Prepared for

ORY

The AI Identity Crisis:

Balancing Innovation with Strict Compliance in the Financial Sector

May 2026 EMA White Paper
By **Ken Buckler, CASP;** Research Director
Information Security, Risk, and Compliance Management

EMA™
IT & DATA MANAGEMENT RESEARCH,
INDUSTRY ANALYSIS & CONSULTING

Executive Summary: Solving the AI Identity Crisis in Finance

In today's high-stakes landscape of AI readiness, financial institutions are navigating a profound strategic tension: the drive to harness the transformative efficiency of AI agents against the non-negotiable mandates of global regulation. While the sector has transitioned from experimental AI to disciplined adoption, a critical identity readiness gap remains. Financial organizations are deploying AI with a culture of prudence, lagging the broader market in full-scale production (29.7% vs. 40.6%) while security teams take an outsized ownership role in oversight.

Introduction and Background

Recently, EMA and Ory conducted a comprehensive global survey titled "**Agentic AI Identities: The Unsecured Frontier of Autonomous Operations.**" The research synthesized data from 1,256 initial participants, resulting in 271 qualified responses from key stakeholders. Representing a cross-section of 20+ industries and five key global markets (including the U.S., UK, Canada, Germany, and France), the study provides a statistically diverse perspective across small, medium, and largescale organizations.

Key Findings from Agentic Survey

- **83% of large and 70% of medium-sized organizations** have already deployed AI agents in production
- **98%** of organizations plan to use AI agents
- **79% of those deployed have no documented policies** for controlling agents
- **66% plan to have "human oversight"** or human in the loop

In addition to the hurdles of deploying AI agents, organizations are finding it difficult to implement the identity and access management (IAM) controls needed to secure them. With 58% of organizations having three or more IAM solutions and 51% citing rising costs as their biggest challenge with IAM deployments, organizations are looking to cut costs while improving functionality. Approximately 68% are looking to bring their IAM deployment in-house instead of relying on SaaS. When it comes to resiliency, compliance, scale, and security, the challenges across all industries are remarkably clear.

Is your existing IAM solution stack robust and ready to solve the challenges coming with handling internal and external agents?

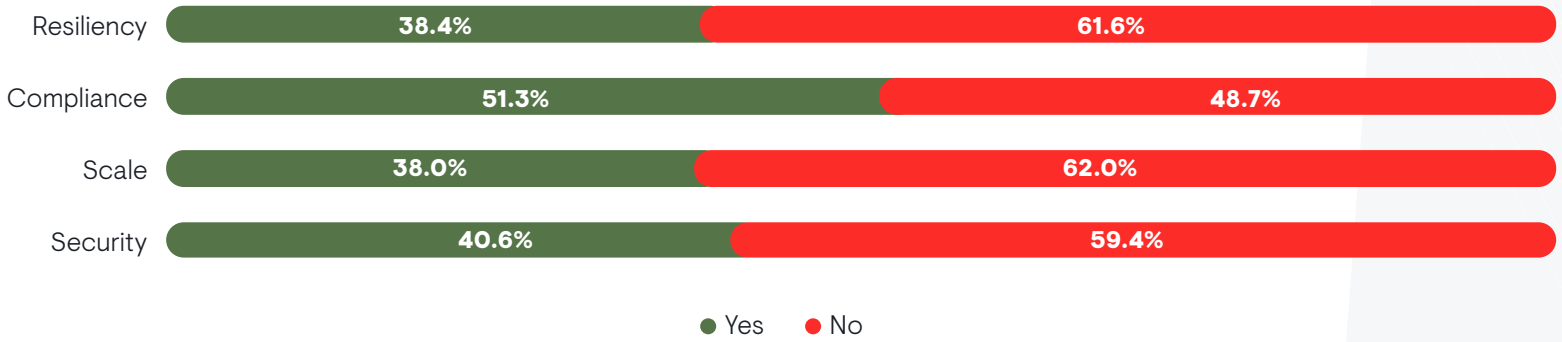

This paper is a view of the survey data from a financial service perspective relative to other industries. Together, these surveys offer a modern perspective on the challenges organizations face today regarding AI agents.

The AI Pivot and the Pacing of Financial Innovation

In the boardrooms of financial institutions today, the posture toward artificial intelligence is undergoing a pronounced and necessary evolution. The earlier imperative to accelerate AI adoption at nearly any cost has started to give way to a more disciplined demand for rigorous accountability, robust risk mitigation, and demonstrable alignment with regulatory expectations. Agentic AI promises systems capable of autonomous action on behalf of employees, customers, and business processes, but carries with it profound implications for security, compliance, intellectual property, and operational resilience.

