Vault CRM Agentic Call Report™

Saying Yes to Commercial Evidence™ in CRM: A Guide for Legal and Compliance Stakeholders in Life Sciences

July 2026

Why this guide?

In March 2026, we published a guide for legal and compliance professionals who had been asked by the business: “Can we put free text call notes in Veeva Vault CRM?” The response told us two things. First, the life sciences industry is excited about the possibilities of Commercial Evidence, so free text is top of mind in nearly every organization. Second, stakeholders like you have more questions about AI features in CRM. The business is not just asking to type notes in the system of record — it is asking to dictate call reports by voice, to let AI assemble the call report, and to use AI to oversee what gets written. Veeva calls this combination the Agentic Call Report.

This expanded guide covers Agentic Call Report: Agentic Voice, which converts a user’s dictated notes into a structured call report; Agentic Text Monitoring for End Users, which screens free text for compliance concerns in real time as end users document their work; and Agentic Text Monitoring for Home Office, which gives compliance teams agentic, retrospective monitoring of free text that was submitted. It also addresses the questions we’re hearing the most from legal and compliance stakeholders when evaluating these features: Do current or pending AI laws and regulations affect them? What about laws governing audio recording and voice data? Where is the data stored, and what logs are kept? How does Veeva test the agents?

A note on what this guide is — and is not. Veeva’s product teams do not provide legal advice to our customers, and nothing in this guide is a substitute for advice from your own counsel. But we do not arrive at these questions empty-handed, either. Our point of view and deep industry knowledge — of how field documentation is regulated, litigated, discovered, and monitored — shape the design decisions behind every feature described here. Where this guide expresses an opinion, it is offered in that spirit: an informed perspective to support your own assessment, not a legal conclusion.

We recognize that each organization has unique needs and risk tolerances, and that decisions about CRM configuration require input from legal, compliance, privacy, and IT teams. Before reading this guide, stakeholders may find it helpful to complete the Agentic Call Report Flightpath. Please reach out if you have questions or would like to discuss configuration options.

Thank you,

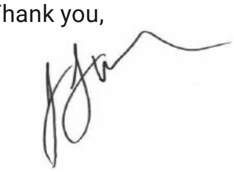

Jenny Stohlmann, RAC
Senior Product Manager, CRM

P.S. The life sciences industry is operating in a time of fast and robust technical change. Regulators are busy trying to keep up. The Vault CRM roadmap has Agentic Call Report features planned in every upcoming release. Conversations about AI will remain lively.

What AI has not and will not change is that Veeva engages with you, our partners in life sciences, to innovate products you can actually use in the regulatory context and provide you with the necessary information to facilitate compliant adoption of those features. We will continue to update this guide as we release enhancements with relevant considerations or we become aware of major changes to laws, regulations, or industry standards.

Contents

What is Agentic Call Report?

Agentic Call Report uses AI to transform natural language into structured call reports, replacing field-by-field data entry with a superior user experience and the ability to capture valuable Commercial Evidence at scale. End users capture their interactions the way people naturally communicate — by speaking, writing, or both — and the agents support the structured work: assembling the call report, mapping details to the right fields and records, and screening free text for compliance concerns.

Three agents power Agentic Call Report:

- **Agentic Voice** converts a user's dictated notes into a populated call report, including information like the account, call channel, products discussed, follow-up activities, and call notes. The user reviews and can edit the call report before saving or submitting it.
- **Agentic Text Monitoring for End Users** screens free text in real time, at the point of entry, and alerts the end user of potential compliance issues before saving the record. Our March 2026 guide covered this feature in depth.
- **Agentic Text Monitoring for Home Office** retrospectively analyzes end users' text inputs for potential compliance concerns, giving monitoring and compliance teams an agentic review capability for documentation that enters the system. This agent became available after our March 2026 guide was published.

