

Observability

Predictions 2024

How AI will revolutionize IT and engineering.





We asked generative AI to write predictions.

Just kidding — as we make our best predictions for 2024, there are many unknown unknowns. And if you know even a little bit about what we do, you know Splunk Observability is built to alert you to the unknown unknowns.

We've designed this year's predictions in that spirit. It's our 20th anniversary. We've been through a tech cycle or two. But this one feels different. Advancements in AI and large language learning models (LLMs) are changing the way we respond to, well, everything. And while AI is here to stay, it doesn't mean our lives are any less complex.

Our observability leaders have made predictions on everything from generative AI to the future of DevSecOps. Things are evolving quickly, but we recommend you take your time and let our observability experts make your journey to digital resilience a little less complex.

"Can we have even more standards?" asks Arijit Mukherji, Splunk distinguished architect. "How can we make it easier to automate sub-parts, like creating an alerting rule, or building a dashboard? As automation takes over many smaller parts of our workflows, open standards can help them interoperate with each other to achieve bigger and better things. Maybe it won't happen immediately, but this has already started." We won't give away all our predictions up front, but a small spoiler alert: Check out the section on tool consolidation and OpenTelemetry.

"There are many new applications for AI, especially in anomaly detection," says Zongjie Diao, Splunk senior director of product management. "As long as you keep customer needs first, you will be on the right track." Many of our experts agree with her. We'll likely use AI to detect trouble spots first, but we may be a few

years away from more advanced applications. Humans will still be in the loop. (We too are hopeful that the robots won't come and take all our jobs.)

Greg Leffler, Splunk observability practitioner director, wonders: "Will AI observability tools keep up with the speed at which people are changing their applications?" And as budgets are mindfully consolidated, how will that show up amidst the enthusiasm for all things AI?

"I think automation is on a meteoric rise," says Cory Minton, Splunk field CTO. "There will be a confluence of automation with machine learning and deep learning. As the sub-disciplines of AI start to work together more effectively, we start to unlock the real, meaningful power of AI."


The way that we instrument modern applications is going to mature in 2024. This will put pressure on observability tooling providers to create out-of-the-box value. Customers will expect more value in the form of automation and AI add-ons.

Even in a year where CIOs and CTOs will have to be mindful of their budgets, observability will continue to drive meaningful outcomes. Analytics will continue to be valuable. And metrics, traces and logs aren't going anywhere.

Read on to explore key challenges.

Contents

- 05 AI and observability
- 07 Tool consolidation
- 09 Anomaly detection
- 11 Edge computing
- 12 Security and observability
- 14 Advice to CTOs and CIOs
- 15 2044
- 16 Contributors



Prediction: AI and observability

AI disrupts observability, though most revolutionary gains won't be seen in 2024.

Unless you live under a rock in a desert on a different planet, you are aware that AI is here. The hype is off the charts. It's early days, the wild west or whatever metaphor you choose to denote an overabundance of enthusiasm. We won't waste your time making a prediction about how AI matters — we all know that the stakes are high.

So, how will AI impact observability in 2024 and beyond? What will you need to know to prepare for the inevitable changes?

Our experts are here with some predictions that may surprise you.

“Generative AI and code copilots will make us superhuman, right? Rising productivity is going to cause an explosion in the number, scale and complexity of things that we build in the coming years,”

says Arijit Mukherji, distinguished architect. “In the next five years, you will have to observe many more things, like environments, applications, microservices, code pushes and clusters. Thanks to AI, observability solutions will have to deal with far more variety and volume of data — open standards will become much more important.”

“AI is very expensive, both from a training standpoint and a computation standpoint to actually do,” says Greg Leffler, observability practitioner director. “And the benefits are small for most use cases. So I think you're gonna see a lot of people saying, ‘Hey, we have AI, but it's evolutionary, not revolutionary.’”

“I think we're still decades away from an AI that can actually think,” Leffler adds. “So everything that it's doing is based on what already exists. Based on how you train the AI and what you teach it to do, observability will probably be a natural outcome of that, right?”

“You will see more adoption of automation for sure. At least for now, cost is still a big concern,” says Zongie Diao, senior director of product management.

With all these massive changes, regulation can't be far behind. When you are training large language models on data generated by humans, you might have a model that reflects the biases of these humans, even subconscious ones. The AI bias might not be inherently discriminatory, but it's important to recognize that the bias is there. There might be a path for observability to help monitor AI.

We also expect to see customers developing their own AI, which will be another mission-critical workload. “I think we're going to have to figure out how to get the metrics, logs and traces from the AI,” says Cory Minton, field CTO. “Is it doing what is expected? And if not, how are we tuning and building that whole machine learning (ML) ops pipeline to make sure that we are maintaining functionality?”

