

When AI Agents Gain Access to Business Money

Agent Identity, Permission, Financial Authority and Liability in the Synthetic Economy

Permission to reach money is not a mandate to spend it

Conceptual framework developed by Steven Alber

Date of first formulation of the programme: 30 August 2026

This version: **v0.1 — 27 October 2026**

Status: public research note, open for criticism

Primary sources: Stripe, *Introducing Stripe Treasury for agents*, 21 October 2026; Stripe, *Treasury for agents*, product documentation, read 27 October 2026; J. Torrey, *Introducing 1Password for Agents*, 22 October 2026

Correspondence and criticism: research@thesyntheticeconomy.org

This note has not been peer-reviewed. No external scholarly reviewer has read it. The critique in Section 11 is internal critique produced within the programme and must not be described as independent external review.

The central proposition is the programme's own and is **not** a claim made by either vendor. The note rests on vendor documents, states no legal conclusion, did not use or test any product, and did not read the products' terms. It reports no incident or loss.

The Synthetic Economy has not experimentally validated the seven-layer authority audit, which is a proposal for research.

1. Abstract and research question

In the week of 21 October 2026 two companies announced products that let software agents act with a business's or a person's financial and credential infrastructure. On 21 October Stripe announced Treasury for agents, which lets agents connected to a business's Treasury account check balances, pay invoices, move funds between accounts and convert currencies. On 22 October 1Password announced 1Password for Agents, which lets a personal agent request a login from the user's vault, with the user approving each request, so that the agent can sign in without being given the password.

The two products are different in kind. One routes financial actions through a payments platform. The other brokers access to the credentials that sit in front of every other service. Taken together they mark a transition that earlier notes in this programme treated as prospective: agents are being given the means to reach money, not merely to advise about it.

The question this note asks is institutional. **When an agent can reach a business's money, which of the things a business normally relies on to control its money are actually in place?** The note separates seven of them — agent identity, access permission, financial authority, human approval, legal delegation, mandate and liability — and asks, for each, what the two announcements show, what is live, what is in preview, what is merely announced, and what is not addressed in the sources read.

The central proposition is the programme's own and is **not** a claim made by either vendor:

Permission to reach money is not a mandate to spend it. The October 2026 announcements supply the first, in increasingly careful form, and leave the second to be supplied elsewhere.

The note is deliberately modest about what it can show. It rests on three vendor documents. It does not show that either product fails, that agents have misused either product, or that any customer has suffered a loss. It shows which questions the published material answers and which it does not.

2. Sources, method and limits

Claim classes. [FACT] a primary source states it · [REPORTED STATUS] a vendor's own statement of what is available, in preview or forthcoming · [INFERENCE] follows from converging facts · [HYPOTHESIS] requires testing · [NORMATIVE] a proposal about what ought to be done. A vendor's statement that a capability exists is a [REPORTED STATUS], not evidence that it works as described. No product described here was used or tested by the programme. The note has not been peer-reviewed, and the critique in Section 11 is internal critique.

Three documents were read directly on 27 October 2026, each by retrieval of the live page:

1. Stripe, *Introducing Stripe Treasury for agents*, newsroom release datelined San Francisco and Dublin, 21 October 2026.
2. Stripe, *Treasury for agents*, product documentation at docs.stripe.com/treasury/agents, as it stood on 27 October 2026. This page is linked from the release and is the only place in the material read that labels individual capabilities by release stage.
3. John Torrey, *Introducing 1Password for Agents*, 1Password blog, 22 October 2026.

Four limits apply and are stated here because they bound every conclusion below.

The sources are vendor documents. They describe intended behaviour. None is an independent test, an audit or a regulator's statement. Promotional claims on the 1Password page are not used here as evidence.

The pages were read through a retrieval tool, not archived. Product documentation changes. Statements about "what the documentation says" are statements about the page on 27 October

26. The documentation page was returned in full text; the release and the blog post were returned as extracts.

Quotations are short fragments as the retrieval tool reproduced them, and where an extract's exact wording could not be confirmed it is paraphrased and not quoted.

