Whitepaper

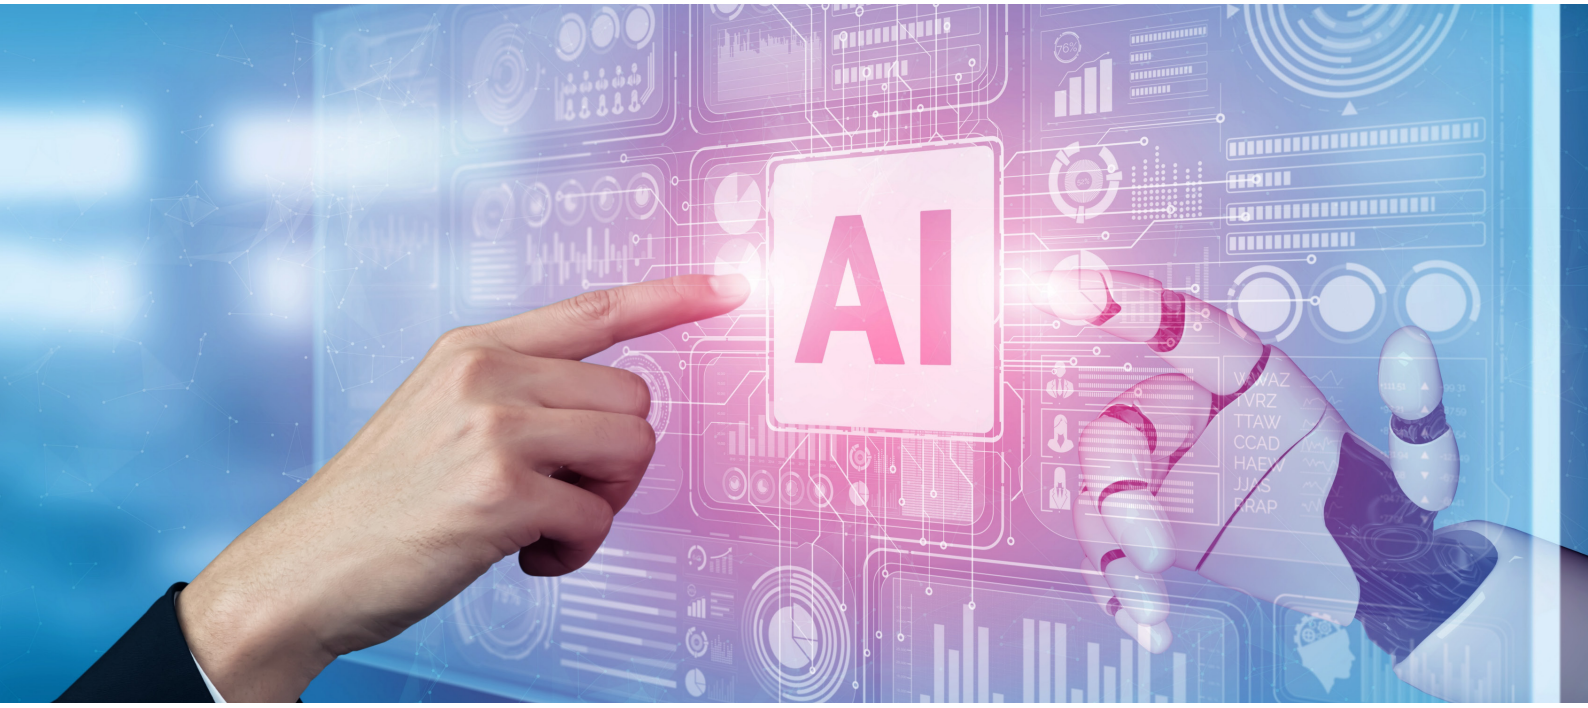

Algorithmic Alert Correlation

Gireesh Sreedhar KP

Senior Technical Manager - Location Zero, GAVS

Abstract

The modern day demands for always-on businesses and 24x7 uptime have sent IT monitoring into overdrive. While constant monitoring of infrastructure and applications is necessary, it can lead to a deluge of generated alerts that can quickly get overwhelming. This often results in a lot of wasted time & effort in just filtering out noise and identifying problem statements. Constantly having to deal with thousands of alerts each day causes alert fatigue, increases the possibility of missing critical alerts, and impacts the overall efficiency of the monitoring process. Billions of dollars are lost by enterprises due to such inefficiencies in IT operations.

Hence, chalking out an optimal strategy for alert management becomes critical. AIOps Platforms leverage Artificial Intelligence & Machine Learning (AI/ML) algorithms to filter out noise, correlate relevant alerts and prioritize them.