The features are modular. You can deploy Agentic Voice alone, with Agentic Text Monitoring for End Users, with Agentic Text Monitoring for Home Office, or with both. You can also enable one or both Agentic Text Monitoring features without Agentic Voice.

The output of Agentic Call Report is twofold: a complete, compliant call report in the system of record and a growing body of what Veeva calls Commercial Evidence — the comprehensive, analyzable record of what is happening in the field.

Why did pharmaceutical companies restrict free text in CRM — and what did it cost?

Following the widespread adoption of electronic SFA and CRM tools in the mid-2000s, pharmaceutical legal and compliance decision makers began limiting the free text fields available to field users in systems of record like Vault CRM. The strategy was motivated by legitimate concerns: user-generated call notes and narrative documentation could be used as evidence against companies in government enforcement actions and private litigation. Picklist-heavy layouts reduced the risk of incriminating or ambiguous written records. For compliance departments with limited monitoring bandwidth, restricting free text also reduced the volume of records requiring review. In the enforcement environment of that era, these were defensible business decisions.

The strategy produced significant unintended consequences for field productivity and for compliance itself. Call page layouts became dominated by structured data — picklists, dropdowns, and checkboxes — with little or no room to document the context, content, or outcomes of engagements. Without adequate space for narrative documentation in CRM, field users increasingly turned to other tools: personal note applications, calendar entries, email drafts, and other unsanctioned systems. Documentation that exists outside of CRM remains discoverable. You should expect it to be subject to oversight by compliance departments. Scattered documentation is harder to monitor and harder to defend.

The practical impact on end users was also substantial. Sales representatives and medical affairs professionals deserve tools that support efficient planning and documentation. Heavily restricted CRM layouts resulted in reluctant user adoption, pushed teams to workarounds, and denied organizations valuable data about field activities.

Is documentation outside of CRM subject to compliance oversight?

Yes. The U.S. Department of Justice (DOJ), HHS Office of Inspector General (OIG), and other enforcement bodies have made clear through corporate integrity agreements, settlement agreements, and compliance program guidance that pharmaceutical companies are expected to exercise oversight over field team activities and documentation, regardless of where that documentation lives.

The DOJ's guidance on effective compliance programs emphasizes systematic monitoring and auditing. If a representative documents a customer interaction in a personal notes application because CRM does not support that use case, that record may still be discoverable in litigation or a government investigation. A company cannot reduce its compliance exposure by making its own system of record less functional. Pushing documentation outside of CRM does not eliminate the documentation — it makes it harder to oversee and creates the kind of audit gaps that regulators and enforcement bodies expect companies to close.

What is Commercial Evidence, and why does it matter now?

Commercial Evidence is the complete, analyzable record of what is happening in field and brand interactions: what healthcare professionals (HCPs) are asking, what concerns them, where patients are encountering barriers to therapy, and how the market is responding to a product. Narrative call notes — captured compliantly in the system of record — are its raw material.

Structured picklist data can tell you that a call occurred and what product was discussed. Free text notes can tell you what the HCP is thinking. Historically, that distinction was meaningless because no organization could read tens of thousands of call notes. AI-driven analytics changed that. Unstructured notes can now be synthesized at scale to surface business-critical information like:

- Treatment barriers and access obstacles as HCPs describe them
- New objection themes and unmet needs
- HCP sentiment and engagement trends that leading indicators like counts of calls cannot capture
- Signals for market access, medical, and brand teams that would otherwise arrive quarters later through syndicated data, if at all

This is the business case in its simplest form: Organizations that allow compliant free text in CRM are building a proprietary data asset from interactions they already pay for. Organizations that do not are leaving that asset unrecorded — or worse, recorded in places where it generates no value and no visibility.

The term 'evidence' is deliberate, and it cuts both ways. Commercial Evidence is evidence for the business: What is working and what is in the way. It is also evidence in the compliance sense: a consolidated, time-stamped, monitored record of field activity — with documented screening at the point of entry — that is far more defensible than fragmented notes scattered across personal tools.

