jumpcloud™

The Agentic IAM Pulse Report:

Closing the Governance Gap to Accelerate with AI

72% of Organizations Have AI Agents in Production.

Those With Governance in Place Are Three Times More Likely to Scale Without Limits.

How much control do you have over your AI agents today?

AI agents are rapidly moving into finance, HR, access management, and customer-facing systems, taking on a growing role in daily operations. Across industries and company sizes, organizations are already deploying them at scale, but governance is not keeping pace.

72% of organizations already have AI agents in production, yet key controls remain inconsistent. AI agents are routinely granted greater access than human users. A significant share of agents lack centralized kill switches, and many are not fully integrated into formal identity and access management (IAM) systems.

When access outpaces oversight and accountability is unclear, security becomes the top barrier to further expansion.

This report examines how organizations are managing the tension between expanding AI agent use and inconsistent control. Covering 250+ IT, security, and identity leaders across the U.S. and U.K., it looks at where AI agents are being deployed, how much access they receive, where governance remains fragile, and why stronger control is becoming a practical requirement for scaling AI agent deployment with confidence.

Control Is Becoming the Main Limit on Scale

- **72%** of organizations have AI agents in production
- **50%** cite security, compliance, or lack of visibility and control as top barriers
- **92%** report limits to scaling AI agents

AI Agents Are Rapidly Scaling Across the Business

Organizations are already relying on agents in core business workflows, and a significant share are well beyond early-stage testing.

72% of organizations have AI agents in production. Of those, **41%** are using them in internal or low-risk workflows, while **31%** have deployed them in business-critical environments spanning financial reporting, HR provisioning, access management, and customer-facing systems. Only **28%** remain in testing.

As agent deployments mature, scale quickly increases. Organizations report an average of **15.8** AI agents actively interacting with internal systems or APIs. Among organizations running business-critical deployments, that figure rises to **20.3**, compared to an average of **10.8** for organizations still in the testing phase.

Agent growth is also reshaping the identity landscape. **53%** of organizations already have more non-human identities than human employees, and **23%** report **six times or more non-human identities than human users**. Among organizations running business-critical AI workflows, **37%** report between **six and 20 times more non-human identities than humans**.

Traditional identity infrastructure was built for a workforce of people. In a growing number of organizations, human identities now make up a smaller share of the overall identity environment.

Deployment Maturity Increases with Company Size

Enterprise-sized organizations are driving the move into production, while smaller companies are proceeding more cautiously. Nearly half of smaller companies (**48%** of those with 200-499 employees) are still keeping their AI agents in testing or sandbox environments, compared with just **22%** of large enterprises (1,000-2,500 employees).

The production gap is even larger at the most advanced level, with only **17%** of small companies running agents in business-critical workflows, versus **37%** of large enterprises.

# of AI agents currently interacting with internal systems or APIs

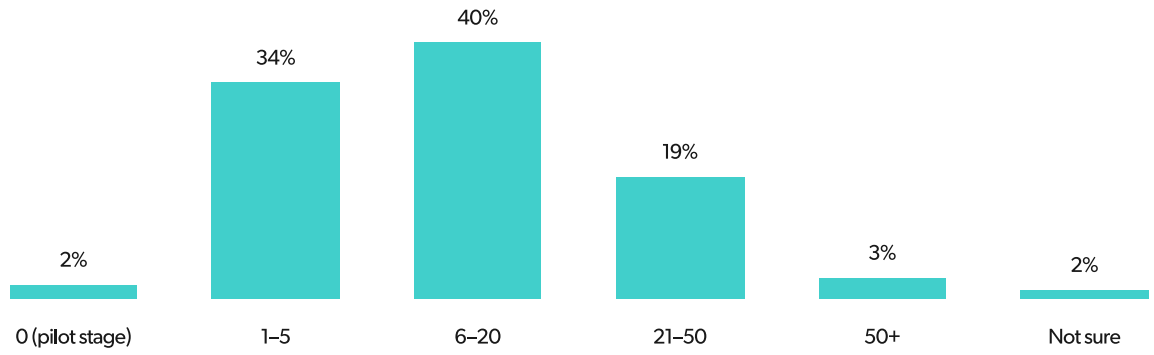

AI Agents Are Gaining Power Without Consistent Oversight

As organizations deploy AI agents in more critical workflows, access and autonomy are expanding faster than governance controls.

