

The AI Philosophy Powering Digital Resilience

How AI will usher in a new era of security and observability



AI spring is here to stay.

Ask anyone to name pivotal moments in history, and a few famous ones are likely to come up. Alexander Graham Bell makes the first-ever phone call. Tim Berners-Lee introduces the internet we know today. Steve Jobs unveils the first-generation iPhone. These inventions fundamentally changed the way we communicate, live and work.

A look back into the decades-long history of AI reveals several peaks and valleys of public interest — respectively, AI springs and winters — since its inception in the mid-1950s. The release of OpenAI's ChatGPT in late 2022 introduced generative AI into the mainstream, marking another pivotal moment that will someday be in the history books. It launched an AI spring that we believe will deeply impact the course of industry and society for decades to come.

Gartner's Hype Cycle for Emerging Technologies in 2023 claims that generative AI will reach the peak of inflated expectations, but we believe in the long term that AI is actually underestimated and the hype cycle represents only a mere sliver of its potential impact. In fact, AI in all of its forms will become as ubiquitous to modern life as electricity and be so ingrained in our daily lives that we'll barely even think about it — like switching on a light.

Today's digital environments are perfectly primed to benefit from the power of generative AI. In the past, a single person could realistically have access to an organization's entire dataset. But as organizations embrace dispersed architecture across hybrid and multicloud, it has become nearly impossible to see all of the pieces of the puzzle. Several confounding factors — including an evolving threat landscape and heavier reliance on digital systems — create a maelstrom of chaos for IT, engineering and security professionals.

Humans simply cannot wrangle, manage and monitor this complexity alone, leaving technology and security teams stressed and resource-strapped. It was hard enough to find a needle in the haystack — but now they must find multiple needles in a whole field of hay. AI steps in as a promising solution. Adoption is soaring; 91% of security teams are currently using public generative AI tools, according to Splunk's [State of Security 2024: The Race to Harness AI](#). And according to [The CISO Report](#), 86% of CISOs believe that generative AI will alleviate skills gaps and talent shortages.

In short, we see AI as a requirement to build digital resilience.



What do we mean by AI?

We define AI as a superset of multiple disciplines that includes machine learning, deep learning and generative AI.



About the author:


Hao Yang, VP of AI at Splunk

As vice president of artificial intelligence, Hao leads Splunk's team of software engineers and data scientists to accelerate the company's innovations in the AI arena. Prior to Splunk, Hao served as VP of artificial intelligence at Visa and has held several AI and big data innovation research and engineering positions with globally-recognized companies, including Google, Nokia and IBM.

It starts with the data.

Yes, AI is critical to digital resilience — but simply incorporating AI into a security and observability platform is the minimum. As LLMs become commoditized, data and domain expertise will be the difference between a good and great AI model. Powering and protecting the AI revolution requires quality data and a lot of it.

The more information a model has to learn from, the smarter and more accurate it will be. An AI model that’s trained on a wide breadth of data can more easily recognize patterns, detect outliers, and adapt accordingly. To accelerate those insights and actions, we’ll bring LLMs to the rich datasets that span an organization’s entire digital footprint — from endpoints to the cloud and everything in between.





The antidote to AI anxiety

While the promise of AI has garnered mainstream interest, it's also at the center of intense scrutiny and distrust, especially for cybersecurity professionals. As the saying goes, "you can't protect what you don't know," and there is still so much to be discovered about AI. Discomfort in the unknown — this AI anxiety — is only natural.

Here's what we do know: AI will be used for good and bad, presenting opportunities for both defenders and malicious actors. It will introduce new social engineering attempts, increased data privacy and security risks, and expanded attack surfaces as teams delve into AI development and lower the barriers to entry for cybercrime. It will also make IT, engineering and security operations faster and easier.

Fueling the uncertainty around AI is a lack of transparency. What is substantiating the decisions that an AI model makes and why? Decision making is particularly high stakes for digital resilience — a wrong decision can result in millions of dollars' worth of downtime or a costly data breach. And while humans can easily discuss their reasoning behind decisions, fully explainable AI doesn't exist yet. Similarly, human biases are well-documented, but biases in AI are baked into its algorithms and training data and are therefore more difficult to detect.