Does restricting free text still reduce risk? A contemporary assessment

The honest comparison is not between 'free text in CRM' and 'no free text anywhere.' Your field teams are already documenting their work in notes apps, calendars, spreadsheets, and email drafts. The comparison a contemporary risk assessment must make is between free text captured in a monitored system of record and the same words captured in systems with no screening, no monitoring, no retention policy, and no audit trail.

Framed in that way, the traditional position rests on a double standard. CRM — the only system with role-based access controls, audit trails, retention management, and now real-time and retrospective AI monitoring — is held to a standard of perfection no other repository is asked to meet. A note that would be flagged, corrected, and logged in Vault CRM is tolerated without any oversight when it lives in a personal app. The risk never disappeared when free text was restricted; it relocated to where compliance cannot see it or respond with Corrective and Preventative Actions (CAPAs).

Allowing free text in CRM with Agentic Call Report is therefore not an exchange of business value for compliance risk. It delivers on both sides at once:

- **Business value:** The organization gains increased analytical value from field documentation — Commercial Evidence that compounds with every interaction.
- **Compliance value:** Monitoring and compliance teams gain visibility into documentation that today is scattered and effectively invisible, plus a structured log of what was flagged and why.

In our view, a risk assessment performed against the current state, rather than an idealized world in which a restricted CRM means no free text documentation exists, is the right thing to do.

What is Agentic Voice and does it record conversations?

Agentic Voice uses AI to convert a user's dictated notes into a call report. The end user describes the interaction in their own words — the account and attendees, the date and channel, the products and key messages discussed, follow-ups, and narrative notes — and the agent transcribes the speech, analyzes the text, and creates a new call report with the information. The user reviews the generated call report and makes any necessary edits before saving or submitting it. Nothing enters the system of record without a human review.

Today, the handling of the audio itself reflects data minimization by design:

- Recording and transcription occur locally on the user's device using native iOS capabilities. A call voice note that is strictly device-bound captures the audio; it does not sync to the Vault CRM database or any other device.
- Only the transcribed text is sent to the cloud-based AI orchestration layer to generate the call report. The audio never leaves the device.
- Call voice note audio is permanently deleted when the user submits the call report, discards the note, or deletes the call report. There is no server-side audio archive to govern, retain, or produce.

Agentic Voice also respects the compliance configuration you already maintain; for example, restricted and not-allowed products are not added to call reports. When Agentic Text Monitoring for End Users is enabled, the dictated call notes are screened just like typed notes. That last point matters: Natural spoken language is more likely than typed notes to include incidental personal details or miss a meaningful 'not' or 'but.' Agentic Text Monitoring for End Users acts as the safety net at the point of entry.

Agentic Voice is not a conversation recorder today. The feature is a dictation workflow. The user speaks their own summary of an interaction after it has concluded. It is not designed to capture HCP conversations. The only voice involved is the voice of the user who initiated the dictation. Outside of Agentic Call Report, Vault CRM does offer the capability to capture consent, then record and transcribe meetings with HCPs with the Medical Interaction feature. We highlight this distinction to avoid confusion for teams that may be reviewing or using both features for different user groups.

What is Agentic Text Monitoring for End Users?

Agentic Text Monitoring for End Users is an AI-powered screening feature that evaluates free text in Vault CRM before it is saved to the system of record. It uses a large language model to analyze free text and identify entries that may present compliance risk or that should be documented in a different system or channel. Out of the box, it screens for Protected Health information (PHI) and Personally Identifiable Information (PII), adverse events and product quality complaints that belong in safety and quality systems, off-label information that belongs in a Medical Inquiry, inappropriate and profane content, and restricted or not-allowed products — plus a configurable Corporate Compliance category grounded in your own plain-language policy documents.

