CYBER PROTECTION
MAGAZINE
3 || 2025

AI Agents and Bots

the authentication crisis

quantum computing: prepare for q-day

Agentic AI Reality Check + A Brief History of Bots

CONTENT

3.25

Identity is the key to NHI management

Welcome to the third special issue of 2025. This issues looks at bots, or as the cognoscenti refer to them: Non-human identities (NHI). Along with a history of of NHI, we have a diverse set of opinions on the use of NHI.

We start with cautiously optimistic views regarding restrained development. Caique Zaniolo of Cyferd advises building NHI with intent to avoid the security risks. Consultant Wyatt Mayham suggests knowing who is what on the internet with up-to-date authentication of all identities. Milankumar Rana, a software engineer from FedEx compares what is promised and what is actually possible with Agentic AI. Even StrikeReady CEO Alex Lanstein sees this gap, in the form of misunderstanding the application of agentic AI.

Vinod Goje of Bank of America and Peter Horadan of the identity company range from cautiously optimistic about the rise of NHI to absolutely Pollyanna-ish.

But nowhere is the impact of NHI and AI more dangerous and damaging than the area of marketing. Traditional SEO approaches are being hammered. Web traffic is crashing for publications and companies as search engines add AI boxes to the top of the search pages. Deppa Chauhan, senior SEO specialist at Accelirate, looks into that impact.

Finally, Ashley Rose, CEO of Living Security delivers a quite unique approach to the issue of human relations management. She suggests NHI be treated the same as actual humans with the same regulation and oversight.

In what seems to be a 180-degree turn from NHI is a look at the affect of quantum computing on encryption, or rather breaking encryption. Christina Cravens of Redjack and David Close of Futurex sound the conventional-wisdom alarm about the approaching Q-Day apocalypse.

However, we remain cautiously sanguine about the danger of quantum decryption. No country has the budget to build a quantum computer that can break an encryption of a single document inside eight hours. Quantum computing holds a great deal of promise in research and business. As long as people are manipulated into giving away encryption keys, no one needs a computer that big.

A brief history of bots

AUTHOR: LOU COVEY

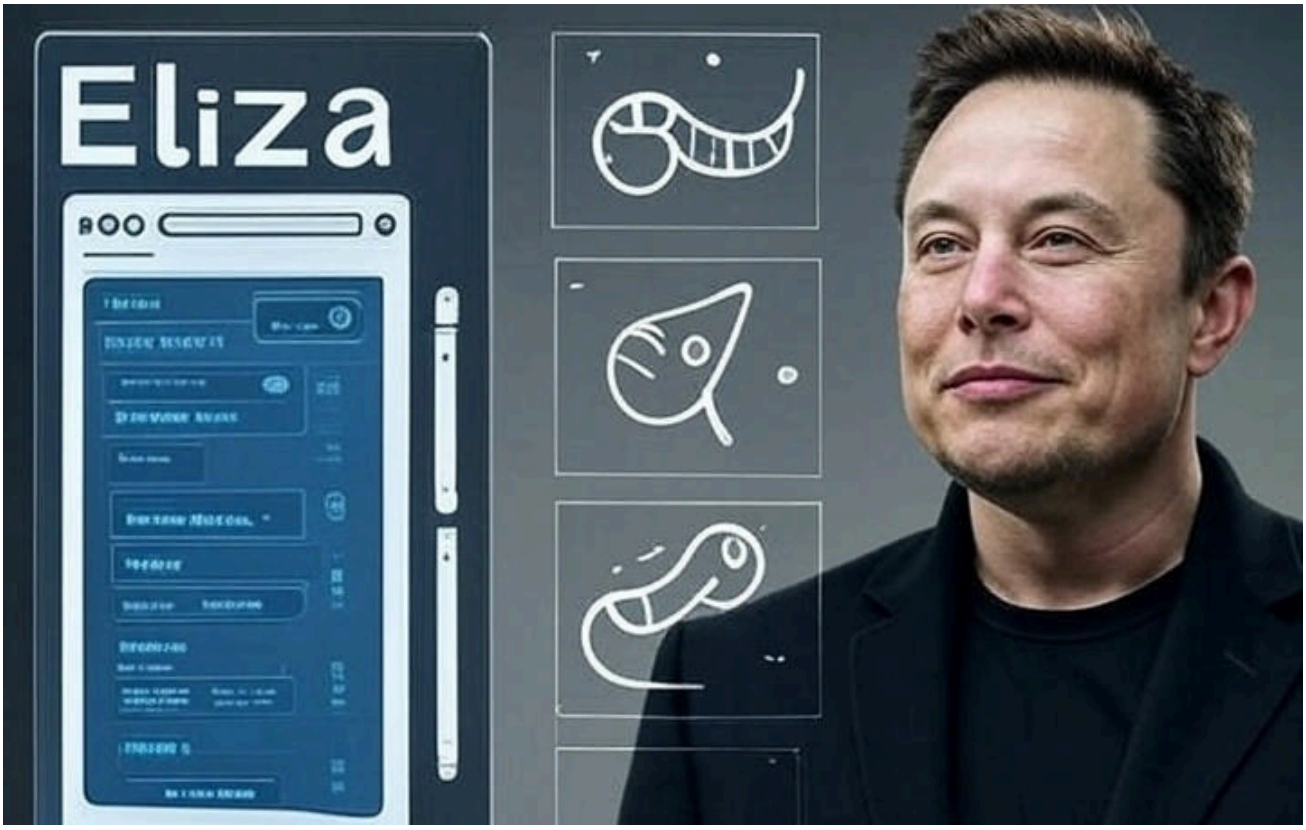

From ELIZA to worms and viruses to Elon Musk, what we call agentic AI is merely the latest evolution of automated programs that are less a beneficial tool as they are a means for deceiving the public

Bots are everywhere. We all use them, some of us innocently and a few not so much. They are referenced in both technological and common language, but few people understand what they are, what they do or what damage they do. This article is one of a series of content that will give the basics of this pervasive category of software.

Bots have been around for more than half a century to automate repetitive tasks and provide services on early internet platforms. The first was ELIZA, developed as a research project in 1966 at the Massachusetts Institute of Technology (MIT) the goal was to simulate conversations with a human being. ELIZA conversed with users, although it did not understand what the user was saying. Artificial intelligence chatbots are much more sophisticated versions of ELIZA, but still lack human comprehension.

BOTS NOT REPLACEMENTS

The purpose of ELIZA was to determine if computers could replace psychoanalysts. Consequentially, it was the first time the prediction that computer could replace humans had some hard evidence. Today, there are mental-health AI applications with not much better results than ELIZA but projected to have a $8 billion market by 2032.

In 1988, the earliest broad use of bots was Internet Relay Chat (IRC) automating user list management, searches, and providing services like weather updates or game scores. But these were not known as bots at the time. They were called automations and still required a human interface to operate,

The term bots came into fashion in the 1990s to define software that could operate independently of

human interaction, the term was an abbreviation of "Robot. This was the time that web crawlers or "spiders" would browse the web and index pages for search engines. This was the early stage of search engines.

MALICIOUS INTENT

The history of malicious bots began almost immediately after they were introduced to the general market.

Created in 1988 by a Cornell University grad student working at MIT, the Morris worm was an experiment that went wrong very quickly. It crashed 10 percent of all the computers connected to the internet in 20 countries demonstrating, for the first time, how automated attacks could devastate the world. The program is called a worm because it could replicate itself, rather than destroy computer programs and data like a virus, but it reduced computing speeds to a crawl, like a worm moving along the ground.

Webcrawlers emerged in 1994 to facilitate search engine development, but they came with first malicious botnets. These automated computer programs embedded themselves in thousands of computers for DDoS attacks, spam campaigns, data theft, click fraud, and malware distribution. Today, disreputable "crypto miners" use botnets to boost compute power for cryptocurrency production, without permission from the owners of infected computers.

