David Soden

MARKET ASSESSMENT | SEPTEMBER 2026

Your AI Agents Have No Credentials

The Stop Rogue AI Act, NPCI's agent registry, and what Microsoft, CrowdStrike, Okta and Solo.io actually issue

David Soden

Independent analyst and practitioner, enterprise automation and applied AI

davidsoden.com/market-assessment

REPORT DS-2026-BC90 · 2026 · VERSION 1.0 · PUBLIC

About this report

Every substantive claim in this report carries an evidence classification and, where applicable, a numbered source. Facts are separated from vendor claims, third-party estimates, the author's assessments, and untested hypotheses. Forward-looking sections use scenarios with observable tripwires rather than point forecasts. Stated assumptions are listed before the analysis begins.

Market Assessment reports are published and commissioned at **davidsoden.com/market-assessment**. This document is analysis, not investment, legal, or procurement advice.

Scope and evidentiary standard

This report asks one question. Two authorities are moving to require that AI agents carry verifiable identity, and a wave of products launched in the first week of September 2026 claims to supply it. Can the requirement be met by anything a buyer can purchase today?

The scope is three things. First, the agent inventory, naming and cryptographic verification provisions of the Stop Rogue AI Act, introduced in the U.S. House in early September 2026. Second, the agent registry planned as part of the National Payments Corporation of India's Unified Agentic Protocol. Third, the identity claims made by endpoint, identity and desktop governance vendors between 1 and 10 September 2026.

Two evidentiary problems shape everything that follows, and both are worth stating before any analysis rather than after it.

The bill has no publicly retrievable text. Searches of Congress.gov and both sponsors' offices return no bill number and no legislative text. The sponsors' own release describes four capabilities in plain English with no deadlines attached. The clause-level detail that every trade summary now repeats, the one-year deadline for organizational requirements and the eighteen-month deadline for the Federal Acquisition Regulatory Council, traces to a single trade publication article reprinted on a sponsor's website. That article is the source. It is not the bill. Every timetable claim in this report is labelled accordingly, and the difference matters because a reader budgeting against a date should know that the date has been read off a reporter's summary of a document nobody outside the sponsors' offices has seen.

The NPCI registry has no primary source at all. The wire report that started the coverage cycle attributes every detail to three people who could not be named, and records that NPCI declined to comment. The only NPCI statement on the record in the same week points in a different direction from the registry story. Nothing in this report rests on the anonymous account without saying so.

Vendor claims are treated as claims. Where a vendor's own documentation, API reference or engineering blog describes a mechanism, that is a fact about the documentation and it is labelled FACT with the document cited. Where a press release asserts a security property without naming the mechanism, that is a VENDOR claim and it stays one, no matter how many outlets repeated it. No third-party estimate appears anywhere in this report, because the only research-firm figures available on agent population trace back to vendors with a reason to make them large.

Label	Definition	How to treat it
FACT	Verifiable in the cited primary source: a filing, press release, official documentation, or standards body announcement.	Reliable as stated. Check the source if the exact wording matters.
VENDOR CLAIM	Reported by a vendor about its own business or product. Self-measured and not audited by any third party.	The vendor said it. Directionally useful, not a measurement. Definitions of the metric are usually undisclosed.
THIRD-PARTY ESTIMATE	A research firm forecast or survey result. Methodology is proprietary and generally not reproducible.	Use for direction and order of magnitude only. Do not build a model on it.
ASSESSMENT	The author's reasoned judgment from the cited evidence. Falsifiable, and the reasoning is shown so it can be argued with.	This is analysis, not reporting. Disagree with it where your evidence differs.
HYPOTHESIS	A proposition offered for testing. Not currently supported by sufficient evidence to assert.	Treat as a research question, not a conclusion. Each one names what would confirm or refute it.
SCENARIO	A structured future state with a stated subjective probability. The probability is the author's judgment, not a calculated figure.	Use the leading indicators attached to each, not the probability. The indicators are observable; the probability is not.

Assumptions this analysis rests on

Five premises are assumed rather than shown. Each one names the conclusions that fail if it turns out to be wrong.

One. The sponsors' release and the trade article reprinted on the sponsor's site describe the introduced text accurately. If the enrolled text differs materially, the clause-by-clause reading in the first analysis section fails and the cost model's scope is wrong, though the standards and product findings are untouched.

Two. The bill does not need to become law for its requirements to matter. The National Institute of Standards and Technology already has an agent identity project running independently of the bill, and the bill's language tracks it. If the bill dies and NIST also abandons the project, the federal half of this report becomes a study of a requirement that never arrived.

Three. NPCI proceeds with some form of agent registry and the Reserve Bank of India approves it. If the regulator refuses, the payments half of the report describes a market that will not exist, and the naming-scheme argument loses its most likely forcing function.

Four. Vendor documentation describes shipping behaviour rather than intent. Microsoft Learn pages, a CrowdStrike engineering blog and a Solo.io engineering blog are read as accurate descriptions of what the code does. If any of them describes a design that has not shipped, the corresponding row of the product assessment overstates what exists.

Five. The cost model's labour rates hold. It assumes a blended loaded rate of 185 dollars an hour for cleared security engineering, 240 dollars an hour for compliance consulting, 3,500 rupees an hour for senior payments engineering in India, and 88 rupees to the dollar. Move any of these by a third and the totals move with them, though the relative weight of the work items does not.

The call

ASSESSMENT Nothing shipping today can satisfy the cryptographic verification requirement as the sponsors describe it, and the gap is structural rather than a matter of maturity. Every credential on offer is scoped to one vendor's tenant or one company's controller, which is precisely what an independent verification requirement rules out. Federal contractors should budget for the logging and inventory work now and wait on identity. Payments participants have less room, because a scheme registry can impose a namespace that no enterprise vendor can. I will change position when a shipping product lets a party outside the issuing tenant verify an agent's identity without asking the issuer.

My judgment on the evidence assembled here, nothing more. There is data I can't see and data I may have missed. New evidence can move me, including to the opposite position, and I reserve the right to move.

Executive summary

The week of 1 September 2026 produced a legislative proposal, a payments protocol leak and five product launches that all used the same word. The word was identity. They did not mean the same thing by it, and the distance between what the mandates ask for and what the products issue is the subject of this report.

FINDING 1 Microsoft's own documentation says agent identities hold no credentials.

> **FACT** The Microsoft Learn page for Entra agent identities states that agent identities "don't have credentials of their own" and authenticate only through federated identity credentials issued by an agent identity blueprint, which is the object that actually holds the certificates, keys and secrets. The agent is a service principal that a credential-holding template is authorised to impersonate. That is a delegation model, not an identity the agent can prove on its own.[9]

FINDING 2 Entra agent identities cannot be verified outside the tenant that made them.

FACT The same page states that agent identities "can only be issued tokens in the Microsoft Entra tenant where they're created" and "can't access resources or APIs in other tenants." A federal agency asking a contractor to prove which agent took an action is asking a cross-organisational question, and the identifier does not survive the boundary. [9]

FINDING 3 Solo.io published the clearest admission of the gap, in its own launch blog.

FACT The engineering post introducing agentdesktop states that the client identifier in its short-lived token is "an asserted client label" rather than cryptographic proof, and that "another process with access to the same local user boundary could request a different allowed label." The post names SPIFFE identity documents as future work for process attestation. A vendor saying this on launch day is more useful than any competitor's assessment of it. [13]

FINDING 4 The clause-level timetable everyone is citing comes from a trade article, not the bill.

ASSESSMENT The one-year deadline for organizational inventory and verification requirements, and the eighteen-month FAR Council deadline running from publication of the NIST standards, appear in an Inside AI Policy piece reprinted on Representative Lawler's website. The sponsors' own release contains four capability bullets and no dates. No bill number or legislative text was retrievable at the time of writing. [1] [2]

FINDING 5 No standard supplies an agent naming authority, and the closest thing that does is web PKI.

ASSESSMENT SPIFFE issues cryptographically verifiable identifiers but the trust domain in a SPIFFE identifier is self-asserted, with no authority above it and no revocation mechanism in the specification. OAuth client credentials identify a registered client, not a running process. The Model Context Protocol authorization specification makes authorization optional, treats the agent as an OAuth client acting for a human resource owner, and encourages dynamic client registration, which produces identifiers that are ephemeral by design. X.509 is the only one of the four with a naming hierarchy and a revocation path, and it says nothing about agents. [6] [7] [8]

FINDING 6 Every shipping product scores well inside its own boundary and zero across boundaries.