All our experts agree that a human still needs to be at the center of the equation. If AI builds an application, humans will still need to keep track of how everything is operating.



Prediction: Tool consolidation

Tool consolidation will be necessary. OpenTelemetry is a rising star and might steal the show.

If you want to avoid downtime and improve customer experience and retention, tool consolidation goes beyond “nice-to-have.” Consolidating your monitoring tools into a single observability system frees up your engineers and system admins to focus on your core business. Reliability matters.

Zongjie Diao, senior director of product management, sees two developing trends: “Because of tool proliferation, customers are saying, ‘I really want to make sure I can take all the different kinds of data together because it’s very fragmented. I want a holistic view.’ On the other hand, they want to consolidate and standardize on one or two vendors. And that’s both for efficiency as well as for cost management concerns.”

The move to consolidate tools is already happening, and it’s picking up steam. Organizations are rightfully concerned that their data is bound to a vendor. They see it as a liability, and they worry that they are not secure or safe. While some organizations may be initially focused on cost, many are also focusing on consolidation to standardize how they collect and consume data.

“Tool consolidation will continue to trend because hybrid cloud environments are not getting less complex,” Diao says.

To manage this complexity, organizations are turning to AI and automation. Our experts break down what automation could mean for your toolset.

“AI will pave the way to more complex systems with greater compliance challenges,” says Arijit Mukherji, distinguished architect. “This will happen through smart AI-based automation that handles patching, auditing, evidence gathering and the like. If I have a system that just does that for me, life becomes much easier.”

Add ChatGPT and auto-personalization tools into the mix, and we're looking at a future where personalized dashboards, portals or developer service platforms can theoretically orient to a particular person's method of troubleshooting or a specific set of guidelines based on the affected users, independent of a technician's credentials.

"I think that the tools have to coalesce at some point — observability tools have to start to coalesce with IT tools, and those tools have to coalesce with security tools. The idea that any of those aren't connected is just disingenuous to where most technology organizations are, or at least want to be," says Cory Minton, field CTO.

OpenTelemetry will continue to gain traction in the coming years. Even if a customer chooses a single vendor, they will retain control over their data. They won't feel the burden of vendor lock-in if their vendor's costs rise — they will retain the freedom to add or switch vendors or tools and take their data with them. That's been true for years, and it will remain true in 2024.



Prediction: Anomaly detection

AI will change the way we detect and identify anomalies — it won't replace manual troubleshooting.

Applying machine learning and AI to security use cases like anomaly detection is already underway. It's a natural first step on the way to solving more complex problems. According to Cory Minton, field CTO, "It just means that AI is applied at a broader scale to a broader set of data — and that's getting easier."

AI is also making it easier for more stakeholders to interact with complex data. "AI in the medium term will begin to break down the laws of technical skill," Minton adds. "You will be able to ask questions and interact in a more human way — so you don't have to write queries."

The big question: How do you build models that incorporate broader datasets? A large portion of observability practitioners will shake their heads and say, "no, that's never going to happen."

Minton predicts that in the long term, AI will bring a more concise understanding of what's going on in an environment. First AI will tackle anomaly detection, next up will be investigation and automated response. We will see automated remediation in the near future.

Still, many developers believe that these are early days and many anomaly detection tools will be branded as AI tools, even though they are not. "There will be real value in enhancing troubleshooting capabilities by helping to continue to refine and present the issue to the correct person," says Greg Leffler, observability practitioner director. "The ML sweet spot is locating the problem and flagging it to the person who is supposed to fix it, because ML and AI platforms are never going to be able to fix every type of problem automatically."

Patrick Lin, Splunk SVP, GM of observability agrees: “No, AI won’t replace manual troubleshooting — what’s more likely is that it will augment the human who still needs to be in the loop. AI/ML will be faster at sifting through all the available data and recommending where to focus, and in what order, to get to problem identification.”

Zongjie Diao, senior director of product management, believes that manual troubleshooting will eventually go the way of typewriters. It’s here to stay for now, because customers are wary of reliability issues.

“For the detection part, total AI detection is going to take some time but it will happen. We’re talking about enterprise usage with many reliability requirements. Some manual troubleshooting settings will always be there,” Diao says. “The issue will be building customer trust in the AI. I don’t think AI is going to replace manual anomaly detection in the next three to five years. It will take time for customers to feel satisfied that AI achieves the reliability levels that they want.”



Prediction: Edge computing

Edge lives up to our high expectations in 2024.

Remember when people were excited about the Nest Learning Thermostat? Neither do we. Feels like a lifetime ago, yet it's only been a decade since the second generation debuted. Uber began offering rides way back in 2010. Time marches forward.

