Open Future Forum

OPEN FUTURE FORUM

CISO AI Leverage Report

NOVEMBER 2026

November 2026 | Edition 5 | Research report | Murray Newlands | Open Future Forum

Prepared 2 November 2026. Evidence cutoff: 31 October 2026, 23:59 America/Los_Angeles. Security-priority and funding results are updated cumulative estimates. Operator access-method results are carried forward. Publication review status appears in the collection audit.

About this publication: Open Future Forum publishes research for executive and investor audiences. Murray Newlands is its founder. This report separates OFF survey observations, external research, author analysis and illustrative scenarios. Vendor-authored guidance is identified as such and does not constitute an independent product assessment. The collection and supporting methods are available through the OFF research index: https://openfutureforum.com/research/

4,468 substantive body words | 5,839 total words
Tables, FAQs, methods and references are excluded from the body count.

In this report

01 The decision in brief
02 Executive summary
03 What changed in this edition
04 What this sample can establish
05 1. Translate concern into a control objective
06 2. Find the budget that owns the control
07 3. Read identity labels as starting points
08 4. Test what survives revocation
09 5. Preserve evidence without collecting everything
10 6. Keep exceptions temporary and actionable
11 External guidance and the boundaries of the comparison
12 A control review the business can use
13 Decisions for the next security review
14 Reader questions
15 Citable facts
16 Methods and limitations
17 References
18 Related research and citation
19 Editable working tool

The decision in brief

The security priority is clear in this sample: 142 of 207 respondents selected securing agents and their access. The control decision is more specific. Before expanding an agent's authority, demonstrate what it can do, how that permission expires, and whether revocation stops downstream actions. A credential inventory alone cannot answer those questions. [SEC-01]

Executive summary

Agent access remains the leading stated security priority in the updated OFF security instrument. It appears in 142 of 207 eligible answers, or 69%. Data leakage appears in 52 answers, shadow AI in 49, attacks using AI in 42 and nothing urgent in 15. Multiple selections are preserved. These results describe priorities among respondents to selective event-registration questions; they are not incident rates or estimates of enterprise-wide exposure. [SEC]

The funding question has the same base but measures something different. Seventy-four respondents selected a dedicated budget line, 68 case-by-case funding, 62 funding from the existing security budget and 22 no AI security spending yet. A respondent can use more than one route. The difference between the prevalence of an access concern and the prevalence of a dedicated budget is not a measured funding shortfall. Existing identity, security and platform budgets may support relevant controls. [SECFUND]

November adds a new security cohort and a more demanding operational question: can the organization prove that authority has ended? The cumulative sample includes repeated respondents, and the additional event's registrations span September and October. The change in the cumulative estimate therefore cannot be interpreted as an independent monthly change in concern. This report keeps the new cohort visible and separates it from the older operator access-method instrument.

The proposed control review follows authority from approval through execution, monitoring, revocation and restoration. It asks the CISO to verify the effective permission at each downstream system, not just the label attached to an agent. A per-agent credential can improve attribution while still granting excessive privilege. Delegated identity can preserve user context while inheriting permissions that exceed the task. Brokered access can narrow credential exposure while concentrating responsibility in the broker.

Too many security gates can slow useful work and encourage unmanaged alternatives. The response should be a small number of reusable control requirements, with evidence proportional to the action's consequence and reversibility. The report supplies an identity lifecycle checklist, a control-review matrix and a worked revocation exercise. Its purpose is to make safe expansion and decisive intervention easier to demonstrate, rather than equating the existence of a policy with the effectiveness of a control.

What changed in this edition

The current security base is 207 for both priorities and funding. The previous October security base was 151. The additional security event contributes 65 eligible answers, nine of which overlap with the previous respondent set. Deduplication therefore produces 207 rather than 216. The event's registration dates run from late September into October; only 56 of the current eligible answers have October registration timestamps. Those timestamps do not establish when respondents last edited their answers. [SEC; SECFUND; change ledger]

Within the additional event cohort, 45 of 65 respondents selected agent access as a priority. That result is consistent with the ordering in the cumulative sample. It does not prove stability across the market: a similar event audience can reproduce a similar priority because the recruitment context is similar. The report consequently uses the new evidence to confirm that access deserves attention within this sample, while avoiding claims about a national trend or a measured rise in attacks.