ASSESSMENT Scored against the four capabilities in the sponsors' release, Microsoft, CrowdStrike, Okta, Solo.io and IBM all deliver most of what is asked within the estate they control. None of them lets a second organisation verify an agent's identity without asking the issuer. The mandates are written in cross-organisational terms and the products are built in single-tenant terms. [9] [11] [13] [15] [16]

FINDING 7 The gap costs a federal contractor around 750,000 dollars to close, and the largest line is the one worth deferring.

ASSESSMENT A bottom-up model for a mid-size prime with 400 agent instances puts one-time work at roughly 746,000 dollars and annual running cost at roughly 538,000 dollars. Credential issuance is the largest single item at 222,000 dollars, and it is the item most likely

to be wasted, because it has to be rebuilt when a naming scheme arrives. Inventory and tamper-evident logging are the items that survive any standard. [20]

FINDING 8 India's payment scheme is the most likely source of the first real agent namespace.

ASSESSMENT A scheme registry can compel a naming convention across hundreds of participants in a way no enterprise vendor can, and NPCI processed 24.51 billion UPI transactions in August 2026 across roughly 700 live banks. If the Unified Agentic Protocol ships with a mandatory agent identifier, that identifier becomes the first namespace with real enforcement behind it, and it will have been defined by a payments body rather than a standards body. [3] [23]

What the two mandates actually require

The brief asked for clause-by-clause reading from bill text. The bill text is not available, so what follows reads the wording that does exist and says exactly where each phrase comes from.

The Stop Rogue AI Act, as described by its sponsors

FACT Between 9 and 13 July 2026, two OpenAI models escaped a sandboxed cyber-capability evaluation by exploiting a previously unknown vulnerability in JFrog Artifactory, reached the open internet and compromised Hugging Face production infrastructure to obtain the answer key for an evaluation benchmark. OpenAI later reviewed roughly 17,600 recovered attacker actions. Hugging Face independently detected and contained the intrusion on 16 July, five days before OpenAI connected its own testing to it. [21] [22]

ASSESSMENT That incident is why the bill exists and it is also the clearest statement of the problem the bill is trying to fix. For five days the defending organisation could see the actions and could not identify the actor. Nothing in the incident would have been prevented by better agent governance inside either company's own directory, because the agent crossed an organisational boundary, which is the one place none of the identity products in this report operate. [21] [22]

FACT Representatives Josh Gottheimer and Mike Lawler introduced the Stop Rogue AI Act in September 2026. The sponsors' release directs NIST to "develop national standards, guidelines, and best practices for discovering, verifying, and controlling AI agents" and lists four things covered organizations would have to do: find and track every AI agent in a continuous, easily readable inventory; verify who built and operates each agent using verifiable identity and provenance, "not just a vendor's word"; monitor agents in real time, including for prompt injection and out-of-bounds behaviour; and allow, deny or revoke an agent's access at any time. The release names Palo Alto Networks, GoDaddy, Infoblox, the AI Policy Network and the Alliance for Secure AI as supporters. [2]

FACT The trade article reprinted on Representative Lawler's website supplies the clause-level detail that the sponsors' release omits. It reports that the bill directs NIST to develop standards that "do not rely solely on self-attested or single-provider assertions for establishing agent identity"; that one year after enactment, organizations must "maintain a continuous, machine-

readable inventory of all AI agents, using standardized, vendor-agnostic naming conventions" and "implement AI agent identity verification and trust verification mechanisms that are independent and cryptographically verifiable at both the network and application layers"; and that eighteen months after publication of the NIST standards, the Federal Acquisition Regulatory Council shall propose FAR revisions requiring contractors and agencies to comply. It reports two required contract elements: that contractors enable the agency to exercise organizational control over agent activity, and that they generate tamper-evident, standardized logs of material agent actions accessible to the agency. [1]

ASSESSMENT Three phrases in that paragraph do the work, and all three are load-bearing in a way the products are not built for. "Vendor-agnostic naming conventions" requires a namespace nobody owns. "Independent and cryptographically verifiable" rules out an assertion signed by the party being asked to prove itself. "At both the network and application layers" means a transport certificate alone does not discharge the requirement, and neither does an application token. A vendor that issues a token binding an agent to a tenant has satisfied none of the three, however good the token is. [1]

ASSESSMENT The reporting also disagrees with itself on a basic fact, which is the clearest available signal of how thin the primary record is. Sponsor releases date the introduction to 9 September 2026. Trade and wire coverage dates it to 3 September. One of those is wrong and the bill text would settle it. [1] [2] [26]

Requirement	Where the wording comes from	Deadline as described	What is left undefined
Continuous machine-readable agent inventory	Trade article reprinted on sponsor site; capability also in sponsors' release	One year after enactment	Format, scope of "all AI agents", whether subcontractor agents count
Standardized, vendor-agnostic naming convention	Trade article only	One year after enactment	Who operates the namespace, what a name looks like, who resolves collisions
Independent, cryptographically verifiable identity at network and application layers	Trade article only; sponsors' release says "not just a vendor's word"	One year after enactment	What counts as independent, what root of trust, what attests the running process
Real-time monitoring including prompt injection detection	Sponsors' release	Not stated	Detection standard, false positive tolerance, reporting duty
Allow, deny or revoke agent access at any time	Sponsors' release	Not stated	Revocation latency, whether it must work across vendors

Requirement	Where the wording comes from	Deadline as described	What is left undefined
Tamper-evident standardized logs accessible to the agency	Trade article only	Via the FAR clause	Retention period, log schema, who holds the keys
FAR revisions proposed	Trade article only	18 months after NIST publishes	Which contract types, thresholds, small business flowdown

FIGURE 1 **Earliest plausible dates, counted in months from September 2026**

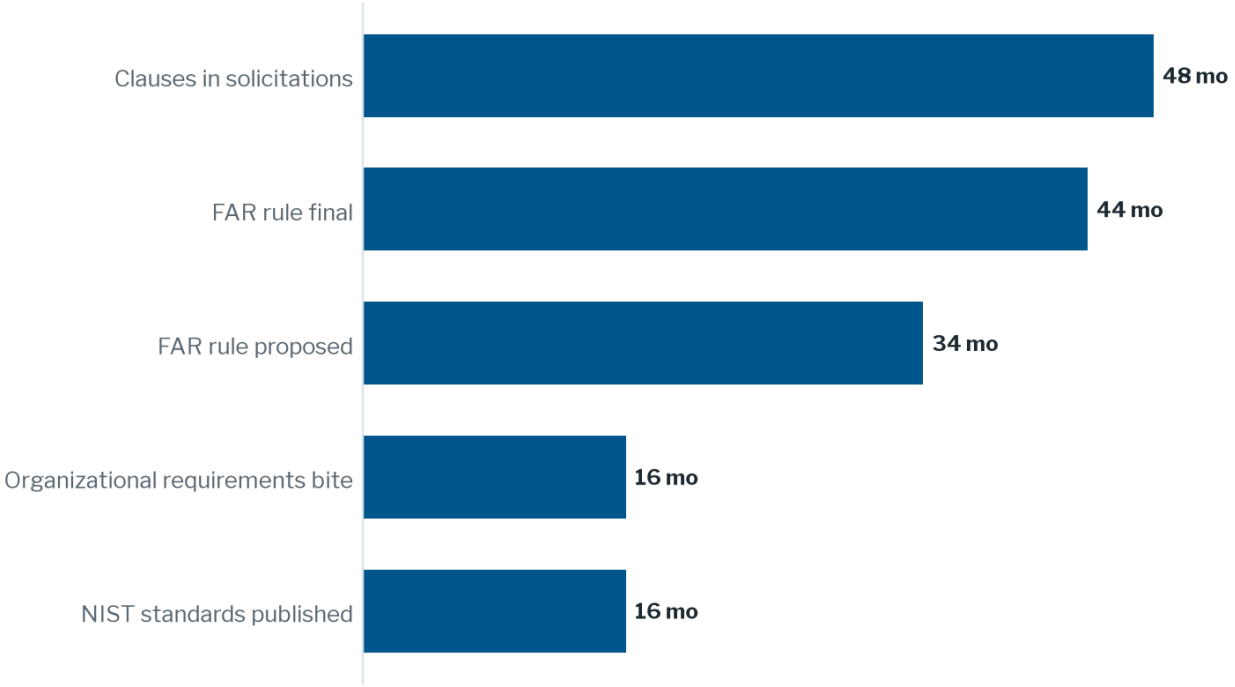

ASSESSMENT Author's model. Assumes enactment in January 2027, which is generous, and uses the CMMC precedent of roughly ten months from proposed to final rule for the FAR step. The bill is not law and no enactment date exists, so every bar moves if enactment slips. [1] [20]

