The Agent-Ready Platform Specification

Version 1.0 | December 2026

Published by VetSoftwareHub.com, in The Veterinary Software Insider, Issue 37. Free to use, adapt, and fork under the same open license as the Veterinary Software Buyer's Addendum v1.0 and the Practice Software Portability Record v1.0 (Creative Commons Attribution 4.0). No email wall.

What this is

A practice manager can hand this page to a software vendor and say: when you tell us your platform is agent ready, this is what we mean, and this is how we'll check.

It assembles the work of Issues 16 and 33 through 36 of The Veterinary Software Insider into five parts and a test. The parts are the door an outside agent comes through, the account it logs in as, the permission it acts under, the record it leaves, and the person it hands work back to. The test is a set of checks a buyer can run in a sandbox in about ninety minutes, with a laptop and a willing technical contact.

Every requirement carries a line that begins "Without it," naming what goes wrong in a hospital when the requirement is missing. That line is the argument. If you disagree with a requirement, disagree with its failure mode, and send the better one.

It isn't legal advice, a security audit, or a certification. Nothing in it binds a vendor who doesn't sign it. Its value comes from being common: a vendor that sees the same page from forty practices is looking at a market.

How to use it

Send it with your request for a demo, before the demo, with one sentence: "Please tell us which items in version 1.0 of the published Agent-Ready Platform Specification

your platform meets, with evidence for each." Then run the test in Part 6 yourself, or have someone run it for you, and keep the sheet.

Evidence carries one of four labels, carried over from Rubric v1.1 in Issue 33. **P** is public: on a page anyone can read. **D** is demonstrated: someone watched it work and wrote down what happened. **A** is attested: the vendor stated it in writing, with nothing public or demonstrated behind it. **U** is unverified: said on a call, or on a roadmap. A claim of conformance to this specification rests on P and D evidence only.

If you can only get three items, get **S2.1** (the agent has its own credential), **S2.5** (a switch that suspends one agent without touching anyone else), and **S4.4** (the record shows what a person reviewed, and says "none" when nobody did). Everything else improves a platform. These three are what let a hospital stop a tool, and then explain what it did.

Part 1. The door

From Issues 13, 16, and 33.

S1.1 A self service write path.

Core records (appointments, clients, patients, clinical notes, reminders, prescription refills) can be created and updated through a documented interface. A hospital or its chosen tool can get credentials without a partner program, an approval queue, or a business development call.

> **Without it:** the agent can only watch, and every useful task waits on a vendor's permission to exist.

S1.2 Writes that survive a retry.

Every create and update accepts an idempotency key. Sending the same request twice with the same key produces one result.

> **Without it:** a network timeout and a well behaved retry double book a client or charge a card twice, and the record makes it look like a person changed her mind.

S1.3 High level actions.

At minimum these five exist as single named operations: book an appointment, retrieve a patient summary with outstanding items, create or update a clinical note, create or update a reminder, and record a prescription refill.

> **Without it:** every task is assembled from a dozen low level reads and writes, each one a place for the agent to get it wrong.

S1.4 Documentation a stranger can use.

The interface documentation is public, versioned, and states for every action what it expects, what it returns, and what each error means.

> **Without it:** the platform is usable only by developers who studied it for weeks, which is a different and much smaller market than the tools a hospital will want to connect.

S1.5 Events that arrive and get retried.

The platform pushes notice of clinically and operationally significant changes, at minimum appointment changes and new results. It signs its messages, retries when the receiver is down, and reports what it couldn't deliver.

> **Without it:** the agent polls the schedule every few minutes, misses what changed between polls, and finds out late about the thing that mattered.

S1.6 Terms that permit real use.

The terms allow continuous automated use, allow third party tools including ones that compete with the vendor's own modules, carry no per connection fee, and offer testing against the practice's own data before go live.

> **Without it:** the engineering opens the door and the contract closes it again.

Part 2. The account

From Issue 34.

S2.1 Its own credential.

A hospital administrator can create a credential for a non human actor from the settings screen, without a support ticket. It can be rotated or revoked without touching any person's access.

> **Without it:** the tool runs on a staff member's login, and the day she leaves, the practice chooses between breaking the tool and keeping a departed employee's account alive.

S2.2 Permissions narrower than a job title.

Each agent credential is limited to named actions, such as "create an appointment, not cancel one." The platform supports "propose only," where a change takes effect when a person accepts it.

> **Without it:** the agent inherits a whole role, and on the bad day the answer to "what was it allowed to do" is "everything she could do."

S2.3 Ceilings.

Each agent credential can carry a limit on volume (actions per hour) and on value (the largest refund or charge it can make before a person approves).

> **Without it:** a loop that goes wrong at 3 PM keeps going until someone notices, and the size of the mistake is set by the clock.

S2.4 Visibly not a person.