The first analytical addition is a distinction between credential attribution and effective authority. October's shared-service-account article tested whether an exception was defensible. November examines what happens after any identity model has been approved: permission changes, delegated actions, queued work, downstream sessions and revocation. The question applies to all access models and is not a rewritten argument for one credential architecture.

The second addition is a control-evidence sequence that connects a release decision to an incident response. A team should know before deployment which record will identify the initiating authority, which mechanism will prevent further action and which owner can restore the service. This provides an explicit handoff between the application team, identity team, security operations and the business owner.

The third addition compares current technical and institutional guidance without treating it as survey validation. The MCP specification addresses authorization behavior in its own protocol context; OWASP describes a runtime-control approach; CISA's announcement recommends bounded autonomy and layered oversight; NIST supplies a broader voluntary risk framework. These sources help formulate questions for a deployment review. They do not certify a vendor, establish the incidence of a threat or prove that the OFF respondents have implemented a particular control.

What this sample can establish

The priority and funding questions come from security-oriented event registrations, but an event label does not verify every respondent's job title or responsibilities. Answers are self-reported, and the sample is selective. Invited-only records are excluded. Eligible application and registration statuses do not imply that a person attended the event. Each question is deduplicated separately by normalized respondent identifier, retaining the latest eligible nonblank answer within the cutoff.

The access-method distribution comes from a different operator instrument with 71 answers. Twenty-two respondents selected not applicable yet, leaving an applicable base of 49. Its results are carried forward and should not be described as the access architecture of the 207 security respondents. No respondent-level

linkage between those instruments is used to infer whether a stated concern resulted in a control or a budget.

For the full access distribution, counts are published so that the denominator remains visible. Percentage comparisons among the very small delegated and brokered categories are not used to draw subgroup conclusions. The applicable base is useful for understanding completion and relevance, not for manufacturing precision from small categories.

The new cohort contains enough answers for descriptive reporting, but it overlaps with the cumulative population and comes from a particular event. The research does not estimate compromise rates, test the effectiveness of identity products or determine legal compliance. The control framework below is a recommendation grounded in these limits and current primary guidance; its thresholds must be chosen for the actual system and business consequence.

1. Translate concern into a control objective

The priority distribution can focus a security conversation, but it cannot perform the risk assessment. Agent access is selected by 69% of this sample, compared with 25% for data leakage, 24% for shadow AI and 20% for attacks using AI. These categories can overlap in a single event: an unmanaged agent with excessive access could disclose information after processing hostile content. Counting each concern separately does not establish separate incident populations. [SEC]

The CISO should begin with the action whose consequence matters. Reading a public document, retrieving confidential customer information and changing a production record create different exposures even if the same model or agent framework is used. The control objective should state the permitted action, the protected resource, the initiating authority and the conditions under which execution is allowed. A generic requirement that an agent be secure is too broad to test.

The business owner contributes the reason for the authority. Security can evaluate whether a permission is constrained, but it cannot decide alone whether the business needs the action. If the requested permission exists only because a connector makes it convenient, the team should first test whether a narrower workflow is possible. Removing unnecessary authority can reduce the number of controls and exceptions needed later.

