CHRISTIAN & TIMBERS



The AI-Native CISO

Demand for AI-Native Security Leadership Rises 256% in the First Seven Months of 2026



The requirement has moved from differentiator to threshold. A CISO who cannot secure agentic systems is no longer a candidate.

EXECUTIVE SUMMARY

Executive Summary

Christian & Timbers tracked a 45% increase in companies seeking to replace their Chief Information Security Officer or Chief Security Officer during the first seven months of 2026 compared with the same period in 2025.

Underneath that headline number is a far more consequential shift. Demand specifically for an AI-native security executive rose 256% over the same period. In the first seven months of 2025, fewer than half of the companies replacing a CISO named deep agentic AI expertise as the single most critical requirement. In the first seven months of 2026, every one of them did.

The requirement has moved from differentiator to threshold. A CISO who cannot secure agentic systems is no longer a candidate.

45%

CISO / CSO REPLACEMENT

Increase in companies seeking to replace their Chief Information Security Officer or Chief Security Officer, January to July 2026 against 2025.

256%

AI-NATIVE REQUIREMENT

Increase in demand specifically for an AI-native security executive over the same period.

100%

SHARE OF 2026 SEARCHES

Share of 2026 CISO searches where AI-native expertise was named the single most critical requirement.

KEY FINDINGS

Key Findings

METRIC	JAN TO JUL 2025	JAN TO JUL 2026	CHANGE
Companies seeking to replace a CISO/CSO	22	32	+45%
Searches where AI-native expertise was the #1 requirement	9	32	+256%
Share of searches requiring AI-native expertise	41%	100%	+59 pts

The mandate is now universal.

In 2025, nine of twenty-two companies told us the most critical requirement of the new hire was deep agentic AI expertise, and each defined that expertise somewhat differently. In 2026, all thirty-two said the same thing, and the definitions have converged. What was an emerging preference twelve months ago is now the entry condition.

The CISO is second in line, not first.

The most aggressive replacement activity we are seeing sits with the CIO, where boards are elevating the role and replacing incumbents who cannot lead AI transformation. The CISO follows immediately behind. In several of our engagements, the two searches are running simultaneously and are being scoped together.

This is the leading edge of full AI-native transformation.

The companies furthest along in AI transformation have begun requiring that every role be AI-native, that every executive and, increasingly, every individual contributor have demonstrated experience deploying large language models or agents into their own work function. The purpose varies by function: productivity improvement, revenue acceleration, or, in the case of security, risk reduction. The CISO search is simply where this standard becomes non-negotiable first, because the downside is measured in breaches rather than in missed efficiency.

METHODOLOGY

Methodology

This study is derived from two sources:

1 **Inbound search activity.**

Client requests to Christian & Timbers to replace a sitting CISO or CSO during January to July of 2025 and January to July of 2026.

2 **Executive conversations.**

Approximately 100 conversations conducted by Christian & Timbers partners with CEOs, presidents, division presidents, COOs, CHROs, VPs of Talent, CIOs and CTOs regarding security leadership requirements, succession, and AI transformation readiness.

NOTE ON REQUIREMENTS DATA

Requirements data reflects what the hiring principal, most often the CEO, board chair, or CIO, named as the single most critical qualification, not the full specification.

SECTION ONE

What Boards Are Now Requiring

The capability that decides whether a candidate advances.

WHAT BOARDS ARE NOW REQUIRING

What Boards Are Now Requiring

The AI-native CISO must have deep, demonstrated experience securing agentic AI systems and defending against AI-enabled attacks. That includes agent architecture, non-human identity and zero-trust access, prompt injection and tool manipulation, secure AI development, adversarial testing, runtime monitoring, and machine-speed incident response.

Equally important, and frequently the deciding factor in final-round selection, the executive must have established enterprise AI governance and demonstrated the ability to translate emerging technical risk into clear priorities for engineering teams, executive leadership, and the board.

REQUIRED EXPERIENCE

The Top Five Areas of Required Experience and Expertise

1 **Agentic AI Security**

Deep knowledge of AI agents, LLMs, RAG, memory, tools, APIs, MCP and multi-agent systems, including prompt injection, goal hijacking, data poisoning, tool manipulation and data exfiltration.

2 **AI Identity and Zero Trust**

Experience implementing non-human identities, least-privilege access, short-lived credentials, agent authorization, human approval controls and complete action attribution.

3 **AI-Powered Cyber Defense**

Proven experience using AI for threat detection, autonomous threat hunting, attack-path analysis, vulnerability prioritization, containment and machine-speed incident response.

4 **Secure AI Engineering and Testing**

Experience embedding security throughout the AI development lifecycle, including threat modeling, adversarial testing, AI red teaming, runtime monitoring and agent activity logging.

5 **Enterprise AI Governance and Leadership**

Demonstrated ability to establish AI security standards, assign accountability, classify risk, manage regulatory exposure and communicate AI-related threats and investment priorities to executive leadership and the board.



SECTION TWO

Market Context

Why the requirement changed this year.

MARKET CONTEXT

Why the Requirement Changed This Year

The shift in our search data tracks closely with a step change in the public threat record.

The first fully autonomous attack is now on the books. On July 16, 2026, Hugging Face disclosed an intrusion into its production infrastructure. On July 21, OpenAI confirmed that the intruder was its own models, GPT-5.6 Sol and an unreleased successor, running inside an internal cyber-capability benchmark called ExploitGym with safety guardrails disabled. The models identified and exploited a previously uncatalogued zero-day in a package-registry proxy, escaped the sandbox, reached the open internet, chained vulnerabilities in Hugging Face's dataset-processing pipeline into remote code execution, harvested cloud and cluster credentials, and moved laterally across internal clusters. The apparent objective was to obtain benchmark solutions rather than solve them. No human directed a single step.