Solving the Agent Identity Crisis: Security and Compliance at Scale

Financial sector leaders now confront a core tension: how to harness the efficiency and scale promised by intelligent agents while safeguarding the sensitive data, customer trust, and regulatory standing that define the sector. Central to this challenge is what can aptly be termed the **AI identity crisis**, which is the urgent need to establish secure, verifiable, and compliant identity and access management frameworks capable of governing machine-scale autonomy without introducing unacceptable levels of exposure.

Agentic AI Survey Results for the Financial Industry: Exhibit Caution

The global EMA research survey, which included a targeted survey of financial sector professionals, reveals a distinctly cautious approach. Overall, AI deployment in the sector stands at 77.8%, slightly below the 82% cross-industry average. More telling is the distribution of maturity: 64.8% of financial organizations remain concentrated in limited deployments or pilot programs compared to 57% across other sectors. Full-scale production deployments lag significantly, with only 29.7% of financial sector respondents reporting agentic AI running externally at scale in production versus 40.6% in the broader market. This measured pace is not hesitation born of inertia, but a reflection of the sector's ingrained culture of prudence. Agentic initiatives in the financial sector are notably more likely to operate under sustained human oversight, and security teams assume ownership of agentic security responsibilities approximately 15% more frequently than in other verticals.

Status of agents working on behalf of external customers and contractors

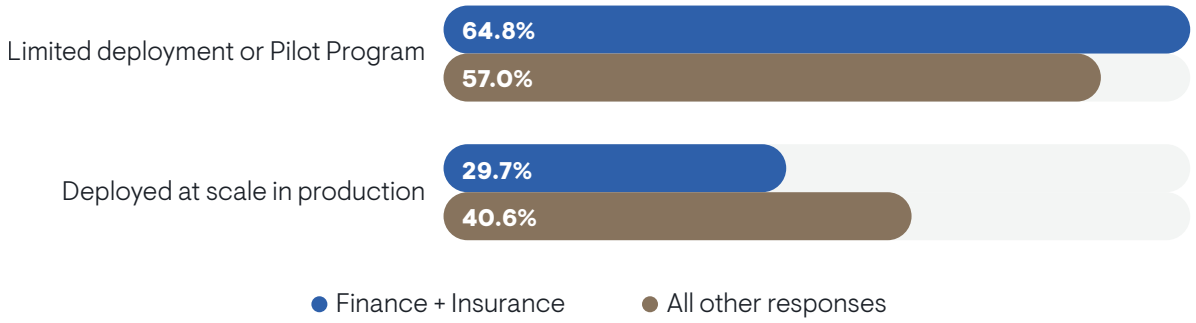

Such caution is well-founded. Agentic systems introduce new attack surfaces and accountability challenges that traditional controls were not designed to address. The potential for prompt injection, data poisoning, or unintended autonomous actions amplifies risks around sensitive financial data and "company jewels," such as proprietary code, models, and customer information. Fearing excessive agency, boards and executives are increasingly requiring clear documentation of AI-generated outputs for intellectual property protection and insisting on auditable governance structures before scaling beyond controlled pilots.

Section Summary

Though considered a responsible approach, pilot programs carry their own limitations. Human oversight models are inherently difficult to scale, and prolonged reliance on pilot programs risks ceding competitive ground in an environment where machine-speed decision-making and customer engagement are becoming table stakes. The path forward demands a foundational layer of identity infrastructure that can support secure, agentic operation at enterprise scale while satisfying the exacting demands of regulators and risk committees alike.

Regulatory Weight and the Agent Identity Readiness Gap

Given the increased sensitivity of data being managed and governed, financial sector organizations operate under one of the most demanding regulatory environments in any sector. Cybersecurity is no longer a back-office concern, but a core element of enterprise risk management and regulatory supervision. While organizations attempt to improve access through open banking and financial-grade API efforts (FAPI), institutions must now demonstrate cybersecurity maturity through more contemporary frameworks, with NIST Cybersecurity Framework (CSF) 2.0 emerging as the de facto standard for U.S. banks and credit unions designed to reduce their risk posture.

