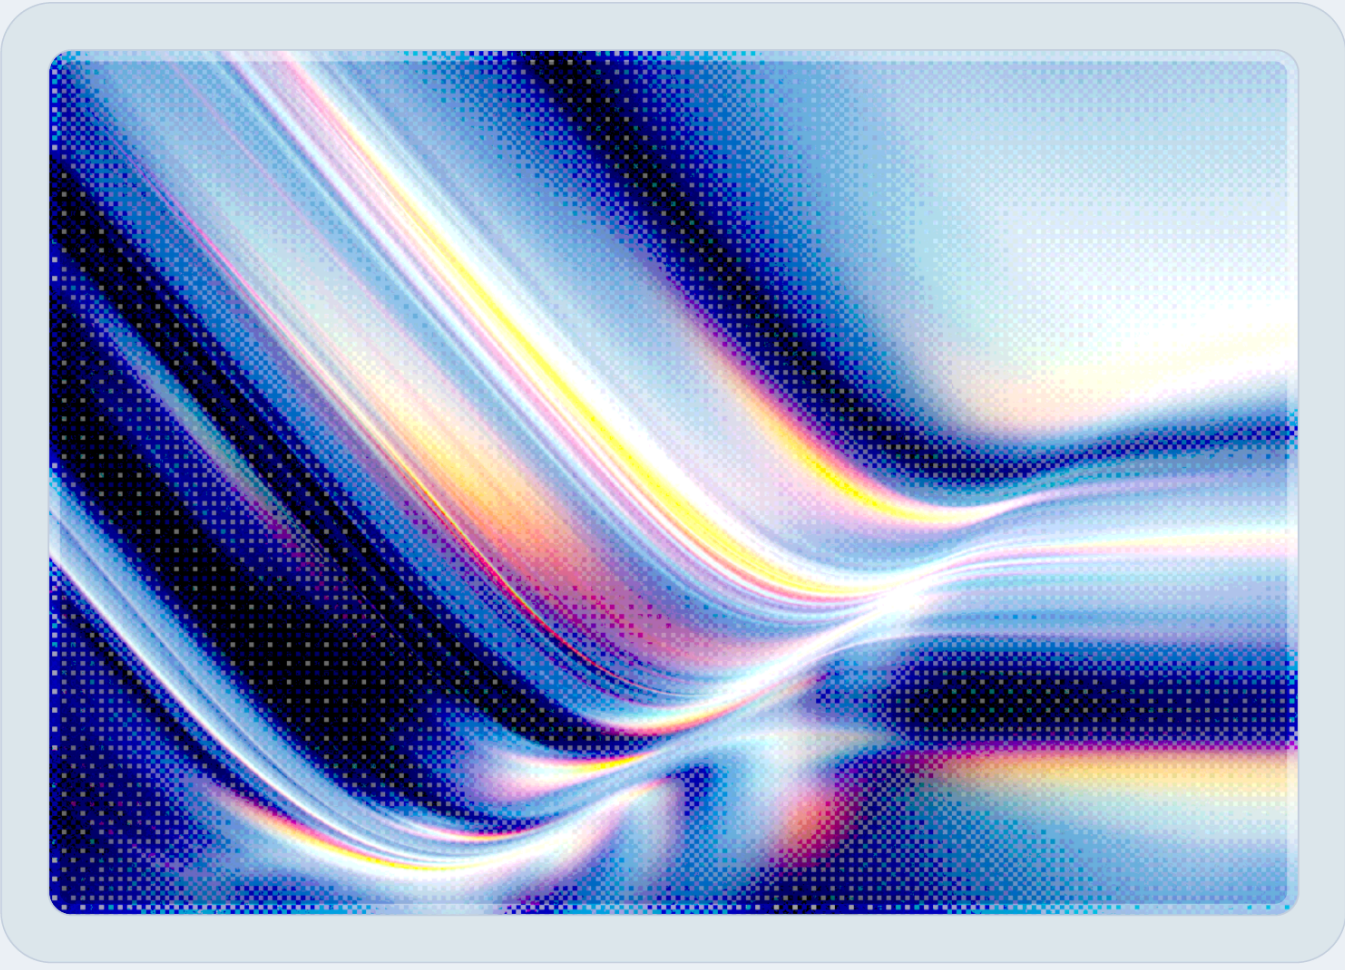
CSAI Foundation | Cloud Security Alliance

NIS2 Article 20: Board Liability for AI Agent Access Decisions

Why management bodies now own the identity and access posture of autonomous agents

