



• WHITE PAPER

AI agents with access but without identity

AI agents, service accounts and technical users in banks and insurers: ownership, runtime control and what DORA requires

ISSUE

October 2026

Version 2.3

AUTHOR


Parsa Jandaghi

CISSP

AUDIENCE

CISO and IT management

IAM architecture



Contents

	Key points	3
1	Non-human identities as the control gap	5
2	The requirements of the DORA RTS	7
3	An AI agent is not just another service account	11
4	The control model at a glance	14
5	Identity and delegation	17
6	Authority and runtime control	21
7	Accountability and lifecycle	25
8	Microsoft Entra and SailPoint in September 2026	31
9	Approach in five phases	34
	Glossary	37
	Sources	40

Key points

Service accounts, workloads and, more recently, AI agents, together with the keys, certificates and tokens they use, access the same systems as employees. Often they hold broader rights, and often nobody can say who is responsible for them.

AI agents are the new class, and they fit none of the existing processes. A service account does what was defined when it was created. An agent receives a goal and decides for itself which tools it calls to reach it. An authorisation concept modelled on service accounts describes something that does not exist for an agent.

In July 2026 BaFin published its analysis of the DORA incident reports for 2025. 733 incidents were reported, and around 1,200 financial entities were affected. In half of the cases the cause lay with a third party, and around 40 percent of the reported attacks targeted a service provider instead of the institution itself.¹



Figure 1 BaFin analysis of DORA incident reports for 2025 · Source: BaFin, article of 16 July 2026

ECB Banking Supervision counts identity and access management among the areas that are insufficiently developed at some supervised banks.² On 7 July 2026 it wrote to all significant institutions and requested an action plan against AI-driven attacks, to be submitted by 31 October 2026. It names least privilege, logging and human oversight as focus areas.³

This paper describes two gaps. The first lies in the rulebook. The regulatory technical standard under DORA covers machines, but it covers them unevenly: it requires unique identification for persons and systems, while the duty of a unique user account per identity applies to staff only.⁴ The second gap lies in the technology. A sign-in confirms who is acting. It says nothing about whether an agent's next posting is permitted, and for an agent exactly this decision arises on every single action.

WertFels addresses both gaps with a control model of six layers. Two of them take effect before an agent is deployed, three on every action and one continuously and at fixed intervals. Each layer is mapped to a provision of the DORA RTS, and each can be implemented with platforms and standards available today, in part with custom work.

This paper substantiates the following statements:

1. In the general framework, the DORA RTS requires access rights to be updated at least once a year, and at least every six months for systems supporting critical or important functions. The text contains no exception for technical accounts.⁵
2. Generic and shared accounts are to be restricted as far as possible, and every action must remain attributable to a user. A shared service account whose actions cannot be attributed to a user in the log fails this requirement.
3. The term machine identity does not appear in the texts of the ECB, EBA, EIOPA, ESMA and BaFin we reviewed. The requirements arise from interpreting general wording.
4. An AI agent is not just another service account, because its set of actions is only determined at runtime. The NIST NCCoE lists the question of how to determine least privilege for agents as an open question.⁶
5. Agents require a decision per action at four points: the task, the data, the tool and the result. The rules belong in one central place; enforcement belongs close to the tool.
6. Whether an agent acts under its own identity with delegated authority or under the identity of a human decides whether its actions can be traced. The standard for this has existed since 2020.⁷
7. Every agent needs exactly one human owner, one deputy and an anchor in an application. When the owner leaves, the agent is transferred along a succession chain that escalates upwards if nobody takes over. The owner's departure is no reason to stop or decommission the agent.
8. Microsoft Entra has handled agents as their own object class with a mandatory sponsor since 2026. SailPoint built agent security for Identity Security Cloud. Anyone running SailPoint IdentityIQ has to add time limits, behaviour comparison, escalation and short-lived credentials themselves.
9. The recertification of an agent checks its authority: whether it is still needed, whether its catalogue matches its actual behaviour and whether stopping it has been rehearsed.

Non-human identities as the control gap

Many attacks on financial institutions today start with a valid credential. For non-human accounts almost every control that has long been standard for humans is missing.

Half of the reported incidents have their cause with third parties

The BaFin analysis for 2025 covers all reported ICT incidents, including outages without an attack. It still shows a pattern that fits poorly with the usual picture of attacks. The first victim is often a service provider with access to the institution's systems. As we observe it, that access rarely runs through a personal user name. It runs through an interface, a certificate, a key or a token.

A valid token needs no sign-in

Between 8 and 18 August 2025 a group of attackers used stolen OAuth tokens of the sales application Salesloft Drift to exfiltrate data from the Salesforce environments of many companies. According to Google Threat Intelligence, the attackers used compromised but valid tokens of the integration. As a result, no sign-in was needed and hence no multi-factor prompt that could have attracted attention. In the stolen data the attackers searched specifically for further credentials, for cloud account keys, passwords and database tokens.⁸ The US self-regulatory organisation FINRA considered the incident relevant enough to warn its member firms separately.⁹



Figure 2 Sequence of the Salesloft Drift incident, August 2025 · WertFels, based on Google Threat Intelligence

What matters for the assessment is what the case does not show. At the affected Salesforce customers no vulnerability was exploited and no password was guessed. A valid machine credential was stolen in one place and used in another.

Credentials are a traded commodity

In its 2025 situation report the German BSI describes the market of access brokers who collect legitimate credentials for corporate networks and sell them on to other attackers. 461 data leaks became known in the reporting period.¹⁰ ENISA reports a 25 percent increase between 2023 and 2024 in keys, tokens and other secrets detected in code repositories.¹¹

In Bitkom's 2026 survey on economic protection, affected German companies name insufficient detection of security incidents as the most frequent cause at 59 percent, followed by misconfigurations at 57 percent and insufficient identity and access management at 55 percent. Technical vulnerabilities only follow at 50 percent. 1,003 companies with ten or more employees and annual revenue of at least one million euros were surveyed.¹²

Agents extend the reach of every access

AI agents add a class that uses valid access and also decides for itself what to use it for. Gartner expects that by 2028 one in four enterprise breaches will be traced back to AI agent abuse, by external attackers as well as by malicious insiders.¹³ The OWASP GenAI Security Project's risk list for agentic applications explicitly names the abuse of identities, privileges and tools as well as goal hijacking.¹⁴

The requirements of the DORA RTS

The regulatory technical standard under DORA covers technical accounts in more places than many institutions assume. One provision leaves them out, and that is where the discussions in audits begin.

This chapter refers to Delegated Regulation (EU) 2024/1774, the regulatory technical standard on ICT risk management under DORA. Articles 20 and 21 apply to financial entities under the standard framework. For small and non-interconnected investment firms and some other groups, the simplified framework in Title III applies, where Article 33 is the provision on access control.