The agent is designed to be supportive, not punitive. When it identifies a concern, it alerts the user by highlighting the text and, where applicable, directs the user to the correct channel — for example, the specified pharmacovigilance system for an adverse event. Nothing is silently suppressed, and the agent can be configured to prevent the record from being saved until the issue is resolved. The goal is good documentation hygiene: notes consolidated in CRM, where compliance teams can monitor them, and other users can leverage them for business analytics.

Our March 2026 guide, Saying Yes to Free Text in CRM, covered the Agentic Text Monitoring for End Users in depth, including how each screening category maps to a specific regulatory obligation. The screening category reference is reproduced in Appendix A.

What is Agentic Text Monitoring for Home Office?

Agentic Text Monitoring for Home Office is the compliance team's side of the Agentic Call Report. Where Agentic Text Monitoring for End Users works with end users in real time at the point of entry, Agentic Text Monitoring for Home Office retrospectively analyzes end users' free text for potential compliance concerns, on behalf of the monitoring, auditing, and compliance functions — the 'home office.'

The feature takes the same fundamental technology used in the end-user-facing agent — a large language model (LLM) that reads text in context — and applies it to the retrospective monitoring workflows that compliance departments already operate. Rather than sampling records manually or relying on keyword searches, compliance teams define their monitoring criteria and let the agent evaluate the free text of the records against them, flagging the appropriate records for human review. Reviewers remain as the decision makers; the agent's role is to focus their attention where it matters.

For organizations subject to corporate integrity agreements or operating risk-based monitoring plans, this is a meaningful capability: it provides documented, systematic, repeatable review coverage across the full population of field documentation — not just the sample of records a human team has bandwidth to read.

How does Agentic Text Monitoring for Home Office improve upon Approved Notes?

Many organizations will recognize the retrospective monitoring concept from Approved Notes, the long-standing Vault CRM feature for monitoring free text fields. Approved Notes works by keyword and phrase matching: Compliance users define monitoring rules containing specific monitored phrases, a scheduled background job evaluates saved text against those phrases, and matches are logged as violations for compliance review.

Approved Notes was an important tool in its era, but phrase matching has a well-known structural weakness: It cannot read context. A rule that flags a drug name or the word 'reaction' will flag it everywhere — including in notes that say a patient did not experience a reaction, in benign administrative uses of a term, and in discussions that are entirely appropriate. The result is a high volume of false positives. In practice, false positives are not merely inefficient. They consume scarce compliance review capacity, train reviewers to expect noise, and increase the odds that a genuine issue is overlooked in a queue full of non-issues. They also make the monitoring program harder to defend, because documented review backlogs are themselves a finding waiting to happen.

Agentic Text Monitoring for Home Office is agentic, not lexical. The LLM evaluates what a note means — handling negation, context, and intent — against monitoring criteria expressed in plain language rather than brittle phrase lists. The practical differences include:

- **Fewer false positives:** The agent understands a note mentioning an adverse event that was already reported in context rather than pattern-matching.
- **Plain-language criteria:** Compliance teams describe what they are looking for the way they would brief a human reviewer instead of attempting to enumerate every phrase variation in advance.
- **Higher-quality queues:** Review queues contain items worthy of a human's judgment, which is a better use of compliance resources and a stronger demonstration of an effective monitoring program.

Appendix B compares the three monitoring approaches side by side.

Where do the rules come from?

A frequent and important question is: Who defines what the agents screen for? The answer differs between the two text monitoring features, and the difference is intentional.

For Agentic Text Monitoring for End Users, the screening rules come from two sources. The out-of-the-box categories are grounded in the regulatory obligations most relevant to each category in the country where the interaction takes place. No global legal research project is required to enable the agent because Veeva has pre-defined them. On top of that, the configurable Corporate Compliance category is grounded in your organization's own plain-language policy documents (e.g., SOPs, work instructions, and handbooks) if you choose to provide them to the agent via retrieval-augmented generation (RAG). User type, country, division, department, and title are considered when evaluating whether text poses a concern. These rules are deliberately field facing: When the agent flags a text string, the user sees it right away to reinforce training in a way silent surveilling never can.