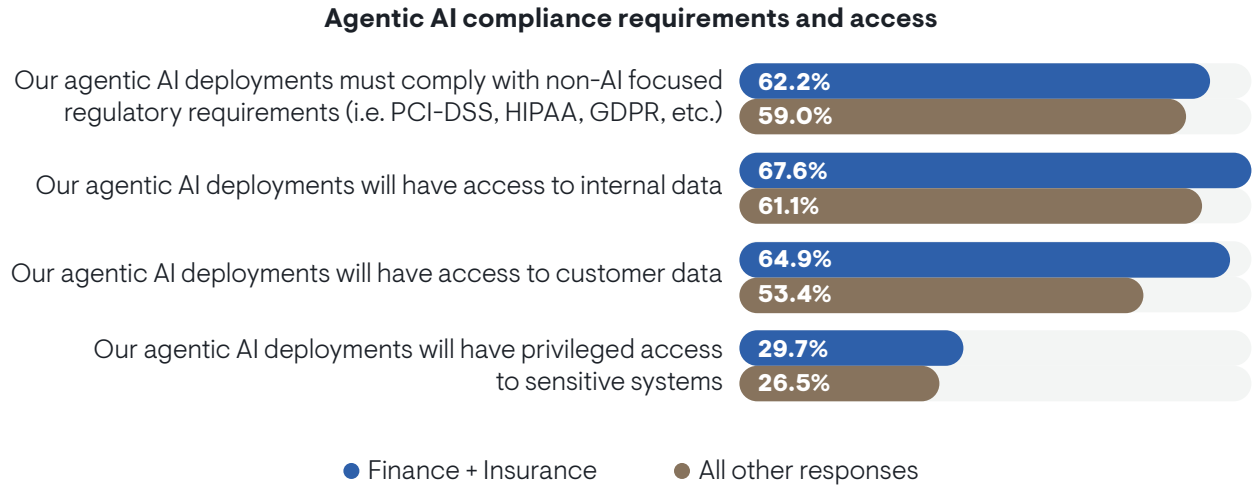

Layered atop this are sector-specific mandates that continue to tighten, such as the Gramm-Leach-Bliley Act (GLBA) Safeguards Rule and its recent amendments, which require timely reporting of certain security events to the Federal Trade Commission. This adds more pressure to ensure agent identities are managed and audited. In the insurance sector, the NAIC Insurance Data Security Model Law – now adopted in the vast majority of U.S. states – imposes uniform requirements for risk assessments, information security programs, and notification of cybersecurity events to the state insurance commissioner within 72 hours.

For institutions with European operations or serving EU clients, the Digital Operational Resilience Act (DORA) represents a significant elevation in expectations, as does the General Data Protection Regulation (GDPR). Add to this Payment Services Directive 3 (PSD3) and Financial Data Access (FiDA) regulations, and those financial sector organizations doing business with the EU find themselves under even more continuously evolving requirements.

These evolving requirements translate into specific technical mandates that directly impact agentic AI deployments. Regulators and cyber insurers increasingly demand phishing-resistant multi-factor authentication, favoring hardware keys or passkeys over legacy methods, such as SMS. Zero trust principles are no longer aspirational, but are expected, mandating continuous verification of every user, device, and now autonomous agent rather than relying on perimeter-based trust.

The Preparedness Paradox

Survey data from financial sector professionals highlights a clear paradox in preparedness for these requirements. While almost half rate their current IAM infrastructure as robust and ready at scale, significant gaps persist in other critical areas. A full 62.2% indicate that IAM resiliency requires additional components (versus 50.4% overall), 59.5% report that security needs strengthening (versus 48.7%), and 43.2% see compliance capabilities as incomplete (versus 39.7%).

