Listicle

# 5 Big Myths of AI and Agentic AI Debunked

Learn how to adopt the next frontier of AI advancements in security and observability operations and position your organization for success

splunk>
a CISCO company

# AI has evolved at an unprecedented pace

It's hard to believe that artificial intelligence (AI) and machine learning (ML) have been around since the 1950s. While ML became mainstream in the early 2010s, more sophisticated types of AI have evolved from academic concepts to everyday business tools over time. The rise of large language models (LLMs) in late 2022 made AI more accessible than ever and drove adoption of generative AI (GenAI) tools. And now, the application and prominence of agentic AI have become the next frontier.

Today's AI can do far more than answer prompts. From summarizing noisy incidents and correlating telemetry to recommending next steps, autonomous agents leveraging agentic AI are becoming a force multiplier for both security and observability teams.

Agentic AI often comes with built-in expertise and can interpret and process information from its environment in a meaningful way, enabling it to make informed decisions and take actions. While this gives teams new advanced capabilities based on an inherent understanding of various domains to get ahead of security and performance issues, it also has the potential to upend traditional monitoring and SOC activities.

Splunk's State of Observability 2025 report indicates only 18% of respondents currently "often" or "always" leverage emerging AI technologies like agentic AI. But the report also predicts a rapid increase in the adoption of these technologies in coming years.

Is your team ready to take full advantage of AI superpowers in a reliable and safe way? In the following chapters, we'll tackle the most persistent myths holding teams back from realizing AI's full potential.

## AI Terms to Know

**Artificial Intelligence (AI)**
The broad field of creating machines and software that can perform tasks requiring human-like intelligence.

**Machine Learning (ML)**
A subset of AI where algorithms learn from data to make predictions or decisions.

**Generative AI (GenAI)**
AI models that can create new content, such as text, images, or code, based on training data and direct input (e.g., prompts).

**Large Language Model (LLM)**
A type of generative AI (GenAI) that is trained on very large text datasets to learn patterns in human language, enabling it to understand, generate, and summarize text.

**Agentic AI**
AI systems designed to take goal-oriented actions, plan steps, and use tools autonomously to achieve outcomes.

**Humans-in-the-loop**
A process where human judgment is a required part of the AI decision-making cycle, with people providing input, approval, or modification before actions are executed.

# Myth #1: Agentic AI is just like generative AI

**Truth:** In reality, it's a big step forward because it can automate and orchestrate tasks to drive outcomes.

It's easy to lump agentic AI into the same bucket as generative AI, as both transform how teams work. However, agentic AI is not simply GenAI in a new wrapper. It represents a fundamental shift from prompt-driven intelligence to a focus on proactive outcomes.

AI agents and GenAI are distinct but complementary technologies. Agents are goal-oriented and center on decision-making steps and executing linked actions, while GenAI focuses on creating content. In a combined system, the agent handles the workflow and reasoning, and the GenAI component provides the creative, human-like output, such as crafting a personalized email or a new recipe.

GenAI, which can be found in the form of a chatbot or assistant, is also more reactive, waiting for you to prompt each step, whereas agentic AI can be more proactive, working autonomously to solve problems. Both types of AI can come together to deliver even greater benefits. For example, agents can research fixes while GenAI creates a report explaining in clear language what it discovered.

In observability environments, this fundamentally transforms how teams handle incidents, troubleshoot, and operate efficiently. It can also help them do more with fewer resources by assigning agents objectives and access to specific knowledge bases and tools, with clear rules and thresholds.

Even when AI is permitted to act autonomously, every decision can be logged and audited, giving teams complete visibility into exactly what was done, why, and with what data. The result? A system that moves quickly yet accountably.

Access controls, audit trails, and continuous monitoring help ensure that "autonomous" never becomes "uncontrolled." Teams can fine-tune the balance between speed and caution, using AI as an extension of their capabilities rather than a replacement for human judgment.

# Myth #2: ChatGPT is all you need

**Truth:** GenAI tools can help automate many tasks but aren't always the best fit for your needs.

Public ChatGPT can be a handy tool for general, non-sensitive work, helping teams write marketing copy, draft sales emails, or summarize publicly available information. For marketing teams: need fresh ad text? Check. For sales: a quick, polished prospecting email? ChatGPT can do that. For general research: create a summary from public sources? Consider it done. But for specialized, business-specific tasks that demand deep knowledge of your systems and operations, public GenAI tools fall short.

All-purpose GenAI tools (like ChatGPT) are not always the best fit. With long or complex prompts, they can stray from the intended goal, while agentic AI is better at staying on track and following defined steps. Out of the box, public GenAI tools don't have access to your organization's proprietary machine data, and they don't know your business. Also, they are not domain-specific, which is a critical piece of the contextual puzzle. They do not understand the unique details of your specific systems, how your network is set up, or what normal performance limits are. They can't effectively handle all the specialized data you use to keep systems secure and running smoothly.

Public GenAI tools can also pose a security risk: using them with your organization's proprietary information can be a recipe for disaster. You put your company at risk of data breaches, intellectual property loss, privacy violations, and reputational damage. Sensitive information, such as client details or proprietary code, can be exposed through the tool's data handling and other potential vulnerabilities. Strong safety measures and clear AI usage policies are critical to safeguarding the information input into ChatGPT and other public generative AI chatbots.

It's important to note that any AI tool may have reliability issues, such as AI hallucinations. It's critical for SOC analysts and engineers to rely on fully trusted information for high-priority tasks, though this level of trust cannot always be guaranteed.

# Myth #3: AI will make runaway decisions without you

**Truth:** You can control how much you want to automate and how frequently you want a human in the loop — such as for business-critical, high-risk decisions and actions.

