GIGAOM

Image credit: bannosuke

Andy Thurai
Apr 30, 2021

GigaOm Vendor Profile: Zebrium v1.0

An Exploration Based on Key and Radar; Cloud Observability

Cloud & Infrastructure

GigaOm Vendor Profile: Zebrium

An Exploration Based on Key and Radar; Cloud Observability

Table of Contents

1. Summary

Zebrium is an Observability/AIOps platform that uses unsupervised machine learning to auto-detect software problems and automatically find root causes, reducing manual labor and speeding incident response. The system requires no manual setup, instead training itself on patterns in logs and metrics to baseline the system, enabling the solution to be ready to perform incident and root cause detection within as little as one day. While most observability tools try to work the whole spectrum—from instrumentation to metrics to logs to incident correlation to root cause analysis—Zebrium concentrates on root cause identification using automated AI/ML to considerably reduce mean time to resolution (MTTR).

HOW TO READ THIS REPORT

This GigaOm report is one of a series of documents that helps IT organizations assess competing solutions in the context of well-defined features and criteria. For a fuller understanding consider reviewing the following reports:

Key Criteria report: A detailed market sector analysis that assesses the impact that key product features and criteria have on top-line solution characteristics—such as scalability, performance, and TCO—that drive purchase decisions.

GigaOm Radar report: A forward-looking analysis that plots the relative value and progression of vendor solutions along multiple axes based on strategy and execution. The Radar report includes a breakdown of each vendor's offering in the sector.

Vendor Profile: An in-depth vendor analysis that builds on the framework developed in the Key Criteria and Radar reports to assess a company's engagement within a technology sector. This analysis includes forward-looking guidance around both strategy and product.

2. Technical Differentiation

Zebrium works almost exclusively with logs and metrics (traces are not supported at this time) to conduct detection and diagnosis. It employs native collectors and fully supports Kubernetes (including variants like OpenShift), ECS/Docker and Linux. It also supports log collection via logstash or syslog for Windows, VMware, and most other environments. A Lambda function can be used to forward logs (for something like Amazon Cloudwatch). The Zebrium platform, which does not rely on sampled data, can handle large-scale data acquisition up to the petabyte level.

The platform can ingest and analyze any data, including unstructured data, which is particularly useful for logs. Its auto-ML functionality learns the structures of log events and looks for hotspots of abnormally correlated anomalous events to detect real incidents rather than just basic anomalies. By correlating logs and metrics from multiple sources, Zebrium can proactively create real incident reports, with details of a potential root cause, without manual intervention.

This automated root cause identification is Zebrium's most unique feature. Most observability solutions only identify anomalies, requiring a SRE to manually review, correlate and determine whether there is an incident or not. Zebrium is able to identify incidents without human interaction and automatically uncover root cause indicators to explain what happened.

Using Zebrium technology, typical customers can drive down the MTTR for a software incident from hours to minutes. Zebrium is delivered as a multi-tenant SaaS or can be deployed on a customer-owned virtual private cloud (VPC) or installed on-premises. It is free to try, and setup is fairly easy and quick. The SaaS solution is priced on data volume and the platform can also be installed on-premises for larger customers. Zebrium offers an augment mode that can be used in conjunction with existing log managers such as the Elastic Stack.