2026-11-04

AI-assisted Rapid Research



CSAI

cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

Article 20 of the NIS2 Directive requires the management bodies of essential and important entities to approve the cybersecurity risk-management measures their organizations take under Article 21 and to oversee how those measures are carried out. It also requires Member States to ensure that management bodies can be held liable for the entity's infringements of Article 21 [1]. Article 21(2)(i) lists "access control policies" among those minimum measures [1]. The Directive does not mention artificial intelligence. Even so, once an organization lets AI agents read data, call APIs, or change systems on its behalf, the decision about what those agents can reach becomes an access control decision, and so falls under the policy the board has approved. This is our reading of the text, not something a regulator has stated. No supervisory authority has yet published guidance or an enforcement decision on AI agents under Article 20.

CSA's 2026 survey on AI agent identity and access, which covers organizations generally rather than NIS2-regulated entities specifically, suggests a substantial gap between that duty and current practice. It found that 43% of organizations use shared service accounts for agents and 31% let agents run under a human user's identity. It also found that 68% cannot clearly tell human activity from agent activity [7]. Where Commission Implementing Regulation (EU) 2024/2690 applies, it requires unique identities and individually assigned access rights, and it allows shared identities only through an explicit approval process [2]. Agent deployments built on borrowed or shared credentials may therefore be hard to defend as compliant when a board is asked to show that it oversaw them.

The practical conclusion is that boards in scope of NIS2 should treat agent access as a named item in the access control policy they approve. They should ask for regular reports on agent identities, privileges, and revocation capability, and they should make sure the training Article 20(2) requires covers delegated autonomy, not only traditional threats. Enforcement against individuals depends on national law, which still differs across the Union. Germany, for example, sends management liability primarily through ordinary company law [4], and in July 2026 the Commission referred four Member States to the Court of Justice of the EU for failing to fully transpose the Directive [5].

Background

NIS2 (Directive (EU) 2022/2555) replaced the original Network and Information Security Directive. It widened the range of covered sectors and introduced a two-tier model of "essential" and "important" entities. Member States had to transpose it by 17 October 2024 [1]. Article 20 is the governance provision. Paragraph 1 requires management bodies to "approve the cybersecurity risk-management measures taken by those entities in order to comply with Article 21, oversee its implementation and can be held liable for infringements by the entities of that Article" [1]. Paragraph 2 requires members of management bodies to follow training so they can identify risks and assess risk-management practices and their effect on the entity's services [1]. Together, these provisions make cybersecurity risk management a governance duty of the management body that cannot simply be delegated to the CISO. How far that duty translates into personal liability for individual directors depends on national law, as discussed below.

The Directive backs that duty with enforcement tools aimed at individuals as well as entities. For essential entities, Article 32(5) lets authorities ask a court or other competent body to temporarily bar a person with managerial responsibility at chief executive or legal representative level from exercising managerial functions, if earlier enforcement measures have failed [1]. Article 32(6) requires Member States to ensure that natural persons who represent or control an essential entity can be held liable for breaching their duty to ensure compliance [1]. At entity level, Article 34 requires Member States to set maximum administrative fines for Article 21 infringements of at least EUR 10 million or 2% of total worldwide annual turnover, whichever is higher, for essential entities, and at least EUR 7 million or 1.4% for important entities [1]. These figures are floors for the national ceilings, not minimum fines; actual fines are set by national authorities within those ceilings.

Article 21 sets out what the board approves. It requires "appropriate and proportionate technical, operational and organisational measures" based on an all-hazards approach. It lists minimum measures that include supply chain security (point d), "human resources security, access control policies and asset management" (point i), and multi-factor or continuous authentication where appropriate (point j) [1]. Commission Implementing Regulation (EU) 2024/2690 turns these measures into detailed technical requirements for digital infrastructure, ICT service management, and digital provider entities such as cloud computing providers, data centre operators, and managed security service providers [2]. ENISA's NIS2 Technical Implementation Guidance, published in June 2025, maps those requirements to examples of evidence [3]. Section 11 of the Regulation's annex covers access control. It addresses need-to-know and least-privilege assignment of rights, revocation when roles change, review of rights at planned intervals, dedicated privileged accounts, unique identification, and authentication strength scaled to the classification of the asset [2].

