okta

Businesses at Work

2026

Contents

Introduction	02
Securing AI agents	05
Security concerns are driving change	14
Securing identity before, during, and after authentication	24
Securing onboarding and digital experience	34
Security-driven governance	39
Securing on-prem and hybrid IT environments	46
Protecting non-human identities	51
Conclusion	60
Methodology	64

Over the last decade, the Businesses at Work report has examined how organizations deploy and secure the software that powers modern work

The report draws on anonymized data from the Okta Integration Network (OIN) of more than 8,000 pre-built enterprise integrations, ranging from SaaS apps and HR systems to security infrastructure and advanced identity tools. Augmented by insights from technology leaders, this foundational data lets us track how enterprise technology adoption continues to evolve, and provides data-backed insights into how identity strategies must evolve alongside it.

A defining shift during this period has been the rapid expansion of SaaS. As organizations adopted hundreds of applications to support their workforce, identity became the foundation for securing access across this growing ecosystem through capabilities such as single sign-on, multi-factor authentication, and lifecycle management.

Today, the enterprise technology landscape is evolving again with the emergence of AI agents, autonomous systems that retrieve information, execute workflows, and interact with enterprise applications and systems. As these agents become embedded in business operations, they introduce a new class of identity that must be governed alongside human users.

While organizations recognize the potential of AI agents to improve productivity and automate workflows, adoption remains measured. For organizations looking to scale agentic AI, the question becomes whether their identity foundations are ready. As organizations introduce AI agents into enterprise environments, those agents must become part of the identity security fabric that governs access, permissions, and activity across systems. The strength of that fabric will ultimately determine how safely and effectively organizations can adopt and scale AI agents.

To frame this analysis, the report draws on two models. The identity security fabric (ISF) defines the architectural approach for securing identities across modern environments, connecting authentication, lifecycle management, governance, and threat detection. The identity maturity model (IMM) measures the maturity and robustness of these capabilities, providing a structured way to assess how prepared organizations are to operate securely as identity complexity grows.

Using these frameworks, this year's report evaluates identity readiness across four critical areas of the identity security fabric:

- **Authentication.** Are your authentication controls strong enough to reliably protect workforce access and secure interactions with enterprise systems as autonomous agents begin operating alongside human users?
- **Lifecycle + access automation.** Can you provision, adjust, and revoke access through automated, policy-driven processes so identities, including autonomous agents, can be evaluated on a continuous basis?
- **Governance.** Is your organization prepared to enforce policies, reviews, and oversight in real time, even as identity volume and activity increase dramatically?
- **Nonhuman identities (NHIs).** Does your organization have high visibility and control over service accounts today—identities autonomous agents will increasingly leverage to interact with enterprise systems?

Together, these capabilities form the foundation of modern identity security and will determine how prepared organizations are to securely adopt and scale AI agents.

Understanding the identity maturity model (IMM)

Organizations improve identity security by deploying capabilities that strengthen outcomes across three areas: **security and compliance, operational agility, and end-user experience**. The **identity maturity model (IMM)** describes how identity programs progress as these capabilities mature.

- **Stage 1: Foundational.** Baseline identity visibility and core protections reduce common credential-based attacks.
- **Stage 2: Scaling.** Expanded MFA adoption and improved lifecycle processes strengthen security while improving operational efficiency.
- **Stage 3: Advanced.** Phishing-resistant authentication, lifecycle automation, and stronger governance significantly reduce identity risk.
- **Stage 4: Strategic.** Identity operates as a continuous control layer, enabling secure access, efficient identity operations, and seamless user experiences across modern digital environments.

Chapter 1

Securing AI agents

Enterprise enthusiasm for agentic AI is high, with 91% of the organizations we surveyed reporting they currently use AI agents. But most organizations remain in early or limited stages of deployment, a constraint that reflects a recognition that governance, compliance and identity risks must be addressed before agents can scale safely in production. The data suggests a foundational truth for today's enterprise:

Agentic AI readiness *is identity readiness.*

As system actions shift from human-driven to autonomous, any legacy IAM models built for static access decisions no longer suffice. Organizations must evolve identity controls to provide continuous visibility and oversight as autonomous agents begin acting across enterprise systems.

As organizations progress along their individual identity maturity journeys, AI agents introduce a new layer of complexity, operating as identities that can independently make decisions and take action across the enterprise. Scaling agentic AI safely requires identity maturity that can continuously govern autonomous agents alongside human and other non-human identities. In this environment, existing identity maturity gaps do not remain contained; they expand.

Importance of IAM when adopting AI

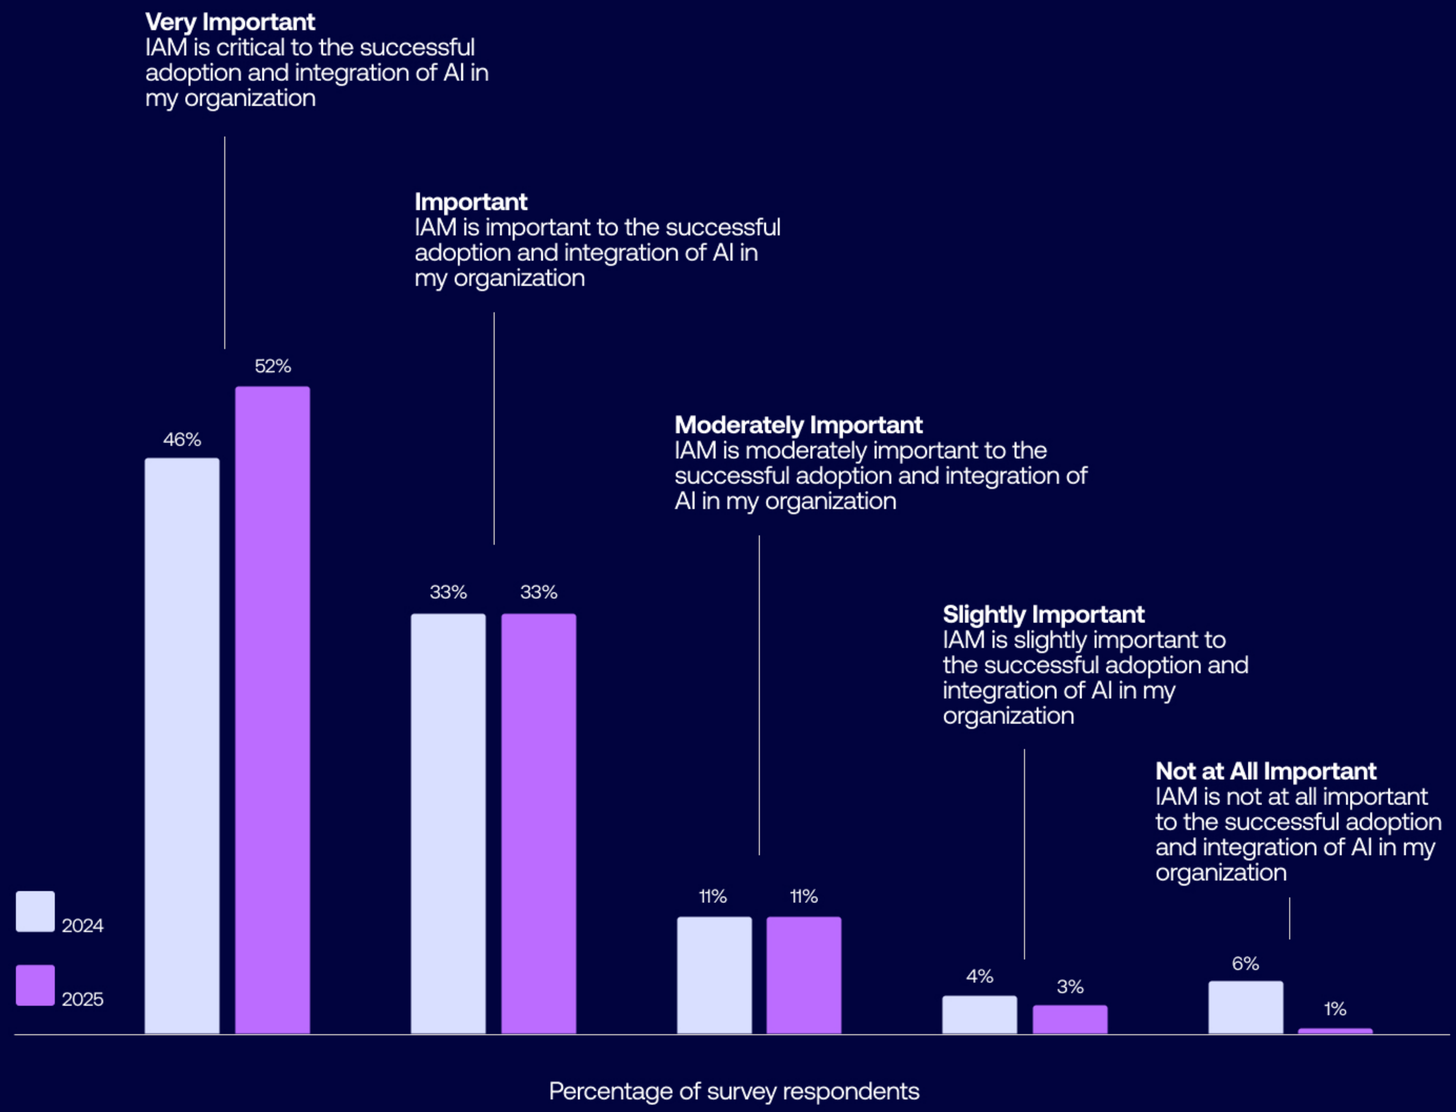

The identity imperative for AI

Okta's survey of C-suite leaders (see Methodology, page 64) reveals a strong consensus: securing AI must begin with securing the identities that interact with it. Nearly all respondents (99%) place some level of importance on Identity and Access Management's (IAM) role in AI adoption. More than half of leaders now view IAM as "very important" to successful AI adoption and integration within their organizations, rising from 46% in 2024 to 52% in 2025.

This upward trend aligns with Gartner's 2026 Cybersecurity Trends[1], which identifies IAM for AI agents as a global priority. Gartner predicts that 40% of enterprise applications will feature task-specific AI agents by the end of 2026[2]—a massive jump from less than 5% in 2025—making the urgency to govern these machine actors paramount. Without a robust identity foundation capable of handling automated credentialing and policy-driven authorization, the theoretical benefits of AI agents cannot be safely realized in production environments.

Primary use cases for AI agents

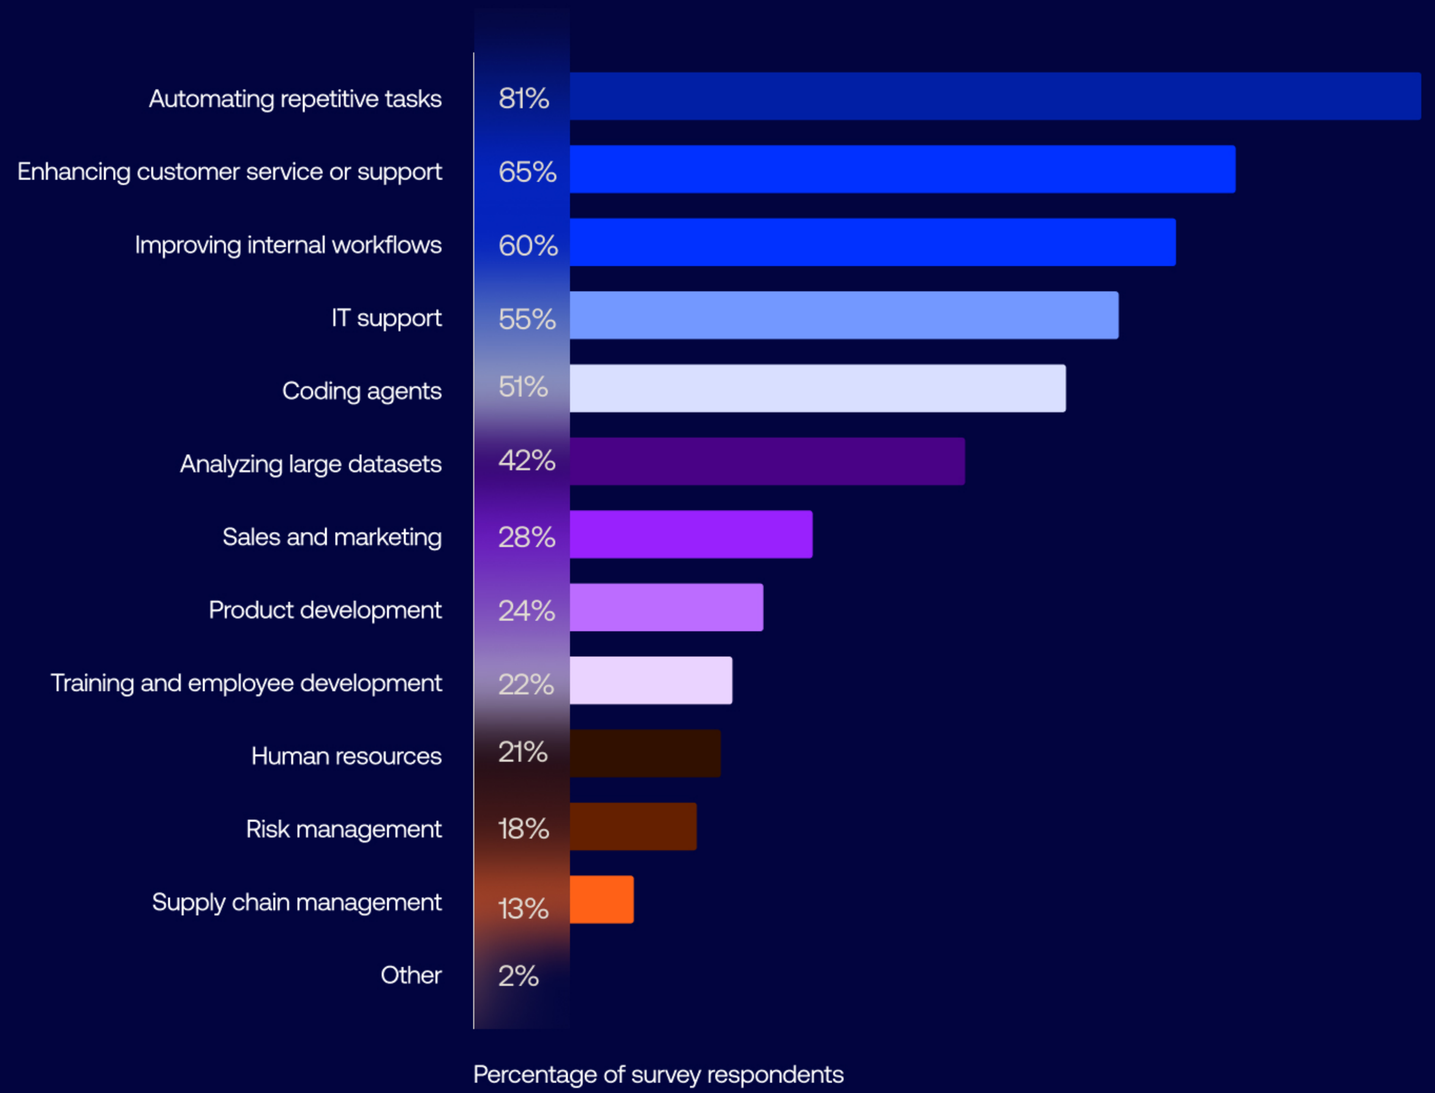

*1 Gartner®, Press Release, February 5, 2026, "Gartner Identifies the Top Cybersecurity Trends for 2026"
*2 Gartner®, Press Release, August 26, 2025, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025"

High-stakes automation: even routine tasks introduce risk

Today's AI agents are primarily being deployed to take on repetitive, operational tasks. According to Capgemini's 2025 Rise of Agentic AI report, "Most AI agents currently operate at low levels of autonomy, primarily as simple agents or semi-autonomous agents." And our survey data confirms that the top use cases include automating repetitive tasks (81%), enhancing customer service or support (65%), and improving internal workflows (60%).

Task lengths are increasing, an indicator that model capabilities are growing. But even for relatively trivial digital tasks, access requirements must be taken seriously. These functions inherently touch privileged actions, sensitive systems, and high-volume operations. As AI agents rapidly expand the non-human identity footprint across enterprises, they require substantially stronger governance.

Failing to secure these workflows can be costly. IBM's 2025 Cost of a Data Breach Report found that 97% of organizations that experienced breaches involving AI models or applications lacked proper AI access controls, and that incidents involving unsanctioned or "Shadow AI" added an average of $670,000 to the total cost of a data breach ($4.44 million is the global average). When agents operate autonomously within sensitive workflows, any compromised credential or API key can result in rapid, systemic exposure.

