

SIGNAL-TO-NOISE RATIO: BRIDGING THE ITSM-ITOM DIVIDE



Over the past couple of years, I've been working with a large financial services organization and its director of IT operations, who has a mandate to improve operational efficiency, reduce costs, and rationalize the organization's tool stack. Of course, they still need to deliver a five-star user experience while doing more with less.

While the organization's digital transformation projects were delivering better customer self-service, the interaction between new (public cloud) and old (mainframe) technology stacks was proving to be a challenge. DevOps was pushing the boundaries with small and frequent releases, but monitoring was showing blind spots in end-to-end user interactions and slow recovery from system failures was impacting customer confidence. Key business stakeholders were getting nervous because increased customer churn could have a direct impact on revenue.

The director was facing three key challenges in:

- Observability at a business service level for prioritizing resources during critical situations.
- Noise reduction and artificial intelligence and machine learning (AI/ML)-based root cause recommendations to automate and speed recovery from poor performance and outages.
- Operating expenditures (OpEx) costs attributed to service and operations management tool sprawl.

The director talked about a recent outage where his staff was overwhelmed with 30,000 events that spanned user-initiated complaints to the global help desk and system-generated alarms from multiple monitoring tools. While AI/ML techniques were reducing alarm noise, it was still difficult to

narrow down root cause, which slowed the resolution time. The organization had no way to correlate between the incidents or to change and monitor events. They had plenty of tools, but it was like having dozens of watches where none of them could give an accurate time.

To put the problem in mathematical terms, he was dealing with a signal-to-noise ratio problem. Mission-critical business services are supporting an omnichannel user experience (mobile, web, voice, API); systems of engagement are hosted on cloud architectures; and systems of record hosted in private data centers are running distributed and mainframe applications. When something goes wrong, like the 30,000-event problem, a lot of noise is created from the events triggered by complex user transactions.

Noise is increased by end users opening service desk tickets, changes from DevOps automation, alarms based on static service level agreements (SLAs), faults from network equipment, and even automatically generated anomalies based on abnormal system behavior. Deciphering the signal from the noise to find the root cause of a system outage is a complex mathematical problem that's best addressed by a machine-based algorithmic approach instead of the traditional approach of putting business users and different IT departments on bridge calls. That just adds more noise to the blame game as disparate teams race to improve their mean time to innocence (MTTI).

For the purposes of this discussion, I am going to focus on the need for a more holistic end-to-end approach to automating noise reduction and root cause analysis. A machine-based algorithmic approach that applies trained AI/ML models can help bridge the divide between different teams and their disparate tools to increase the odds of quicker resolution and overall cost savings.

I will use the example of a mobile application where a user transaction that starts on a mobile device depends both on modern microservices running on a public cloud and monolithic applications hosted in private data centers. Slow response time on the mobile device can be difficult to triage because it can be attributed to code execution, infrastructure resource constraints, or network congestion anywhere on this long and complex execution path, as shown in Figure 1.




Figure 1. End-to-End Mobile Transaction Flow

The two areas that we need to focus on to more quickly decipher the signal from the noise are **IT service management (ITSM)** and **IT operations management (ITOM)**. Per our example above, a mobile application experiencing slow response times can potentially generate thousands of events. Figure 2 below shows how ITSM and ITOM systems exist in silos, with multiple tools that generate tickets, metrics, logs, and alarms, etc.




Figure 2. Event Noise from ITSM and ITOM

To address this, advanced analytics should be applied to ITSM and ITOM datasets to build situational awareness of impacted users and business services. **ITSM** systems deal with incident and change management and are a goldmine of historical and real-time data of user issues, change/work orders, and knowledge base articles. **AI service management (AISM)** applies AI/ML techniques like natural language processing (NLP), clustering, and classification to reduce incoming ticket noise, group major incidents/change events, recommend knowledge base articles, or automatically assign support groups.

Figure 3 illustrates how AISM can automatically group and create situational awareness, represented by the cluster of solid circles, diamonds, and squares. This automatic grouping of incidents and change events based on text, time, and relationships can help service desk agents focus their efforts on higher priority issues.