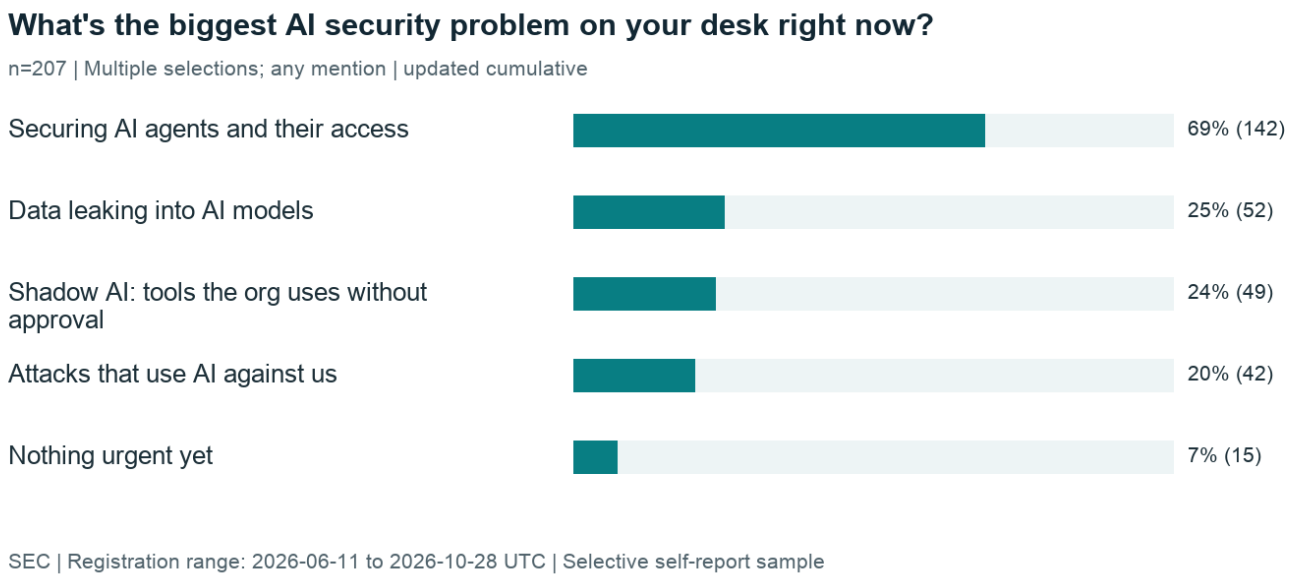
The review also needs a plausible failure condition. A legitimate request might be transformed by an untrusted tool result, a stale document or an ambiguous instruction. A credential might remain usable after the associated workflow has been paused. A queued task might execute after its approver has withdrawn consent. These are scenarios for testing the design, not claims that OFF measured their frequency.

For each scenario, ask which system enforces the boundary. Instructions to a model may help guide behavior, but the decisive permission should be enforced where the resource or action is controlled. The evidence should show the actual rejected request, limited scope or blocked action, rather than only a policy statement. This keeps the security review connected to observable behavior.

The practical output is a concise control objective for each material action class. It should be understandable to the business owner and specific enough for an engineer to demonstrate. Once that objective is agreed,

the CISO can evaluate competing implementations without assuming that one named identity pattern is always sufficient. The priority survey identifies the area of attention; the action-level objective determines what the organization must prove.

Exhibit 1. Security priorities in the updated cumulative sample



Exact selection	Count	Base	Any mention
Securing AI agents and their access	142	207	69%
Data leaking into AI models	52	207	25%
Shadow AI: tools the org uses without approval	49	207	24%
Attacks that use AI against us	42	207	20%
Nothing urgent yet	15	207	7%

Exhibit note: SEC; updated cumulative security instrument through the cutoff. Multiple answers permitted. This is concern prevalence among eligible respondents, not incident prevalence. The added cohort spans September and October registrations.

2. Find the budget that owns the control

The funding answers show several possible routes to implementation. A dedicated line is selected by 74 of 207 respondents, or 36%; case-by-case funding by 68, or 33%; existing security budget by 62, or 30%; and no AI security spend yet by 22, or 11%. Because combinations are allowed, these categories do not partition the sample. A company can have a dedicated initiative and still fund specific controls through its existing identity program. [SECFUND]

The CISO should resist subtracting the dedicated-budget share from the access-concern share and calling the difference underfunding. The questions do not measure required investment, available control capacity or spending adequacy. A company with no separate AI line may already have usable identity, logging and response infrastructure. Another company may have a dedicated budget that is poorly matched to its actual exposures.

