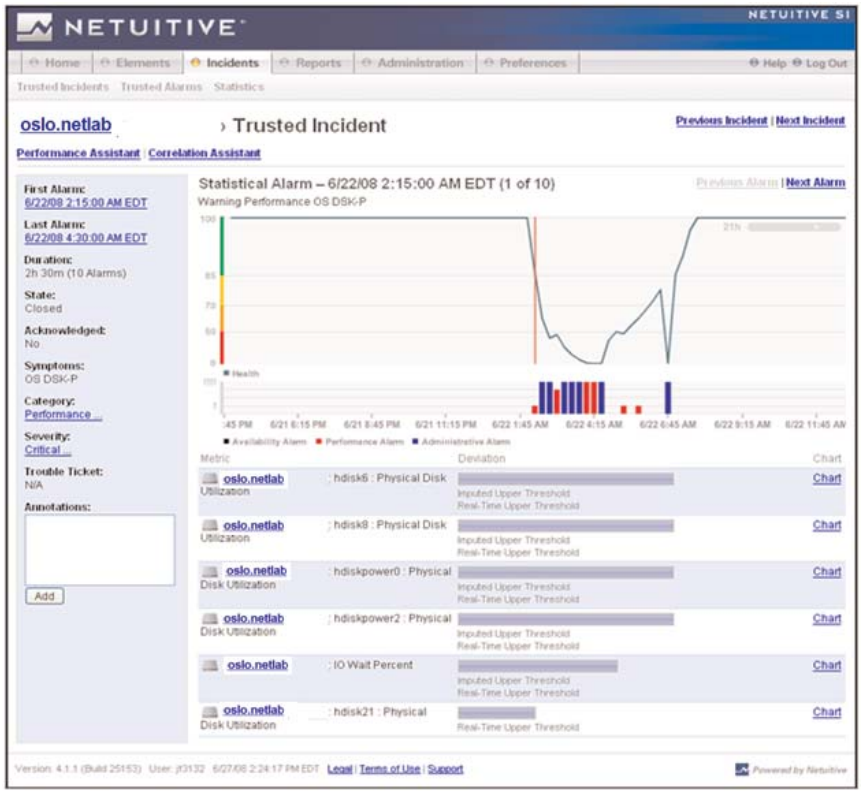

End-to-End Service Monitoring

Business services and mission critical applications have many stakeholders - from the application owners who report to line-of-business executives, to the varied IT Operations staff belonging to different technology silos. This includes server, database, and network administrators who monitor the infrastructure, and software engineering staff who manage the applications themselves.

To manage the complexity of true end-to-end monitoring of a business service, Gartner suggests that companies "consolidate results from multiple availability and performance monitoring tools into a single event console and business service view."[2] This dashboard view has many consumers, the most important of which is the relatively new role of IT Service Manager or Application Manager–the person responsible for the complete "health" and performance of an application or business service from multiple perspectives.

Seeing multiple IT performance metrics side-by-side with other business related metrics is extremely important to be able to proactively manage the overall health of an application.

- If a performance metric such as application response time or "orders per hour" begins to degrade or deviate from the norm, it is important to be able to analyze the IT infrastructure to determine if it is the cause of the degradation.
- Conversely, if IT metrics such as database record locks or I/O wait time begin to degrade–indicating an impending application problem–application owners would want know if the problem has impacted users or the business yet.

This type of application dashboard is often referred to by Forrester as a "bottoms" up or performance based approach to Business Service Management that can yield quick results to bridge the business-IT gap and provide measurable value of IT's contribution to a business service[1].

The Emergence of Self-Learning Performance Management

Given the dynamics of today's IT environments, leading analysts have come up with an IT Operations "wish list" of sorts which outlines the needs for performance management in today's data centers.[2]

- Supplement traditional fault management tools with proactive performance management and predictive learning tools.
- Reduce complexity, effort and costs. Consolidate results from multiple availability and performance monitoring tools into a single event console and business service view.

So is there a technology solution to the IT performance management chaos that seems to persist in data centers today? The answer lies in the marriage of sophisticated mathematics and readily available, cheap computing power.

Both Gartner and Forrester have recognized this emerging technology (which they categorize as "Behavior Learning" technologies) as absolutely crucial to help automate performance management operations in data centers. In fact, after a period of gestation in the traditional technology "hype curve", Gartner feels that this technology is now ready for prime time.

