

Improving APM Data Usability with Machine Learning Predictive Analytics

Important lessons learned from the world of Big Data analytics exposes information hidden in huge APM data stores enabling higher service levels at a lower cost.

Prelert is a venture-backed startup with considerable expertise in successful IT management companies and a serious pedigree in machine learning and computation mathematics.

Executive Summary

Despite decades of vendor improvements to infrastructure and application management tools, IT executives are still challenged by:

- The embarrassing failure of proactive management efforts to resolve more than 50% of performance issues before users report impact;
- The troublesome diversion of expert resources away from strategic projects to support problem diagnosis efforts; and
- A worrisome growth in recurring, unresolved performance issues now estimated to exceed 20% of reported incidents.

The good news is that most organizations now have an overabundance of management data describing the behavior of critical applications and their underlying infrastructure. The bad news is that, 'rolled up' to minimize storage requirements and scattered across domain specific 'silo' repositories, this data is largely unusable for diagnostics efforts.

As a result, escalation teams are assembled that drain precious resources from strategic projects. These teams then waste precious time manually searching and cross-correlating disparate data sources in an attempt to resolve performance issues. And because it is often more practical to restart affected servers than to spend days troubleshooting a problem's root cause, the number of recurring issues grow over time to the detriment of service levels and business risk management.

Due to advances in Artificial Intelligence and computational mathematics, the technology now exists to address these challenges. It is the same technology that shapes our web experience, forecasts our credit worthiness and helps authorities identify the speech or facial features of terrorists. For the first time in the history of IT management, cost effective solutions are available that allow IT organizations to quickly analyze tremendous volumes of data distributed across disparate data stores and identify patterns that foretell good or bad outcomes.

This paper explores the advances in Big Data Analytics and how they can be easily applied to existing application performance data stores to significantly improve proactive management, slash troubleshooting times and reduce the support drain on expert development resources.

The technology now exists to analyze huge volumes of data, of different types, in disparate data stores to foretell negative performance outcomes and diagnose them as they develop while reducing the need for human configuration and maintenance of existing management systems.

The Usability of Application Performance Data

A recent survey conducted by TRAC Research illustrates APM frustration levels. There remains a huge gap between the proactive management goals of the majority of IT teams and reality. When asked to identify key challenges faced in managing application performance:

- 63% of respondents said that expert staff were spending too much time troubleshooting problems;
- 61% said that users were still identifying performance issues before operations were aware of them; and
- More than 40% identified the usability of data output from management tools as a problem.

Another alarming trend is the increasing percentage of performance problems that are reoccurrences of ‘known problems’. In recent briefings with Forrester, Gartner and the Enterprise Management Group analysts it was agreed that more than 20% of escalated performance issues, those handled by the experts who are supposed to be working on new projects, are the recurrence of a known problem that was never fully resolved.

On closer examination of the challenges cited by APM users it becomes obvious that issue of ‘the usability of APM data’ is likely causal to every other challenge mentioned. Naturally, if the data produced by a monitoring system is of questionable value, diagnostics will take too long, users will be calling the operations team to notify them of problems, visibility to end-user experience and transaction flow will be limited and as a result management costs will be too high.



In actuality, the data collected in most IT environments is highly usable. Properly analyzed, it can provide the behavioral insight required for true proactive management.

The problem is that this wealth of data is scattered across multiple tools belonging to different organizations and usually ‘rolled up’ to minimize storage requirements at the cost of the rich detail needed for diagnostic purposes.

Due to the field of Big Data analytics, the technology now exists to leverage this data. Advances in analytics technology now enable us to quickly analyze disparate data sets and identify patterns that foretell good or bad outcomes.

These advances can be applied to existing management environments through a layer of advanced analytics in products such as Prelert that processes raw data from each tool in real-time to provide true proactive alerting and diagnostics.

Many vendors are now claiming advanced or predictive analytic capabilities so it is important for the compare vendor claims against the abilities of their underlying technology. To that end, Prelert recommends the following 3 core capabilities, or lessons learned from the world of Big Data analytics.



Lesson 1: Pattern based predictive analytics expose behavior activity patterns that foretell negative outcomes.

The more advanced performance analysis products recognize the fact that single variable thresholds are of little use in and of themselves. Instead they analyze the behavior of multiple, related variables. A classic example is the relationship between server CPU utilization and response times. CPU utilization spikes may be of no consequence unless they are related to a slow down in performance.

The most advanced analytics systems, however, are based on a more accurate understanding of the nature of problems in complex application environments. They are not based on simplistic relationships between a few metrics. Each transaction or outcome is the result of a causal chain of component interactions. A browser interacts with a front-end server that communicates via middleware to a back-end server. A database lookup then returns information back up the chain to a user. Prelert refers to these causal chains as ‘activity patterns’.