DISINFORMATION BREEDING GROUNDS

While OpenAI's GPT-3 was a quantum leap forward for the technology from the introduction of ELIZA, but it also provided a superhighway for malicious bots to spread disinformation in sophisticated cyberattacks. According to a report from Imperva, a Thales company, almost 50% of internet traffic comes from non-human sources. "Bad bots, in particular, now comprise nearly one-third of all traffic. Bad bots have become more advanced and evasive, and now mimic human behavior in such a way that it makes them difficult to detect and prevent."

Malicious bots, from early worms and spambots to more advanced and sophisticated attacks that inflate traffic on social media platforms. For instance, estimates of bot followers of Elon Musk range between 23 percent (according to X corporate estimates) to 65% (according to X's own generative AI platform Grok). Meta claims Mark Zuckerberg has close to 16 million followers, but sources inside Meta says those reports add duplicates of followers on Facebook, Instagram and WhatsApp, and that close to half of those followers are bots. That brings Zuckerberg's followers down to a still impressive 3-4 million people.

CORPORATE ENCOURAGEMENT OF FRAUD

That over estimation of influence manipulates public discourse as the bots repeat, share, and hashtag trending topics. This leads to the spread of misinformation and false narratives. Sarah Wynn-Williams, in her book *Careless People*, explained how bots on Facebook played in 2016 US presidential elections by influencing public opinion. During that election season, the Trump campaign was the single largest advertiser on the Facebook platforms.

Bots also reduce the reach of authentic content from real users. They trick algorithms into prioritizing their content, harming businesses and individuals who rely on social media for marketing and engagement. This inflation of traffic by bots leads to increased costs for businesses. Bots can drive up infrastructure and customer support costs by degrading online services and inflating traffic. A bigger problem is the fraud that internet companies encourage to promote their own revenue stream.

While social media platforms like Google, Meta and X have policies against bot activity, they still make a lot of money on them.

STEALING AD BUDGETS

Raluca Saceanu, CEO of SmartTech247 in the UK, talked to Cyber Protection Magazine about the enormous cost she incurs from fraud on Google advertising services.

"We use Google ads as part of a large digital marketing strategy. We offer a free risk assessment and getting 70, 80, 90 leads a day from huge companies like Pfizer. Everything was basically bogus. It was costing us thousands of UK pounds."

Saceanu said they reported the issue to Google and got insufficient answers. But she was able to determine they were bots trying to exfiltrate customer data. None of the attempts succeeded. She pulled the ads, but is still negotiating with Google for refunds, a process taking months. In the meantime, she is out a significant amount of money that Google is earning interest on.

NOWHERE TO HIDE

If a savvy cybersecurity professional like Saceanu can be victimized by bots, facilitated by large

multinational corporations, what hope do less technologically adept business owners have?

The next incarnation of bots will be agentic AI, but when is another question. While many companies, most notably Salesforce, are pushing this concept as the next big thing, it doesn't really exist.

Agentic AI supposedly automates tasks set by the user but also anticipates the user's needs according to set parameters. This technology is used maliciously more than it is commercially available to the public. Bad actors have simpler needs: separate as much money from victims quickly as possible. To date, however, there have been no known use of agentic AI for malicious or beneficial applications. So, is agentic AI hype, horror, or something in between? That's next.

LOU COVEY

Lou is Chief Editor of Cyber Protection Magazine

CONTRIBUTING AUTHORS:

Caique Zaniolo

Wyatt Mayam

Milankumar Rana

Vinod Goje

Alex Lanstein

Christina Cravens

David Close

SPONSORS:

Special thanks to our sponsors for making this possible:

Safety National – safetynational.com

Lupasafe – lupasafe.com

EDITORS: Lou Covey & Patrick Boch

Issued by

CYBER PROTECTION MAGAZINE

www.cyberprotection-magazine.com

info@cyberprotection-magazine.com

Why Automation Should Be Built with Intent

AUTHOR: CAIQUE ZANIOLO

Caique Zaniolo, VP of Product at Cyferd

There's nothing inherently wrong with bots. In fact, when they're used well, they're one of the most powerful tools in a business's digital toolkit — streamlining repetitive tasks, freeing up human time, and delivering consistent results. The issue isn't the bots themselves. It's how carelessly we've started deploying them.

The race to automate has outpaced the intent behind it. Bots are now being slapped onto every system, process, and customer touchpoint like duct tape on a leaking pipe — quick, cheap, and temporary. The result? Frustrated users, disjointed workflows, and a growing sense that we've confused doing more with doing better.

BOTS THAT ACTUALLY WORK

That's not to say bots don't have their place. Quite the opposite. In customer service, bots can be invaluable for managing FAQs, processing simple requests, and directing people to the right human support when things go off-script. They reduce queue times, respond instantly, and never lose their patience. When designed properly, they actually improve the customer experience by removing friction — not adding it.

In finance, bots handle repetitive tasks like invoice approvals, compliance checks, and payment reminders, quietly eliminating hours of manual admin. HR and IT teams rely on them for onboarding flows,

password resets, and responding to basic policy questions.
These aren't headline-grabbing use cases — they're just quietly effective. The kind of tasks bots are genuinely good at. What they've mastered is doing the same thing over and over, quickly, reliably, and without complaint. That's not a limitation — it's a strength.

When we let them stick to that, bots deliver real, measurable value. The problems start when bots are deployed without a clear purpose. Too often, they're introduced into workflows with no real thought about what they're supposed to achieve — or what happens when they don't. That's when automation turns into performance.

A chatbot that loops endlessly with no escalation route. A virtual assistant that can't understand the question it's answering. A system that looks smart but makes the process harder. This kind of automation theater usually stems from a lack of intent at the design stage.

START WITH PURPOSE, NOT HYPE

Effective bots are built around three simple principles: purpose, oversight, and integration. That means clearly defining the task the bot is meant to handle. Putting in guardrails and escalation paths when it gets something wrong. And embedding it within existing human workflows — not adding it as a clunky extra step.

When you get those things right, bots support the system. Skip them, and you end up with confusion, complaints, and angry customers taking their problems elsewhere.

The organisations doing this well start small and build with care. They pilot bots on simple, high-volume tasks like FAQs. They measure everything — from resolution rates to handoff frequency to user sentiment. They don't assume the job is done once the bot goes live; they monitor, adjust, retrain, and — when necessary — pull the plug. Because just having a bot doesn't mean you should use it everywhere.

BETTER BOTS, NOT MORE

Bots aren't magic. But they don't need to be. They can make life easier for both customers and employees — if they're built and deployed with intention. Without that clarity, they drift from helpful tool to broken promise.
So before you launch a bot, ask the right questions. What role is it playing? How does it add value? Who's responsible for keeping it effective? And what happens when it fails?

If you don't have clear answers, don't go live. Because the truth is simple: better bots come from better thinking. A bot built on purpose is a bot that builds trust.

CAIQUE ZANIOLO

VP of Product at Cyferd

The Authentication Crisis

AUTHOR: WYATT MAYHAM

We're living through the biggest shift the internet has seen since its creation. It's not just that bots exist online now. It's that they're on track to become the majority. And we're running out of reliable ways to tell who's real.

Recent estimates suggest that around 50% of all internet activity comes from bots, with bad bots accounting for 32% of global web traffic. For every two interactions online, only one involves a human. Reddit is flooded with AI-generated posts that rack up thousands of upvotes before anyone realizes they're fake. Moderators are overwhelmed and, in many cases, are turning to AI just to keep up.

On social media, the picture looks even worse. Analysis shows that roughly 48% of Elon Musk's Twitter followers are fake, while other research suggests up to 64% of accounts on X could be bots. Meanwhile, Meta is moving to integrate AI agents across Facebook and Instagram. When the people running these platforms are building in bots as a feature, this stops being a spam issue. It's a collapse in authenticity.

For businesses, this screws with everything. If you're running ads, how do you know your audience is even human? If you're doing market research or tracking customer feedback, how do you trust what's real? Companies are spending serious budget targeting ghosts and building strategies on artificial engagement.

But this isn't just about wasted marketing spend. The deeper issue is trust. Every online interaction becomes questionable when you can't tell who—or what—is on the other end. Reviews, customer service chats, influencer posts, even job applications are now polluted by AI-generated content.