66% of organizations say AI agents have equal or greater system access than human users:

- **29%** give agents equal access
- **20%** give them more access than humans
- **17%** give them significantly more access than humans

AI agent access is even more significant in business-critical environments, with **38%** of organizations giving AI agents significantly more system access than the humans working alongside them.

Among testing-stage organizations, **48%** give AI agents less access than human users, but this restraint largely disappears for business-critical deployers, where agents are granted far greater control over sensitive systems.

AI agents system access compared to humans
by Deployment stage

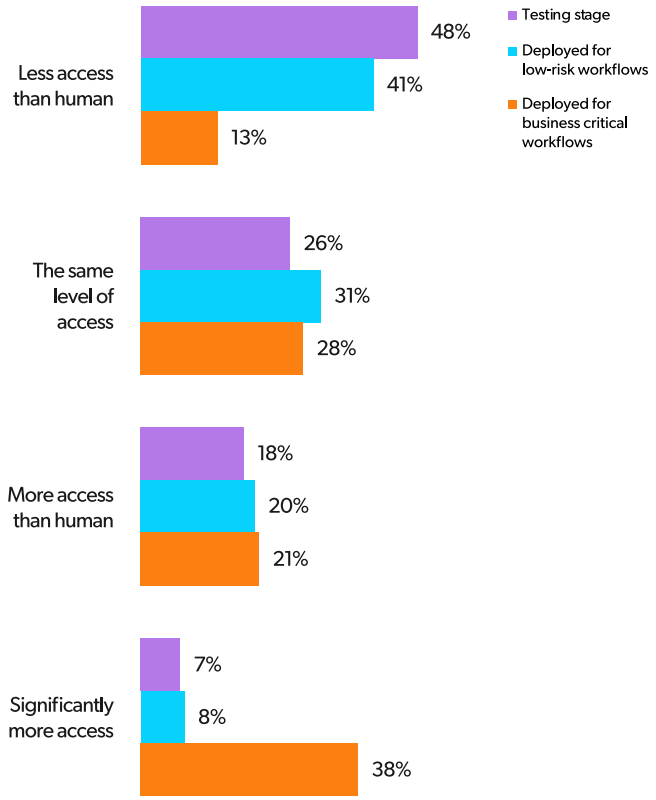

The same holds true for oversight practices. **46%** of organizations allow high-risk actions, including automated with post-action review (35%) and fully autonomous with no human oversight (11%) to proceed without prior human approval.

Handling high-risk AI agent actions

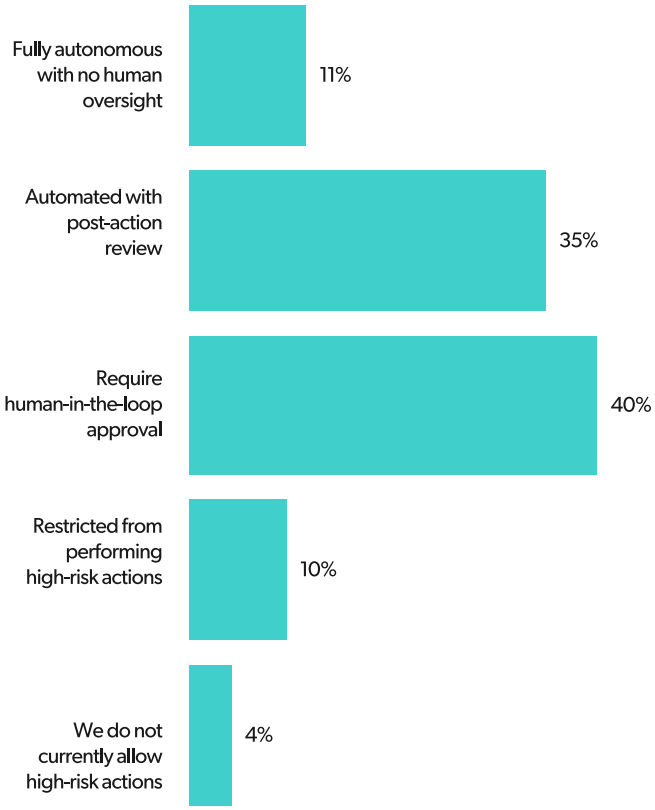

Meanwhile, the share requiring Human-in-the-Loop (HITL) approval before high-risk actions proceed falls from 48% in testing to just 29% in business-critical environments — replaced by automated post-action reviews and fully autonomous actions (39%) with zero human supervision (24%). Oversight is declining precisely where the stakes are highest.