Article 20 covers systems, unique accounts only for staff

Paragraph 1 requires policies and procedures that ensure the unique identification and authentication of the natural persons and systems accessing the financial entity's information. Paragraph 2 point (a) then narrows this: each member of staff of the financial entity and of ICT third-party service providers is to be assigned a unique identity corresponding to a unique user account. The point applies without prejudice to Article 21 point (c). The duty to keep records of these assignments, also after a reorganisation or the end of a contract, refers explicitly to point (a) in the same paragraph.⁴

Here lies the core of the problem. An institution that runs a service account without a named owner does not breach paragraph 2 point (a), because that provision speaks only of staff. An auditor can, however, demand a named owner through Article 21 point (e)(i), explained below.

The lifecycle applies to all accounts

Article 20 paragraph 2 point (b) requires a process for the entire lifecycle of identities and accounts: creation, change, review and update, temporary deactivation and termination. The text refers to all accounts. Where appropriate and possible, this process is to be automated.

Rights by necessity, generic accounts only with accountability

Article 21 point (a) requires access rights to be assigned by necessity and according to the principle of least privilege, explicitly including remote and emergency access. Point (c) requires generic and shared user accounts to be restricted as far as possible and actions in ICT systems to be attributable to a user at

all times.⁵ Point (d) requires access to ICT assets to be restricted through controls and tools against unauthorised access. This is the basis of enforcement at runtime.

For service accounts, point (c) is the most important provision in practice. Generic accounts are permitted, but only on the condition that accountability is preserved. An account shared by three teams whose actions cannot be told apart in the log does not meet this condition.

Update yearly, every six months for critical functions

Point (e) defines the account management procedures and names user and generic accounts explicitly. Point (i) requires roles and responsibilities to be assigned for granting, reviewing and revoking access rights. This is the obvious provision for a named owner, including for technical accounts. Point (iii) requires access rights to be withdrawn without undue delay when they are no longer needed. Point (iv) requires access rights to be updated when changes are needed, and at least once a year for all ICT systems and at least every six months for systems supporting critical or important functions. In practice institutions implement this duty through recertification. At this point the text makes no distinction by type of account.

Logging names access control explicitly

Article 12 paragraph 2 point (c)(i) requires events of logical and physical access control under Article 21 and of identity management to be logged. Point (d) requires the logging systems themselves to be protected against tampering, deletion and unauthorised access. Point (f) requires the clocks of all ICT systems to be synchronised with a documented, reliable reference time source.¹⁵

Without synchronised timestamps a chain of actions across several systems cannot be reconstructed. For machine accounts acting within milliseconds this is not a formality.

DORA RTS		DELEGATED REGULATION (EU) 2024/1774	
REQUIREMENT	PROVISION	Staff	System, agent
Unique identification and authentication	Art. 20(1)	✓	✓
Unique user account per identity	Art. 20(2)(a)	✓	⊖
Lifecycle of all accounts	Art. 20(2)(b)	✓	✓
Rights by necessity and least privilege	Art. 21(a)	✓	✓
Restrict generic accounts, accountability	Art. 21(c)	✓	✓
Update of rights, yearly, critical every six months	Art. 21(e)(iv)	✓	✓
Logging of access control and identity management	Art. 12(2)(c)	✓	✓
⊖ worded for staff only		✓ applies by its wording regardless of account type	

Figure 3 Requirements of the DORA RTS by type of identity · WertFels, based on Delegated Regulation (EU) 2024/1774

For the audit

Anyone who does not recertify a service account because it is not a human finds no argument for this in the regulation. Article 21 point (e) names user and generic accounts explicitly. Only the duty of a unique user account per identity and the record of these assignments distinguish between staff and other identities.

Gaps in the regulation

The terms machine identity, service account, technical user and non-human identity do not appear anywhere in Delegated Regulation (EU) 2024/1774. They do not appear in the publications of the ECB, EBA, EIOPA, ESMA and BaFin we reviewed either. BaFin's former circulars BAIT and VAIT contained requirements on technical users. Since 17 January 2025 they no longer apply to entities under DORA.¹⁶

This is the real difficulty in audits. No provision names technical accounts. The duties follow from generally worded articles, and both sides argue with that wording.

Further requirements

For high-risk AI systems the AI Act requires automatic logging, effective human oversight and appropriate cybersecurity (Articles 12, 14 and 15). Article 14 explicitly scales the oversight measures to the risk, the level of autonomy and the context of use. Where appropriate and proportionate, the person overseeing the system must be able to intervene and to stop it through a stop button or a similar procedure.¹⁷ These three articles address the design of the system. As the deployer of a system, whether built in-house or bought, an institution is mainly bound by Article 26 (human oversight by competent persons, retention of logs) and Article 27 (fundamental rights impact assessment). For banks and insurers the obvious cases are in Annex III: creditworthiness assessment of natural persons as well as risk assessment and pricing in life and health insurance. After the postponement through the Digital Omnibus, these obligations apply from 2 December 2027.¹⁸

For significant institutions there is also the action plan that ECB Banking Supervision requests by 31 October 2026. The letter names least privilege, logging and human oversight, controls this paper makes concrete for agents. It is addressed to significant institutions only. Insurers are not affected by it; for them the DORA requirements apply unchanged.

An AI agent is not just another service account

All access management rests on the assumption that we know, when granting access, what an account will do. For AI agents this assumption no longer holds.

The set of actions is only known at runtime

A classic service account has one property on which everything builds: its actions are known when access is granted. The nightly job reads from two tables and writes to a third. The interface to payments calls four endpoints. In this case least privilege is arithmetic: all required rights and nothing else.

For an AI agent this does not apply. It receives a goal instead of a sequence. Which tools it calls, in which order and how often, follows from the goal, the state of the systems and what it finds along the way. The set of actions emerges anew in every run.




Figure 4 Service account and AI agent: when the set of actions is known · WertFels

The NIST National Cybersecurity Center of Excellence describes exactly this situation as unresolved. In its concept paper of February 2026, the question of how least privilege can be determined when the required actions are unknown

at deployment is on a list of open issues. Next to it are the questions of how an agent proves its authority to act and how non-repudiation of its actions is achieved.⁶

We draw a practical conclusion from this. An authorisation concept that has the same form for agents as for service accounts lists rights that at runtime either fall short or go too far.

Three kinds of non-human identity

The discussion suffers from putting everything that is not a human into one bucket. In practice three kinds must be separated, because they have nothing in common in issuance, review and removal.

The static service account is long-lived, has a password or a key stored somewhere and a fixed purpose. Large institutions often run such accounts in the thousands. Its control consists of inventory, owner, rotation and removal.

The workload is short-lived and receives its identity anew at every start, for example a container that obtains, through workload identity federation, a token that expires after a few minutes. Here the issuance itself is controlled: who may identify as this workload, against which trust anchor and with which lifetime.