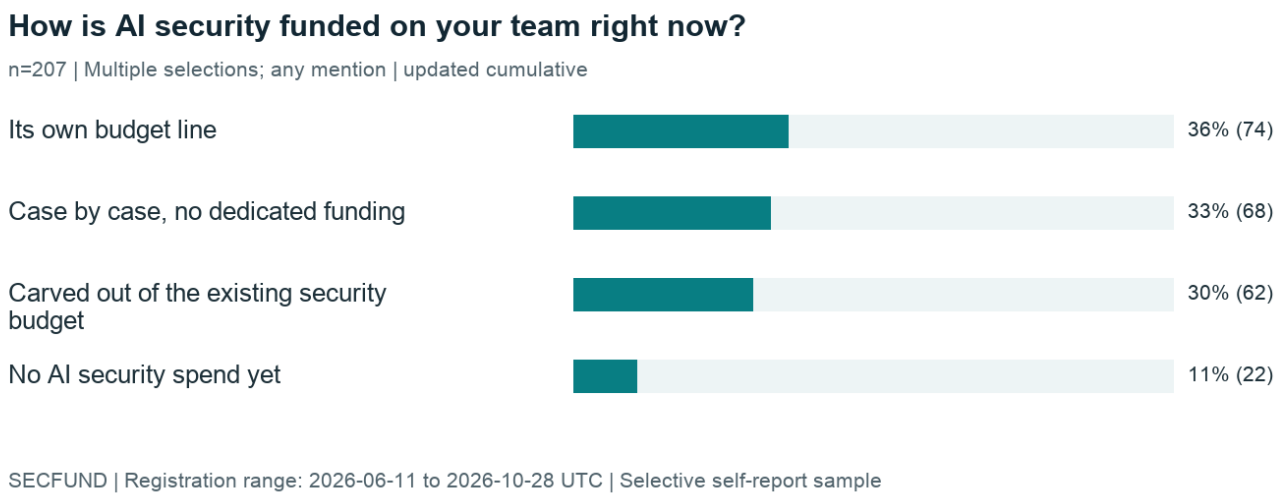
A better funding request starts with the missing capability and the team that must operate it. If an agent cannot be revoked reliably, the proposal should identify the identity work, application changes, downstream dependencies and response testing needed to close that gap. If logs cannot connect an action to an approval, the request should specify the evidence path and the systems that must contribute. This makes the investment review more concrete than asking for an undifferentiated AI security allocation.

Recurring responsibilities matter as much as implementation. Someone must review exceptions, update permissions when the workflow changes, maintain detections and test recovery. A one-time purchase may leave those tasks unfunded. The request should therefore separate setup work from ongoing operation and identify which team absorbs the latter. Capacity already available in an existing program can be recognized without pretending that the work has no cost.

The business owner should also understand the tradeoff. A broader autonomous action may require more expensive controls or a narrower service commitment. A limited approval step may be cheaper and sufficient for a low-volume workflow. The CISO can present alternatives with their operating consequences rather than treating security as a single fixed surcharge on any deployment.

Finally, funding should follow the control's actual owner. The application team may need to change its execution path; the platform team may need to support enforceable policy; security operations may need a response integration. A central budget can sponsor these changes, but it cannot substitute for the teams' delivery commitments. The review should end with named work, evidence of completion and a person responsible for keeping the control effective.

Exhibit 2. Reported AI-security funding routes



Exact selection	Count	Base	Any mention
Its own budget line	74	207	36%
Case by case, no dedicated funding	68	207	33%
Carved out of the existing security budget	62	207	30%
No AI security spend yet	22	207	11%

Exhibit note: SECFUND; updated cumulative security instrument. Multi-select. Funding route does not establish adequacy, effectiveness or the amount spent.

3. Read identity labels as starting points

The carried-forward operator instrument records 24 respondents using per-agent credentials, 18 shared service accounts, four delegated user identity, three credential-free or brokered access and 22 not applicable yet. The full base is 71; the applicable base is 49. The small delegated and brokered categories are reported as counts without comparative claims. This is an architectural description supplied by respondents, not a technical inspection of their systems. [ACCESS]

Per-agent credentials can make it easier to distinguish one agent from another, provided the surrounding logs and credential lifecycle preserve that distinction. They do not, by themselves, prove least privilege, task-level authorization or safe handling of credentials. An individually named identity with broad access can still execute an excessive action. The reviewer must inspect what the credential authorizes and how that authorization changes.

Shared service accounts can make attribution more difficult because several callers may appear under one principal. That does not establish that every shared account is unmanaged or unlawful. A compensating architecture may preserve caller identity, narrow allowed actions and record approvals outside the downstream account. The question is whether that evidence remains trustworthy and whether one caller can exceed or borrow another's authority. October's exception-test article examines that choice in more detail.

Delegated user identity can connect an action to a person, but the presence of a user identity is not the same as informed approval of every later action. The system needs to distinguish the user's general permissions from the specific task delegated to the agent. It should also account for permission changes, user departure and the expiry of the original request. Otherwise, a valid identity can conceal authority that is stale or broader than intended.

Brokered access can keep raw credentials away from an agent and centralize enforcement. That shifts important questions to the broker: how it authenticates callers, evaluates policy, records decisions and limits downstream requests. The phrase credential-free should not be read as authority-free. Some component still mediates access to protected resources, and its failure or misconfiguration can affect multiple workflows.

The comparison therefore belongs at the level of control properties. Can the system attribute an action, enforce a narrow permission, expire authority, detect abnormal use and stop execution? Different

architectures may satisfy those properties differently. The CISO should ask for demonstrations against the same material actions, so that a product label does not determine the outcome of the review before the evidence is examined.

