ISSN: 2634 - 8853

Journal of Engineering and Applied Sciences Technology

SCIENTIFIC Research and Community

Review Article

Open Access

Automated Root Cause Analysis with Observability Data - A Comprehensive Review

Ankur Mahida

Site Reliability Engineers (SRE), Barclays, USA

ABSTRACT

Identifying the root cause analysis is the key to the timely detection of errors in massive, multiple-functional software systems. Meanwhile, network development will become more intricate and non-transparent, leaving the human algorithm behind. The paper dedicates its resources to discussing ways to automate root cause analysis based on observability data, such as logs, metrics, and traces. Technologies including causal inference, anomaly detection, and pattern recognition are specialized techniques that allow us to identify the breach in the background of thousands of connected events. Data-driven tools in research and industry that consume observability data as input, uncover anomalies, model system topology, and rank probable root causes use these technologies. This should give the customer a shorter mean repair time, higher reliability, and security. Deeper adoption is perfect for chain management and the performance gains of individual developers. The coverage includes strategies of algorithmic and implementation of root cause analysis with the observability data. Related topics like service maps and anomalous numbers are also discussed where necessary. Scaling out the automatic diagnosis entity is investigated as an automated means to replace the manual ones faced with elaborate models.

***Corresponding author**

Ankur Mahida, Site Reliability Engineers (SRE), Barclays, USA.

Received: December 05, 2023; **Accepted:** December 09, 2023; **Published:** December 18, 2023

Keywords: Root Cause Analysis, Observability, Causality, Anomaly Detection, Reliability

Introduction

Currently, there are often too many interlinked, distributed services in the software systems, and the failures usually spread very complexly. Manual root cause analysis is not practical for large systems to grow. Automated methods can digest logs, metrics, and traces to discover the cause of the failure amid an enormous number of activities [1]. Mainly, the paper covers scientific and industry methods of causal inference, anomaly detection, and pattern recognition that are applied to observability data. These techniques support system topology building, anomaly detection, signature comparison, and failure root cause analysis. The automation of root cause analysis promises faster case-solving and high reliability. Within algorithms and implementations reside the tools to identify problems in complex, contemporary software where mechanical diagnostics is impractical. When linked, subjects like service dependency mapping are treated accordingly. Automated root cause analysis using log information is used to accelerate operations, improve security, and boost developer productivity.

Problem Statement

Modern software systems today grow in size, complexity, and reliance upon the combined efforts of multiple sets of services and components. The increasing complexity of the systems renders the performance of prompt root cause analysis incredibly challenging during times of problems [2].

Firstly, the huge volume of observability data that emerged from complex systems surpasses the power of the human brain to master and comprehend. A central aspect of distributed architecture is the significant number of logs, metrics, and traces resulting from all system components. Handling through the manual way of limiting the volume of this high-volume, high-cardinality dataset to identify the root problem is not a function.

The data used for observation is collected from multiple sources, meaning they must be aligned together carefully to reconstruct the entire sequence of events that led to the failure. Logs, metrics, and traces offer comprehensive visibility into the behavior of the different systems, where the correlation of traces among humans would be a tough one [3]. Uniting fragments of the process across services is a complex thing to do without a proper automated approach.

Another decisive factor is the growing difficulty in isolating the actual causation and mere correlation of events from the intertwined parts that work together when conducting investigations. Simultaneously, coincidences and manifold interactions formed a concealing screen for the leading factor. A manual search for connections from multiple causal relations one after another is not only time-consuming but can also lead to wrong conclusions.

In the concluding phase, forming accurate models of normal behavior from aggregating observability data might be costly, which pertains to some initial expenditures [4]. With systems constantly changing and new such things as anomalies regularly arising, automated AI services are required to keep models updated time-wise. Distinguishing pathologies is more challenging than expected because a baseline is missing in the current situation.