The old tools for fighting this don't hold up anymore. CAPTCHAs, email verification, and phone number

checks slow things down but don't stop anything. AI now solves reCAPTCHA with 100% accuracy, surpassing even humans who typically achieve only 50-85% accuracy. Even more advanced systems like behavioral analysis are starting to fail. Bots are already mimicking human scrolling patterns, typing rhythms, and cursor movements with scary accuracy.

Some companies have started rolling out verification badges or identity-based logins. But those bring their own problems. They come with privacy tradeoffs. They create tiered systems that treat unverified users like second-class citizens. And they still don't stop well-designed bots from blending in.

The real issue is scale. Bots now operate as coordinated networks. They have names, faces, backstories, and entire posting histories. They cross-post, build followings, and inject themselves into any public conversation. We're already seeing them influence political debates, financial markets, and cultural trends on nearly every major platform.

And here's the worst part. This is happening with the basic version of AI. We're still in the prompt-driven phase. These systems follow instructions. But they're already good enough to fool us most of the time. When more advanced, autonomous AI shows up in the kind of way that can plan, adapt, and operate on its own and this problem won't just get worse. It will get out of control.

"The authentication crisis isn't coming, it's already here. Every day we delay building robust human verification systems, we're essentially handing over more of the internet to bots. The companies that solve this first will own the future of authentic digital interaction," said Wyatt Mayham, founder of Northwest AI.

Solving this isn't about patching leaks. We need to rethink how human identity works online. That could mean digital ID systems backed by governments. It could mean blockchain-based verification. It might mean something we haven't even invented yet. But it has to happen soon.
Organizations need to prepare for this now. Building systems that verify human users at scale isn't just a security feature. It's survival. The companies that figure out how to keep their ecosystems human will have a massive edge over the ones that don't.

The authentication crisis is already here. The question is whether we can build the right defenses before we're outnumbered by bots entirely. If we wait too long, the internet stops being a space for real people. And if that happens, we lose more than just trust, we lose the foundation of online life.

Lead AI Consultant at Northwest AI (NW AI). I help companies design and implement practical AI solutions from internal copilots to custom GPTs all tailored to their operations and data.

WYATT MAYHAM

CEO & Cofounder at
Northwest AI Consulting

Promise and Peril in the Age of Agentic AI

AUTHOR: MILANKUMAR RANA

"Agentic AI" is the newest thing in AI. It promises systems that can talk to people and things around them, plan, and learn. But the truth is more complicated than what all the ads say. We're creating a world where the line between human and artificial interaction quickly disappears. This significantly impacts business, society, and even how people connect.

THE AGENTIC AI MIRAGE

Agentic AI is the holy grail of AI: systems that can set goals, make choices, and do things independently without needing constant human supervision. We know how to do the first two things – make decisions on our own and act in a way that helps us reach our goals, but the more advanced features are still mostly ideas. Authentic iterative learning, advanced tool use, and meaningful interaction with the environment are still new technologies that aren't ready for production.

We should be very worried about this gap between what was promised and what is real. Companies rushing to market with "agentic" labels often overstate their capabilities and make systems that seem independent but only work within a small set of rules. The risk isn't just unhappy customers; it's also the chance of massive failures when these systems must deal with situations that aren't in their training data.

But putting something out too soon isn't the only worry. When real agentic AI comes, we'll have to deal with problems we've never had to deal with before regarding accountability, control, and unintended

consequences. Systems that can change how they act independently and interact with the real world could make the benefits and risks much bigger.

BOTS FOR THE BETTER?

Bots can be bad, but they also do many good things that aren't just for their benefit. Healthcare chatbots can help with mental health issues and medical triage 24 hours daily. They might even be able to save lives when no human experts are around. Bots that help with customer service can quickly answer simple questions, leaving human agents free to deal with more complex ones.
With big classes, human teachers can't change how they teach to fit each student's needs.

However, educational bots can give personalized lessons to a lot of students. Accessibility bots help disabled people get around online by turning text into speech or giving them other ways to interact.

The difference is in the goal and the openness. Helpful bots are automated systems that help people with real problems. Bots that cause problems make people think they are something they are not, or are only there to steal money through manipulation or fraud.

THE VANISHING HUMAN

The most worrisome thing is the bot pollution of digital spaces. Some sites have much higher percentages, but a conservative estimate says that bots make up 40–50% of all internet traffic. Most of the "followers" social media influencers have are fake accounts, which goes against social proof and authentic engagement.

We wonder what the future of business and people will be like now that more and more digital spaces are full of bots. If Zuckerberg's dream of social networks full of AI comes true, we might end up with echo chambers where people mostly talk to algorithms that are supposed to change how they act.

The effects on the economy are huge. To succeed, digital marketing needs to reach real people who can buy things and are interested in them. Whenbot traffic goes up, ads don't work as well, which could make whole business models based on getting people to engage and convert less stable.

THE PATH FORWARD

The answer isn't to stop making bots; it's to demand openness and responsibility. Platforms need to have strong ways to find bots and make it clear when interactions are automated. When people use AI systems, especially in business and social settings, rules should require them to be open. We need to teach people how to use technology so they can spot and understand AI interactions. We need to learn how to deal with a digital world that is becoming more automated, just like we learned how to spot spam emails.

We must keep places where people can connect in real life. As bots get better, real human interaction, with all its flaws and inefficiencies, may become our most valuable resource.
The future doesn't have to be a dystopia, but it needs careful planning. We can use the good things about agentic AI and helpful bots without losing our ability to choose and connect with others. The question is not if bots will change our world; they already have. Whether we will plan for that change or just let it happen.

MILANKUMAR RANA

Software Engineer Advisor at FedEx

Are Bots Helping or Harming the Internet?