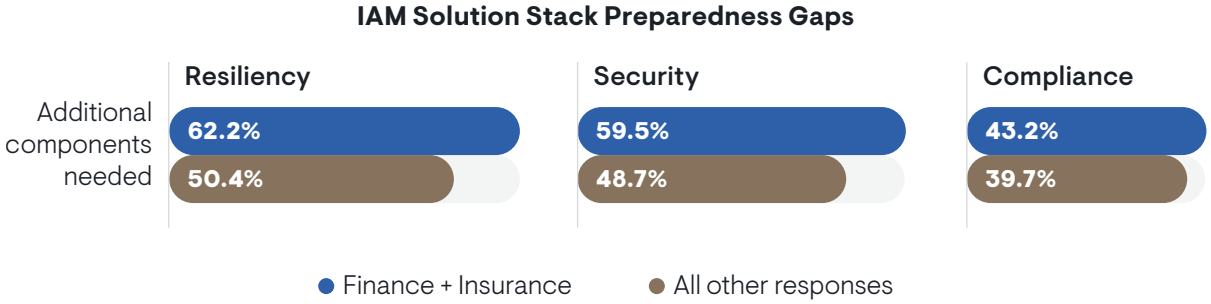

Compounding these technical gaps is a distinct organizational dynamic: security teams in the financial sector assume ownership of agentic security responsibilities approximately 15% more often than in other sectors. This centralized control culture underscores prudent risk governance, but also reveals the scalability challenge. Agentic systems, by design, operate at machine speed and scale, often with broad access privileges to execute tasks on behalf of employees or customers. Traditional human oversight models provide short-term comfort, but cannot sustainably govern at scale the volume and velocity of autonomous actions. The risks are not theoretical. Agentic accounts, with their potential for elevated privileges and direct interaction with sensitive systems, present attractive targets.

Attackers increasingly bypass credential theft entirely, exploiting prompt injections, data poisoning, or adversarial inputs to manipulate AI behavior and extract proprietary information or execute unauthorized actions. In an environment where users have moved beyond “attackers don’t break in…they log in” to “attackers don’t log in…they politely ask agentic AI for access,” the identity layer must now govern and protect in real time not only human users, but also non-human identities whose decisions and actions carry regulatory, financial, and reputational consequences.

Section Summary

The regulatory weight and identity readiness gap creates a material barrier to moving agentic AI from controlled pilot programs into full production. Financial institutions recognize the competitive necessity of scaling intelligent automation, yet they cannot do so until identity infrastructure can deliver verifiable assurance, resilient operation, and seamless compliance at machine scale. **The tension between strict oversight and innovation velocity is now the central strategic issue for security and technology leaders in the sector.**

Solving the SaaS Paradox with Composability

Financial sector organizations face a clear strategic tension in their identity and access management strategies. Survey findings reveal that 35.1% of respondents in the sector cite limited ability to customize or extend their current IAM solutions as a top challenge, slightly higher than the 30.8% cross-industry average. At the same time, 67.6% express a strong preference for an all-in-one SaaS platform, far exceeding the 50.4% recorded among other sectors. This apparent paradox stems from a practical reality: institutions require the operational simplicity, rapid deployment, and unified experience that a mature SaaS offering provides, yet they cannot afford the rigidity that often accompanies monolithic platforms when operating under stringent regulatory scrutiny and evolving agentic workloads.