Top security concerns related to AI agents in next 3 years

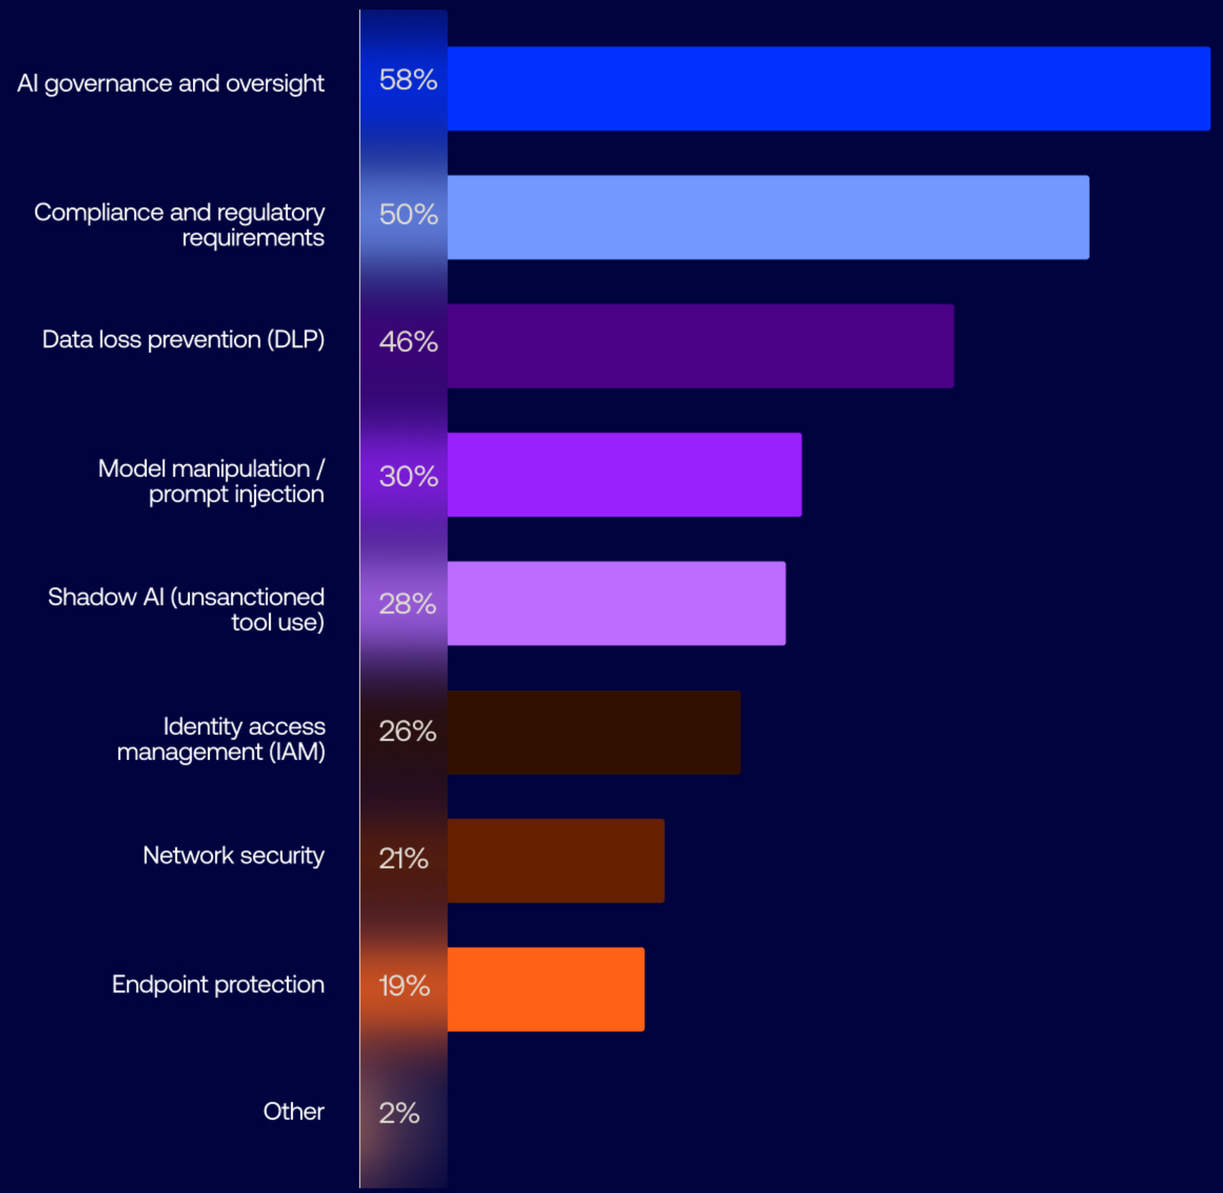

Governance tops the list of enterprise concerns

AI agents do not sit neatly behind corporate firewalls. They constantly execute API calls across internal and external systems, interacting dynamically with enterprise data and services. Because an agent can itself become the threat vector—whether through an overly permissioned workflow, a hallucination, or a prompt injection attack—every agent action must be continuously authenticated and governed by least privileged access controls.

For our survey respondents, this necessary focus on governance is already an operational reality. AI agent governance and oversight ranks as the top security concern (58%) related to AI agent adoption, followed closely by compliance and regulatory requirements (50%). Tellingly, identity and access management (IAM) concerns rank higher, at 26%, than both network security (21%) and endpoint protection (19%).

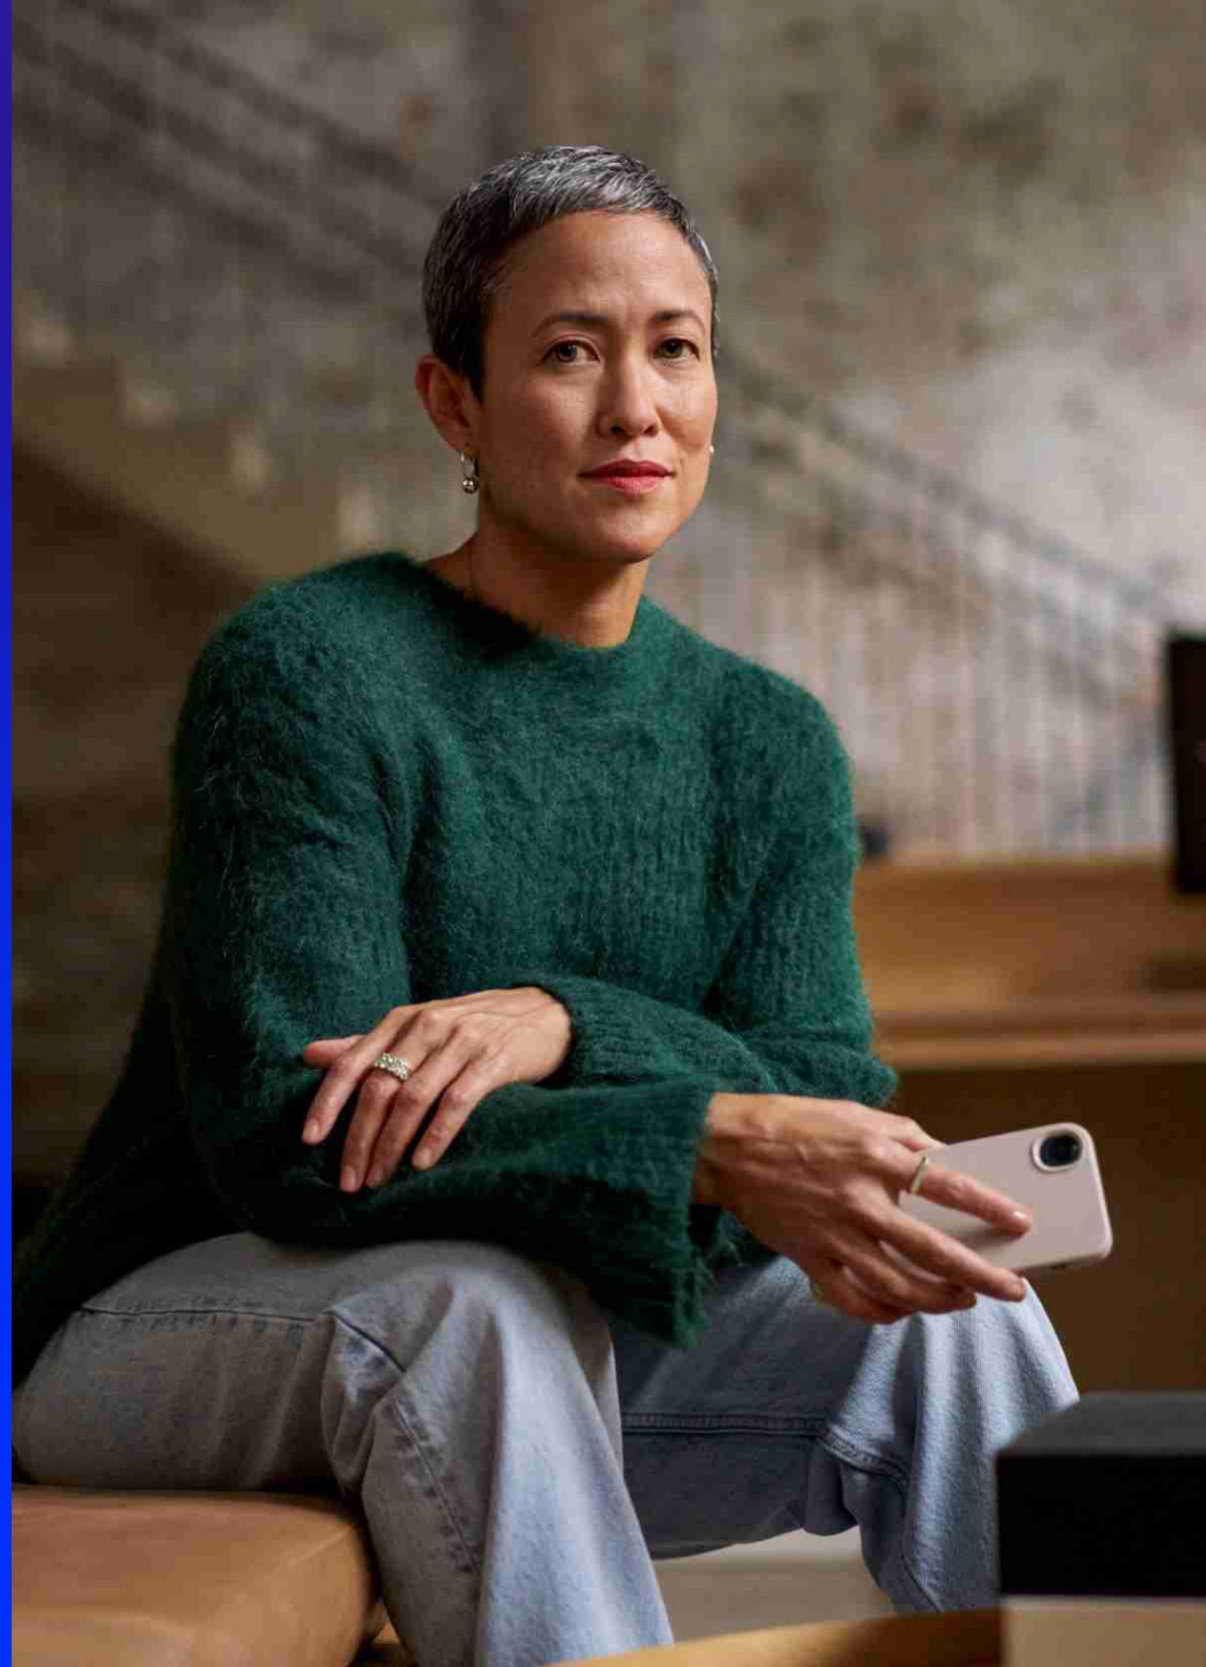

Despite this awareness, a dangerous governance gap remains. AI agents authenticate to applications and APIs using credentials and tokens, often acting on behalf of users to execute workflows, yet only 32% of organizations currently govern AI agents with the same rigor as human identities. This blind spot could be extremely dangerous: According to the 2025 Verizon Data Breach Investigations Report, 88% of basic web application attacks involve stolen credentials. Identity and access management remains the key control plane, and weak authentication controls and over-privileged access, as autonomy scales, can quickly and dramatically expand the credential-based attack surface.

Extent of use of AI agents

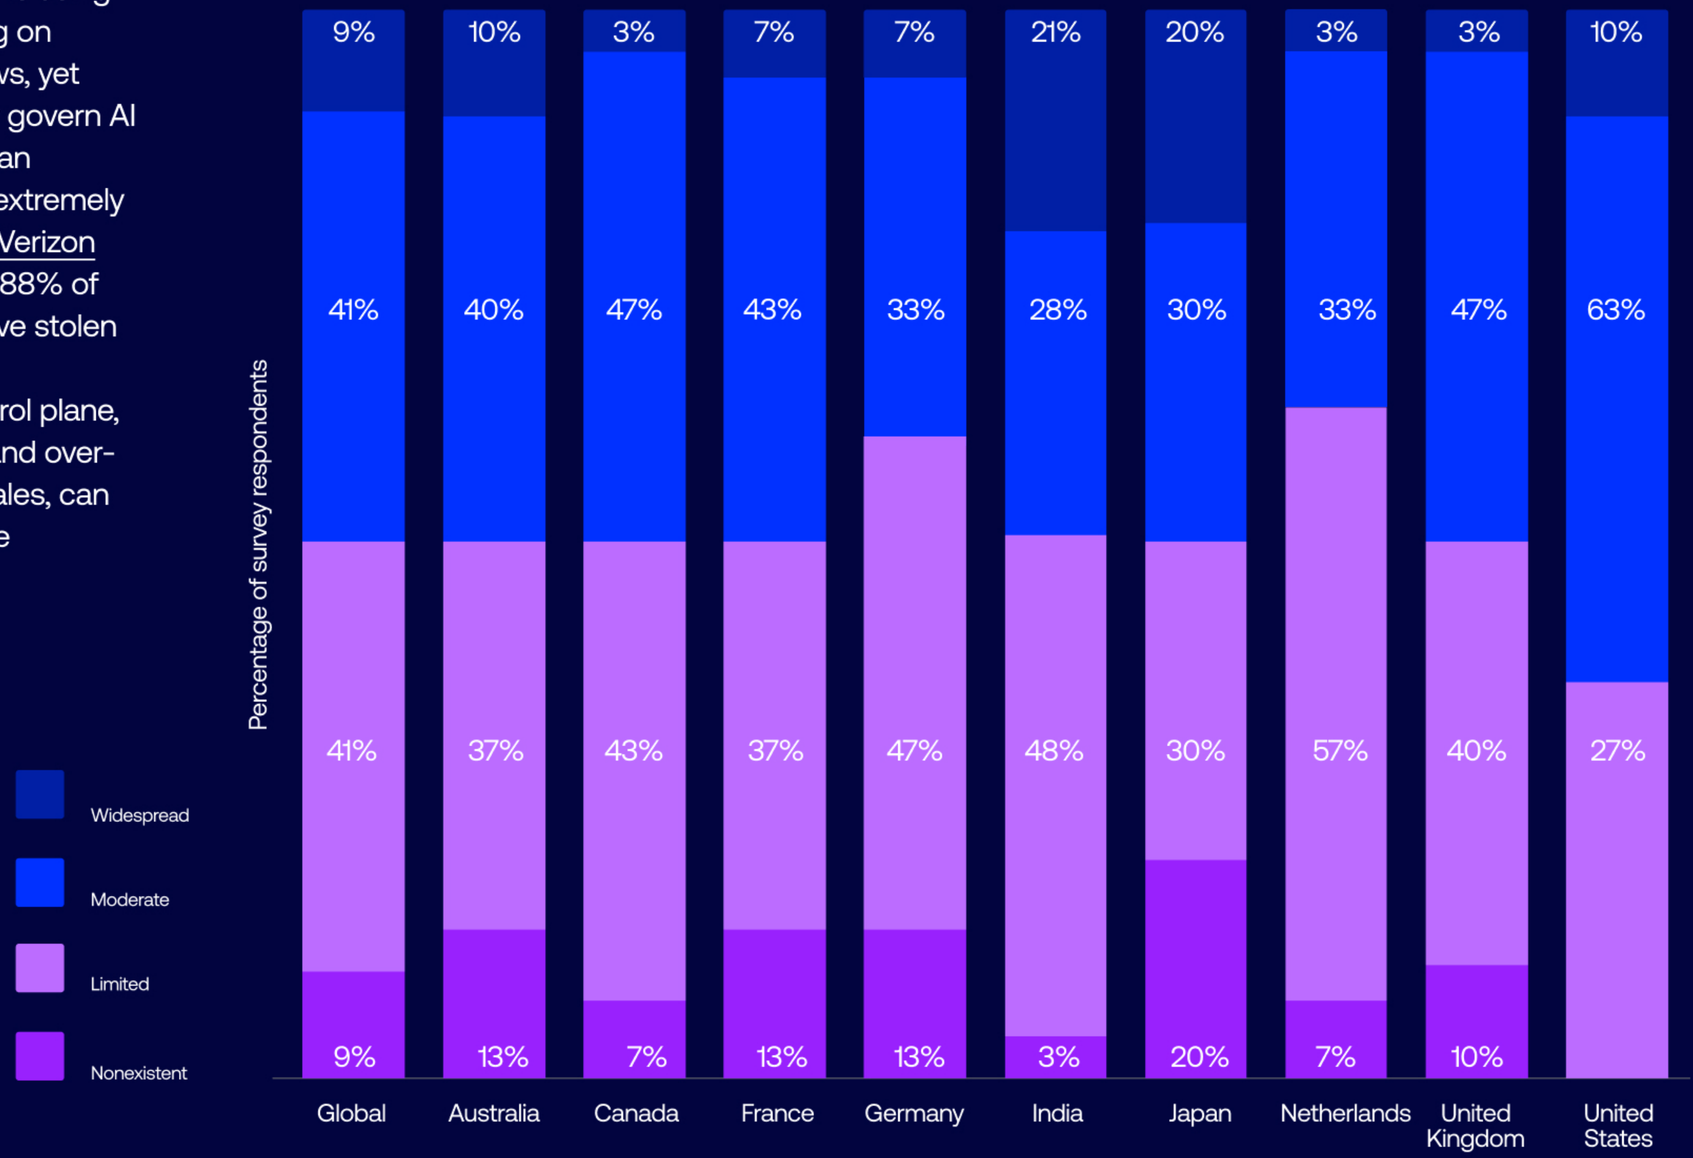

The global divide in AI deployment

While the tension between AI ambition and identity security is a universal challenge, the pace at which organizations are navigating it varies around the globe. The mandate to govern these machine actors is clear, but how quickly businesses are moving forward depends heavily on their local regulatory climates.

Globally, AI agent adoption remains an emerging phenomenon, with usage split evenly, among our survey respondents (see Methodology), between "moderate" (41%) and "limited" (41%). However, country-level data reveals meaningful differences. Heavily regulated markets like Germany and France show more cautious adoption, while the United States shows the highest combined moderate and widespread adoption. In India and Japan, widespread use is already elevated relative to other regions.

These differences reflect varying levels of regulatory pressure and identity readiness. As scrutiny of automated decision making increases and executive accountability for compliance failures tightens, organizations face growing expectations to govern autonomous systems with the same rigor as human identities.

What was once a regional best practice is transitioning into a universal mandate: scaling agentic AI safely depends on identity governance maturity.

Scaling agentic AI safely depends on identity governance maturity.

Identity maturity model focus: Assessing enterprise readiness for agentic AI

Current maturity signal: The data from the C-suite representatives Okta surveyed shows that while enthusiasm for deploying agentic AI is high, it is primarily used for relatively trivial use cases like automating repetitive tasks or enhancing customer service. These non-holistic, incremental solutions signify a lower identity maturity stage, inadequate to govern an agentic AI's entire identity lifecycle.

Capabilities required to advance: For organizations still in the early stages of identity maturity, this caution is appropriate: Safely deploying agentic AI for more complex tasks requires companies to move beyond static access granting. They need to be able to provide continuous visibility and real-time auditing of agentic behavior, with controls spanning the full identity lifecycle, including discovery, identity assignment, authorization, and access to credentials and secrets. To create an AI-enabled ecosystem that mitigates the risk of AI sprawl while optimizing its ability to be a force multiplier for business, customers need to advance their identity maturity to Stage 3 (Advanced).