Some may fear that leveraging agentic AI across SOC, IT Ops, and Engineering teams could lead to decisions gone rogue. The truth is that with proper oversight, teams can mitigate this risk.

As AI agents work alongside SOC analysts and engineers, they surface insights the moment they are needed — or even before. In a security setting, AI agents can triage, conduct malware analysis, author playbooks, and execute other tasks. On the observability side, AI can automatically detect issues, troubleshoot, and find root causes — and run low-risk remediations, either autonomously or with a human in the loop.

A robust framework of controls is essential to ensure that human analysts and engineers remain central to key decisions and have the mechanisms for review and override. This allows AI agents to handle routine cases with decision autonomy, while more high-stakes decisions necessitate human approval. Agentic systems also allow for the implementation of well-defined policies and workflows exactly as we want them. Through this and other means, they allow us to introduce domain knowledge.

Establishing clear operational roles for AI agents and defining the level of AI autonomy are measures teams should take to ensure agentic AI is not running unchecked. Levels of autonomy can range from manual to fully autonomous, with specific roles defined as human in the loop and trust thresholds aligned to specific tasks.

Continuous feedback loops are also critical for AI agents to learn and refine their approach. To support this, human analysts and engineers should have the ability to rate AI performance and output. This human intervention and validation serve as a training opportunity, ensuring the system continuously improves over time and helps mitigate the risk of any rogue decisions.

# Myth #4: AI will take your job

**Truth:** AI works best as a partner, extending your capabilities, accelerating innovation, and relying on human judgment in the high-risk situations where oversight is essential.

Some may fear that, without their involvement, AI could run rogue. Others worry that AI might replace them altogether. For the best business results, it's critical for teams to retain human talent to supervise and sign off on recommendations from AI agents. Domain-specific intelligence from agentic AI dramatically reduces manual work, filters noise, and accelerates investigations and remediation. People bring the judgment, oversight, and accountability that keep operations aligned to business needs and technical realities.

For resource-constrained teams, the key to success is finding the right balance. While the autonomy of agentic AI may seem daunting to some, these systems work best within human-defined boundaries such as approval points, thresholds, and escalation rules.

This working model also builds trust. SOC and IT Ops teams know what data was used, why a recommendation was made, and where controls are in place. Because every decision can be traced back through its signals and logic, teams can confirm not only what AI agents have surfaced, but why the agent surfaced it. That transparency turns system outputs into actions people can rely on.

Think of AI as a power tool that expands what you can achieve. Routine, low-risk tasks can be handled with minimal oversight, giving you more time to focus on complex, mission-critical challenges where your experience matters most. The real advantage comes when human ingenuity and AI speed work together to solve problems, uncover new opportunities, and deliver results that neither could achieve alone.

Speed and scale matter, but so do operational context, informed judgment, and ownership of decisions. Domain expertise ensures accuracy, relevance, and safety, while AI agents broaden visibility and accelerate outcomes. There's plenty of work to do now and in the future, and competitive advantage will belong to the organizations whose teams combine their expertise with tools designed to work alongside them.

# Myth #5: Leveraging AI agents requires a data science dream team and deep pockets

**Truth:** Implementing agentic AI tools is less complex and resource-intensive than you may think.

Historically, implementing advanced AI and ML solutions required significant expertise in data science. Given the complexity of security and observability systems — enormous volumes of diverse and often unstructured data — it's easy to see why this was the case. Add into the mix the unique environments, specific threats, and custom applications in which practitioners work, adopting and implementing emerging AI tools may seem daunting.

Today, that's no longer the case. Imagine having smart assistants already trained and ready to go out of the box. Unlocking value from these tools will still require some AI skills, but it is more straightforward than in the past. Many such tools use open standards, meaning you can easily add to and personalize AI features — like adjusting specific alerts or rules — without reinventing the wheel.

This means you don't need to be an AI expert, or spend weeks setting things up. Powerful AI features are now often part of the security and observability tools you already use, making it easy for anyone on your team to use and benefit from them.

One example of how open standards are making AI easier to use is Modal Context Protocol (MCP), a flexible framework for sharing context between many different tools and systems. MCP isn't limited to agentic AI, and it can help connect and coordinate AI capabilities across a wide range of applications, but it's especially valuable when building or expanding agents.

By giving these agents a common, ready-made way to "talk" to other platforms, MCP cuts down on custom integration work and speeds up the time it takes to get them doing real, valuable tasks. It's another sign of how modern AI adoption is shifting from heavy, bespoke projects to quick, outcome-focused wins, often by pairing flexible standards like MCP with ready-to-use assistants.

Aside from the major costs of paying to train a model from scratch and potential delays when building one, you can also save yourself a lot of hassle with out-of-the-box assistants and agents. AI agents can deliver the intelligence you need and cut costs, without the DIY headache.

# Reality: AI is a journey

As we have seen with the widespread adoption of GenAI tools, the teams that can smartly leverage agentic AI capabilities into their workflows to make their jobs easier will derive the greatest value from emerging AI capabilities.

Today, SOC, IT operations, and Engineering teams alike are accelerating detection, investigation, and response, using both GenAI and agentic AI. It is a big step forward from GenAI tools to agentic AI, and the teams that are evolving from working with chatbots to implementing agentic AI with human supervision are poised to hone their competitive advantage. Finally, out-of-the-box AI agents make it easier than ever to do this, with flexibility and minimal training.

Learn more about how your security and observability teams can take their operations to the next level and do more with less by incorporating AI into their workflows.

splunk> a cisco company



25_CMP_listicle_5-big-myths-of-AI-and-agentic-AI-debunked_v6