Analysts share 8 ChatGPT security predictions for 2023

Tim Keary

@tim_keary

March 23, 2023 7:07 AM

Image Credit: Image by Louis Rosenberg via Midjourney

The release of ChatGPT-4 last week shook the world, but the jury is still out on what it means for the data security landscape. On one side of the coin, generating malware and ransomware is easier than ever before. On the other, there are a range of new defensive use cases.

Recently, VentureBeat spoke to some of the world's top cybersecurity analysts to gather their predictions for ChatGPT and generative AI in 2023. The experts' predictions include:

- ChatGPT will lower the barrier to entry for cybercrime.
- Crafting convincing phishing emails will become easier.
- Organizations will need AI-literate security professionals.
- Enterprises will need to validate generative AI output.
- Generative AI will upscale existing threats.
- Companies will define expectations for ChatGPT use.
- AI will augment the human element.
- Organizations will still face the same old threats.

Below is an edited transcript of their responses.

1. ChatGPT will lower the barrier to entry for cybercrime

“ChatGPT lowers the barrier to entry, making technology that traditionally required highly skilled individuals and substantial funding available to anyone with access to the internet. Less-skilled attackers now have the means to generate malicious code in bulk.

“For example, they can ask the program to write code that will generate text messages to hundreds of individuals, much as a non-criminal marketing team might. Instead of taking the recipient to a safe site, it directs them to a site with a malicious payload. The code in and of itself isn’t malicious, but it can be used to deliver dangerous content.

“As with any new or emerging technology or application, there are pros and cons. ChatGPT will be used by both good and bad actors, and the cybersecurity community must remain vigilant to the ways it can be exploited.”

— Steve Grobman, senior vice president and chief technology officer, McAfee

2. Crafting convincing phishing emails will become easier

“Broadly, generative AI is a tool, and like all tools, it can be used for good or nefarious purposes. There have already been a number of use cases cited where threat actors and curious researchers are crafting more convincing phishing emails, generating baseline malicious code and scripts to launch potential attacks, or even just querying better, faster intelligence.

“But for every misuse case, there will continue to be controls put in place to counter them; that’s the nature of cybersecurity — a neverending race to outpace the adversary and outgun the defender.

“As with any tool that can be used for harm, guardrails and protections must be put in place to protect the public from misuse. There’s a very fine ethical line between experimentation and exploitation.”

— Justin Greis, partner, McKinsey & Company

3. Organizations will need AI-literate security professionals

“ChatGPT has already taken the world by storm, but we’re still barely in the infancy stages regarding its impact on the cybersecurity landscape. It signifies the beginning of a new era for AI/ML adoption on both sides of the dividing line, less because of what ChatGPT can do and more because it has forced AI/ML into the public spotlight.

“On the one hand, ChatGPT could potentially be leveraged to democratize social engineering — giving inexperienced threat actors the newfound capability to generate pretexting scams quickly and easily, deploying sophisticated phishing attacks at scale.

“On the other hand, when it comes to creating novel attacks or defenses, ChatGPT is much less capable. This isn’t a failure, because we are asking it to do something it was not trained to do.

“What does this mean for security professionals? Can we safely ignore ChatGPT? No. As security professionals, many of us have already tested ChatGPT to see how well it could perform basic functions. Can it write our pen test proposals? Phishing pretext? How about helping set up attack infrastructure and C2? So far, there have been mixed results.

“However, the bigger conversation for security is not about ChatGPT. It’s about whether or not we have people in security roles today who understand how to build, use and interpret AI/ML technologies.”

— David Hoelzer, SANS fellow at the SANS Institute

4. Enterprises will need to validate generative AI output

"In some cases, when security staff do not validate its outputs, ChatGPT will cause more problems than it solves. For example, it will inevitably miss vulnerabilities and give companies a false sense of security.

"Similarly, it will miss phishing attacks it is told to detect. It will provide incorrect or outdated threat intelligence.

"So we will definitely see cases in 2023 where ChatGPT will be responsible for missing attacks and vulnerabilities that lead to data breaches at the organizations using it."

— Avivah Litan, Gartner analyst

5. Generative AI will upscale existing threats

"Like a lot of new technologies, I don't think ChatGPT will introduce new threats — I think the biggest change it will make to the security landscape is scaling, accelerating and enhancing existing threats, specifically phishing.

"At a basic level, ChatGPT can provide attackers with grammatically correct phishing emails, something that we don't always see today.

"While ChatGPT is still an offline service, it's only a matter of time before threat actors start combining internet access, automation and AI to create persistent advanced attacks.

"With chatbots, you won't need a human spammer to write the lures. Instead, they could write a script that says 'Use internet data to gain familiarity with so-and-so and keep messaging them until they click on a link.'