Outcomes unlocked: Foundational gaps in visibility and control over autonomous systems can impede an organization's ability to safely manage AI agents at scale. But with smart investments targeting security & compliance and operational agility, they can safely set a course for reaping the benefits of the power of AI.

Chapter 2

Security concerns are driving change

Agentic AI adoption is accelerating within an enterprise environment where security investment and scrutiny have already been rising for years

While endpoint security spending is rising, **organizations see identity and access management as more critical to securing AI agents.**

The digital landscape has radically transformed since our first Businesses at Work report more than a decade ago. Today, organizations are balancing hybrid work, ever-expanding SaaS ecosystems, and now the rapid introduction of AI-powered tools and agents into everyday work.

The charts in this section of the report each draw on a year's worth of data from the Okta Integration Network. The OIN's pre-built integrations enable thousands of global enterprises to equip their workforces with seamless access to best-of-breed apps, IT infrastructure, and other technology. Millions of daily OIN authentications and interactions around the world supply the aggregated, anonymized data that has fueled these global insights for more than a decade now.

Unsurprisingly, as the threat landscape grows more complex and organizations introduce AI-powered tools and autonomous agents, security remains a central concern. Half of the fastest growing apps this year are security tools, covering endpoints, networks, and password management, with technology, retail, and education leading the charge. NinjaOne far outpaces the pack as this year's fastest-growing app overall, exhibiting massive 240% YoY growth. However, when it comes to AI agents, the organizations we surveyed rank identity and access management (IAM) as a greater security concern than endpoint and network security.

Half of this year's fastest-growing apps are security solutions

Security solutions have claimed an outsize share of the fastest-growing apps since at least 2021. This year, security tools account for half of the fastest-growing apps, up from 40% last year.

Endpoint management tools saw particularly strong growth. NinjaOne appeared out of nowhere to lead the pack with a 240% YoY growth in number of customers, and CrowdStrike Falcon took second place with 66% YoY growth. Notably, we have not seen an endpoint management tool in the fastest-growing apps since Kandji took the top spot in 2022 data, indicating how companies are shifting their priorities to contend with new security challenges.

How organizations balance these security priorities also varies by region, often reflecting differences in how quickly enterprises are adopting AI agents.

Highly regulated countries and regions, for example, may tend to prioritize security, and consequently may adopt AI agents more cautiously. Companies around the world are trying to balance their security investments with network and collaboration tools as well: GitHub and AWS feature prominently around the world, and this year's most popular apps list (ranked by number of customers) includes a quick rise in adoption of the collaborative design tool Figma, and the entrance of Zoom to the top 15 among Fortune 500 companies in the OIN.

50% of the fastest-growing apps are security tools, **up from 40% last year.**

Fastest-growing apps, by number of customers

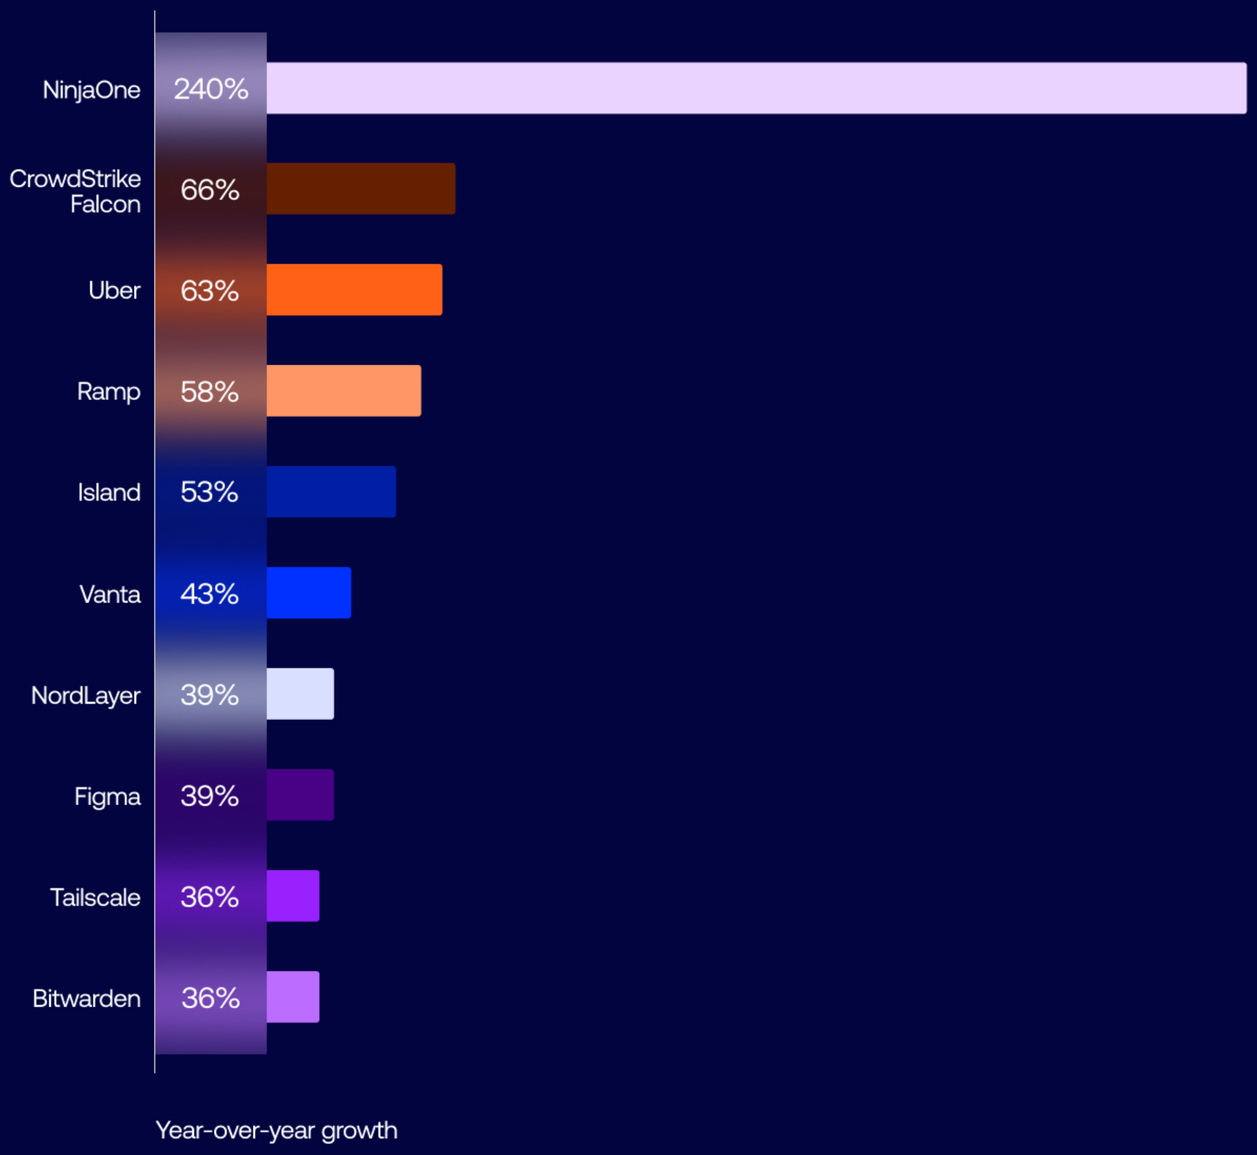

Another newcomer to the fastest-growing apps list is Nordlayer, a network security tool enjoying a solid 39% YoY growth in adoption. Two security products from last year's list – network security tool Tailscale and password manager Bitwarden – have managed to remain on the list this year, each growing at 36% YoY. (Last year Tailscale grew at 59% YoY and Bitwarden clocked in at 63%.)

Further underlining the overwhelming importance of security to companies today, we also see two tools that are "security-adjacent." Island, an enterprise browser with embedded security which joined the list this year, and data compliance tool Vanta.

The highest-ranking non-security app on this year's list is recurring stalwart Uber. Last year, a foray into premium travel with Uber Business Black helped drive 56% YoY growth; this year Uber managed to increase that to 63% with further innovation, like a new pilot program designed to streamline airport dropoffs. Expense management app Ramp also extended its stay on the list, showing 58% YoY growth (compared to last year's 56%).

Global apps-per-company adoption rate slows, but not everywhere

After several years of slow but steady growth, the total number of apps each company utilizes has settled a bit, from 101 last year to 98 this year. That's the global figure across all business sizes: Small businesses (those with < 2000 employees) use an average of just 72 apps, while larger businesses (with 2000 or more employees) average over three times that, at 259.

From a regional perspective, EMEA and APAC are still growing their toolkits, with total apps per company in those regions up an average of 5% and 3% YoY, respectively. Individual countries where the number of apps per company is still growing strong include standouts Italy and France (up 21% and 17% YoY), and Japan and South Korea, where companies are each using 15% more apps than they did last year.

The security tool 1Password showed the highest industry-level growth, notching a **370% YoY increase in the technology sector.**

By industry, the fastest-growing apps include 1Password (tech) and GitHub (healthcare)

When we take an industry-level look at the fastest-growing apps across 11 key industry sectors, we find security and collaboration tools dominating. The continued focus on security makes sense, given the new risks of autonomous agents. 1Password showed the highest industry-level growth (in the technology sector), notching a 370% YoY increase in unique users. Design tool Canva takes the collaboration prize by number of customers with 34% YoY growth in finance and banking.

Last year the nonprofit sector ranked design app Canva at the top, both by number of customers and by unique users. This year, nonprofits are emphasizing project management (monday.com is up 19% YoY by number of customers) and cloud platforms (AWS is up 143% by unique users). Healthcare and pharmaceuticals made GitHub their top pick for both metrics this year, rising 13% YoY by number of customers and 68% by unique users.

In education, decisionmakers have been building diverse toolkits, favoring security apps two years ago, then cloud and developer tools last year. This year they're coming back to security (Palo Alto Networks grew 36% YoY by unique users) and collaboration (Zoom, with respectable 4% YoY growth in customers).

Fastest-growing apps, by industry

Industry	By number of customers	By number of unique users
Arts, Entertainment, and Recreation	**GitHub** ▲ 15%	**Microsoft 365** ▲ 55%
Education	**Zoom** ▲ 4%	**Palo Alto Networks** ▲ 36%
Finance and Banking	**Canva** ▲ 34%	**Postman** ▲ 70%
Healthcare and Pharmaceuticals	**GitHub** ▲ 13%	**GitHub** ▲ 68%
Insurance	**GitHub** ▲ 15%	**GitHub** ▲ 52%
Manufacturing	**Figma** ▲ 21%	**Adobe Creative Cloud** ▲ 121%
Media and Communications	**Datadog** ▲ 19%	**Snowflake** ▲ 39%
Nonprofit	**monday.com** ▲ 19%	**AWS** ▲ 143%
Professional Services	**Ramp** ▲ 66%	**Ramp** ▲ 293%
Retail	**Datadog** ▲ 19%	**Jamf** ▲ 43%
Technology	**Ramp** ▲ 32%	**1Password** ▲ 370%

Year-over-year growth

By country, the fastest-growing apps include NinjaOne in America and Notion in Japan

Looking at the fastest-growing apps by country reveals an interesting mix of app success stories and big shifts from the prior year. Last year was a great time for developer tools, with GitHub the fastest-growing app across Japan, Australia and Israel. This year the picture is more diverse. Japan's fastest-growing app is now the collaboration tool Notion (43% YoY growth), and Slack takes the top spot in Israel (up 23% YoY). But GitHub continues to be the hot app in Australia (14% YoY), and Germany has caught the GitHub bug as well (up 17% YoY).

Meanwhile, in North America, there is an increasing emphasis on security, with NinjaOne (240% YoY growth) and 1Password (25% YoY) taking the US and Canada, respectively, by storm. Last year, neither country had a security tool as its fastest-growing app: The US was focused on adding data compliance, and Canada's top choice was the collaborative design tool Figma.

In the UK, last year's fastest-growing app was the data platform Snowflake, while this year it's the CRM tool HubSpot, with 24% YoY growth. France, in contrast, was all in for cloud-based CRM tool Salesforce last year, and this year their favorite tool is AWS (up 8% YoY). AWS is the fastest-growing app in South Korea, too (up 11% YoY).

Fastest-growing apps, by country

United States	NinjaOne	▲ 240%
Japan	Notion	▲ 43%
Canada	1Password	▲ 25%
United Kingdom	HubSpot	▲ 24%
Israel	Slack	▲ 23%
Germany	GitHub	▲ 17%
Australia	GitHub	▲ 14%
South Korea	AWS	▲ 11%
France	AWS	▲ 8%

Year-over-year growth

This year's most popular apps show companies laying the groundwork for agentic AI deployment

Each year we rank the year's most widely deployed apps, also known as most popular apps, as measured by number of customers. We look at this three ways: across companies of all sizes, across tech startups (defined as companies with 100 or fewer employees in the tech sector), and across Fortune 500 companies. If the fastest growing apps represent new momentum, the most popular apps are the familiar platforms companies rely on, year after year.

By its nature, the most popular apps list is traditionally more static than the fastest-growing apps, and this year's no different. What's changed is that the most widely deployed apps (AWS, Google Workspace, and Microsoft 365) are becoming the primary systems through which AI agents will operate. As a result, deploying identity and access controls that allow agentic AI to operate safely across these platforms is now a top concern. Further down the list, collaborative design tool Figma leaps onto this year's chart at an impressive No. 11 among all companies, collaboration phenomenon Slack rises one spot to No. 6, and developer darling GitHub leapfrogs Docusign to land at No. 8.

▲ Higher than overall

▼ Lower than overall

Most popular apps

	Overall	Startups	Fortune 500
1.	Microsoft 365	AWS ▲	Microsoft 365
2.	Google Workspace	Google Workspace	Salesforce ▲
3.	Amazon Web Services (AWS)	Slack ▲	AWS
4.	Salesforce	Microsoft 365 ▼	GitHub ▲
5.	Zoom	Atlassian Product Suite ▲	ServiceNow ▲
6.	Slack	GitHub ▲	Atlassian Product Suite ▲
7.	Atlassian Product Suite	Zoom ▼	Snowflake ▲
8.	GitHub	Salesforce ▼	Palo Alto Networks ▲
9.	Docusign	HubSpot ▲	Google Workspace ▼
10.	KnowBe4	Datadog ▲	Workday ▲
11.	Figma	Zendesk ▲	Lucid ▲
12.	Cisco Meraki	Docusign ▼	Zscaler ▲
13.	Lucid	1Password ▲	Slack ▼
14.	Jamf	Figma ▲	Splunk ▲
15.	Palo Alto Networks	PagerDuty ▲	Zoom ▼

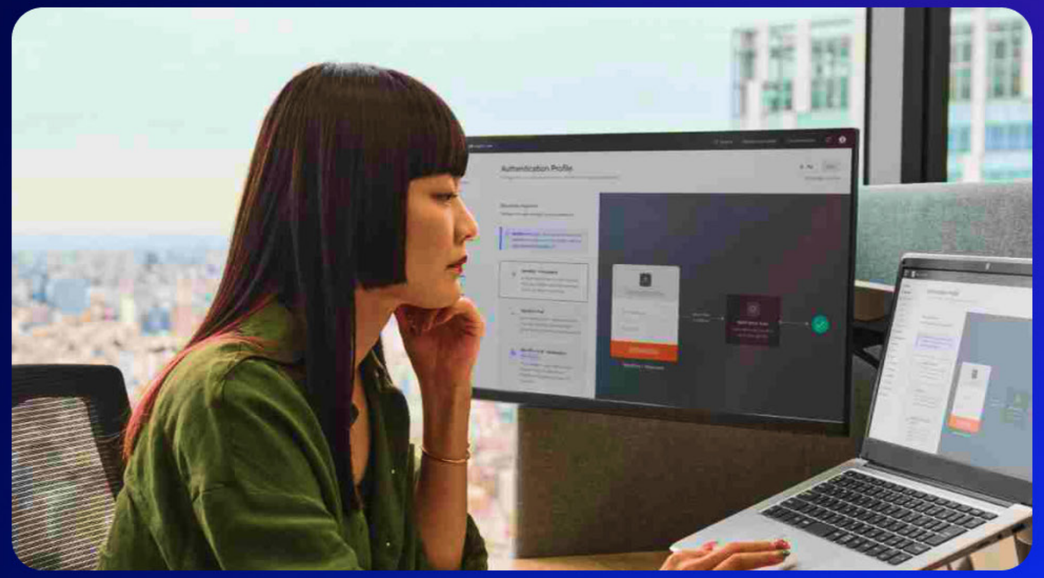

Unsurprisingly, there is some churn in the scrappy world of startups. On that roster, the CRM platform Hubspot, the developer tool Datadog, and the password manager 1Password each moved up two spots. Among Fortune 500 companies, a variety of apps each climbed one notch this year, with GitHub rising to No. 4, and Snowflake, Google Workspace, Lucid Software, Zscaler, and Slack each taking a step up. Finally, Zoom which showed a respectable popularity overall (No. 5) and among startups (No. 7) gets some respect at last from larger companies, joining the chart at No. 15 among Fortune 500 companies in the OIN.

More than half of companies now supplement Microsoft with best-of-breed apps

Okta customers that deploy Microsoft 365 often supplement their suite with various best-of-breed solutions for nominally “included” functions like video conferencing, messaging, and file storage. We define “best-of-breed apps” as those that dominate their categories (by number of customers) for a particular stand-alone functionality.

For the first time ever, we’re seeing more than half of Okta’s Microsoft 365 customers supplement their suite, prioritizing the best functionality even when this means redundancy.

Companies arm employees with top tools: Increasingly, Okta's Microsoft 365 customers are supplementing out-of-the-box Microsoft 365 functionality, such as video conferencing, messaging, and file storage, with redundant functionality from Zoom (48%), Slack (42%), and Box (21%).

Parallel processing: Companies aren't just adding extra apps: Today, over half (51%) of Okta's customers deploying Microsoft 365 also deploy Google Workspace — a 3 percentage point increase over last year, and 6 points over the prior year.

Best-of-breed reigns over one-stop shopping: Instead of choosing tools such as Microsoft Dynamics 365, Azure, and Power BI, we see the increased deployment of Salesforce, AWS, and Tableau. In fact, among Okta customers deploying Microsoft 365, 41% now deploy four or more of our eight featured best-of-breed apps — up from 39% last year.