The agent differs from both in that it acts on behalf of someone and makes choices while doing so. It has its own identity; its authorisation derives from a principal. For an agent, control therefore centres on its authority.




Figure 5 Three kinds of non-human identity and their control points · WertFels

Connections turn a model into an actor

A language model on its own produces text. It can only act once it is connected to tools, to mailboxes, customer files, ticketing systems and internal interfaces. The Model Context Protocol is currently one of the most widespread ways to make these connections. Every new server extends what an agent can reach, often faster than information security learns about it.

The protocol's Authorization specification treats a server as an OAuth resource server. It requires the client to name the target resource under RFC 8707¹⁹ in authorisation and token requests, the server to accept only tokens issued for itself and never to pass received tokens on to downstream services.²⁰ This way a token can only be used at the service it was issued for.

Two limitations should be stated openly. Authorization is optional in the protocol, and for local connections over standard input and output the credentials come from the process environment. In addition, the standard does not prescribe any check of authority; it only provides the means. An institution without a policy gets none from the protocol.

The decision arises on every action

This leads to the most important difference from a service account. A successful sign-in establishes who the agent is. Then the actual questions begin: whether it may read this data, call this tool with these parameters and output this result. Each of these questions arises anew in every run. A control model for agents must therefore work before deployment, on every action, and continuously and at fixed intervals.

The control model at a glance

Governance asks who is responsible for an agent. Runtime control asks what it may do at this moment. Both questions belong in one model.

We arrange the controls for agents in six layers. Visibility and accountability must be in place before an agent goes into production. Identity, authority and enforcement take effect on every single action. Evidence closes the loop, continuously through logging and at fixed intervals through recertification.




Figure 6 The WertFels control model for agent identities · WertFels

The first layer is visibility. A register lists every agent and every technical account with purpose, owner, target systems and last use. It also shows the relationships: which agent reaches which tools, interfaces and data stores.

Accountability comes second. Every agent has exactly one human as owner, one deputy and an application as anchor. If the owner leaves the organisation, a succession chain applies and the agent keeps running.

Identity follows as the third layer. Every agent has its own identity, issued without a stored secret, short-lived and bound to a key, so that a stolen token is useless elsewhere.

Authority forms the fourth layer. It defines the autonomy level at which an agent works, which tools and operations are in its task catalogue, how long the authority is valid and on whose behalf the agent acts.

Enforcement is the fifth layer. At four checkpoints a central policy decides on every action: at the task, the data, the tool and the result.

Evidence is the sixth layer. The log records the entire chain with one correlation ID, recertification checks the authority, and stopping is rehearsed.

Six principles

The technology in this field changes within months. The principles should be stable enough to outlast changing platforms and protocols.

1. Identity first, but not alone. Every decision starts with the agent's identity and includes the person, the purpose and the data.
2. No standing privileges. The authority sets the frame for months. The agent receives the actual rights per task and for minutes, never in advance.
3. Delegation instead of impersonation. The log always shows the agent and its principal, the person or organisational unit it acts for.
4. Central policy, distributed enforcement. The rules are maintained and tested in one place and applied close to the tool and the data.
5. Evidence from the first call. What is not logged can neither be reviewed nor reconstructed.
6. Easy to stop, hard to reactivate. The owner may stop an agent without further consent, and so may the function designated for stopping it. The technical custodian may reactivate it only with the owner's consent.

Our fourth principle corresponds to the separation of policy decision point and policy enforcement point that NIST describes for zero trust architectures.²¹ What is new with agents is only how often decisions have to be made.

The model in an example

An insurer deploys an agent in accounts payable. The agent reads incoming invoices from a shared mailbox, matches them with supplier master data and enters payments in the accounting system. Each payment only takes effect once a clerk approves it. For amounts of 10,000 euros or more, two approvals are required. The figures are examples.

LAYER	IN THE EXAMPLE
Visibility	register entry with purpose invoice processing, target systems mailbox, master data and accounting, last use
Accountability	owner is the head of accounts payable, deputy is the payments team lead, anchor is the accounting system
Identity	own agent identity, token valid for 15 minutes, bound to the agent's key
Authority	level 3 for payments and replies to third parties, level 1 for master data, catalogue of four operations, time limit of twelve months
Enforcement	payments only to bank details on file for known suppliers, approval per payment, two approvals from 10,000 euros, no bank details in results sent to third parties
Evidence	every posting logged with the correlation ID of the run, recertification every six months, stop rehearsed once per period

Identity and delegation

An agent needs its own identity. It is issued without a stored secret, expires after minutes and shows in the target system on whose behalf the agent acts.

Issuance without a stored secret

Every API key, every client secret and every simple bearer token has to be stored, distributed and rotated somewhere. For agents whose instances appear and disappear faster than any rotation cycle, this arithmetic does not work.

The alternative derives the identity from a property of the runtime environment that a third party can confirm. The platform attests, for example, the checksum of the image, the service account in the Kubernetes cluster or the state of a TPM. A trust anchor checks this attestation against a policy and only then issues a short-lived identity.

SPIFFE describes this pattern across vendors. The identity of a service is called a SPIFFE ID and is issued in a short-lived, cryptographically verifiable document, the SVID, as an X.509 certificate or as a JWT. The production-ready implementation SPIRE attests nodes and workloads and issues the document only if the confirmed properties match the policy.²² The large cloud providers offer the same principle as workload identity federation. Microsoft also explicitly advises against client secrets in production for Agent Identity Blueprints in Entra and recommends federated credentials with managed identities.²³




Figure 7 Issuing an agent identity without a stored secret · WertFels, based on SPIFFE, RFC 9449 and RFC 8705

Short-lived and bound to a key

The lifetime of an agent determines the model of its credentials. Long-lived agents, such as a permanently running service bot, sign in through federated identities and are re-attested at fixed intervals. Short-lived agents are created for one task, receive tokens valid for minutes and lose their identity as soon as the container ends.

A short expiry alone is not enough. Until it expires, a stolen token is as good as the original. Demonstrating Proof of Possession (DPoP) under RFC 9449 binds an access token to a key that only the client holds, so an intercepted token alone cannot be used elsewhere.²⁴ Certificate-bound tokens over mutual TLS under RFC 8705 achieve the same at the transport layer.²⁵ In the Drift incident, this binding would have made the stolen tokens worthless to the attackers, provided the private key had not been stolen as well.

Delegation instead of impersonation

When an agent creates a posting on behalf of a clerk, there are two technical routes. If the agent uses the clerk's token, the target system logs the clerk's name. If the agent has its own identity and carries the clerk's authority with it, the log shows that agent A acted on behalf of person B.

