FUTURIOM
FUTURE OF CLOUD TECH

Cybersecurity Leadership Brief

Sponsored by:

Teleport

Identity Is the Foundation of Agentic AI Security

Highlights:

- **Identity is ready for its spotlight.** AI agents' speed and scale overwhelm current approaches. They represent an essentially infinite influx of new "users" that force a more thorough approach toward nonhuman identity.
- **Agents magnify security risks on multiple fronts.** Their autonomy, their unpredictability, and their scale all present new challenges that security platforms must adapt to.
- **Principles such as zero-trust remain intact but require new rigor.** What's necessary is a new approach, not necessarily a newly invented security technology.
- **Unification is crucial.** Identity has evolved piecemeal in most organizations, a patchwork of disparate logins and authorization schemes. A necessary first step is to unify the enforcement of identity across the organization.

1. Intro: AI Generates a New Threat Landscape

While AI helps users generate text, images, and video, it has also generated a raft of justifiable security concerns. For example, it's well recognized that large language models (LLMs) are black boxes; it can be difficult or impossible to discern how the model arrived at its choices. The model can also drift subtly until, like a ship lost at sea, it winds up severely off target. Worse, the model can be manipulated by bad actors to deliver untruths or sneak, unauthorized, into sensitive data resources.

Agents multiply these risks many times over. Model builders and enterprises pursuing agentic AI envision automated bots generating their own queries and tapping multiple models (large or small). Think of a factory that can build its own assembly lines for different products. You wouldn't want it running unsupervised for too long.

Agents widen the blast radius. They won't trace the same paths through the network every time, and they won't follow predictable courses of action. Root cause analysis becomes a tangle; the scale, speed, and seemingly random actions of agents could easily overwhelm existing methodologies.

Further, both agents and models need access to vast amounts of data if they're to do sophisticated work. Advancements such as the Model Context Protocol (MCP) imply that some of these visitors will come from outside the organization, or from distant departments within the same company. Either way, it's important to know who's accessing that data.

Clearly, this is a multifaceted problem. But a rethinking of the threat landscape can go a long way toward curbing these issues and reining in agents.

As a starting point, we must think of agents like human users. Treat them with the same zero-trust, least-privilege principles that we apply to people. Limit their access only to what's necessary, and diligently track what they've done, in case signs of malfeasance pop up later. It's all easy to say but challenging to implement, due to the scale and speed of agentic AI.

In this Leadership Brief, we'll explore this deeper perspective on identity and access. We believe identity is due for a new phase, one that's suitable for tracking nonhuman users and can act quickly enough to monitor their unpredictable interactions with enterprise resources.

2. Enterprises Have Big Plans for Agents

More powerful than fixed-function bots or scripts, agents are AI-driven tools that scan their surroundings, accessing information from databases, documents, even Internet sources. They then use AI models as their "brainpower" to decide how to complete a task. An agent is goal-oriented, and that goal could range from a one-time task to long-term monitoring.

Some examples of agentic work have become familiar by now:

- In healthcare, monitoring a patient's vital signs continuously armed with knowledge of the patient's health history and current medications. One goal here would be to watch for early warning signs of sepsis or other complications.
- In retail, managing inventory across a fleet of stores. An agent can use computer vision to examine store shelves and coordinate necessary shipments from warehouses.
- Customer service bots. We've all experienced this one. They interpret natural-language questions and requests, seek answers that go beyond a static script, and escalate as necessary.

Security Gets Squeezed, As Always

This is just the beginning, as the aspirations for agents are sky-high: agents performing entire processes end-to-end, agents acting as always-on assistants, and so forth. Organizations like PwC envision agents fully transforming the nature of work. Tech companies are striving to make that come true—Windows PCs shipping with Copilot installed, for example. And many enterprises have latched onto this thinking eagerly.