Figure 3. AISM Real-Time Incident Cluster

In our mobile application example, if the business service is impacted by slow response times, real-time incident correlation will help accelerate the triage process by eliminating the manual work that would have been done by multiple service desk agents. Figure 4 shows a real-time incident correlation dashboard where AISM automatically correlates and groups related incidents into a single view.

If the mobile application slowness is caused by a known issue, then this would trigger a runbook for quick remediation. However, in many cases this will require further investigation and correlation with the ITOM systems for further diagnosis. Speeding resolution requires an ITSM and ITOM integration strategy, which can be complex.

Our [research](#) shows only 23 percent of organizations have integrated the two disciplines. Without such integration, there are many hand-offs between teams and inefficient, error-prone manual processes that result in delays and customer dissatisfaction.

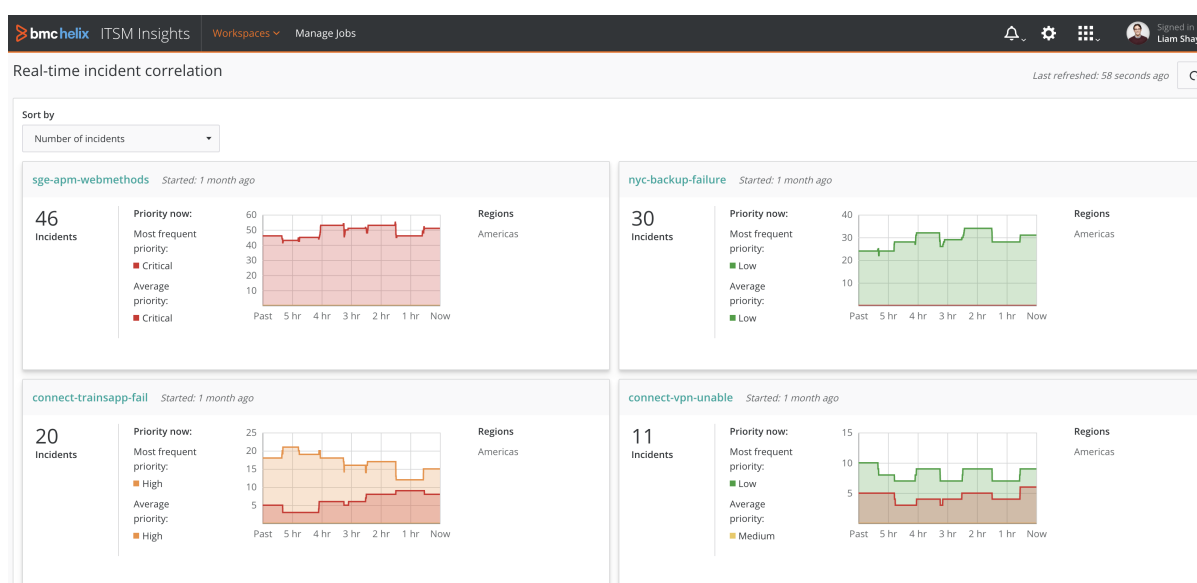


Figure 4. Real-Time Incident Correlation

ITOM systems deal with observability and automation to improve operational efficiency across applications, infrastructure, and networks. A typical mobile application journey is monitored by many tools that collect performance metrics, alarms, anomalies, logs, and topology information. Additionally, modern DevOps practices have increased the volume and frequency of changes in this dynamic landscape. Triageing and diagnosing production issues is akin to trying to find a needle in a haystack when dealing with very large and diverse datasets.

AI for IT operations (AIOps) applies AI/ML to these datasets to reduce noise and find root cause more quickly. AIOps creates situational awareness by applying algorithms for noise reduction, anomaly detection, clustering, and graph theory to automatically assess impact and arrive at a root cause, represented by the cluster of crossed circles, diamonds, and squares shown in Figure 5. Automatic grouping of alarms, metrics, logs, and topology for an impacted business service accelerates root cause analysis and recovery from outages.




Figure 5. AIOps Probable Root Cause Clusters

In our mobile banking application example, AIOps will accelerate the diagnosis process by identifying a probable root cause for the slow response time. Figure 6 below shows a service monitoring dashboard for an impacted business service where metrics, topology, events, and change are grouped together with an identified root cause. This saves support staff hours, if not days, of triaging and collecting evidence. A probability score for the determined root cause can also increase confidence and lead to automation for self-remediation in the future.