While there are many industry solutions for alert correlation, most of them do not go beyond basic alert correlation to intelligently address challenges in the determination of 'relatedness' in alerts. In this white paper, we take a close look at how these challenges can be addressed.

Determining Alert Relationships

There are many ways to determine alert relatedness. One way is through similarity definitions in the context of the IT landscape. A definition for instance, would group together alerts generated from applications on the same host, or connectivity issues from the same data center. This implies that similarity definitions depend on the physical and logical relationships in the environment. The mapping of relationships between the different IT components is interchangeably referred to as relationship map, topology map, or enterprise blueprint.

Commonly Available Functionality in AI-led Alert Correlation

Large volumes of data - consisting of process syslogs, alerts, event logs etc. - from the AIOps platform and other enterprise monitoring tools are ingested and processed in real-time at very high speeds. Then, a high-performance correlation engine uses AI/ML and navigates through the flood of alerts, eliminates noise, correlates relevant alerts based on underlying relationships, groups related alerts into high-level incidents, and prioritizes them depending on impact. This intelligent alert correlation drastically reduces the number of alerts to be acted upon and drives precision in root cause analysis.

The relationship map of the IT landscape is one of the key inputs for alert correlation. Some AIOps platforms have an Auto-Discovery & Dependency Mapping (ADDM) feature that automatically discovers infrastructure & applications, maps their physical & logical dependencies, and keeps the topology map updated by continuously monitoring the environment for changes.

The Challenges in Alert Relationship Determination

In real business scenarios, quite often there are a lot of restrictions – due to security and regulatory compliance requirements - on scanning internal networks to 'discover' IT components and their relationships. This eliminates the possibility of constructing a dynamic topology map to aid in alert correlation.

Secondly, a sizeable chunk of alerts that are actually related, are generated from entities that are neither physically nor logically linked. This implies that relying only on the relationship map for alert correlation would be akin to addressing only the visible portion of the alert iceberg!

To give a hypothetical example, let's say application A accesses a server S which is responding slowly, and so A triggers alert A1. This slow communication of A with S eats up host bandwidth, and hence affects another application B in the same host. Due to this, if a third application C from another host calls B, alert A2 is fired by C due to the delayed response from B. Now, although we see the link between alerts A1 & A2, they are neither physically nor logically related, so how can they be correlated? Such situations could imply thousands of individual alerts that cannot be combined.

AI Intelligence to Address These Challenges

These are some of the many challenges in IT operations that need to be solved by leveraging AI/ML algorithms.

While most AIOps platforms will utilize any available information on the IT components & their relationships obtained either through their own discovery & monitoring features or through other tools in the environment, ***it is very important for algorithmic alert correlation to be completely self-reliant and to function independently with no 3rd party inputs like discovery data, CMDB, topology maps etc.*** Ideally, the algorithm should not need any access to the infrastructure components and should be able to work with only the data supplied by alert monitoring tools, as it's input. This feature is critical in client environments where access restrictions are enforced.

What's required is an AI-based solution that independently discovers alert associations based on the various critical details available in the alerts. The relationship discovery needs to be continuous and should reflect the changing dynamics of the systems. Pure play AI algorithms for alert correlation will not need the support of external discoveries and inputs of source & problem mappings and will be able to derive relational mappings on their own using event input data from monitoring tools. This will be a unique alert correlation proposition in the industry.

Algorithmic alert correlation will need to leverage ML algorithms & Natural Language Processing (NLP) techniques, and additionally tap into the knowledge of Subject Matter Experts (SMEs) to address this problem. This is a classic example of the need for Human-Machine Harmony.

Let's take a deeper look into a powerful two-pronged approach to resolve this through the correlation algorithm. Each alert carries in it the 'source' information and 'problem' information. The core of the algorithm builds an association map based on these captious aspects of alert data such as the source and the problem.
The correlation algorithm delivers the powerful proposition of 'learning' relationships between the various sources of the alerts and 'learning' relationships between the different problems of the alerts. The algorithm calculates the score of each such relationship and builds logical maps of the sources and problems based on the scores.