Legal and contractual documents were not read. The Stripe release refers to "applicable product terms" and to a bank partner. Those terms, 1Password's terms, and any account or customer agreement were not examined. Where this note says that a source "does not state" who bears a loss, it means the announcement does not, not that no document anywhere does.

Product names on the 1Password page could not be independently confirmed. Several integration names on the page are unfamiliar. This note relies on their stated status and not on their identity, and names none of them.

3. What Stripe announced

[FACT] On 21 October 2026 Stripe announced Stripe Treasury for agents, described in the release as aimed at startups. The release says business finances "can now be executed by agents from inside the tools they use". It states that agents "can check balances, pay invoices, transfer funds between accounts, and convert currencies."

[FACT] The release describes two things as not yet available. "Businesses will soon be able to use agents for more tasks, including creating cards." And businesses will "be able to set spending limits for their agents in the coming weeks."

[FACT] On controls, the release says that startups using the product choose "which accounts and data agents can access", "whether they can read information or take action, and when money movement requires human approval." It states that "Stripe is not a bank" and that checking accounts are provided exclusively by Fifth Third Bank, National Association, Member FDIC. It also says that the product includes "various distinct Stripe products" and refers the reader to the applicable product terms.

[FACT] The release's illustrative customer is a business whose founder uses his own agent to check the business's balances, subscriptions, failed payments and refunds each morning, among other things, a use described in the present tense. The same release says that he "plans to use" the product to let the agent transfer money between the business's accounts, with "all payments approved" by him, and that he will be able to ask the agent to prepare financial actions and "approve them before any money moves." The money-moving use in the featured example is therefore described as a plan.

[FACT] The product documentation adds detail that the release does not. Stripe's model context protocol server supports two authentication methods: OAuth, "when a client accesses Stripe on behalf of your Stripe user account", and agent API keys, "when an AI application needs independent, persistent access to Stripe." Permissions are set per capability, access can be granted to live-mode accounts or sandboxes separately, authorised sessions can be listed and revoked, and an administrator can enable or disable this access for a whole team. The documentation gives setup instructions for several agent clients, among them Claude, ChatGPT and Codex. It lists five Treasury-specific tools — a balance summary, currency conversion, sending money to a recipient, setting up a recipient, and transferring money between supported Stripe balances — and labels **each of them "Public preview"**. It states: "By default, you need to review and approve any action that your agent takes that results in live money movement."

[INFERENCE] The release and the documentation can be reconciled but they are not the same statement. The release speaks in the present tense about what agents "can" do and what finances "can now" be executed. The documentation, as read on 27 October 2026, labels the individual tools as public preview. A reader who read only the release might reasonably conclude that the capability is generally available. A reader who read the documentation would conclude that it is available as a preview. Neither source, as read, describes it as generally available.

4. What 1Password announced

[FACT] On 22 October 2026 1Password announced 1Password for Agents, a platform which the page opens by saying: "Personal AI agents should work for you without taking over your digital life." It is presented for personal agents. The page states that 1Password for Agents "is currently available at no additional cost for 1Password Individual and Family accounts" on 1Password.com. It also says: "Accounts on 1Password.eu or 1Password.ca aren't currently supported." And: "Availability, platforms, and capabilities vary by partner integration."

[FACT] The credential model, as described, works per request. An agent reaching a login asks to use a credential from the user's vault; the user reviews and approves the request in 1Password; 1Password then fills the credential into the destination site "so the agent doesn't have access to the credentials." The user chooses which logins an agent may use and can disconnect the integration at any time.

[REPORTED STATUS] The page lists integrations with several named agent products. Two are described as "rolling out over the coming days across all platforms"; others are "coming soon", one of them as "web and iOS integrations coming soon"; a further group of names is listed as coming soon without detail. The page also says: "The safeguards and controls vary by integration." Several of the page's descriptions of how integrations behave are in the present tense; its status lines are not. The platform is available. Which agents can use it on a given day is a question the page answers only as "soon".