Okta's Microsoft 365 customers with best-of-breed apps

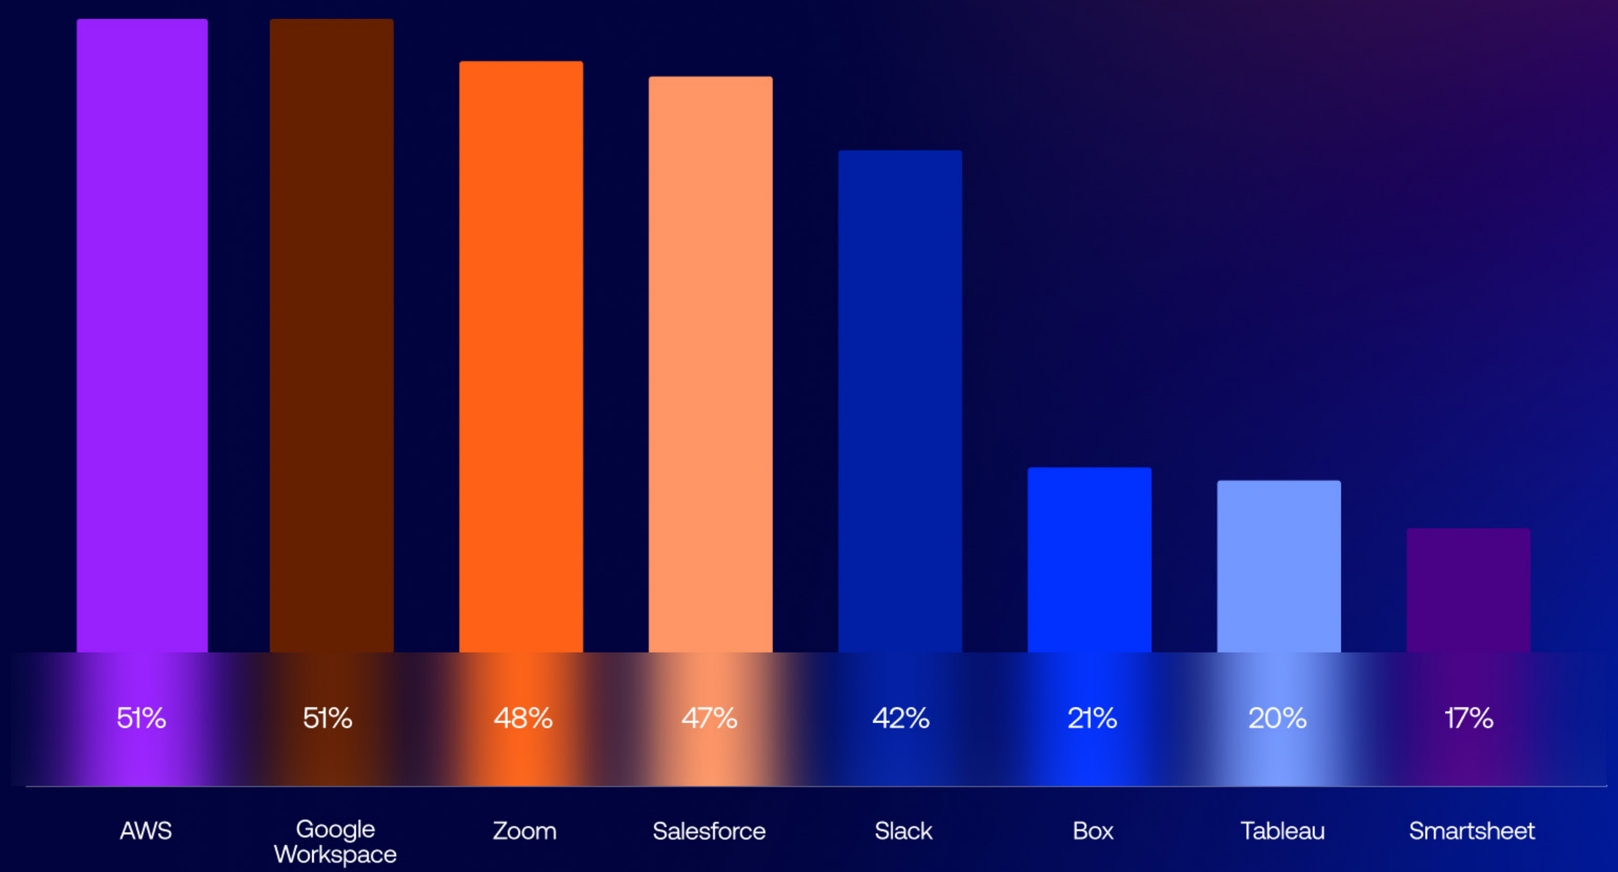

Percentage of Okta's Microsoft 365 customers with app

Chapter 3

Securing identity before, during, and after authentication

As attackers automate credential-based attacks at scale, authentication has become both the first line of defense and a growing point of failure

To securely deploy agentic AI, organizations must start by strengthening the first and most critical checkpoint in the identity lifecycle: authentication. Organizations can no longer rely on legacy, low-assurance factors like SMS or security questions that are easily intercepted by today's fraudsters.

High-assurance MFA adoption rose 8 percentage points, but threats across the landscape are accelerating 6.3x faster, as we'll detail in this section. This creates a widening exposure gap, where the keys to autonomous agents remain vulnerable to a single successful phish, shifting the MFA narrative from factor quality to defensive velocity.

The scale of the threat is exploding, as attackers weaponize AI to automate credential-based attacks, which now drive 60% of all security incidents and 88% of web app breaches, translating to an average global breach cost of approximately $4.67 million. Without the right guardrails, companies risk deploying AI agents in a digital minefield.

To navigate this era, authentication must evolve from static, one-time events into continuous, adaptive evaluation of identity and intent, anchored by phishing resistant authentication methods. Transitioning to phishing-resistant, high-assurance workflows is no longer just a defensive upgrade—it is a prerequisite for identity maturity. Only by closing this phishing gap can organizations ensure that when an agent requests access to sensitive data, the underlying identity is backed by cryptographic certainty.

Adoption of MFA is still growing, slowly (up 8 percentage points YOY), but the threats are growing faster.

MFA factors by assurance level

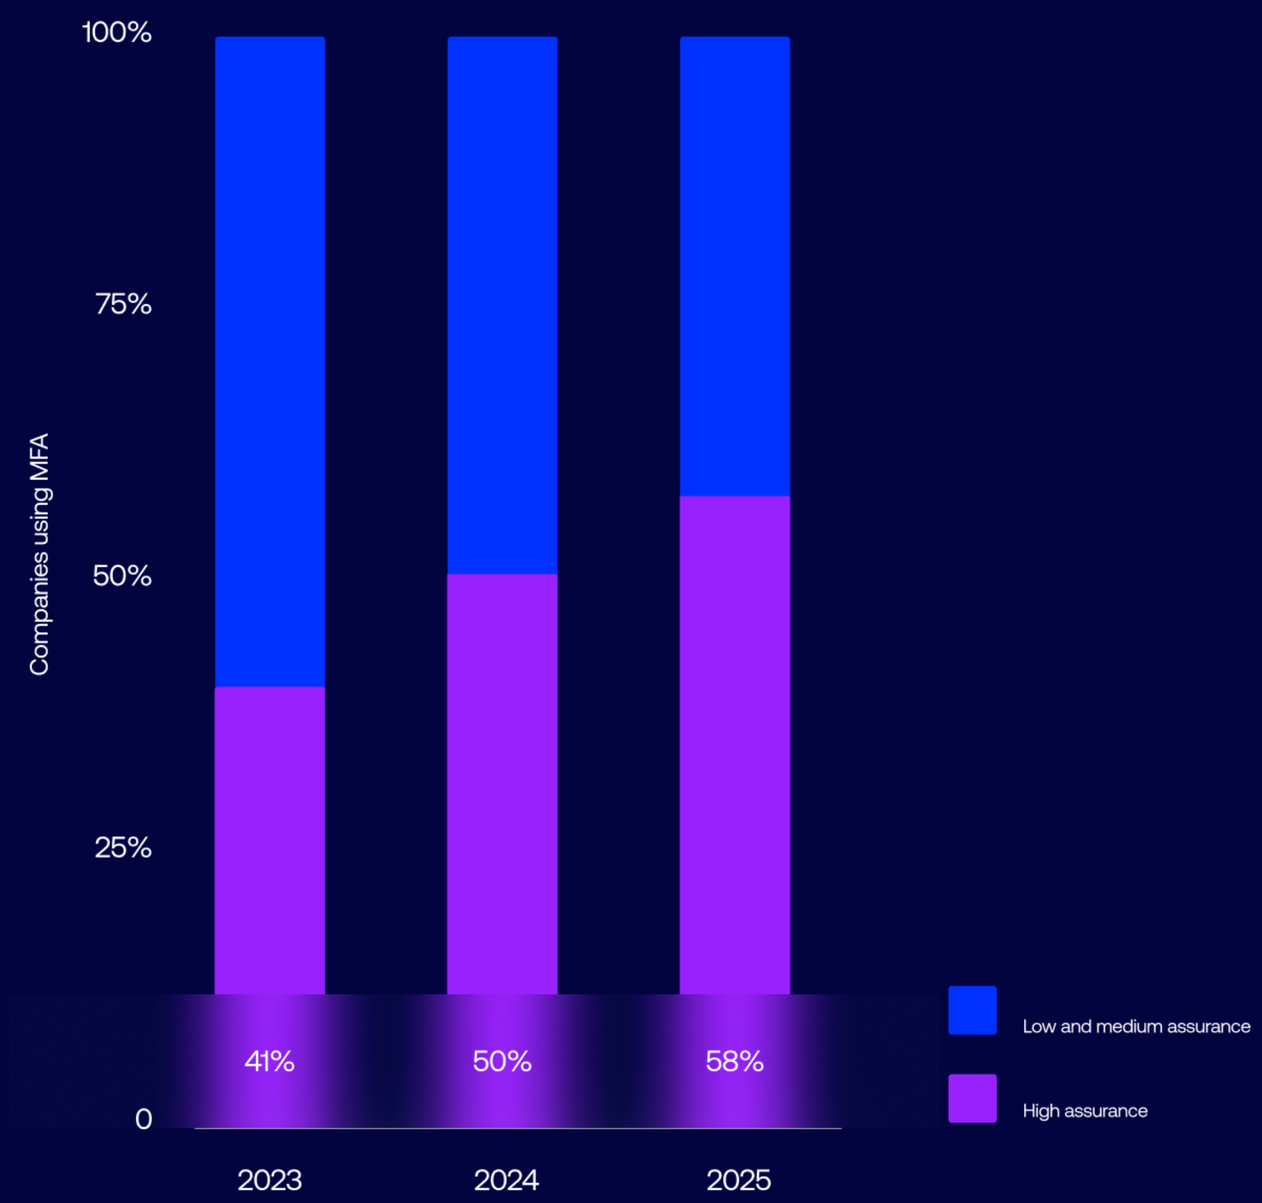

The push for high-assurance phishing resistance

Multi-factor authentication (MFA) improves security, but not all factors provide the level of assurance required as credential-based attacks continue to rise. To safely deploy agents that act on behalf of users at scale, organizations must move away from low and medium assurance factors like SMS and security questions, toward high-assurance, phishing-resistant authentication. When authentication is built on legacy factors that today's fraudsters can intercept, any agentic architecture built on top of this shaky foundation becomes a liability.

The good news: Over the last few years, we have seen an impressive migration toward high-assurance factors—including Okta FastPass. High-assurance factors are now being used by 58% of companies, an impressive step up from 41% just two years ago. At the same time, low-assurance factors like SMS, voice, and security questions are losing ground. The Verizon 2025 Data Breach Investigations Report (DBIR) says that sophisticated phishing and "pretexting" (creating a false scenario to get a victim to share sensitive information) are the main techniques fraudsters leverage to trick employees, further validating the move away from low- and medium-assurance factors.

Factor groupings

Factor grouping	Factors included	Assurance level	Phishing-resistant
Okta FastPass	Okta FastPass (with and without Biometrics)	High	Yes
Security Key or Biometrics	WebAuthn (FIDO2), YubiKey (FIDO2), Touch ID	High	Yes
Okta Verify (OTP and Push)	Okta Verify (OTP and Push) includes Okta Verify with One-Time Passcodes and Okta Verify with Push Notifications	Medium	No
Other OTP	Duo Security, Google Authenticator, HOTP, RSA SecurID, Symantec VIP, YubiKey OTP	Medium	No
Email	Email	Low	No
Security Question	Security Question	Low	No
Phone (Voice Call/SMS)	SMS or voice call	Low	No

