AI Provider Information Framework for Veterinary Practice

Developed by the Veterinary AI Transparency Alliance

Draft working paper

Prepared from stakeholder workshop themes on safe and responsible use of AI in veterinary practice. **Last updated: June 2026.**

Executive summary

This framework sets out the information AI providers should make available to veterinary professionals so that AI tools can be adopted safely, proportionately and with informed judgement. The detailed expectations are set out in Tables 1 and 4, with risk-banding and proportionate disclosure requirements set out in Tables 2 and 3.

Four basic themes

1. Match the depth of scrutiny to the risk. The higher a tool's influence on clinical records, reasoning, triage, diagnosis or treatment, the more disclosure and oversight it warrants. Low-risk administrative tools should not be held to the same documentary burden as high-risk clinical ones.

2. Keep a qualified human in charge. AI can support, but does not replace, professional judgement. The person exercising oversight must be competent to challenge the output, and where a tool's output leads toward a diagnosis, oversight should rest with a veterinary surgeon

3. How a tool can go wrong matters as much as how well it performs at its best. Limitations, contraindications and failure modes should be as visible as intended uses.

4. Be transparent with clients and with data. Practices should be equipped to explain AI use to clients clearly and in a timely way, and to understand what data is processed, where it goes, and whether it is reused.

This version (v3) is intended to be understandable to all parties — developers, purchasers and end-users alike — rather than split into separate documents. A principle running throughout is that one size does not fit all: the framework should be read with sensitivity to the user type, the species concerned (small-animal, farm and equine contexts differ), and the profession of the user (veterinary surgeons and veterinary nurses). Future iterations may tailor the framework further; v3 keeps a single, shared document.

1. Purpose of this document

This document sets out a proposed framework for the information that AI providers should make available to veterinary professionals as a matter of best practice. It is intended to support safe, proportionate, and informed adoption of AI tools in veterinary settings.

The framework has been written in a formal, practical style so that it can be used in several ways: as a discussion draft for stakeholders, as the basis for a provider's disclosure template, or as an annex to future guidance for the profession. It is accompanied by Table 4, which sets out the questions veterinary professionals may wish to ask when assessing provider information.

The emphasis throughout is on transparency, human oversight, data governance, training, and the need for communication that is clear enough to support professional judgement and maintain public trust.

The framework is intended to be applied in a risk-proportionate way. Providers and practices should scale the depth of disclosure and oversight to the level of risk a tool carries: a simple administrative tool and a tool that influences diagnosis should not attract identical documentation. A suggested approach is to band tools by risk — for example low, moderate and high — and to expect the detail and the safeguards to increase with the band.

This framework will itself be kept under review so that it remains current as technology, products and relevant legislation change. Providers should update their disclosures whenever their product changes and should record a “last updated” date so that users can see how current the information is.

2. How to use this framework

Table 1 is designed to describe the categories of information that a provider should disclose. It is structured so that each category includes: the information expected from the provider, the reason that information matters in practice, and a minimum expectation. This is intended to keep the framework useful both for developers preparing documentation and for veterinary teams reviewing it.
Table 4 is aimed at veterinary professionals, practice leaders, and other decision-makers. It does not simply ask whether information has been provided. Instead, it helps the reader assess whether the information is sufficiently clear, credible, and relevant to support safe use in the proposed context.

3. Information AI providers should provide to veterinary professionals

AI providers should provide veterinary professionals with clear, accessible, and sufficiently detailed information to enable safe and proportionate use of AI tools in practice. The information provided should be appropriate to the level of risk associated with the tool, particularly where the tool may influence clinical records, clinical reasoning, triage, diagnosis, treatment planning, or other aspects of patient care. This is not a marketing tool and is not intended to be the place for developers to host information about specific features, benefits or USPs.