The challenge of a truly advanced analytics product is to learn these activity patterns to understand normal system behavior and identify problems. A slow end-user response time may be impossible to diagnose without understanding the full chain of component relationships that caused the problem.

Why Activity Pattern Analysis is the Best Approach to Proactive Application Performance Management

A typical application performance problem

W = Service Failover (log message = 'Success')

X = Network Spike as load shifts (threshold alert)

Y = Server NIC has buffer overflows (no threshold)

Z = App Users are Disconnected (Users call help desk)

Log Files

T

W

Metrics


Y

P

X

Help Desk

Z



Without pattern analysis, escalation teams must manually search through disparate data sets to find causal relationships!

A branch of Predictive Analytics is dedicated to uncovering anomalies in activity patterns that foretell a positive or negative outcome. You may have experienced this software if you’ve ever used your ATM card in say San Francisco for breakfast, Denver for lunch and then not been able to withdraw money from an ATM in New York City that evening. In the last decade, considerable development has occurred in the field of pattern recognition particularly as it applies to uncovering patterns in huge data volumes.

Pattern based predictive analytics are employed by Prelert to learn these causal chains or activity patterns. Prelert learns the relationships between components as well as the typical behavior of each related component. By scoring occurrences of repeating patterns according to their anomalousness or rarity it flags potential problems. Prelert provides both the components of the pattern and their individual behaviors, as the problem is unfolding. Support teams are alerted proactively to a problem at its earliest phase of development and diagnosis is reduced to a matter of minutes.

Often the behavior of individual components in a problem activity would not, in and of themselves, generate alarm. Since Prelert identifies anomalies in activity patterns, whether or not the metric values would typically trigger an alarm, it is particularly adept at catching the earliest warning signs of developing issues. Often these anomalies develop hours before the system degrades enough to affect end-users giving support teams ample time to remediate them proactively.

It is in this manner that Prelert fits the classic definition of Predictive Analytics for it’s ability to identify anomalous activity patterns that foretell a negative outcome.

Lesson 2: 3rd generation machine learning technologies can infer system behaviors from descriptive data without human input

Today's application performance management tools collect huge amounts of detailed data. Unfortunately, most APM systems require users to identify key performance indicators, inter-metric relationships and sometimes even threshold values.

The APM system alone, behind one regional bank's on-line banking system, monitors close to one million metric measurements per minute. How likely is it that a team of experts can define all of the potential causal relationships in such a complex environment? Even if, through an improbable amount of time and patience, they could, then how likely is it that they can keep this relationship model updated to accommodate change?

In actuality, much of this configuration work is never completed. So naturally, a significant percentage of the problems the system experiences will remain tremendously difficult to diagnose. In fact analysts from such diverse groups as Gartner, Forrester Research and Enterprise Management Associates, agree that more than 20% of the incidents handled by escalation teams are repeat occurrences of previously unresolved 'known problems'.

Even among the camp of APM analytics vendors that offer multi-variant analysis the majority require human experts to define KPI and inter-metric relationships. This requirement that humans define the monitored environment consistently results in lengthy deployment times, an inaccurate or partial definition of the actual behavior of the system and the requirement that this model be continually updated as the monitored system changes.

Machine Learning, a branch of Artificial Intelligence, has evolved from the early days of Expert Systems, through the troubled years of Neural Networks and arrived at a 3rd generation with capabilities of which many otherwise savvy technologists are still unaware.

The machine learning systems developed in the last decade are fully capable, given sufficient data, of inferring the behavior of a system without human intervention.

If, like many others upon first hearing, you struggle to believe this assertion, ponder the following example. Google shapes your very web experience from the ads you see to the content you are served based on a careful analysis of your surfing habits and those of millions of others that exhibit partially similar habits around the world.

Do you think Google employ's hundreds of psychologists that set metric value thresholds, or write rules to accomplish this? If the answer is yes, think of the complexity involved in just the one example involving whether or not to show ads for bathing suits to someone booking travel in the winter. If they are from northern states booking travel heading to the Caribbean that might be a good idea. But if you think about it for a few minutes, you will arrive at a rather tedious decision tree that would be required to accommodate people in other hemispheres, cultures with strict bans on exposing skin, people traveling north, etc., etc.

The fact is that advanced machine learning analytics already shape your everyday experience in countless ways of which you are largely unaware.

Prelert is the only Predictive Analytics solution for IT operations that is 100% self-learning. The software installs in hours and integrates with existing monitoring systems to process raw data as it is collected. Prelert can self-learn behavior patterns based on two days of data (which can be historical). It then begins providing value accurately identifying problems and their causal activity patterns.

It is impractical to assume that human IT experts can develop and maintain models sufficient to describe all of the likely failure modes that exist in today's complex application environments.

At the same time, Machine Learning technology has advanced to the point of inferring and continually updating that model given sufficient data in the form of monitoring metrics and logs.