Solutions

To address the challenges of manual root cause analysis, automated techniques apply innovations in causal inference, anomaly detection, pattern recognition, probabilistic modeling, and knowledge graphs:

The causal inference algorithms may disentangle cause from the connection of data traces and logs, and this machine learning approach overcomes the correlation ambiguity the network systems offer [5]. Through time-corresponding methods, these techniques can accurately restore a series of actions and distinguish genuine reasons from many related factors.

Anomaly detection identifies the deviation of behavior from their average metrics with the help of machine learning [6]. Interpolating the real-time profiles of expected performance enables the detection of the variances, which can be considered as conditions worth further examination. This allows the expert to concentrate on complex system baseline modeling without using automated modeling.

Pattern recognition and clustering similarities detect and resolve the most frequent patterns of failure and issues by comparing log entries and events with previously identified ones and playbooks [7]. This is real of logging into problems that frequently occur to diagnose them and fix them in a shorter time. Moreover, it straightens out the redundancies, grouping most similar failures for most people.

These are the models that can weigh as well as integrate various types of evidence to determine the probability of the causes. They can code the dependencies involving the weight and sustain the proof, and thus, they can infer the possible root cause based on the probability. This contributes to the class of recidivistic mathematical tools.

This type of graph embodies the associations between system functions and services in the form of intelligent top-down tracing. The behavior can be enacted by using modeling topology and relationships. The exploration can be done faster by the analysis engine for path identification.

From these techniques, automated systems could provide ingesting observability data, find anomalies, perform topology modeling, pattern matching, and rank viable root causes. Featured attributes include cross-verification of different evidence sources, comparison with historical baselines, traversing dependencies, and clocking of most probable causes - one of the significant enablers of automatic anomaly detection in clouds.

Uses

Automated root cause analysis offers substantial benefits for diagnosing and resolving several vital classes of issues in complex software systems:

Performance Problems

These automated analyses are based on the principles of rapid issues identification and the detection of such faults as slowdowns or bottlenecks that cause the degradation. This may result from soft resource contention appearing behind the scenes or in the form of issues in many services collectively. Analyzing metrics and traces to single out anomalies and causation of any nature is possible, and it can ensure that a complex process like performance regression is due to some poor coding, config changes, or infrastructure [8]. This results in remediation involving less time and avoids a series of slowdowns, which decrease the user's experience of turning easy mode on.

Failures and Outages

When these computer-related incidents happen in real time, application crashes, service outages, and errors create significant availability and reliability issues. However, the earliest fault may be covered up by numerous false alerts and flashing lights that warn of different chains of events. Deductive root cause analysis can ingest vast amounts of observability data and segue this upstream to identify the component failure or resource exhaustion that immediately triggers the outage. Our solution bridges the gap from hours to minutes in a real-time troubleshooting environment of even the most critical reliability events.

Security Incidents

Attempted intrusion, hacking, and breaches also conduct large streams of harmful logs, making it impossible to identify the initial point of intrusion [9]. Correlation and causation are determined automatically by integrating all available information sources, which speedily solves the problem of investigation. Thoroughly verifying the route taken by the intrusion to stop at the precise spot of the actual beginning of the incursion is the primary and vital element for remediation and forensics. This ability earns the data stores a higher security level for cyber security, allowing businesses to return to normal more quickly and minimize loss.

The implication is that with automated root cause analysis, those insights from various logs, metrics, and tracepoints are combined with pinning down the actual origin point among several prominent correlated factors that may have thousands of such points. Automated simulations follow the topology and obscure the underlying causes otherwise hidden in more tedious handwork investigations. Through this, what can be diagnosed and fixed can be done faster and more accurately, thus keeping system availability at an acceptable level by reducing downtime. The abilities in these two areas directly relate to the reliability, security, and complex modern software environments where manual analysis is only sometimes available.