[FACT] As read, the page says nothing about audit trails or logging, nothing about agent identity, and does not offer business-account availability. It mentions workforce agents in connection with one integration but does not describe controls for business accounts.

5. Live, in preview, announced: a status ledger

The table classifies each capability by the strongest statement the sources make about it. **Available** means the announcement describes it as available now. **Preview** means the vendor's documentation labels it "Public preview". **Forthcoming** means the source says "soon", "in the coming weeks" or "coming soon". **Plan** means a customer's stated intention. **Not stated** means the sources read do not say.

CAPABILITY	WHAT THE SOURCE SAYS	STATUS
Stripe: agents check balances	Release: agents "can check balances". Documentation: balance summary tool labelled Public preview	Available (release); Preview (documentation)
Stripe: agents convert currencies	Release: "convert currencies". Documentation: tool labelled Public preview	Available (release); Preview (documentation)
Stripe: agents transfer between accounts	Release: "transfer funds between accounts". Documentation: transfer tool labelled Public preview	Available (release); Preview (documentation)
Stripe: agents send money to a recipient	Documentation: send-money and recipient-setup tools, each labelled Public preview. Not itemised in the release	Preview
Stripe: agents pay invoices	Release: "pay invoices". No tool for it is named among the five Treasury-specific tools; the wider tool list was not examined	Available (release); Not stated (documentation read)
Stripe: agents create cards	Release: "will soon be able"	Forthcoming
Stripe: spending limits for agents	Release: "in the coming weeks". Not described in the documentation as read	Forthcoming
Stripe: approval before live money movement	Release: the business sets when approval is required. Documentation: required by default	Available
Stripe: featured customer moves money by agent	Release: "plans to use"	Plan
1Password: platform for Individual and Family accounts on 1Password.com	"Currently available at no additional cost"	Available
1Password: .eu and .ca accounts	"Aren't currently supported"	Not supported
1Password: two named agent integrations	"Rolling out over the coming days"	Forthcoming
1Password: other named integrations	"Coming soon"	Forthcoming
1Password: business or workforce accounts	Not offered on the page	Not stated
1Password: audit, logging, agent identity	Not mentioned	Not stated

[INFERENCE] Read at its strongest, the evidence supports the following. As of 27 October 2026 a business could, according to the vendor's own documentation and subject to availability and eligibility conditions that were not examined, give an agent preview access to a defined set of Treasury operations with human approval by default. Cards and spending limits were not described as available. A person could use a credential broker with a personal 1Password account, but the page does not establish that any particular agent could yet use it. Nothing in the sources shows what happens when the controls fail.

6. Seven things that are not one thing

It is easy to treat "the agent is authorised" as a single fact. It is at least seven. They are supplied by different parties, fail in different ways, and can each be present while the others are absent. The definitions below are the

programme's working usage and are not statutory or standards-body definitions.

1. **Agent identity** — which actor this is, distinguishable from its principal and from other agents, in a form a counterparty or an auditor can use.
2. **Access permission** — what the actor is technically able to reach: which accounts, which operations, which credentials.
3. **Financial authority** — the economic effect the actor is permitted to produce: how much, to whom, for what purpose, within what period.
4. **Human approval** — a person's confirmation of a specific action before it executes.
5. **Legal delegation** — the principal's authorisation of the actor to bind it, in a form a legal order recognises. This note does not decide how any jurisdiction characterises a software agent, and states no legal conclusion.
6. **Mandate** — the bounded purpose the delegation serves, expressed so that it persists across changes to the actor and can be checked against an executed action afterwards. Research Note No. 1 calls the reconstructable link between an executed action and the mandate, scope and approval under which it was taken *mandate provenance*.
7. **Liability** — who bears the loss when an action is wrong or unauthorised. This is allocated by law and contract, not by product features.

