splunk> + VictorOps

Optimizing Incident Response

Through Automation & Machine Learning

VictorOps is purpose-built incident management software for DevOps-focused teams. Sign up for a 14-day free trial to start making on-call suck less.

TABLE OF CONTENTS

What is Machine Learning?

Machine learning and automation are revolutionizing the way IT teams approach myriad tasks. Incident response is no exception. By providing new opportunities for addressing the inefficiencies that have traditionally plagued incident response teams, machine learning is making it much easier to resolve more incidents, at greater speed, and with less effort on the part of team members.

Indeed, just as the rise of technologies like cloud computing and containers over the past decade has allowed IT teams to deploy applications much more efficiently and with less manual effort, the pairing of machine learning with automation is emerging as one of the next great opportunities for optimizing a task – incident response – that has conventionally demanded a great deal of manual effort, and was prone to oversights and inconsistencies borne of human error. Not only that, but given the ever-increasing complexity of modern software environments, machine learning and automation are becoming must-haves for helping IT professionals to see through all of the noise, and approach incident response effectively.

Yet, like most innovative IT techniques, automation and machine learning require careful planning in order to deliver benefits for incident response. Teams must understand the specific problems that automation and machine learning help to resolve within the discipline of incident response, and they must learn best practices for applying these technologies to incident response operations.

With those needs in mind, this whitepaper details why and how IT teams should take advantage of machine learning and automation within the context of incident response. It explains the efficiencies and opportunities that these techniques introduce; then discusses the specific within the realm of incident response to which they can be applied.

The Meaning of Automation and Machine Learning for Incident Response

Before delving into a discussion of the problems that automation and machine learning solve within the context of incident response, a concise definition of these terms is in order.

Incident response refers to the processes and tools that IT teams use to identify, interpret and resolve performance or availability issues with hardware and software. For years, incident response has been a critical component of delivering reliable, high-performing application experiences to end users – although, as this whitepaper explains, new technologies are fundamentally changing the way IT teams approach incident response.

Machine learning is a technique that leverages artificial intelligence, or AI, to interpret data and make decisions based on it. When combined with automation, which makes it possible to execute actions automatically and systematically, machine learning eliminates the need for IT professionals to interpret complex data sets manually – a process that is slow and often inconsistent when performed by hand, since different engineers may interpret the same data in different ways.

Challenges for Efficient Incident Response

Traditionally, IT teams response for performing incident response have faced several challenges that make it difficult to identify, interpret and resolve incidents quickly and efficiently.

Determining Who Should Respond

Perhaps the greatest efficiency challenge for incident response is determining which individual engineer or engineers should respond to a given incident. This is a challenge not only because it requires identifying which staff member has the necessary skills to resolve an incident, but also because the decision must reflect engineers' availability. The most qualified engineer to handle a given incident might not be on call when the incident arises, and incident response strategies that do not take staff members' availability into account run the risk of undercutting morale.

For most teams, this has been a difficult challenge to resolve in an efficient way. Spreadsheets or calendars go some way toward coordinating engineers' availability and skills with incidents, but they are not a fully-automated solution.

Detecting Incidents

Teams have long relied on automated monitoring tools to help detect incidents. Sometimes, those tools use machine learning to help analyze data from hardware and software environments to make a determination about whether certain behavior reflects an incident that the IT team should investigate and address.

However, given the complexity of modern software environments, traditional techniques for detecting incidents often no longer suffice. Today, applications are commonly deployed on distributed architectures that span multiple servers, some of which may exist on-premise, while others exist in the cloud. Networking configurations typically involve multiple layers of internal and external-facing endpoints. Applications are frequently packaged inside containers, which add another layer of complexity for IT teams to manage.

Given such complex infrastructure, detecting incidents using traditional techniques has become very challenging, because there is simply so much data to parse.

For most teams, this has been a difficult challenge to resolve in an efficient way. Spreadsheets or calendars go some way toward coordinating engineers' availability and skills with incidents, but they are not a fully-automated solution.

Analyzing and Responding to Incidents Quickly

Given infinite time to respond, most IT teams could resolve every incident that arises. But in real-world production environments, of course, time is of the essence when identifying and responding to incidents. Without automated, AI-driven techniques for understanding incidents and coordinating who will respond to them, understanding the nature of an incident and responding to it in a timely fashion is very difficult.

Understanding Relationships Between Incidents

Determining whether and how one incident relates to another is critical for responding efficiently. In some cases, multiple symptoms may result from the same underlying cause. For example, a decrease in response time from two applications might be the result of a problem with a network load-balancer that supports both applications. Or, each application could be failing for separate reasons, requiring a different remediation solution for each.