Every action carries the identity that took it and its type, person or agent, in an audit view the hospital opens itself.

> **Without it:** the record names the person whose login was used, and she gets asked about something she didn't do.

S2.5 A switch.

A hospital can suspend one agent in minutes, see everything it did in the last day, and turn it back on, without changing anyone else's access.

> **Without it:** the only way to stop a program is to lock a colleague out of her own computer.

S2.6 A stated price.

The vendor states in writing whether an agent credential takes a seat, carries a flat fee, or is included.

> **Without it:** the practice finds out at renewal, and a postponed pricing decision tends to show up as a borrowed login.

S2.7 Retries scoped to the sender.

An idempotency key belongs to the identity that sent it. The same key from a different identity is a different request.

> **Without it:** two tools that happen to generate the same key collide, and one of them silently does nothing.

Part 3. The grant

From Issue 35.

S3.1 Every grant has four things.

A named owner, a date, a scope written as specific actions rather than categories, and a ceiling.

> **Without it:** a switch labeled "payments: write" gets switched on in March, and in November nobody can say who decided it or how far it reaches.

S3.2 A review date.

Every grant carries a date on which its owner has to look at it again, and the platform reminds the owner. A grant with no end date is shown as such.

> **Without it:** permissions granted for reasons that made sense in one season keep running in a hospital that has changed.

S3.3 A grant covers what it covered.

The platform records what a grant included on the day it was made. When a release adds a capability under an existing category, the grant does not widen on its own, and the grant's owner is told.

> **Without it:** a hospital that said yes to refunds in March has said yes to store credit in October without knowing it.

S3.4 History that doesn't rewrite itself.

When a grant changes, earlier actions still show the terms they were taken under.

> **Without it:** every change to the settings edits the past, and a record that edits itself isn't a record.

S3.5 No orphans.

When the owner of a grant is deactivated, the platform flags every grant she owned for reassignment.

> **Without it:** the permission outlives the person who could have withdrawn it, which is the departed employee's login problem one level up.

Part 4. The action record

From Issue 35.

For every action taken by an agent, one screen the hospital can open shows six lines.

S4.1 What was done, in plain words.

The action as the hospital would describe it, with values before and after. "Refund of $212.00 on this invoice, to this card," not "payment object updated."

> **Without it:** the manager needs an engineer to read her own record.

S4.2 Which actor.

The agent's own identity, labeled as an agent (S2.4).

> **Without it:** attribution falls back to whoever's name is nearest.

S4.3 Under what authority.

The grant the action was taken under, with its owner, date, scope, ceiling, and expiry as they stood at that moment (S3.4).

> **Without it:** the record says who acted and not whether it was supposed to, and the second question is the one a client asks.

S4.4 What a person reviewed.

One of four values: approved before, reviewed after, sampled, or none. "None" is a value shown on the screen in plain words, never a blank.

> **Without it:** an empty field reads as a review that didn't happen, and nobody can count how many actions had no person in the loop.

S4.5 What triggered it.

A link to the input that prompted the action: the client's email, the voicemail, the rule that fired.

> **Without it:** the reason is reconstructed afterward from the surrounding timeline, by people under pressure.

S4.6 How to undo it.

Whether the action can be reversed, by whom, and whether it has been.

> **Without it:** a manager grants the next permission without knowing that a message sent to a client can't be unsent.

S4.7 The hospital holds it.

The record is in a view the hospital opens without asking the vendor, and it is included in the practice's data export.

> **Without it:** the accountability record is a report the practice has to request, and the vendor decides how long it lasts.

Part 5. The handoff

From Issues 15, 35, and 36.

S5.1 A named receiver.

When an agent hands work to a person, the hospital names who receives it and who backs her up, and sets what happens if neither responds within a stated time.

> **Without it:** the handoff rings on a screen and waits, and it lands on whoever is already busiest.

S5.2 Uncertainty goes to a person.

When the agent isn't sure, the item goes to the named receiver, and the record says the system wasn't certain.

> **Without it:** the urgent message that sounds like a routine refill gets handled as a routine refill.

S5.3 A pause list the practice sets.

The hospital chooses which actions require approval before they happen. The default list, until the practice changes it, is money leaving the hospital, anything that deletes something, and anything sent to a client in the hospital's name.

> **Without it:** either everything needs a click, so the clicks stop meaning anything, or nothing does, by default and not by decision.

S5.4 A baseline view.

The platform reports what the practice needs to see whether handing work over helped: wait times, outstanding callbacks, and rebooked slots, from before the agent started.

> **Without it:** the practice has a feeling that things got calmer, and a vendor's dashboard of calls handled.

Part 6. The test

From Issues 33, 34, and 35. About ninety minutes in a sandbox, with no help from the vendor's engineering team beyond credentials.

The door (Rubric v1.1, scored 0 to 14 with evidence labels).