Table 1. Information AI providers should provide to veterinary professionals

Section	Information providers should disclose	Why this matters for veterinary professionals	Minimum expectation
1. Tool purpose and intended use	A clear description of what the tool is designed to do, the settings in which it is intended to be used, and the user groups for whom it is designed.	Veterinary professionals need to understand whether the tool is appropriate for the proposed use case and whether it fits within their role, competence and workflow.	Plain-language statement of intended purpose, intended users and intended setting.
2. Scope of use, limitations and contraindications	A clear description of what the tool can and cannot do, including whether it supports administration, record drafting, decision support, triage, diagnosis, treatment planning, or operational tasks. This should include an explicit caveat on off-label use — a clear statement of what the tool should *not* be used for, with exclusions given equal prominence to intended uses. Known limitations, excluded use cases, contraindications, circumstances where the	The level of professional, ethical and governance risk depends heavily on the function the tool performs. Safe adoption depends not only on understanding capability, but also on recognising when the tool may be unsuitable, unreliable or unsafe.	Clear categorisation of the tool's function and boundaries of use, including an explicit statement of off-label and excluded uses. Explicit list of limitations, exclusions and higher-risk scenarios. Contraindications should be presented alongside indications, with the same prominence and weight.

	tool should not be used, and situations requiring caution or escalation should also be described.		
3. Degree of autonomy	Whether the tool only generates suggestions, whether it influences decisions, and whether it can trigger actions or outputs automatically without the user's knowledge.	Greater autonomy increases the need for safeguards, review and governance.	Clear statement of whether the tool is assistive, advisory, semi-automated or capable of autonomous action.
4. Integration and interoperability	What systems the tool integrates with, what other data or platforms it can access, and any limitations or dependencies affecting implementation.	Integrations may improve utility, but also create additional governance, security and workflow risks.	Clear list of integrations, dependencies and access implications.
5. Workflow impact and implementation considerations	How the tool is expected to fit into practice workflows, what changes it may introduce, and where new risks, burdens or safeguards may arise.	Even a technically capable tool may create operational or governance issues if poorly embedded in practice.	Summary of expected workflow impact and key implementation considerations.
6. Version control and updates	The current version, date of release, how updates are communicated, whether outputs may change between versions, and whether revised versions are revalidated. This should cover how and how frequently updates happen, and for diagnostic or predictive tools whether predictive ability has changed after an update, since tools may "look the same but be different underneath." This links to post-deployment monitoring for model drift (see sections 14-15).	Veterinary professionals need confidence that changes are controlled and that safety-critical information is not obscured by silent updates.	Version identifier, date, release notes and summary of the update process, including how users are notified of changes and whether performance has changed.

7. Human oversight requirements	The level of human review expected, which tasks must remain under veterinary control, and what the provider considers appropriate human oversight. The provider should also detail the design features intended to prevent *automation bias,* the tendency to over-trust a confident-looking machine output and under-apply one's own judgement. Relevant design features include how uncertainty is surfaced, whether deliberate review or pause points are built in, whether outputs can be auto-accepted, and how low-confidence or out-of-scope cases are flagged.	Veterinary professionals remain responsible for professional judgement and need to understand how oversight should be exercised in practice.	Clear statement that the tool does not replace professional judgement, together with defined review expectations, and a description of the design features that counter automation bias.
8. Intended users and competency assumptions	Who should use the tool, who should not use it, and what baseline knowledge, training or competence is assumed. This should specify the minimum competency required of the person providing oversight. Where a tool's output leads toward a diagnosis, oversight should rest with an appropriately qualified clinician rather than a lay person.	Safe use depends on whether the user has the right knowledge, role and authority to interpret and act on outputs.	Description of intended user groups and any required training or competency assumptions, including the minimum competency for the oversight role.
9. User guidance and training	Practical guidance on how to use the tool well, how not to use it, how to input information appropriately, and how to interpret or review outputs.	Safe use depends not only on the tool itself, but on users being trained to use it responsibly.	Accessible user guidance and training materials appropriate to the role and risk level.
10. Support and escalation	How users can obtain support, raise concerns, report	Users need clear reporting pathways if the tool behaves	Clear support contacts and reporting pathway for concerns and incidents.