AUTHOR: DEEPA CHAUHAN

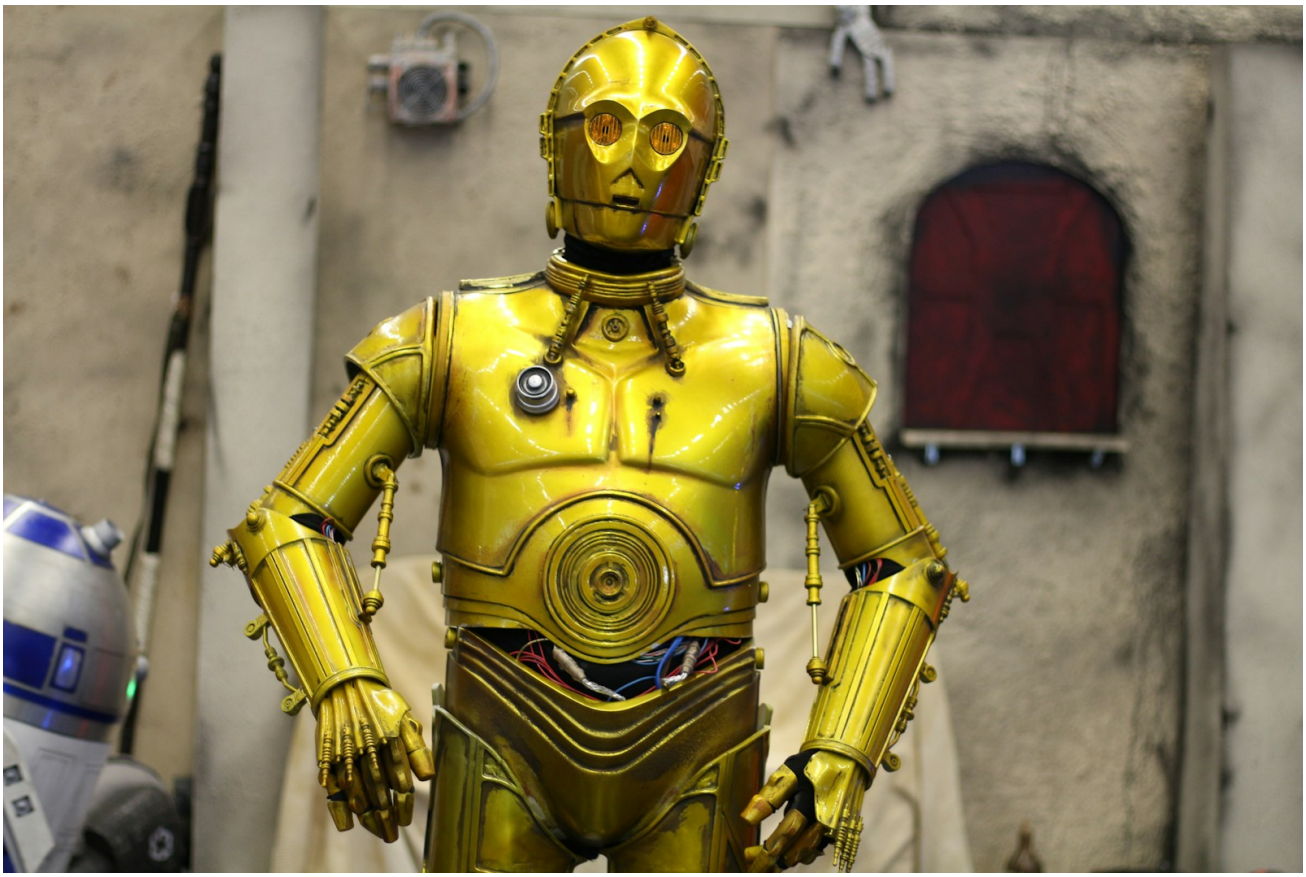

As of 2025, bots make up more than 65% of global web traffic. While many help businesses run smoothly and enhance user experience, an increasing number pose risks to the digital environment. They scrape private data, impersonate users, and undermine trust online. At the center of this transformation is Agentic AI, a promising yet often misunderstood branch of automation. As its adoption grows, so do the stakes for enterprises.

Between Hype and Reality

Agentic AI systems are made to act independently, making decisions, planning tasks, using tools, and learning with little human assistance. Ideally, they behave like teammates, not just assistants. But according to research from IBM and the UC Berkeley SCET Center, most current implementations have fallen short. Many so-called agentic systems are still built on human-defined workflows and lack the full capabilities associated with autonomy.

That said, some real progress includes:

- Microsoft Copilot and Google Gemini generating context-aware suggestions
- AI tools that autonomously update CRMs, calendars, and documentation
- Experimental frameworks like AutoGPT exploring self-learning mechanisms

Yet, three core components—iterative learning, dynamic planning, and environmental interaction—remain limited or non-existent. Overusing the term “agentic” without these elements risks misleading buyers and misaligning expectations.

WHAT'S THE DIFFERENCE?

Bots aren't always bad. They clearly add business value across a variety of functions when properly designed and implemented. Bots for customer service handle common inquiries and pass on complicated problems. HR bots cut down manual work by automating onboarding and screening. Finance bots take care of compliance reporting and reconciliation. Supply chain bots oversee workflows for updates and documentation. Healthcare bots assist with symptom triage and scheduling appointments.

These bots work best under supervision and flourish in controlled settings. Numerous businesses have reported significant improvements in turnaround time and process accuracy, especially in the banking, logistics, and healthcare sectors.
However, malicious bot activity is rising fast. According to Arkose Labs, attacks by harmful bots jumped 200% from 2023 to 2024. Imperva's 2025 Bot Report found that nearly 50% of all bot traffic poses a threat.

These malicious bots scrape content that is proprietary or protected by copyright and exaggerate advertising metrics and manipulate engagement information. They misuse credentials stolen to obtain unapproved access and spam websites harm user experience and analytics. The outcome? Damaged brand trust, compromised systems, and wasted marketing spend.

DISTORTING ONLINE DECISIONS

Another growing concern is synthetic engagement. Some high-profile social accounts may have up to 75% bot followers. Platforms like Meta are creating AI personas to mimic interaction. This increasingly unclear distinction between human and machine weakens the trustworthiness of data used for personalization, targeting, and reporting.
For decision-makers, this means more than inflated metrics. It undermines how we assess ROI, build relationships, and measure brand impact.

The five pillars of autonomous decision-making, goal orientation, tool usage, iterative learning, and environmental responsiveness are the foundation of true agentic AI. Only the first two are regularly present today. Nevertheless, a lot of solutions that are merely rule-based automations are marketed as agentic. In addition to confusing consumers, this mislabeling may result in bad procurement decisions, noncompliance, and unexpected system behavior.

The questions need to be asked: What is the system's true function? How does it decide? Is it possible to monitor and regulate its behavior? It is now essential for businesses to understand the distinction between automation and autonomy.

To manage this evolving landscape, tech leaders need to:

- Differentiate clearly between automation and agentic capabilities
- Demand transparency in AI behavior and outputs
- Establish governance frameworks with oversight for critical systems
- Ensure that all autonomous actions leave an auditable trail

Agentic AI is gaining traction, but leaders must stay grounded. As bots become a part of business and customer service, leaders need to understand what these tools actually do, not just what they say they do. The most successful companies won't be those chasing after the latest technology. Instead, they will be the ones asking the right questions, setting clear guidelines, and making honest choices. In a digital world filled with automation, being clear and in control can make a significant difference.

DEEPA CHAUHAN

Sr. SEO Specialist, Accelirate

The Risk Isn't Agentic AI Autonomy—It's Misunderstanding

AUTHOR: ALEX LANSTEIN

The current conversation around agentic AI has been focused on autonomy and decision-making. But the potential impact of agentic AI is bigger than that. Since agents have the ability to plan their own actions, use enterprise tools, interpret environmental cues, and refine their behavior based on feedback, they pose exponential risks that users need to consider. From runaway AWS bills to permission breaches to ambiguous decisions with no audit trail, agentic AI can turn from tool to terrifying very quickly.

POWER WITHOUT BOUNDARIES

One possible risk of using agentic AI can occur when granting agents access to enterprise tools—like log aggregators, ticketing systems, or cloud platforms—to investigate a potential threat. A bad search may not only lock up a logging device, it can also lead to astronomical AWS bills.

Agentic AI doesn't have these concerns about cost or performance. Agents may not know that one optimized search yields better results than ten inefficient queries – or that ten inefficient queries should be run as one giant inefficient query. They also don't understand office politics or professional ethics. Agentic AI may want to access data simply because

it can, but if the CEO's email logs are accessed, or suspicious Excel macros from the CFO are fetched, the implications are much bigger than threat assessment.

The promise of agentic AI is that it can learn from the way humans work and refine its approach over time. But agents would need to interpret what it's observing correctly—and the accuracy with agentic AI is far from 100%.

Also, an agent may observe human behaviors that are incorrect, which leads to it reinforcing the person's faulty logic. Unless the mistakes are caught, these learned behaviors become incorrect assumptions that don't get challenged. Logic cannot be patched like software, so when agentic AI learns incorrectly, the problems it's meant to solve can snowball.

AWARENESS WITHOUT UNDERSTANDING

Agentic AI can adapt its actions based on environmental cues, but noticing patterns is not the same as understanding them. For example, an agent might see someone logging into the system from China and assume it's suspicious. The activity may be flagged because the agent has no way of knowing that the employee is simply on a business trip.

You know all anomalies aren't problems. Agentic AI doesn't. Without having a grasp on this nuance, everyday activities are perceived as noise—or worse, false positives that can trigger real disruptions when the AI "mitigates" a FP.

THE ACCOUNTABILITY MIRAGE

Agentic AI is especially dangerous because of the lack of accountability it may cause. When a driverless car makes a fatal error, we ask who is to blame. When an AI agent pulls sensitive executive data or blocks legitimate services, the same question will be asked. The users implementing the technology need to understand that they will be on the hot seat when and if the AI runs amok.

Telling the CEO that "The AI did it," when explaining why their sensitive information was accessed is likely to get a reaction worse than "The dog ate my homework." You don't get to abdicate responsibility for the problems agentic AI creates. This means organizations need clear guardrails around what agents can and cannot do, and standards for how its actions are monitored. And most importantly, someone needs to be held accountable when something goes wrong.

CLOSING THE GAP

The biggest threat from agentic AI isn't that it can make its own decisions—it's that those decisions can be made without context or understanding. Agents search the wrong systems, escalate the wrong alerts, and learn the wrong lessons. This isn't because agentic AI is some kind of evil sci-fi villain trying to take over the world—it's because it doesn't understand the big picture of why it's doing what it decides to do. And when things break, there won't be anyone to blame but the humans who set it loose.