Now, consider the power of edge computing. An inexpensive quartz watch is more powerful than the computers on Apollo 11. Most people in San Francisco have not hailed a driverless vehicle ride, but they will by the end of 2024. The Nest Thermostat could only regulate temperature — imagine if it could do upwards of a million more things. These are early days.

Computational capabilities are deployed at the edge daily. Look no further than the proliferation of video content creation, which you no longer need to do on a clunky desktop computer. These next generation computational environments will require AI because in the future, many of the problems we observe will no longer be at human scale.

“Across industries, we’re pushing more and more computational power to the edge,” says Cory Minton, field CTO. “We’re gonna have to learn how to properly instrument, identify and observe challenges fast enough that we can handle the remediations in an appropriate time. There’s going to be a complete paradigm shift.”

“I think for edge computing, we are seeing more data streams than we’ve ever seen before. Everything that we do generates some sort of data, some sort of telemetry, and it gets sent somewhere. There is going to be an ever increasing volume of data, which lets you get richer insights into what’s going on — helping you to figure out what your customers are doing,” says Greg Leffler, observability practitioner director. “There are potential data risks that people will need to consider as they embrace the edge. I think our customers are well-positioned to deal with those risks, but it’s still going to be a concern for many people.”

“I think edge computing has lived up to my expectation, but not to the industry’s expectation. I feel like IoT has been taught for between 10-20 years. The industry was expected to be a lot further along than we are today with leveraging ARP, AR and VR. I think there is still a lot of room to grow there,” says Zongjie Diao, senior director of product management.

It’s just like the start of the internet, when people primarily used websites for education and entertainment. Soon, the secondary and tertiary use cases will emerge. Those use cases eventually seep into enterprise technology, which always lags behind consumer tech. First, you’ll see an explosion of consumer use cases. Then, very quickly — especially in retail, banking and media — the edge will become a hot topic for IT development and security departments.

Prediction: Security and observability

Observability becomes a meaningful signal to security operations.

Remember the character Bertram Gilfoyle, of *Silicon Valley* fame? He was the ops guy. He set up servers for the fictional startup, Pied Piper, in their garage. Those servers live in the cloud for most software companies, which makes sense if you're a SaaS company. If your servers are in your garage, more power to you.

If you're building cloud-native applications, there are services that can help you get a fuller picture of your infrastructure — including application security. For many vendors, observability products are completely separate from security products. Customers are often frustrated by their lack of interoperability. They use too many tools, or at best, two completely separate platforms that can't speak to each other. Get ready for 2024, because it's going to change — but your DevSecOps perspective will still depend on what team you're on.

You can define DevSecOps from a developer perspective. Or, you can think of DevSecOps from a security perspective. There are many security practitioners who are interested in observability data because it gives them context on security issues. They want to see metrics and potentially trace data correlated with events. With a unified security and observability platform, they will be able to decide if an incident is security or performance relevant.

Recognition of the increasingly threatening threat landscape will drive the shift, as will compliance and regulatory needs. The increasing prevalence of cloud-first or cloud-native organizations also lends a natural gravity to DevSecOps.

“Yes, I think DevSecOps will happen,” says Zongjie Diao. “I don't think everything will be about security and ops together all the time. There will continue to be individual domain areas that are very ops-centric over security-centric. But there will be shared areas in the near future that will be considered DevSecOps.”

“Much of the chatter on the topic of DevSecOps focuses on ensuring that security can be done in a way that is able to keep up with the pace of development,” says Patrick Lin, SVP, GM of observability. “For a lot of organizations, the answer today is a hard no, and it’s going to take some amount of innovation in this area. But there are some simpler problems at the intersection of security and observability that can be tackled in the nearer term — for example, showing service owners more context, and whether there are security issues affecting their service.”

“The idea [of integrating operations and security] is great, where we have everybody who is involved in supporting the application, you know, working together. But modern application security

is so complex, there’s no way that somebody can be [both] an expert in application security and expert enough in operations to adequately serve the application, right? You need specialists,” says Greg Leffler, observability practitioner director. “Is this an ops problem? Is this just traffic? There’s never going to be a DevSecOps engineer. I just don’t predict that that’s possible.”

So, DevSecOps has already arrived or it will be coming soon, depending on which angle you look at it and what team you’re on. In the meantime, stay observable. Stay secure. Implement tools that will help you do both. Whether your servers live in the cloud or a back corner of your garage, a DevSecOps mindset will lead your organization — big or small — toward digital resilience.



Prediction: Advice to CTOs and CIOs

CIOs and CTOs will cut back on their architecture and infrastructure spending, making this the year of mindful budgets and massive disruption.