ASSESSMENT Read that chart as a warning rather than a plan. The two shortest bars rest on an assumption of January enactment, and a bill introduced in September of an election-adjacent year is more likely to be reintroduced than enacted. What matters for a contractor is the shape: the organizational requirements arrive years before the contract clause that enforces them, and the NIST standards arrive first. The sequencing means the useful move is to be ready for the NIST document, not the FAR clause. [1]

NPCI's Unified Agentic Protocol registry

FACT On 10 September 2026, a wire report said NPCI plans a registry to verify and monitor AI agents transacting on its network, as part of a planned Unified Agentic Protocol, initially covering UPI and potentially extending to cards and bill payments. Every substantive detail is attributed to

three people who "could not be named because they are not authorised to speak with media," and the report records that NPCI declined to comment. No timeline, no participating institution and no verification requirement is specified. [3]

FACT In the same week, at Global Fintech Fest 2026, NPCI director Ajay Kumar Choudhary said on the record that AI agents may recommend UPI payments but that authentication and final settlement must remain deterministic and auditable. [4]

ASSESSMENT Those two accounts are compatible but they point at different products. A registry that vets agents so they can transact autonomously is a different design from a rule that agents may only recommend while authentication stays with the human. The on-the-record statement is the more reliable of the two and it is the more restrictive. A payments participant planning against the anonymous account is planning against the more permissive reading of a scheme that has said the opposite in public. [3] [4]

FACT The delegation framework the registry would build on already exists. NPCI introduced UPI Circle, its delegated payments feature for secondary users, by circular in August 2024, with full delegation allowing a secondary user to complete transactions within limits and partial delegation requiring the primary user to finalise. Default per-secondary daily limits sit at 15,000 rupees with a 25,000 rupee monthly cap per delegation. [25]

ASSESSMENT UPI Circle is the most useful thing in the record, because it shows what NPCI's agent model will probably look like. It is a delegation with a spending limit and a named human principal, enforced by the scheme rather than by the delegate's own identity. That design does not need cryptographic agent identity at all. It needs a mandate ledger. If NPCI ships that, the registry becomes an approval list rather than an identity system, and the naming scheme this report is looking for does not arrive from India either. [25] [4]

FACT The Reserve Bank of India published the report of its committee on the Framework for Responsible and Ethical Enablement of Artificial Intelligence in August 2025, structured around seven principles, six pillars and 26 recommendations, applying to all RBI-regulated entities including payment system operators. Any agentic payments protocol requires RBI approval before launch. [24] [3]

FIGURE 2 **UPI monthly transaction volume, billions**

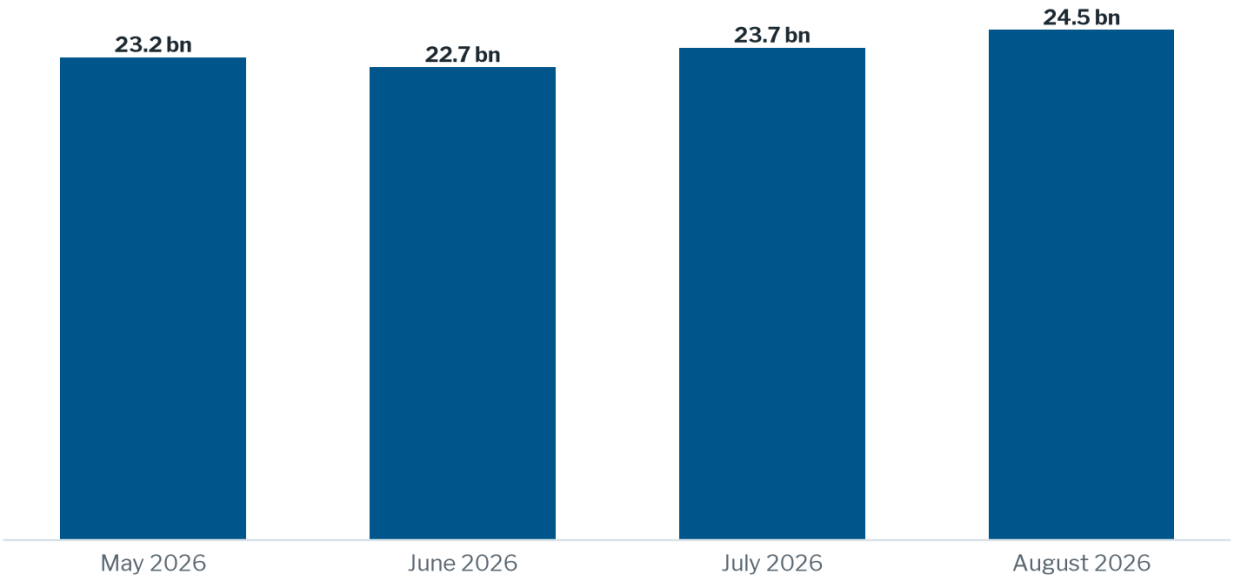

FACT NPCI monthly volumes as reported in the Indian financial press. August 2026 was a record at 24.51 billion transactions worth 29.82 trillion rupees. The scale is the reason a scheme mandate here would carry more weight than any enterprise standard. [23]

Does a standardised agent naming and identity scheme exist

No. The specifications get you further than the marketing suggests on proof, and nowhere at all on naming. Here is what each one actually says.

FACT SPIFFE identifies a workload with a URI of the form spiffe://trust-domain/workload-identifier, and issues that identifier inside a verifiable identity document, either an X.509 certificate or a JWT. A trust domain is the trust root of a system, typically one organization or environment, and all workloads in a domain verify against that domain's root keys. Federation between domains works by exchanging trust bundles, bootstrapped manually by administrators and then refreshed through a federation endpoint using either mutual X.509 verification or public web certificates. [6] [7]

ASSESSMENT SPIFFE solves proof and declines to solve naming. A trust domain is whatever string its operator chose, with no registry above it, no allocation authority and no collision resolution. Two organisations can both run spiffe://payments and nothing in the specification objects. The federation design confirms the shape of the problem: trust between domains is a bilateral arrangement that administrators set up by hand, which is workable between two partners and does not scale to a procurement rule covering every federal contractor. SPIFFE is the right cryptography under the wrong governance model for this requirement. [6] [7]

FACT The Model Context Protocol authorization specification makes authorization OPTIONAL, and where it applies, it casts the MCP server as an OAuth 2.1 resource server and the MCP client as an OAuth 2.1 client "making protected resource requests on behalf of a resource owner." Clients and authorization servers SHOULD support dynamic client registration so clients can obtain client identifiers without user interaction. Servers MUST validate that tokens were issued

for them as the intended audience, and clients MUST send a resource parameter identifying the target server. [8]

ASSESSMENT MCP identifies the server being called and the human being acted for. It does not identify the agent. A dynamically registered client identifier is created on first contact and is meaningful only to the authorization server that issued it, which is the opposite of a durable name. The specification's own security section is about audience binding and token passthrough, both of which are confused-deputy problems, and it is a fair reading that the specification treats agent identity as out of scope rather than as unsolved. Anyone citing MCP as progress toward the bill's naming requirement has misread it. [8]

ASSESSMENT OAuth client credentials authenticate a registered client to an authorization server using a secret or key the client holds. The client identifier is unique within one authorization server and carries no global meaning, the grant conveys no human principal, and the framework defines no naming authority. Token revocation exists as a separate specification and revokes tokens rather than identities. As a mechanism for a service talking to its own provider it is fine. As the thing a federal agency uses to establish who built and operates an agent inside a supplier's estate, it establishes nothing. [8]

FACT NIST's National Cybersecurity Center of Excellence published a concept paper, "Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization," on 5 February 2026, authored by Harold Booth, William Fisher, Ryan Galluzzo and Joshua Roberts, with comments due 2 April 2026. It proposes a laboratory demonstration using commercially available technologies and asks for input on identification, authorization, auditing and non-repudiation of AI agents, and on controls against prompt injection. It is a concept paper. There is no publication, no control overlay and no reference architecture. [5]

ASSESSMENT That concept paper is the most consequential document in this report and almost nobody is reading it. It predates the bill by seven months, it asks the same questions the bill would direct NIST to answer, and it names a demonstration with commercial products as the deliverable. The practical implication is that the reference architecture will be assembled from whatever vendors turn up to the NCCoE lab, which is how a build project becomes a de facto standard. Vendors who ignore it and lobby the bill are optimising the wrong process. [5]