The difference has been standardised since January 2020. RFC 8693 on OAuth 2.0 Token Exchange separates both cases explicitly. With impersonation, the acting party receives all rights of the principal and cannot be told apart from it. With delegation, it keeps its own identity, and it remains visible that it acts for someone else. Technically this happens through two values in the request: the subject token names the party on whose behalf the action is taken; the actor token names the acting party. The issued token then carries an actor claim that records the delegation.⁷




Figure 8 Impersonation and delegation in the log of the target system · WertFels, based on RFC 8693

We consider the choice between delegation and impersonation the single most important decision when deploying agents in supervised institutions. Article 21 point (c) of the DORA RTS requires actions to be attributable to a user at all times. If an agent acts under a human's identity, this requirement is met to the letter but missed in substance. The log shows a human, while software acted. In an audit this cannot be explained, and after an incident it cannot be reconstructed.

Platforms implement delegation differently. Microsoft Entra maps it for agents through its own on-behalf-of flow, which builds on the JWT bearer grant and not on token exchange under RFC 8693.²³ An institution should verify in its own tenant which information about the acting agent the token carries into the target system before basing its logging on it.

In chains the principal and the last link count

As soon as an agent instructs another agent, a chain forms. RFC 8693 represents it through nesting, and the outermost actor claim names the current actor. The standard is unambiguous: whoever evaluates a token may only consider the top-level values and the current actor for the access decision. Earlier actors in nested entries are information only.⁷




Figure 9 Delegation chain and evaluation of the actor claim · WertFels, based on RFC 8693, section 4.1

For an institution this leads to two rules that no standard tool has built in. The length of the chain has to be limited, and every link has to be logged with the same correlation ID, because the target system does not evaluate the chain when deciding. If the depth is left open, the origin of an action can only be determined by merging several logs.

Standards now emerging

Two IETF drafts extend the pattern beyond the boundary of a single service. Transaction Tokens carry user identity, workload identity and authorisation context as short-lived, signed tokens through the call chain within one trust domain.²⁶ The Identity Assertion JWT Authorization Grant allows an application to obtain an access token for another application's API through the shared identity provider. It is the basis of the pattern that has become known as Cross App Access.²⁷

Neither draft has been adopted yet. Institutions should treat the two drafts as a direction and design their token issuance so that they can adopt them later. Rich Authorization Requests under RFC 9396, by contrast, are already a standard and allow a token request to describe precisely which access is needed, for what and to what extent.²⁸ The order is what matters now: attested sign-in without a secret and delegation belong in the agent before it goes into production. Retrofitted into agents already running with shared secrets, they cost many times more.

Authority and runtime control

Authority defines what an agent may do in principle. Runtime control decides whether the concrete action is permitted at this moment.

Rights by autonomy level

Gartner distinguishes four autonomy levels: observe, advise, act with approval and act autonomously.²⁹ For each level an institution defines which systems, data classes and operations are permitted. An agent at level 1 may only read. At level 2 it may create proposals that a human adopts, but make no final writes. At level 3 it may write, but every write operation waits for human approval. For systems supporting critical or important functions, we recommend not approving level 4 for the time being.




Figure 10 Four autonomy levels as the basis for granting rights · Levels according to Gartner, mapping by WertFels

Rights per task instead of in reserve

The authority defines what an agent may do at all over a period of months. The agent receives the concrete authorisation only for the duration of one task and returns it afterwards. Technically these are short-lived tokens bound to one target with a narrow scope. Access packages with an expiry date in Entra additionally limit the assignment itself, but they work in days and not per task. The Model Context Protocol provides a step-up at runtime for this: if a token lacks a scope, the server names the required permissions, and the client requests exactly those.²⁰

Approval per action for anything that moves money

For actions that move money, change rights or delete data, approval per action is mandatory, whatever the agent's level. With the Access Request and Approval Profile, the OpenID Foundation is working on a standard that puts a request into a waiting state until a named person has approved it. The example in the draft is an agent that triggers supplier payments above a threshold.³⁰ Until the standard is adopted, approval per action remains vendor-specific, for example as an escalation in an authorisation gateway.

Four checkpoints on the path of an action

An agent action passes four points where a decision is needed. At the task, the check is whether the assignment fits the agent's purpose and level and whether it comes from an authorised source at all. Instructions hidden in documents or emails the agent has read are not an assignment. At the data, the check is which sources the agent may read and which parts of them reach its context, with filters per row, column and document. At the tool, the check is whether the operation is in the catalogue and whether parameters and amounts are within limits. At the result, the check is what leaves the agent and to whom.




Figure 11 Four checkpoints on the path of an agent action, decided in one central place · WertFels

In the accounts payable example this means: an assignment that originates from the text of an invoice is discarded. Only the shared mailbox and the supplier master data are read. A payment to bank details not on file for the supplier is rejected. Every other payment waits for approval, and payments of 10,000 euros or more wait for two. A reply to an external recipient contains no bank details.

Together with the identity and authority layers, the four checkpoints mitigate the first three risks the OWASP GenAI Security Project names for agentic applications: goal hijacking, tool misuse and the abuse of identities and privileges.¹⁴

Person and agent count together

There are three patterns for every decision. If only the agent counts, it accumulates the rights of all its users over time. If only the person counts, the agent may do everything the person may do, including what never belonged to its purpose. The composite decision is the one that holds up: only what the person may do and what is at the same time in the agent's catalogue is permitted.



Figure 12 Composite decision from user rights and agent authority · WertFels

The principle is familiar from delegated permissions. An application acting on behalf of a user can reach nothing the signed-in user could not reach.³¹ With agents the second boundary is added, the task catalogue.

Central decision, distributed enforcement

The rules for these decisions belong in one place. If authorisation logic is built into every agent separately, every team builds its own model, and nobody can say any longer which rule applies where. Enforcement, by contrast, belongs close to the action: in the agent's runtime, in a gateway in front of the MCP servers, in the API gateway, in the data access layer and at the exit.

The connection between the two has been standardised since January 2026. The OpenID Foundation adopted the AuthZEN Authorization API 1.0 as a final specification. It describes how an enforcement point requests a decision from the decision point, independent of the vendor on either side.³² The policies themselves should be maintained as code, versioned and tested before every change. Every decision is logged with request, rule and result.

For the audit

Article 21 point (a) requires rights by necessity. For an agent, necessity is a property of the individual action. An institution that makes the decision per action centrally and logs every decision can prove for every action which rule permitted it. The rules themselves have to be reviewed regularly for necessity.

Accountability and lifecycle

An agent cannot be held liable. Every control meant to work in an emergency ultimately needs a human who decides, confirms and takes responsibility.

No named human, no accountability

As the first and simplest rule, we recommend that every agent has exactly one human owner. A team, a department or a functional mailbox is not enough. The reason is the same as for service accounts, only sharper: an account that belongs to a group belongs to nobody. An agent, in addition, makes decisions that nobody approved in detail beforehand. Whoever answers for them afterwards has to be determined in advance.