> This approach is an innovation in the application of Machine Learning and is a non-intrusive way of relationship mapping that is built entirely using only alert data, which makes it a stand-alone process that is completely independent of external inputs. The relationship discovery is continuous, enabling the correlation process to mature with constant learning and unlearning from new data.

i) Relationship Mapping of Alert Sources

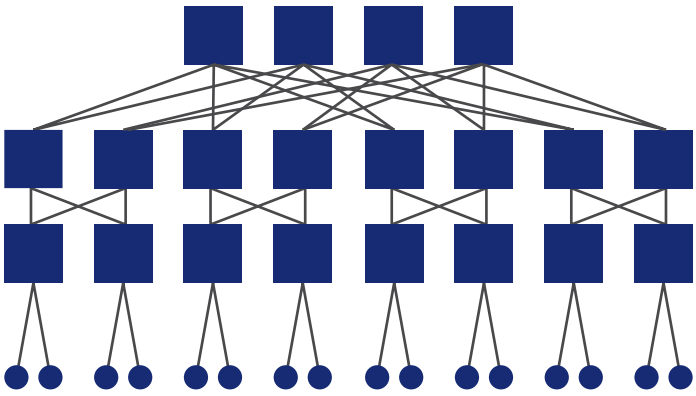

Here, the algorithm uses the source information in alerts to build a relationship map of devices in the environment. This eliminates the need for external discovery.

ii) Relationship Mapping of Alert Problems

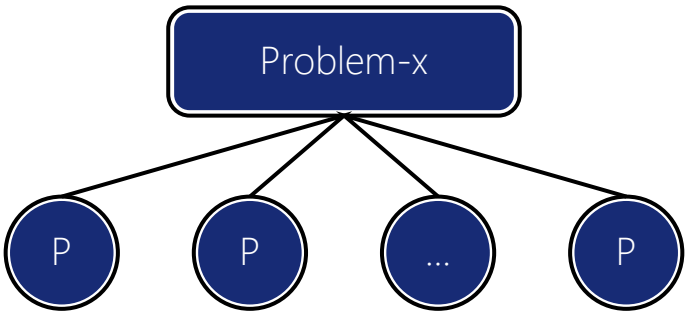

Here, the algorithm uses the problem information in alerts to build a relationship map of problems in the environment. This gives correlation the powerful dimension of 'problem' relationships.

Further, the correlation algorithm works on the principle of three dimensionality, taking three functional inputs into consideration; these are the WHEN, WHERE, and WHAT of the alerts that help determine the time, device relationship and problem relationship respectively.

The algorithm can also be customized by fine-tuning parameters based on SME inputs.
These parameters assign weightages to the different dimensions, in order to combine the critical alert aspects in the right proportions for correlation.