Transposition has been uneven, and that affects how much exposure any individual director actually faces. On 8 July 2026, the European Commission referred Ireland, Spain, France, and the Netherlands to the Court of Justice of the EU for failing to fully transpose the Directive [5]. A referral is an escalation step, so it does not establish that every other Member State had completed transposition by that date. Germany's implementing act entered into force in December 2025. Section 38 of the revised BSI Act (BSIG) obliges management to implement and oversee the required risk-management measures. It makes management liable to the entity for culpably caused damage under the company-law rules for the entity's legal form, and it requires management to attend training regularly [4]. Separately, on 20 January 2026 the Commission proposed targeted amendments to NIS2 as part of a wider cybersecurity package, aimed mainly at simplification and harmonizing technical measures [6]. Public analyses of the proposal do not indicate changes to Article 20, but the legislative process was still under way when this note was written.

Security Analysis

Agent access is an Article 21 access control decision

The central analytical point is a matter of categories. An AI agent that retrieves records from a CRM, files tickets, changes cloud configuration, or triggers payments is exercising access rights over network and information systems. How those rights are granted, scoped, reviewed, and revoked falls under the access control policies of Article 21(2)(i), and the identity under which the agent authenticates falls under point (j) and, where the Implementing Regulation applies, Sections 11.5 and 11.6 of its annex [1][2]. Nothing in the Directive exempts automated actors, and machine identities fall within the scope of network and information system security. What is different about agents is that they often combine broad, dynamically used access with behavior that depends on instructions and inputs, which raises the stakes of getting the policy wrong.

That framing has a direct governance consequence. If agent access sits inside the access control policy, it is arguably covered by the board's approval, whether or not anyone told the board that agents were in use. A management body that approved an access control policy written for human users and conventional service accounts, and was never told that agents were running with inherited human privileges, may find it hard to show that it "oversaw" implementation in the sense Article 20(1) requires. We cannot predict how courts and authorities will read "oversee" in practice. However, the language suggests that a board cannot discharge the duty by approving a policy once and not asking how it applies to new classes of actor.

Where current practice diverges from the regulatory baseline

CSA's survey data shows how far current agent deployments sit from the baseline the Implementing Regulation describes. The survey, conducted online in January 2026 and commissioned by Aembit, gathered responses from 228 IT and security professionals; several of its identity questions allowed multiple answers, so the percentages for a given question can sum to more than 100% [7]. The table below compares selected requirements from the Regulation's annex with the survey findings on AI agent identity and access [2][7].

Requirement area (CIR 2024/2690, Annex)	Regulatory expectation [2]	Observed AI agent practice [7]
Identification (§11.5)	Unique identities; identities linked to a single person; shared identities only where necessary, with explicit approval and documentation	43% use shared or generic service accounts for agents; 31% let agents operate under a human user's identity; 52% use workload identities (multiple answers permitted)
Management of access rights (§11.2)	Rights granted on need-to-know and least privilege; reviewed at planned intervals	74% agree agents often receive more access than necessary; 52% report agents at least sometimes inherit access intended for humans or other systems
Monitoring and logging (§3.2; identity logging in §11.5)	Logging of activity on network and information systems, and of identity management	68% cannot clearly distinguish human from AI agent activity
Ownership and consistency (CSA interpretation of §11.1)	Access control policy applied consistently across the actors it covers	Only 22% apply access frameworks very consistently to agents; agent identity is owned by security in 28% of organizations, development or engineering in 21%, IT in 19%, and IAM teams in 9%
Credential life cycle (§11.6)	Managed authentication credentials, scaled to asset	One-third are unsure how often agent credentials are rotated

Requirement area (CIR 2024/2690, Annex)	Regulatory expectation [2]	Observed AI agent practice [7]
	classification	

Table note: The annex requirements on identification are written primarily with human users in mind. Applying them to AI agents, including the expectation that each identity trace to an accountable person, reflects the authors' reading rather than explicit regulatory text.

Two further caveats apply. First, the Implementing Regulation binds only the digital-sector entity types it names. For other essential and important entities, national authorities may draw on it as a reference point for what "appropriate and proportionate" means, but it is not directly binding. Second, the survey describes organizations in general, not NIS2-regulated entities specifically, so it shows a likely pattern rather than a measured compliance rate. Taken together, the evidence still points toward a structural mismatch. The identification model in the Regulation assumes that every identity traces to an accountable person, while the most common agent deployment patterns break that link. The survey also found that organizations lean on governance-based, procedural safeguards to make up for missing identity controls [7], a point that matters for Article 20 because such safeguards provide evidence of oversight only if the board actually reviews them.