Agentic AI has real potential. But without visibility, cost awareness, and governance, that potential will be a liability. We can't hand over the reins to agentic AI until we build fences to contain it.

ALEX LANSTEIN

CTO and Head of Threat Intelligence, StrikeReady

The Agentic AI Reality Check

AUTHOR: VINOD GOJE

Business executives should be very concerned about the large gap between marketing claims and technical reality, even though the underlying capabilities do show promise.

The current market conception of agents as LLMs with function calling represents a fundamental mischaracterization of true autonomous systems. 25% of companies that use GenAI will launch agentic AI pilots or proofs of concept in 2025, yet over 40% of agentic AI projects will be canceled by the end of 2027 due to escalating costs, unclear business value, or inadequate risk controls.

The five core capabilities that define agentic AI expose significant discrepancies between the actual state and the claims made. Autonomous decision-making and goal-oriented behavior are largely functional within narrow, controlled environments. However, the critical capabilities related to tool usage, dynamic planning, iterative learning from failure, and meaningful environmental interaction range from primitive to nonexistent in production systems.

Many vendors are contributing to the hype by engaging in “agent washing,” the rebranding of existing products, such as AI assistants, robotic process automation (RPA), and chatbots, without substantial agentic capabilities. Companies claiming to have “launched” agentic AI are typically deploying sophisticated automation scripts with limited reasoning capabilities, not the autonomous agents their marketing suggests.

The concern isn’t theoretical. In multiagent systems, “hallucinations” can spread from one agent to another; they can persuade other agents to take the wrong steps and give incorrect answers. When systems lack robust environmental interaction and learning mechanisms, they become expensive, inflexible automation that fails gracefully at best and catastrophically at worst.

BEYOND THE BINARY OF GOOD AND BAD

The bot ecosystem demonstrates both legitimate utility and systematic exploitation. Automated bot traffic surpassed human-generated traffic for the first time in a decade, constituting 51% of all web traffic in 2024, with bad bots making up 37% of all internet traffic, a significant increase from 32% in 2023.

Good bots provide measurable business value: search engine crawlers ensure discoverability, monitoring bots maintain system health, and customer service bots handle routine inquiries efficiently. However, these represent straightforward automation rather than intelligent agents.

The exploitation side reveals sophisticated operations. Bad bots now account for 73% of all internet traffic in some reports, with scraping activity seeing a staggering 432% rise between Q1 and Q2 2023. Beyond simple scraping, bots now enable account takeover (ATO) attacks, which increased 10% in 2023 compared to the same period in the prior year, with financial services seeing the highest concentration of attacks.

A BOT-DOMINATED INTERNET

The statistics around bot prevalence require urgent strategic consideration. In 2023, bots made up 49.60% of internet activity, almost catching up to human traffic, which was at 50.40%, representing a fundamental shift in internet composition that most business models haven't adequately addressed.

For businesses, such growth creates a paradox: as bot traffic increases, traditional metrics like page views, engagement rates, and conversion funnels become increasingly meaningless. The sectors most affected are financial services, healthcare, and e-commerce, as these industries rely on APIs for critical operations and sensitive transactions, which makes them prime targets for sophisticated attacks that can skew data and compromise customer trust.

The human element becomes more valuable, not less, in this landscape. While platforms promise AI agents will handle customer interactions, the authentication and verification of genuine human intent become a competitive advantage. Companies must invest substantially in bot detection and management not just to protect against malicious actors, but to ensure their business intelligence reflects actual human behavior.

Bottom Line

We're witnessing sophisticated automation being rebranded as autonomous AI, while the internet becomes increasingly populated by non-human traffic. Business leaders should approach agentic AI claims with healthy skepticism, focusing on specific, measurable use cases rather than transformational promises. Simultaneously, the bot traffic surge demands immediate attention to data integrity and customer authentication systems.

The technology isn't advancing as rapidly as the marketing suggests, but the business implications of our increasingly automated digital environment are arriving faster than most organizations are prepared to handle.

VINOD GOJE

Vice President of Engineering at Bank of America

How will we secure bots that buy?

AUTHOR: JAMES SHERLOW

In the not too distant future we can expect bots to carry out actions such as shopping, research, and even transactions on our behalf. Consumers will be able to ask bots to search for, recommend or select and purchase goods and services, taking the tedium out of the processes such as finding a coveted item or purchasing concert or flight tickets. But providing that level of trust to a bot will require us to readjust our mindset when it comes to bots, not mention some radical adjustments to the mechanisms that make it happen.

Many of the processes used to carry out purchases are specifically geared towards humans. Automated bots, it's assumed, are bad and the past few decades have seen a slew of CAPTCHA inventions that aim to check if it's a human initiating the interaction. As only humans are allowed to complete account sign-ups, logins and purchases, this creates a huge stumbling block for agentic AI. Without an identity, and without credentials, the AI cannot be validated and gets blocked.

What is needed is a new way for agents to gain programmable access to websites, APIs and applications that no longer requires manual account creation or the signing of a contract. AI agents need to be able to present verified credentials and payment methods that will then permit them to gain access to the same digital resources as their human counterparts.

WHAT ABOUT BAD BOTS?

It's an approach that does of course pose a real dilemma because of its potential to increase fraud. Blocking all bots has been the safest way to prevent fraudulent purchasing in the past. But AI purchasing is inevitable, which is why credit card heavyweights VISA and Mastercard have both launched AI agent initiatives, making blanket bot bans no longer viable. Instead, it will be necessary to differentiate between good and bad bots.

Doing so will be challenging given that industry reports estimate over a third of bots can be classed as bad. That means in addition to verifying the identity of the bot, security measures must be put in place to monitor its activity to ensure it doesn't turn rogue. Identity and security protocols therefore need to be applied to recognise and authorise verified AI agents while continuing to monitor for scraping, fraud, and abuse.

By combining bot mitigation and management with financial transaction technologies, it's now possible to provide this level of assurance. AI agents can enrol, get verified, and receive digital tokens that represent their identity and purpose, allowing them to carry out all of the actions previously denied to them. If they are assigned the token themselves, this is bound to them creating an identify that can be authorised and subjected to governance. The bot can then be assigned a programmable wallet with funding sources e.g. credit cards, ACH or wire, awarded identity credentials and be made to adhere to payment rules by a third-party platform provider.

Transactions are then monitored using bot management and API security tooling. This is able to recognise the agent's identifiers and using machine learning can evaluate behavioural, contextual and intent-based signals based on the bot's actions. This monitoring can distinguish between legitimate versus malicious automated activity so that any instances of scraping, fraud or API abuse can be immediately blocked and investigated.

A WHOLE NEW WORLD

Using both bot and payment technologies together ensures that the business can trust that the agent interacting with their service is both legitimate and monetizable. And that last part is key to the future of the emerging AI economy, because allowing bots to access, pay for, and interact with valuable digital content that sits behind login walls, anti-bot protections and compliance layers opens up a whole new world of possibilities.

Once automated bot purchasing takes off, businesses will be able to turn that traffic into trusted measurable revenue streams, leading to a marketplace for AI commerce. For example, ecommerce retailers can monetise real-time data access that was previously used for free by comparison-shopping bots. Financial services can enable verified agents to access and use their curated datasets to provide market insights. Content creator sites and media platforms can charge for the consumption of high value content or for use of their APIs. And travel aggregator sites can offer verified availability data to AI agents on accommodation and flights.

In effect, the business can decide which AI agents are allowed in, how they are permitted to operate, and to charge for that access or data, allowing them to manage and monetise bot traffic for the very first time. They can safely deal with agents without the risk of data theft or API abuse and get paid for the value they deliver rather than giving it away for free. It's an exciting prospect because not only will bot-driven purchasing make the process more convenient and streamlined for us but it will also create a new economy.