	incidents, and escalate suspected errors or adverse events.	unexpectedly or gives rise to safety concerns.	
11. Client communication considerations	Information the practice may need in order to explain the use of the tool to clients, including possible wording on transparency, privacy and use of data.	Trust may depend on a practice's ability to explain AI use clearly and honestly.	Practical wording or guidance to support transparent, timely communication with clients.
12. Client consent and client choice considerations	The provider should equip the practice to decide whether client notice or client consent is required. The provider should flag the circumstances in which notice or consent considerations are likely to arise for this tool, set out what client data is processed and where it moves, flag any sensitive data or special category / protected characteristics involved, and supply adaptable client-facing wording the practice can edit and own. The provider should also state whether it is technically possible for a client to opt out of this tool.	Some uses of AI may require more than general transparency, especially as autonomy and clinical impact increase.	Clear indication of when consent or notice considerations arise, what data is processed and where it goes, any sensitive or protected data involved, editable client-facing wording, and whether client opt-out is possible with the decision and accountability resting with the practice.
13. Veterinary involvement in development and current ownership	Whether veterinary surgeons, RVNs or other relevant subject matter experts were involved in design, testing, validation or review. The provider should also make clear who currently has responsibility for clinical oversight of the tool, and providers should disclose who owns or controls the tool, including any parent company or commercial interest	Clinical relevance and safety are strengthened where veterinary expertise has informed development.	Disclosure of whether veterinary professionals contributed and in what capacity, and identification of the party currently responsible for the tool and its clinical oversight.

	whose products or services the tool might recommend, so that users can judge whether outputs may be influenced by commercial incentives. The emphasis is on present accountability rather than a historical list of past contributors.		
14. Validation and testing	How the tool has been tested, what it has been tested against, in what settings, and for which species, breeds, ages, or case types. The provider should specify the nature of the validation: internal versus external validation, retrospective versus prospective use, and whether real or synthetic data was used. The provider should also describe post-deployment monitoring for model drift and ongoing performance over time.	Users need to know whether evidence for the tool reflects real veterinary use and whether it is applicable to their practice context.	Summary of testing approach, validation setting and major applicability boundaries, specifying internal/external, retrospective/prospective, real/synthetic data, and ongoing monitoring.
15. Evidence base	Whether there is peer-reviewed evidence, third-party review, benchmarking, independent assurance, or other external validation.	Independent evidence improves confidence and helps distinguish robust products from unsubstantiated claims. It is recognised that a formal evidence base is harder for small companies, as peer review is slow and expensive; this is a fairness consideration in how evidence is weighed, not a reason to relax the expectation of transparency.	Summary of available evidence and whether any independent review has taken place.
16. Training data and provenance	What data sources were used to train or configure the tool, including provenance, representativeness,	Users need to understand whether outputs may be affected by bias, incompleteness or	High-level description of training or reference data and any major representational limitations, including any

	relevant species coverage, and known data gaps. The provider should explicitly caution against data and monitoring bias, for example a tool trained predominantly on one breed, species or country, and state how representative the data is for the intended population.	lack of relevance to veterinary practice.	breed, species or geographic bias.
17. Use of third-party models	Whether external or foundation models are used, including whether the system is closed, open, hybrid, deterministic, stochastic, or dependent on third-party providers where relevant. Where a large language model or foundation model sits underneath the tool, this should be made transparent, and any third-party tools or services relied on should be disclosed.	Use of external models may affect explainability, control, data flows, versioning and accountability.	Clear statement of reliance on any external models or services, in plain language, including any underlying Large Language Model or foundation model.
18. Accuracy and performance	What is known about performance, accuracy, consistency, error rates, confidence limits or known failure patterns, and how these have been measured. Where the tool is diagnostic, accuracy should be reported in epidemiological terms - sensitivity, specificity, positive and negative predictive value (PPV/NPV) together with the prevalence and population the tool was tested in, and guidance on when it is and is not appropriate to use. This is an open field: the provider describes performance in terms	Veterinary professionals need realistic expectations and must be able to judge whether claimed performance is meaningful for the intended use.	Clear performance summary in understandable terms, including important caveats - with sensitivity, specificity, PPV/NPV and tested prevalence/population where the tool is diagnostic.