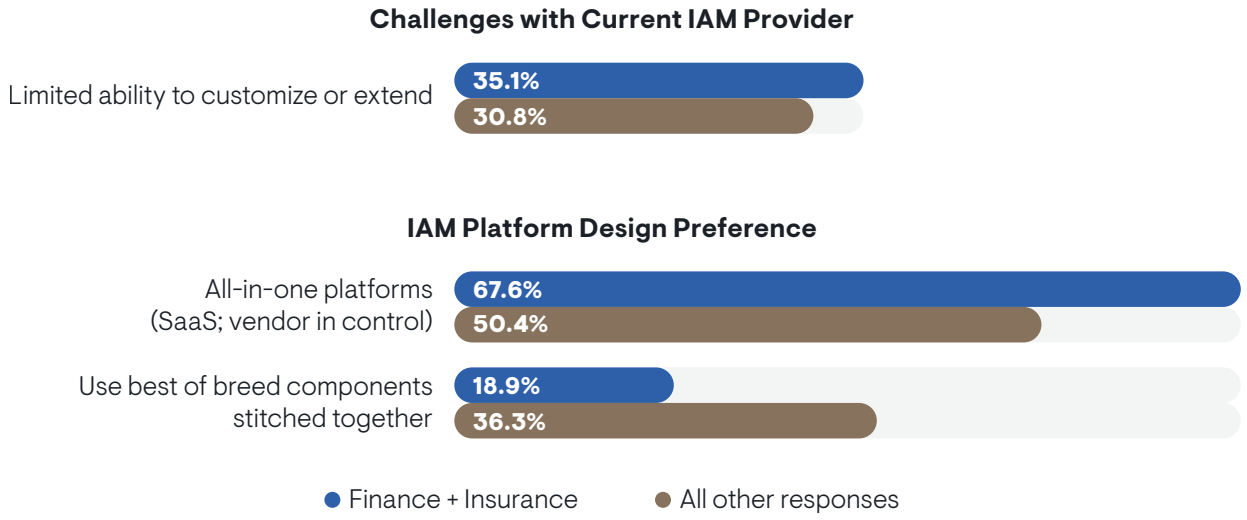

Traditional all-in-one IAM platforms frequently deliver convenience at the expense of flexibility. Once implemented, extending capabilities to address new resiliency requirements, integrating specialized security controls, or adapting to updated compliance mandates can involve lengthy vendor roadmaps, custom development projects, or costly professional services. In an environment where NIST CSF 2.0, DORA, and GLBA demand continuous proof of resilience, third-party oversight, phishing-resistant MFA, and zero trust enforcement, such constraints create friction that slows progress from pilot to production.

Solving the Problem through Composability

A composable approach to IAM resolves this conflict by delivering the best of both worlds. It provides the frictionless, unified management experience associated with leading SaaS platforms while preserving true modularity. Organizations can select and integrate precise components for authentication, authorization, fine-grained permissions, consent management, and auditability – mixing and matching them as needed without sacrificing operational simplicity or centralized visibility. This architectural model enables financial institutions to plug in specialized resiliency mechanisms, enhanced compliance tooling, or agent-specific governance features exactly where and when regulatory or threat landscapes require them.

For agentic AI initiatives, composability offers particular advantages. Autonomous agents functioning as "trusted coworkers" or customer-facing "virtual assistants" require identity models that treat non-human actors as first-class entities, with scoped privileges, delegated authority, dynamic policy enforcement, and comprehensive audit trails. A flexible and composable foundation supports these requirements natively: it allows security teams, which already own agentic security responsibilities more frequently in the financial sector, to enforce least-privilege access, real-time verification, and machine-scale controls without rebuilding core infrastructure. It also accommodates the sector's preference for human oversight during early deployments while providing a clear pathway to higher autonomy as confidence and controls mature.

Section Summary

A composable model aligns directly with the sector's risk culture. Institutions gain the ability to maintain rigorous customization and control for compliance reporting, data residency, and supply chain oversight while benefiting from the scalability, availability, and managed operations characteristic of modern SaaS environments. The result is an identity layer that is both enterprise-grade in simplicity and versatile in how you deliver it.

Conclusion

The financial sector will safely advance agentic AI from limited pilots into secure, regulated production environments only when identity infrastructure can deliver verifiable assurance at machine scale. A composable IAM architecture provides that essential foundation: it reconciles the demand for all-in-one operational excellence with the non-negotiable need for customization, extension, and resilience. By adopting this approach, institutions can confidently verify identities, whether human or agent, and strengthen infrastructure resiliency, satisfy regulators, and unlock the competitive advantages of autonomous intelligence without compromising the controls that define trust in the sector.

Report Sponsor

Ory, the modern choice for builders, offers customer identity and access management (CIAM), Agent IAM, B2B IAM, and one of the world's most widely adopted IAM platforms, managing more than 2.5 billion identities across open source and commercial deployments. Ory's infrastructure powers 10 percent of the top 40 websites and serves leading enterprises in financial services, technology, media, and other sectors requiring flexible, high performance identity solutions. With over 45,000 GitHub stars and 700 million downloads, Ory delivers enterprise grade security with developer friendly flexibility. Ory is backed by investments from Insight Partners, Balderton Capital, PHX Ventures, and IQT. For more information specific to financial services, visit **https://www.ory.com/industry/financial-services**.

Methodology and Demographics

Out of 1,256 respondents, there were 271 qualified responses from over 20 different industries, five different countries (United States, United Kingdom, Canada, Germany, France), and organizations with 500 or more employees. Targeted job roles included IT and security directors, managers, and C-level executives. In the research, we classified 500-999 employees as "medium" organizations, and 1,000 and above as "large" organizations.