LAYER	QUESTION IT ANSWERS	IN THE SOURCES READ
Agent identity	Which actor is this?	Stripe: a client acting for a user (OAuth) or an application with its own key; documentation says agent keys give "independent, persistent access". Neither source says an agent is identified as an agent to counterparties. 1Password: not mentioned
Access permission	What can it reach?	Stripe: per-capability, per-account, live and sandbox separately, revocable. 1Password: per login, per request, revocable
Financial authority	What economic effect may it have?	Stripe: spending limits forthcoming. 1Password: not addressed; a login carries whatever the destination account allows
Human approval	Did a person confirm this action?	Stripe: by default before live money movement. 1Password: approval of a credential request
Legal delegation	Has the principal authorised it to bind them?	Not addressed in either announcement
Mandate	What is it for, and can that be checked afterwards?	Not addressed in either announcement
Liability	Who bears a wrong or unauthorised loss?	Not stated in either source; terms not read

[INFERENCE] The table has a pattern. The two announcements are strongest on layers 2 and 4, partial on layer 3, thin on layer 1 and silent on layers 5, 6 and 7. That is not a criticism of the vendors. A payments platform and a credential vault can reasonably be expected to supply permission and approval. Legal delegation, mandate and liability are not things either firm could supply alone, because they belong to the principal, the law and the contract.

The pattern matters because a business that adopts the products may experience the strong layers as the whole of control.

An **illustration**, which is hypothetical and does not describe a reported event: an agent holds permission to send money to recipients already set up; no spending limit is yet in force; it prepares a payment to an existing recipient on the strength of an invoice it was given; a person approves it after seeing the agent's summary. Every layer the products supply is satisfied. Whether the invoice was genuine, whether the payment fell within any purpose the business intended, and who bears the loss if it did not, are decided by none of them.

7. Two models of delegation

The two products delegate differently, and the difference matters for what can go wrong.

Stripe routes actions through the platform that operates the account. The agent calls defined operations. Each operation has a named permission in the documentation. **[INFERENCE]** The platform can therefore refuse an operation that the granted permission does not cover. The unit of control is the operation.

1Password brokers entry to a service, whether or not that service holds money. The agent asks to use a login; the vault fills it in. The unit of control, as described, is the credential.

[INFERENCE] A credential is a coarse instrument. A password to a bank or a supplier portal normally carries the full authority of the account it opens, not a subset of it. The page describes approval of the request to use a credential. It does not describe a constraint on what the agent may do once the session is open. Approving a login is an approval of access, not of any later action. The page is silent on whether any further control exists, and the destination service may provide one. This is a statement about what the page describes, not about what 1Password or the destination services can do.

[INFERENCE] For the same reason the destination service will usually see an authenticated session belonging to the person, not an agent. Whether services can or will distinguish agents from their principals is the identity question in its simplest form, and the 1Password page does not address it.

[INFERENCE] The Stripe model has an advantage the credential model lacks: the platform that moves the money is also the party that defines and enforces the operations. That concentrates the controls, and also concentrates the evidence. The documentation as read describes sessions that can be revoked but does not describe a record of what an agent did, and the release does not mention one. Whether such a record exists is not established here.

[INFERENCE] A real deployment may combine both products with a model from a third vendor. The business would then hold permission from one firm, credential custody with another and the reasoning agent from a third. No one of the three holds a view of the whole delegation, and none is described as being responsible for the mandate. The programme's recurring concern — that authority in agentic systems is assembled from parts supplied by different parties, and that nobody holds its provenance — is here a description of an ordinary product combination, not a speculation.

8. Human approval: what it can and cannot do

Both products place a human approval at their centre, and for good reason. Approval is the most legible control available, and the Stripe documentation makes it the default for live money movement. This section asks what an approval establishes.

[FACT] In the Stripe release's example the person who approves a payment is the person whose agent prepared it. In the 1Password model the person who approves the credential is the person whose agent requested it.

[INFERENCE] Approval is therefore not a second control in the sense a finance function means. Where one person instructs the agent and also approves its prepared actions, there is no separation between preparing and authorising. The control works if the person examines the underlying facts. It works less well if the person approves what the agent has chosen to show. This is a point about a structure, not about the featured customer's conduct.

[HYPOTHESIS] Approval prompts lose protective value as their number rises. A business that adopts agents to save time on routine payments will tend to produce many routine approvals, and routine approvals may be given without examination. Whether this occurs, and at what rate, is an empirical question that neither source addresses and the programme has not tested.