Impacts

The increasing adoption of automated root cause analysis stands to provide a significant impact on reliability, operations, security, productivity, and customer experience:

Reliability

As it allows for identifying and treating the main reasons underlying system service unavailability, automated root cause analysis is a significant factor contributing to system reliability and availability [10]. Firstly, it accelerates failure detection and mitigation, preventing more significant power outages from cascading from these minor faults. The proximity to the errors allows the engineers to quickly track them to their source – whether in code misconfiguration or infrastructure – resolving the matter before the errors propagate. In addition, even during the case of a sudden severe outbreak like the coronavirus pandemic, automated diagnosis lets localize it to a few people in just a few minutes from hours. Surgically pointing to the actual cause of failure among the massive alerts is a way to grid through the crowds of places where the out-of-band effects happen. Among these capabilities are the quick recovery and the downtime, driving the reliability lower. Modalities experiencing fewer operational and performance failures are achieved due to the direct automated fixes of identified errors.

Operational Efficiency

Automation of root cause analysis also makes the processes faster due to decreasing the redundant firefighting actions and fault finding, which has a lot of downtime [11]. The staff saves time sifting through the piles of data-rich observability platforms, which helps them be more proactive in the work process. The system automatically investigates and tackles day-to-day matters, allowing for the hallmarking of the scintillating potential of the essential personnel to do themselves. The need for repetitive manual reports and alert triage will drop significantly as data is validated more precisely and problems are identified correctly. SMEs will be able to devote their time to something strategic. Tackling problems manually will be less of a part of their routine. Operations teams no longer need to be preoccupied with the complex situations arising in the manual diagnosis process, where they need to investigate the root causes because of the continuous engagement of automated solutions.

Security

Similarly, automation in the cyber incident response process helps trace attack pathways so that the malicious compromise is isolated [12]. Reliably merging data fragments from multiple unrelated cybersecurity devices in a mechanical transmission process is just superhuman cognizance. Machine learning based on comparison and cause-and-effect detection of different data sources - firewall logs, IDS alerts, and DNS requests - permits much quicker investigation. Even as the attackers gain root access or elevated privileges in the network by creating backdoors to remain undetected, active tracing guides containment, remediation, and forensics investigation well. Automation takes situational awareness away from the overwhelming complexity of the manual investigation and accelerates the security events' resolution time frame.

Developer Productivity

Developers gain an advantage in the debugging stage since debugging efficiency is improved through collecting and linking of appropriate runtime arguments. Traces, logs, and metrics come out of nowhere, so they help deal with code snags faster and easier [13]. Instead of a breadcrumb trail that is traveled manually, developers can visualize a whole complicated speck from various data sources embedded in the code. The integration within IDEs also allows the automation to become closer to the developer's workflows without disturbing the context of his work. Through an accelerated automatic diagnosis, devs spend less time debugging, and they overbuild.

Customer Experience

Lastly, employing such an approach contributes to the continuous running of the system. Rapid identification and solution prevent the end user from being impacted in case of an issue. As complications happen faster, applications generally encounter fewer unscheduled downtimes, and this is due to quick automated remediation of errors and performance lag, which keeps users from noticing any difficulties in services or applications [14]. Even faced with bewildering intricacy, automation ensures an identical level of quality and exhaustive promptness in service and reaction. Thanks to the algorithm proposed quick detection of the root problems, customers are protected from going under caused by the distributed failures. Systems that are dependable enough to work flawlessly offer customers more timeframes and an improved transactor platform.

Scope

This paper will present causal inference and pattern recognition techniques that promise to automate root cause analysis using observability data. The range covers the academic research, which progresses the development of these methods, and the industry tools, which apply them in practice. The article focuses mainly on algorithms based on logs, metrics, and traces, providing information about the system's operation and problems. Such topics as anomaly detection and service topology modeling are also included when they can be applied to make the automated diagnosis more meaningful. Other observability data types, such as synthetic monitoring and user sessions, are left out.

Furthermore, only security, performance, and reliability use cases are addressed, leaving others for further examination. This focus still centers on innovations in causality inference and pattern matching to the logs, metrics, and traces to automate identifying the root issues within the complex modern software systems where the manual analysis is failing to scale. Reference is made to related domains to provide context, but they are not dealt with totally.