Netuitive calls this technology "self-learning performance management", and it is at the core of its Netuitive SI and Netuitive Service Analyzer products.

The Technology Behind Self-Learning Performance Management

At its core, self-learning technology is "math on steroids", using sophisticated mathematical algorithms to analyze data from dozens of metrics and automatically calculate how they behave in relation to one another over time.

This underlying technology enables software to:

- Self-learn and adapt to the normal, rhythmic behavior of IT components and business services over time.
- Automatically determine how the behavior of one metric impacts another, based on the unbiased mathematical certainty of statistical correlation.
- Forecast where the metrics may trend in the near future and how service may be impacted.

All of these capabilities allow the software to get an accurate and objective picture of whether or not the component is behaving "normally" and to alarm when behavior deviates from normal.

What has made this technology more feasible in recent years is the marriage of complex mathematical algorithms and cheap computing technology–the same phenomenon that has revolutionized Computer Generated Imagery (CGI), put accurate GPS receivers in every car, and complex stock trading applications on computer desktops. Previously, the technology used in learning algorithms was limited to the realm of super-computers. Now, the algorithms use sophisticated methods to be computationally more efficient and are deployable to basic enterprise-class servers.

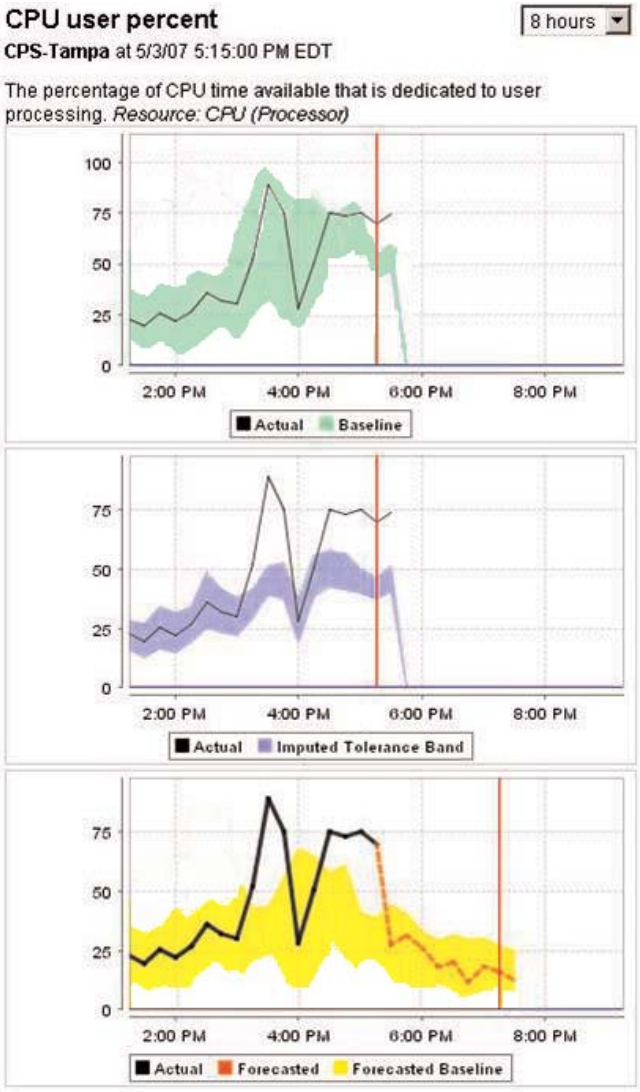

One of the important but often overlooked advantages of a mathematical approach is the total lack of bias with which these algorithms operate. Fixed rules and thresholding models - while they may work in many instances–inherently reflect human biases of the users who create them. Self-learning technologies replace this "guesswork" modeling with unbiased mathematical certainty. The software has no preconceived notions of how metrics or components are expected to behave–it simply observes the data and produces objective results.

Netuitive Self-Learning Technology Works Today

The words "automation" and "forecasting" may raise an eyebrow with skeptics that self-learning technology is too new or complex to really work in IT environments. But proof points exist today to indicate that the technology is already working–and very well.

The first indication is that the mathematics behind self-learning performance management is already being used successfully in other applications.

- For example, process manufacturing plants - such as oil refineries, chemical plans, power plants–regularly use similar technology for factory automation to monitor manufacturing processes.
- Weather forecasting today is more accurate than ever–especially for short term forecasts–through the use of statistical math models.
- The healthcare profession uses diagnostic tools to analyze multiple health metrics - pulse, respiratory rate, blood pressure etc.–and examine them holistically to help diagnose the probable root cause of certain ailments.

