

From Chaos to Clarity Leveraging GenAI to Amplify Machine Learning in Observability

Cracking the code of unpredictability
with smarter AI

splunk>
a CISCO company



The crime

A crime is committed. And you saw it happen. In fact, there are many witnesses in different locations with different assumptions and unique suspects. It's chaotic, with clues scattered everywhere. Which ones matter and which are irrelevant?

The crime? This time, it's a database outage. Last time, a failed Kubernetes node. And the time before that? An application server in a race condition. These tech "crimes" are common. They jeopardize uptime and performance. Like a fingerprint, all crimes are unique in some way — almost as if a master villain were orchestrating the chaos through digital sabotage.

Every observer sees something unique based on their role, and each may solve these tech "crimes" differently. The Site Reliability Engineer (SRE) ensures systems are up and running. The Platform Engineer focuses on keeping underlying platforms performing optimally. The Software Engineer tries to build the best and highest-performing apps. So, when you ask them what happened, they may each have distinct points of view.

IT and engineering issues need prompt resolution when an organization's health, reputation, and digital resilience are at stake. Outages, application slowdowns, and even too many alerts can directly impact the bottom line, affect the brand, and make or break customer loyalty. Every minute it takes to uncover clues, solve the mystery (aka resolve the issue), and meet SLAs, can hurt a business and ruin customer experiences.

So, where do you turn? Forensic tools exist that can see through the chaos. Observability solutions can help engineers and IT professionals pinpoint potential causes.



Different experts bring specialized toolsets, perspectives, skills, and approaches tailored to gain insights and chase clues. Issues need to be solved quickly, and accurately. When experts use different data and different tools to solve the same crime, resolving their points of view can hinder resolving the root cause of the problem.

84% of respondents have explored generative AI within observability platforms, specifically for chatbots and embedded AI assistants. They envision using these capabilities for data analysis (66%) and recommendations to resolve issues (60%).

— [State of Observability 2024](#)



If this were an actual crime, you might turn to the authorities to solve it. To expedite an investigation, you may enlist the help of a detective agency specializing in that particular crime.

Ready for a fictional adventure in crime solving? Let’s travel back to the late 1800s in London, England, and the world of famous author Sir Arthur Conan Doyle. When the authorities were stumped on a masterful crime, usually committed by the evil Professor Moriarty, they enlisted the help of a detective agency of two primary employees: Sherlock Holmes and Dr. John H. Watson.

So, what do Sherlock, Dr. Watson, and Moriarty have to do with observability? A lot, actually. Especially if you think about things in these ways:

- Moriarty, as the chaos within observability,
- The detective agency, as the assistance of Artificial Intelligence (AI) and,
- Sherlock and Dr. Watson, as powerful, complementary aspects of AI — machine learning (ML) and generative AI (GenAI), respectively.

As we navigate the shadows of observability chaos, we’ll uncover the challenges engineers, application developers, and IT professionals face and explore how new AI technologies help solve IT’s daily mysteries.

 Chaos within observability (Professor James Moriarty)	 Artificial Intelligence (Detective Agency)	
	Organizing chaos through machine learning (Sherlock Holmes)	Simplifying through GenAI (Dr. John H. Watson)
Professor Moriarty represents chaos within observability. While you can collect and observe telemetry data, false positives, performance degradations, code changes, failing infrastructure, and unexpected outages disrupt the business. Coupled with overwhelming data volume, fragmented logs, and the complexity of modern systems, Moriarty throws organizations into chaos.	Sherlock’s uncanny ability to recognize patterns in the chaos and filter through the noise is essential to solving any crime. ML can parse through massive amounts of data to detect anomalies, reduce alert noise, and make complex analytics more understandable. Sherlock is an expert at finding clues that others would miss.	The crime-solving of Sherlock wouldn’t be complete without the insights into human nature of Dr. Watson. Through GenAI, complex insights become more accessible when you can ask questions in plain English, expediting resolutions for a broader set of “victims.” GenAI provides the ability to go beyond the raw details, navigate the patterns, and craft understandable summaries of how to solve crimes.

The chaos of observability and the influence of Moriarty

Digital environments have become a cacophony of data points and metrics, where even the sharpest minds struggle to find problems within observability chaos. Observability experts can become confounded by metrics, logs, events, and traces and struggle to identify the root cause of the chaos.

This is where Professor Moriarty triumphs.