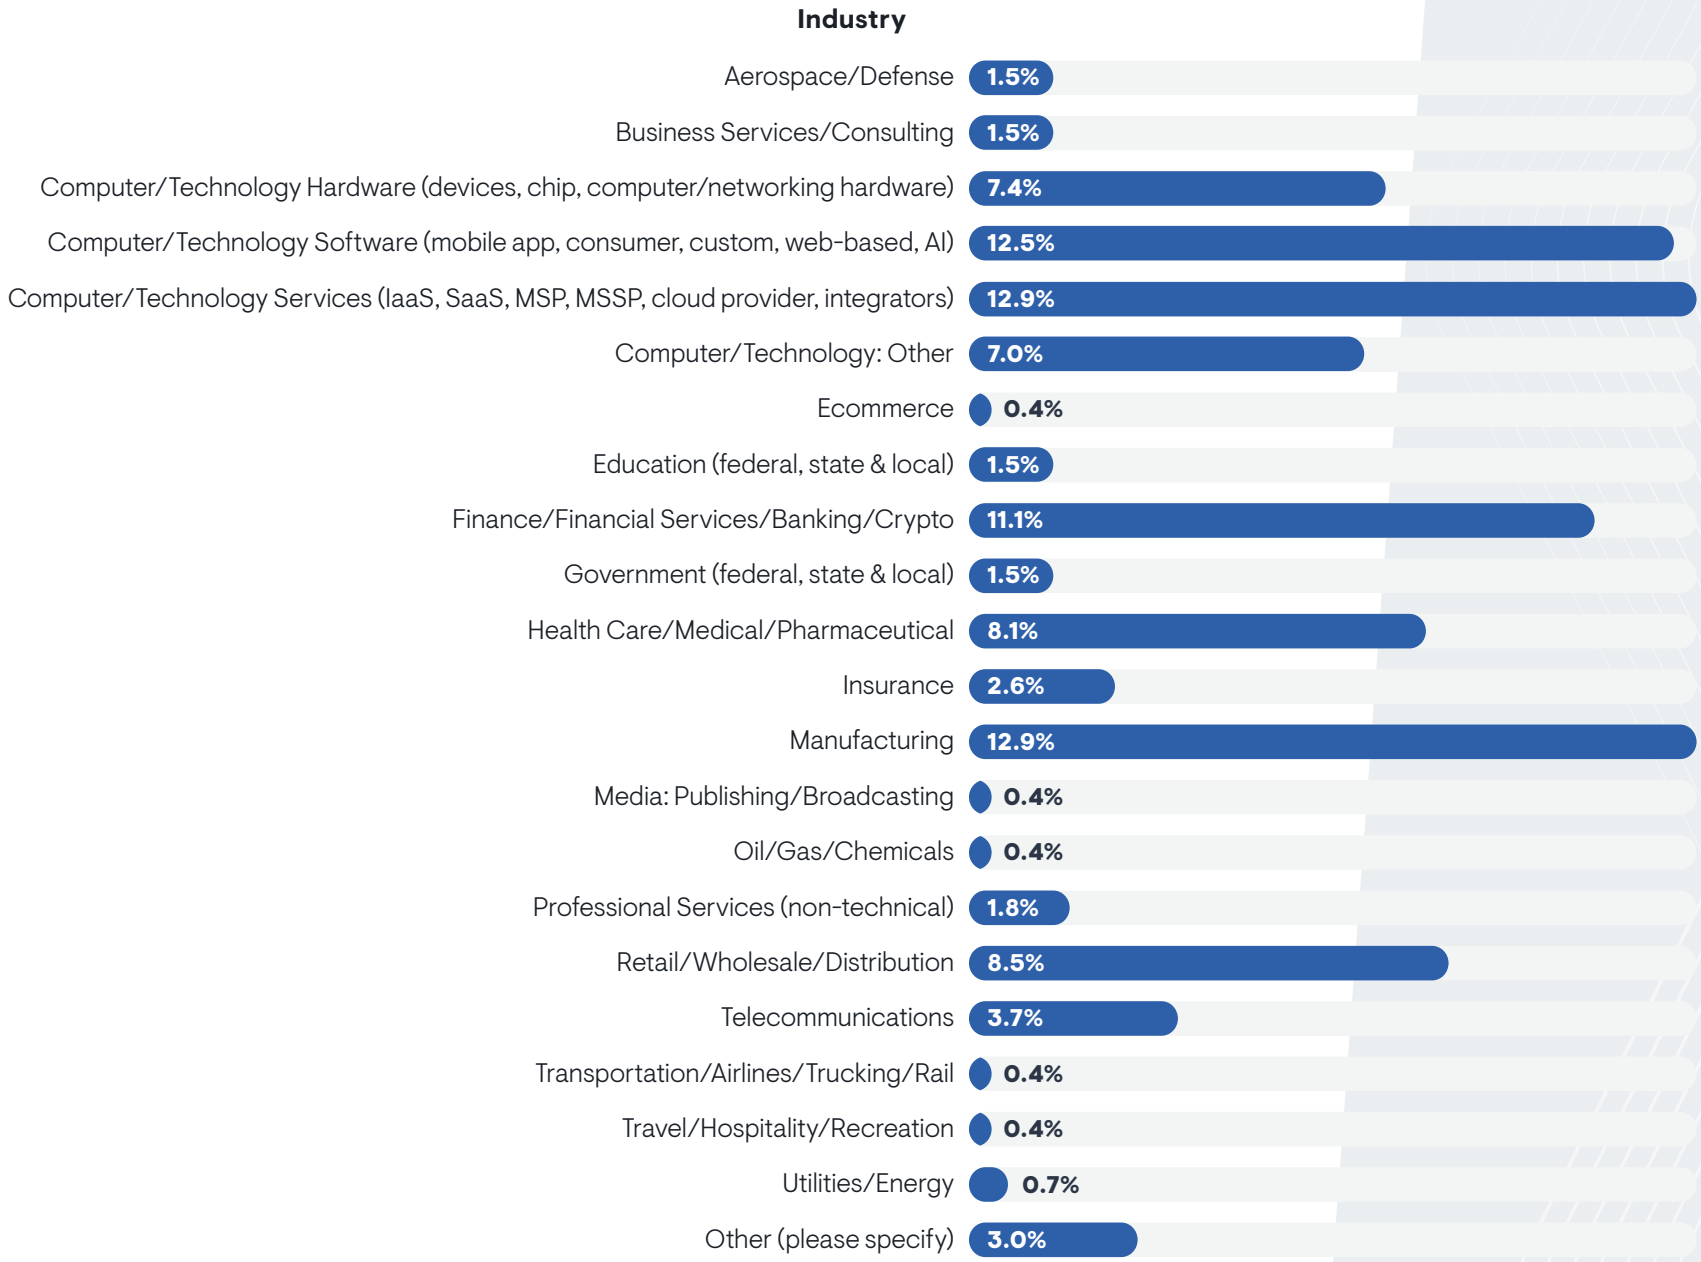
Industry
Aerospace/Defense 1.5%
Business Services/Consulting 1.5%
Computer/Technology Hardware (devices, chip, computer/networking hardware) 7.4%
Computer/Technology Software (mobile app, consumer, custom, web-based, AI) 12.5%
Computer/Technology Services (IaaS, SaaS, MSP, MSSP, cloud provider, integrators) 12.9%
Computer/Technology: Other 7.0%
Ecommerce 0.4%
Education (federal, state & local) 1.5%
Finance/Financial Services/Banking/Crypto 11.1%
Government (federal, state & local) 1.5%
Health Care/Medical/Pharmaceutical 8.1%
Insurance 2.6%
Manufacturing 12.9%
Media: Publishing/Broadcasting 0.4%
Oil/Gas/Chemicals 0.4%
Professional Services (non-technical) 1.8%
Retail/Wholesale/Distribution 8.5%
Telecommunications 3.7%
Transportation/Airlines/Trucking/Rail 0.4%
Travel/Hospitality/Recreation 0.4%
Utilities/Energy 0.7%
Other (please specify) 3.0%