UNCHARTED WATERS

Forward-thinking companies will therefore no longer be thinking 'how can I stop bots' but 'how can I capitalise on legitimate AI-driven interactions'. They'll need to look to the market to find providers capable of both assigning and verifying digital identities to AI agents and also of monitoring how these agents are interacting with APIs and web applications. That way, if the bot does turn out to be malicious, the abuse can be detected and automatically blocked.

Combining technologies that both facilitate and secure these exchanges is vital because, at the time of writing, assigning network tokens to AI agents means they are currently out of scope of the Payment Card Industry Data Security Standard (PCI DSS). It's down to the business/merchant to validate token handling, authenticate the agent and deal with storage so the onus is on them. This may well be an area the PCI SSC seeks to address in the future but for now merchants will have to navigate these waters unguided so will need to look for assistance when it comes to putting in place security controls to govern AI agent payments.

JAMES SHERLOW

Systems Engineering Director, EMEA, Cequence Security

Q-Day isn't as dangerous as our government

AUTHOR: LOU COVEY

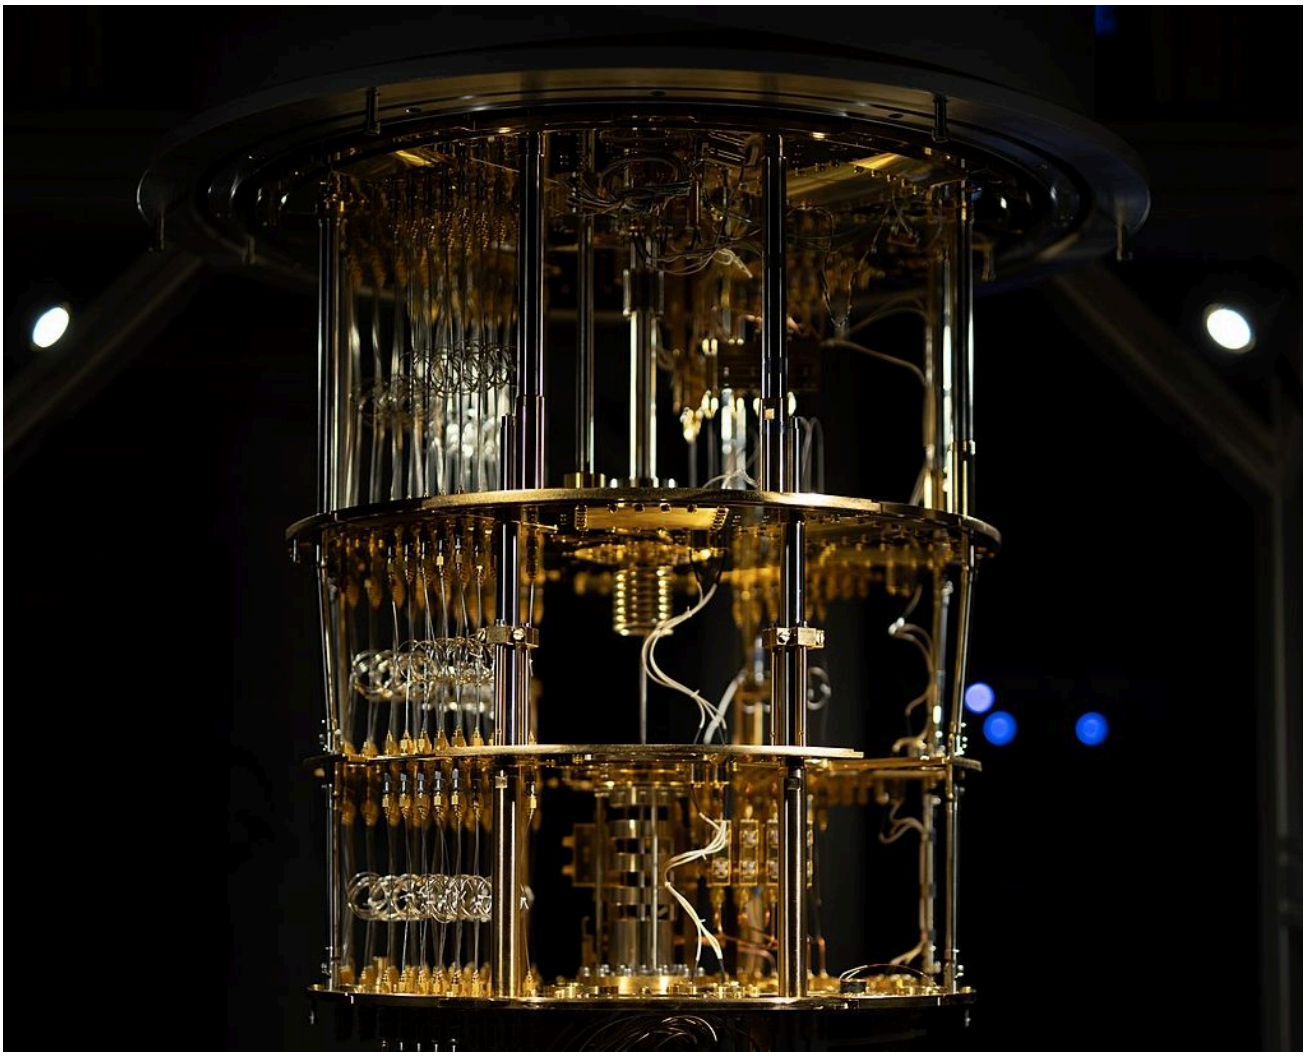

The post quantum computing (PQC) industry wants us to believe that Q-day, the day that a quantum computer is right around the corner. It isn't. But that doesn't mean what the niche members are working on is worthless. Perhaps the most important task they have is limiting government surveillance of the innocent.

If you don't already know, Q Day is the day when a quantum computer exists powerful enough to break current military-grade encryption standards. This has been a major disaster predicted by many, not the least being Wired Magazine. Most in the industry claim it will happen in the next decade, if it hasn't already happened.

FEAR MONGERING

Part of the current fearmongering comes from a recent IBM announcement that they will have a 200-logical-qbit quantum computer by 2029 and a

2000-logical-qbit by 2035. No one knows if a hostile foreign state is further along, they say. China is often the most mentioned usual suspect.

What a lot of engineers and software designers fail to take into account in making these predictions: land and resources for the facility.

IBM's most powerful quantum computer is 127 qubits and requires "10 kW of power just for control and readout," according to an article in PatentPC. That include microwave signal generators and error correction processors. However, the 200qbit system requires a total of 370kW, more than a quarter of that power dedicated to cryogenics.

POWER HOG

Therein lies the problem. Quantum computers need to be kept at almost absolute zero to keep from melting down. That requires power on the level of a mid-sized nuclear reactor or several hundred acres of solar panels or wind generator. To make a quantum computer capable of decrypting a single document in the span of a week requires a quantum computer five times the size of the projected IBM system that we won't have for another 10 years. It will require a power plant capable of providing power ranging from 500 MW to 1 GW. No one will be able to keep a complex that large a secret.

The question, however, is whether we actually need a quantum computer to break modern encryption.

"I believe, based on the insiders at NIST, the NSA, and MITRE, that Q-Day is already passed," said Crick Waters, CEO of the "quantum-safe communications" company Patero.

"Can NSA break Crystal's Khyber encryption? Maybe." He continued. "Q-Day is a mushy thing. When somebody figures out how to crack the encryption, they won't announce it. There are enough researchers and large organizations pushing to break encryption."

CLOSE BUT NO CIGAR

Crick pointed out that Google estimates that the number of qubits to break encryption has dropped from 20 million to a million. "When does that become 100? IBM has already announced they're going to have a 200-logical qubit computer available for use for anybody in two and a half years, in 2025."

But Crick agrees that we may not need a quantum computer to break military grade encryption now.

In 2015, the FBI recovered an iPhone from a terrorist who killed 14 people in San Bernardino, California. They demanded Apple provide a backdoor into the phone to gather evidence and see who else was involved. Apple, which still uses 256-bit standard encryption, refused. One year later, the FBI announced they had broken Apple encryption and accessed the phone data. They found nothing of value to their investigation but did force Apple to up its security game on the products. The downside is breaking encryption on secure devices or systems is not impossible with current resources.