ITOps and Engineering teams face various observability challenges, such as alert fatigue and data overload. Moriarty makes managing multiple systems overwhelmingly complex. Here are a few of his favorite methods:

Alert fatigue – he generates a barrage of alerts to numb even the best teams, causing critical signals to be lost in the noise.

- **False positives** – he delights in sending false signals — alerts that cry wolf when there is no danger.
- **Cascading alerts** – he understands the power of a chain reaction where a single disruption triggers alert streams across multiple interconnected systems, obscuring the root cause.
- **Flapping alerts** – he knows how to turn alerts on and off, perhaps where thresholds are incorrectly defined, to create a relentless loop of uncertainty, distracting from actual issues.
- **Duplicate alerts** – he amplifies noise and obscures genuine signals by echoing the same alert multiple times across different tools.

Data overload – he floods engineers with an avalanche of telemetry data across dispersed systems, hoping to drown them in information without insights.

Professor James Moriarty

THE CHAOS WITHIN OBSERVABILITY



ALIAS The Master of Chaos

ROLE Chaos Personified in Observability

STRENGTHS

- **Data Disruptor:** Thrives on fragmented logs, noisy data, and unmanageable telemetry.
- **Alert Overloader:** Orchestrates alert fatigue with false positives, cascading alerts, and data overload.
- **Master of Disguise:** His disruptions take many forms — unexpected outages, slowdowns, and performance degradation.

PERSONIFICATION The relentless villain that throws unpredictability at digital environments, leaving teams struggling to identify and distinguish relevant signals from the noise.

SPECIAL ABILITIES

- **Alert Flood:** Generates flapping, false positives, and duplicate alerts to confuse engineers.
- **Misdirection:** Creates performance issues that look like multiple root causes, making it hard to pinpoint the real problem.



FAVORITE QUOTE Chaos is the perfect smokescreen. Who can solve a mystery when they can't even see what's real?

Uncovering the root cause of a particular issue often requires advanced expertise and an intimate understanding of a tool or suite of solutions. This domain know-how is often scarce and expensive.


66% of engineers experience burnout or missed issues due to alert fatigue.

— State of Observability 2024


As chaos grows, traditional approaches to observability may not be enough. Teams need to filter the noise, detect anomalies, and swiftly identify root causes before the business is hurt.

Master the art of alert management and equip ITOps teams, engineers, and managers with tools needed to weather alert storms.

5 Ways to Avoid Alert Storms



Machine learning (ML) and generative AI (GenAI) bring order to the chaos of observability by revealing hidden patterns and transforming them into actionable insights. ML excels at uncovering patterns within vast amounts of data, such as anomaly detection. GenAI makes these findings accessible and easy to act upon. Together, they help engineers see through the noise, maintain stability, and achieve digital resilience.



CASE #1

The mystery of the failing Kubernetes node

SCENARIO

A major e-commerce platform is experiencing sporadic performance degradation, and several critical services are running slowly. After much manual investigation, the team suspects the issue lies with one of their Kubernetes nodes, but they can't identify the root cause.

CHALLENGE

The team is overwhelmed with data from multiple clusters, and they're unsure if the problem is infrastructure-related, caused by resource limitations, or something more subtle. False alarms are making it difficult to pinpoint the real issue, and valuable time is being wasted on troubleshooting the wrong areas.

HOW SPLUNK HELPS

With ML, Splunk Observability Cloud quickly surfaces the anomaly in the Kubernetes node's resource usage, helping the team isolate the root cause of the problem — an out-of-memory error. GenAI/AI Assistant simplifies the findings by providing step-by-step troubleshooting so the engineers understand and resolve the issue promptly.

The multi-layered power of machine learning as Sherlock unravels complexities

Digital infrastructures — complex, layered, and rapidly evolving — are the best environments for Moriarty to inflict observability mysteries. Problems arise suddenly when least expected and without explanation, impacting digital resilience.

Imagine a large-scale cloud deployment experiencing unexpected slowdowns, which result in user complaints and potential SLA violations. Engineers struggle to identify the root cause amid the overwhelming data volume and conflicting performance metrics. You guessed it — Moriarty broke in, and a digital crime scene is the result.

But as Moriarty scatters observability chaos across digital topologies without rhyme or reason, it's time to employ the services of Sherlock Holmes. His keen ability to understand hidden complexities, patterns, and nuances by deciphering digital clues in a logical, organized, and structured manner makes him indispensable in an investigation.