Handling high-risk AI agent actions
by Deployment stage

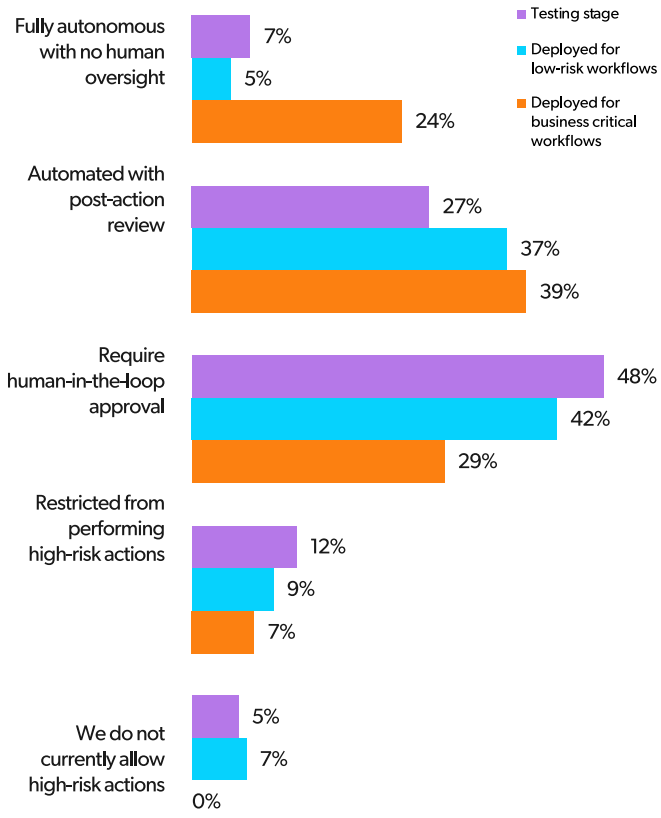

The Compounding Risk of Scale

Greater deployment depth produces a larger, more complex environment, and the governance gaps scale along with it. Business-critical organizations report nearly twice as many AI agents, on average, as testing-stage organizations. Advanced deployers are also significantly more likely to report much higher ratios of non-human identities to employees.

When an enterprise's fastest-growing identity group is also its least supervised, scaling deployment means scaling the organization's security blind spots.

AI agents system access compared to humans
by Deployment stage

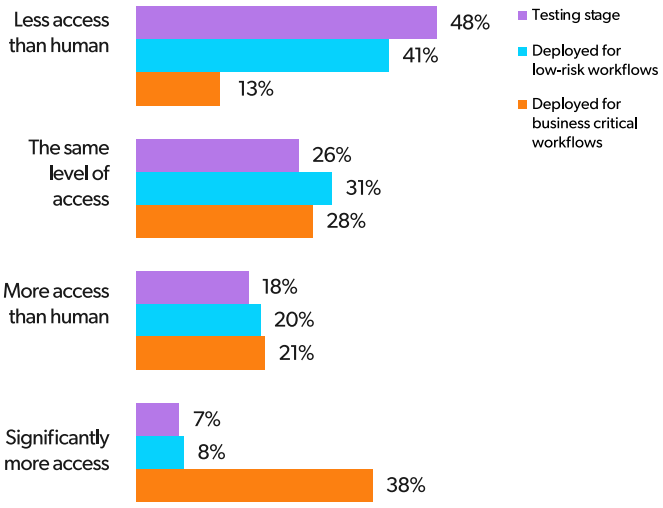

AI Agent Access Has Already Outpaced Governance

- **66%** of organizations grant AI agents equal or greater access than human users.
- **38%** of organizations in business-critical environments grant significantly more access than human users.
- **24%** of organizations with business-critical deployments allow agents to execute high-risk actions with fully autonomous, zero-human oversight.

AI Agent Identity Is Not Yet Standardized

Organizations are bringing AI agents into their IAM processes, but the majority are in transition rather than operating under a clear and consistent model. Barely one-third (**37%**) have fully integrated AI agents into IAM.