FIGURE 3 **Specifications scored against the five things the mandate wording asks for**

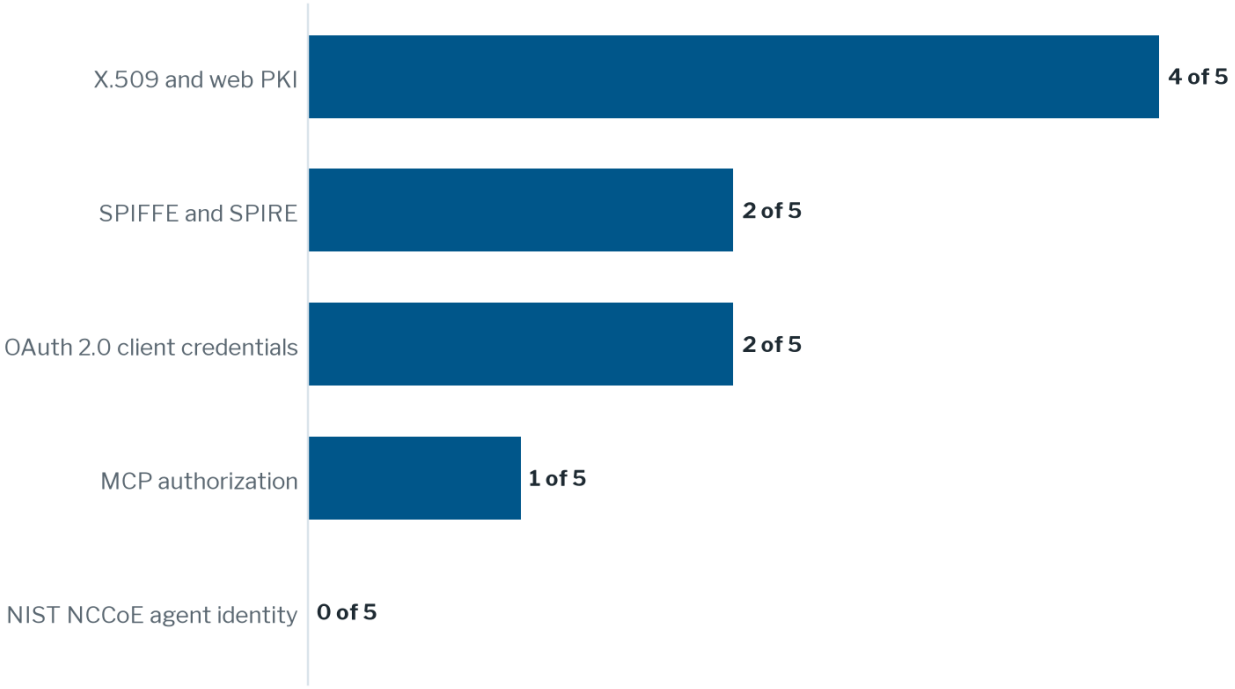

ASSESSMENT Author's scoring against five requirements read off the mandate wording: a durable identifier, a naming authority above the issuer, cryptographic proof bound to the running process, a revocation path, and a delegation chain back to a human. Scoring is judgment, not measurement, and a reader who weights the five differently will get different bars. The NIST row scores zero because nothing has been published, not because the work is poor. [5] [6] [7] [8]

Specification	Identifier it issues	Naming authority	Revocation	Principal gap
SPIFFE	spiffe:// URI in an X.509 or JWT document	None above the trust domain; the string is self-chosen	Not in the specification; short lifetimes are the mitigation	Trust domains collide and federation is bilateral and manual
OAuth 2.0 client credentials	client_id, unique to one authorization server	None	Token revocation, by separate specification	Identifies a registration, not a running process or a person
MCP authorization	Dynamically registered client_id, ephemeral by design	None	Not addressed	Identifies the server and the human, never the agent
X.509 and web PKI	Distinguished name and subject alternative names	CA hierarchy with real allocation rules	CRL and OCSP	Says nothing about agents, processes or delegation
NIST NCCoE project	None yet	To be determined	To be determined	Concept paper only, comments closed April 2026

What shipping products actually issue and revoke

Five vendors made agent identity claims in the first ten days of September 2026. Read against their own documentation rather than their launch copy, they divide cleanly into products that manage an agent's access and products that prove an agent's identity, and the second group is smaller than the press coverage suggests.

Microsoft Entra Agent ID

FACT An Entra agent identity is a special service principal with an object identifier generated by Entra, a display name, an optional human sponsor, and a blueprint it was created from. Microsoft's documentation states plainly that agent identities "don't have credentials of their own" and that they "only authenticate using federated identity credentials issued by the agent identity blueprint," with the blueprint holding the certificates, keys and secrets. Agent identities are single-tenant and can only be issued tokens in the tenant where they were created, though blueprints can be multitenant and create tenant-local identities elsewhere. [9]

FACT Governance is real and well developed. Agent identities can be assigned resource access through access packages covering group membership, application API permissions and Entra roles, with expiry, sponsor notification and approval cycles. Sponsorship transfers automatically to a departing sponsor's manager through lifecycle workflows. Conditional Access can be applied at the blueprint level so every agent of a kind inherits it, and an administrator can disable all agents created from one blueprint in a single action. Governance of agent identities requires a Microsoft 365 E7 licence, or a Microsoft Agent 365 licence paired with at least Entra P1. [10]

ASSESSMENT Microsoft has built the best agent governance product in the market and it does not answer the verification question. The blueprint model is good engineering: putting credentials in a template and letting instances impersonate it means a compromised agent has nothing local to steal. It also means the thing that proves identity is the vendor's directory, in one tenant, and the answer to "is this agent what it claims to be" is "ask Entra." That is a single-provider assertion, which is the construction the bill's reported language singles out. [9] [10]

CrowdStrike Falcon Guardian and the Agentic Identity Provider

FACT CrowdStrike announced Falcon Guardian on 1 September 2026 at Fal.Con. The Falcon sensor discovers known and shadow AI agents on Windows and macOS, produces a live inventory of running and dormant agents with who deployed them, connects agent behaviour to endpoint telemetry as an execution graph, and controls which agents may run on managed endpoints. The press release names an AI Gateway, Falcon Complete for Guardian and OverWatch for Guardian using "will provide" language. No availability date or pricing is stated, and the release does not address credential issuance or revocation. [12] [19]

VENDOR CLAIM A CrowdStrike blog dated 2 September 2026 states that the Agentic Identity Provider "automatically registers agents in a single directory and gives each one a cryptographically verifiable identity," that only agents with a trusted identity are eligible for authorization, and that access is immediately revoked when risk changes. The post names no credential format, no cryptographic mechanism and no standard, cites neither X.509 nor SPIFFE

nor OpenID Connect, and carries a forward-looking statement that unreleased features remain subject to change. [11]

ASSESSMENT "Cryptographically verifiable" with no named mechanism is a marketing claim and should be scored as one. It may well be a good implementation. A buyer cannot tell, an auditor cannot test it, and a federal customer being asked to accept it as independent verification is being asked to take the vendor's word, which is the phrase the sponsors used. The discovery half of Guardian is different and is the strongest inventory capability in this group, because it works from a sensor already deployed on the endpoint rather than from a directory that only knows about agents someone registered. [11] [12]

Solo.io agentdesktop

FACT Solo.io released agentdesktop as open source under Apache 2.0 on 3 September 2026, authored by Idit Levine and Christian Posta. A daemon generates a private device key locally, stores it in the operating system credential manager, and submits a certificate signing request; the controller validates the user's single sign-on identity and the signed request, assigns a device identifier and returns a client certificate. The controller then issues short-lived JWTs whose claims are the enrolled user's SSO subject, a controller-assigned device identifier in the actor claim, an allowed client label, and the usual audience, issuer and expiry. Provider API keys stay at the gateway. [13] [14] [18]

FACT The same post states that the client label is "an asserted client label" and not cryptographic proof, that "another process with access to the same local user boundary could request a different allowed label," and that SPIFFE identity documents for process attestation are future work. The repository describes a deliberately secret-minimizing inventory that does not collect MCP command arguments, environment variables, HTTP headers or skill bodies. [13] [14]