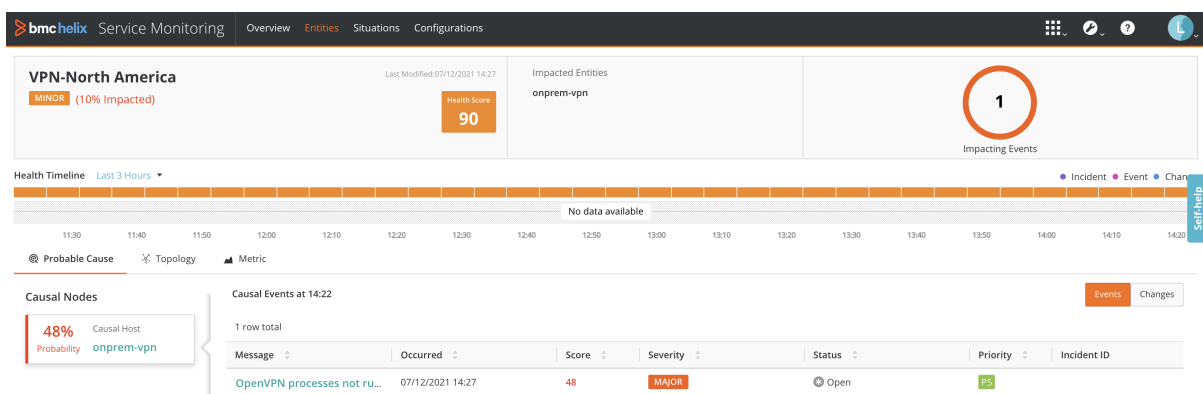


Figure 6. AIOps Service Monitoring Dashboard

Recognizing the lack of an integration strategy between ITSM and ITOM systems, there's still a need to correlate and bridge the gap between the insights discovered by AISM and AIOps to avoid error-prone hand offs and inefficient manual processes. Figure 7 illustrates how solid and crossed cluster groups can be automatically correlated using advanced analytics and common relationships. Automatic correlation between real-time incident clusters and probable root cause recommendations requires a service-centric approach that shares a common data model and datastore to correlate across shared resources.




Figure 7. AISM and AIOps Correlation

In our mobile example, the system would automatically correlate the reported slow response time incidents and change requests with the probable root cause, as shown in Figure 8 below. This empowers the service desk agent with the context to intelligently engage the right support groups, initiate tasks, or automate runbooks for quick remediation.

The screenshot shows the BMC Helix ITSM interface. At the top is a navigation bar with 'bmc helix ITSM' and various menu items. Below is a header for an incident: 'INC000000032176' with status 'Critical' and 'Major incident'. The incident title is 'I cannot connect to VPN'. Below this are fields for 'Created' (12/15/2021 12:19:03 AM), 'Updated' (12/15/2021 2:20:57 PM), 'Status' (In Progress), and 'Major incident: Yes'. A customer profile for 'Corey Cooper' is shown. The incident details section includes 'Incident type' (User Service Restoration), 'Reported date' (12/15/2021 12:19:03 AM), 'Operational category' (None set), 'Reported source' (External Escalation), 'Product category' (None set), 'Description' (I cannot connect to VPN), 'Resolution category' (None set), 'Service' (VPN-North America), 'CI' (None set), and 'Resolution product category' (None set). At the bottom, a 'Top probable causes nodes' section shows 'Impacted Server ONPREM-VPN' with a '48% Probability'. A dropdown menu is open next to it, showing options: 'View recent change requests', 'Create change request', 'Create task', and 'Relate configuration item'.

Figure 8. Incident and Probable Root Cause Correlation

The ability to improve service delivery with integrated ITSM and ITOM capabilities is what BMC refers to as ServiceOps. It brings technology and data together in one central platform ([BMC Helix Platform](#)) that spans organizational silos and has a common data store with open integrations to third-party tools, further strengthened by our recent [StreamWeaver acquisition](#). The entire solution is designed to help your organization reduce the signal-to-noise ratio, improve response time, and provide a better customer experience.