We promise it's not all doom and gloom — far from it, actually. The future of AI is incredibly bright, and it will be a positive force for digital resilience.

Regulation and clear governance will likely address some of the problems currently associated with AI. But in the meantime, it's on companies pioneering the usage and development of AI — like Splunk — to help organizations navigate AI safely and effectively.

Since embracing AI as a discipline in 2015, we've had to answer hard questions: How do we keep our customers safe in this era of uncertainty? How can we lean on AI to enhance digital resilience, not erode it?

Companies leading this new era of AI must take thoughtful, deliberate approaches to its development. That's why we created our AI philosophy: Splunk's set of beliefs around AI that will never waver. These tenets will deliver value to organizations around the world, inform our long-term strategy and most importantly, drive the responsible use and development of AI.

With this philosophy acting as a North Star, our AI adventure continues.

There is beauty and grace in AI, in its ability to automate and reduce friction. And while this will be incredibly productive, there are still some unknowns. We'll question the security and reliability of AI, asking: Do we trust this? Is this accurate? Has the output been poisoned?

– Paul Kurtz, Splunk Chief Cybersecurity Advisor and Field CTO

Purpose-built, embedded AI drives the best digital resilience outcomes.

Did you know that the founders of Splunk, Erik Swan and Rob Das, first tried to cure cancer? While their initial failure can be attributed to a variety of reasons, one is that they didn't have domain expertise. "I still have my bookmarks of genomics research in my browser and I never made it through the first paper," Swan said once in a Silicon Angle interview.

Instead, they decided to solve a problem they deeply understood through experiences at their previous company: the challenge of searching through logs to troubleshoot why IT infrastructure was down. Having been personally frustrated by this tedious process in the past, they decided to make Splunk the solution.

Similarly, experts in security and observability should design and build AI models for security and observability because they have the best insights derived from real-world experience and data. Models should also be trained on data sets based on specific use cases and environments to deliver relevant context. For AI to be truly useful in the context of digital resilience, it must be domain-specific — in other words, purpose-built — and tightly embedded into everyday workflows.

Let's delve into what it means for AI to be purpose-built. Generic AI large language models (LLMs) are trained on massive datasets, enabling them to have some knowledge about just about everything. Using generic AI to solve specific problems, however, is like trying to build a fence with a Swiss Army knife rather than a power drill. A generic LLM can analyze incidents at a high level, for example, but that analysis may not be accurate or relevant to your environment.

Domain-specific AI depends on relevant and high-quality data, and as such, delivers much more useful and accurate outcomes. For specialized fields such as security, IT and engineering, this is vital.

Equally important is the ability for those domain-specific models to work seamlessly in your existing detection, investigation and response workflows. Currently these teams are forced to switch to another tab (or another app entirely), type their prompt into an AI service, switch back to their dashboard and copy and paste the response — while also creating greater data exposure and privacy concerns.

Security and observability practitioners are particularly exhausted with cumbersome and inefficient workflows. Organizations operate and maintain an average of 165 internally developed business applications, according to Splunk's [2023 State of Observability report](#). On the security side, 65% of security operations center (SOC) teams complained about pivoting among too many disparate tools and management consoles, inhibiting comprehensive and timely investigation and response, according to Splunk's [State of Security 2024: The Race to Harness AI](#). When it comes to digital resilience, every second of detection or response time matters. And when teams are forced to use multiple tools, they start to develop multiple sources of truth. Purpose-built AI integrates into existing workflows and surfaces at the right time, with the right context. Simply put, it works the right way when it counts.

Humans belong in the driver's seat, with AI as a trusted copilot.

People once envisioned the future as a world conquered completely by machines, a cartoon-like fantasy involving self-driving (or even flying) cars, friendly robot butlers and carefree vacations to outer space. But some of those predictions didn't quite pan out. The self-driving cars that do exist rely on deep learning algorithms trained on large datasets, enabling the vehicle to recognize road signs, traffic and pedestrians and make decisions based on that intelligence. And unfortunately, self-driving cars have made headlines for causing traffic jams and crashes — some of which have been fatal.