These findings extend an older problem. CSA's earlier survey on non-human identity security reported that only 20% of organizations had formal processes for offboarding and revoking API keys. Respondents named lack of credential rotation (45%), inadequate monitoring (37%), and overprivileged accounts (37%) as common causes of non-human identity incidents [8]. Agents inherit all of these weaknesses and add one more. An agent's actual behavior depends on its instructions and inputs, so permissions that look safe at provisioning time can be exercised in ways nobody anticipated.

Why the liability question lands on agent access specifically

Several features of agent deployments make them a likely focus in a post-incident review of whether the board met its obligations. Agents are often deployed by business units or engineering teams outside formal identity governance, and the survey found that ownership of agent identity is fragmented across security, engineering, IT, and IAM teams, with no function owning it in a majority of organizations [7]. That pattern resembles the "shadow IT" problem, on which the asset management obligations in Article 21(2)(i) bear directly. Agents also tend to sit at points of concentrated access. One agent may hold tokens for several SaaS applications, internal APIs, and cloud infrastructure, so a single compromised or misdirected agent can cause an incident that is significant under Article 23 reporting rules [1].

Inherited human privileges blur the evidence trail. When an agent acts under an employee's identity, logs attribute the action to that person. This complicates incident reconstruction, and it may obscure whether access control policy failed or a person misbehaved. If an agent operating under an employee's credentials exfiltrated records, for instance, the initial investigation could focus on the employee rather than the agent. The Article 23 incident reporting process, which runs on 24-hour, 72-hour, and one-month timelines [1], requires entities to understand what happened quickly. An entity that cannot tell agent actions from human ones may struggle to meet those deadlines, and that difficulty could itself become evidence of weak oversight.

Finally, the training obligation in Article 20(2) is relevant. Board members must gain "sufficient knowledge and skills" to identify risks and assess risk-management practices [1]. Germany's Section 38(3) BSIG requires regular training with a similar aim [4]. As agents become operationally significant, training that leaves out delegated machine autonomy may not give board members the knowledge the provision contemplates. Whether a regulator would treat that omission as a deficiency is an open question.

Limits of the liability exposure

The analysis should not overstate personal exposure. Article 20(1) works through national law, and Member States have chosen different mechanisms. Germany's Section 38(2) BSIG makes management liable to the entity, not directly to third parties, for culpably caused damage under the company-law rules for the entity's legal form; only where company law contains no applicable liability rule does the BSIG supply one itself [4]. In practice, that means directors' existing duty of care, now framed by an explicit statutory oversight obligation. The power in Article 32(5) to bar managers temporarily applies only to essential entities, only after earlier enforcement measures have failed, and only until the entity corrects the deficiency [1]. How Member States implement the individual-liability provision in Article 32(6) will vary as well. As of this writing we found no published enforcement action that ties Article 20 liability to AI agent access. Our conclusion is therefore a forecast of where scrutiny is likely to fall, not a description of settled enforcement practice.

Recommendations

Immediate Actions

Management bodies of entities in scope should ask, as a formal agenda item, for an inventory of AI agents with production access. For each agent, the inventory should list the identity it authenticates under, the systems it can reach, and its accountable business owner. Many organizations will find that no single

function can produce this inventory, which is itself a finding worth recording [7]. In parallel, the CISO should confirm that agents are explicitly covered by the access control policy the board has approved under Article 21(2)(i). If they are not, the CISO should prepare an amendment for board approval, so the minutes show the board acted on agent access rather than being unaware of it.

Organizations should also identify any agent operating under a named human user's credentials or a shared service account. Each such case should either be moved to a dedicated identity or documented as an approved exception with a business justification, matching the exception process the Implementing Regulation describes for shared identities [2]. Finally, security teams should confirm they can revoke every production agent's credentials quickly. Disabling identities or revoking tokens is the most common containment action organizations report for agents [7], and it works only when agents have distinct, revocable identities.

Short-Term Mitigations

As a suggested planning horizon, over roughly the next two quarters or by the next board reporting cycle, organizations should give agents first-class identities that are distinct from human users and from generic service accounts. Entitlements should be scoped to the agent's task, not inherited from the person who deployed it. Agent access rights should go into the same periodic review cycle used for privileged human access, with review intervals set by policy as Section 11.2 of the Implementing Regulation requires [2]. Logging should tag agent-initiated actions so they can be separated from human activity during incident reporting, which directly addresses the attribution gap found in the survey [7].