To assess the relationship between surface-level incidents, IT teams must interpret large amounts of data and consider how the multiple layers of application environments interact. This is a difficult task to perform quickly and consistently, especially in today's highly complex microservices environments.

Feedback Loops

Achieving effective incident response over the long term requires not only responding to issues when they arise, but also communicating information back to developers and system architects so that problems can be prevented from recurring. Providing this type of feedback can be difficult, especially within larger organizations where developers, system architects and incident response teams do not work in close collaboration.

Consistency in Incident Response

Organizations also commonly suffer from the challenge of keeping incident response practices consistent, so that all team members approach incident analysis and resolution in the same way. Consistency is important for ensuring that incident response is predictable and systematic. It also helps to meet special requirements, such as those imposed by compliance rules that mandate certain procedures to be followed when working with a given system.

Achieving this type of consistency is difficult, particularly within larger organizations that maintain multiple systems. Different teams might be responsible for incident response within different systems. Responsibility for incident response may also be handed off from one team to another when, for example, on-call schedules change. Homogenizing the behavior of so many individuals and teams working on multiple systems is difficult.

Applying Machine Learning to Incident Response

By embracing incident response tools and strategies that leverage machine learning and automation, IT teams can resolve the challenges described above and achieve a faster, more predictable and more efficient approach to incident response.

Assigning Incident Response Duties

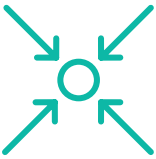

First and foremost, machine learning can automatically make suggestions regarding which engineers should respond to an incident on the basis of the nature of the incident; these suggestions are then assessed based on staff members' expertise and availability.

By using machine learning to make these recommendations automatically, IT teams no longer have to puzzle over spreadsheets and calendars in an effort to figure out who is best fit to handle a given issue, and whether that person is actually available. Nor do engineers have to worry about being asked to respond to an incident when they are not on call, due to their team's inability to align response needs with on-call scheduling.

Faster, More Accurate Incident Detection

Machine learning simplifies and automates the process of determining which behavior within an application environment constitutes an incident. It can also help to categorize the seriousness of the incident and make recommendations about which problems IT teams should prioritize. With machine learning, no amount of data becomes too much to interpret when identifying and assessing incidents within highly complex environments.

Categorizing Related Incidents

Likewise, machine learning can help IT teams parse through all of the data generated by their environments in order to understand whether and how two or more incidents might be related to each other. No longer do engineers need to assess complex data sets manually or parse complex configuration files, to assess whether multiple issues share a common underlying cause.

Automated Remediation

Although incident response will always require some level of human oversight and intervention, machine learning is opening up new opportunities to automate remediation entirely. By using machine learning to analyze the cause of an incident and determine how to resolve it, tools can take action automatically, especially in cases where the necessary fix is relatively simple and easily automated. For example, a failed container instance could be automatically restarted, or a database that is running out of disk space could be migrated to different infrastructure with more storage.

Unifying Incident Response Strategies and Techniques

When machine learning and automation become the basis of incident response, all teams gain a common approach to resolving incidents, no matter which systems they are working with. In this way, machine learning and automation help to achieve consistency within incident response.multiple issues share a common underlying cause.

Identifying Opportunities for Long-Term Improvement

Although incident response will always require some level of human oversight and intervention, machine learning is opening up new opportunities to automate remediation entirely. By using machine learning to analyze the cause of an incident and determine how to resolve it, tools can take action automatically, especially in cases where the necessary fix is relatively simple and easily automated. For example, a failed container instance could be automatically restarted, or a database that is running out of disk space could be migrated to different infrastructure with more storage.

Conclusion

Incident management has long been a process that is time-consuming and rife with inconsistency. It has also typically imposed an outsized burden on IT engineers, whose ability to enjoy time off has often been compromised by the need to respond to incidents.

With machine learning and automation, however, a better world is possible. Not only can incident response staffing be overhauled in a way that is more efficient and convenient for engineers, but the insights required to understand and address incidents – over both the short term and the long term – become much easier to gain with the help of machine learning and automated analysis.

In today's world, defined by complex technologies like containers and cloud-native architectures, machine learning has become an essential component of effective incident response. It is the only way for IT teams to remain effective in their incident response activities, and for businesses to retain an edge over competitors by minimizing downtime and keeping IT operations efficient and agile.

Chris Tozzi

(Journalist and Linux System Administrator)

Chris Tozzi has worked as a journalist and Linux systems administrator. He has particular interests in open source, agile infrastructure and networking. He is Senior Editor of content and a DevOps Analyst at Fixate IO. His latest book, For Fun and Profit: A History of the Free and Open Source Software Revolution, was published in 2017.