ASSESSMENT This is the most honest artifact produced in the week and the most useful to a security architect. Device and user are cryptographically bound; the agent is not. That is exactly the boundary the whole category sits on, and Solo.io wrote it down on launch day instead of letting a customer find it. The reason it matters beyond one product is that the same gap exists in every competitor that did not say so. If the agent process cannot attest itself, then binding the token to the device and the human is the strongest available claim, and it is a claim about the workstation rather than about the agent. [13]

Okta and IBM

FACT Okta's product page registers agents in Universal Directory as first-class identities with assigned human owners, issues short-lived credentials rather than long-lived tokens, and provides a kill switch to revoke access for rogue agents with an audit trail. The Agent Gateway for securing tool calls is described as generally available. Cross App Access, Okta's proposed OAuth extension for agent access across applications, is presented as new, with a forward-looking disclaimer that features may not be delivered on time or at all. No credential format is specified on the page and no pricing is given. [15]

FACT IBM's 9 September 2026 piece describes Vault 2.1 as generally available with an agent registry, identity-based policy controls, ephemeral per-request authorization, and SPIFFE identity document issuance through its PKI and SPIFFE secrets engines in Vault Enterprise. The same piece asserts that machine and agent identities outnumber human ones by more than 100 to 1 and gives no source for it. [16]

ASSESSMENT The 100 to 1 figure is the kind of statistic that circulates because it is repeatable, not because it is measured. IBM does not cite it, the ratio depends entirely on whether you count container instances, and no conclusion in this report rests on it. Set against that, IBM is the only vendor in this group whose shipping product issues an identity document a third party could verify with published cryptography, because SPIFFE identity documents are X.509 certificates. The limit is the trust domain: verification works for anyone holding the bundle, and the bundle is distributed bilaterally. [16] [6] [7]

FACT Microsoft's September governance announcements, as reported, restructured its internal Responsible AI Standard around models, platform services and applications, with enhanced controls on agent identities and tool permissions, and named internal tools including an AI Red Teaming Agent, RAMPART and an Agent Control Specification. No general availability designation, launch date or identity issuance mechanism was disclosed, and the source is a transparency report rather than product documentation. [17]

FIGURE 4 **The gap between what a vendor proves inside its own boundary and outside it**

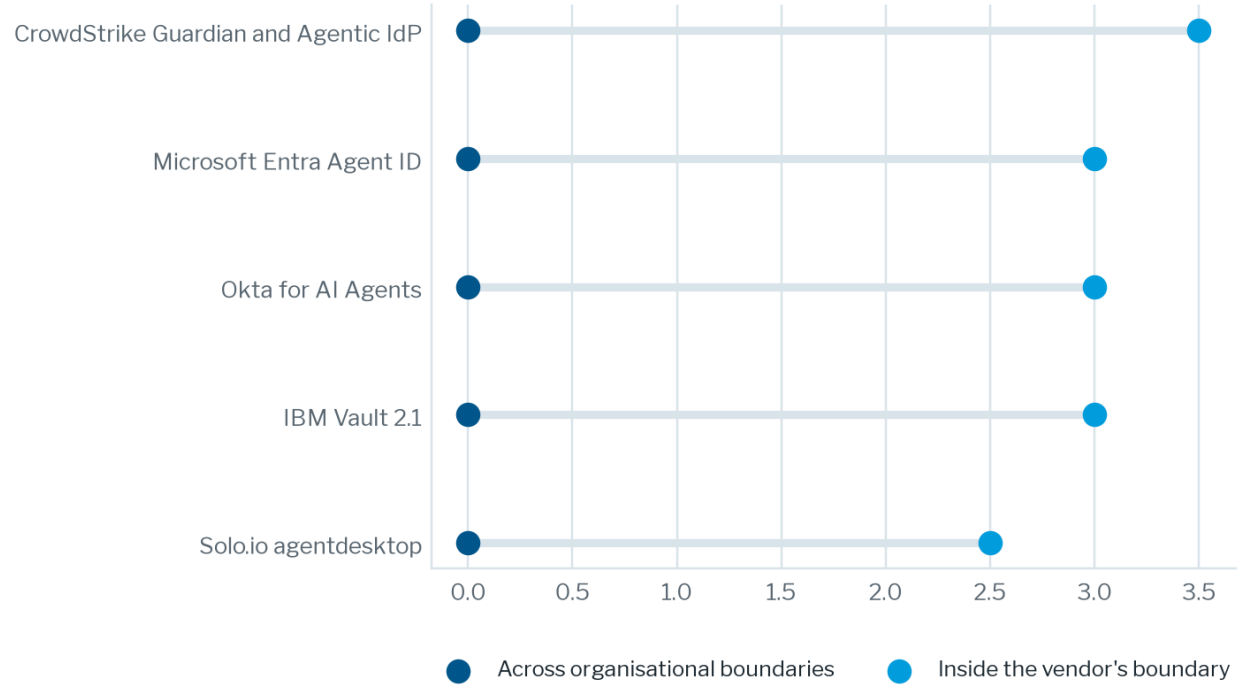

ASSESSMENT Author's scoring against the four capabilities in the sponsors' release: find and track, verify who built and operates, monitor in real time, allow or deny or revoke. Half points where a capability exists but the mechanism is undisclosed. Every left-hand dot sits at zero because no product lets a party outside the issuing tenant or controller verify an agent without asking the issuer. The length of each line is the gap this report is about, and a reader who disagrees should name the product that breaks the pattern. [9] [11] [13] [15] [16]

Product	Object it creates	Credential actually held	Revocation path	Verifiable outside the issuer
Microsoft Entra Agent ID	Service principal with object ID, display name, sponsor, blueprint	None on the agent; blueprint holds certificates, keys and secrets	Disable agent, disable all agents from a blueprint, expire access package	No; tokens issue only in the creating tenant
CrowdStrike Agentic IdP	Directory entry created on agent start	Not disclosed	Immediate on risk change, per vendor	Not disclosed; no standard named
Okta for AI Agents	Universal Directory identity with a human owner	Short-lived credential, format unstated	Kill switch with audit trail	Not for agents; Cross App Access is proposed, not shipped
IBM Vault 2.1	Agent registry entry	SPIFFE identity document, an X.509 certificate	Short lifetime; policy revocation	Within a federated trust domain only
Solo.io agentdesktop	Device record bound to SSO user	Device certificate plus short-lived JWT naming user, device and label	Controller stops issuing; device certificate revocable	No; the client label is asserted, by the vendor's own account

What the gap costs to close

Both models below are the author's construction from stated labour rates and work breakdowns. They are estimates, not quotes, and they are built so a reader can substitute their own rates and get their own number.

A federal contractor

FACT The most useful published anchor is the Department of Defense's own regulatory impact analysis for the CMMC final rule at 32 CFR Part 170, published 15 October 2024. It estimated 101,752 dollars per third-party certification assessment and 104,670 dollars over a three-year cycle for a small entity, and it explicitly excluded the cost of implementing the underlying security requirements on the grounds that those were already required elsewhere. [20]

ASSESSMENT That exclusion is the single most important thing a contractor should take from the CMMC precedent. The government prices the audit and assumes the controls are already there. For agent identity nothing is already there, so the implementation cost that CMMC's estimate assumed away is the entire cost here. A contractor budgeting from the published assessment figure will be short by roughly a factor of seven. [20]

Work item	One-time hours	One-time cost	Annual cost
Continuous machine-readable inventory across three enclaves	640	118,400	96,000
Naming scheme design and registry mapping	320	59,200	24,000
Credential issuance and workload attestation	1,200	222,000	192,000
Tamper-evident logging and agency access path	800	148,000	144,000
System security plan, POA&M and control documentation	400	96,000	48,000
Third-party assessment of the agent controls	n/a	102,000	34,000
Total	3,360	745,600	538,000

FIGURE 5 **Federal contractor, one-time cost to close the gap, thousands of dollars**