[HYPOTHESIS] What an agent shows an approver is itself a product of the agent. If the agent is manipulated by the documents it reads, the summary it shows may be shaped by the manipulation. No incident is reported in the sources; the hypothesis is that approval of an agent-prepared action is only as reliable as the agent's account of the action.

[NORMATIVE] An approval should be recorded as an event with its own evidence: who approved, what was shown to them, and what was executed. An approval whose content is not retained establishes that a person pressed a button, not that a person authorised a payment.

9. Economic and institutional consequences

The consequences below follow from the structure of the products, not from observed outcomes. Each is labelled accordingly.

[INFERENCE] Delegated spending becomes a distinct class of transaction. A payment initiated by an agent under a human's standing approval differs from one a person initiates, in who decided, on what information and at what speed. The sources describe no mechanism by which a counterparty or a bank could tell the two apart.

[INFERENCE] The small business is the primary user and has the fewest internal controls. The release frames the product as for startups. A start-up may have one person who instructs, approves and reconciles. The control the products offer, human approval, is the control such a business is least able to separate from the agent's own instructions.

[HYPOTHESIS] Attacks will target what the agent reads, not what it holds. An agent that pays invoices reads invoices. If a fraudulent or altered invoice can be made to look ordinary to the agent, the permission and the approval may both be exercised correctly against a bad input. The programme has no evidence of this occurring with either product. It is stated because the controls described address who may act and not whether the input was genuine.

[INFERENCE] Evidence becomes the contested resource. When a payment is disputed, the question is who decided and under what authority. The sources describe controls that prevent actions. They do not describe records that reconstruct decisions. Research Note No. 1 identified this asymmetry in one enterprise agent-governance platform: controls that confer authority were specified in more detail than records that allow its exercise to be contested. The two October announcements, as far as they were read, fit that pattern and do not establish it.

[INFERENCE] Authority is assembled across firms. See Section 7. The institutional significance is that the arrangements which allocate loss when authority fails are, as far as the sources show, bilateral terms between each vendor and the business. The sources describe no arrangement that covers the combination.

10. What remains institutionally unresolved

Each item below is a question to which the three sources do not supply an answer, stated with the evidence that would settle it.

Whether an agent is identifiable as an agent. Stripe's documentation distinguishes acting on behalf of a user from independent access by an application, which is an operational distinction. Whether counterparties, banks and other services can see that a payment or login was agent-initiated is not stated. Evidence: documented fields in the payment or authentication record that mark agent-initiated actions.

Whether permission scopes can express a purpose. The permissions described are operation-level. The spending limits Stripe announced would add an amount. Neither expresses what the delegation is for. Evidence: a documented control that binds an agent to a stated purpose, counterparty class or period, and refuses actions outside it.

Who bears loss from an agent-initiated error. Neither source says. The terms were not read. Evidence: the product terms and account agreements, read in full and compared across the vendors a business combines.

What the law makes of the delegation. The programme has stated no legal conclusion here and the vendors make none. Evidence: legal analysis by qualified practitioners in each relevant jurisdiction.

Whether records reconstruct decisions. The sources describe session revocation and approval prompts. They do not describe a retained record of what the agent was shown, what it did and under what approval. Evidence: documentation or testing of the logs a business could produce after the event.

How controls compose across vendors. Evidence: a deployment, tested, combining a Treasury integration, a credential broker and a third-party model, in which the effective authority at each step is traced.

Whether the preview becomes the product. The Treasury tools are in preview and cards and spending limits are forthcoming. The state of the controls on the day a business adopts them may differ from the state described here.

11. A proposed instrument, and internal critique

[NORMATIVE] The programme proposes, as a research instrument and not as a standard, a **seven-layer authority audit**: for any deployment in which an agent can reach business money, record for each layer in Section 6 who supplies it, what evidence shows it is in place, and its status in the vocabulary of Section 5. A deployment that cannot answer for layers 5, 6 and 7 has not thereby failed; it has located the question that remains. The instrument has not been applied to any real deployment, has not been tested and has not been validated.