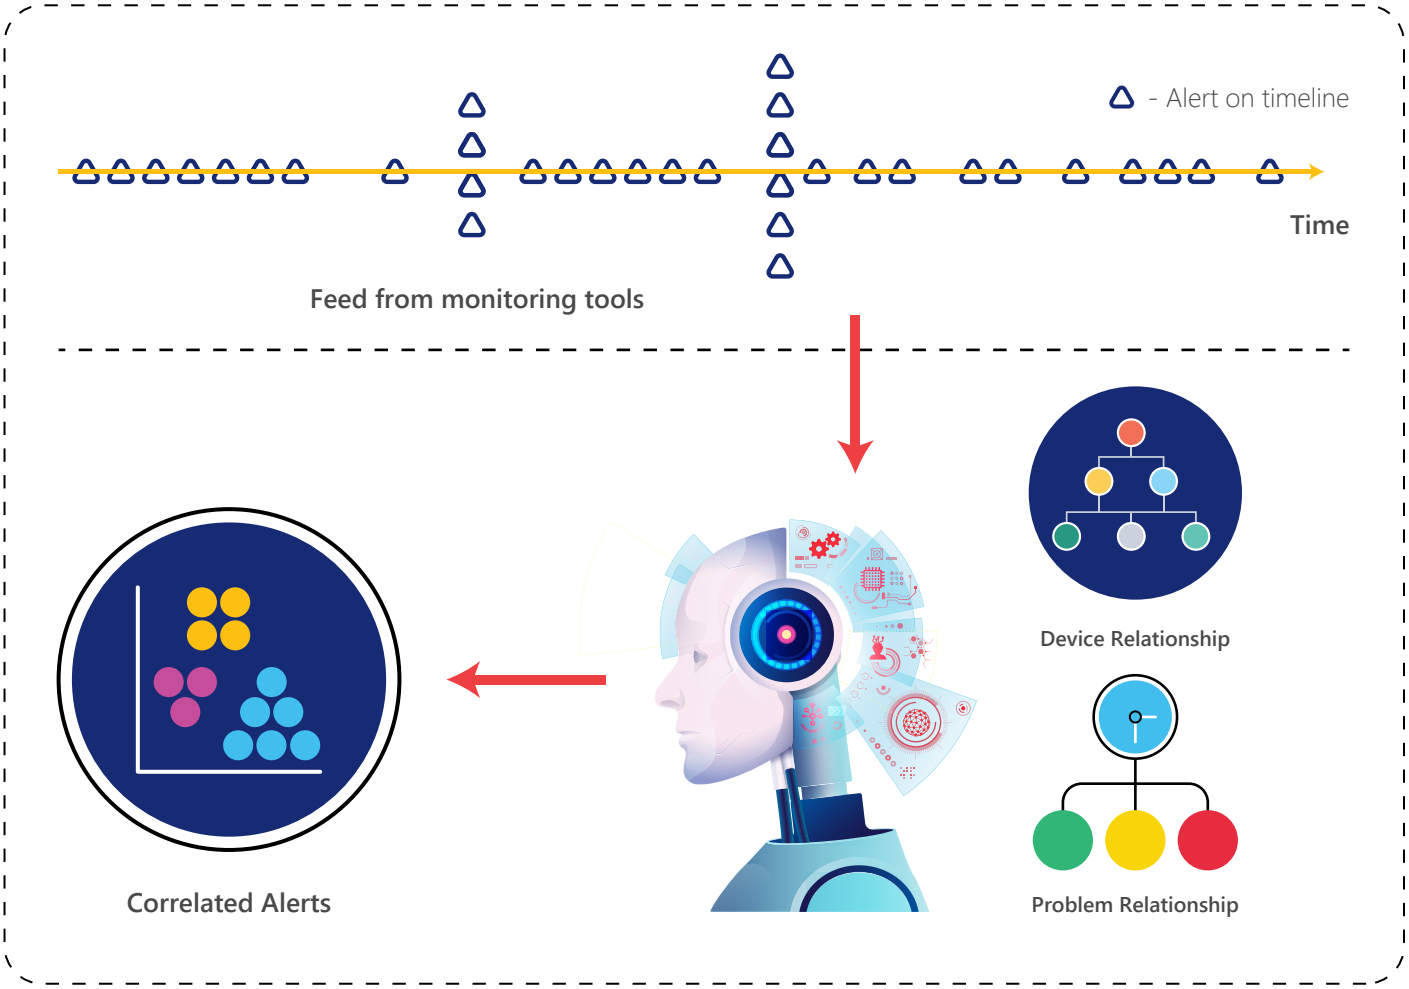

Business Benefits of Algorithmic Alert Correlation

This unique alert correlation approach is AI-enabled and holistic, and intelligently correlates alerts generated from entities that are physically, logically, or contextually related. This intelligence in alert correlation drastically increases the productivity of IT engineers and translates to significantly better revenue for organizations.

GAVS' AIOps Platform, Zero Incident Framework™ (ZIF™) implements such AI/ML algorithmic approaches for 360o alert correlation. ZIF filters out alert noise by identifying and eliminating duplicates & false positives. Noise reduction lowers the probability of missing a critical alert by at least 90%. With noise out of the way, ZIF aggregates actionable alerts, prioritizes them based on business impact. This drives precision in Root Cause Analysis (RCA), which in turn helps reduce mean time to resolve (MTTR) by at least 60%. With faster resolutions to IT incidents, business services gain predictable uptime and consistent performance.

The stand-alone capabilities of ZIF's alert correlation helps it function without dependence on external inputs and irrespective of access restrictions in the client environment.

Going beyond problem solving, ZIF continuously updates its knowledge base with contextual data to leverage for future decisions. This helps ZIF mature over time, with contextual learning and unlearning from more and more data, driving continuous improvements in business service performance.

Conclusion

Traditional alert correlation has not been able to scale up to handle the volume and complexity of alerts generated by the modern-day hybrid and dynamic IT infrastructure. We have reached a point where our ITOps needs have surpassed the limits of human capabilities, and so, supplementing our intelligence with Artificial Intelligence and Machine Learning have now become indispensable.

ZIF leverages AI/ML to derive its own relational mappings using event input data from monitoring tools, and hence does not need any access to the infrastructure components. This helps ZIF meet stringent security and regulatory compliance requirements. ZIF's alert correlation algorithms drive complete self-reliance and eliminate the dependence on external discovery or relationship data. This, we believe, is a truly unique and one-of-a- kind proposition in the AIOps industry.

ZIF

ZIF (**Zero Incident Framework™**), is an award-winning AIOps platform for IT Operations. ZIF delivers business outcomes by leveraging unsupervised pattern-based machine learning algorithms. Infrastructure and application telemetry data are aggregated, correlated, and potential failures are predicted. To enable faster resolution and better user experience, ZIF deploys intelligent bots for proactive remediation. Developed by GAVS Technologies (www.gavstech.com), ZIF is available as an on-premise and SAAS solution.

To find out how ZIF can help your organization, please visit **www.zif.ai** or write to **inquiry@zif.ai**