“With great power comes great responsibility.” We all know this famous quote from none other than Spider-Man. (Yes, the comic book superhero, and unfortunately still not an employee of Splunk.)

This quote rings true with AI — people are excited, but they are also nervous. This is normal. These innovations are arriving in a difficult macroeconomic environment where CIOs and CTOs are mindful of keeping their budgets in check.

Our experts are divided between excitement and hesitancy. They all agree that we are on the verge of major disruption, and most will say that it can't possibly be all good news. Systems will break down. New threats will emerge. New game plans will be implemented. As systems become more resilient and more sophisticated, threat actors also become more sophisticated and more capable of hacking even better systems.

“We're at the beginning of the hype cycle for AI — this is a good time to slow down and recognize that,” says Zongjie Diao, senior director, product management. “It's a journey and it will be complex.”

CIOs and CTOs will feel the demand to get more from less. They will need to adopt observability platforms to ensure that their customer experience doesn't decline — or customers might leave. “The best advice is: If you need to cut back, you'll need observability to know if you've cut back too far,” says Greg Leffler, observability practitioner director.

“It's no longer just about detecting, investigating and responding to things in cloud-native apps,” says Cory Minton, field CTO. “A robust observability practice will build a unified view of your organization, including security risk, compliance, governance and financial impact.”

Arijit Mukherji, distinguished architect, believes that this is the year of AI. “By 2024, CIOs and CTOs better form an understanding of how AI might disrupt them, how it is an opportunity for them — or both. They will need to develop a point of view, and quickly. This is not just isolated to the tech industry, because AI will disrupt a wide range of areas from healthcare, law, finance, travel and more in ways we haven't yet imagined.”

2044

We're going back to the future — you'll want to stick around for the credits.

We made some bold predictions about AI and observability in this year's Predictions. The main takeaway: It's impossible to anticipate how AI will disrupt our lives. Managing complex hybrid multicloud infrastructures isn't going to get easier, and keeping systems safe and resilient is a tall order. As we embrace the edge, it's clear that observability will play as big a role as ever.

This year, as we celebrate Splunk's 20th anniversary and enter into the era of AI, we thought it only fitting to ask our senior leadership team one final question: What do you envision for the world in 2044?

"Hybrid computing gets a whole new definition. Right now, when we talk about hybrid, most of what we talk about is cloud versus on prem. Hybrid also kind of means 'edge.' I think what's going to happen in the next 20 years is that we're going to have a completely new set of computational capabilities come online that significantly challenge many of the things we hold true in technology," says Cory Minton, field CTO.

"Humans will start designing our lives around automation, as opposed to the other way around," says Arijit Mukherji, distinguished architect. "In 20 years, houses will be designed to accommodate automatic vacuums. Your backyard will be designed for an automatic lawn mower."

The world looks very different today than it did when Splunk launched twenty years ago. As governments mandate digital resilience and as new compliance standards for AI evolve, we predict that a unified security and observability platform will be more important than ever.



Contributors




Zongjie Dao

Zongjie, aka "Z", is a senior director of product management on the Splunk observability team. Prior to Splunk, she led the strategy and new product incubation product team on Cisco's data center business unit. Prior to Cisco, she worked at McKinsey as an engagement manager.



Cory Minton

As field CTO for Splunk, Cory Minton leads strategy and innovation alignment with Splunk's largest customers, including the technical strategy for Splunk's partnership with McLaren's F1 and Esports Teams. Cory is also known as the "Big Data Beard" from the eponymous podcast he leads.



Greg Leffler

Greg heads the observability practitioner team at Splunk, and is on a mission to spread the good word of observability to the world. Greg's career has taken him from the NOC to SRE to SRE management, with side stops in security and editorial functions.



Arijit Mukherji

Arijit is a distinguished architect at Splunk leading architecture for Splunk's observability portfolio. Previously he was CTO at SignalFx (acquired by Splunk), where he was instrumental in building their observability solution from ground up. He holds a master's degree in computer science and is the author of 10 technology patents.



Patrick Lin

Patrick Lin is the SVP and GM, observability at Splunk. He joined Splunk in 2019 through the acquisition of SignalFx, where he was the chief product officer. Patrick is the former VP of product management at VMware, where he led the teams responsible for the virtual infrastructure platform.

For more 2024 predictions, read our executive and security reports.

Read now



splunk>

Splunk, Splunk> and Turn Data Into Doing are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names or trademarks belong to their respective owners. © 2023 Splunk Inc. All rights reserved.

23-347911-Splunk-Observability Predictions 2024-EB-116

Keep the conversation going with Splunk