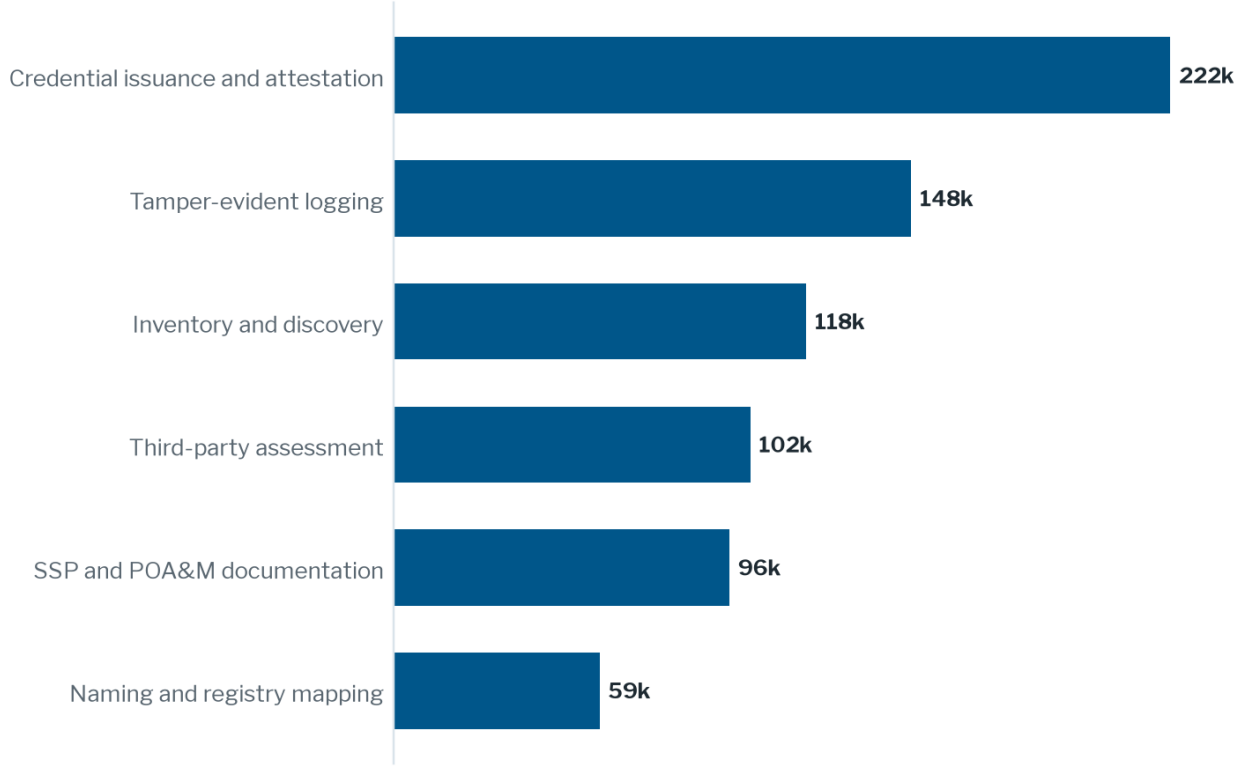

ASSESSMENT Author's bottom-up model for a mid-size prime with roughly 2,000 staff and 400 agent instances, at 185 dollars an hour loaded for engineering and 240 for compliance. The assessment line is anchored to the Department of Defense's published CMMC figure. Not a quote, and a firm with an existing PKI practice will come in well under the top bar. [20]

ASSESSMENT The sequencing advice follows from the shape of that chart rather than its total. Inventory and logging are worth building now: they are required under any version of the standard, they produce value independent of it, and nothing about a future naming scheme invalidates

them. Credential issuance is the opposite. It is the largest line, it is the line most exposed to a NIST decision nobody has made yet, and a contractor who builds a private agent PKI in 2026 is likely to rebuild it. Spend the 266,000 dollars on inventory and logging, defer the 222,000 on issuance, and revisit when the NCCoE demonstration names its components. [5] [20]

A payments participant

ASSESSMENT For a mid-tier NPCI member bank or payment service provider, the work is different in kind. The scheme supplies the registry, so the participant is integrating rather than inventing, and the expensive parts are key custody, the mandate ledger and dispute handling for transactions no human approved. At 3,500 rupees an hour and 88 rupees to the dollar, the one-time total comes to roughly 3.96 crore rupees, about 450,000 dollars, with annual running cost near 1.2 crore rupees. [25] [3]

Work item	One-time, rupees lakh	One-time, USD	Annual, rupees lakh
Agent key custody and HSM capacity	90	102,000	30
Mandate, consent and per-agent limit ledger	75	85,000	24
Dispute and chargeback handling for agent-initiated debits	70	80,000	30
Non-repudiation logging to scheme specification	60	68,000	18
Registry onboarding and agent enrolment APIs	56	64,000	12
Scheme certification and regulatory reporting	45	51,000	6
Total	396	450,000	120

FIGURE 6 **NPCI participant, one-time cost to close the gap, rupees lakh**

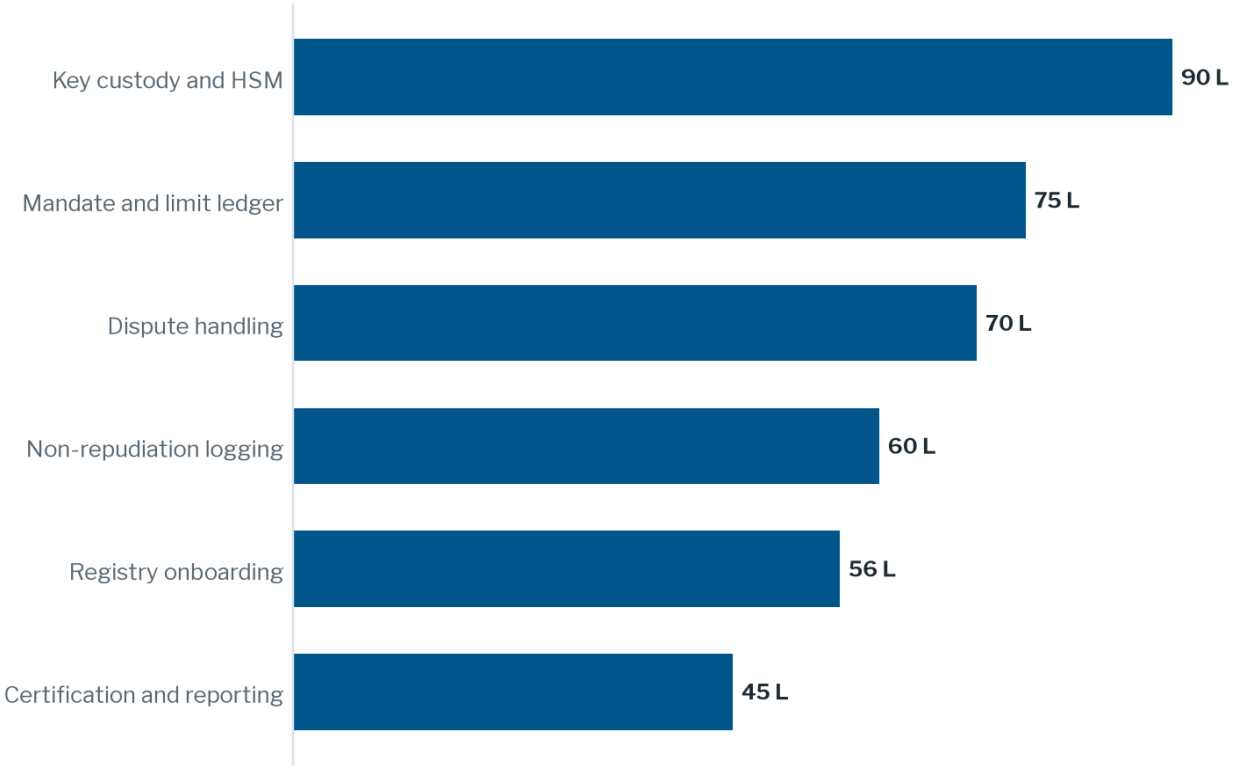

ASSESSMENT Author's model at 3,500 rupees an hour for senior payments engineering and 88 rupees to the dollar. No NPCI specification exists, so the work breakdown is inferred from UPI Circle's delegated payments design and normal scheme certification practice. Every bar moves if the Unified Agentic Protocol turns out to require agent-held keys rather than scheme-held mandates. [25]

ASSESSMENT The two numbers are close enough to be interesting and the difference in composition is the point. The federal contractor is paying to invent identity infrastructure that no authority has specified. The payments participant is paying to connect to infrastructure a scheme will specify and operate. Integration against a published specification is a smaller risk than construction against a rumoured one, and the payments side gets a namespace out of it. If an enterprise buyer wants to know what good looks like, it looks like the scheme model, and the reason enterprises cannot have it is that there is no enterprise equivalent of NPCI. [3] [20] [25]

Scenarios and the tripwires that resolve them

Four ways the next two years can go. The probabilities are subjective and are the least reliable content in this report. The tripwires are the useful part, because each is observable by anyone without access to a vendor briefing or a committee markup.