Exhibit 3. Access methods reported by a separate operator sample

Selection	Count	Full base	Applicable base
Per-agent credentials	24	71	49
Shared service accounts	18	71	49
Delegated user identity	4	71	49
Credential-free or brokered access	3	71	49
Not applicable yet	22	71	Not included

Exhibit note: ACCESS; carried-forward operator instrument, single selection. Applicable base excludes 22 not-applicable answers. Small-category percentage comparisons are suppressed. These are not the access methods of the 207-person security sample.

4. Test what survives revocation

Revocation is often described as an identity operation, but a material agent workflow can continue through several downstream mechanisms. An access token may remain valid until expiry. A task may already be queued. A downstream service may have created its own session or scheduled action. A copied dataset may remain available after the original connection is removed. The review needs to identify those possibilities in the actual architecture rather than assume that disabling one identity stops everything.

Begin with a map of the authority path for one consequential action. Identify the initiating request, approval, agent execution, tool invocation and resource change. Record where authorization is checked and where work can persist independently. This is a scoped control map, not an instruction to document every internal function in the software. Its purpose is to expose the points at which a stop decision might fail to propagate.

The test should distinguish stopping future requests from reversing completed actions. Revocation may prevent another transaction while leaving a completed record change intact. Some actions can be rolled back; others require compensation, notification or manual correction. The business owner should understand that distinction before granting autonomous authority, because the acceptable level of oversight depends partly on reversibility.

Test evidence should include timing. Record when the stop decision was made, when each enforcement point changed state and whether any further action succeeded. The organization should choose an acceptable interval based on consequence and operating needs. A universal response-time claim would be misleading: a read-only assistant and an agent capable of altering critical records do not have the same tolerance.

Restoration deserves equal attention. A team that can stop an agent but cannot safely resume service may face pressure to bypass the control during an incident. Define who verifies remediation, which credentials or sessions are replaced, what backlog is reviewed and which authority is restored first. A staged restoration may be appropriate where the cause is uncertain or queued actions need examination.

The November contribution is the evidence of authority ending, including the limits of that evidence. A successful test should record the specific components and version tested. A connector change, new tool or broader scope can invalidate the result. The CISO should make retesting part of material change review, so that a control demonstrated during the pilot does not become an unsupported assurance after the workflow expands.

5. Preserve evidence without collecting everything

Security operations needs enough information to reconstruct consequential activity, but indiscriminate logging can create its own privacy, confidentiality and operational burdens. A useful evidence plan begins with the question the organization may need to answer: who authorized the action, what scope applied, which system enforced it, what occurred and whether intervention worked. The plan should then identify the records needed for those questions.

The initiating request and the executed action may differ. An agent can translate a broad request into several tool calls, and a tool can trigger additional downstream behavior. Logs should preserve the relationship between the approved task and the material actions without treating every internal model token as necessary evidence. The application team should be able to show how an investigator would follow that relationship across systems.

Approval records need context. A timestamp and an approver name are weak evidence if the approver could not see the relevant data, action or consequence. The review should examine what was presented for approval and whether the eventual execution stayed within that scope. If the system changes a material parameter after approval, the organization needs a rule for renewed review or rejection.

Retention should follow the organization's legal and operational requirements. The GC and privacy team should determine applicable obligations and preservation needs; the CISO should ensure that the architecture can meet them. Access to sensitive evidence should be restricted, and routine telemetry should avoid unnecessary raw confidential content where a less revealing record supports the control objective. This is a design recommendation, not a claim that one retention period fits every jurisdiction or workflow.

Evidence quality also depends on time and identity consistency. If systems cannot correlate records, an incident team may see several valid logs without a reliable sequence. If a shared principal obscures the initiating caller, the team may need additional trustworthy attribution. The review should test an actual reconstruction using the records available to responders, rather than relying only on a diagram that assumes all information can be joined.

The final requirement is ownership. Someone must maintain the evidence path when the application, connector or identity provider changes. Security operations should know whom to contact when a record is missing. A control that depends on a log no team operates is fragile even if the initial demonstration

succeeds. The evidence plan should therefore be part of the release record and the material-change checklist.