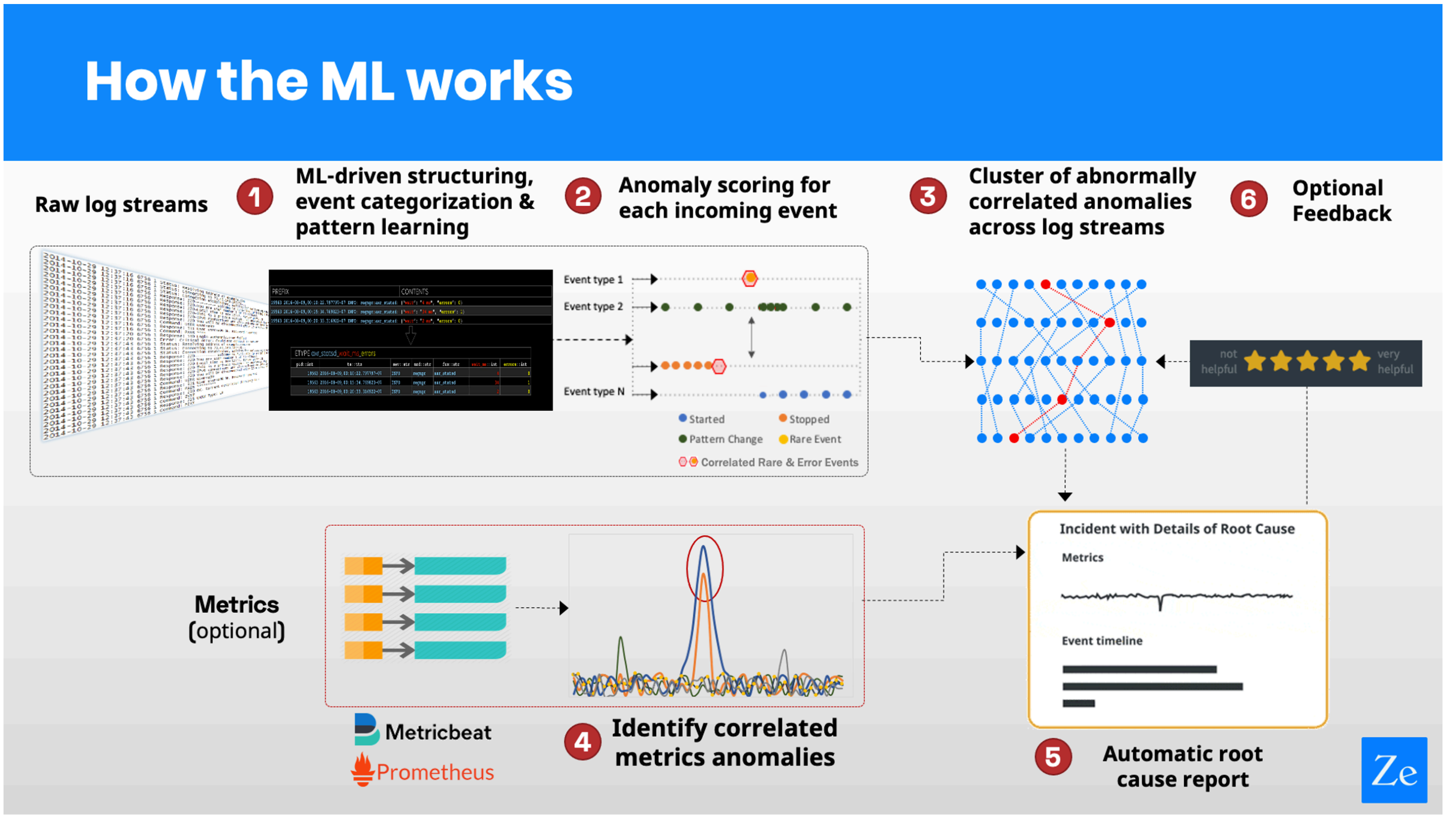
How the ML works
Raw log streams
1 ML-driven structuring, event categorization & pattern learning
2 Anomaly scoring for each incoming event
3 Cluster of abnormally correlated anomalies across log streams
6 Optional Feedback
Event type 1
Event type 2
Event type N
Started
Stopped
Pattern Change
Rare Event
Correlated Rare & Error Events
not helpful
very helpful
Incident with Details of Root Cause
Metrics
Event timeline
Metrics (optional)
Metricbeat
Prometheus
4 Identify correlated metrics anomalies
5 Automatic root cause report
Ze

Figure 1. The Zebrium Workflow

The Zebrium platform can work with any IT monitoring product currently on the market, plugging into its solution and automatically reporting problems and details of root cause to the NOC or SOC.

This is a great first step from an intriguing start-up—using AI/ML to differentiate real incidents from the anomaly "noise." Currently, it only uses logs, metrics and signals based on incidents or alerts, but Zebrium is planning to add other data inputs and signals in the future.