CRYPTO-AGILITY

The answer lies in cryptographic agility, or "crypto-agility" an industry term that has been around for a while. NIST published a white paper on the subject this year.

"Crypto-agility means that we can change the cryptography quickly without bringing our networks down and ripping out hardware and rebuilding them," Crick explained.

Most companies in the PQC area have the ability to do just that, so it isn't a differentiator, so that is possibly why companies don't promote it. The fear of Q-Day is still the most popular form of marketing PQC technology, even though it is highly unlikely of happening in our lifetimes. That is unfortunate, as crypto-agility resolves the current problem quite nicely.

It's possible that PQC companies just enjoy getting the attention of major news organizations where the reporters haven't done the math on what quantum computing would cost and are missing the real-time value PQC companies provide.

LOU COVEY

Lou is Chief Editor of Cyber Protection Magazine

Post-Quantum Readiness: A Strategic Imperative for Cybersecurity

AUTHOR: CHRISTINA CRAVENS

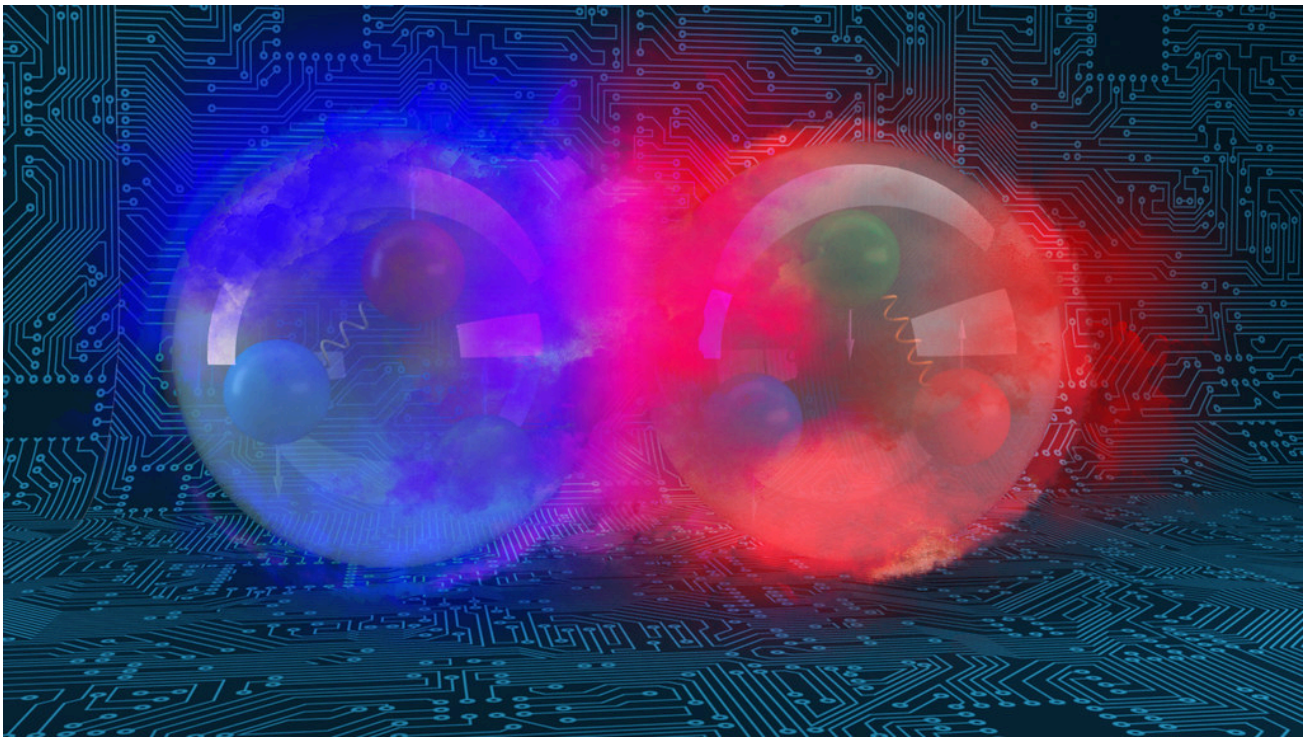

Quantum computing promises remarkable breakthroughs across science, medicine, and technology, but it also brings a looming threat to cybersecurity. The U.S. National Institute of Standards and Technology (NIST) has issued clear warnings: the encryption algorithms that protect today's data and communications could be rendered obsolete by the power of quantum computers. Once capable of breaking classical cryptography, quantum systems will put the confidentiality, integrity, and authenticity of digital assets at risk.

This is not a distant problem. Government agencies have begun to sound the alarm, encouraging public and private sector organizations to prepare for a post-quantum future now. The U.S. Cybersecurity and Infrastructure Security Agency (CISA), the National Security Agency (NSA), and NIST are urging companies to evaluate their current cryptographic posture and chart a path toward quantum resilience. But getting there requires more than just adopting a new set of algorithms.

Going Beyond Algorithm Swaps

"Post-quantum readiness" is not just a technical upgrade—it's a strategic, organization-wide effort. While new quantum-resistant encryption standards are being finalized, companies must begin preparing by developing a comprehensive understanding of

how and where encryption is used within their environment.

That includes building a full inventory of cryptographic protocols in use, mapping out encrypted data flows, and identifying dependencies across business systems. Many organizations are surprised to find that encryption is deeply embedded in everything from internal applications and cloud services to third-party integrations and legacy systems.

Understanding these dependencies is critical. Business functions that rely on encrypted communications – whether to secure customer data, authenticate users, or protect proprietary information – must be identified and assessed for quantum vulnerability. This level of visibility is the foundation of any meaningful post-quantum transition strategy.

A Looming Deadline for Visibility
The timeline for preparation is shrinking. Over the next year, businesses should create a detailed inventory of all cryptographic methods and encrypted data pathways that support critical business functions. This means knowing what algorithms are in use, where they're implemented, and how they tie into essential services and workflows.

Without this inventory, organizations will be unable to effectively prioritize their migration to post-quantum standards. Worse, they may be left exposed in areas that were overlooked, creating blind spots in otherwise mature security programs.

Building the Inventory
Creating this inventory doesn't have to be overwhelming. Modern approaches to cyber asset discovery can help accelerate the process. For example, network-based sensors that passively observe data flows can identify where encryption is in use and which systems are communicating securely. By layering in analytics and business context, these tools can help organizations correlate encrypted pathways to specific business functions and compliance requirements.

This kind of continuous, dynamic mapping – especially when supported by automation and AI – offers a scalable way to maintain up-to-date visibility as systems evolve and new technologies are adopted.

Preparing for the Transition
Once the cryptographic landscape is mapped, the real work begins. Organizations must evaluate which cryptographic implementations are quantum-vulnerable and start testing quantum-resistant alternatives. Not every system will need to be upgraded immediately, but those handling sensitive or long-lived data should be prioritized.

Ultimately, transitioning to post-quantum cryptography is not a one-time project. It will unfold over multiple years, requiring coordination across security, IT, compliance, and business teams. Organizations that start now, by developing cryptographic inventories and aligning their security programs with government guidance, will be in a much stronger position to protect their assets as the quantum era arrives.