An agent that can be approved therefore has three bindings: an owner with a deputy, an application in the configuration management database as anchor and an authority with catalogue, level and time limit.




Figure 13 The three bindings of an approvable agent · WertFels

Seven attributes per agent

In inventories we check seven attributes per agent. If one of them is missing, the agent cannot be approved.

1. The owner is a named human with a deputy.
2. It is defined on whose behalf the agent acts: permanently for an organisational unit, in which case the authority is generic and must be narrow, or each time for the triggering person, in which case delegation must be implemented technically.
3. A catalogue describes the permitted tools and operations. System rights describe what would be possible, the catalogue describes what is intended.
4. The authority is time-limited and must be renewed explicitly. For service accounts an authority without an expiry date is the norm and the reason they turn into legacy liabilities.
5. The identity comes from a controlled trust anchor and is valid for minutes, not years.
6. The chain is logged: triggering person, agent, downstream agents, tool, target system and result, with a consistent correlation ID and synchronised clocks.
7. A named function can stop the agent, with a rehearsed procedure and a measured time from decision to standstill.

The lifecycle in six stages

The lifecycle of an agent corresponds to what institutions know for humans as joiner, mover and leaver, with two differences: the agent has an owner from the start, and it has an expiry date.




Figure 14 Lifecycle of an agent in six stages · WertFels

At registration the agent is created as its own object, with purpose, owner, deputy, catalogue and time limit. An agent that exists only in a development environment and was never registered is shadow IT. Approval is confirmed by the owner, the business unit and information security. If the agent accesses systems supporting critical or important functions, the approval belongs in the same process as privileged rights. At issuance the agent receives short-lived credentials. New static keys and passwords are not permitted for it. In operation every action is logged, and anomalous behaviour leads to an intervention. At recertification the owner confirms need, catalogue and behaviour. Finally, decommissioning deactivates the agent, revokes its credentials, removes its rights and retains its log. Triggers are the end of the time limit without renewal, the end of its purpose or an incident after which the agent is not to return to operation. In an incident the agent is stopped first and decommissioned after the assessment. While ownership of the agent is passing through the succession chain, an expiring time limit is not a trigger. The authority stays frozen at its current scope until a new owner confirms.

Assigning ownership: one person, one deputy, one deadline

From the platform models and the standards, an assignment pattern can be derived that works everywhere. The business owner is a natural person, called sponsor in Entra, primary owner in SailPoint Identity Security Cloud and administrator in SailPoint IdentityIQ. A deputy is named. Technical support is separated from business accountability: the business owner may stop and transfer, the technical custodian may reactivate and restore with the owner's consent.

Ownership itself should be time-limited. The platforms support expiry dates for permissions, but not for ownership. We recommend linking both through recertification: whoever does not confirm the agent on time loses ownership to the next link of the succession chain. The agent keeps running, and the question of its owner moves upwards.

When the owner leaves: transfer and escalate instead of stopping

Microsoft Entra and SailPoint respond to the owner's departure with a transfer, Entra only for the sponsor. Neither platform decommissions the agent for that reason. We consider this right and go one step further: even at the end of the succession chain after expired deadlines, stopping or decommissioning an agent in production solely because of missing ownership is not a permissible measure. A departure or a suspension often comes without notice. An agent deactivated at that moment stops a business process nobody wanted to stop. A payment run or an interface to the core banking system must not depend on whether HR names a successor in time.

The justification comes from a basic principle of security governance: information security is meant to enable business processes. A control that paralyses operations to close an administrative gap has lost its proportionality. A security incident is a reason to stop an agent, a missing owner is not.

The pressure to find a successor therefore comes from a succession chain of three links that ends in an escalation. The first link is the named deputy, without whom an agent cannot be created at all. The second link is the application owner according to the configuration management database. This person knows the agent's purpose because they know the application, and when the role holder changes, the successor is set through the change process, without involving HR. The BSI implementation notes on module ORP.4 explicitly permit a role holder as the responsible party for group accounts³³, and CIS Controls safeguard 5.5 requires an inventory of service accounts with responsible department, review date and purpose.³⁴ CIS is satisfied with a department; we require a person. The third link is the line manager of the departed person as interim owner with a deadline. When it expires, ownership passes to that manager's superior, with the same deadline, and so on upwards. Within a few weeks, an ownerless agent thus lands with a division head who will want to know exactly who is taking care of it.



Figure 15 Succession chain when the owner leaves · WertFels

Every transfer triggers three measures. The new owner confirms purpose, catalogue and behaviour of the last period in an unscheduled recertification. Until then no extensions can be approved, so no new tools, no new target systems and no extension of the time limit. Finally, any secret the departed owner knew or held is rotated, and the departed owner's rights on the registration or blueprint are removed.

An attack path that Ridgeline Cyber described for Entra in May 2026 shows why this care is needed. When a user account is deactivated, its ownership of service principals remains. Whoever later reactivates or takes over the account inherits it and can create new credentials.³⁵ Orphaned ownership is therefore an entry point. The Cloud Security Alliance ranks the termination of an agent identity on a par with its creation.³⁶ Microsoft reports on ten roundtables with around 150 identity security professionals at Identiverse 2026. In nine of them the ownerless agent was named as a problem.³⁷ Gartner expects that by 2027 four in ten enterprises will downgrade or switch off autonomous agents because of governance gaps.²⁹

Recertification checks the authority

The six-month deadline for critical functions applies to agents just as it does to human accounts. In its usual form, however, it misses the point for them, because a ticked list of system rights says little about an agent. We propose five review questions the owner answers:

1. Is the agent still needed?
2. Does its catalogue match what it actually did?
3. Is the authority limited to what is necessary?
4. Do the runs of the last period contain only actions the owner would answer for?
5. Was stopping the agent rehearsed at least once in the period?

The second question is the most revealing. It requires the actual calls to be analysed and compared with the catalogue. An institution that does this usually finds two things: tools that were approved and never used, and calls nobody anticipated.

Treat stopping like an emergency drill

The joint guidance on adopting agentic AI published in May 2026 by CISA and five partner agencies recommends human control points throughout an agent's workflow, explicitly including the ability to interrupt a task while it runs.³⁸ It does not describe a fixed stop procedure with deadlines. We recommend rehearsing the stop before go-live and then once per recertification period on each agent in production, in a planned maintenance window. The time from decision to stand-still is measured and documented. As long as the procedure has never been rehearsed, nobody knows whether it works.

Microsoft Entra and SailPoint in September 2026

In German banks and insurers, Microsoft Entra ID as directory and SailPoint as governance layer usually sit side by side. Both have added agent capabilities since 2025, with clear differences.