The solution is ML model driven and so it can easily adjust for seasonality and periodic changes without detecting them as an anomaly. The solution also has an Open API that allows for partner integration. Existing integrations include tools such as PagerDuty, Opsgenie, Slack and more.

The platform also includes a log manager. While not as sophisticated as some of the mature enterprise log offerings, it has adequate features and capabilities worth considering for small-scale operations. Notably, Zebrium offers a wide variety of integrations with existing enterprise log systems for existing users.

3. Business Impact

The most impressive feature about Zebrium is its Time-To-Value (TTV), especially when compared with other AIOps platforms. Deployment in cloud-native environments can be accomplished in a matter of hours, and the time to value can be achieved in a matter of days. This condensed time frame is possible because installing the Zebrium collectors is the only manual operation. The rest of the configuration is automatic, and the learning is unsupervised (although user feedback can be used to customize ML-generated reports). This level of sophistication speaks volumes about the power of AI and ML to enhance any technology. No extended POCs, no extended contract negotiations, no trial, and error to figure out which models will work. To prove this claim, they also offer "free POCs" which is very compelling. Their publicly available list pricing is one of the cheapest among the Observability vendors.

As a SaaS/multi-tenant solution, the software is updated automatically and the latest functions and features are available to the customers always. They also have private SaaS (VPC based), or On-prem versions for customers who are worried about security, compliance, and isolated environments.

Another interesting aspect is that there is no required pre-configuration of rules, parameters, and models. The models self-train in an unsupervised mode.

Cool Feature

Zebrium has added a feature that uses GPT-3 NLP processing for root cause summarization. While it is not needed to identify and explain the root cause of the incident, it is a welcome addition when teams are racing against time to identify the real cause of the incident. For example, in a recent demo, when the underlying PostgreSQL database was forcefully stopped, the Jira (Atlassian) app and related components threw thousands of error messages. While all these errors were critical and provided related logs/metrics information, by distilling them using Zebrium's ML and feeding them to the GPT-3 engine as a prompt, Zebrium was able to produce a cool message suggesting that the database server was stopped by the administrator – thereby identifying the real culprit! (**Figure 2**) This feature can not only help lower level support (L1/L2) teams identify problems quickly, it can also prevent L3, engineering, and DevOps teams from being dragged into the fray, which can save a lot of costs.

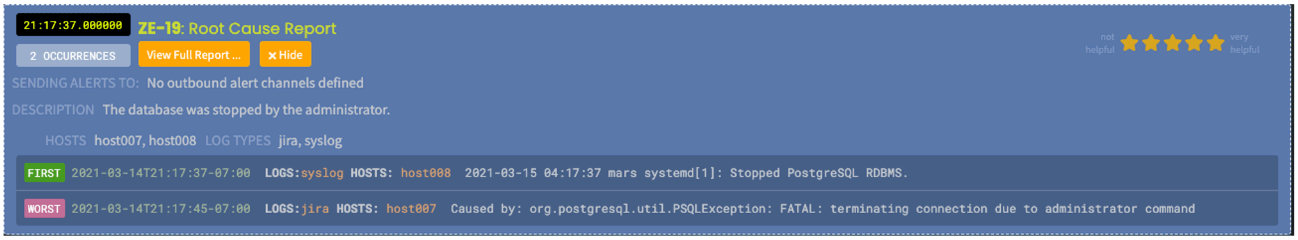

21:17:37.000000 ZE-19: Root Cause Report

SENDING ALERTS TO: No outbound alert channels defined

DESCRIPTION The database was stopped by the administrator.

