PROBABLE CAUSE ANALYSIS: A KEY VALUE DRIVER OF AIOPS

In their report entitled "IDC FutureScape: Worldwide CIO Agenda 2019 Predictions," IDC predicts that, by 2021, 70 percent of CIOs will aggressively apply data and AI to IT operations, tools, and processes. Driven by demands for improving service levels, efficiency, and agility, IT leadership is clearly counting on the promise of artificial intelligence for IT operations (AIOps).

As CIOs set out to pursue their AIOps initiatives, it will be critical to establish the near-term wins that are vital in demonstrating value and fueling longer-term buy-in, support, and investment. Toward that end, many IT leaders are struggling to determine the best near-term use cases to start with.

This is the third in a series of posts I'm publishing on AIOps use cases. In these posts, I've been focusing on those use cases that offer organizations some of the most significant near-term potential. In our last post, we examined how AIOps enables teams to establish event noise reduction and predictive alerting. In this post, we'll look at some of the near-term gains AIOps can provide in the area of probable cause analysis.

The Demand for Effective Probable Cause Analysis—and the Penalty for the Lack Thereof

Today, the ability to deliver reliable, optimally performing digital services plays an increasingly influential role in an organization's performance and competitiveness. Even small hiccups can delay critical operations, frustrate customers, and erode revenues. Business service downtime has an immediate, significant budgetary impact. For example, one survey found that, for more than 80

percent of businesses, an hour of downtime costs more than $300,000.

Given these realities, the pressure continues to mount on IT organizations, who are tasked with ensuring that required service levels are attained. To deliver reliable, optimized services, IT Ops must be able to identify issues that arise across the environment, quickly determine the cause of that issue and remediate it to maintain performance and service levels. In these efforts, effective capabilities for probable cause analysis are an imperative.

What is Probable Cause Analysis and Root Cause Analysis?

Probable cause analysis is the ability to understand relationships between infrastructure, applications and services and correlate millions of monitoring data points including performance metrics, events, logs, anomalies and baselines to deliver a scored and ranked list of the most likely causes for any problem in the environment. Root cause analysis refers to the next stage of problem identification. Log analytics enable deep analysis of log files to troubleshoot, recognize patterns, detect anomalies and identify the root cause.

Issues can be anything across the environment from a server being down to a slowdown in application response to network latency problems or CPU capacity levels. In some cases, probable cause analysis and root cause analysis can be a straightforward exercise. For example, a server administrator could receive an alert that a server is down, inspect the server's performance metrics, see that that CPU utilization is maxed out, and take steps to reallocate workload to another server. In other cases, particularly in today's interrelated, complex environments, probable cause analysis may require the investigations of several domain experts and different types of data from a number of monitoring tools.

At a high level, probable cause analysis requires a combination of effort, expertise, and data. Ultimately, the more complete, current, and targeted the data, the better equipped the administrator will be to assess the probable cause. In other words, the better the data, the less expertise and effort that will be required for probable cause analysis.

The Problem: Data Volumes are Overwhelming, While Insights are in Short Supply

In pursuing probable cause analysis, too many teams currently have the odds stacked against them. First, as outlined in our last post, that's because environments continue to grow more complex. Whether it's due to the proliferation of microservices, DevOps, or multi-cloud approaches, the reality is that environments continue to grow more dynamic, interrelated, and composite in nature.

Exacerbating these realities is the fact that teams tend to be battling with very limited visibility. Operators rely on technology- or domain-specific tools that only provide a fragmented view of the environment. With these tools, it's difficult, if not impossible, to assess how a given system-level issue affects the business services that run on top of the infrastructure. At the same time, if an issue is discovered at the business service level—for example, users are calling to complain about a service being down—it's very difficult and time consuming to determine where the actual issue is.

These obstacles have been significant, and only continue to be compounded as IT environments keep growing in scale and complexity. Ultimately, these obstacles leave IT teams plagued by staff inefficiency, high costs, and poor service levels.

How can AIOps improve Probable Cause Analysis?

Today, AIOps represents a key approach for teams looking to more quickly, accurately, and efficiently diagnose issues. By harnessing AIOps capabilities, teams can:

- Employ service modeling to gain critical context around how users and business services are affected by issues.
- Correlate millions of monitoring data points, including metrics, events, logs, anomalies, and baselines to automatically identify the causes of issues.
- Score and rank causal metrics to quickly reveal the most likely source of issues.
- Use analytics on log files to drill down into root cause.
- Employ event analytics to track event patterns within the context of applications.

How Have Organizations Benefited from AIOps?

By employing AIOps solutions, organizations can gain the advanced probable cause analysis capabilities they need to maximize their speed and efficiency in troubleshooting and remediating issues. We've seen organizations across diverse industries achieve 30 – 75% reductions in the time it takes to diagnose issues and similar reductions in mean-time-to-repair (MTTR).

Following are examples of two TrueSight customers that have are harnessing AIOps to fuel enhanced probable cause analysis:

- **Park Place Technologies.** Park Place Technologies has employed an AIOps solution to power more intelligent probable cause analysis. By doing so, their teams have been able to gain better, more timely insights, which has enabled staff to resolve incidents 31 percent faster and achieve a first time fix rate of 99%.
- **Brazil Ministry of Education.** The Brazil Ministry of Education deploys AIOps to understand event service impact across their environment and uses log analytics to identify root cause. Now, when major infrastructure events do occur, the team is able to do root cause analysis 50 percent faster than it could before.

The Potential of AIOps

For today's IT operations teams, delivering optimized service levels is a vital imperative. However, tracking and managing service levels only seems to get more difficult as environments grow ever more complex and dynamic. This is a key reason why the promise of AIOps is so compelling. With the right AIOps platform, teams can far more quickly and intelligently detect the probable and root causes of issues.

Be sure to keep an eye out for our next blog post in this series, which will provide a detailed look at another use case in the AIOps value chain: automatic remediation, incident, and change management to link IT Operations and the Service Desk. In the meantime, to learn more about our AIOps offerings, be sure to visit the TrueSight AIOps page.