For Agentic Text Monitoring for Home Office, the monitoring criteria are 100% specific to your organization. The compliance and monitoring organization or its designee authors and maintains them, reflecting the monitoring plan, risk assessments, and audit priorities. The agent manages these criteria independently of the end-user-facing configuration and does not display them to field users. That separation preserves the integrity of the monitoring function: The field sees a standard set of rules that are directly tied to top-level risks of documenting text in CRM, while the home office retains the ability to monitor for patterns and risks it does not wish to telegraph.

In short: The end-user agent enforces the rules you publish to the field while the home office agent applies the criteria your monitoring program keeps to itself. Various sources, different audiences, and one consistent system of record.

Should we deploy one text monitoring feature, or both?

Agentic Text Monitoring for End Users and Agentic Text Monitoring for Home Office can be used individually or together. They are complementary by design, because they answer different questions:

- Agentic Text Monitoring for End Users is preventive. It operates at the point of entry, in front of the end user, and keeps non-compliant content from reaching the system of record where it is visible to other users and flows to analytics engines downstream as Commercial Evidence. This supports your field representatives in getting their documentation right.
- Agentic Text Monitoring for Home Office is detective. It operates after the fact, in front of the compliance team, and provides systematic oversight of what field teams documented. This supports your monitoring function to scale as field representatives document more free text Commercial Evidence.

Together, they map directly onto the prevention and detection structure expected of an effective compliance program: a published guardrail that helps users get it right, and an independent monitoring capability that verifies outcomes. Organizations that deploy both get defense in depth and a documented, end-to-end story of a functioning control environment.

It is important to understand that the features stand on their own. Some organizations will start with the end-user agent to improve the field experience and stop documentation errors at the source. This can help build trust with end users as they grow accustomed to adding free text to the Agentic Call Report. Others, particularly those with established monitoring programs built on Approved Notes, may begin with Agentic Text Monitoring for Home Office to modernize retrospective review before changing the end-user experience.

All three options are legitimate paths; the right choice for your organization depends on your monitoring obligations, employee training and communication considerations, and your compliance team's capacity.

Do current or pending AI laws and regulations affect these features?

What follows is Veeva's informed point of view on the landscape as of July 2026. This represents the perspective that shaped how we designed these agents, not legal advice. Developments around AI usage in life sciences are fast-moving, and your counsel should validate how each instrument applies to your organization's specific deployment.

European Union: the AI Act

The EU AI Act entered into force in August 2024, with obligations phasing in over several years. In May 2026, EU negotiators reached provisional agreement on the ‘Digital Omnibus’ package, which postpones the obligations for most high-risk (Annex III) AI systems from August 2026 to December 2027 and introduces targeted simplifications. The European Commission also published draft guidelines on high-risk classification in May 2026.

Our assessment of where Agentic Call Report features sit in that framework:

- No feature comes near the AI Act’s prohibited practices. There is no social scoring, no manipulation, and — importantly for workplace tools — no emotion recognition: The agents evaluate the compliance content of text, not the emotional state of the person who wrote it.
- The high-risk category we are most attentive to is Annex III’s employment category, which covers AI systems used to make or materially influence employment decisions and to monitor or evaluate worker performance and behavior. Whether a text monitoring deployment falls in scope depends not on the product features we have delivered, but rather on how your organization uses the outputs. For example, screening documentation quality is not the same as evaluating employees, but outputs that feed disciplinary processes deserve a closer look. Organizations should make this assessment deliberately, and the December 2027 deferral provides time to do it well.
- The design of the agents anticipates the AI Act’s direction: human review of consequential outputs, transparency to users when text is flagged, logging and observability, and narrow, documented purposes. Obligations attached to the underlying general-purpose model sit upstream with the model provider. We make fit-for-purpose selections of the models we use for out-of-the-box agents, and EU AI Act compliance is a part of those decisions. Today, all standard Vault CRM AI features use Anthropic models, hosted on AWS Bedrock. Agentic Voice additionally uses Apple models for transcription and Siri commands on iOS devices.