HOSTS host007, host008 LOG TYPES jira, syslog

```
FIRST 2021-03-14T21:17:37-07:00  LOGS:syslog HOSTS: host008  2021-03-15 04:17:37 mars systemd[1]: Stopped PostgreSQL RDBMS.
WORST 2021-03-14T21:17:45-07:00  LOGS:jira HOSTS: host007  Caused by: org.postgresql.util.PSQLException: FATAL: terminating connection due to administrator command
```

Figure 2. Output from the GPT-3 Engine

Primary Use Cases

One compelling aspect of the Zebrium approach is its ability to work with existing telemetry, rather than set up yet another stack of monitoring/observability tools to capture data. Most enterprises are already well-instrumented enough that introducing a new toolset, with a new dataset, can actually delay the process of determining root cause instead of improving it.

Two use cases in particular stand out:

Automated Identification of Root Cause for Real Incidents
This enterprise use case has two components. The first challenge is dealing with all the noise in the signal–an issue particularly important for large enterprises. To help distinguish valid signals from noise, Zebrium looks for clusters of correlated log and metric anomalies across multiple sources, using its proprietary technology to identify real incidents.

The second aspect is integration with alerting tools such as PagerDuty, Opsgenie, Slack, and the like. Zebrium offers robust integration here, which allows the tool to pull together evidence supporting an issue investigation. When an alerting tool identifies a problem and notifies an incident, Zebrium can pick it up from the stream and its ML engine can then create an ML root cause report that is automatically added to the incident. So when an engineer/DevOps technician gets the alert, they will also receive a fully detailed RCA report that pinpoints the time of interest when the incident occurred without manual intervention.

Finding the Simmering Bug that Is Not Yet an Incident
When change cycles shorten, particularly in an agile environment, DevOps teams almost inevitably push faster changes to the production environment. While the specific unit of change itself is mostly tested, the interconnecting intricacies can cause disruption over time. What's worse, the changes may not trigger a detection by the installed anomaly or rules-based incident detection systems. But the change could be strong enough that system performance may degrade over time or become more disruptive as other changes are made.

The ability to flag simmering issues can be valuable where there are a lot of daily changes, with most of them flying under the radar without triggering an incident by other observability/monitoring tools. A daily/weekly/monthly proactive RCA report can be analyzed to detect the correlation of system behavior with specific changes (such as in config, code, system, infrastructure, and the like). This correlation detection can help identify latent bugs before they manifest as P1 incidents.

4. Bottom Line

In our research, Zebrium stands out as one of the most innovative solutions in the observability space. Instead of concentrating on instrumentation or data collection, the company opted to use the existing data and derive meaning out of it in the most automated way.

While the primary use case is quick incident root cause identification, the secondary use case (proactive detection) is very compelling as well—unearthing underlying problems that have not yet made it to the anomaly level. Bugs, misconfigurations, and other issues can smolder undetected if the system changes that trigger them are not significant enough to deviate heavily from the baseline. These issues can be hard to detect and can deteriorate over time. By helping analyze the newer observable data against the original baseline before change, Zebrium helps identify the hidden, slow-simmering problems that might become an issue later.

As an innovative player in the observability market, Zebrium is using AI/ML capabilities to the fullest extent to solve the “fastest time to root cause” problem in an automated way. It is a refreshing approach compared to the complex, old-school way in which teams explore metrics to see if something went wrong, then dig into traces to see where things went wrong, and finally venture into logs to determine what it was that actually went wrong.

If you are in the market for an observability solution, and one of the above use cases describes your primary need, Zebrium is worth a look.

5 About Andy Thurai

Andy Thurai is an accomplished IT executive, strategist, advisor, and evangelist with 25-plus years of experience in executive, technical, and architectural leadership positions at companies such as IBM, Intel, BMC, Nortel, and Oracle. He also advises many start-ups. He has been a keynote speaker at major conferences and served as host for many webcasts, podcasts, webinars, and video chats. Andy has written more than 100 articles on emerging technology topics for publications such as Forbes, The New Stack, AI World, VentureBeat, and Wired magazine.

Andy's topics of interest and expertise include AIOps, ITOps, observability, artificial intelligence, machine learning, cloud, edge, and other enterprise software. His strength is selling technology to the CxO audience with value proposition rather than a technology pitch.

You can find more details and samples of Andy's work on his website at www.thefieldcto.com

6. About GigaOm

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.

7. Copyright