Because this area changes within months, the information reflects the state of September 2026. It comes from vendor documentation, dated where the vendor provides a date. The same assessment scheme applies to further platforms such as Keycloak, One Identity Manager or Omada Identity Suite, but this paper does not evaluate them.

Microsoft Entra Agent ID

Since spring 2026 Microsoft has handled agents as their own object class in the directory, with general availability announced on the what's new page as of 1 May 2026.³⁹ Since May 2026 Copilot Studio automatically creates such an identity for every new agent, and this can no longer be switched off at environment level.⁴⁰

The object model defines four types. The Agent Identity Blueprint is the template, comparable to an app registration. The Agent Identity is the running instance, comparable to a service principal. The Agent User is a user account for agents that need a mailbox, licence or group membership. The Blueprint Principal represents the blueprint in a tenant, for multitenant agents also in other tenants.⁴¹ In the audit log the types appear under the base types application, service principal and user, identifiable by the attribute `agentType`.⁴²

Ownership is split in two. In Entra the owner is the technical role. It is optional and can be a user or a service principal, but never a group. The sponsor is the business role and mandatory at creation. Sponsors may deactivate, change the sponsor and delete; only owners may reactivate. A user who creates an agent without naming a sponsor becomes the sponsor, unless the user holds an Agent ID admin role. Without user context, the creating application must set the sponsor.⁴³ In practice this separation is the most valuable decision in the Entra model, because it leaves accountability for the purpose with the business unit and for operations with IT. Entra does not enforce the consent rule of principle 6: a single Entra owner, the technical custodian in our terms, can reactivate an agent alone.

Conditional Access is evaluated on every token issuance for agents. Policies can target individual agent identities, blueprints and, through custom security attributes, groups of them. The same documentation names the limits: policies

for all users do not include agent users, and access through API keys bypasses the control entirely.⁴⁴ Whether continuous access evaluation applies to agent identities is not documented. For service principals it is limited to Microsoft Graph and single-tenant scenarios.⁴⁵ An institution should therefore assume that an issued agent token stays valid until it expires, and choose a short validity.

Identity Protection provides a report on risky agents with eight detections, such as unfamiliar resource access and suspicious use of credentials. All of them run after the fact. If an agent acts on behalf of a user, the risk is attributed to the user.⁴⁶

Entra ID Governance provides access packages for agent identities, with an expiry date per assignment and notification of the sponsor.⁴⁷ Microsoft recommends a recertification by the sponsor every six to twelve months.⁴⁸ On departure, a lifecycle workflow transfers sponsorship to the manager.⁴⁹ There is no such automation for the technical owner. Microsoft provides for several owners as a safeguard.⁴³ Entra does not enforce a time limit on the agent identity itself, only on the assignments, and it does not see agents that do not use Entra tokens.

SailPoint Identity Security Cloud and SailPoint IdentityIQ

SailPoint built agent security in two steps, both for Identity Security Cloud. Agent Identity Security has been available since autumn 2025 and handles AI agents as a kind of machine identity.⁵⁰ Agentic Fabric has been generally available since 4 August 2026 and brings a register of non-human identities, discovery through sensors on endpoint and browser, activation and deactivation with an approval workflow, owner assignment with succession and access reviews.⁵¹

The owner model of the cloud suite is the most mature of the platforms assessed here. The primary owner must be a human, and up to ten additional owners are possible. If the primary owner becomes inactive, ownership passes to the first active additional owner, otherwise to that owner's manager. Failing that, the agent has no primary owner until one is added manually.⁵²

For SailPoint IdentityIQ, which many German banks and insurers run on premises, neither product exists. SailPoint offers a plugin that supplies identities from IdentityIQ to Agentic Fabric as owners and users of non-human identities.⁵³ The release notes of version 8.5 from July 2025 contain no function for agents.⁵⁴ Nor does the announcement of IdentityIQ 9.0 in October 2026.⁵⁵

With built-in features IdentityIQ can do more than the missing products suggest. The identity attribute Type has values such as service account and RPA/bots and can be extended by a value for AI agents. The attribute Administrator names owner and certifier of a non-human account. Targeted certifications can be filtered by type.⁵⁶ On departure, Rapid Setup can reassign the identities the

person owned to the manager, by rule or to an alternative identity.⁵⁷ There is no automatic escalation without a successor. Such an escalation has to be built with lifecycle events.

For institutions on IdentityIQ our conclusion is clear. Agents as their own identity type with an owner, targeted certification and reassignment on departure can typically be implemented with built-in features within a few weeks. The institution has to build four things itself: the time limit with a renewal duty, the comparison of catalogue and behaviour, the escalation when no successor is named and the integration of short-lived credentials. We do not expect agent functions in IdentityIQ itself, because SailPoint is betting on the cloud with Agentic Fabric.




Figure 16 Coverage of lifecycle functions by platform · As of September 2026, vendor documentation

Operations and evidence

For traceability, Entra provides the attribute agentType and the blueprint identifier in its audit logs, and the agent type as a filter in its sign-in logs.⁴² For Agent 365 instances, Microsoft Purview captures input, output and referenced files for the interactions between agent, human and tool.⁵⁸

A cross-vendor standard that links a chain of actions across Entra, SailPoint and the target systems with one correlation ID is still missing. Candidate building blocks are Shared Signals, AuthZEN and Transaction Tokens. Until then the institution establishes the correlation itself: with its own correlation ID that the agent carries on every call, and with synchronised clocks.

Approach in five phases

Sequence matters. An institution that starts with the authorisation concept without knowing which agents exist writes rules for an unknown set.




Figure 17 Approach in five phases · WertFels

Phase 1: Inventory

The first question is how many agents and technical accounts exist and who they belong to. An institution that asks it systematically for the first time usually finds considerably more accounts than expected, and for a substantial share no owner can be named. There are no reliable industry-wide figures. All published ratios between machine and human identities come from vendors with a sales interest or were commissioned by them. Depending on the method, they range from 20:1 to over 100:1.⁵⁹

The inventory covers the three kinds from chapter 3 and their credentials in five places: service accounts in directories and databases, keys and certificates in configurations and vaults, OAuth applications with their consents, workloads in cloud and container environments, as well as agents, including those that employees created in tools such as Copilot Studio without IT knowing. A plain inventory is not enough. The register shows which agent reaches which tools, APIs and data, because only these relationships are the actual access paths. An entry without use in 90 days is a candidate for decommissioning; an entry without an owner goes to the application owner in the configuration management database, or to the responsible division head if there is no application.

Phase 2: Establish ownership

Only with the register can accountability be assigned. Every account gets a human owner, a deputy and an anchor in an application. This phase is organisational and the least popular, because it forces business units to accept accounts they never knowingly ordered. In our experience willingness rises as soon as an audit report names an ownerless account with access to a critical system. An institution that completes this phase before the audit sets the tone of the conversation.

