Ovum

On the Radar: Grok AIOps platform Grok aims to eliminate the noise and uses ML to assess the health of all metrics, events, and logs

Artificial intelligence technology for solving operational problems and their root cause

Publication Date: 27 Nov 2019 | Product code: INT003-000409

Michael Azoff

Summary

Catalyst

The service assurance teams in network operations and command centers now need to sift through a massive stream of signals to determine which events might be the root cause of an incident. Over the years enterprises have invested millions of dollars in tools for root-cause analysis systems, topology-based discovery, business service management, and application performance monitoring (APM). They have tried to solve incident root cause in many ways to increase service assurance, but the trouble is that there are more and more surveillance systems being deployed in the environments, which are driving more signals into these centers. The questions revolve around how to reduce that noise, how to get to the root cause, and then how to correlate the root cause into something that's meaningful for the organization so the problem can be fixed.

Grok AIOps has evolved by combining multiple layers of intelligence based on machine learning (ML), to drive improvement in noise reduction and predict when incidents are going to happen and do that prediction significantly ahead of when a human or an agent can do that with the tools in place today.

By eliminating noise and matching patterns of event clusters and anomalies to historic incident data, Grok delivers accurate root cause analysis but also gives an indication of the priority and impact of potential failures. Grok makes the lives of support teams far easier and their current research is moving past automated ticket creation into automatic remediation of faults.

Key messages

- Grok AIOps provides a significant reduction of noise, often up to a 90% overall reduction, suppressing and correlating events to predicted incidents.
- Predicted incidents are surfaced hours ahead of current manual or rule-based methods and offer prioritized notifications that classify the service impact and allocate appropriate incident tickets.
- The solution is easy to adopt and integrates well into current environments, bringing a machine learning intelligence layer without organizations having to change their tooling. The self-learning algorithms in the unsupervised layers bring immediate results.
- The next stage of Grok's development is to work on the automated remediation of predicted and actual incident issues.
- Grok's technology is notable for combining machine learning techniques with AI licensed from Numenta called Hierarchical Temporal Memory (HTM).

Ovum view

Traditional operations and alerting platforms produce too much noise and do not eliminate the vast amounts of false positives and uncorrelated events that frequently overwhelm support teams. Grok aims to eliminate the noise and use machine intelligence to assess the health of all metrics, events, and logs on every app or service.

Grok solves the ML context problem by combining unsupervised ML algorithms with supervised classification algorithms that provide the critical context needed to identify incidents before they happen. The anomaly detection that underpins the platform makes use of HTM technology licensed from Numenta, an organization that researches AI technology based on neuroscience.

Ovum is impressed with the way Grok uses multiple sophisticated correlation and classification mechanisms as well as machine learning to identify patterns that map events and anomalies to previously identified forms of incidents. This enables quick and accurate identification of the actions required to resolve problems before they become critical. Automating this resolution is the next step in Grok's path.

Recommendations for enterprises

Why put Grok on your radar?

If you have an Ops team that struggles to resolve a high volume of incidents and is swamped by the noise created by the monitors placed on your streams, Grok could significantly reduce your problems.

Because Grok ingests metrics and events from existing APM tools such as Datadog or Dynatrace, it's easy to incorporate data from a range of other sources, including the network infrastructure, applications, servers, IoT-connected sensors, solar arrays, and customer experience metrics. It can also accept data from any on-premises or cloud resources with its out-of-the-box connectors and REST API.

Grok's incident prediction capability detects, correlates, and routes 60% of the incidents with associated events and anomalies after minimal unsupervised self-learning. The predicted incident analysis can be easily integrated with other tools that are already in use, such as ServiceNow, to create awareness for support teams of an issue, as well as what the particular service impact is, and which group needs to be assigned to prevent service disruption. In many cases it is also possible to automatically open a ticket within the ITSM system and have it assigned to the correct group.

The remaining events or anomalies not automatically assigned are clustered, providing a reduced set of items to view and process by level 1 support. Any incidents that subsequently are opened manually are used by Grok as a new training set to be applied as part of future predictions.

The resolution of incidents requires accurate analysis of the root cause of the issue. Grok collects all the events, logs, and metrics that are available in the environment, correlates them into patterns, and the patterns are trained against historical tickets to be able to say that, for example, "when this pattern emerges within the underlying data, 99% of the time it relates to this specific type of ticket in your environment." Once Grok makes this connection it can pull all the information into the ticket before presenting it to level 2 support.