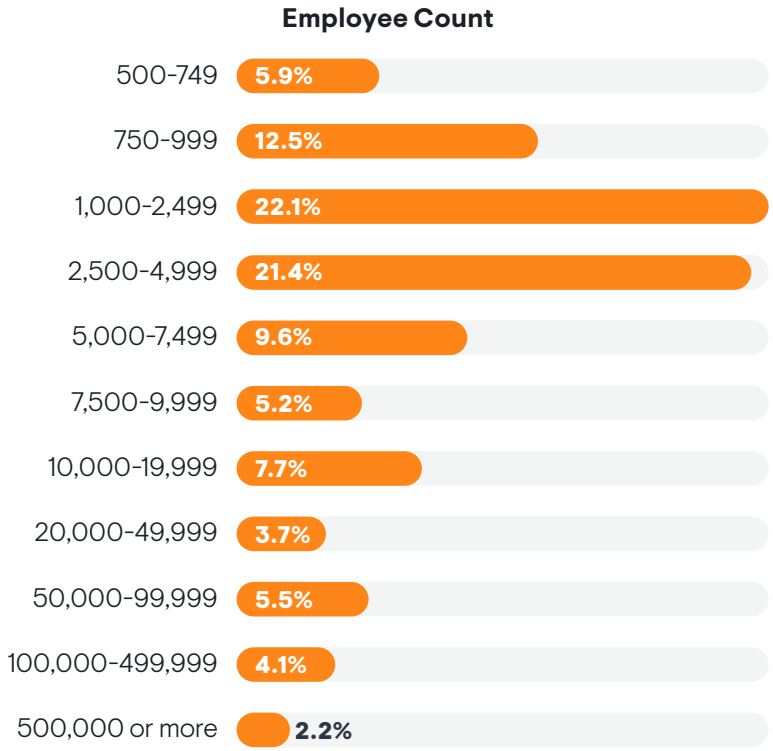
Employee Count
500-749
5.9%
750-999
12.5%
1,000-2,499
22.1%
2,500-4,999
21.4%
5,000-7,499
9.6%
7,500-9,999
5.2%
10,000-19,999
7.7%
20,000-49,999
3.7%
50,000-99,999
5.5%
100,000-499,999
4.1%
500,000 or more
2.2%

