



IAM Buyer's Guide

Evaluating Agentic AI Solutions for Disconnected Applications



Introduction

Why the Right Architecture Matters

The rise of agentic AI introduces new potential for identity security, especially in automating integrations with disconnected apps. The promise is compelling: an AI agent that can build and run workflows on your behalf, with no need for costly custom development or professional services.

But while AI demos may showcase fast, seamless integrations, real enterprise environments are rarely that predictable. Identity workflows span apps that change frequently, fail unexpectedly, and behave inconsistently. The real challenge isn't building a workflow once, it's developing, running, and maintaining it reliably, securely, and at scale.

This guide outlines what to evaluate when assessing agentic AI solutions for disconnected apps. It focuses on where a purely agentic approach falls short, when human-in-the-loop support is still needed, and how to choose an architecture designed for the real-world complexity of enterprise app environments.

5

Critical Areas to Evaluate in Agentic AI Solutions

A purely agentic AI approach gives a single agent responsibility for both creating and executing workflows, without human-in-the-loop support. While appealing in theory, this architecture introduces risk, fragility, and operational inefficiencies when applied to real-world identity security environments.

Below are five key areas buyers should evaluate to understand whether an agentic AI solution is ready for enterprise scale.



01 • EXECUTION METHOD

Probabilistic ≠ Reliable

Agentic AI systems are probabilistic by nature. Rather than following a fully mapped sequence of steps, the agent “reasons” through workflows at runtime, guessing the next best action based on training data.

This approach is especially risky in identity lifecycle management (LCM), where tasks like provisioning, deprovisioning, and access updates require accuracy and predictability. When a probabilistic agent encounters an unpredictable application UI—a renamed button, a missing field, a pop-up error—workflows can break silently or take the wrong action entirely.



02 • APPLICATION COVERAGE

Browser-Based Solutions Aren't Enough

Many agentic solutions are limited to automating interactions in a web browser environment. While this may cover a portion of your app landscape, it excludes a wide range of critical systems, including on-premise thick clients (desktop apps), mobile apps, and custom-built internal apps. If the agent can't interact with these systems, it leaves critical identity and governance gaps.

A modern identity security solution must reach the full ecosystem—not just the apps that are easiest to automate.



03 • SCALABILITY & COST

AI Agents Can Strain Budgets and Performance

AI agents require compute cycles, inference costs, and API token usage to process large datasets or run frequent sync jobs. These demands can be significant, and in many cases, the costs are passed directly to customers, making the solution cost-prohibitive at scale. Costs can add up fast:

Running daily syncs across hundreds of apps and thousands of users can cost tens—or even hundreds—of thousands of dollars per year.

Unlike deterministic systems, agentic agents must “reason through” each step, which slows execution and increases infrastructure overhead.

Latency, timeouts, and slow throughput (e.g., paginated user lists) can turn routine identity operations into multi-hour or multi-day jobs.

A solution that can't operate cost efficiently under load becomes a blocker—not a bridge—for identity workflows. You need confidence that performance is predictable and that costs remain sustainable as you scale.



04 • PERFORMANCE

Slow Agents Can't Keep Up

Agentic AI isn't just expensive—it's often slow. Because of their non-deterministic nature, AI agents must “reason through” each step of a workflow in real time, which adds latency at every turn. The time it takes to process thousands of records can easily stretch into hours—or even days—making daily synchronization, a core requirement for identity management, impractical. Timeouts, partial failures, and inconsistent performance are common, especially at enterprise scale.

If your agent can't operate with predictable speed, it stops being a value driver and starts becoming an operational bottleneck.



05 • SECURITY READINESS

Authentication, MFA, and Bot Detection Can Break Agentic Workflows

Authentication is where most agentic systems run into real-world friction. A solution must do more than execute a workflow—it must do so securely, in a way that respects your existing controls and avoids detection by anti-bot systems.

MFA and Authentication

The agent needs to authenticate itself to carry out its task. Can the solution handle proprietary MFA methods, such as a virtualized soft token, or a multistep authentication flow? If the agent can't log in securely, the workflow is dead on arrival.

Navigating Bot Detection

As bots become more common, sophisticated bot detection solutions (like those from Cloudflare) are increasingly blocking non-human identity traffic. A basic agent running from a headless browser will be detected and blocked. The solution must intelligently run the workflow in a way that convincingly appears to be human.

A scalable identity solution must be able to log in, authenticate, and run workflows without compromising controls or raising red flags with detection systems.

Agentic AI Solution Due Diligence Checklist

Use this checklist to pressure-test any vendor relying on a purely agentic approach.

Critical Consideration	Question to Ask Your Vendor	Your Business Risk
☐ Execution Architecture	Is the AI agent used for both creation and direct execution? How predictable and safe is the AI's behavior during workflow execution, especially for sensitive identity flows like LCM?	Unpredictable Outcomes: Probabilistic execution introduces risk to critical LCM operations.
☐ Edge Case Coverage	How is the agent guaranteed to handle the non-happy paths (e.g., existing user, permission conflict) without human intervention?	Operational Downtime: Workflow execution stops at the first unhandled error, requiring manual recovery.
☐ Reliability Under UI Changes	What happens when the application UI changes? How does the agent ensure it does not take an unintended, destructive action?	Security Risk/Compliance: Unintended LCM operations (e.g., failed de-provisioning) could violate security policy.
☐ Support for Varied App Types	Can the solution support more than just web apps? What about on-prem, thick clients, or mobile apps?	Incomplete Coverage: Critical disconnected applications are left unsecured and ungoverned.
☐ Scalability & Cost	• How does the system handle large-scale syncs and high-volume jobs? • How long will it take your agent to synchronize 10,000 records without a CSV file? • Are compute/token costs sustainable at enterprise scale? • Does the vendor or customer pay this?	Cost-Prohibitive Operations: Routine, high-volume operations become too slow and too expensive to run daily.
☐ Security & MFA	What logic is built to handle multi-factor authentication (MFA) methods like multi-step flows, proprietary soft tokens, etc.? Can it complete complex login sequences involving captchas or pop-ups?	Security Gaps: You are forced to lower the security bar by disabling MFA to enable automation.
☐ Bot Detection Evasion	What specific infrastructure ensures the agent can reliably bypass bot detection systems like Cloudflare or Akamai without being blocked?	Integration Failure: The agent is blocked from logging in, making the integration unusable at scale.

About Cerby

Cerby is the identity automation platform purpose-built to secure disconnected applications—those that fall outside the reach of traditional identity security tools. By integrating with existing IAM, IGA, and PAM systems, Cerby brings centralized access controls, automates manual security tasks, and extends governance across your entire application ecosystem. IT and security teams gain complete visibility and control, reducing risk and operational overhead. Founded in 2020, Cerby is backed by leading investors and trusted by enterprises across the globe.

To learn more, visit www.cerby.com