In these situations, machine learning as part of the broader AI toolset within an observability solution becomes the best detective you could ask for: the “Sherlock Holmes” to bring order and clarity to the scene, uncovering anomalies, connecting the dots across complex layers, and digging deeper into areas where humans alone might struggle.

Over half (55%) of respondents use AI/ML tools to determine root causes during investigations. Grouping and correlating data are seen as major challenges, with AI/ML tools alleviating these pains.

— [State of Observability 2024](#)

Sherlock Holmes

MACHINE LEARNING



ALIAS The Pattern Seeker

ROLE The ML Detective of Observability

STRENGTHS

- **Pattern Recognition:** Identifies patterns amid complexity and connects the dots across disparate datasets.
- **Deductive Mastery:** Spots subtle deviations that no human could easily detect.
- **Data Sleuth:** Examines huge datasets with precision, discovering hidden connections and anomalies no human could easily spot.

PERSONIFICATION Machine learning embodied as the brilliant detective. With his impeccable reasoning, Sherlock sifts through mounds of data with precision.

SPECIAL ABILITIES

- **Anomaly Detection:** Uses his magnifying glass to identify unusual behavior and outliers in metrics and logs.
- **Signal in the Noise:** Knows which clues are critical to solving the case, just as ML identifies and prioritizes relevant alerts, reducing noise and distractions.
- **Root Cause Analysis:** Pinpoints the cause by connecting every clue, drilling deep into the chaos to solve the mystery.
- **Predictive Analytics:** Looks into the data “crystal ball” to use patterns from the past to forecast future issues and prevent them.



FAVORITE QUOTE The game is afoot — data speaks the truth when you know how to listen.

However, Sherlock's deductive reasoning and systematic approach, like skillfully navigating an observability solution powered by ML, requires technical depth and expertise. These ML-powered tools analyze telemetry data to pinpoint underlying issues by detecting patterns, identifying anomalies, and surfacing insights that would be impossible for a human to catch alone.

Machine learning is, in essence, Sherlock's toolkit. Let's see what's in the kit:


- **Anomaly Detection** – Sherlock's magnifying glass, ML identifies unusual patterns, just like Sherlock spots suspicious behavior in a crowd.
- **Alert Noise Reduction** – Separating the signal from the noise, Sherlock knows which clues are relevant and which are not. ML helps teams reduce noise from false positives to allow them to focus on meaningful alerts.
- **Event Correlation** – Seeing the big picture, ML connects seemingly unrelated events, much like Sherlock does, to reveal the broader mystery and prioritize what's important.
- **Root Cause Analysis** – By connecting all the dots, ML bubbles up probable root causes by exhaustively correlating data, just like Sherlock links every clue to identify the culprit.
- **Predictive Analytics** – Like a crystal ball, historical data can foreshadow future issues, enabling engineers to address potential disruptions before they escalate proactively. Just as Sherlock anticipates his adversary's next move, ML allows teams to foresee and mitigate risks.

Machine learning enables engineers to overcome Moriarty's chaos enablers (data overload, alert fatigue, and false positives), discover meaningful clues, and tackle the chaos dramatically faster and in more structured ways.

ML uncovers insights with precision, efficiency, and speed — and with far less human toil and effort.

However, what is uncovered can be too complex for those without deep forensic experience. What Sherlock or ML uncovers doesn't always guarantee quick action without interpretation.

Sometimes, the best detective needs a translator, a trusted partner who makes complex findings accessible and actionable.



CASE #2

The case of the disappearing transactions

SCENARIO

A payment service is experiencing intermittent failures, and customer complaints are piling up. The team suspects the problem is an obscure networking issue, but even after sifting through log data from multiple services, they can't pinpoint the cause.

CHALLENGE

The team is struggling to correlate telemetry data from different services. The issue is distributed across multiple network layers, making it difficult to diagnose with traditional tools.

HOW SPLUNK HELPS

Using **ML**, **Splunk Infrastructure Monitoring** detects a pattern in network latency spikes that coincides with payment failures, which traces back to an overloaded network switch. **GenAI/AI Assistant** helps surface the problem and guides the engineers through the resolution process, speeding up incident response and avoiding downtime.

Dr. Watson's approach: using GenAI to bridge the gap between complexity and human understanding

Moriarty leaves chaos in his wake — volumes of disjointed telemetry, conflicting logs, and elusive clues.