Meanwhile, security teams are doing what they always do: trying to keep all this activity safe while also letting developers maintain the pace that's demanded of them. It's a push-pull dynamic that has always existed, but the urgency of AI has intensified the struggle on the security side. The nature of agents only magnifies the potential security risks, as we'll examine below.

Don't Discount the Power of AI Enthusiasm

To be clear, enterprise AI enthusiasm isn't entirely driven by overzealous executives. Some of it comes from genuine developer enthusiasm. Futuriom attended the inaugural MCP Developers Summit in 2025, which focused on the Model Context Protocol (including security implications) but also became a symposium on developers' dreams for agentic AI. The energy was palpable. Developers are excited to make these new toys do complex tricks.

This means, however, that a wave of shadow IT has likely emerged, analogous to the way developers used their public clouds of choice early on. Quite a few ad hoc agent experiments have already reached production, possibly without receiving the security scrutiny they deserve.

3. Identity Is Already an Enterprise Issue

Even before considering AI, we can see plenty of potential weak spots in the average enterprise's identity framework. Certainly, identity and access management (IAM) and privileged access management (PAM) tools are available to help, but depending on how they're applied, they can still leave gaps:

- **Overprivileged accounts and long-lived secrets.** Some organizations, particularly older ones, carry volumes of accumulated technical debt. One likely consequence: a forgotten "superuser" identity, with unfettered access to everything within an organization, gets revived by the wrong hands. Technical debt creates prime back-door targets for intruders.
- **Human-scale thinking.** The days of identity being assigned by a human sitting at a workstation are long gone. Even newer tools are arguably descended from that mode of thinking, even though we know the public cloud has pushed scale and speed beyond the reach of human hands. Agentic AI expands this issue exponentially.
- **Fragmented platforms.** When identity and access security are applied diligently but through multiple platforms, it creates an invisible vulnerability. Disparate logins and multiple points of entry lead to exploitable vulnerabilities, such as those forgotten overprivileged accounts. The situation is common; many existing enterprise frameworks are really a combination of multiple platforms. Some were implemented for certain groups or certain projects; some arrived via an acquisition.

Adding agents to the mix certainly complicates the picture. The volume of agents already tapping web sources—think of Wikipedia's pleas about AI bot scrapers overwhelming its servers—shows how quickly agents can surpass our usual ideas about scale. These machine users, which we must consider equivalent to human users, will appear in huge numbers at essentially random times.

Moreover, agents themselves provide juicy hacking targets. They're highly autonomous and are often built expressly for the purpose of accessing data. Even if an agent is legitimate, it can unintentionally stumble onto loopholes in security policy. In a zero-trust world, agents truly deserve zero trust.

Agentic Advancements Can Mean Even More Trouble

Now think about MCP and similar protocols, which provide a common way for AI agents to connect to data sources and to one another. Potentially, this creates more junctions among systems and agents, paths that a security team might not have considered, or—considering agents that are created sporadically—might never have known were possible. We must consider that agents will always find new ways to wander amidst an organization's data sources.

On a more sophisticated level, agents can also create their own data structures. This doesn't mean full-blown databases built from scratch (yet). Still, in our conversations, companies in the data and storage world have brought up the idea of having agents create their own short-lived vector indexes or temporary tables. Think of them as mini reference cards or quick-lookup resources. They're applied only to the job of the moment and are intended to disappear once that job is done.

These subsets of existing data resources should be harmless—but, as we know from decades of cybersecurity, it's always worthwhile to be paranoid. Consider, for example, the danger of "policy drift." If the agent has tapped both sensitive and not-so-sensitive information, can we ensure that an ephemeral index or table still carries the proper level of security? And, again, an agent doing this kind of data manipulation would be an awfully tempting target for hackers.

The point is: Agents have the potential to run rampant, either by accident or by design. What's required is an *automated* and *scalable* approach to identity and privilege.

More importantly, the foundation underlying that approach cannot be based on multiple, independent identity systems. "Automated" and "scalable" are necessary, but they must also be backed by *unity*.