Lesson 3: Probabilistic computational mathematics enable real-time analysis of huge volumes of data of different types, time granularity and quality

When evaluating the analytics either embedded in existing APM products or offered by Prelert competitors, the following challenges inherent in analyzing huge volumes of data of disparate sources, types and quality should be considered.

Dimensionality and the size of the data store

Previously we referred to an On-Line Banking system that generates 1,000,000 performance metrics every minute. Prelert analyzes each metric every 15 seconds. Understanding relationships and activity patterns in data requires detailed analysis of metrics and messages. The analytics must compare the occurrence of metric A to message B, message C, metric D etc. to identify the relationships between messages.

One Prelert customer is dealing with a system that generates 1,000,000 log messages per second. In ten minutes of analysis this involves 108 log messages requiring approximately 1016 comparisons. Comparing each message to every other to identify relationships is not practical.

To overcome this challenge of dimensionality, most analytics systems limit the number of comparisons it is possible to make to a few hundred. While this may be computationally efficient, it will not suffice to find unknown relationships in even modest data sets.

Modern inference techniques allow Prelert to reduce the dimensionality of the problem and solve it more efficiently. Using this technology, Prelert can effectively scale far beyond the limits of competitive solutions and identify unknown relationships among the even the largest IT performance data sets.

Analyzing data of different types

Time series data, which represents the value of performance metrics sampled at regular time, is synchronous in nature. Notification data (e.g. log file messages, traps, events/alerts from management systems) on the other hand, is asynchronous as it reflects events within the system. Both data types must be analyzed to facilitate accelerated troubleshooting.

The normal troubleshooting process followed by escalation teams is to observe the KPI metric values that prompted an alert, look at other metrics known to be associated with the KPI and then turn to an log files to search for notifications that would shed more light on the context of the problem. An analytics product should automatically analyze and group all of the notifications and metrics associated with an event and provide them to the diagnostics team.

Unlike competitive offerings, Prelert self-learns relationships between data types by leveraging the full attribute value in log files and event notifications. If your analytics vendor cannot do the same, your support and escalation teams will be left to search terabytes of data trying to establish those relationships manually.

Analyzing data of different qualities

An even larger challenge lies in the fact that data arriving from different sources can be of significantly different quality. Time series data can be of different granularities (e.g. 1s, 15s, 5min intervals). Data can be missing or contain significant gaps. The analytic approach taken must assume uncertainty in order to provide the most accurate results.

As an example, multi-variant times series analysis products must assume values for missing data and choose the longest time interval amongst their sources as the analysis time window. This means data may be incorrect and then statistically correlated as if all system activities fit neatly in to the time buckets of the monitoring system. This approach fails to detect causal activity patterns that span multiple time buckets (a significant percentage of activities in the real-world).

Prelert uses 'time warping' and probabilistic mathematics, two unique and highly sophisticated approaches, to maximize the accuracy of our analyses. The probabilistic approach assumes inaccuracies, gaps and uncertainties in the underlying data to yield more accurate results than standard statistical correlation approaches. 'Time warping' is a methodology that allows the system to accommodate events that are related across time even if that relationship changes temporally.

Summary

A 3rd generation of analytics technology is now available that is typified by three advances in computational mathematics and artificial intelligence. When applied to the goal of proactive management, these advances fundamentally alter the value equation of performance analytics.

If your organizations goals include improving proactivity, reducing troubleshooting times and getting more answers from existing data, you should evaluate adding advanced analytics to your current management tools. The following capabilities are key in your competitive evaluation and selection process:

- *Do the analytics allow the cross-correlation of data from multiple management tools (e.g. vCenter, network management and APM)?*
- *Can the analytics code identify and diagnose failures regardless of whether the data is 100% time series metrics, 100% notification (log files and events) or a mix of both?*
- *Are the analytics capable of detecting the 'activity patterns' (the causal sequence of events) behind application behaviors and performance issues?*
- *Does the analytics offering maximize the use of machine learning to reduce implementation times and ongoing 'tuning'?*
- *If it requires humans to pre-define KPI and or relationships between data metrics, who is going to maintain these definitions as your monitored environment evolves?*
- *What is the limit to the number of metrics and notifications that can be compared to identify or baseline inter-relationships? What % is that of the total number of monitored metrics?*
- *How do the analytics handle data from multiple sources that collect data at different time intervals? If its 'time buckets' are set to the longest interval, how informative will the resultant analyses be in your unique environment?*

Because Prelert leverages the latest advances in machine learning and predictive analytics, we surpass any competitive offering in our ability meet these critical considerations.

Prelert is provided as an easily deployed layer of analytics that leverages the data collected by your existing management tools.

Prelert provides insights that enable orders of magnitude improvements in services levels while increasing expert resource productivity and slashing the amount of time spent troubleshooting and maintaining management tools.