"Phishing is still one of the top causes of cybersecurity breaches. Having a natural language bot use distributed spear-phishing tools to work at scale on hundreds of users simultaneously will make it even harder for security teams to do their jobs."

— Rob Hughes, chief information security officer at RSA

6. Companies will define expectations for ChatGPT use

"As organizations explore use cases for ChatGPT, security will be top of mind. The following are some steps to help get ahead of the hype in 2023:

1. Set expectations for how ChatGPT and similar solutions should be used in an enterprise context. Develop acceptable use policies; define a list of all approved solutions, use cases and data that staff can rely on; and require that checks be established to validate the accuracy of responses.

2. Establish internal processes to review the implications and evolution of regulations regarding the use of cognitive automation solutions, particularly the management of intellectual property, personal data, and inclusion and diversity where appropriate.

3. Implement technical cyber controls, paying special attention to testing code for operational resilience and scanning for malicious payloads. Other controls include, but are not limited to: multifactor authentication and enabling access only to authorized users; application of data loss-prevention solutions; processes to ensure all code produced by the tool undergoes standard reviews and cannot be directly copied into production environments; and configuration of web filtering to provide alerts when staff accesses non-approved solutions."

— Matt Miller, principal, cyber security services, KPMG

7. AI will augment the human element

"Like most new technologies, ChatGPT will be a resource for adversaries and defenders alike, with adversarial use cases including recon and defenders seeking best practices as well as threat intelligence markets. And as with other ChatGPT use cases, mileage will vary as users test the fidelity of the responses as the system is trained on an already large and continually growing corpus of data.

“While use cases will expand on both sides of the equation, sharing threat intel for threat hunting and updating rules and defense models amongst members in a cohort is promising. ChatGPT is another example, however, of AI augmenting, not replacing, the human element required to apply context in any type of threat investigation.”

— Doug Cahill, senior vice president, analyst services and senior analyst at ESG

8. Organizations will still face the same old threats

“While ChatGPT is a powerful language generation model, this technology is not a standalone tool and cannot operate independently. It relies on user input and is limited by the data it has been trained on.

“For example, phishing text generated by the model still needs to be sent from an email account and point to a website. These are both traditional indicators that can be analyzed to help with the detection.

“Although ChatGPT has the capability to write exploits and payloads, tests have revealed that the features do not work as well as initially suggested. The platform can also write malware; while these codes are already available online and can be found on various forums, ChatGPT makes it more accessible to the masses.

“However, the variation is still limited, making it simple to detect such malware with behavior-based detection and other methods. ChatGPT is not designed to specifically target or exploit vulnerabilities; however, it may increase the frequency of automated or impersonated messages. It lowers the entry bar for cybercriminals, but it won’t invite completely new attack methods for already established groups.”

— Candid Wuest, VP of global research at Acronis

Daily insights on business use cases with VB Daily

If you want to impress your boss, VB Daily has you covered. We give you the inside scoop on what companies are doing with generative AI, from regulatory shifts to practical deployments, so you can share insights for maximum ROI.