6. Keep exceptions temporary and actionable

A production organization will encounter systems that cannot immediately support the preferred control. An exception may be justified by a limited business need, a constrained action set and credible compensating measures. It should not become a permanent category called legacy AI. The decision needs a specific scope, an accountable owner, an expiry and a condition that triggers earlier review.

The CISO should distinguish acceptance of residual risk from acceptance of missing information. If the team does not know what the agent can access, it cannot describe the residual exposure credibly. A bounded discovery period may be appropriate, but it should not silently authorize consequential production activity. Unknown scope is a reason to narrow the deployment while the facts are established.

The exception record should identify what will close the gap. That might be an application change, a narrower connector, a stronger approval mechanism or a different operating model. A statement that the vendor will improve the product later is insufficient unless the organization has an interim limit and a decision date. Procurement and the business owner should understand the dependency rather than leaving security to carry an open-ended promise.

The review cadence should reflect change and consequence. A stable, low-impact read workflow may need less frequent attention than an agent whose tools and permissions change weekly. Material events should trigger review even when the scheduled date is distant: a new data class, external destination, financial action, customer population or downstream integration can alter the original risk decision.

Exceptions should also connect to incident readiness. The response team needs to know which controls are weaker, which alternative evidence exists and whether a stop action has known limitations. Hiding exceptions in a separate governance file can leave responders assuming a control is present when it is not. The operating record should expose the relevant limitation without disclosing unnecessary sensitive detail to every user.

The goal is a smaller, better-understood set of exceptions. Report their age, scope changes, remediation progress and failed closure dates to the appropriate executive. Counts alone can mislead: ten narrow exceptions may create less exposure than one broad, untested permission. The CISO's report should explain the material decisions and required resources, so that management can act instead of merely observing a rising total.

External guidance and the boundaries of the comparison

The MCP authorization security-considerations document for specification 2026-07-28 requires validation that tokens are intended for the receiving server and prohibits passing the client's token through to upstream APIs. [E-MCP] This is a concrete protocol requirement, not a general certification that any MCP deployment is secure. The November control review uses it to illustrate why a valid credential must be evaluated in context: the intended resource and authorization path matter, as well as possession of the token.

OWASP's September 2026 description of the Agent Control Standard presents middleware hooks and runtime policy enforcement as a basis for portable controls. [E-OWASP-ACS] That provides a useful implementation direction for asking where an action can be intercepted or denied. It does not establish that a particular platform implements the standard completely, that every downstream action passes through the hook, or that a deployed policy matches the business owner's intended scope. Those remain demonstration questions.

CISA's May 2026 joint-guidance announcement recommends limiting agent autonomy, using layered defenses and maintaining strong identity and oversight. [E-CISA] The announcement supports a bounded deployment approach, but it is not an incident survey. This report relies on the accessible agency announcement for those high-level recommendations; detailed provisions of the linked guidance were not independently retrieved and are not represented as reviewed here.

NIST's AI Risk Management Framework organizes voluntary risk work around Govern, Map, Measure and Manage. [E-NIST] The framework is useful for placing technical tests inside an accountable organizational process. A revocation demonstration without an owner who can order it is incomplete; a governance policy without evidence that it works is equally incomplete. This report's lifecycle record connects the two. The record is an OFF recommendation and is not presented as a NIST-prescribed template.

These sources address different layers, which is why agreement among them should not be overstated. A protocol specification governs particular implementations. A runtime-control approach concerns enforcement integration. Agency guidance frames deployment risk. A voluntary framework organizes management responsibilities. None measures whether the respondents selecting agent access have implemented adequate controls, and none turns a concern percentage into an investment requirement.

The practical comparison is therefore a set of questions for the actual deployment. Does its protocol behavior meet the applicable specification? Can its runtime enforce the intended policy? Is the autonomy bounded by the business need? Can management demonstrate ownership, measurement and intervention? Answering those questions with evidence is more useful than listing several respected frameworks on a slide and assuming that the list establishes readiness.

Source	What it contributes	What the company must still prove
MCP 2026-07-28 authorization considerations	Specific token-handling requirements	Its actual implementation and downstream authorization path
OWASP ACS description, September 2026	Runtime enforcement approach	Coverage, policy behavior and limits in the deployed system

Source	What it contributes	What the company must still prove
CISA joint-guidance announcement, May 2026	Bounded autonomy and layered oversight	The control design for the particular action
NIST AI RMF	Voluntary management framework	Effective ownership, testing and intervention

A control review the business can use

Select a material agent workflow and define its consequential actions. Use the lifecycle record below to connect the initial approval to current permissions and the evidence needed for intervention. The application owner should prepare the record; identity and security teams should verify their parts; the business owner should confirm the permitted purpose and acceptable service consequence. The GC joins where contractual, data-use or preservation obligations affect the design.