The remaining **63%** of organizations lack complete IAM coverage: **42%** are only partially integrated, **17%** manage agents entirely outside formal systems, and **4%** lack formal policies entirely.

Authentication practices are equally fragmented, with organizations relying on a mix of methods rather than a standardized approach:

- **53%** use individually scoped machine or workload identities
- **49%** rely on long-lived API keys
- **37%** use shared service accounts
- **36%** use secretless access through a centralized gateway
- **26%** use just-in-time or temporary authorization for sensitive actions

On average, organizations use **2.2** authentication methods overall, rising to **2.7** among those with business-critical deployments. The mix of authentication methods shows many organizations are still prioritizing speed and practicality over standardization, and among business-critical deployers, long-lived API keys climb to **71%**, making the most sensitive environments the most exposed.

The IT Leadership Disconnect

There is a clear gap in perception between executives and the teams responsible for day-to-day implementation. **68%** of CIOs believe their AI agents are fully integrated into formal IAM policies. Among IT Managers and IT Team Leads, only **35%** report full integration.

When leadership believes IAM integration is complete and practitioners know it's not, the result is a fragmented view of AI agent governance.

AI agents authentication when accessing internal systems

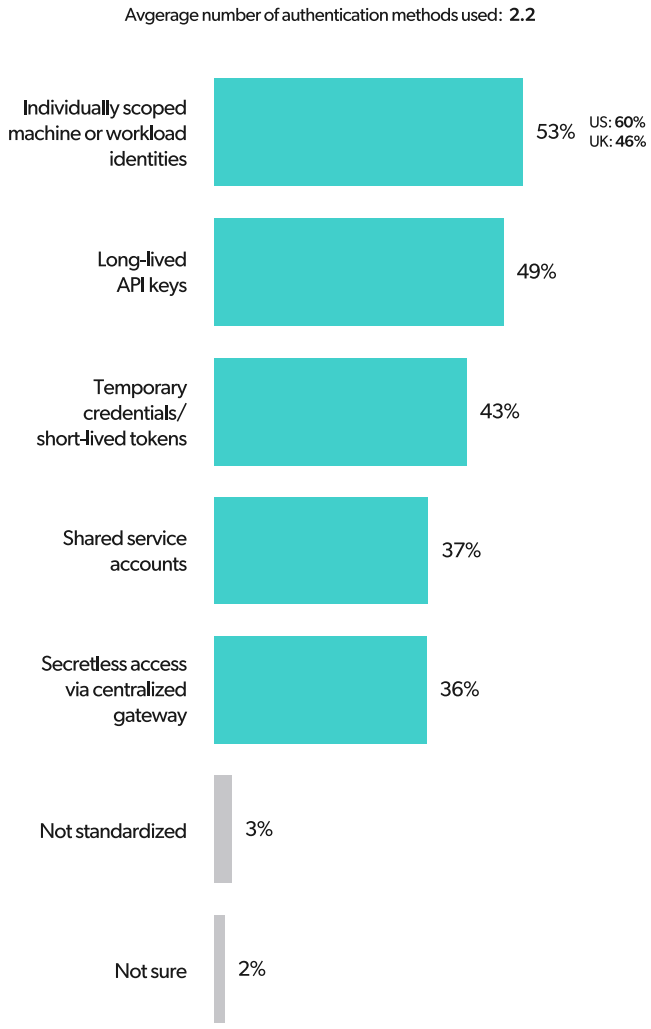

Larger Companies Rely More on Risky Authentication

As organizations grow, they need to rely on more fragmented and less secure authentication methods, trading security for deployment speed.

Organizations with 1,000-2,500 employees are significantly more likely to use long-lived API keys (**56%**) than smaller organizations with 200-499 employees (**31%**). Larger companies also juggle more authentication methods on average, with **2.5** methods compared to **1.8** for smaller orgs.

Most Organizations Lack Clear Control and Ownership

Access is only part of the governance problem. Managing AI agents confidently over time requires visibility into agent activity, audit trails, the ability to shut them down quickly, and clear lines of accountability.

There are significant gaps on each of these fronts for organizations:

- **59%** lack centralized visibility into agent activity
- **59%** don't maintain full audit trails
- **55%** don't have a centralized kill switch for AI agents
- **33%** say their only option is to disable agents manually, system by system

The ability to see what agents are doing — and stop them if needed — remains inconsistent even after they've moved into production.