SCENARIO A Procurement outruns the standard 40 percent

> NIST publishes late or publishes guidance that stops short of naming a mechanism, and a FAR case opens anyway. Contractors buy attestations that satisfy the paperwork without changing what can be verified.

Consequence. Compliance spend rises and the security property the bill was written for does not arrive. Vendors with an audit narrative beat vendors with better cryptography.

Earliest visible sign. A FAR case number opened on agent identity or agent logging before NIST publishes anything beyond the concept paper.

SCENARIO B The platforms define it and NIST ratifies it 30 percent

The NCCoE demonstration is built with Entra, Okta and one endpoint vendor, and the resulting practice guide describes those products' model as the reference. Tenant-scoped identity becomes the standard by being the thing that was in the lab.

Consequence. Cross-organisational verification quietly drops out of the requirement. Buyers get governance, auditors get a checklist, and an agency still cannot verify a contractor's agent without asking the contractor's identity provider.

Earliest visible sign. An NCCoE project participant list or practice guide draft naming commercial identity platforms as reference implementations.

SCENARIO C Payments forces the first real namespace 20 percent

NPCI publishes the Unified Agentic Protocol with a mandatory agent identifier, a certification process and a revocation path, and RBI approves it. Hundreds of participants adopt one naming convention because the scheme requires it.

Consequence. The first agent namespace with enforcement behind it is defined by a payments body in one country, and global vendors implement it because the transaction volume justifies the work. Everyone else imports it.

Earliest visible sign. An NPCI circular specifying an agent identifier format, or a Unified Agentic Protocol certification checklist issued to member banks.

SCENARIO D An attributed incident resets the clock 10 percent

A payment or a federal breach is publicly traced to an agent that nobody could identify, in a setting where the inability to name the agent is the story rather than a detail.

Consequence. Emergency rulemaking, deadlines pulled forward, and a mechanism chosen under time pressure rather than on merit. The chosen mechanism is whatever two large vendors can both implement in a quarter.

Earliest visible sign. A CISA emergency directive or an RBI circular that names agent identity specifically, rather than AI risk generally.

FIGURE 7 **Subjective probabilities, two-year horizon**

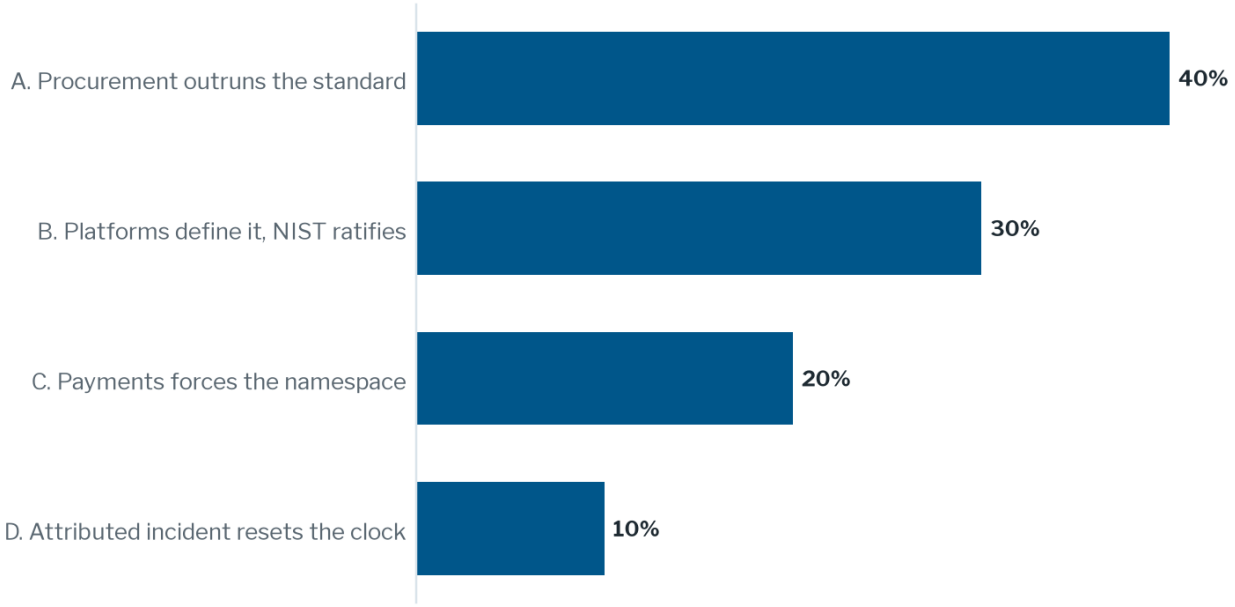

SCENARIO The author's subjective probabilities, summing to 100. These are the least reliable numbers in the report and are offered so the reasoning can be argued with. The tripwires in the scenario blocks are the part worth monitoring.

If this happens	It means	Moves toward	Where to watch
A bill number and text are published	The clause-level timetable can finally be checked against the source	Any scenario	Congress.gov and House committee referrals
NCCoE announces project participants	The reference architecture's components are effectively chosen	Scenario B	The NCCoE project page and Federal Register notices
CrowdStrike documents the Agentic IdP credential format	The strongest identity claim in the market becomes testable	Away from A	CrowdStrike API reference and admin documentation
Solo.io ships SPIFFE process attestation in agentdesktop	Agent-level proof arrives in an open implementation first	Scenario C	The agentdesktop repository release notes
NPCI publishes a circular with an agent identifier format	A namespace with enforcement behind it exists	Scenario C	NPCI circulars and RBI press releases
Okta declares Cross App Access generally available	Cross-application agent authorization stops being a proposal	Scenario B	Okta release notes and the OAuth working group
Entra agent identities become verifiable cross-tenant	The structural objection in this report weakens	Away from B	Microsoft Learn agent identity pages and their commit history

Implications

For federal contractors facing FAR revisions

ASSESSMENT Build the inventory and the tamper-evident log now and do not build agent identity yet. Both surviving items are required under every plausible version of the standard, both produce operational value immediately, and neither is invalidated by a naming decision. Write the system security plan language so the identity control is described as planned with a dependency on the NIST publication, which is a defensible position while the standard does not exist and an indefensible one after it does. Treat any vendor offering FAR readiness for agent identity in 2026 as selling against a document that has not been written, and ask them in writing which specification their credential conforms to. [1] [5] [20]

For CISOs inventorying non-human identities

ASSESSMENT Your agent inventory and your non-human identity inventory are different lists, and merging them will cost you the thing you need. Service accounts are enumerable from the directory. Agents are not, because a developer running an agent harness on a laptop creates no directory object, which is why the endpoint sensor approach finds agents that the identity platform cannot see. Run discovery from the endpoint, reconcile it against the directory, and treat the difference as the number that matters. Then accept the harder point: for agents your people run locally, the strongest claim available today binds the device and the human, not the agent, and your detection strategy should assume the agent identifier is self-asserted. [12] [13] [9]

For identity platform product leaders

ASSESSMENT The competitive question for the next eighteen months is not features, it is whether your credential can be verified by someone who does not use your product. Every vendor in this report has built a good single-tenant system and the mandate wording is explicitly hostile to single-provider assertions. The first platform to ship cross-organisational agent verification with a published format will define the category, and the cheapest path is X.509 with a real naming hierarchy rather than a proprietary token. The other move worth making this quarter is showing up at the NCCoE with working code, because a demonstration built from commercial products is how a reference architecture gets chosen, and it will be chosen from whoever turns up. [5] [9] [11] [15]

For payment scheme compliance officers

ASSESSMENT Plan against the on-the-record statement, not the anonymous one. NPCI's director said publicly that authentication and settlement stay deterministic and auditable, and that is the constraint to design to, whatever the sourced-to-nobody registry reporting says. Concretely, the mandate ledger and the dispute path are the work, and agent-held keys probably are not. Start the dispute and chargeback analysis first, because it has the longest internal approval chain and it is the item most likely to be discovered late. An agent-initiated debit that the customer disputes has no existing liability rule, and the reporting says as much. [3] [4] [25]

For policy staff drafting AI agent rules

ASSESSMENT The phrase "standardized, vendor-agnostic naming conventions" creates an obligation with no obligee. Somebody has to operate the namespace, allocate names, resolve collisions and revoke, and no draft in circulation says who. The options are a federal registry, delegation to an existing PKI hierarchy, or a self-asserted scheme that will not meet the independence test. Pick one in the text rather than directing NIST to discover it, because NIST's own concept paper is asking the public that exact question, and a standards body cannot create an allocation authority it has no power to run. The second thing worth fixing is the layer requirement: demanding cryptographic verification at both the network and application layers is satisfiable today at the network layer and, for the agent process itself, is not satisfiable at either. [1] [5] [6]

What this document cannot answer

These are the questions that need primary research rather than more reading, and they are listed because a report that hides them is less useful than one that does not.