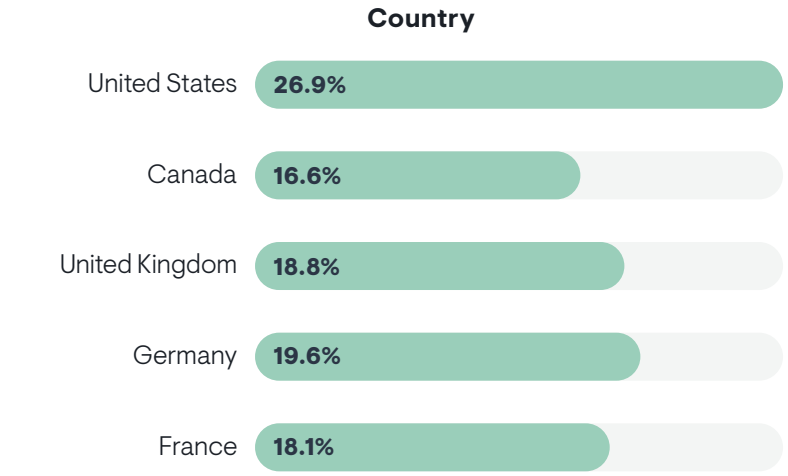
Country
United States
26.9%
Canada
16.6%
United Kingdom
18.8%
Germany
19.6%
France
18.1%

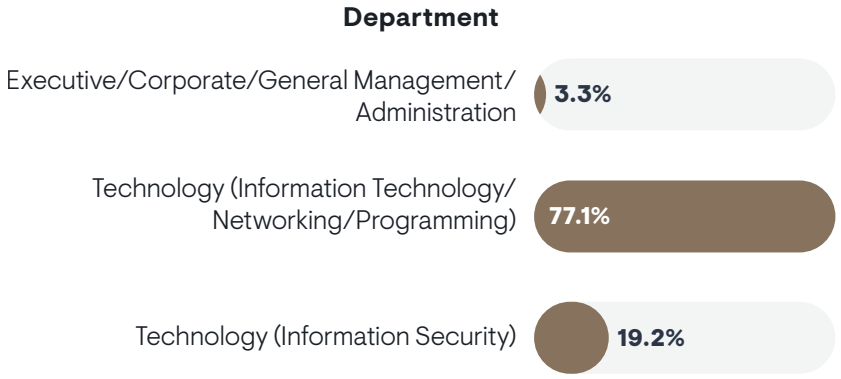
Department
Executive/Corporate/General Management/ Administration
3.3%
Technology (Information Technology/ Networking/Programming)
77.1%
Technology (Information Security)
19.2%

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT research and consulting firm dedicated to delivering actionable insights across the evolving technology landscape. Through independent research, market analysis, and vendor evaluations, we empower organizations to make well-informed technology decisions. Our team of analysts combines practical experience with a deep understanding of industry best practices and emerging vendor solutions to help clients achieve their strategic objectives. Learn more about EMA research, analysis, and consulting services at **www.enterprisemanagement.com** or follow EMA on **X** or **LinkedIn**.

2770 Arapahoe Road, Ste 132 - 544, Lafayette, CO 80026 +1 303.543.9500 **www.enterprisemanagement.com**

4722.050526