Phase 3: Target architecture and authorisation concept

With known agents and named owners, the institution defines how agents are controlled. The target architecture describes which building blocks of the control model are provided as shared services: issuance of identities, policy and decision point, enforcement points and logging. If they are available early, no development team builds its own identity model. The authorisation concept defines autonomy levels, task catalogue and delegation. Where target systems do not yet support delegation, the gap is recorded in the register and taken into account in the risk assessment.

Phase 4: Pilot and operation

The first implementation is meant to test the target architecture. The pilot is small enough to run quickly and typical enough to stand for later agents, for example the accounts payable agent. It covers all six layers of the control model and passes through the lifecycle from registration to the first recertification. What works becomes the pattern for the next agents. Recertification runs on the same deadline as for human accounts, and every change of owner triggers an unscheduled one.

Phase 5: Mapping and evidence

Finally the approach is mapped against the requirements that apply to the institution: the DORA RTS with Articles 12, 20 and 21; from December 2027 the AI Act for institutions with high-risk AI systems, as deployer in any case Articles 26 and 27, for their own systems also Articles 12, 14 and 15; and for significant institutions the action plan of ECB Banking Supervision. Because the action plan is due by 31 October 2026, it should be prepared in parallel with phases 1 and 2. The mapping shows where the institution meets the requirements, and it shows where the regulation stays general and the institution has to justify its own standard. For machine and agent identities this is almost always the case. Most agents in institutions are not high-risk systems, and the accounts payable agent is not one either. This classification, with its justification, belongs in the documentation before the supervisor asks for it.

Ten questions a CISO should be able to answer today

We ask the following questions in initial meetings. Every question you cannot answer with certainty and evidence is a task.

1. Do you know the number of non-human accounts with access to systems supporting critical or important functions?
2. Does each of these accounts have a named owner and a deputy?
3. Do you know how many AI agents run in your organisation, including those created by employees themselves?
4. Do your agents act under their own identity or under the identity of the triggering person?
5. Can you reconstruct the triggering person, the authority and all intermediate steps for any action of an agent in under an hour?
6. What happens to the agents of an employee who is suspended tomorrow?
7. When was an agent in production last deliberately stopped, and how long did it take?
8. Are service accounts recertified on the same deadline as human accounts?
9. Are all credentials of your agents valid for minutes only?
10. Does your documentation state which standard you apply to agent identities and why?

Glossary

The terms are defined as this paper uses them.

Actor claim

Value in a security token under RFC 8693 that names the acting party and distinguishes it from the principal. Technical basis of delegation.

Agent, AI agent

Software that receives a goal and decides for itself which tools it calls in which order. Its set of actions is only determined at runtime.

Attestation

Confirmation of properties of a runtime environment by a third party, such as the checksum of an image. Basis of issuance without a stored secret.

Authority

Framework within which an agent may act: autonomy level, task catalogue, principal and time limit. The concrete rights arise per task within this frame.

Autonomy level

Classification of agents by their degree of independence: observe, advise, act with approval, act autonomously.

Composite decision

Access decision that checks the person's rights and the agent's authority at the same time. Only the intersection is permitted.

Configuration management database (CMDB)

Register of an organisation's applications and systems with their responsible parties. The application owner recorded there is the second link of the succession chain.

Control model

The six layers in which this paper arranges the controls for agents: visibility, accountability, identity, authority, enforcement and evidence.

Decommissioning

Permanent retirement of an agent: deactivate, revoke credentials, remove rights, retain the log. Triggered by expiry of the time limit without renewal, the end of its purpose or an incident, never by a missing owner alone. During succession an expiring time limit triggers nothing.

Delegation

An agent acts under its own identity with a person's authority carried along. The log shows that agent A acted on behalf of person B.

DPoP

Demonstrating Proof of Possession under RFC 9449. Binds a token to a key of the client, so a stolen token alone cannot be used.

Impersonation

An agent acts with a person's token and cannot be told apart from that person in the target system.

Least privilege

Principle that an identity receives only the rights it needs for its purpose. For agents it can only be implemented through a task catalogue, rights per task and a decision per action.

Model Context Protocol (MCP)

Open protocol through which agents address tools and data sources. Its Authorization specification treats servers as OAuth resource servers and prohibits passing tokens on.

Non-human identity

Umbrella term for identities that do not represent a natural person. This paper distinguishes three kinds: service accounts, workloads and agents. Keys, certificates, tokens and OAuth consents are their credentials.

Owner

Natural person with business accountability for purpose, catalogue and behaviour of a non-human identity. In Entra Agent ID the business role is called sponsor, in SailPoint Identity Security Cloud primary owner, in SailPoint IdentityIQ administrator.

Policy decision point

Central function that checks requests against the policy and decides. The AuthZEN Authorization API standardises the request to it.

Policy enforcement point

Function that permits or rejects an action, for example a gateway in front of tools or APIs. It requests the decision from the decision point.

Recertification

Regular confirmation by the owner that an identity is still needed and its rights are appropriate. The usual way to meet the DORA RTS duty to update access rights, at least yearly, for critical or important functions at least every six months.

Revocation

Invalidating the credentials of an identity. Part of stopping and of decommissioning.

SPIFFE

Cross-vendor standard for short-lived, cryptographically verifiable workload identities, with SPIRE as a widely used implementation.

Stop

Immediate interruption of an agent, for example after an incident, including revocation of its current credentials. To be rehearsed and measured before go-live and then once per recertification period.

Succession chain

Order in which ownership of an agent passes when the owner drops out: named deputy, application owner, line manager with a deadline, then escalation upwards. The agent keeps running throughout.

Task catalogue

List of tools and operations an agent may use for its purpose. For agents it replaces the list of rights and is the basis of recertification.

Technical custodian

Role responsible for operating an agent, separate from the business owner. Re-activates and restores, only with the owner's consent. In Entra this corresponds to the owner role.

Transaction token

IETF draft for short-lived, signed tokens that carry identity and context through a call chain within one trust domain.

Trust anchor

Function to which a workload or an agent proves its identity in order to receive credentials.

Workload

Short-lived execution unit such as a container or a function that receives its identity anew at every start.