Highlights

Grok's AIOps platform ingests events, logs, and performance metrics and applies proprietary ML algorithms for event/log clustering and anomaly detection. This creates issue signatures that feed proprietary incident prediction and remediation classification algorithms resulting in automatically detected incidents.

Multiple Layers of Intelligence

Figure 1: Grok AIOps architecture showing intelligence layers

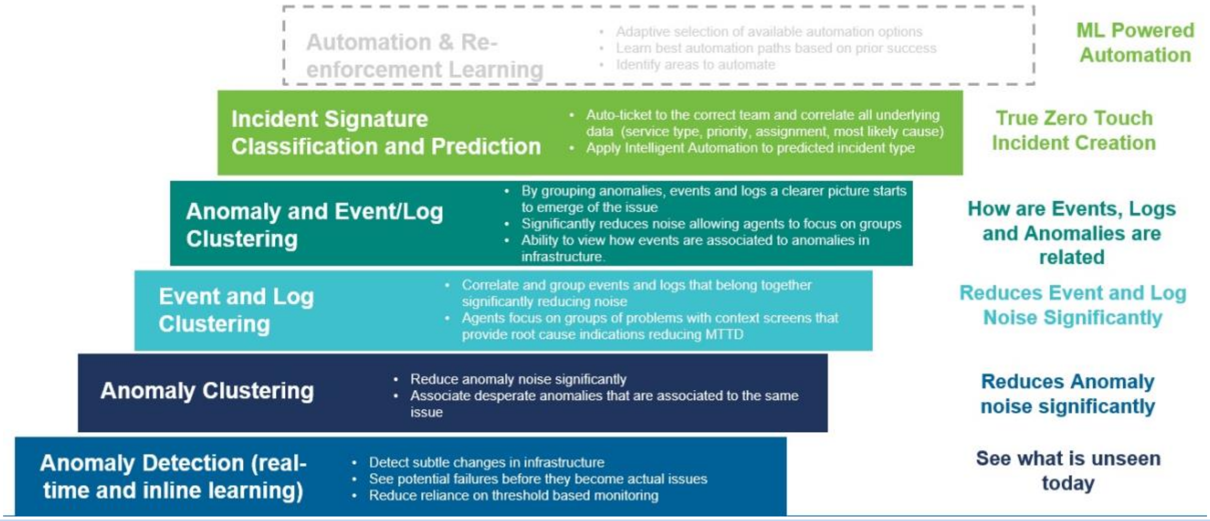

Automation & Re-enforcement Learning

- Adaptive selection of available automation options
- Learn best automation paths based on prior success
- Identify areas to automate

ML Powered Automation

Incident Signature Classification and Prediction

- Auto-ticket to the correct team and correlate all underlying data (service type, priority, assignment, most likely cause)
- Apply Intelligent Automation to predicted incident type

True Zero Touch Incident Creation

Anomaly and Event/Log Clustering

- By grouping anomalies, events and logs a clearer picture starts to emerge of the issue
- Significantly reduces noise allowing agents to focus on groups
- Ability to view how events are associated to anomalies in infrastructure.

How are Events, Logs and Anomalies are related

Event and Log Clustering

- Correlate and group events and logs that belong together significantly reducing noise
- Agents focus on groups of problems with context screens that provide root cause indications reducing MTTD

Reduces Event and Log Noise Significantly

Anomaly Clustering

- Reduce anomaly noise significantly
- Associate desperate anomalies that are associated to the same issue

Reduces Anomaly noise significantly

Anomaly Detection (real-time and inline learning)

- Detect subtle changes in infrastructure
- See potential failures before they become actual issues
- Reduce reliance on threshold based monitoring

See what is unseen today

Source: Grokstream

Grok uses multiple layers of intelligence to identify issues and support their resolution (see Figure 1).

Anomaly detection

Through its partnership with Numenta, Grok leverages the HTM algorithm for anomaly detection. Grok has found that HTM has consistently been the best performing anomaly detection algorithm in the industry resulting in less noise, fewer false positives, and more accurate detection. It can detect anomalies in signals beyond low and high thresholds, such as signal frequency changes that reflect changes in the behavior of the underlying systems.

Anomaly clustering

To help reduce noise, Grok clusters anomalies that belong together through the same event or cause.

Event and log clustering

Grok ingests all the events and logs from the integrated monitors and then applies to it event and log clustering algorithms. It applies two different algorithm types: one is a semantic grouping and the other is a dynamic time-warping algorithm.

Semantic grouping (clustering) will group events across multiple fields that fire at close proximity over time (name, domain), grouping events together and reducing noise.