Working tool: agent identity lifecycle checklist

Stage	Required record	Owner	Demonstration / failure condition
Request	Purpose, business owner, resources and actions	Business and application owners	Reject unexplained or unbounded authority
Provision	Principal, scopes, expiry, downstream dependencies	Identity / platform team	Demonstrate effective permissions
Release	Approved version, tests, oversight and exceptions	Application owner; security review	Hold consequential actions with untested boundaries
Operate	Action attribution, alerts, permission changes	Application and security operations	Investigate missing evidence or scope drift
Revoke	Stop mechanism, downstream propagation and timing	Identity / incident owner	Test queued and persistent authority
Restore / retire	Remediation, backlog handling, credential closure	Business and application owners	Do not restore unsupported scopes

An illustrative exercise concerns an agent that reads supplier correspondence and prepares changes to vendor records. Its approved scope permits draft preparation; a human must approve a bank-detail change before the system writes it. The exercise introduces an untrusted instruction inside an incoming document, asking the agent to bypass the approval step. This is a defensive test scenario, not an observed OFF incident or a report about any named company.

The first test asks whether the write can occur without valid approval. The expected evidence is an enforced rejection at the relevant action boundary, together with a record linking the attempted change to the initiating

task. A reassuring explanation generated by the agent is not sufficient. The business owner needs to know whether the protected action occurred, and the investigator needs evidence that can be examined independently of the agent's description.

The second test withdraws the agent's authority while work is in progress. The team checks new calls, queued changes and any downstream session that could continue independently. It records which actions stop immediately, which remain pending and which require a separate control. If a queued write can still execute, the exercise has identified a specific remediation need. It has not demonstrated successful end-to-end revocation merely because the main credential was disabled.

The third test concerns restoration. Before returning the workflow to service, the owner confirms that pending changes are reviewed, the cause of the failure is addressed and the permitted scope is restored deliberately. The team records any temporary limitation, such as read-only operation. This makes the exercise useful to operations: the output is a recoverable service with known boundaries, rather than a security finding disconnected from the work.

Working tool: evidence-based control review

Review question	Evidence input	Decision owner	Review trigger
Can the action exceed the approved purpose?	Negative tests and actual permissions	Business owner with security	New tool, scope or data class
Is approval bound to the executed action?	Approval display and execution record	Application owner	Changed parameters or approval flow
Can authority be ended?	Timed revocation test across dependencies	Identity and incident owners	Connector or session change
Can the event be reconstructed?	Correlated records available to responders	Security operations	Missing or changed telemetry
Is the exception still justified?	Scope, expiry and remediation evidence	Risk owner	Expiry, incident or material expansion

The release decision should record unresolved gaps and their consequences. A narrow read-only deployment may proceed while a write-path control is improved, provided the system actually enforces the narrower scope. A consequential action with no reliable intervention path should remain constrained. These are recommended decision principles, not universal legal requirements or survey-derived thresholds.

After the exercise, give each remediation a delivery owner and a retest condition. Close the issue only when the relevant demonstration succeeds or an authorized executive accepts a clearly described residual risk within the company's process. Preserve the tested version and scope so that later teams can determine whether the evidence still applies. The next change review can then build on a known result rather than restarting from a generic questionnaire.

Decisions for the next security review

The CISO should first identify which material agents can take actions that are difficult to reverse. Ask the business owner to justify each action and identify the system that enforces its limit. A complete inventory is useful, but the first review should concentrate on consequential authority rather than treating every assistant as equivalent.

Second, choose a revocation demonstration that crosses the relevant downstream systems. Require a clear record of what stopped, what remained possible and who can correct the gap. The identity team, application owner and incident lead should agree on the result. A partial stop should be described as partial, with its operational consequence visible.

Third, connect the remediation to an actual budget and delivery commitment. The survey does not establish that a separate AI-security budget is necessary. It does show several funding routes in use. The company should choose the route that can deliver and sustain the needed control, with an owner responsible for the result.

The test of November readiness is whether an executive can ask a specific question about an agent's authority and receive evidence rather than a product label. Which actions are permitted? Where are they constrained? How does authority end? Who can intervene and restore service? Clear answers make expansion more defensible and make a pause less dependent on improvisation.