Number of FastPass authentications

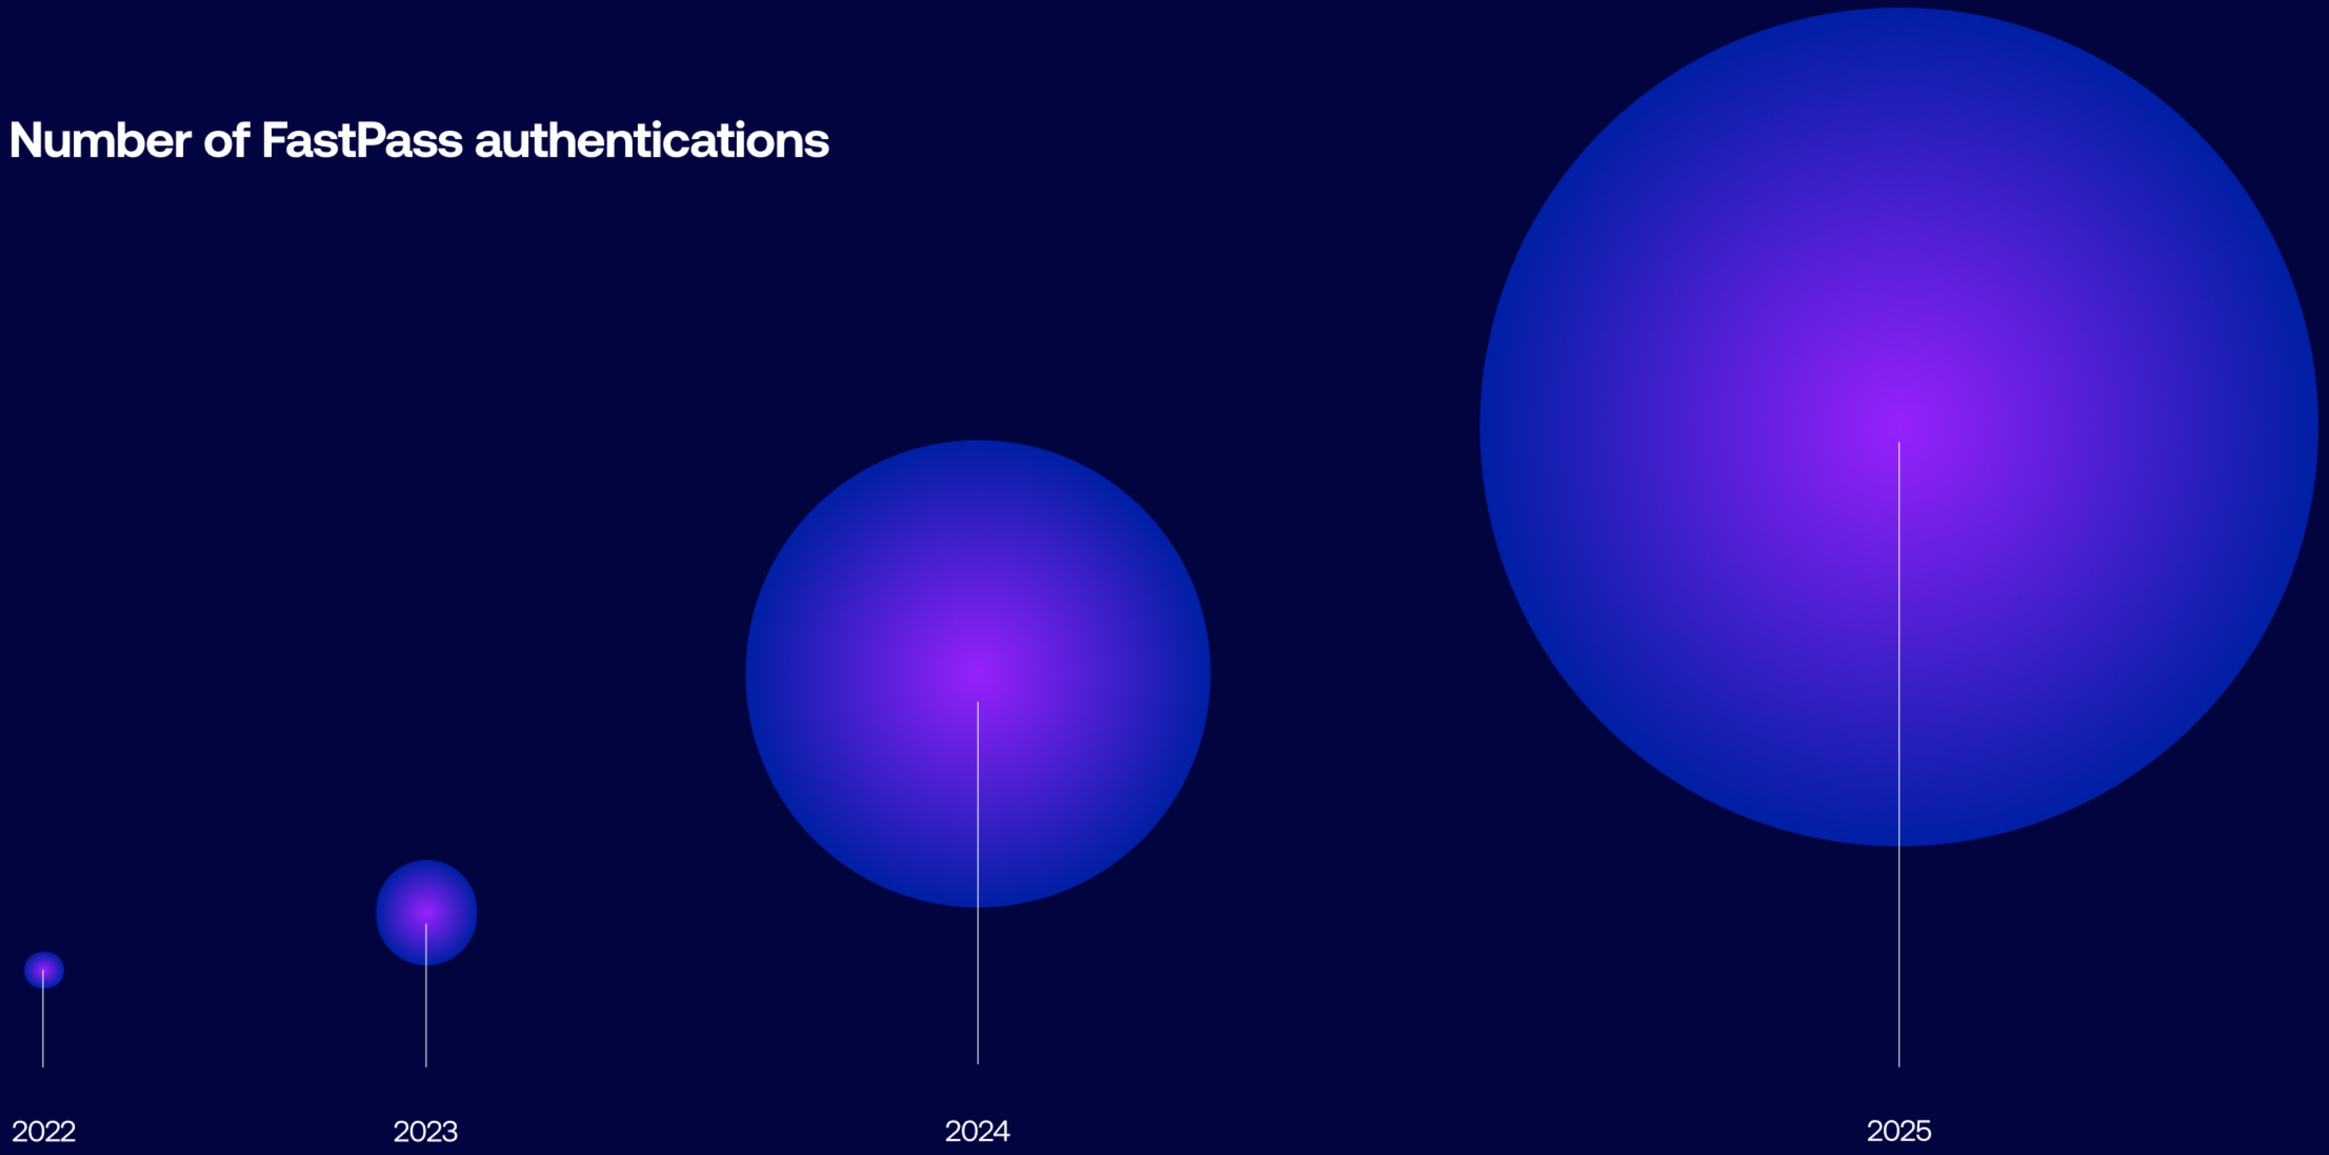

The rise of agentic AI brings this risk into sharp focus: Organizations can't hope to confidently authenticate autonomous agents if their authentication protocols can be easily intercepted. Deploying phishing-resistant authentication at scale is a critical step toward identity maturity, and helps ensure that when an agent requests access to a sensitive system, the underlying identity has been verified using high-assurance methods.

Based on 1 year of cumulative data

Accelerating the move to passwordless

Our data show strong momentum building behind phishing-resistant, high-assurance, passwordless authentication. Across the Okta Integration Network, the total volume of Okta FastPass authentications grew 81% YoY, continuing a four year trend of rapid adoption. The average number of authentications per account also rose sharply, increasing by 43% signaling a broader shift toward passwordless, high-assurance authentication across enterprise environments. This momentum is helping strengthen the identity foundation needed to safely deploy agentic AI.

A closer look at the data reveals shifts in enterprise usage. For the first time, Windows devices account for a larger share of FastPass authentications than MacOS, increasing their share of the total by 8% this year. Mobile behavior is shifting as well, with 1.9x more FastPass authentications occurring on Android than iOS. Passwordless is enjoying a moment but there's still room for improvement. Among companies using FastPass, only 16% use biometrics (same as the year prior), which provide the highest level of assurance. Android leads the way in biometric adoption at 24%.

Adaptive MFA events have **grown by 129% over the past two years.**

By region

Regional regulations (like NIS2 in Europe) are beginning to require stronger authentication controls, and cyberinsurers are increasingly requiring phishing-resistant MFA, accelerating a shift toward phishing-resistant authentication. In our data, the Netherlands has now surpassed France with the highest average number of FastPass authentications per account, and the US and Australia show strong adoption of biometric authentication, accounting for 27% and 22% of biometric authentications respectively.

AMFA evaluation events

Authentication can no longer rely on static policies alone. As credential-based attacks grow more sophisticated, organizations are increasingly turning to risk-based authentication where contextual signals such as device, location, and user behavior determine when step-up or re-authentication is triggered.

This massive recent growth in adaptive MFA (AMFA) events is most aggressive in sectors where digital and physical operations are rapidly converging. Retail leads the way, with 33% YoY growth in AMFA, helping organizations secure massive transaction volumes and diverse user populations against automated fraud and credential-stuffing attacks.

Transportation and warehousing follow closely (up 32% YoY), as companies digitize supply chains and connect trusted suppliers and freight forwarders across a highly distributed workforce. In these environments, AMFA helps protect distributed systems and workforce access from account-takeover.

Meanwhile, the insurance industry has seen a similar 29% YoY increase, using adaptive authentication to protect sensitive personal and financial data from increasingly sophisticated credential-based attacks.

Ratio of detected threats versus authentications, by industry

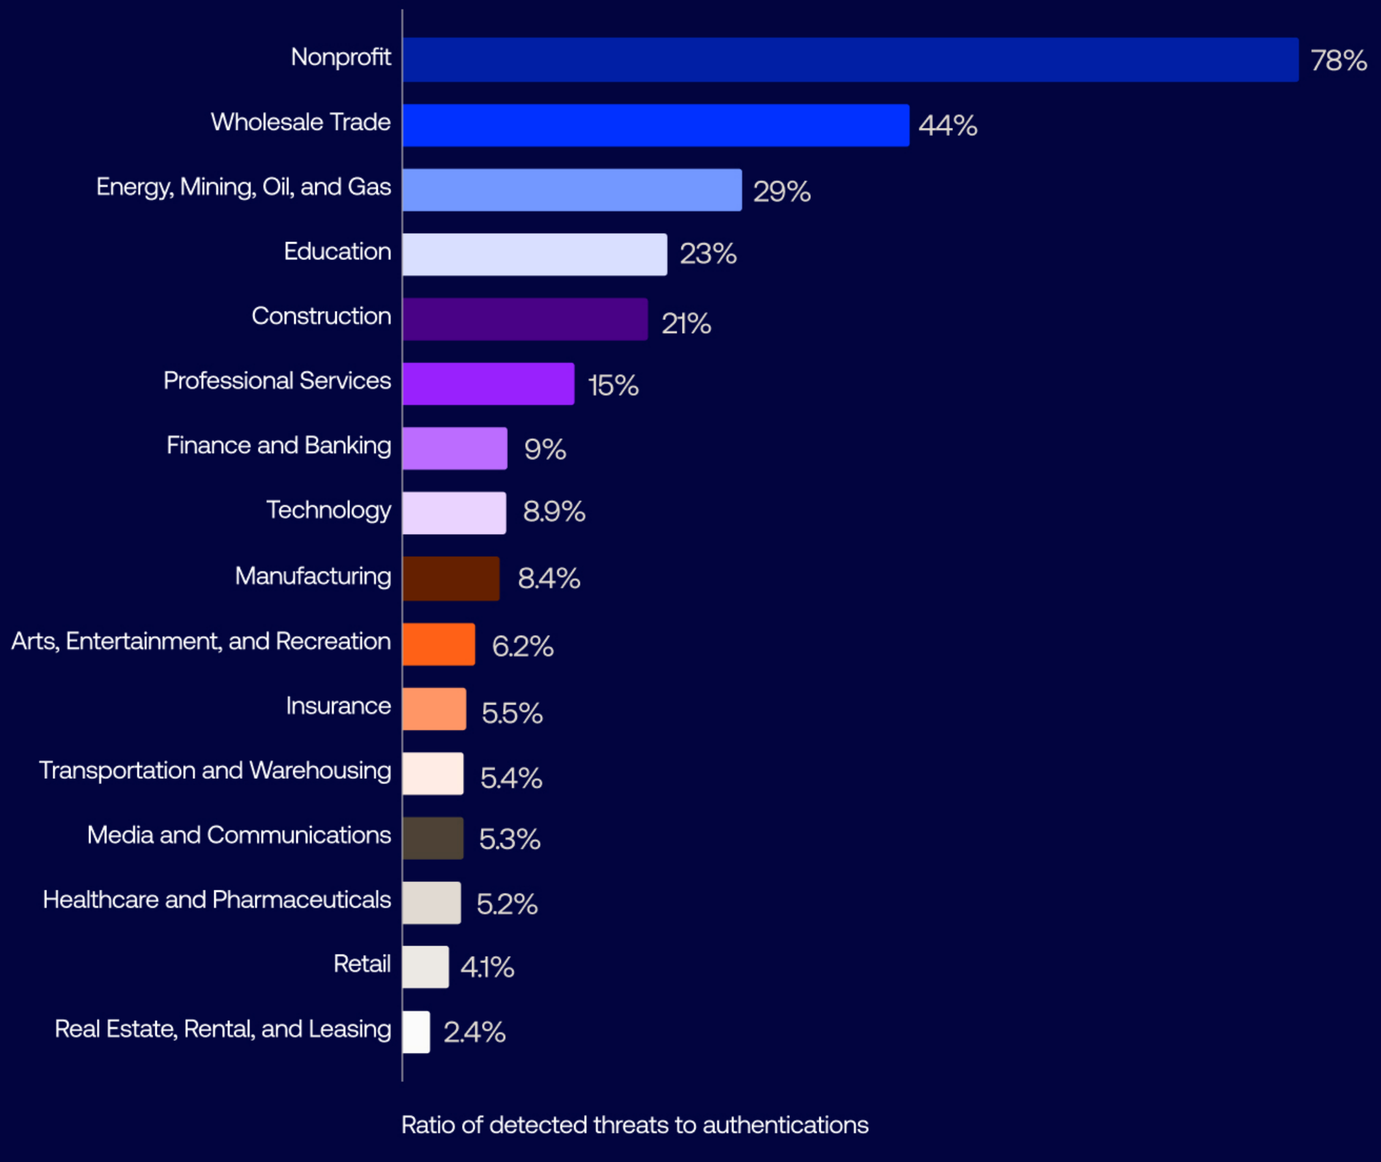

Based on 1 year of cumulative data
Percentages over 10% have been rounded

The front line: targeted industries and global trends

Our data shows that the volume of detected threats is not just increasing—it's exploding. Scaled, automated attacks are testing organizations' ability to harden their defenses fast enough to secure their enterprises, even with best practices like adopting higher-assurance MFA.

Threats in this landscape are rising 6.3x faster than the rate at which organizations are adopting high-assurance protections, creating a growing gap between attackers' ability to compromise identities and organizations' ability to defend them. This surge is driven in part by attackers weaponizing AI to automate brute-force and credential stuffing attacks at a scale legacy defenses cannot match.

Okta ThreatInsight evaluates legitimate and malicious sign-in activity across our customer base to detect risky IP addresses and thwart credential-based attacks such as password spraying, credential stuffing, and brute-force cryptographic attacks. Aggregated ThreatInsight data comparing the total volume of detected threats to the total volume of authentications provides a useful proxy for the state of cyberthreats across industries and countries.

This year, nonprofits have seen their threat-to-authentication ratio soar to 78%, an incredible leap from "only" 18% a year prior, when they ranked second behind energy, mining, oil and gas. (More detail on the specific challenges nonprofits face can be found in our third annual Nonprofits at Work Report.) Last year, the energy industry led the pack at a ratio of 32% detected threats vs authentications; that sector comes in third this year with a tamer 29%. Wholesale trade ranks second this year at 44%, compared to only 11% a year ago.

Threats in this landscape are rising 6.3x faster than the rate at which organizations are adopting high-assurance protections.

The rise of nonprofits and wholesale trade in these ranks suggests that attackers may be weaponizing AI to launch massive, automated credential-stuffing campaigns in sectors with less robust identity infrastructure than highly targeted sectors like finance and tech. Yet even finance and banking, which ranked fifth for two years, saw its threat ratio rise from 2.9% to 9% this year. And healthcare and pharmaceuticals, which had the lowest ratio among industries last year, has seen its threat ratio triple to 5.2%, showing that no industry is safe from the rise in identity-based attacks.

As identity-based attacks accelerate, the implications extend beyond human access alone. When agents begin handling more sensitive data and executing tasks on behalf of users, a single compromised identity can enable automated actions across enterprise systems. Verizon's 2025 Data Breach Investigations Report identifies credential abuse as the #1 access vector for breaches—reinforcing that phishing-resistant, adaptive authentication is essential as organizations enter the agentic era.

Ratio of detected threats versus authentications, by country

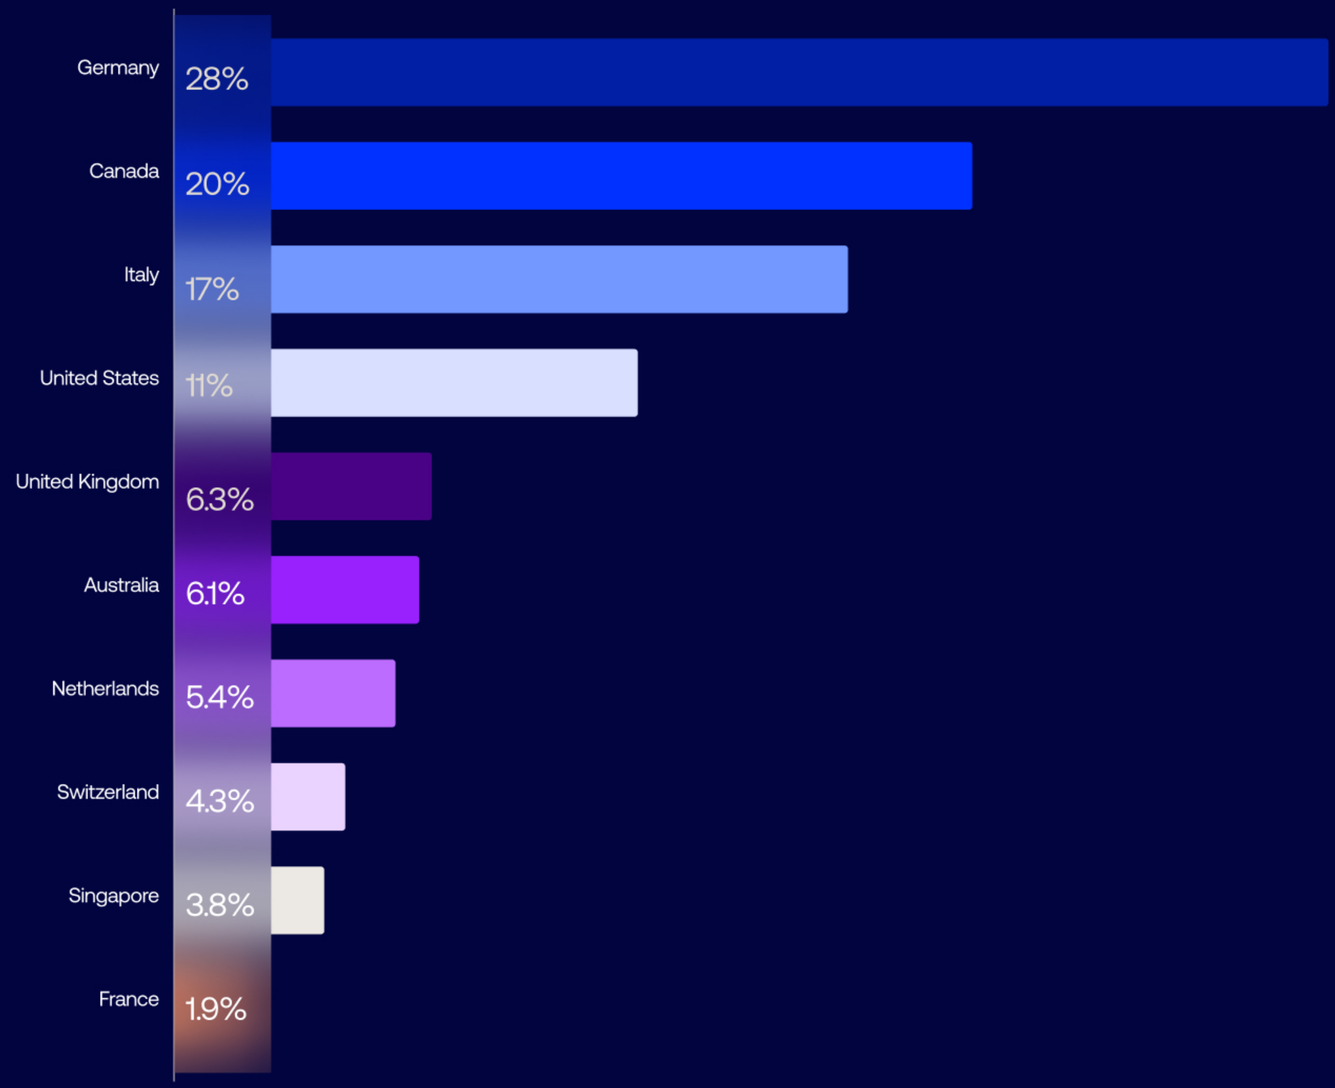

Ratio of detected threats to authentications

Based on 1 year of cumulative data

Percentages over 10% have been rounded

The global escalation of identity attacks

Geographically, the rise in threats per authentication pressure is most intense in Germany, where the threat ratio jumped to 28% (up from 19% last year). Canada moves from the No. 5 spot to No. 2, with detected threats now accounting for one in five authentication attempts. Italy joins the top ranks with a 17% ratio, while the United States, which led all countries last year at 6.6%, rises to 11% this year to claim the No. 4 spot. The 2026 Unit 42 Global Incident Response Report confirms that identity weaknesses play a material role in nearly 90% of all threat investigations, as attackers increasingly bypass traditional defenses by simply logging in with compromised credentials. This data highlights that the authentication foundation must be strong enough to withstand a global environment where nearly a third of all login attempts in some regions are malicious. As agentic AI comes online, the consequences of compromised credentials expand beyond individual accounts, enabling automated access across enterprise systems.