European Union: General Data Privacy Regulation (GDPR)

GDPR is not new, but it remains the prevailing data privacy framework for most processing that the Agentic Call Report performs. The contents of Call Voice Notes, transcripts, and free text could be considered personal data (information relating to an identifiable individual) of the user and, potentially, of the HCP. Three points matter most. First, Article 22, the provision giving individuals the right not to be subject to a decision based solely on automated processing that produces legal or similarly significant effects: No feature in the Agentic Call Report makes such a decision — a human reviews every call report before submission, and a human reviews every monitoring flag. Second, transparency: Employees should be informed about how monitoring works, and in jurisdictions with work councils or similar employee representation, consultation may be required before deployment. Third, accountability: A Data Protection Impact Assessment, the structured evaluation GDPR expects before higher-risk processing, is a sensible step that most life sciences companies already take when enabling new features in Vault CRM, and the structured logs the agents generate can make that assessment easier to document.

United States: state AI and privacy laws

The U.S. state landscape remains turbulent, but as of June 2026, the practical picture for Agentic Call Report is manageable. For example:

- **Colorado:** The Colorado AI Act was revised in May 2026 (SB 189): Its effective date moved to January 1, 2027, and its original risk-based duty-of-care framework was scaled back to a narrower disclosure-and-transparency approach for certain automated decision-making technologies. If 'consequential' employment decisions are being made in an automated fashion based on a CRM agent's outputs, Colorado's compliance obligations should be scoped with your counsel.
- **Texas:** The Texas Responsible AI Governance Act (TRAIGA) took effect January 1, 2026. It is intent-based — prohibiting AI developed or deployed with intent to unlawfully discriminate — with enforcement reserved to the Attorney General and a cure period. There is no intent to discriminate in the Agentic Call Report features.
- **Illinois:** HB 3773 amended the Illinois Human Rights Act effective January 1, 2026, to cover AI-mediated discrimination across the employment lifecycle, including discipline, and requires notice to employees when AI is used in such decisions. If Agentic Text Monitoring outputs inform disciplinary processes, employee notice may belong in your deployment plan.

- **California:** The California Consumer Privacy Act (CCPA) regulations for automated decision-making technology require compliance by January 1, 2027, for technologies that replace or substantially replace human decision-making in significant decisions, including employment. The human-in-the-loop design of Agentic Call Report keeps humans deciding — but California's risk assessment requirements for profiling based on work performance should be scoped with your counsel.

Our point of view on the trend: Across every jurisdiction, the common denominators are transparency, human oversight, documented testing, and accountability records. These are design principles the Agentic Call Report already embodies. Where these laws reach discrimination — Texas' intent-based prohibition, Illinois' coverage of AI-mediated employment decisions, and Colorado's algorithmic-discrimination framing — a meaningful share of the responsibility for bias originating in the underlying model sits upstream with the model provider, as it does under the EU AI Act. Veeva's fit-for-purpose model selection accounts for this, and the deployer-facing duties that remain with your organization are familiar ones. For most Vault CRM customers, the compliance effort these laws create is governance already in effect — notices, impact assessments, vendor diligence.

What about laws governing audio recording and voice data?

Stakeholders reasonably ask whether Agentic Voice implicates recording consent laws and biometric privacy statutes. In our conversations, we find that a clear understanding of the agent's design can help to quickly rule out a lot of your concerns.

Recording consent laws

Federal law and most U.S. states require the consent of one party to record a conversation; 11 states require the consent of all parties. These statutes are aimed at recording communication between people. Agentic Voice records a dictation: the user deliberately speaking their own summary of an interaction to their own device. The only speaker is the consenting user. The feature is not designed to record HCP conversations, and the dictation workflow naturally takes place after the interaction has concluded.