Reader questions

Does 69% mean 69% have experienced an agent-access incident?

No. It is the share selecting agent access as a security priority: 142 of 207 eligible respondents, multi-select. [SEC-01]

Is the November increase in sample size a monthly trend?

No. The new event cohort spans September and October and overlaps with previous respondents. The cumulative base is deduplicated, not an independent monthly panel.

Are shared service accounts automatically unacceptable?

The survey cannot establish that. Review attribution, scope, enforcement, revocation and compensating controls for the actual workflow. The architecture label alone does not prove effectiveness or noncompliance.

Does a dedicated budget prove adequate security investment?

No. Funding route and control effectiveness are different measures. Existing programs may fund relevant capabilities; a dedicated line may still leave important work incomplete. [SECFUND]

What does the access-method applicable base mean?

Forty-nine of 71 operator respondents selected an access method; 22 selected not applicable yet. This separate sample is carried forward. [ACCESS]

Why test queued work during revocation?

Because a stop at the initiating identity may not cancel work already accepted downstream. The test must establish the actual system's behavior; it should not assume propagation.

Does following a framework certify an agent?

No certification is claimed here. The referenced sources provide requirements or guidance within their own scope; the deployment still needs implementation evidence.

Citable facts

- Agent access: 142/207, 69%, updated cumulative security-priority sample, multi-select. [SEC-01]

- Dedicated security budget line: 74/207, 36%, reported funding route, multi-select. [SECFUND-01]

- Access-method question: 71 answers, including 22 not applicable; applicable base 49, carried forward. [ACCESS]

Methods and limitations

Calculation-v1 uses authorized exports dated 1 November 2026 and the cutoff stated above. Byte-identical exports and invited-only records are excluded. Latest eligible nonblank answers are retained per normalized respondent identifier per question. Registration timestamps order records but do not verify answer-edit times. The source manifest records file-level provenance, dates and duplicate status.

SEC and SECFUND are multi-select, n=207 each. The added event has 65 eligible respondents; nine overlap with the previous 151-person base. Its separate results are descriptive cohort context. ACCESS is a carried-forward single-selection operator question, n=71, applicable n=49. The instruments are not joined to infer individual concern, spending or control relationships. Small access categories are shown as counts without subgroup percentage comparisons.

No incident prevalence, market representativeness, causal effect, product ranking or legal conclusion is inferred. Current technical guidance is cited within its stated scope. The accessible CISA announcement was reviewed; the linked full guidance could not be retrieved, so detailed claims from that document are excluded. No private respondent quotations or identifying free text are published.

References

[OFF-DATA] Open Future Forum. Executive AI Statistics, November 2026; aggregate data, source manifest and change ledger.

[E-MCP] Model Context Protocol. Authorization Security Considerations, specification 2026-07-28. https://github.com/modelcontextprotocol/modelcontextprotocol/blob/main/docs/specification/2026-07-28/basic/authorization/security-considerations.mdx

[E-OWASP-ACS] OWASP Gen AI Security Project. Agent Control Standard (ACS), 1 September 2026. Official description. https://genai.owasp.org/resource/agent-control-standard-acs/

[E-CISA] CISA and partners. Guidance on Adopting Agentic AI Services, official announcement, 1 May 2026. https://content.govdelivery.com/accounts/USDHSCISA/bulletins/41544ff

[E-NIST] NIST. AI Risk Management Framework 1.0; official resource page consulted 2 November 2026. The page notes that a revision is in progress. https://airc.nist.gov/airmf-resources/airmf/

[OFF-EXCEPTION] Open Future Forum. AI Agent Shared Service Account Exception Test, October 2026. https://openfutureforum.com/blog/ai-agent-shared-service-account-exception-test

Related research and citation

The November AI Leaders report covers release and production evidence; the General Counsel report covers authorization and legal records; the Board Preview covers materiality and escalation. They are complementary analyses, not independent surveys confirming this report.

Suggested citation: Newlands, Murray. CISO AI Leverage Report, November 2026. Open Future Forum. Evidence cutoff 31 October 2026.

Open Future Forum convenes executive discussions on enterprise AI. Research and methods: https://openfutureforum.com/research/

Editable working tool

Authority termination test: download the editable CSV worksheet. The worksheet implements the report's recommendation; it is not a validated score, mandatory standard or completed assessment.