Reduced Actionable Events by
90% ⬇

Reduced Event Processing Time by
70% ⬇

Annual OPEX Reduction
23% ⬇

Annual OPEX Cost Avoidance
60% ⬇

CUSTOMER ENVIRONMENT

The world's largest service providers have many challenges in common, but perhaps one of the most pressing on the mind of the CIO is how to effectively manage the massive volume and velocity of the data that they collect daily. This large telecommunications service provider is no exception. With 1,000's of customers and billions in revenue focused on B2B services ranging from hosted data services, secure interconnectivity and internally built hosted applications, they needed to figure out how to continue to expand their services and customer base but at the same time deliver the highest levels of service uptime and reliability.

This was no easy task as the customer has multiple follow the sun command centers that manage their large complex multi-national infrastructure with tens of thousands of devices ranging from network, servers, databases, applications and mobile infrastructure. They use many new tools, but also rely on many legacy tools such as IBM Netcool for Event Management and BMC Remedy for ticketing.

The customer puts significant importance on ensuring resiliency and responsiveness of the services that they offer to their customers. Any downtime could have severe financial and reputation impact to their business and their customers.

KEY OBJECTIVES

Improve Visibility of Infrastructure and Cross-domain Issues

Through the last couple of decades the company has gone through several acquisitions and divestitures resulting in many new and legacy monitoring and event management tools to manage the fault and performance of their growing infrastructure. The company was looking for an AI driven intelligence platform that could help them to consolidate event and data feeds across the various systems. The platform needed to improve visibility and insight across services and applications without having to rip and replace systems that were already deeply embedded with their operational teams.

Reduce Event Noise to Focus on Service Impacting Issues

As the company continued to grow, they rapidly expanded the types of services they offered and the number of customers they needed to support. In order to support the new demands of the business, there was a substantial adoption of new technologies and an increased frequency of change needed in their environment to support these services. The result was a massive increase in the volume of event and performance data that needed to be managed by their existing teams (millions of events per month) – more than they could manage with existing teams. They needed an approach that would allow them to intelligently organize, group, prioritize, and respond to critical service impacting incidents more effectively and efficiently.

Reduce Time Spent on Manual Validation, Correlation & Routing of Events

In addition, when incidents arose, operational tools and Netcool would simultaneously show hundreds of red events. Determining root cause through this "sea of red" was difficult, time consuming and resource intensive. It was difficult to distinguish false positive, sympathetic or causal events related to the disrupted service. They needed to more efficiently determine root cause of problems so they could quickly address issues, improve response times and get their services operational again.

Improve Operational and Business Metrics

While the result of the initiatives helped to improve overall performance of their teams, the customer also needed to show improvements to their operational metrics. They needed to reduce overall Mean Time To Repair (MTTR), reduce Mean Time To Identify (MTTI), reduce Time To Escalate (TTE), and ultimately reduce overall OPEX costs.

GROK SOLUTION

Grok AIOps was implemented to provide better visibility and to enable the customer to improve their ability to manage increasing event volumes and improve service levels with their customers. Grok seamlessly integrated with existing infrastructure monitoring tools and ticketing tools and added a machine learning intelligence layer to reduce noise and improve their ability to respond and take action on incidents. After an initial proof of concept, Grok was fully implemented within weeks to begin delivering value to the organization. Several components of Grok's machine learning platform were deployed to address their key challenges.

- **Event Clustering:** Provided significant event compression by using machine learning to automatically build a representational model of the customer's infrastructure. This enabled Grok to automatically correlate and cluster similar and related events to reduce noise and the manual work required by NOC technicians to find root cause.

- **Anomaly Detection:** Quickly determined anomalous behavior in key portions of the infrastructure and reduced the reliance on traditional noisy threshold-based performance monitoring. Teams were able to proactively identify and address issues that potentially would have resulted in an outage or performance problem.

- **Incident Classification:** With machine learning algorithms, Grok created incident classification signatures which automatically classify events and detections into incidents with known actions and workflow such as automated enrichment, assignment and auto-ticketing without the need to create or maintain complex rules.

- **Intelligent Event Ingestion:** Out-of-box connectors with IBM Netcool and BMC Remedy enabled Grok to bi-directionally receive and update event and incident information with existing tools. Event status updates and newly found detections are sent back to Netcool and notes are updated with Remedy tickets. By allowing the technicians to leverage Grok information within the context of their existing event management and ITSM systems, Grok allowed the customer to drive huge operational improvements with minimal impact to their existing processes.

KEY RESULTS

Grok drastically reduced noise, correlated disparate events across multiple systems and helped to proactively detect anomalous behavior. The Grok platform enabled the customer to significantly improve their ability to respond to issues and get one step ahead of problems and outages which could have a huge impact on their customers. The Grok implementation was completed in weeks and the customer started to see event clustering, patterns and positive results within days and continued to improve as Grok self-learned the environment and continuously updated the machine learning models.

By complementing and integrating with their existing tools, Grok was easily adopted into overall operational workflows and required minimal time in onboarding with existing teams. The customer was able to use the existing work queues but gained powerful insight and intelligence with Grok's AI and machine learning capabilities.

The Grok platform was successful in driving measurable benefits and savings into operations. Grok was able to reduce noise in the customer environment by 90% allowing the teams to focus their resources on critical, service-impacting incidents. In addition, with Grok's event correlation and probable root cause capabilities, the customer was able to manage disparate events from multiple systems and reduce their event processing time by 70%. In the end the business was able to reduce annual OPEX costs by 23% with a 130% Return on Investment within the first 6 months.

The customer was able to use their existing tools, but gained powerful insight and intelligence with Grok's AI and machine learning capabilities

Noise Reduction

Probable Root Cause

Anomaly Detection

Intelligent Event Ingestion

Plug 'n Play Infrastructure Modeling

www.grokstream.com

Sales@grokstream.com