In a high-stakes situation like driving, machines shouldn't be trusted as the sole decision makers. No matter how large the training model, a vehicle can't — and shouldn't — comprehend the emotional significance of choosing between a dog retrieving a ball or a squirrel darting in the road.

Similarly, a single decision can make a big impact on observability and security outcomes, with a ripple effect into the business. AI should work alongside humans, acting as copilots to quickly deliver facts, summarize events and bring priority alerts to the forefront. But humans should always be in the driver's seat.

In [*The Role of Emotions in Military Strategy*](#), professor Samuel Zilincik argues that emotion can be beneficial in war, whether used as a motivator or method of manipulation. Blue teamers in particular must be more creative, empathetic and motivated to outsmart their adversaries, who often have a leg up due to the lawlessness of their strategies. Tacit knowledge, instinct and good old common sense are all underestimated skills in any security operations center (SOC), engineering or IT operations team that AI will never be able to replace.

Today's digital battlefield is fraught with complex environments and attack methods, and the stakes are incredibly high. Navigating it requires nuance that only humans have. Determining the right way to respond to an incident isn't always as straightforward as understanding what happened, who was involved and where it occurred. It requires untangling a web of context and peeling back layers of logic and reasoning. Researching the threat tactics of a new ransomware group also requires nuance and human ingenuity. Without existing reference material, an AI system would struggle to keep up with evolving types of attacks and incidents.

Human-in-the-loop feedback also contributes to more responsible models, which enhances trustworthiness. When it comes to matters of digital resilience, trust is paramount.



Today's digital battlefield is fraught with complex environments and attack methods, and the stakes are incredibly high. Navigating it requires nuance that only humans have.

– Hao Yang, VP of AI



Openness and extensibility will supercharge AI innovation.

Powerful AI requires the right data on a large scale. AI already relies on and generates a massive amount of data, and that will only expand as organizations train and run models with billions of parameters. Building AI models that will power the next 20 years and beyond requires forethought, planning and collaboration.

Open, extensible AI has a few different meanings. Extensible AI means flexibility for developers, data scientists and partners to build, extend or connect AI models in ways that work best for them. They can use [Splunk's proprietary models](#) or bring their own models to align with their organization's policies and risk tolerance.

If history has taught us anything, it's that people want to be met where they are. Just as hybrid and multicloud isn't going away anytime soon, neither is a desire to tailor and adapt to each organization's specific needs. Rigidity will only pose limitations and prevent innovation. To that end, AI systems should be easily integrated, customized and extended — and they should work seamlessly across the many different tools already within an observability and security stack.

Not only does extensible AI enable developers to build models faster and easier, but it also promotes community and collaboration. Just as open source has ushered in innovation within software

development, we believe that extensibility is the cornerstone of building AI to scale, which is a true team effort. Our portfolio is made more robust not only by internal engineering teams, but also by partners and third-party apps — we believe building AI is an equally collaborative process. And as customers opt in to share aggregated and anonymized data, we'll be able to fine-tune and continuously improve models based on their actual environments.

Open, transparent AI also means peeling back the curtain of mystery that often shrouds it. This is important for a few reasons. AI is like Pandora's box in more ways than one. Data privacy issues aside, it's difficult to grasp how an AI model makes decisions — and that's only getting harder as AI becomes smarter and more widespread. Scaling, creating and adapting anything requires some knowledge of the current framework. We believe that transparency for how things work further enables developers and data scientists to adapt models and build atop them.

That transparency promotes more responsible use of AI, too, revealing biases, inaccuracies and hallucinations. Future-proofing is also about building on a solid foundation and addressing problems so that cracks don't form later on.



Bringing our philosophy to life to catalyze digital resilience

In the pursuit of building a safer, more resilient world, our AI strategy will focus on four areas designed to build value and trust — and in turn, drive better outcomes for our customers.

Find the needles and know what to do about them.

For security and observability professionals, wading through hundreds of alerts every day doesn't exactly bring joy. It's also not the best use of time for a professional with expertise and domain knowledge, who's better off strategizing or problem solving instead.