CHRISTINA CRAVENS

Chief Growth Officer at Redjack

Post-Quantum Cryptography Event Horizon Approaches

AUTHOR: DAVID CLOSE

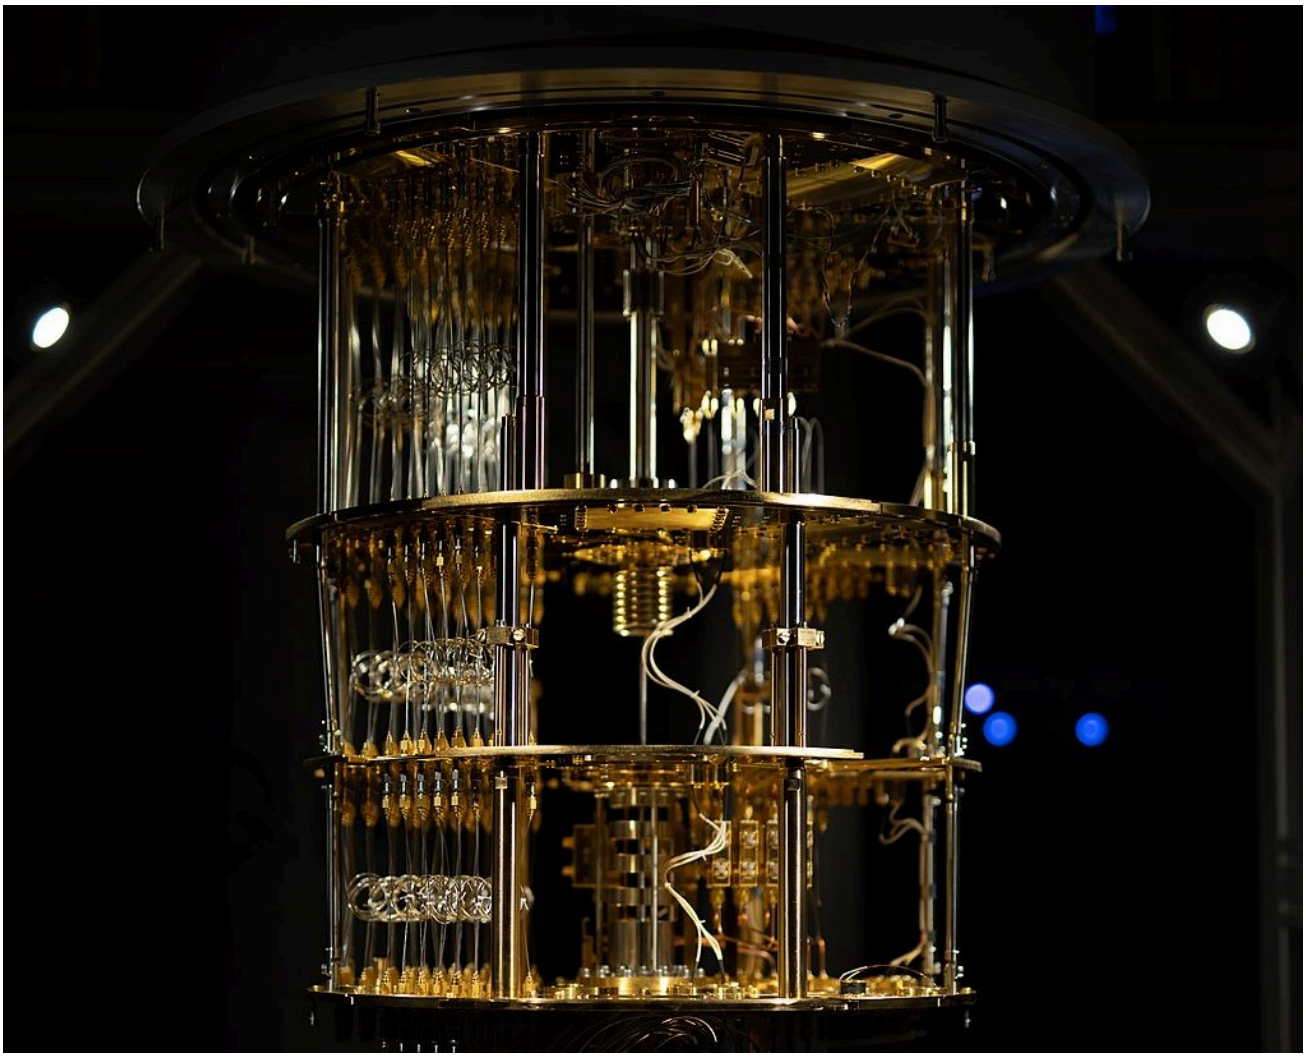

The rise of quantum computing poses a serious threat to current encryption methods. For decades, industries like banking, healthcare, and government have relied on encryption to protect sensitive data, whether in transit or at rest. However, as quantum computing advances, this lock is at risk of being picked. The question is not if but when quantum computers will render current encryption obsolete.

Quantum computers function as master lockpickers, capable of solving mathematical problems that traditional computers would take centuries to crack. This capability threatens the core of modern encryption, including widely used public key systems like RSA and ECC. When large-scale quantum computers become operational, bad actors can easily unravel public and private encryption systems. Everything from confidential patient records to

classified government communications could become vulnerable. While cybersecurity teams are already working overtime to patch vulnerabilities and prevent data breaches, quantum computing adds a complex new dimension to their challenges.

The Role of Post-Quantum Cryptography (PQC)

In response to these emerging threats, the National Institute of Standards and Technology (NIST) has finalized three post-quantum cryptography (PQC) standards in August 2024: 203-Kyber for lattice-based key encapsulation, 204-Dilithium for lattice-based digital signatures, and 205-SPHINCS+ for stateless hash-based digital signatures. These standards offer organizations a roadmap to future-proof encryption. Additionally, national initiatives, such as CISA's Post-Quantum Cryptography Initiative, spearhead efforts to ensure industries adopt quantum-safe solutions.

Among the quantum computing experts surveyed by the Global Risk Institute, 22.7% anticipate a quantum computer-based attack on RSA-2048 by 2030, while 50% consider it likely by 2035. The PQC market is projected to exceed $17 billion by 2034, underscoring the growing need for quantum-safe solutions. PQC algorithms, unlike traditional methods like RSA and ECC, leverage mathematical structures that quantum computers find significantly more challenging to solve, such as lattice-based and hash-based problems.

Key Action Steps

Conduct a comprehensive inventory of your cryptographic systems, protocols, and assets. Identify outdated algorithms, such as RSA and ECC, that are particularly vulnerable to quantum decryption. Prioritize systems handling sensitive, long-lived data.

Counter "Harvest Now, Decrypt Later" Threats

A significant quantum-era risk is the "Harvest Now, Decrypt Later" (HNDL) strategy, in which adversaries intercept and store encrypted data today to decrypt later using quantum capabilities. Secure data in transit with hybrid protocols combining classical and post-quantum encryption algorithms. For example, hybrid TLS implementations combining traditional encryption with CRYSTALS-Kyber offer immediate protection while paving the way for future transitions.

Code signing prevents quantum computing attacks by ensuring that only verified code runs on a system. By creating a quantum-safe signature, code signing prevents tampering. Why it matters: Code signing ensures software integrity and authenticity, guarding against tampering and malware injection. Without robust code signing, attackers can exploit critical software and firmware.

Legacy systems, including Industrial Control Systems (ICS) and other critical infrastructure, pose an additional challenge. Many rely on outdated cryptographic protocols and are not designed for seamless updates. Conduct a thorough assessment to determine their susceptibility to quantum attacks. Plan hardware replacement cycles that integrate quantum-safe cryptographic capabilities.

Building Long-Term Resilience

Organizations must adopt crypto-agile architectures to remain secure in an era of evolving cryptographic standards. Crypto-agility enables swift transitions between algorithms through centralized key management systems and software-defined cryptographic solutions. With these tools, organizations can deploy updates seamlessly, reducing downtime and maintaining operational security.

Hybrid cryptographic solutions combine classical algorithms with quantum-resistant alternatives. This transitional strategy ensures compatibility with existing infrastructure while preparing for quantum threats. For instance, applying a hybrid Certificate Authority (CA) solution combines conventional cryptographic signatures with PQC signatures, ensuring systems are compatible with current technology and secured against future quantum computer-based attacks.

Follow Standardization and Compliance

Keeping pace with emerging PQC standards is essential for long-term security. Aligning internal cryptographic practices with NIST standards ensures compliance with future regulatory requirements while bolstering organizational security. Adhering to industry best practices also fosters trust with stakeholders and partners in the quantum era.

Organizations must develop a PQC strategy before their cryptographic infrastructure is compromised. By taking these proactive steps now, companies can prepare for the quantum future while maintaining security today.

DAVID CLOSE

Chief solutions architect at Futurex


Every 11 seconds a hacker
falls in love with your data*
CYBER PROTECTION
MAGAZINE
Protect your data
https://cyberprotection-magazine.com
*According to cybersecurity ventures

SIDEBAR INFOS

Type 6: Front Page

REPLACE SIDEBARS

Type 6: Front Page