	appropriate to the tool type		
19. Known risks and failure modes	The ways in which the tool may go wrong, including misleading outputs, omissions, bias, unsafe suggestions, or over-reliance risks. This should include whether the developer has a transparent error-recording process, actively invites such reporting from users, and has a plan to act on substantive returns.	It is often more important to understand how a tool fails than how it performs at its best.	Summary of known or foreseeable failure modes and their practical implications.
20. Data inputs	What data is entered into the system, what categories of client, clinical or operational data are processed, and whether all such data are necessary for the function provided.	Practices need to judge whether data use is proportionate and appropriate to the task.	Clear description of data inputs and categories of data processed.
21. Data storage, access and retention	Where data is stored and processed, who can access it, how long it is retained, and how deletion is managed.	Data handling is central to trust, confidentiality, legal compliance and client expectations.	Clear statement covering storage location, access controls, retention period and deletion arrangements.
22. Data ownership and secondary use	Who owns the inputted data and outputs, whether customer data is reused for product improvement or model training, and whether client data may be used beyond the immediate service provided. The provider should be specific about what the data is used for for example selling, model training or product improvement and whether it is anonymised before any such use.	Veterinary professionals need to understand whether practice or client data is being reused in ways that create legal, ethical or trust concerns.	Clear statement of ownership rights and any secondary data use, specifying the purpose of any reuse and whether data is anonymised.
23. Data protection and compliance	Relevant data protection, privacy and information	Practices need assurance that the provider has	Clear summary of relevant compliance

	governance standards met by the provider, including whether the provider states compliance with applicable legal requirements.	considered privacy, security and governance properly.	position and governance commitments.

3a. Proposed risk-banding of tools

This section sets out how the framework's risk-proportionate principle could be applied in practice by banding tools into three levels and indicating which disclosures are expected at each level.

Table 2. Proposed risk bands

Risk band	What it covers (proposed)
Low	Administrative or operational tools with no meaningful clinical influence — for example appointment booking, reminders, or drafting of non-clinical text.
Moderate	Tools that influence clinical workflow or the content of clinical records, but where a veterinary professional clearly makes the decision — for example clinical scribes or decision-support tools.
High	Tools that influence or produce triage, diagnosis or treatment, or that can act with a degree of autonomy.

Who assigns the band: the provider should state the band they consider applies to their tool and the reason for it. The practice should verify that the claimed band is appropriate using Table 4, since the band drives how much scrutiny the tool receives.

Table 3. What each band requires (proposed)

Key: ✓ = expected; ~ = expected in proportionate or lighter form; blank = not required at this band. All entries are proposals for the group to adjust.

Framework section	Low	Moderate	High
1. Tool purpose and intended use	✓	✓	✓
2. Scope of use, limitations and contraindications	✓	✓	✓
3. Degree of autonomy	✓	✓	✓
4. Integration and interoperability	✓	✓	✓
5. Workflow impact and implementation	~	✓	✓
6. Version control and updates	✓	✓	✓
7. Human oversight requirements	✓	✓	✓
8. Intended users and competency (incl. oversight competency)	~	✓	✓
9. User guidance and training	✓	✓	✓
10. Support and escalation	✓	✓	✓
11. Client communication considerations	✓	✓	✓
12. Consent and choice considerations	~	✓	✓
13. Veterinary involvement and current ownership	~	✓	✓
14. Validation and testing	~	✓	✓

15. Evidence base		~	✓
16. Training data and provenance	~	✓	✓
17. Use of third-party models	✓	✓	✓
18. Accuracy and performance (epidemiological where diagnostic)	~	✓	✓
19. Known risks and failure modes	✓	✓	✓
20. Data inputs	✓	✓	✓
21. Data storage, access and retention	✓	✓	✓
22. Data ownership and secondary use	✓	✓	✓
23. Data protection and compliance	✓	✓	✓

4. Notes for veterinary professionals on assessing provider information

The framework above is intended to support responsible adoption, not simply disclosure for its own sake. In reviewing provider information, veterinary professionals should consider not only whether information has been supplied, but whether it is sufficiently clear, specific, and credible to support safe use in the proposed context.