Your Email

Subscribe Now

Read our Privacy Policy

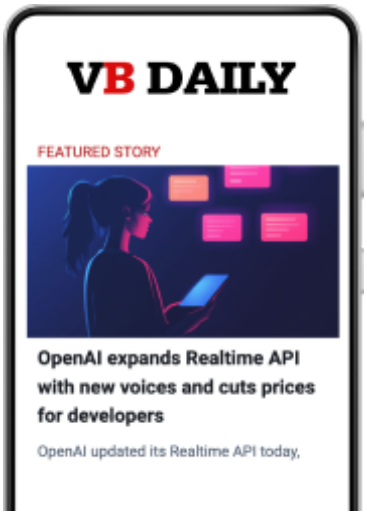

Find Your Place In The World

Senior Actuarial Associate
Forecasting - Remote
DW Simpson
San Francisco

See Job

Quality Assurance Specialist
Inceed
Houston

See Job

Manager Small to Medium Business Sales
Optimum
Bethpage

See Job

Product Engineer
Thrivve by intelletec
City of London

See Job

Search More Roles

Partner Content

Cisco explains why the network is AI's looming bottleneck

VB Staff

June 10, 2025 06:00 AM

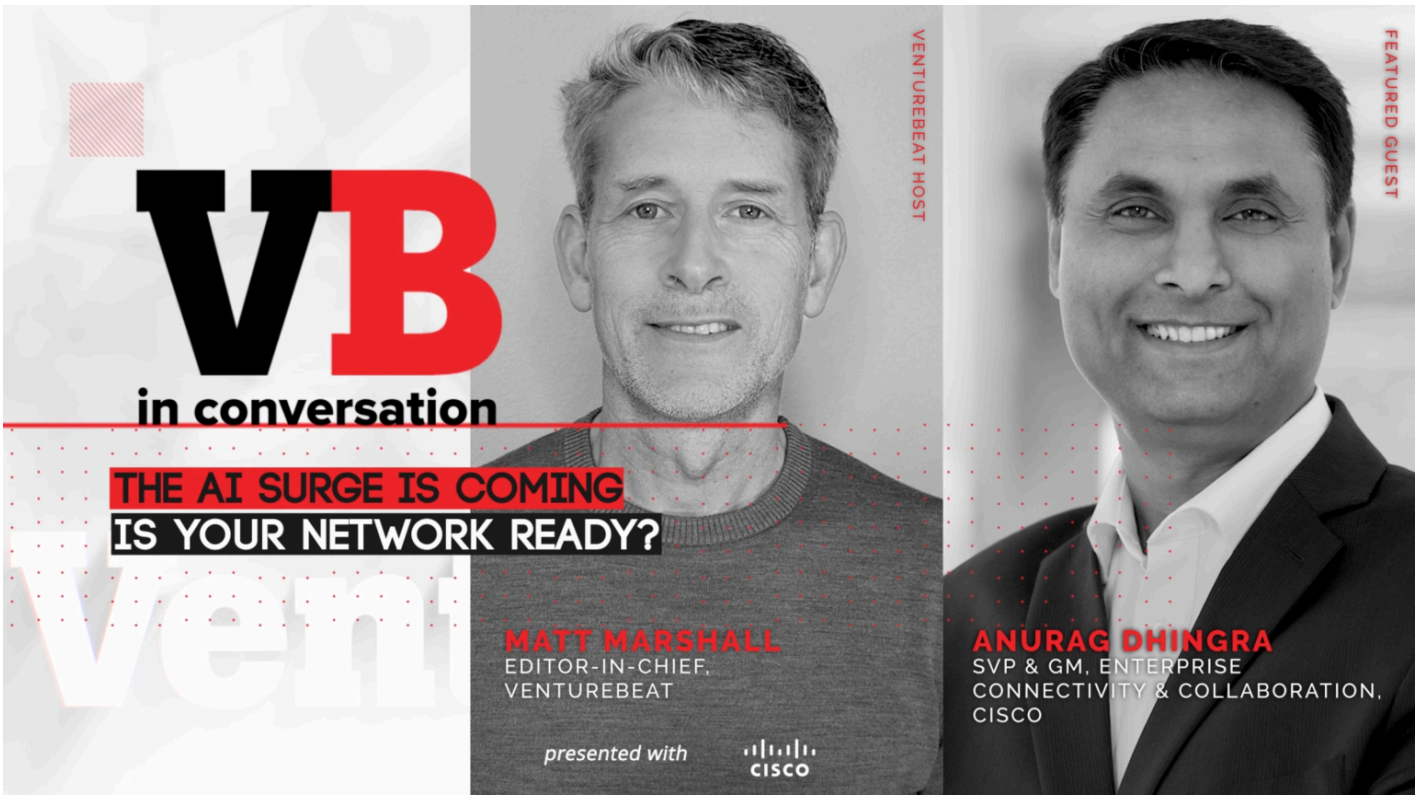

Presented by Cisco

Few enterprise technology shifts are more significant than the rise of AI agents. This means a leap in digital headcount that will force every CIO to revisit infrastructure once taken for granted: the network.

Matt Marshall, editor in chief of VentureBeat, welcomed Anurag Dhingra, SVP and GM of enterprise connectivity and collaboration at Cisco, to the latest VentureBeat in Conversation. They talked through where networks fall short, what AI-ready really means and how companies can turn connectivity into a competitive edge.

“When agents — whether they’re embodied AI or software agents — become ubiquitous, it’s going to feel like the workforce went through a step function increase,” Dhingra said. “Of course, network traffic is going to go through the roof as a result. So, you have to be very sure you have the right infrastructure to be ready for that AI future. Unfortunately, not many organizations today do just yet, despite the fact that this is not an optional technology going forward.”

New architecture for a new era

“Connectivity from the office environment to the internet had to be reimagined and rewired for the transition to cloud and SaaS, and I believe we’re at the cusp of a similar technology transition,” Dhingra says. “It’s not just about agentic AI. It’s around taking a step back to start thinking about what the right architecture for this new era — AI-ready and secure — looks like.”

There are three elements to consider: first is ensuring you have scalable and secure devices that are ready for all the traffic and the security exposure that’s expected in this new world. Second is about fusing security into the fabric of the network, so all of the new threats that are surfacing are manageable. And then finally, AgenticOps, or leveraging AI to help simplify operations, which is a boon to IT departments that are struggling with shortages of skilled people.