Dynamic Time Warping also finds events that fire at similar time intervals over time, not only events that fire close in time proximity, but also those that regularly occur hours or even days later. Over time, the algorithm will find these events and start to group them into clusters as well provide additional noise reduction and correlation over and above methods that only leverage semantic grouping.

Semantic grouping combined with Dynamic Time Warping therefore provides better correlation and event clustering.

Anomaly and event/log clustering

The components identified and grouped at this layer of anomaly, event, and log clustering create a pattern for each incident, creating a signature that is then mapped to an incident type.

Incident prediction

Grok has developed proprietary ML techniques to ingest the output from both the anomaly detection and clustering algorithms to accurately predict IT incidents with significant advanced notification over current service assurance methods. It's the incident signature classification that does the linkage to the actual specific incident that Grok is looking to predict in the environment.

Use of ML in the solution

The first four layers in Figure 1 all use unsupervised learning algorithms. Once set in motion, Grok lets the algorithms run, typically going into a customer's site and delivering up to 80% to 90% noise reduction. The supervised learning portion occurs in the top layer (incident signature), where support staff need to train the solution on historical incidents data. Significant additional value comes from training on a couple of months of historical incidents.

Automation and reinforcement learning

The very top layer is a work in progress. This an area of research in which Grok is aiming to answer questions about how to automatically decide which automation to apply once the incident type has been identified. The idea is to give the system a space of options so that incidents are not explicitly tied to an automation resolution on a one-to-one basis. Grok will need a set of options because the permutations of options are too large to manage with a strict rule.

Background

Grok was formed as a partnership between the AI research organization Numenta and the founders of Avik Partners, who have successfully started and grown multiple technology companies in service assurance and automation, notably Resolve, before establishing Grok. The main players include the brothers Josh and Casey Kindiger, CEO and president respectively; Paul Scully VP sales; and Roumen Antonov, CTO. The Kindiger brothers, Scully, Antonov, and others at Grok have worked as a team for some 15 years.

Their first product to market as a team was Resolve, a tool to automate processes and IT operations, when they transitioned from a consulting company into a software company. Five years ago, Casey saw where ML and artificial intelligence was going, and he left Resolve two years before they sold it and started Grokstream.

Current position

Grok is deployed in three major production environments with version 3 of the product. In each version, Grok has been iterating its approach and strategy. Version 1 was a test phase, version 2 (beta) involved much larger market testing and was released to mid-market companies with 50 signups to the program. In Grok version 3, the company is focusing on large enterprises.

Data sheet

Key facts

Table 1: Data sheet:

Product name	Grok AIOps platform	**Product classification**	AIOps, APM
Version number	3.4	**Release date**	January 2020
Industries covered	All verticals	**Geographies covered**	Global
Relevant company sizes	Large enterprise	**URL**	grokstream.com

Source: Ovum

Appendix

On the Radar

On the Radar is a series of research notes about vendors bringing innovative ideas, products, or business models to their markets. Although On the Radar vendors may not be ready for prime time, they bear watching for their potential impact on markets and could be suitable for certain enterprise and public sector IT organizations.

Authors

Martin Gandar, Associate Senior Analyst

martin.gandar@ovum.com

Michael Azoff, Distinguished Analyst

michael.azoff@ovum.com

Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at consulting@ovum.com.

Copyright notice and disclaimer

The contents of this product are protected by international copyright laws, database rights and other intellectual property rights. The owner of these rights is Informa Telecoms and Media Limited, our affiliates or other third party licensors. All product and company names and logos contained within or appearing on this product are the trademarks, service marks or trading names of their respective owners, including Informa Telecoms and Media Limited. This product may not be copied, reproduced, Whilst reasonable efforts have been made to ensure that the information and content of this product was correct as at the date of first publication, neither Informa Telecoms and Media Limited nor any person engaged or employed by Informa Telecoms and Media Limited accepts any liability for any errors, omissions or other inaccuracies. Readers should independently verify any facts and figures as no liability can be accepted in this regard – readers assume full responsibility and risk accordingly for their use of such information and content.

Any views and/or opinions expressed in this product by individual authors or contributors are their personal views and/or opinions and do not necessarily reflect the views and/or opinions of Informa Telecoms and Media Limited.

CONTACT US

ovum.informa.com

askananalyst@ovum.com

INTERNATIONAL OFFICES

Beijing
Boston
Chicago
Dubai
Hong Kong
Hyderabad
Johannesburg
London
Melbourne
New York
Paris
San Francisco
Sao Paulo
Shanghai
Singapore
Sydney
Tokyo