The questions below are designed to help practices, clinical leads, educators, and other decision-makers assess provider claims in a practical way. They may also help identify where additional internal safeguards, training, or governance arrangements are required before a tool is adopted.

Table 4. Questions veterinary professionals may wish to ask when reviewing provider information

Question for the practice to ask	Why it matters	Signs of assurance	Potential red flags
Is the tool's purpose clear and appropriately limited?	Tools with vague claims or overly broad descriptions are harder to govern safely.	Clear intended use, clear exclusions, plain language.	Marketing language without practical boundaries.
Is the proposed use low risk, moderate risk or high risk? What does the risk impact, e.g. data protection, or patient wellbeing and outcomes?	The level of scrutiny should increase with the level of clinical influence or autonomy. Different risk types may need different safeguards.	Risk-proportionate documentation and safeguards.	The same level of documentation regardless of risk.
Is there meaningful human oversight?	Human review only reduces risk if it is real, not tokenistic.	Clear review steps, deliberate pause points and explicit accountability.	Outputs likely to be accepted without scrutiny.
Does the tool’s presentation of outputs genuinely support challenge, or quietly invite rubber-stamping?	Automation bias is most dangerous when a confident-looking output is accepted by default. How outputs are presented either supports independent	Uncertainty surfaced, low-confidence cases flagged, deliberate review points, no silent auto-accept.	Confident single answers, no indication of uncertainty, design that nudges passive acceptance.

	judgement or undermines it.		
Has the tool been tested in relevant veterinary contexts?	Evidence from unrelated settings may not transfer safely into practice.	Veterinary-specific testing, species relevance and transparent limitations.	No evidence, unclear validation, or non-veterinary examples only.
Is the performance information meaningful for this type of tool?	A single accuracy figure is rarely meaningful. What “good” looks like depends on the tool type, so the question should be matched to the tool.	For a diagnostic tool: sensitivity, specificity and the prevalence it was tested in. For a scribe: information on omissions, invented or incorrect detail, and how much correction is required. For triage/booking: how it handles the emergency it must never miss.	A single “% accurate” figure with no context, prevalence or population stated.
Are limitations and failure modes described honestly?	The profession needs to understand how a tool can go wrong, not only what it does well.	Specific limitations, practical warnings and examples of edge cases.	No limitations stated or claims that imply near-infallibility.
Is the data governance explanation adequate?	Weak governance may undermine trust even where the tool is otherwise useful.	Clear storage, retention, ownership and access information.	Unclear data location, unclear ownership, or vague privacy wording.
Does this tool process any personal data (client, staff, or third party)? If yes, has the practice confirmed which data protection obligations apply?	This is the gateway question for GDPR and other data protection obligations. If personal data is processed, a chain of legal and governance requirements follows.	Clear confirmation of whether personal data is processed, with categories specified and lawful basis identified.	No mention of personal data, or assumption that data protection does not apply because data is "only" clinical or operational.
Is the use of data proportionate?	Not all tools require extensive client or clinical data to function.	Limited and relevant data input requirements.	Excessive or unnecessary data collection.
Are version control and updates transparent?	Silent or poorly explained changes may alter risk without users realising.	Dated release notes and update summaries.	No version history or unexplained changes.
Is training and support adequate for the intended users?	Tools are only as safe as their deployment and use in practice.	Clear training, support materials and escalation routes.	Assumption that the tool is self-explanatory.
Could the tool encourage over-reliance or reduce critical thinking?	Plausible outputs may be given undue weight in fast-moving clinical settings.	Strong emphasis on review, challenge and professional judgement.	Framing that encourages passive acceptance of outputs.

Does the provider help the practice explain AI use to clients?	Transparency and trust will depend on communication as well as governance.	Clear client-facing guidance or suggested wording, which the practice can edit and own.	No support for client explanation despite material AI involvement.