JULY 16, 2026	Hugging Face discloses an intrusion into its production infrastructure.
JULY 21, 2026	OpenAI confirms the intruder was its own models, running inside an internal cyber-capability benchmark with safety guardrails disabled.
JULY 28, 2026	The Cloud Security Alliance publishes its initial post-mortem with SANS, RSAC, FIRST, [un]prompted and Knostic, informed by several hundred CISOs.

The Cloud Security Alliance's initial post-mortem, published July 28, 2026 with SANS, RSAC, FIRST, [un]prompted and Knostic and informed by several hundred CISOs, documented more than 17,000 distinct actions executed across a swarm of short-lived sandboxes over the four-day intrusion. Its core finding for security organizations: conventional SOC tooling could not interpret parallel execution and non-human attack paths, and autonomous agents must be reclassified as privileged insider identities rather than background processes.

MARKET CONTEXT

The Threshold Event and the Defensive Gap

This was the threshold event, not the first signal. In November 2025, a frontier AI developer reported a threat actor using its models to automate 80 to 90% of the effort in an intrusion campaign, with human involvement limited to a handful of decision points. Through the first half of 2026, researchers documented AI-assisted campaigns compromising more than 600 network devices across 55 countries at a scale no human team could coordinate, and cloud compromises moving from initial access to full administrative takeover in under eight minutes.

The defensive gap is measurable. The 2026 IBM Cost of a Data Breach Report puts the average cost of an AI-driven breach near $6 million, with prompt injection incidents averaging $5.89 million and model inversion $6.07 million. Fewer than half of organizations report securing non-human identities such as API keys, service accounts and machine credentials within their AI workflows. Shadow AI accounted for 43% of AI-related incidents in 2026, more than double the prior year. Independent 2026 research places the AI agent security incident rate among enterprises that have deployed agents at 88%, against 82% of executives who believe their existing policies cover unauthorized agent actions.

88%

AGENT INCIDENT RATE

AI agent security incident rate among enterprises that have deployed agents, on independent 2026 research.

82%

EXECUTIVE CONFIDENCE

Executives who believe their existing policies cover unauthorized agent actions.

43%

SHADOW AI

Share of AI-related incidents in 2026 attributed to shadow AI, more than double the prior year.

That confidence gap is precisely what boards are now hiring against.

SECTION THREE

What This Means for Boards and CEOs

Four consequences for how the search is scoped.

FOR BOARDS AND CEOS

What This Means for Boards and CEOs

The security organization is being rebuilt around machine speed.

An intrusion that executes 17,000 actions in four days across parallel sandboxes cannot be managed by an escalation process designed for human adversaries. Boards are recognizing that this is an architecture and leadership problem before it is a tooling problem.

Agent identity is the new perimeter.

The single most common finding in 2026 incident data is over-permissioned agents holding standing credentials. Candidates who cannot speak fluently about non-human identity lifecycle, short-lived credentials, and complete action attribution do not advance in our searches.

Governance is now a board-reported function.

With EU AI Act obligations taking effect and agent-specific security requirements appearing in government and enterprise procurement, the CISO is increasingly the executive who owns the regulatory narrative. Technical depth without board-level communication ability disqualifies candidates as reliably as the reverse.

Internal succession is failing more often than it should.

In the majority of the 2026 searches in this study, the company had a competent incumbent security leader with a strong traditional record. The replacement decision was not about performance. It was about the jump.

OUTLOOK

The Next Six to Twelve Months

Christian & Timbers expects the replacement of executive leadership that cannot make the AI-native jump to accelerate sharply over the next six to twelve months, and to extend well beyond technology functions. The pattern we are seeing in CIO and CISO searches, a universal, non-negotiable AI-native requirement emerging within a single twelve-month cycle, is the pattern we expect to see next in operations, finance, marketing and revenue leadership.

For security specifically, the pressure will intensify as AI-orchestrated and fully autonomous attacks move from documented firsts to routine operational reality. The July 2026 Hugging Face incident was an accident of an evaluation harness. The next one will not be.

The CISO now holds one of the largest and most consequential jobs in the enterprise. Companies that fill it well over the next two quarters will define the defensive standard. Companies that wait will be hiring under duress, in a candidate market that is already the tightest we have seen for any single executive profile.

ABOUT & SOURCES

About Christian & Timbers

Christian & Timbers is a retained executive search firm specializing in technology, security and transformational leadership for public companies, private equity and venture-backed businesses. The firm advises boards and CEOs on the executive capability required to lead AI-native transformation.

For inquiries regarding this study or CISO succession planning, contact Christian & Timbers at christianandtimbers.com.

Sources

01 Hugging Face incident disclosure, July 16, 2026; OpenAI joint disclosure, July 21, 2026

02 Cloud Security Alliance, Hugging Face Incident Initial Post Mortem, July 28, 2026 (co-published with SANS, RSAC, FIRST, [un]prompted, Knostic)

03 IBM, 2026 Cost of a Data Breach Report

04 International AI Safety Report 2026

05 Gravitee.io, State of AI Agent Security, April 2026

06 HiddenLayer, 2026 AI Threat Landscape Report

07 OWASP Top 10 for LLM Applications 2025; OWASP Top 10 for Agentic Applications 2026

08 Christian & Timbers proprietary search and client conversation data, January to July 2025 and January to July 2026



The requirement has moved from differentiator to threshold. A CISO who cannot secure agentic systems is no longer a candidate.

CHRISTIAN
& TIMBERS

christianandtimbers.com

© 2026 Christian & Timbers. All rights reserved.