That's why we've been innovating with ML for anomaly detection since 2018. We see the potential to take this further by detecting anomalies across security and observability — two sides of the same coin. Both are data problems. An anomalous spike in server utilization could be noisy code or an attack.

Our aim is to build a centralized AI service that will surface insights across our security and observability applications so humans can approach problems from multiple angles and remediate them more quickly. This can be applied across different use cases — we're building for centralization, not silos.

AI will allow us to better understand time-series, correlations and anomalies, identify and analyze the root cause, and recommend ways to remediate the problem.

Embed AI-powered assistive experiences into workflows.

We recognize that security, IT and engineering professionals are overwhelmed with the complexities in their environments — whether that takes the form of attack methods, data, tools or alerts. That's why our next pillar is to make Splunk more efficient and effective for our users. We'll surface AI in the right times and places to lighten the load, completely transforming time-to-value and daily outcomes customers can achieve. When users rely on Splunk, AI will help them get to better answers and actions faster.

We'll build, train and embed domain-specific AI techniques to enhance investigation and response workflows. AI will help write searches, summarize incidents and recommend investigation steps — all personalized for a customer's environment and use cases.

Support unique use cases at incredible scale.

Resilience at scale requires granularity — the ability to access large sets of detailed data without restriction. That’s why our goal is to enable our highly complex customer base to create and train their own models at massive scale.

We’ve been enabling our customers to do this through the [Machine Learning Toolkit \(MLTK\)](#) — the most downloaded app in Splunkbase today. But since its inception, the landscape has evolved. We understand the need to future-proof those capabilities for the exponentially increasing data and workload volumes customers expect today and tomorrow.

We will continue to support the broad use of MTLK while simultaneously architecting AI services for the future needs of data at a massive scale. This will enable organizations to run models with billions of parameters and get the granularity needed to deliver resilience at any scale.

Deliver differentiated AI outcomes with unparalleled data.

Each component of an organization's tech stack holds important context for an AI model. Network traffic captures usage trends, such as bandwidth peaks or latency issues. Endpoint logs contain a wealth of information, from incoming requests, errors, health and performance metrics, and sensor data. Security data can point to a wide range of threats within an environment, and observability data can track performance across IT systems. Applying AI to this vast dataset will improve detection, investigation and response.

Siloed data only reveals a fraction of the truth — but taken together, this data tells the complete story of a digital enterprise. Effective use of the right data at massive scale will be a critical currency to enabling AI success, helping organizations drive outcomes never before possible. Over time, Splunk will bring together aggregated and anonymized security and observability data and the right domain expertise to drive stronger capabilities, continuously improve AI model performance, and ultimately deliver better outcomes for our customers.

The future forecast: AI everywhere

The power of AI has barely begun to scratch the surface. McKinsey & Company predicts that generative AI will automate half of today's work activities between 2030 and 2060. The spectrum of AI technology ranges from building distributions and finding anomalies to now building LLMs with hundreds of billions of parameters.

Embracing the full spectrum of AI comes with both anxiety and excitement. We welcome the excitement of driving better, faster outcomes for digital resilience and making our customers' lives easier. We will seize the hope and possibility that accompanies this AI spring. We will also acknowledge the anxiety of this new season so we can recognize the unknown and find ways to alleviate it.

That means taking a step back and examining the big questions that AI brings: How can organizations trust it? How does it scale without breaking things? What does responsible, fair AI use look like? We don't claim to know everything yet, but with our philosophy acting as a compass, we can shed light on how AI can drive digital resilience forward.

Delve deeper into all things AI.

Learn how to catalyze your digital resilience with Splunk AI.

Get started



Generative AI elicits both fear and excitement in CISOs, according to our report. Discover more emerging trends, threats and strategies facing today's security leaders in [The CISO Report](#).



Get more executive insights and viewpoints on AI, security and observability from thought leaders and experts at [Perspectives by Splunk](#).



Learn how AI and ML can bolster your organization's [observability practice](#) and [security program](#) in our e-books about use cases.



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. © 2024 Splunk Inc. All rights reserved.

24-601501-Splunk-AI Philosophy Powering Dig Res-EB-116