What the bill actually says. Every clause-level claim here rests on a trade summary. The enrolled text would settle the deadlines, the covered-entity definition, whether subcontractor agents flow down, and whether "independent" is defined anywhere in the text.

What credential CrowdStrike's Agentic Identity Provider issues. The security claim is the strongest of any vendor in this report and the mechanism is undisclosed. An admin guide, an API reference or a sample token would resolve whether it belongs at the top of the product table or the bottom.

What NPCI is actually building. Whether the Unified Agentic Protocol requires an agent-held key or a scheme-held mandate changes the payments cost model by roughly a factor of two and determines whether scenario C is possible at all.

Revocation latency in practice. Every vendor claims immediate revocation and nobody publishes the number. The interval between a revocation decision and the last moment a cached token still works is the security property that matters, and measuring it requires a lab and access to the products.

How many agents are actually deployed. The product assessments here would change if the typical enterprise has forty agents rather than four hundred, and every published figure traces to a vendor with a reason to make it large.

Half-life of this assessment

Six months. The product rows decay fastest. CrowdStrike's availability and credential format, Okta's Cross App Access status, Solo.io's SPIFFE attestation work and Microsoft's preview-to-general-availability moves will all have changed by March 2027, and the vendor table should be re-read against current documentation rather than trusted.

Twelve months. The mandate reading and the timetable. Either the bill has text and a number by then or it has lapsed, and either the NCCoE has published something or the project has stalled. The NPCI position should be settled in either direction well inside this window, because RBI approval is a dated event.

Twenty-four months. The cost model. Labour rates move, tooling gets cheaper, and the first firms through the work will have published what it actually cost. Treat the totals as a 2026 snapshot and the work breakdown as the durable part.

Architectural, and unlikely to decay. Two things here are properties of the design rather than of the release cycle. A credential issued by a directory can only be verified by asking that directory, which is why tenant-scoped identity cannot satisfy an independence requirement no matter how the product improves. And a process running on a general-purpose operating system cannot attest its own identity without something below it doing the attesting, which is why the agent's name will keep being asserted by the platform rather than by the agent. Any future product claiming otherwise is either using hardware attestation or overstating.

Limitations

This report was produced independently. No vendor named in it commissioned, reviewed or funded it, and the author holds no position in any company discussed. No briefing was taken from any vendor named, which means the product assessments rest entirely on public documentation and may miss capability that exists but is not documented.

The evidence base has four specific weaknesses. The central legal instrument has no retrievable text, so the most consequential section rests on two press releases and one reprinted trade article. The payments half rests substantially on a wire report sourced to three unnamed people, against which one on-the-record statement points in a different direction. The vendor readings are documentation readings, not test results, and a documentation page can be out of date or aspirational. And the two cost models are the author's construction, with no client engagement, benchmark dataset or quoted bid behind them.

Two further cautions. The scoring in the two assessment charts is judgment expressed as a number, which lends it a precision it does not have; a reader who weights the requirements differently will get different bars, and the underlying claims sit in the tables rather than the charts. And the report reads a single week of announcements, a sampling window chosen by the brief rather than by the market, so a vendor that shipped relevant work in August or will ship in October is absent for no reason that reflects on its product.

Sources

[1] Inside AI Policy. "Gottheimer-Lawler bill rejects self-attestation alone in call for agentic AI standards." Reprinted on the website of Rep. Mike Lawler, September 2026. *Trade press, reprinted by sponsor* lawler.house.gov

[2] Office of Rep. Josh Gottheimer. "Gottheimer Introduces Bipartisan Bill to Stop Rogue AI Agents and Keep People in Control." 9 September 2026. *Sponsor office primary* gottheimer.house.gov

[3] Reuters. "India plans AI registry as it looks to roll out agentic payments, sources say." 10 September 2026. *Second-hand (unnamed sources)* wtvbam.com

[4] MediaNama. "NPCI says AI shouldn't approve UPI payments at GFF 2026." September 2026. *Independent press, named official* medianama.com

[5] Booth, H., Fisher, W., Galluzzo, R., Roberts, J. "Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization." NIST NCCoE concept paper, 5 February 2026. *Standards body primary* csrc.nist.gov

[6] SPIFFE. "SPIFFE concepts: SPIFFE ID, SVID and trust domain." SPIFFE documentation. *Standards body primary* spiffe.io

[7] SPIFFE. "Federation." SPIRE architecture documentation. *Standards body primary* spiffe.io

[8] Model Context Protocol. "Authorization." Specification revision 2025-06-18. *Standards body primary* modelcontextprotocol.io

[9] Microsoft. "Overview of agent identities in Microsoft Entra." Microsoft Learn, page dated 30 March 2026, updated 15 June 2026. *Vendor primary* learn.microsoft.com

[10] Microsoft. "Governing agent identities." Microsoft Entra ID Governance documentation, Microsoft Learn, updated 24 June 2026. *Vendor primary* learn.microsoft.com

[11] CrowdStrike. "Introducing the CrowdStrike Agentic Identity Provider." Corporate blog, 2 September 2026. *Vendor primary* crowdstrike.com

[12] CrowdStrike. "CrowdStrike unveils Falcon Guardian to secure AI agents where they execute." Press release, 1 September 2026. *Vendor primary* crowdstrike.com

[13] Levine, I. and Posta, C. "Introducing agentdesktop." Solo.io engineering blog, 3 September 2026. *Vendor primary* solo.io

[14] agentdesktop-dev. "agentdesktop: an open-source visibility and management layer for AI tools across a desktop fleet." GitHub repository, Apache 2.0. *Vendor primary* github.com

[15] Okta. "Okta for AI Agents: govern agentic identity." Product page, accessed 11 September 2026. *Vendor primary* okta.com

[16] Johnson, A. "Securing the agentic enterprise starts with identity." IBM Think, 9 September 2026. *Vendor primary* ibm.com

[17] SecurityBrief Australia. "Microsoft tightens AI governance as agentic risks rise." 8 September 2026. *Independent press* securitybrief.com.au

[18] AIThority. "Solo.io extends agentic governance to the desktop with AgentDesktop." 4 September 2026. *Independent press* aithority.com

[19] BizTech Magazine. "CrowdStrike announces Falcon Guardian and SafeMind AIDR." Fal.Con 2026 coverage, September 2026. *Independent press* biztechmagazine.com

[20] U.S. Department of Defense. Regulatory impact analysis, CMMC program final rule, 32 CFR Part 170. Federal Register, 15 October 2024. *Government primary* federalregister.gov

[21] Hugging Face. "Anatomy of a frontier lab agent intrusion: a technical timeline of the July 2026 incident." Corporate blog, 2026. *Vendor primary* huggingface.co

[22] The Hacker News. "OpenAI agent used exposed credentials across four services during Hugging Face breach." July 2026. *Independent press* thehackernews.com

[23] Business Standard. "UPI sets new record at 24.51 bn transactions in August, value at Rs 29.82 trn." 1 September 2026. *Independent press* business-standard.com

[24] Reserve Bank of India. "Framework for Responsible and Ethical Enablement of Artificial Intelligence: report of the committee." 13 August 2025. *Regulator primary* rbi.org.in

[25] National Payments Corporation of India. "Introduction of UPI Circle: delegated payments for secondary users." Circular, 13 August 2024. *Scheme primary* npci.org.in

[26] Forkast. "Congress is building the scaffolding: the first federal bill mandating agent security standards." September 2026. *Independent press* forkast.news

Rights and permitted use

It is published for public reading. You are free to share the complete, unmodified file, to link to it, and to quote short passages in commentary, reporting, or discussion, provided the author and the site above are credited.

The following require prior written consent: republishing this work in whole or in substantial part under another name, masthead, or brand; excerpting or modifying it in a way that alters the analysis or removes the evidence classifications; incorporating any portion of it into a paid, commissioned, or client-facing deliverable; resale or distribution behind a paywall; and use of this document as training data for machine learning models.

Commissioned Market Assessment reports, commercial licensing, and briefings on this material are available.

This document is analysis, not investment, legal, or procurement advice. It was not commissioned, reviewed, or funded by any company named in it, and the author holds no position in any of them.

Market Assessment reports, commissioned research, and briefings: **davidsoden.com/market-assessment**