


TELECOM: TRANSFORMING SERVICE ASSURANCE WITH AIOPS FOR THE ERA OF NETWORK AUTOMATION



In competitive and fast-changing telecommunications markets, service quality is a critical differentiator. Customers have demanding expectations for their experience, and they won't accept disappointing performance, unreliable availability, or a sluggish response to problems. Meeting their high standards was challenging enough in earlier eras. Now, as autonomous technologies increase the speed of network operations across software-driven environments, operators need to make sure that the service management layer becomes faster and more automated as well. In my earlier blogs, I discussed evolving approaches to Service Assurance and the growing role of artificial intelligence (AI), analytics, and automation in technology operations. Now, I'll drill down on the use of AIOps to transform Service Assurance for the era of network automation and support key use cases such as fault and problem prediction, zero-touch network operations, change, and dynamic inventory.

What network automation means for Service Assurance

The drive toward autonomous networking is advancing at full speed, as initiatives such as the Open Network Automation Platform (ONAP) seek to enable real-time, policy-driven orchestration and automation of physical and virtual network functions. While enabling operators to respond more quickly to customer requests, and to optimize operations across their increasingly software-driven network environment, this increase in automation will also place new importance on service assurance.

Network automation can be thought of in terms of in-band and out-of-band use cases. For in-band use cases (including many performance or capacity issues), things follow a relatively predictable set of patterns, enabling end-to-end automation. Here, Service Assurance doesn't need to do much more than record the processes performed. For out-of-band use cases, (physical infrastructure failures or rare alarm conditions) however, it can be either less clear what should happen next or require boot-on-the-ground interaction. As exceptions arise, an operator may need to get involved to make decisions. At that point, we need to ensure a level of governance across operational processes such as change—though without reverting to a fully manual approach.

AIOps offers a solution. AI, big data analytics, and machine learning make it possible to augment, guide, and increasingly replace human decision processes so that operators can ensure service quality more efficiently at scale to meet customer expectations.

Applying AIOps to key use cases

Service Assurance offers a variety of suitable use cases for AIOps, with reasonable large data sets to which AI and machine learning can be applied to cluster related faults, identify underlying network problems, prioritize resolution, and so on. High-value use cases include the following.

Improve prediction of faults and problems – A classic AIOps use case is to shift from proactive to preventive remediation by predicting problems before they've arisen. In a network automation context, as operators seek to increase the level of automation in an area that can't be fully automated, AIOps makes it possible to support this more predictive approach by automatically detecting anomalies based on established, dynamic baselines. Once a potential problem has been identified, machine learning enables fault clustering of related problems with the same root cause to speed troubleshooting and prioritize resolution. In some cases, it may even be possible to fix the automatically without human intervention.

Evolve toward zero-touch network operations – While the full zero-touch network operations center (NOC) remains for now an aspirational goal, with the network automatically assessing events and then making and acting on its own decisions, AIOps can already help operators achieve a higher level of automation. At the current level of solution maturity, machine learning can be used to guide operator actions based on what's been done in the past, how well it worked, and its chance of success on the current case. In situations where an engineer needs to go on-site to make a repair, the system can identify the right person to send to the right location with the right equipment to fix the problem effectively. By spending less time investigating before initiating a repair, the operator can reduce MTTR and touchpoints, deal with more problems concurrently, and improve their fix rate. If a given problem is found to be non-critical, the NOC can choose to wait to send an engineer until a more efficient time, such as a trip when multiple repairs can be combined on a single visit.

Manage change within an autonomous network

As network automation advances, a key step will be removing the human element in the way changes are handled. Having machines take over that function calls for a new way of thinking about change management, including the assessment, planning, approval, scheduling, and execution of changes. With AIOps, the assessment of the impact of a change can be automated, including not just its effect on SLAs, but also how its execution can be scheduled to minimize disruption to the customer—and whether the customer needs to be informed at all. Today, operators generally err on the side of caution to tell customers about situations with even a relatively small risk of disruption. If

AIOps allows us to reduce that risk below a certain threshold, this notification can become unnecessary, sparing customers the need to focus on contingencies with minimal likelihood of occurring.

Support Service Assurance and dynamic inventory

What does an increasingly dynamic world mean for service assurance? In the past, change was relatively infrequent, with predictable impacts. As virtualization increases, things move around more quickly, and fixes are made automatically, operators may not have such a clear understanding of the way things are connected and affected by each other. AIOps is now necessary to keep up with this more ephemeral environment. That's even more true as customer workloads move to the network edge, leading to a tighter relationship between the two and expanding the scope of service assurance accordingly. Changes to the network can now directly affect the performance of edge workloads, and in more significant ways—for example, an autonomous vehicle workload for which slower performance can be a matter not just of convenience, but of safety.

Simply put, in a world where everything is in motion, it may no longer be possible or even necessary to have a complete view of the entire environment at a given moment. Instead, we should focus on understanding specific impacts well enough to support faster decision-making—and AIOps will make that possible.