1. Get self service credentials and create an appointment. Send the identical request again with the same idempotency key. Count the appointments (S1.1, S1.2).
2. List the five front desk actions and count how many exist as single named operations (S1.3).
3. The cold read. Give the public documentation to an AI model with no other context and ask it to perform three of the five actions on the test system. Don't help it (S1.4).
4. Subscribe to appointment changes. Change one and time the notice. Turn your receiver off, change another, turn it back on, and see whether the platform retries (S1.5).
5. Read the terms. Four yes or no questions (S1.6).

The account (the identity annex, pass or fail).

6. Create an agent credential. Time it, and note who had to approve it (S2.1).
7. Open its permissions. Can it create without cancelling? Can it be capped by volume and by value (S2.2, S2.3)?
8. Have it create an appointment and open the audit view. Is it labeled as an agent, visible without asking the vendor (S2.4)?
9. Send the same request twice under this identity, then once under another. Expect one appointment, then a second (S2.7).
10. Suspend it. Confirm it stops within minutes, nobody else's access changed, and its history is still there (S2.5).

The record (the record annex, pass or fail).

11. Pick one automated action. From the screen alone, answer the six lines: what was done, which actor, which grant, what review, what triggered it, how to undo it. Time yourself (S4.1 to S4.6).
12. Lower the grant's ceiling and open the earlier action again. It should show the ceiling that applied at the time (S3.4).
13. Find a grant with no review state set on its actions. The screen should say "none" (S4.4).
14. Check the data export for the action record (S4.7).

What a result means. The rubric score stays a gradient, out of fourteen, with its four bands. Conformance to version 1.0 is a separate statement: a platform conforms when every item in Parts 2 through 4 passes on P or D evidence and its rubric score is in the agent capable band or higher. Part 5 is reported but not yet required, because the handoff depends as much on a hospital's setup as on the software. A vendor may state "conforms to the Agent-Ready Platform Specification v1.0" only with the completed sheet published alongside the claim.

The sheet. Platform. Date. Tested by. For each numbered check: result, evidence label, one line of what was seen. Rubric total and band. Annex results as yes or no lines. Then one line: which single item, if it changed, would change the result.

Proposed additions to the Buyer's Addendum (v1.1 draft)

Drafts for counsel to read, not final text. They join the three proposed in the Portability Record.

Clause 7.1, non human actors. Vendor will provide, at no additional charge or at the fee stated in Exhibit A, the ability for Practice to create credentials for non human actors that meet items S2.1 through S2.5 and S2.7 of the Agent-Ready Platform Specification v1.0, and will label every action taken under such a credential as automated in an audit view available to Practice without request.

Clause 7.2, action records. For every action taken under a non human credential, the System records the six lines described in Part 4 of the Specification, retains them for the term of the Agreement, and includes them in any export of Practice Data under Section 1.

Clause 7.3, grant scope on release. Vendor will not expand the scope of a permission granted to a non human actor by means of a product release. Where a release adds capability under an existing permission category, Vendor will notify Practice in writing at least thirty (30) days before the release, and the existing grant will not include the new capability until Practice extends it.

What is deliberately not in version 1.0

- How good any model is. The specification is about the account a model logs in as, not about how well it listens, drafts, or summarizes.
- The vendor's own AI features. A platform's built in assistant should meet the same requirements as an outside one, and how good it is belongs to a demo.
- Liability. Who bears the cost when an automated action goes wrong is a question for contracts and counsel. The specification is about what the software writes down so they have something to work from.
- Security certification. The test is a screen, not an audit. A hospital connecting an outside tool to patient and financial records should have someone competent review how it's done.
- What an agent credential should cost. S2.6 asks for a stated price, not a fair one.
- Clinical judgment. Nothing here says what an agent should or shouldn't decide about patient care.
- A shared record format across vendors. Version 1.0 asks for the six lines on a screen and in the export, not for an exchange format. That is a candidate for a later version, and only once somebody has published the six lines at all.
- Staffing and employment. What a hospital does with time a tool frees is the hospital's decision.

Governance and versioning

Version 1.0 is maintained by VetSoftwareHub.com. Proposed changes, completed test sheets, and any language a vendor has actually signed can be sent to Adam Wysocki through VetSoftwareHub.com, and every disposition will be recorded publicly. A vendor that wants to be tested and named in print can ask; the procedure is the one in Issue 33, and nothing about it is for sale. Version 2.0 will remove or deprecate at least one item, for the same reason the Record will: a specification that only ever grows becomes unusable by the people it was written for.

Adam Wysocki is the founder of VetSoftwareHub.com, an independent, vendor-neutral veterinary software directory and consulting practice, and principal of the PIMS Selection Navigator engagement service. VetSoftwareHub accepts no referral fees or affiliate compensation from any software vendor.

Related documents

The Practice Software Portability Record

The Veterinary Software Buyer's Addendum