Internal critique

C1 — The note rests on vendor statements. Conceded. Three documents were read and none is independent. The note's claims about status are claims about what the vendors say.

C2 — The terms were not read. Conceded, and it is the largest limitation. The liability and delegation questions may be answered in documents this note did not examine. The note does not claim they are not.

C3 — Existing law may already cover agents. Possible and unaddressed. Established principles of agency and authorised signatories might resolve much of what the note calls unresolved. The note therefore states open questions as questions and not as gaps in the law. This is the objection the programme could least answer.

C4 — The seven layers may be over-divided. Partly conceded. Layers 3 and 6 are close. They are separated because a spending limit can exist without a purpose, and the products show exactly that.

C5 — Approval by default may be adequate. Not refuted. The Stripe default is a serious control. The note's point is narrower: approval does not supply identity, mandate or liability.

C6 — The note did not use the products. Conceded. Behaviour in practice may differ from the documents.

Falsification conditions

The note's central proposition would be weakened by any of the following.

- Product terms, read in full, that allocate liability for agent-initiated actions, and that are consistent across the vendors a business would combine.
- A documented control that binds an agent to a stated purpose and refuses actions outside it, in production and not in preview.
- Documented, retained records of what an approver was shown and what was executed.
- Evidence that counterparties can distinguish agent-initiated payments from others.
- Independent testing showing that approval by default is effective in practice at scale.

12. Conclusion

The announcements of 21 and 22 October 2026 are careful and they are real steps. They place a human approval at the centre of agent access to money and credentials, they scope permissions to operations and to logins, and they allow revocation. The Stripe documentation labels its tools as preview and the 1Password page labels most of its integrations as soon. The detailed pages are more precise about status than the headlines are.

What they supply is access, and a layer of confirmation on top of it. What a business needs in order to say that an agent was authorised to spend its money is more: an identity for the agent, a limit on its authority, a legal basis for the delegation, a mandate against which what it did can be checked, and an agreed allocation of the loss. The sources read do not address those, and no single vendor in the examples could supply them.

The programme's proposition is that this is a structural gap and not a temporary one. It lies between firms and between institutions: payments, credentials, models, contracts and law. The first generation of products supplies the layers that sit inside one company's control. Whether the remaining layers are supplied, by whom and in what form, is the research question this note leaves open.

References

Primary sources. All read on 27 October 2026.

1. Stripe, "Introducing Stripe Treasury for agents", newsroom release, 21 October 2026. <https://stripe.com/en-ee/newsroom/news/treasury-for-agents>
2. Stripe, "Treasury for agents", product documentation, as it stood on 27 October 2026. <https://docs.stripe.com/treasury/agents>
3. John Torrey, "Introducing 1Password for Agents", 1Password blog, 22 October 2026. <https://1password.com/blog/1password-for-ai-agents>

Programme documents. Research Note No. 1 v0.1, *When 10,000 Agents Act* (12 October 2026); Research Note No. 2 v0.1, *When the System Becomes Its Own Judge* (13 October 2026); Revision & Evidence Ledger v0.3 with Supplement S; Working Paper No. 1 v0.3-RC1 (private review candidate).

No other literature is cited. The programme does not cite works it has not read.

Research assistance and attribution

Conceptual framework developed by **Steven Alber**. Date of first formulation of the programme: 30 August 2026.

Research assistance was AI-assisted: retrieval of the sources, extraction of stated facts, drafting, and the internal critique in Section 11 were performed with AI assistance under the author's direction. Selection of propositions, the decision as to what may be claimed, and final responsibility are the author's.

This note has not been peer-reviewed. It states no legal conclusion and does not establish any legal principle of agent authority. The Synthetic Economy has not experimentally validated the seven-layer authority audit, which is a proposal for research. The programme is an independent research programme and claims no sponsorship, accreditation or institutional affiliation in connection with this note.