Sources

1. BaFin, Vernetzung vervielfacht IT-Risiken (analysis of the 2025 DORA incident reports), article of 16 July 2026.
2. ECB Banking Supervision, Supervisory Newsletter of 13 November 2024.
3. ECB Banking Supervision, letter SSM-2026-0301 to all significant institutions, 7 July 2026.
4. Commission Delegated Regulation (EU) 2024/1774, Article 20(1) and (2).
5. Commission Delegated Regulation (EU) 2024/1774, Article 21 points (a), (c), (d) and (e).
6. NIST NCCoE, Accelerating the Adoption of Software and AI Agent Identity and Authorization, concept paper, draft of February 2026, comment period from 5 February 2026.
7. RFC 8693, OAuth 2.0 Token Exchange, January 2020, sections 1.1, 2.1 and 4.1.
8. Google Threat Intelligence, Widespread Data Theft Targets Salesforce Instances via Salesloft Drift, 26 August 2025.
9. FINRA, Cybersecurity Alert: Salesloft Drift AI Supply Chain Attack, undated.
10. BSI, Die Lage der IT-Sicherheit in Deutschland 2025, chapter on data leaks.
11. ENISA Threat Landscape 2025, Version 1.2, page 11.
12. Bitkom, Wirtschaftsschutz 2026, press release of 26 August 2026.
13. Gartner, Top Predictions for IT Organizations and Users in 2025 and Beyond, press release of 22 October 2024.
14. OWASP GenAI Security Project, OWASP Top 10 for Agentic Applications for 2026, 9 December 2025, risks ASI01 to ASI03.
15. Commission Delegated Regulation (EU) 2024/1774, Article 12(2) points (c), (d) and (f).
16. BaFin, DORA kommt: Änderungen bei den aufsichtlichen Anforderungen an die IT, notice of 10 January 2025.
17. Regulation (EU) 2024/1689 (AI Act), Articles 12, 14, 15, 26 and 27 and Annex III point 5.
18. Regulation (EU) 2026/1744 amending the AI Act (Digital Omnibus), in force since 27 July 2026.
19. RFC 8707, Resource Indicators for OAuth 2.0, February 2020.
20. Model Context Protocol, Specification 2026-07-28, Authorization, including scope challenge handling.
21. NIST Special Publication 800-207, Zero Trust Architecture, August 2020.
22. SPIFFE and SPIRE, project documentation, spiffe.io, retrieved September 2026.
23. Microsoft Learn, Agent OAuth flows: On-behalf-of flow, Microsoft Entra Agent ID, updated 10 August 2026.
24. RFC 9449, OAuth 2.0 Demonstrating Proof of Possession (DPoP), September 2023.
25. RFC 8705, OAuth 2.0 Mutual-TLS Client Authentication and Certificate-Bound Access Tokens, February 2020.
26. IETF, Transaction Tokens, draft-ietf-oauth-transaction-tokens-08, 2 March 2026, since superseded by revision 11 of 30 July 2026.
27. IETF, Identity Assertion JWT Authorization Grant, draft-ietf-oauth-identity-assertion-authz-grant-04, 21 May 2026.
28. RFC 9396, OAuth 2.0 Rich Authorization Requests, May 2023.
29. Gartner, Gartner Says Applying Uniform Governance Across AI Agents Will Lead to Enterprise AI Agent Failure, press release of 26 May 2026.
30. OpenID Foundation, AuthZEN Working Group Drafts: Access Request and Approval Profile, 15 June 2026.
31. Microsoft Learn, Overview of permissions and consent in the Microsoft identity platform, updated 18 March 2025.
32. OpenID Foundation, Authorization API 1.0, Final Specification, approved in January 2026, announced on 12 January 2026.
33. BSI, implementation notes on module ORP.4 Identitäts- und Berechtigungsmanagement, edition 2022.

34. Center for Internet Security, CIS Controls v8.1, Safeguard 5.5, Establish and Maintain an Inventory of Service Accounts, retrieved September 2026.
35. Ridgeline Cyber, Service Principal Ownership Is the Attack Path Nobody Governs, 12 May 2026.
36. Cloud Security Alliance, Agentic AI Identity and Access Management: A New Approach, 18 August 2025.
37. Microsoft Entra Blog, AI agents are everywhere, are your access controls ready, 15 July 2026.
38. CISA and partner agencies, Careful Adoption of Agentic AI Services, 1 May 2026, page 15.
39. Microsoft Learn, What's new in Microsoft Entra Agent ID, as of 1 May 2026.
40. Microsoft Learn, Use Microsoft Entra agent identities in Copilot Studio, updated 15 September 2026.
41. Microsoft Learn, Microsoft Entra Agent ID governance overview, 5 June 2026.
42. Microsoft Learn, Microsoft Entra Agent ID logs, updated 3 September 2026.
43. Microsoft Learn, Agent owners, sponsors, and managers, 16 April 2026, updated 21 July 2026.
44. Microsoft Learn, Conditional Access for agents, updated 30 September 2026.
45. Microsoft Learn, Continuous access evaluation for workload identities, 22 April 2026.
46. Microsoft Learn, Risky agents in Microsoft Entra ID Protection, 17 June 2026.
47. Microsoft Learn, Access packages for agent identities, 1 May 2026.
48. Microsoft Learn, Best practices for Microsoft Entra Agent ID, retrieved September 2026.
49. Microsoft Learn, Lifecycle workflow tasks for agent sponsors, updated 3 September 2026.
50. SailPoint Developer Community, New capability: Agent Identity Security, 6 October 2025.
51. SailPoint, press release of 4 August 2026, and Developer Community, New product: SailPoint Agentic Fabric, 4 August 2026.
52. SailPoint Documentation, Managing AI agents, Identity Security Cloud, retrieved September 2026.
53. SailPoint Developer Community, Agentic Fabric, section on the IdentityIQ plugin integration, 4 August 2026.
54. SailPoint Community, IdentityIQ 8.5 Release Notes, July 2025.
55. SailPoint, announcement of IdentityIQ 9.0 at Navigate 2026, 6 October 2026.
56. SailPoint Documentation, IdentityIQ, RPA/Bot identities, retrieved September 2026.
57. SailPoint Documentation, IdentityIQ Rapid Setup, Leaver configuration, retrieved September 2026.
58. Microsoft Learn, Microsoft Purview for Agent 365, 1 May 2026.
59. Palo Alto Networks, 2026 Identity Security Landscape, May 2026 (109 to 1, survey); Entro Security Labs, 2025 State of Non-Human Identities (92 to 1, own data and surveys); Cloud Security Alliance, commissioned by Astrix, State of Non-Human Identity Security, September 2024 (20 to 1, survey). All figures collected or commissioned by vendors.

We answer the questions from chapter 9 together with you.

WertFels GmbH advises banks, insurers and other regulated companies on identity and access management independent of vendors, from the inventory of non-human identities to control models and authorisation concepts for AI agents and the mapping against DORA and the AI Act. We work with Keycloak, Microsoft Entra ID, Omada, One Identity Manager, SailPoint IdentityIQ, SailPoint Identity Security Cloud and Saviynt, as well as PAM solutions such as CyberArk.



WertFels GmbH
info@wertfels.com
+49 89 244102321
wertfels.com

This paper reflects the state of September 2026, with corrections from October 2026. Legal statements are interpretations by WertFels GmbH and do not replace legal advice. Product information comes from vendor documentation as of the date stated in each case.