Reporting to the board should include a small set of agent-specific indicators. Suitable measures include the number of production agents, the share running on dedicated identities, the share with least-privilege entitlements reviewed in the current cycle, and the time needed to revoke an agent's access. The next round of Article 20(2) management training should include a module on agentic AI risk, covering delegated authority, privilege inheritance, and prompt-driven misuse of legitimate access.

Strategic Considerations

In the longer term, organizations should assign a single accountable owner for non-human and agent identity. Article 20 places accountability at the top, and that accountability is easier to discharge when there is one clear line of responsibility underneath it. Entities that procure agentic capabilities from third parties should also treat agent access as a supply chain matter under Article 21(2)(d). Contracts should cover how vendor-operated agents authenticate into the entity's systems and how their access can be suspended [1].

Legal and compliance teams should track two moving targets. The first is national transposition and supervisory guidance in each Member State where the entity operates, since the mechanics of individual liability vary [4][5]. The second is the progress of the Commission's January 2026 NIS2 amendment proposal [6]. Organizations that are also subject to the EU AI Act should map where its governance requirements, including the accuracy, robustness, and cybersecurity requirements for high-risk AI systems in Article 15 [11], overlap with NIS2 access control obligations, so that agent oversight is evidenced once and reused across both regimes.

CSA Resource Alignment

CSA's survey report *Identity and Access Gaps in the Age of Autonomous AI* [7] provides the main evidence base for this note. It documents the practices most likely to put boards at odds with NIS2 access control obligations: shared and borrowed identities, inherited privileges, fragmented ownership, and weak attribution of agent activity. Its finding that organizations rely on procedural safeguards to make up for missing identity controls is especially relevant to Article 20. Procedural safeguards that the board never reviews give little evidence of oversight.

CSA's *State of Non-Human Identity Security Survey Report* [8] provides the baseline for credential life cycle failures, such as weak key offboarding and lack of rotation, that agents inherit. It is useful for entities building their Section 11 evidence, because the same controls that fix service-account hygiene form the foundation for agent identity.

The *AI Controls Matrix (AICM) v1.1* [9] provides a control framework for turning board-level policy into auditable controls. Its Identity and Access Management domain and its Governance, Risk, and Compliance domain can be mapped to the Article 21(2)(i) and (j) obligations and the Section 11 requirements. Its shared-responsibility model across model providers, orchestrated service providers, application providers, and AI customers helps clarify who is responsible for agent access when agents come from third parties. Organizations can use AICM control mappings to show regulators that agent access decisions are governed by a structured, reviewable control set.

Finally, CSA's *AI Organizational Responsibilities: Governance, Risk Management, Compliance and Cultural Aspects* [10] addresses the board and executive roles that Article 20 makes legally relevant. It also covers shadow AI prevention, which bears directly on the inventory gap identified above. Read together, these artifacts give management bodies both the evidence of where the risk sits and a control framework for showing that they oversaw it.

References

[1] European Parliament and Council of the European Union. "Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2 Directive)." Official Journal of the European Union, December 2022.

[2] European Commission. "Commission Implementing Regulation (EU) 2024/2690 laying down rules for the application of Directive (EU) 2022/2555 as regards technical and methodological requirements of cybersecurity risk-management measures." Official Journal of the European Union, October 2024.

[3] ENISA. "NIS2 Technical Implementation Guidance." European Union Agency for Cybersecurity, June 2025.

[4] Federal Republic of Germany. "§ 38 BSIG – Umsetzungs-, Überwachungs- und Schulungspflicht für Geschäftsleitungen besonders wichtiger Einrichtungen und wichtiger Einrichtungen." Gesetze im Internet, Bundesministerium der Justiz, accessed November 2026.

[5] Hunton Andrews Kurth. "European Commission Refers Four Member States to CJEU Over NIS2 Transposition Delays." Privacy and Cybersecurity Law Blog, July 2026.

[6] Freshfields. "EU Cybersecurity Package: Commission Proposes Targeted NIS2 Amendments and a Revised Cybersecurity Act." Technology Quotient, January 2026.

[7] Cloud Security Alliance. "Identity and Access Gaps in the Age of Autonomous AI." Cloud Security Alliance, March 2026.

[8] Cloud Security Alliance. "State of Non-Human Identity Security Survey Report." Cloud Security Alliance, 2024.

[9] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2025.

[10] Cloud Security Alliance. "AI Organizational Responsibilities: Governance, Risk Management, Compliance and Cultural Aspects." Cloud Security Alliance, 2024.

[11] European Parliament and Council of the European Union. "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)." Official Journal of the European Union, July 2024.

© 2026 Cloud Security Alliance. AI-assisted Rapid Research.