Beyond the question of wired vs. wireless

Cisco has delivered on its promise to converge its full portfolio across switching, routing, wireless, and even industrial networking, into a single, unified platform. Through this platform – manageable on-premises, in the cloud, or a mix of both – customers benefit from end-to-end network assurance, rich telemetry and insights — even into unowned infrastructure online that connects you to cloud services. All of that is accessible through a conversational agentic AI interface that helps users make sense of the data and can also take actions on the user’s behalf.

The Cisco C9350 and C9610 Smart Switches enable a more efficient and extensible architecture. These new devices can run networking and security processes in parallel, without impacting the performance of the switch, by integrating services directly with the network, rather than bolting them on top. By combining their own Cisco Silicon One ASICs with this co-processing capability, they’re not only weaving security directly into the fabric of the network, they’re making it easy for IT teams to scale and adapt networks as business needs change, without adding expensive new compute or new hardware. And that’s where cost savings emerge: hardware consolidation, reduced power consumption, and operational simplicity.

This approach builds on the journey Cisco began with smart switches in the data center, where security was directly embedded into the switching layer with AI-native, hardware-accelerated architectures like Hypershield. By integrating security into the fabric itself, smart switches reduce the need for additional appliances and enable data center operators to create micro perimeters around every service that makes up a workload.

But now that traffic is increasingly impacting the campus network, Cisco is extending this same concept to campus environments. The switches incorporate two processing engines: a high-performance network processor for stable data transfer, and a network services co-processor for agile security processing, with traffic steered between them intelligently, to ensure performance remains fast and stable.

Boosting network security

With new attacks emerging that target both network infrastructure and encrypted data flowing across the switch, security for the campus and branch networks requires a multilayered approach. It starts with securing the device itself, and then one layer up is connectivity, and the top layer is securing users, applications, and devices.

“When you look at the core of this layer stack, how do we build trustworthy systems?” Dhingra says. “This is where our custom ASIC, Silicon One, has some amazing security capabilities. We can provide tamper-proof devices that are ready for post-quantum cryptography. One layer up, you can also encrypt data flows with post-quantum crypto. On top, where capabilities like segmentation and universal zero trust network access sit, you can define policies in one place, and you push those policies out to every network element.”

Each device has a context that carries across user identity, device identity and machine identity. With a policy in that context, each device can apply the right type of security policy so that enforcement becomes truly distributed, but policy is managed in one place. It’s a multilayered approach to address new risks that are emerging.

Security policies are automatically updated to the right enforcement points to keep security posture up to date. Plus, organizations can make policy lifecycle management work at scale by using self-qualifying policy updates before deployment, and extend consistent policy enforcement across multiple domains. Because policies can be managed across a library of enforcement points in the cloud, on-prem and on traditional next-gen firewalls, customers now have a single management system with Cisco Hybrid Mesh Firewall.

The solution, seamlessly integrated into existing processes, supports common and separate workflows for NetOps, SecOps or NetSecOps teams using a single solution to maintain connectivity and security.

AI assistants for the network

Today many developers are using AI to write code, and in some cases the agents writing that code are completely autonomous — but in all cases, a human evaluates that code before it’s put into production. Cisco’s new agentic assistants for network management take inspiration from those workflows, by offering configuration change suggestions based on what they’re observing in the network, which are reviewed before they’re deployed to put guardrails against the possibility that these models will hallucinate.

Cisco has also introduced a collaborative workspace for NetOps, SecOps and DevOps teams with AI Canvas.

“This is how we’re going to create a next step forward in AgenticOps with humans in the loop, simplifying the workflow to deliver amazing experiences to people,” Dhingra explains.

An admin who’s using an AI assistant with a conversational model to troubleshoot a network can pull in another coworker into a collaborative canvas that incorporates all the charts and graphs being dynamically generated by the AI assistant and invite them to take a look and help unknot any issues.

Ensuring reliability and service across the whole network

“We think of assurance as a promise,” Dhingra says. “It’s a promise to deliver amazing experiences to people who are connecting to the network. That’s why we build capabilities that are not just about monitoring what’s going on, but proactively and predictively taking steps to not have any incidents or any outages to begin with. That’s the holy grail, the journey we’re on.”

In the meantime, it’s about getting visibility across every piece in your infrastruct6ure, and every piece that is outside of your infrastructure. Cisco takes a layered approach, with rich telemetry coming out of devices across the switching, routing and wireless domains, which can be stitched together to make sense of the data. Layered on top are synthetic tests, which are enabled by ThousandEyes.

The product offers a single time-correlated view of critical metrics, based on all the networks and services that make up the user experience. Its cross-correlation algorithms produce cross-layer telemetry as well as interactive visuals to make it easy to plan service rollouts, isolate problems, take action, and resolve issues faster.

“We’re taking a step forward in the ability to apply AI to simplify operations,” Dhingra added. “It’s really about operations at scale. I’m very excited about the new capabilities both within Cisco and as agentic AI evolves.”