The second, more important indication is hundreds of companies in banking, telecommunications, and retail are already deploying the technology today on a large scale and seeing results and great return on their investments.

Jim Hirschauer, Technology Architecture Manager and technical expert for Wachovia's Corporate and Investment Banking Group, trusts the technology implicitly to warn him of impending problems. "After two weeks…the baseline thresholds and algorithms get very, very accurate for us… We take it very seriously when it predicts we'll have a problem with something."[6]

A large, multi-national Telco was investigating Netuitive to monitor critical business systems which had, in the past, experienced problems during the rollout of new services. After a successful trial, during which Netuitive forecasted a performance problem that legacy monitoring tools missed completely, the Telco embarked on an aggressive rollout. In less than 9 months, a small but dedicated team deployed Netuitive software to 8 data centers to monitor application performance for 300 mission-critical applications. This rollout consisted of 5,000 servers, 457 service views, and 700 databases.

One of the rollout Team Leads noted "Netuitive has the best application and system performance correlation engine in the market. The Netuitive alerts are the best performance alerts because they are composite alarms and the performance metrics are ranked by contribution."

Companies that deploy Netuitive invariably see the same core economic benefits:

- Automation of time-consuming administrative tasks for performance monitoring
- Early notification of impending performance problems
- Improvement in the service levels for mission-critical, revenue generating applications

Conclusion

Three consecutive years of bad to mediocre satisfaction ratings certainly reinforces the point that traditional, threshold-based performance management technologies have failed to inspire confidence that they can adapt to the dynamic environment in today's data centers.

In the long and winding path that is the evolution of systems management technologies, leading analysts such as Gartner and Forrester both agree that automation and "behavior learning" technologies are needed across multiple levels of fault management processes for IT Operations staff to deal with the challenges of tight resources and increased service level expectations for business services.

Based on technology and statistical methods already proven in other applications and industries, self-learning performance management software is making its mark in the industry. And the proof is in production case studies where the technology has allowed operations staff to proactively prevent outages.

The promise of self-learning performance management and the role Netuitive as a leader in the space is best summarized by the comments of Joseph Ambrose, VP of Operations and Engineering for a major grocery retailer: "Netuitive brings us a unique, automated intelligence layer for our IT environment."

More Information on Netuitive and Self-Learning Technology

For more information on how Netuitive self-learning technology products that address the challenges discussed in this white paper, please visit our web site at www.netuitive.com

For more on Netuitive's technology please contact a Netuitive sales representative for a copy of the white paper "Netuitive Self-Learning Technology - A Technical Overview."

References

1 Evelyn Hubbert, "BSM: a New Approach for Quick Wins", Forrester webinar, March 19, 2008

2 David Williams and Debra Curtis, "Climbing the Maturity Mountain from Event Management to Business Service Management", STR24_111, 6/08, AE, Gartner IT Operations Summit, June 2008

3 Denise Dubie, NetworkWorld, "Big 4 management vendors squeak by with a passing grade, garner little confidence from Gartner attendees ", Network Workd, June 23, 2008 http://www.networkworld.com/community/node/29188

4 Cameron Haight and Donna Scott, "Big Four Vendors Face Competitive Threats", Gartner ID Number: G00146927, April 13, 2007

5 Denise Dubie, "Vendors Scramble to Rein in Virtual Environments", Network World, September 4, 2008 http://www.networkworld.com/news/2008/090408-virt-man.html

6 Kevin Fogarty, "Wachovia Reaps Rewards with Active Systems Monitoring", Baseline Magazine, July 8, 2008 http://www.baselinemag.com/c/a/IT-Management/Wachovia-Reaps-Rewards-With-Active-Systems-Monitoring/

Self-Learning Performance Management

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, Basingstoke
RG24 8AL United Kingdom
Main : +44 (0)844.546.5001

Web : www.netuitive.com

About Netuitive, Inc.

Netuitive helps IT organizations reduce operating costs and increase productivity by automating the most difficult, time-consuming and costly processes in managing systems performance and business services. Major companies and organizations such Dow Jones, XM Satellite Radio, BlueCross BlueShield, Coca Cola Enterprises, ADP and the National Institutes of Health rely on Netuitive SI and Netuitive Service Analyzer software to increase the value of their existing performance management systems and assure the availability of business-critical applications.