Sherlock (ML) finds the hidden clues.

However, Dr. Watson (GenAI) explains these clues to the rest of us, making sense of complexity and providing practical solutions.

Some challenges engineers encounter when attempting to understand observability chaos:

- **Data interpretation** – Even when issues are surfaced, understanding them often requires deep technical expertise.
- **Root Cause Analysis** – Engineers struggle to identify the true root cause within vast volumes of data and the complexity of modern environments, leading to prolonged troubleshooting and delayed resolutions. Manual data correlation across multiple sources and tools makes pinpointing issues difficult.
- **Faster MTTI/MTTR** – Engineers need faster MTTI and MTTR to avoid business- or customer-impacting issues.
- **“It’s buried in the manual”** – Observability experts know the tools they use. However, not everyone is an expert, and detailed documentation is daunting.

Dr. Watson provides a more human and approachable method for interpreting and acting on huge volumes of data organized within Sherlock’s brain. He knows how to ask the pertinent questions to get actionable responses.

84% of respondents have explored generative AI within observability platforms, specifically for chatbots and embedded AI assistants. They envision using these capabilities for data analysis (66%) and recommendations to resolve issues (60%). — [State of Observability 2024](#)

Dr. John H. Watson

GENERATIVE AI



ALIAS The Translator

ROLE The GenAI Assistant in Observability

STRENGTHS

- **Humanizer:** Makes complex data insights accessible to everyone, regardless of expertise.
- **Provides Guidance:** Analyzes available ML data and turns it into actionable recommendations.
- **Insights Facilitator:** Provides the code and guidance necessary to develop dashboards that bring data insights to the forefront.

PERSONIFICATION Dr. Watson, the indispensable partner to Sherlock. GenAI, like Dr. Watson, translates complex findings into easy-to-understand explanations for everyone involved.

SPECIAL ABILITIES

- **Natural Language Queries:** Ask for and receive observability insights in human terms, not jargon.
- **Guided Troubleshooting:** Walks teams through problems step-by-step, expediting Mean Time to Resolution (MTTR).
- **Upskilling Mentor:** Provides summaries and tool how-to's, helping engineers learn faster and onboard more easily.



FAVORITE QUOTE Sometimes, a little clarity is all that stands between chaos and resolution.

GenAI assists in making ML's discoveries understandable and actionable for every engineer, regardless of expertise. Simply stated, GenAI brings order to observability chaos through natural language interpretation in a variety of ways:


- **Natural Language Interactions** – Dr. Watson's language superpowers lie in asking Sherlock simple questions and conveying answers in a plain, natural language. GenAI mirrors this by enabling natural language queries and delivering clear, jargon-free responses.
- **Guided Troubleshooting** – Picture Dr. Watson guiding the user through each deduction step, ensuring everyone understands the clues and their significance. GenAI provides engineers with a clear, step-by-step troubleshooting guide, turning findings into concrete actions that can be taken without hesitation, reducing MTTR.
- **Dynamic Charts** – Dr. Watson might sketch a crime scene to help Sherlock and others see the connections. Similarly, GenAI assists engineers in creating dynamic charts or dashboards, bringing key metrics into focus.
- **Accelerating Learning** – Just as Dr. Watson's presence in the field helps others understand Sherlock's methods, GenAI acts like a mentor. It provides contextual summaries and walkthroughs to engineers, enabling them to learn much faster. GenAI not only solves the case but also equips the team with knowledge, simplifies onboarding, and maximizes the value of observability investments.

While Sherlock and Dr. Watson can operate independently, they are most powerful together. They are, after all, from the same detective agency.

Machine learning brings deep analysis, and GenAI bridges the understanding, making complex insights digestible. Together, they align knowledge with intuition, translating complexity into actionable results. This partnership equips teams — from Software Developers to Site Reliability Engineers — with the tools they need to meet challenges head-on, enhancing their ability to keep systems running smoothly and efficiently.



GenAI capabilities aren't solely optional enhancements — they are transformative tools that revolutionize how observability challenges are addressed, aiming to achieve a digitally resilient state. By providing a natural language interface, guided troubleshooting, and dashboarding and summary support, GenAI empowers engineers of all expertise levels to tackle complex system issues effectively, driving faster resolutions and enhancing overall reliability.