4. Identity Requires a Unified Front

First, note that many potential pitfalls of an identity and privilege approach come from simply having too many systems—too many ways for things to fall through the cracks. This alone suggests a unified approach is warranted.

Now remember our recurring theme: Agents magnify problems. Even if we assume no bad actors, the sheer volume of activity means agents will stumble onto vulnerabilities. To paraphrase Amazon, an event with million-to-one odds isn't so unusual if you're doing 10 million things a day.

Think back to the push-pull dynamic, where executives press for fast AI adoption while security teams struggle to keep their systems safe. If agents are crawling unpredictably through many corners of an organization's data sources, then that security team benefits from having one source of information. That means one system to oversee activity, spot anomalies and alerts, and trace problems back to their causes. It's counterproductive to ask this team to lock down privileges and enforce policy while also maintaining multiple platforms.

Here's what a unified approach requires:

- **A unified identity layer.** This doesn't mean starting from scratch but instead unifying all identities for all platforms and securing them cryptographically. This identity layer then allows fleets of agents to interact with zero-trust principles and access control guardrails.
- **A unified approach to users.** People and machines must be treated under the same identity rules.
- **"Life cycle" thinking.** It's not just about assigning credentials. A rigorous approach to identity includes ongoing verification, *auditability*, and the continued vigilance/challenge-presenting of a truly zero-trust environment.
- **Real-time reactivity.** With agents arriving and disappearing at blinding speeds, the platform must be able to recognize inappropriate actions immediately. When possible, this includes automated response and remediation. If human intervention is required, the platform must provide immediate relevant information—which must be high-level enough for human comprehension but granular enough for root-of-cause analysis.
- **A standards-driven approach.** Technically, this isn't required, but given the amount of change being wrought by AI, now is not the time to deviate from foundations such as the NIST Cybersecurity Framework or OAuth.

Modern Identity Includes New Thinking Around Old Ideas

Expanding on that last point, we certainly don't want to throw out proven good practices. Here are a few key examples:

- **Zero trust** remains a treasured principle. All users must be verified and their permissions kept to a minimum.
- The same goes for **zero standing privileges (ZSP)**: Nobody's access privileges should be permanent. When a task is complete, the associated privileges should evaporate.
- **Hardware roots of trust** for all actors, human or machine. This includes biometrics, trusted platform modules (TPM), and hardware security modules (HSM). People and agents need identities within this framework, but so too do endpoints: laptops, mobile devices, and other client devices.

5. Conclusion

Many of our arguments have been in the cybersecurity zeitgeist for a long time, but they're applied with varying degrees of rigor. Realistically, very few enterprises have the resources and time to lock down cybersecurity as well as they'd like. For many companies, though, the wounds are self-inflicted. Many enterprises, given the tradeoff between tighter security versus faster product development, choose the latter.

Agentic AI is a forcing function, as the organizations that don't reconsider identity and access leave themselves particularly vulnerable to breaches and snafus. Agents have the potential to cause trouble even in the absence of any malicious actors.

Certainly, technology needs to advance to keep up with the speed and scale of agents. But new technology won't do the job unless it's applied with a new philosophy around identity. Nonhuman users must be treated as if they were human, and all users must be treated with skepticism. Privileges need to be limited to a minimum, both in the level of permissions and the amount of time they're allowed to persist.

Unification, though, must be the philosophy underlying all these principles. This is arguably the most radical of our suggestions, in that it contrasts with current practices. We're not saying that enterprises enjoy having fragmented identity frameworks. It's just that in a pragmatic sense, many organizations evolved that way, and the risk/reward calculation hasn't prodded them to improve the situation.

With agentic AI having arrived, that calculus has changed. Cybersecurity must now be anchored by a robust identity platform, one that applies to all users across all systems. Scale and speed are important, but gaps will remain until an enterprise embraces identity unification.