Ownership Is Fragmented

Accountability is not being designed deliberately; it's falling by default to whoever owns the adjacent systems. Only **17%** of organizations have a designated security leader accountable for AI agent actions, **47%** default that responsibility to IT, and just **6%** have a cross-functional governance committee in place.

As organizations move from testing to full production, shared responsibility declines. In the testing phase, accountability is split between IT (**38%**) and a joint IT-security function (**33%**). In business-critical deployments, shared responsibility drops to **17%**, and IT alone absorbs that accountability for **51%** of organizations.

The more advanced the deployment, the more concentrated and fragile the ownership structure becomes.

Control Gaps Remain Wide

- **83%** don't have clear security ownership
- **59%** have no centralized visibility
- **55%** lack a centralized kill switch

Speed vs. Control: U.S. and UK Take Different Paths

Organizations in the U.S. and UK are approaching AI agent deployment differently. U.S. organizations are moving faster and granting broader access, while their UK counterparts are more likely to apply structured oversight and control.

Full autonomy for high-risk actions:
16% US vs. 8% UK

Greater-than-human system access:
44% US vs. 29% UK

Human-in-the-loop approvals:
52% US vs. 64% UK

Full audit trails:
34% US vs. 47% UK

Business-critical deployment:
34% US vs. 29% UK

Governance Is Becoming the Limit to Growth

AI agent deployment is no longer the challenge. Scaling it safely is. The data shows that governance, not budget or skills, is now the primary constraint on expansion.

Security concerns are the single biggest barrier to expanding AI agent deployment:

- **26%** cite security concerns
- **16%** cite integration complexity
- **13%** cite budget constraints
- **12%** cite skills gaps

A total of **92%** of organizations report some limit to scaling AI agents, and only 8% report no limits on scaling. Reaching a more advanced stage of deployment doesn't ease these concerns, as organizations running business-critical workflows are the most likely to cite security as their top expansion barrier (**28%**).

A small group has cleared the governance hurdle, and their scaling capacity reflects this. Organizations with business-critical deployments are **more than three times more likely** than testing-stage organizations to report no current limits to scaling: **13%** of business-critical deployers report no limits, compared to **4%** still in testing.

That same group is also more likely to cite security as their top remaining barrier — not because governance has failed them, but because deeper deployment creates a clearer picture of where gaps remain. Getting further along doesn't eliminate risk awareness; it exposes where control is still incomplete.

US vs. UK: Speed vs. Control Shapes Scaling Outcomes

How organizations approach governance has a direct bearing on how far they can scale. U.S. organizations cite security as their top barrier at **32%**, compared to **21%** in the UK, where barriers are more evenly distributed across integration complexity (**19%**) and budget constraints (**17%**).

10% of UK organizations report no current limits to scaling AI agents, compared to **5%** in the U.S. Stronger governance does not eliminate scaling barriers, but it does reduce their impact.

Top barriers to expanding AI agent deployment

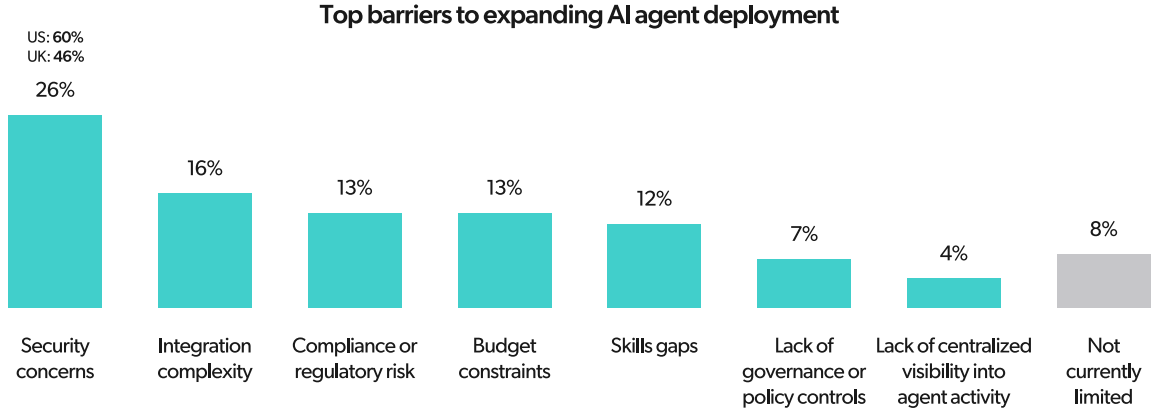

A Practical Framework for AI Agent Governance

Your organization needs a clearer way to govern how agents are identified, what they can access, what they are permitted to do, and how their actions are reviewed over time.