Identity maturity model focus: Authentication maturity is still scaling

Current maturity signal: The data suggests most organizations remain in the Scaling stage (Stage 2) of identity maturity for authentication, where MFA adoption is widespread but still relies heavily on phishable authentication factors such as SMS and OTP. As credential-based attacks continue to scale, this gap between adoption and assurance leaves many authentication flows vulnerable.

Capabilities required to advance: Moving to Advanced maturity (Stage 3) requires deploying phishing-resistant authentication, including passkeys and hardware-backed authenticators. Organizations also need automation around policies, privileges, and remediation, all in conjunction with risk-based authentication and with enough granularity to support attribute-based decisions.

Outcomes unlocked: Phishing-resistant authentication reduces credential-based attacks and account takeovers while enabling a seamless passwordless experience for users. It also creates the high-assurance identity foundation required to securely deploy AI agents and automated workflows. However, organizations in this “stepping-stone” stage typically lack higher-grade real-time remediation as well as the posture management and practices they need to combat today’s most common threats.

Chapter 4

Securing onboarding and digital experience (LCM)

The era of manual provisioning is fading, as the human identity lifecycle becomes increasingly automated across business, collaboration, and HR applications

This shift is no longer just an efficiency play for HR departments; it has become a foundational requirement for securing the modern enterprise, and is laying the groundwork for a wholesale enterprise pivot to agentic AI.

However, this rapid influx of AI agents could create a governance vacuum. The established, HR-driven model developed for human identities is breaking; there is no "Workday for agents" to serve as the authoritative source of truth. Identity management must evolve into a comprehensive lifecycle engine capable of governing both human and non-human identities across systems and diverse deployment platforms.

Unless their automated lifecycle management can enforce least privilege at every single joiner, mover, and leaver moment, organizations risk building a weaponized, overprivileged army of nonhuman agents operating at unmanageable speeds.

Identity management must evolve into a strong and safe automated lifecycle engine for **human and non-human identities.**

Apps with provisioning used, by app category

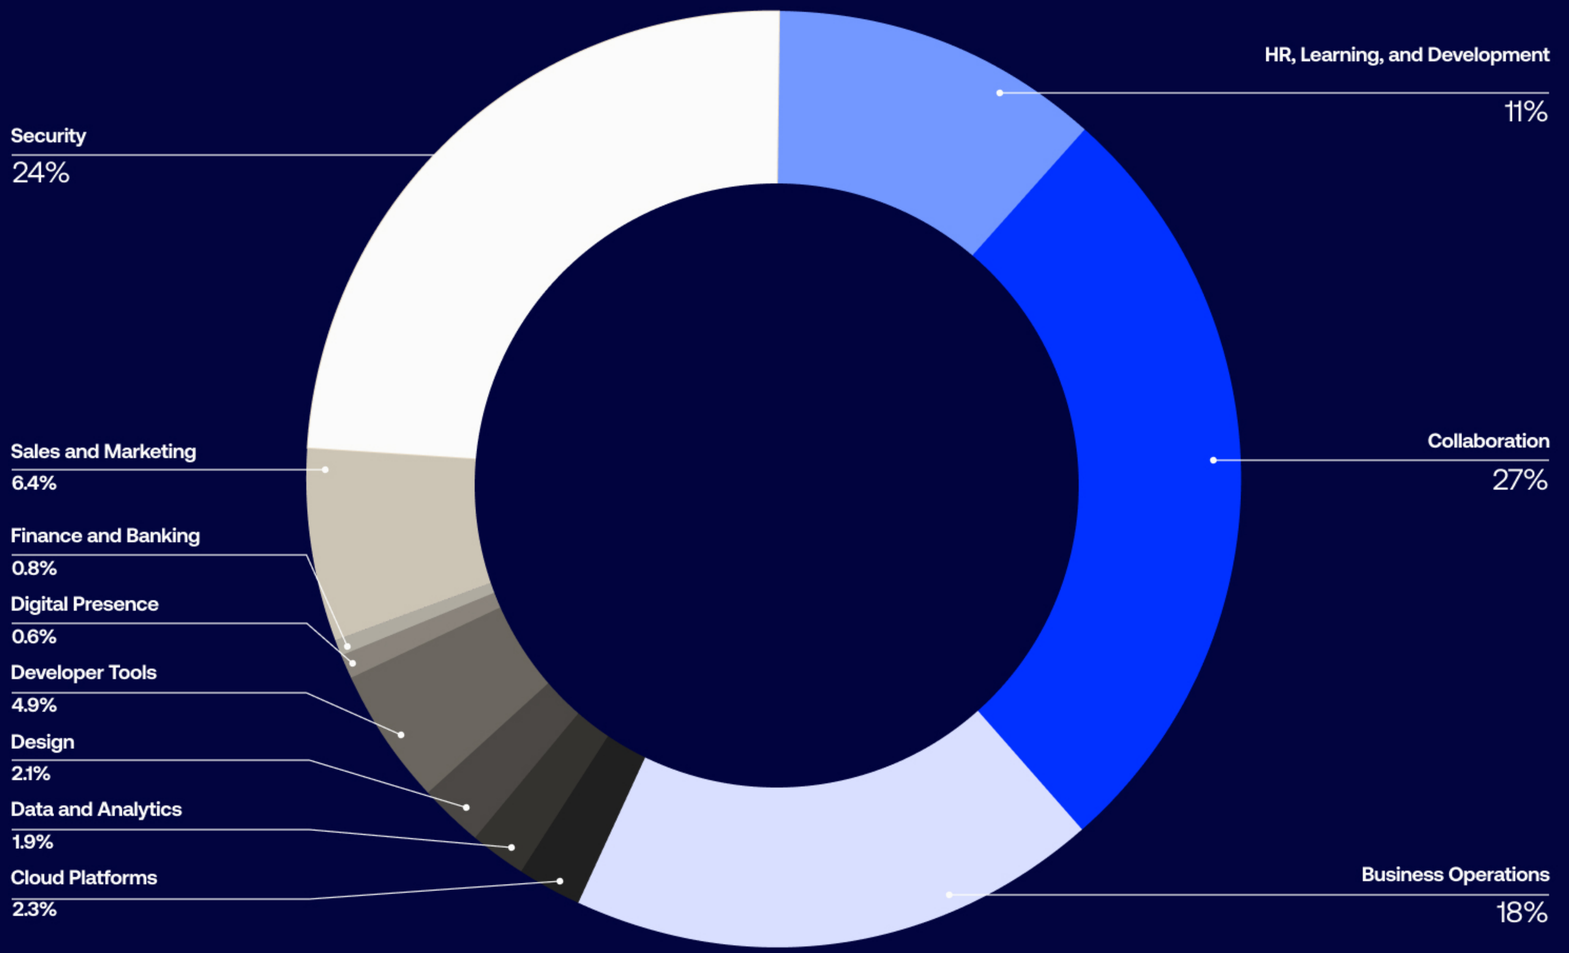

Percentages over 10% have been rounded

The new architecture of onboarding

Currently, 56% of all automated provisioning activity occurs within core business pillars: HR, collaboration, and business operations. These are the traditional starting points for a human's digital journey, but they are increasingly becoming the birthplaces of digital co-workers, too. And according to Forrester's 2026 predictions, the top five HCM platforms will offer digital employee management capabilities. HR tech will play a major role in the integration of digital workforce management. This evolution is critical, now that identities can be spun up by the thousands in seconds. If onboarding remains tethered to slow, manual processes, these agents will operate entirely outside the view of security teams.

The most telling signal in our data is that security-focused tools now account for nearly a quarter (24%) of all automated provisioning. This indicates that organizations are no longer treating access and protection as separate phases of identity governance. Instead, they are automating security tool provisioning (such as network and device protection) from Day One, by taking a "secure-by-design" approach that embeds security directly into the automated lifecycle.

Industry momentum in lifecycle maturity

The push toward automated identity readiness is not uniform across the global economy. Currently, the insurance and transportation sectors are setting the pace for year-over-year growth when we look at the number of companies investing in identity lifecycle automation. For these industries, the move toward agentic workflows is a matter of survival. As agents take over complex tasks like real-time inventory rerouting and claims processing, the underlying identity engine must be robust enough to handle the sheer volume of "joiner, mover, leaver" events that occur without human intervention.

When we look at automated lifecycle events per customer, the outperforming industries might not be the ones you might expect. The arts, entertainment, and recreation sector, along with manufacturing, show the highest average automated lifecycle events per customer. Whether an organization is in a high-growth phase like insurance or a high-volume phase like arts, entertainment, and recreation, the imperative is the same: linking automation with governance to enforce least privilege even at scale.

Identity maturity model focus: Extending lifecycle management to the agentic world

Current maturity signal: Most organizations are in the Scaling stage of lifecycle (Stage 2) and access automation maturity, with automated onboarding, role changes, and offboarding for human identities connected to core business systems. Organizations in this stage are typically focused on establishing “golden sources of truth” in the identity landscape, and may lean more heavily on automations and/or federation.

Capabilities required to advance: Progressing to Advanced maturity (Stage 3) requires a gradual but purposeful move to more automation, more federation, and better centralized “command-center” style administration.

Outcomes unlocked: A unified lifecycle foundation that governs both human and nonhuman identities, reducing unmanaged access and enabling organizations to operate automated agents safely at scale. Ultimately, this leads toward increased employee productivity and reduced administrative costs.

Chapter 5

Security-driven governance

Across industries and organizations, as agentic AI moves from pilot to production, traditional governance is hitting a breaking point

When autonomous agents can make independent decisions and call APIs between periodic human reviews, visibility gaps become systemic risks. Our data shows governance activity surging: access requests per company have more than doubled in the past year (up 1140% over two years), while access certifications continue to climb. At this scale, any process still reliant on a human reviewer becomes a critical security bottleneck. In this high-velocity environment, identity governance must evolve from a repetitive administrative process into a continuous, machine-speed orchestration layer. Without this shift, the speed of autonomy that makes agents powerful will quickly outpace the governance systems designed to control them.

Autonomous agents are **moving faster than traditional governance can keep up.**

Identity governance

Note: High growth in Access Request volume reflects the 2023 expansion of the OIG platform. Usage scaled as the customer base grew and organizations fully integrated the feature into their workflows.

The governance surge

The speed of modern business is increasingly dictated by automated workflows, and identity governance activity is rising just as quickly. Across the Okta Integration Network, the average number of access requests per company grew 158% YoY and 1140% over two years. At the same time, access certifications increased 76% YoY and 810% over two years, reflecting a sharp rise in the number of access decisions organizations must review and validate.

This surge reflects the growing pace and complexity of modern digital environments. As organizations automate more business processes and connect more applications, systems, infrastructure, and services, the number of access changes across enterprise systems increases accordingly. Each of those changes – whether an access request, approval, or certification – creates a governance event that must be tracked and reviewed.

The broader identity landscape is also expanding rapidly. Non-human identities are widely estimated to outnumber human identities by as much as 45:1 in modern cloud environments, dramatically increasing the number of access relationships organizations must manage. At the same time, automated systems and machine-driven workflows are increasing the speed at which changes to those access relationships occur, compounding the governance pressure organizations must manage.

Governance pressure is not uniform across industries. Manufacturing leads the charge, with an 86% YoY increase in access certifications, while finance & banking and technology follow with 67% and 56% growth, respectively. These patterns suggest governance activity is rising fastest in sectors where digital operations and system integration are most advanced.

As governance activity accelerates, organizations are discovering that manual review processes cannot keep pace. What was once an administrative control is becoming an operational bottleneck. Increasingly, organizations are turning to workflow-driven automation to manage access reviews, approvals, and certifications at the scale required by modern digital environments.