CASE #3

The database performance bottleneck

SCENARIO

An online retail platform is experiencing delays in processing customer orders. Transactions are taking longer than usual, leading to customer complaints about order confirmation delays. The development team suspects a database issue, but with multiple application instances and complex logs, it's difficult to identify the exact cause.

CHALLENGE

The team is overwhelmed by telemetry data from the database environment and struggles to determine if the bottleneck is related to disk usage, CPU performance, or resource contention. Manually analyzing logs is proving too slow to keep up with the problem.

HOW SPLUNK HELPS

ML within **Splunk Observability Cloud** detects CPU and disk usage anomalies on database instances, surfacing performance degradation tied to resource exhaustion. **GenAI/AI Assistant** presents the data in a clear format, revealing the correlation between transaction delays and underperforming database resources. GenAI also gives guidance on how to set up ongoing performance monitoring for faster identification of future bottlenecks.

Solving the observability mystery: the dynamic partnership of machine learning and generative AI

It's been quite an investigative journey. First, we were thrown into the chaos of observability, overwhelmed by data overload, disparate logs, alert fatigue, the challenges of rapidly evolving modern IT infrastructures, and an erosion of digital resilience.

Next, witnesses had suspicions and ideas about “whodunnit,” but these were often conflicting and based on different pieces of evidence.

And while fictional murder victims typically can't complain, end users of your applications certainly can and do. As they suffer, so do those trying to understand what is wrong and how to resolve it.

Servers, software, transactions, and infrastructure leave full stacks of clues in the form of events, logs, and traces, but without AI, you may need to bring in highly-trained, expensive experts or leverage resource-intensive tools to get some initial clarity.

56% of respondents use AI/ML to correlate events and prioritize alerts, and 53% use it to recommend solutions. Among leading organizations, 91% rely on AI/ML for recommending solutions.

— [State of Observability 2024](#)

Machine learning helps detect patterns in observability chaos.

GenAI bridges the skills gap, allowing engineers to both query observability tools in natural language and receive results in plain, accessible terms — making sophisticated observability solutions attainable regardless of their expertise.

Observability tools coupled with AI (specifically ML and GenAI) allow engineers to focus on what matters — solving the crime fast with minimal toil. MTTR is faster and less costly.


When it comes to observability, leveraging ML and GenAI isn't just about maintaining uptime or high-performance systems. It is also a strategic advantage that allows your engineers to focus on high-value issues, innovate, prevent burn-out, and improve their skill sets.

And, of course, you want to defeat Moriarty.

It's your case now. Which detective agency will you call, and who will make up your sleuthing dream team to solve the crime of observability chaos?

Baker Street AI Agency

ARTIFICIAL INTELLIGENCE



ALIAS The Unified Artificial Intelligence Bureau

ROLE The Foundation of Observability Intelligence

STRENGTHS

- **Holistic Intelligence:** Combines the power of pattern recognition with human-centric explanation.
- **Unified Approach:** Unites deductive analysis and interpretive guidance, ensuring comprehensive coverage of every observability challenge.
- **Rapid Response Team:** A seamless synergy of ML and GenAI, working together to quickly identify, understand, and resolve complex IT and engineering mysteries.

PERSONIFICATION The Detective Agency, the command center behind the great detective duo, Sherlock and Dr. Watson. It is AI itself — harnessing both brilliant deduction (ML) and effective communication (GenAI) to solve the toughest observability cases.

SPECIAL ABILITIES

- **Collaborative Analysis:** Leverages Sherlock's ML-driven insights and Dr. Watson's human-centric simplification to ensure no stone is left unturned.
- **Integrated Investigation:** Combining in-depth analysis with clear interpretation to address every aspect of the digital mystery.
- **Cohesive Oversight:** Synthesizes fragmented data into cohesive insights, providing engineers with a comprehensive view of their environments.



FAVORITE QUOTE A case may begin with a clue, but the combination of insight and intuition ensures it's solved.

Turn detective work into actionable insights with Splunk’s AI solutions for observability. Whether it’s pinpointing anomalies with ML or simplifying complex data with GenAI, Splunk equips your team with the tools to stay ahead of disruptions and ensure digital resilience.

Learn more



Splunk, Splunk> and Turn Data Into Doing are trademarks and registered trademarks of Splunk LLC. in the United States and other countries. All other brand names, product names or trademarks belong to their respective owners. © 2025 Splunk LLC. All rights reserved.