Correspondence, and in particular criticism: **research@thesyntheticeconomy.org**

Appendix A — Source and Citation Register

Level 1 — read directly

L1-RN3-01 · Stripe, "Introducing Stripe Treasury for agents", newsroom release, 21 October 2026.

<https://stripe.com/en-ee/newsroom/news/treasury-for-agents>

Read 27 October 2026, by page retrieval, as extracts. **Verified as stated by the page:** date and dateline (San Francisco and Dublin); the capabilities agents "can check balances, pay invoices, transfer funds between accounts, and convert currencies"; creating cards described as something businesses "will soon be able" to do; spending limits "in the coming weeks"; business choice of accessible accounts and data, of read versus act, and of when money movement requires human approval; "Stripe is not a bank"; checking accounts provided exclusively by Fifth Third Bank, National Association, Member FDIC; a reference to applicable product terms; a featured customer whose agent checks balances each morning, who "plans to use" the product for transfers, with payments approved by the founder.

Not verified. Exact wording of passages beyond the fragments quoted in the note. Eligibility and country availability (the release refers to an availability page that was not read). The release's own use of the words "generally available", "beta" or "preview": the retrieval did not find them, and absence in an extract is not proof of absence in the page.

L1-RN3-02 · Stripe, "Treasury for agents", product documentation, as it stood on 27 October 2026.

<https://docs.stripe.com/treasury/agents>

Read 27 October 2026, in full text. **Verified as stated by the page:** two authentication methods, OAuth ("on behalf of your Stripe user account") and agent API keys ("independent, persistent access"); per-capability permissions; separate grants for live mode and sandboxes; session listing and revocation by the user and by an administrator; team-wide enable and disable of access; five Treasury-specific tools each labelled "Public preview" (balance summary, currency conversion, sending money, recipient setup, transfer between supported Stripe balances) with the permissions each requires; the default that live money movement actions require review and approval; setup instructions for several agent clients.

Not verified. The "other MCP tools" page, which may list tools (including for invoices) not named on this page. Any audit or logging feature. Whether approval can be configured away for a given action. Documentation changes without notice; the note describes this page on the day it was read.

L1-RN3-03 · John Torrey, "Introducing 1Password for Agents", 1Password blog, 22 October 2026.

<https://1password.com/blog/1password-for-ai-agents>

Read 27 October 2026, by page retrieval, as extracts, in three passes. **Verified as stated by the page:** author and date; the opening sentence; availability "at no additional cost for 1Password Individual and Family accounts"; 1Password.eu and 1Password.ca accounts "aren't currently supported"; "availability, platforms, and capabilities vary by partner integration"; "the safeguards and controls vary by integration"; per-request credential approval in 1Password; auto-fill by 1Password so the agent does not have access to the credential; disconnection at any time; two integrations "rolling out over the coming days", others "coming soon"; no statement on audit, logging or agent identity; no business-account offering.

Not verified. The identity and accuracy of the integration product names on the page, which include unfamiliar names. The note relies on their stated status only and names none. The linked security documentation, which was not read.

Level 2 — incorporated by reference

The Level 1 and Level 2 registers of the Revision & Evidence Ledger v0.3 and of the Source and Citation Registers of Research Notes No. 1 and No. 2 are incorporated by reference and were not re-verified for this note. RN-003 does not restate any of their findings in stronger terms than they record. The note's one use of a prior finding is the characterisation, in Research Note No. 1, of mandate provenance and of the asymmetry between controls that confer authority and records that allow it to be contested, which concerned one enterprise agent-governance platform and is not extended here beyond that.

Level 3 — verification pending; non-load-bearing by design

L3-RN3-01 · Secondary coverage of the two announcements. Not consulted, not relied on.

L3-RN3-02 · Stripe and 1Password terms of service, account agreements, bank-partner agreements. Not read. Every statement in the note about who bears loss, or about legal delegation, is a statement about what the three announcements do and do not say.

L3-RN3-03 · Incidents, misuse or losses involving either product. None is reported in the sources read and none is asserted in the note. The hypotheses in Sections 8 and 9 are hypotheses.

All Level 1 content was read on 27 October 2026 and extracted before being used in argument.