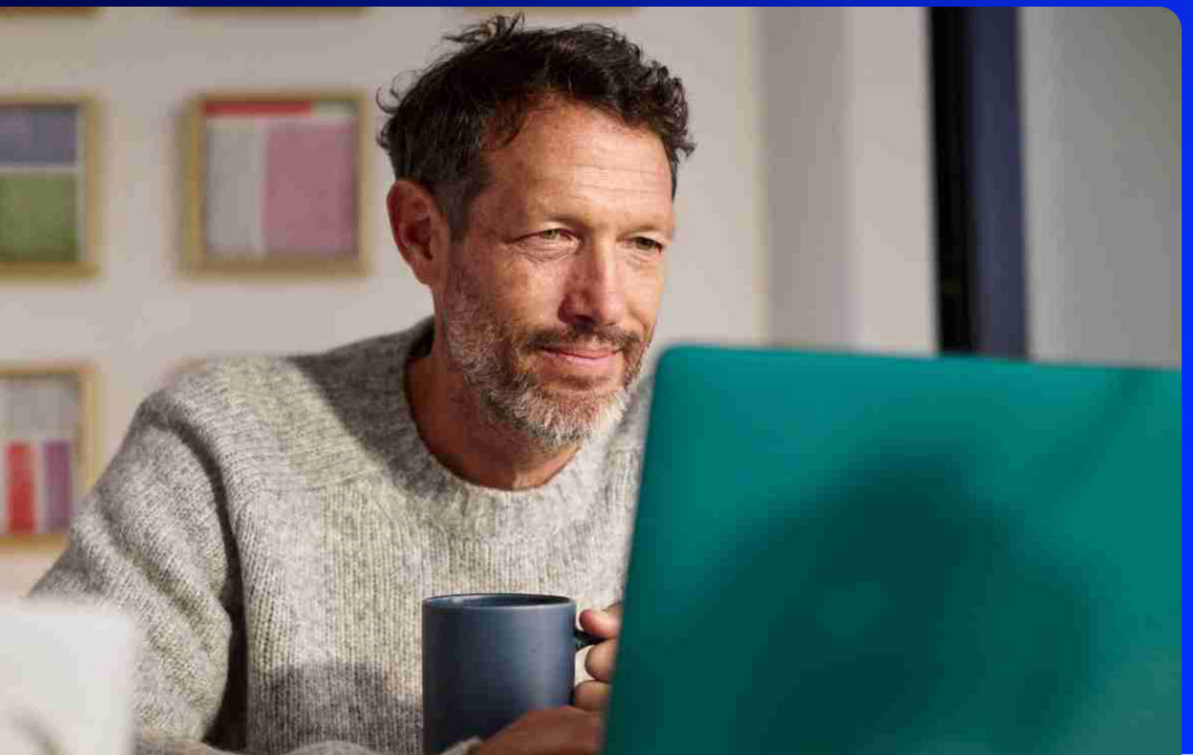

Access requests per company surged **1140% in the past two years.**

Industries where organizations use three or more automated governance workflows

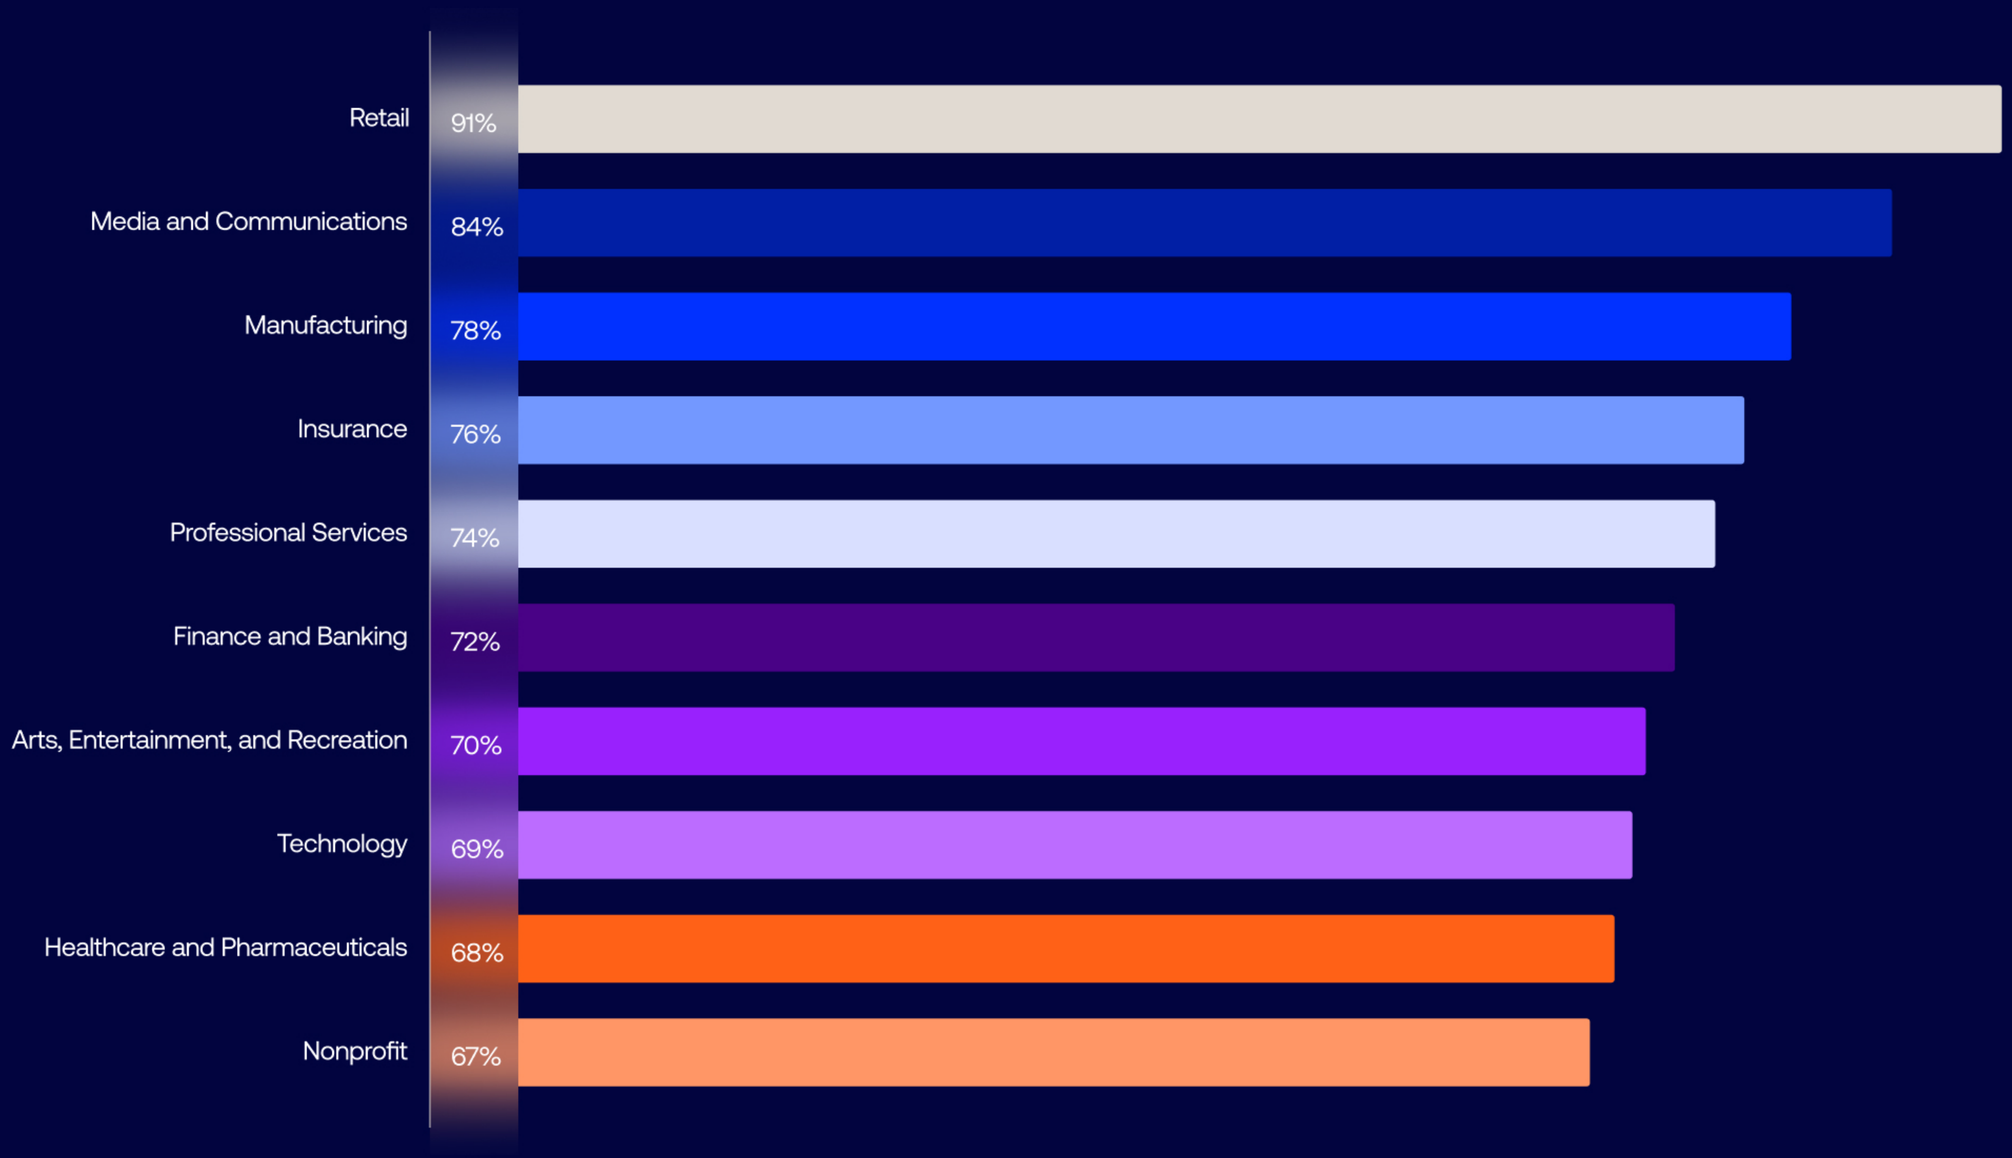

Percent of Okta's OIG customers with three or more active workflows

Security-driven governance

Automating identity governance at scale

As governance activity scales, organizations are turning to workflow-driven automation as an operational requirement rather than an administrative convenience. Retail currently leads this shift, with nine out of ten organizations utilizing three or more automated governance workflows. This trend is global: 87% of organizations in the Netherlands have adopted these automated workflows, followed by Germany at 79% and Canada at 78%. These adoption rates reflect how quickly enterprises are embedding automation into governance processes to keep pace with the scale and speed of modern identity activity.

The transition to security-driven governance represents a critical step in scaling identity operations. By embedding automation directly into the identity lifecycle, organizations can move from periodic, manual reviews to continuous, event-driven oversight. This shift allows governance processes to operate at the same speed as modern digital environments while maintaining strong security controls.

Identity maturity model focus: Continuous governance at machine speed

Current maturity signal: Many organizations have reached Advanced identity maturity for governance (Stage 3) by automating access reviews and approvals through governance workflows. These capabilities improve visibility and help enterprises manage the growing volume of access decisions across modern digital environments.

Capabilities required to advance: Reaching Strategic maturity (Stage 4) requires enabling continuous governance across the identity lifecycle, where identity changes automatically trigger policy enforcement, approvals, and oversight. As organizations adopt AI agents, these same governance controls must extend beyond human identities to include AI agents and other non-human identities (e.g., service accounts).

Outcomes unlocked: Organizations strengthen security posture and operational agility by enforcing consistent governance and least-privilege access across both human identities and AI agents. A primary goalpost: enforcing authorization and least-privilege governance.

Chapter 6

Securing on-prem and hybrid IT environments

Most enterprises still operate complex hybrid environments where legacy systems and on-prem applications sit outside modern identity controls

As organizations adopt agentic AI, this fragmented identity coverage becomes increasingly risky. Our data shows that while organizations are extending modern access controls across more systems, inconsistent identity coverage continues to create security gaps.

In the age of autonomous agents, these gaps become especially dangerous: any application outside the unified identity security fabric can act as an ungoverned pathway to sensitive assets. Achieving true agentic AI readiness requires extending consistent identity policy and authentication across both cloud and on-prem environments.

Hybrid and on-prem authentications

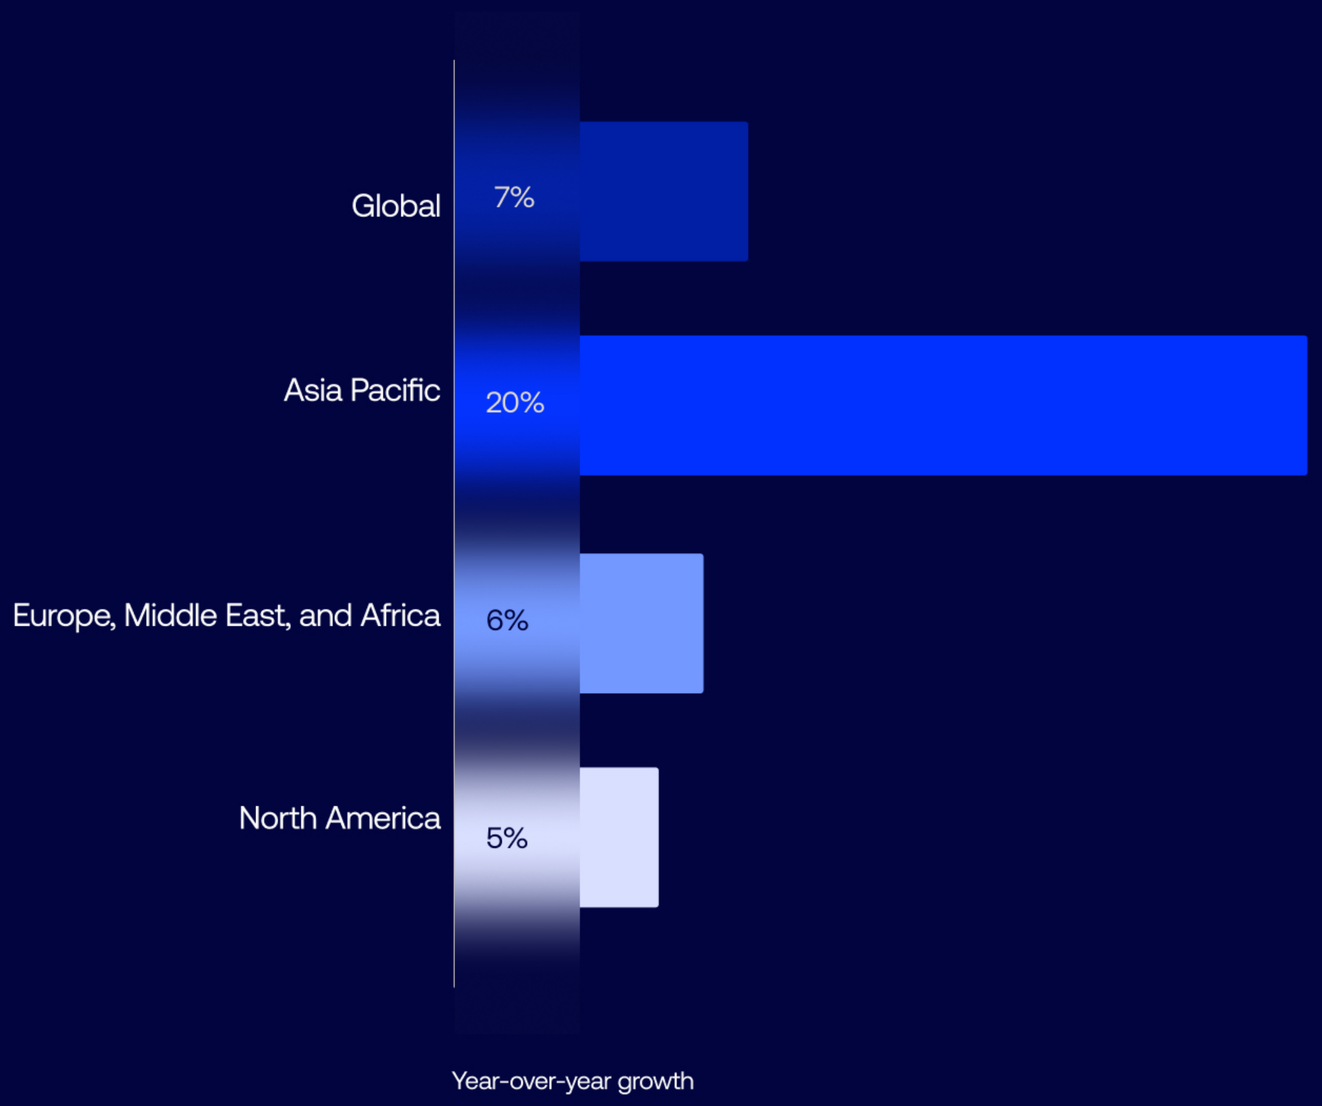

Hybrid and on-prem authentications in APAC grew 20% YoY.

The expanding hybrid frontier

The "cloud-only" narrative isn't reality for many modern enterprises. As identity programs mature, organizations are increasingly pulling legacy and on-prem systems inside their security perimeters to eliminate the blind spots that attackers—and now autonomous agents—can exploit. The shift to autonomous decision engines makes identity the ultimate control plane, ensuring that permissions for agents moving between on-prem databases and cloud-native apps are governed by a single, consistent source of truth.

Globally, the number of hybrid and on-prem authentications is up 7% YoY, reflecting this accelerated digital transformation. The APAC region is leading the charge, with a 20% YoY increase in hybrid authentications. IDC projects that AI-related spending in the Asia-Pacific region will grow 1.7x faster than overall digital technology investments over the next three years. Conversely, highly-regulated regions like EMEA (6%) and North America (5%) are moving more cautiously, likely hampered by the weight of compliance mandates and the long tail of legacy debt.

Stolen credentials remain the top entry point for 22% of all breaches, according to Verizon's 2025 Data Breach Investigations Report. In hybrid environments, inconsistent identity coverage creates new attack paths between legacy systems and cloud applications. Findings from the 2025 IBM Cost of a Data Breach Report also show that breaches spanning multiple environments are among the most expensive and hardest to contain, averaging $5.05 million per incident and taking an average of 276 days to identify and fully contain.

Identity maturity model focus: Getting to universal authentication in hybrid environments

Current maturity signal: Most organizations remain in the Scaling maturity stage (Stage 2) for authentication, where modern authentication is deployed for many cloud applications but legacy and on-prem systems still operate outside consistent modern identity controls. This fragmented authentication landscape creates security gaps and uneven protection across the application estate, and every landscape with a large dependency on on-prem infrastructure will be inherently more vulnerable to rogue AI agents.

Capabilities required to advance: Reaching Advanced maturity (Stage 3) requires extending modern authentication standards and consistent identity policies across cloud, on-prem, and legacy applications so all access requests are verified through a unified identity control plane. Move use cases to the cloud when possible, consolidate security administration, use on-prem adapters to bridge the gap, and keep on-prem AI tooling focused on minimizing blind spots.

Outcomes unlocked: Organizations strengthen security posture and deliver more consistent access experiences while ensuring high-assurance authentication for users and for AI agents acting on their behalf across hybrid environments.

Chapter 7

Protecting non-human identities

The era of agentic AI is rapidly expanding the role of non-human identities (NHIs) across enterprise environments

Today, these identities are most commonly represented by service accounts that enable automated systems, applications, and workflows to interact with infrastructure and data without human intervention. As organizations move from simple automation toward agentic systems capable of executing tasks across cloud and on-prem environments, these machine identities become the operational layer through which autonomous systems access enterprise resources.

90% of organizations lack a comprehensive strategy to govern non-human identities.

Our data signals a critical governance gap: while 74% of organizations report widespread or moderate use of NHIs, a staggering 90% still lack a comprehensive strategy to govern them. This gap represents more than just an agentic AI readiness challenge; it's an alarming security risk. Autonomous agents can already operate with administrative privileges that often exceed those of their human creators. Without a unified governance model for service accounts today, organizations will struggle to control the agentic workforce of tomorrow.

Adoption of non-human identities

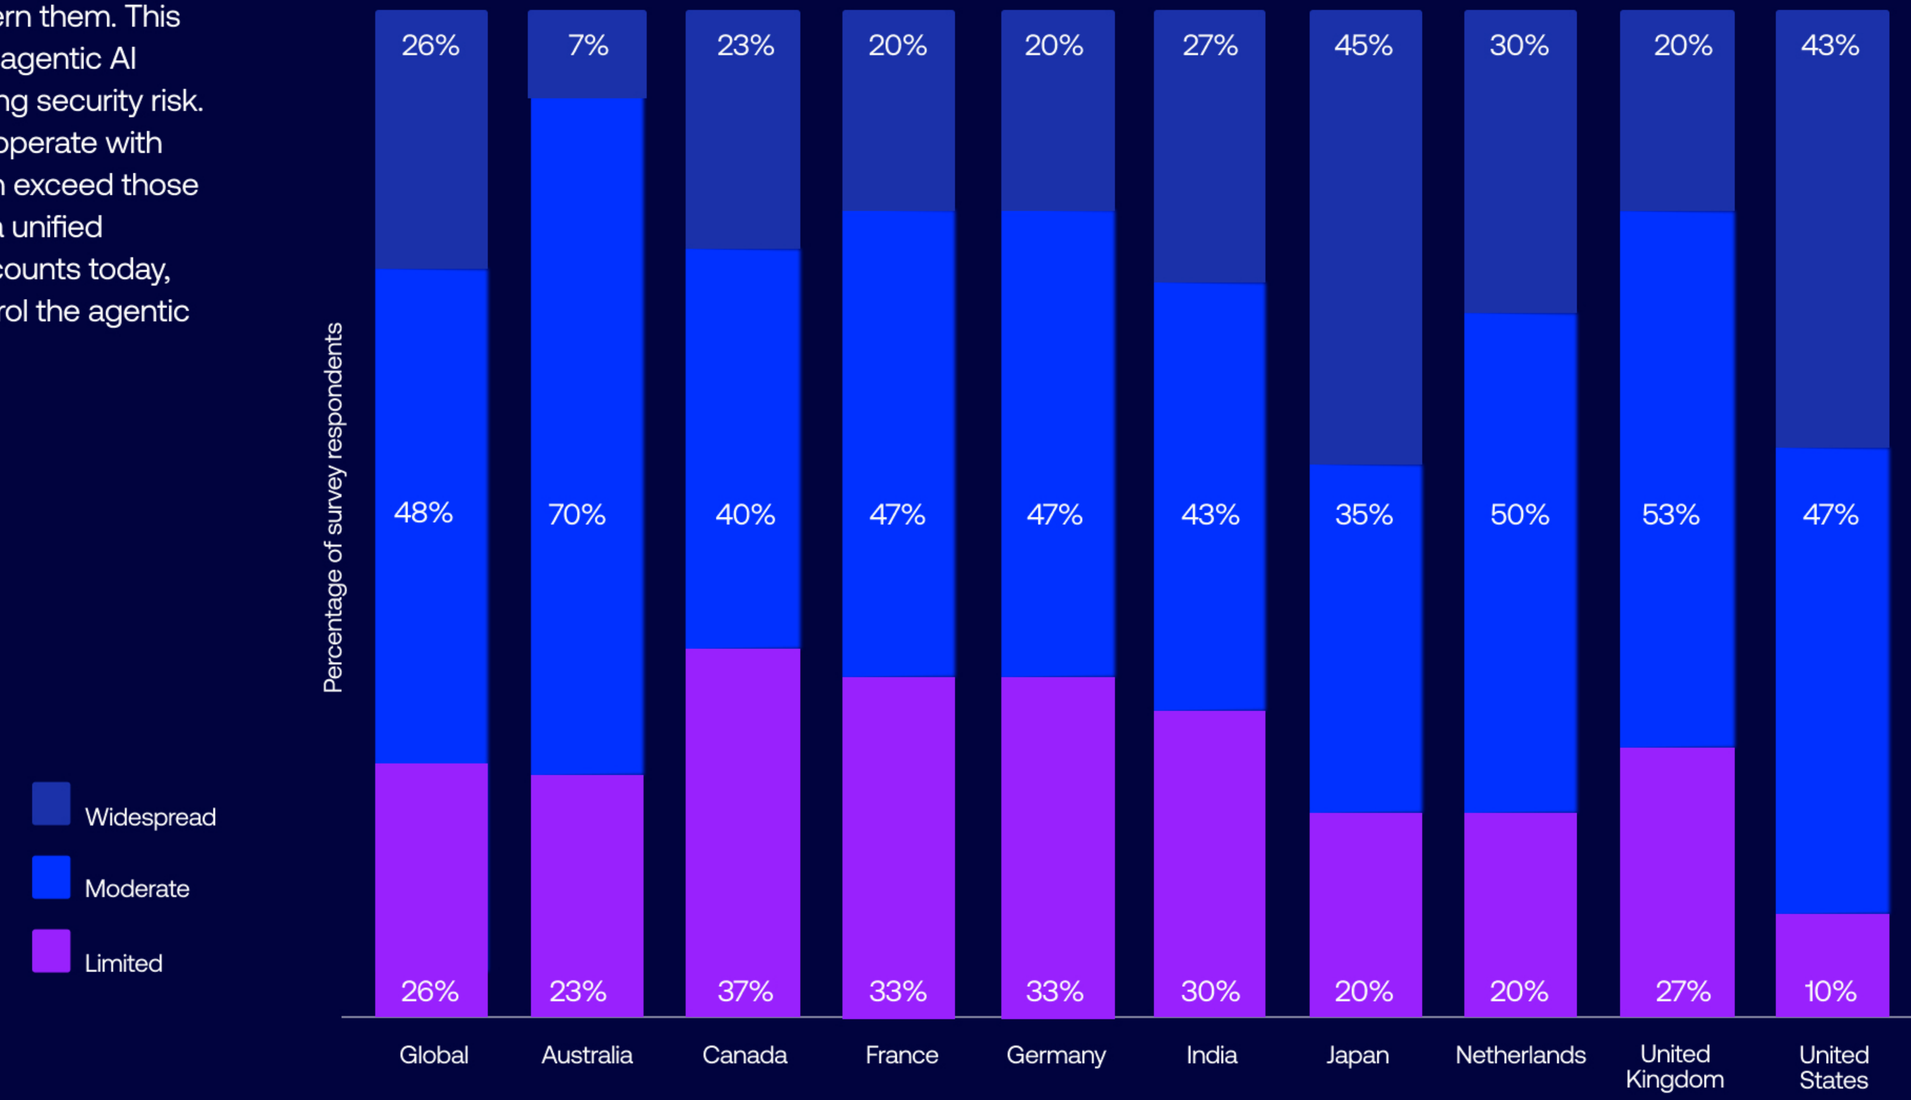

The non-human workforce is already here

Non-human identities are no longer a niche tool for DevOps teams; they have become the foundational layer for many modern business functions. AI agents build on this same foundation, acting through the service accounts and machine credentials that allow software to access applications, data, and infrastructure. As a result, agents will act as a multiplier, accelerating both growth and the criticality of machine identities. Globally, adoption of NHIs is already robust, with 74% of organizations we surveyed reporting at least moderate use of NHIs. Japan and the United States lead the charge in "widespread" adoption, at 45% and 43% respectively, reflecting deeper maturity in automation and the machine identity infrastructure that AI agents will rely on. Countries like Canada, France, and Germany show more muted adoption, with a third or more of companies within each country reporting only limited adoption.