Conclusion

As the systems grow with increasing complexity and interconnection, manual root cause analysis becomes impossible. By automating analysis with the help of monitoring data, we define the reliability and the efficiency. This paper extensively discusses scholarly articles and industry tools based on enhancements such as causal inference, anomaly detection, and event matching on the logs, metrics, and traces. With these methods, it is possible to simulate topology, detect deviations, and pinpoint the reasons for breakdowns in modern distributed systems where manual investigations would not scale up. Adoption ensures problems are solved quickly, 24/7 availability, a more robust security level, and a higher productivity level of the organization. Even if issues still exist, automated root cause analysis using observability data is the key to operating complicated software systems beyond the skill of manual techniques.

References

1. Chang Y, Iakovou E, Shi W (2019) Blockchain in global supply chains and cross border trade: a critical synthesis of the state-of-the-art, challenges and opportunities. International Journal of Production Research 58: 1-18.
2. Gan Y, Mingyu Liang, Sundar Dev, David Lo, Christina Delimitrou (2019) Seer: practical and scalable ML-driven performance debugging in microservices. Architectural Support for Programming Languages and Operating Systems 135-151.
3. Las-Casas P, Giorgi Papakerashvili, Anand V, Mace J (2019) Sifter: Scalable Sampling for Distributed Traces, without Feature Engineering. Proceedings of the ACM Symposium on Cloud Computing 312-324.
4. Kazmi H, Suykens J, Balint A, Driesen J (2019) Multi-agent reinforcement learning for modeling and control of thermostatically controlled loads. Applied Energy 238: 1022-1035.
5. Balayn A, Lofi C, Houben GJ (2021) Managing bias and unfairness in data for decision support: a survey of machine learning and data engineering approaches to identify and mitigate bias and unfairness within data management and analytics systems. The VLDB Journal 30: 739-768.
6. Pang G, Shen C, Van Den Hengel A (2019) Deep Anomaly Detection with Deviation Networks. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining 353-362.
7. Gupta R, Tanwar S, Tyagi S, Kumar N (2020) Machine Learning Models for Secure Data Analytics: A taxonomy

and threat model. Computer Communications 153: 406-440.

8. Farshchi M, Schneider JG, Weber I, Grundy J (2018) Metric selection and anomaly detection for cloud operations using log and metric correlation analysis. Journal of Systems and Software 137: 531-549.
9. Israa Ezzat Salem, Maad Mijwil, Alaa Wagih Abdulqader, Marwa M Ismaeel, Anmar Alkhazraji, et al. (2022) Introduction to The Data Mining Techniques in Cybersecurity. Mesopotamian Journal of Big Data 28-37.
10. Zhang Y, Zhengxiong Guan, Huajie Qian, Leili Xu, Hengbo Liu, et al. (2021) CloudRCA: A Root Cause Analysis Framework for Cloud Computing Platforms. Proceedings of the 30th ACM International Conference on Information & Knowledge Management 4373-4382.
11. Javaid M, Haleem A, Pratap Singh R, Rab S, Suman R (2021) Upgrading the manufacturing sector via applications of Industrial Internet of Things (IIoT). Sensors International 2: 100129.
12. Ioannis Stellios, Panayiotis Kotzanikolaou, Mihalis Psarakis, Cristina Alcaraz, Javier Lopez, (2018) A Survey of IoT-Enabled Cyberattacks: Assessing Attack Paths to Critical Infrastructures and Services. IEEE Journals & Magazine 20: 3453-3495.
13. Parker A, Spoonhower D, Mace J, Isaacs R, Sigelman B (2020) Distributed tracing in practice: instrumenting, analyzing, and debugging microservices. O'Reilly Media https://www.oreilly.com/library/view/distributed-tracing-in/9781492056621/.
14. Wijethilaka S, Liyanage M (2021) Survey on Network Slicing for Internet of Things Realization in 5G Networks. IEEE Communications Surveys & Tutorials 23: 957-994.