A clear governance model moves through four stages, from visibility to identity and access control to ongoing human review. Here are the four key stages that define this approach:

1. Discover

Governance starts with knowing what you have. Your organization needs a complete, accurate picture of the AI agents operating across your environment, including which tools they connect to, which APIs they call, which systems they can reach, and which workflows they touch. Without that inventory, consistent policy enforcement is impossible and incident response becomes guesswork.

Only 33% of organizations in the testing phase report centralized visibility into agent activity, a figure that rises to just 43% among business-critical deployers. The starting point for any governance program is a full audit of what is running and what it can access — sanctioned and unsanctioned agents alike, regardless of how they were deployed or who owns them.

2. Register

Once identified, every AI agent needs a formal identity record. Registration means placing each agent in a central directory, documenting what it is and what it does, mapping the systems and tools it interacts with, and assigning a named human owner who is accountable for its behavior and responsible for reviewing its access over time.

As deployment deepens, accountability becomes less shared and more concentrated in IT by default. A named custodian responsible for each agent from day one is the foundation that organizations need to make every subsequent governance step workable.

3. Manage

Registration establishes what an agent is and who owns it. Management determines what it's allowed to do. Apply least-privilege access as a baseline, ensuring agents can only reach the systems and perform the actions their function genuinely requires.

Long-lived credentials and shared service accounts create unnecessary exposure and should be replaced with time-bound or secretless access wherever possible.

High-risk actions like financial transactions, identity changes, and access provisioning should require explicit human approval before execution. Every agent deployment also needs a single point from which access can be revoked immediately across all connected systems if something goes wrong.

Control Should Be Instant

Only **45%** of organizations have implemented role-based access policies to specific AI agents, and more than half lack a centralized kill switch entirely. When agents are executing consequential, irreversible actions, the ability to act in seconds is crucial.

4. Govern

Access granted is not access permanently justified. You need an ongoing review process that gives humans meaningful visibility into what agents are doing, how their behavior compares to what was intended, and whether their current permissions still make sense. Logs, audit trails, and regular access reviews are the operational backbone of that process.

The data suggests that organizations are moving in the wrong direction. In the testing phase, 48% rely on proactive human-in-the-loop approvals. For low-risk workflows, the number drops to 42%, and in business-critical deployments, it falls to 29%, with organizations opting for automated post-action reviews instead.

Post-action review leaves organizations permanently one step behind. With agents taking on more sensitive work, the review process should become more rigorous, built around scheduled access reviews, clearly defined escalation thresholds, and a clear process for updating or revoking permissions when an agent's behavior or role changes.

The Four-Stage AI Agent Governance Framework

The following framework outlines how organizations can close the gap.

Stage	Core Question	What It Requires
Discover	What AI agents are operating in the environment?	Visibility into agents, tools, integrations, connected systems, and workflows.
Register	How is each agent identified and governed?	Directory entry, identity record, system mapping, and human owner.
Manage	What can each agent access and do?	Least privilege, entitlements, role-based policies, time-bound or secretless access, approvals, and shutoff capacity.
Govern	How are actions monitored and evaluated over time?	Logs, audit trails, human review, investigation, and policy updates.

AI agent deployment has outrun the controls needed to manage it safely. Agents are operating in your most sensitive workflows with more access than your human employees, fragmented identities, and no clear owner.

The organizations that close that gap first — on identity, access, visibility, and accountability — will be the ones that can scale with confidence. The ones that don't will find that their most powerful tools have become their greatest liability.

jumpcloud™

JumpCloud® is the AI-powered unified IT management platform designed to secure the modern workforce. By consolidating identity, device, and access management, JumpCloud provides intelligent, secure IT that scales from human users to autonomous AI agents. We help organizations around the globe eliminate complexity and turn AI risk into an optimized advantage, ensuring the right people and agents have secure access to the right resources at all times.

Jumpcloud.com | Blog | Resources | X | in | ▶