What exactly qualifies as an NHI?

As adoption of NHIs grows, there can be some definitional creep as well. We recently asked C-suite respondents to tell us what constitutes an NHI from their perspective. Respondents pointed to a wide range of identities that operate without direct human interaction across enterprise systems. These include traditional machine credentials as well as newer forms of automated and agent-driven identities used across applications and infrastructure.

In ranked order:

1. Digital credentials
2. API keys and tokens
3. Bots and agents
4. Machine-to-machine interactions
5. Digital identities assigned to software apps & services
6. IoT and edge devices
7. Automated workflows
8. AI models and algorithms with decisionmaking capabilities
9. Digital personas and avatars
10. Securely stored sensitive information, such as passwords or encryption keys

Securing the "keys to the kingdom"

As NHIs proliferate rapidly, organizations are scrambling to intelligently prioritize their security efforts. According to our data, API keys and tokens have emerged as the most critical non-human identities to secure, followed closely by service accounts. These priorities align with the operational "points of impact" where agentic AI is most likely to operate. Service accounts are particularly vulnerable, because they often hold shared identities used for long-term integrations—the exact type of persistent access an AI agent requires to, say, monitor cloud infrastructure or execute system-level tasks autonomously.

When NHI creation and deployment outpaces an enterprise's ability to put appropriate guardrails in place, organizations can quickly face a governance deficit. AI agents often inherit the permissions of their human creators, turning any instance of "excess privilege" into instantaneous, automated exposure. To combat this, organizations are increasingly turning to centralized authorization control systems like Privileged Access Management (PAM). Our data reflects this shift, with a massive 650% YoY growth in the number of service accounts being centrally managed by enterprises.

Our data shows a **650% YoY growth in the number of service accounts being centrally managed** by enterprises.

Most pressing security concerns related to NHIs

Controlling NHI access and permissions
Challenges in effectively managing the access rights and permissions granted to non-human identities, which may lead to unauthorized actions or data exposure.

78%

Governing the NHI lifecycle
Issues in overseeing the entire lifecycle of non-human identities, including creation, maintenance, and deactivation, leading to dormant or orphaned accounts that present security vulnerabilities.

69%

Visibility into NHI sprawl
Difficulty in tracking and managing the proliferation of non-human identities across various systems.

57%

Remediating risky NHI accounts
Difficulty in identifying and addressing non-human identities with risky behavior or improper configurations that could pose security threats to the organization.

53%

Other*

2%

Percentage of survey respondents

*Other concerns include regulatory compliance, emergency scenarios, and no simple MFA equivalent

The friction of machine-speed security

The move toward Agentic AI is intensifying security challenges around non-human identities, with 78% of the organizations we surveyed citing the control of access and permissions as their top NHI challenge. As NHIs proliferate at machine speed, driven increasingly by automation and AI agents, maintaining the principle of least privilege becomes significantly harder. Organizations are also grappling with the lifecycle of these identities; 69% report concern about orphaned or dormant accounts that remain active long after their associated project or agent has been retired, potentially providing backdoors for attackers.

These concerns reflect broader trends in the modern threat landscape. Verizon's 2025 Data Breach Investigations Report found that while the human element remains a factor, third-party involvement has doubled, from 15% to 30%. The median time to remediate leaked secrets in public repositories was found to be 94 days—an unthinkably long window in a time when an AI-driven attacker can exploit a vulnerability in seconds. For C-suite leaders, the message is clear: visibility and governance are no longer nice-to-haves for machine identities. They are the primary defense against the scale and speed of agentic threats.

Centrally managed service accounts per organization

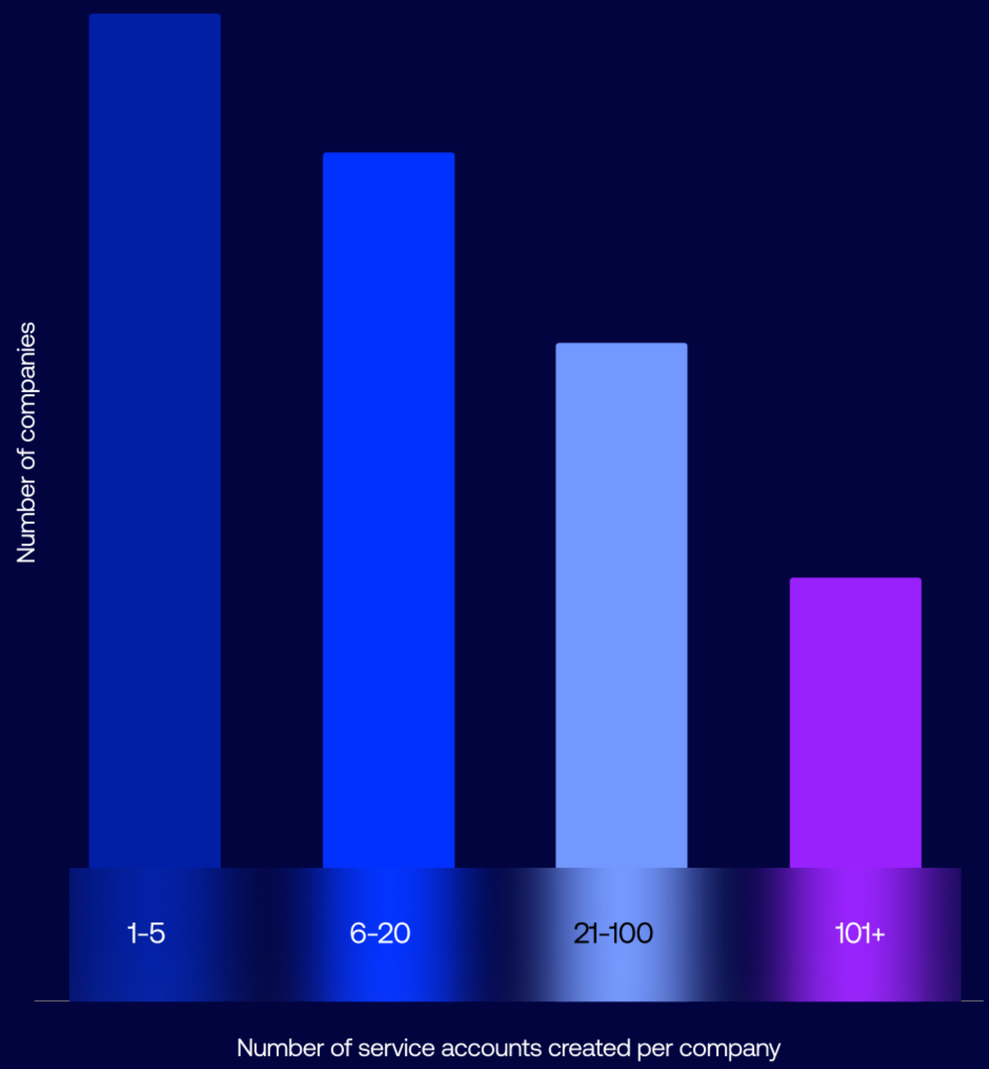

Mapping the NHI frontier by industry

A look at the current distribution of service accounts by company (OIN data) shows that most organizations currently centrally manage relatively small numbers of service accounts, with 1-5 accounts being the most commonly cited range. Because NHIs still operate outside centralized governance and telemetry, these figures reflect only a portion of the total service account footprint, highlighting how early many organizations are in centralizing governance of non-human identities. A meaningful subset of organizations already report managing 20 or more centrally governed accounts, an early indicator of growing non-human identity complexity as enterprises continue to add machine identities. However, the complexity of the non-human identity landscape varies significantly by industry and organizational maturity. By industry, media and communications firms manage the highest number of centrally governed service accounts (averaging 78 per company), followed by technology (28 per company). These high-volume, data-intensive sectors are among the earliest adopters of agentic AI for content orchestration and software development.

78% of organizations say controlling access and permissions is their top NHI challenge.

Even in traditional sectors like manufacturing and finance, we are seeing rapid growth. Manufacturing, in particular, saw a 494% YoY increase in centrally managed service accounts, while finance and banking grew by 353% YoY. This surge suggests that these industries are "cleaning house," centralizing control of their machine identities in preparation for more autonomous operations. As AI agents become the new workforce, the ability to see and control these accounts across sectors will determine which organizations can safely innovate and which will fall victim to the machine-speed risks of the new frontier.

Manufacturing and finance are rapidly centralizing control, with **managed service accounts growing 494% and 353% YoY, respectively.**

Identity maturity model focus: Establishing governance of service accounts as the foundation for managing non-human identities

Current maturity signal: Most organizations remain between the Foundational (Stage 1) and Scaling (Stage 2) stages of identity maturity for non-human identities (NHIs). Service accounts are widely used across enterprise systems, but visibility, lifecycle management, and centralized governance remain inconsistent.

Capabilities required to advance: Reaching Advanced maturity (Stage 3) requires extending identity governance to service accounts through centralized visibility, lifecycle controls, and consistent policy enforcement across environments so machine identities can be managed with the same rigor as human identities. Service account rotations, vaulting, and similar elevated governance mechanisms are critical to close security gaps here.

Outcomes unlocked: This will help companies strengthen security and compliance, improve operational agility, and deliver more reliable end-user experiences while establishing the governance foundation needed to securely support AI agents operating through machine identities across modern environments.

Chapter 8

Conclusion

For over a decade, Okta's Businesses at Work report has tracked how enterprises adopt and secure modern applications across the Okta Integration Network

As organizations build increasingly complex digital ecosystems, identity remains central to enabling secure and seamless access for users while protecting against persistent threats.

This year, however, identity complexity has taken a significant step forward, with the emergence of agentic AI. As the definition of a user expands to include autonomous agents acting on behalf of humans, applications, and other systems at machine speed, identity governance becomes the critical control plane for securing this new digital workforce.

Widespread use of nonhuman identities to carry out authorized digital tasks presents both enormous opportunity and significant risk for enterprises. Before they can safely scale automated systems and AI agents, they must first reassess and strengthen the identity foundations they already rely on: authentication, governance, and nonhuman identity management across hybrid environments.

Conclusion

The most forward-looking organizations are embedding security into their processes from day one and automating governance to operate at machine speed. But how do you translate these principles into practice?

Based on continuous work with leading organizations, Okta has developed the blueprint for the secure agentic enterprise. These frameworks and tools are designed to help companies manage AI risk, providing a foundation for IT and security leaders to assess their readiness and answer the three critical questions that will define the next era of security:

Where are my agents?

What can they connect to?

What can they do?

Deploying AI agents safely and at scale isn't merely about understanding and adopting a new technology. It requires reckoning with any identity security debt you've already accumulated—identity isn't the entire solution, but is a critical part of it. This is a shared responsibility between technology providers, customers, and the developers of AI agents. Any gaps across authentication, governance, and non-human identity management across hybrid environments are no longer just best practice gaps. They are blockers to safely deploying the AI agents and nonhuman identities that are sure to become a key part of remaining competitive as business evolves in the months and years to come.

The conclusion is unavoidable: Agentic AI readiness *is* identity readiness. **Are you ready?**

Take our AI Readiness Assessment to find out.

This report analyzes aggregated, anonymized data from across the Okta platform. This dataset reflects identity activity across thousands of organizations, applications, and IT infrastructure integrations worldwide. The analysis draws on signals related to authentication, application access, identity governance, and non-human identities to understand how enterprises manage and secure modern digital ecosystems. Because this report reflects identity interactions occurring within customer environments, the findings highlight trends across organizations that rely on Okta to manage workforce access to systems, applications, and infrastructure.

Unless otherwise noted, all data discussed in the Okta Businesses at Work 2026 Report is derived by analyzing data from activity observed on the Okta platform, and the applications and services connected to it, between November 1, 2024 and October 31, 2025. References to “this year” reflect this reporting period, while comparisons to prior years use equivalent time windows. Additionally, some insights in this report incorporate findings from Okta’s AI at Work 2025 survey of 261 C-suite and VP-level respondents. This research examines enterprise adoption of artificial intelligence and autonomous systems and helps contextualize how emerging technologies such as AI agents and automated workflows are influencing identity usage and security requirements across organizations.

To interpret these findings, this report references two analytical frameworks. The identity maturity model (IMM) describes how organizations evolve their identity capabilities over time, moving through four stages, from Foundational (consolidate and simplify), through Scaling (layer security controls) and Advanced (automate and elevate experience) to Strategic (optimize and extend). The report also references the identity security fabric (ISF), an identity architecture concept that outlines how identity capabilities work together to secure modern enterprise environments.

Unless otherwise specified, the data included in the Okta Businesses at Work 2026 report is limited to Okta customers who have deployed at least one app to users through the Okta Integration Network, looking only at apps deployed for corporate use. Throughout the report, we use the terms “app” and “tool” to refer to applications, services, and integrations available through the Okta Integration Network.

- The trends we describe for Okta’s Microsoft 365 customers may differ from those of Microsoft 365 customers who don’t use Okta (i.e., those using Azure Active Directory or other identity platforms that don’t provide strong cross-app integration support.)
- FastPass data is based on one year of cumulative data from Okta Workforce Identity customers authenticating through Okta FastPass.
- MFA usage data comes from all production Okta Workforce Identity organizations.
- ThreatInsights data is based on one year of cumulative system logs, which show important activities such as failed login attempts or traffic spikes.

Safe Harbor Statement

This report contains "forward-looking statements." These statements are not historical facts but rather are based on Okta's current expectations, estimates, and projections about our industry, business, and the technology landscape. Forward-looking statements can often be identified by words such as "believes," "expects," "anticipates," "estimates," "intends," "plans," "seeks," "will," "may," "should," "predicts," "projects," "targets," "will likely result," or the negative of these terms or other comparable terminology.

These statements include, but are not limited to, predictions and observations regarding:

• The future adoption, growth, and impact of agentic AI and other artificial intelligence technologies.

• Trends in application usage, security practices, and identity management.

• The future evolution of the threat landscape and cybersecurity responses.

• Projections about customer behavior and technology investment priorities.

• The expected performance and integration of various technologies and platforms.

These forward-looking statements are subject to a number of risks, uncertainties, and assumptions that could cause actual results to differ materially from those described in the forward-looking statements. These risks and uncertainties include, but are not limited to:

• The pace of technological change and market acceptance of new and developing technologies, including AI.

• Changes in the competitive landscape and the strategies of our competitors.

• Uncertainty in the global economic environment.

• The evolution of and changes to the regulatory and legal landscape, particularly concerning data privacy and artificial intelligence.

• Our ability to accurately interpret and analyze the underlying data from the Okta Integration Network and other third-party sources.

Okta undertakes no obligation to update any forward-looking statements in this report, whether as a result of new information, future events, or otherwise. In light of these risks and uncertainties, we caution you not to place undue reliance on these forward-looking statements.

Third-Party Trademark Notice

The "2026 Business at Work" report mentions various third-party companies, products, services, and applications for identification and illustrative purposes only.

All third-party trademarks, service marks, trade names, product names, and logos appearing in this report are the property of their respective owners. The use of any third-party trademarks does not imply an affiliation with, endorsement by, or sponsorship of Okta by such third parties.

Okta does not claim any ownership rights in these third-party marks. Any reference to third-party products, services, or companies is for informational and analytical purposes only and should not be construed as a recommendation or endorsement.

How to cite the Okta Businesses at Work 2026 report

We love it when people share Businesses at Work insights. Here's how to properly cite data, statistics, and any other information found in the Okta Businesses at Work 2026 report:

Give us credit: Please cite the source as "Okta Businesses at Work 2026 report" when referencing any content.

No modifications: Content must be cited exactly as it appears in the report. If you wish to paraphrase, we'd appreciate it if you contact us for approval.

Please share: If you'd like to share the report with others, please provide a link to our download page: www.okta.com/businesses-at-work

We appreciate your helping us keep our insights accurate and accessible to everyone.

About Okta

Okta, Inc. is The World's Identity Company™. We secure identity, so everyone is free to safely use any technology. Our customer and workforce solutions empower businesses and developers to use the power of Identity to drive security, efficiencies, and success—all while protecting their users, employees, and partners. Learn why the world's leading brands trust Okta for authentication, authorization, and more at okta.com.