Our position is that our customers can address this through policy match design without a mechanism to capture one-party consent in the user interface. Acceptable-use policies and field training should state that Agentic Voice is for dictating the user's own notes and must not be used to record other people. This is the same instruction many organizations already give for any voice memo or dictation tool, and it keeps the all-party consent question from arising in the first place.

We again note that Vault CRM offers consented in-meeting audio recording on Medical Interactions. This is not currently a part of Agentic Call Report and therefore outside the scope of this guide.

Biometric and voice data laws

Biometric privacy statutes — the Illinois Biometric Information Privacy Act (BIPA), the Texas Capture or Use of Biometric Identifier Act (CUBI), and the Washington Biometric Identifier Law and Biometric Privacy Protection Act (WBPA/WBIPA) — regulate biometric identifiers such as voiceprints: representations of a voice used to identify an individual. Today, Agentic Voice performs speech-to-text transcription. It does not create voiceprints, does not enroll or identify speakers by voice, and does not perform speaker diarization to attribute speech to identified individuals. The same distinction runs through GDPR: Voice data is 'biometric data' in the special-category sense only when processed for the purpose of uniquely identifying a person, which transcription is not.

Retention compounds the point. Because audio is processed on the device, never logged or synced to the server, and permanently deleted upon submission or discard, there is no standing audio archive to secure, retain, subject to access requests, or produce in discovery. The audio transcript is purged after 120 days, and the transcript's useful content survives only as user-reviewed records in Vault CRM — where it is governed by the controls you already operate. And if incidental personal details do surface in dictated notes, Agentic Text Monitoring for End Users screens for PHI and PII at the point of entry, notifying users to prevent the subject data from leaking into the Vault CRM database.

Where is data stored, and what logs are kept?

Agentic Call Report features are built into Vault CRM and follow the Vault Platform's security and data governance model. The agents use Anthropic's Claude as the LLM, hosted on AWS Bedrock, connected to a specialized Vault CRM orchestration and data layer that grounds the agents in your organization's own Vault data. End-user-facing agents inherit the initiating user's data access permissions and sharing rules, so AI-generated outputs stay within your established governance boundaries.

Where each category of data lives:

- **Voice Note audio:** captured and transcribed locally on the user's device; never synced to the Vault database or other devices; permanently deleted when the user submits the call report, discards the voice note, or deletes the call report.

- **Voice Note transcripts:** The transcript produced from the user's audio input is sent to the cloud-based AI orchestration layer to generate the call report. Like other ordinary logs that Veeva maintains for the purposes of providing product support, monitoring the health and performance of the product, billing, etc., any logs that contain the text of the transcription are stored in accordance with your Master Services Agreement and purged after 120 days based on Veeva's current retention schedule.
- **Call report records:** The user-reviewed, saved call report — including free text call notes — resides in your Vault CRM instance as the system of record, subject to Vault's data residency, security, audit trail, and retention controls. The Vault platform's Litigation Hold feature can be used on these records.
- **Monitoring results:** Agentic Text Monitoring for End Users generates a structured log of what was flagged and why. This can be compared to the value of the call notes field on the call report records to understand how the user responded. Today, these logs are provided in an ad hoc report; however, they can soon be reviewed as Monitored Text records in Vault (roadmap estimate: August 2026). Agentic Text Monitoring for Home Office results are likewise available to compliance reviewers as Monitored Text records in Vault today. This monitoring data supports trend analysis, risk-based monitoring decisions, and training priorities. It also serves to document that a real-time detection capability is in place and functioning.
- **Agent activity logs:** Administrators can monitor agent activity through agent activity logs and an observability portal covering usage patterns, adoption, success rates, and LLM token usage. Agent instance and agent action execution data for standard agents is refreshed on a regular cycle throughout the day in Vault.

The result is an unusual and useful inversion: The AI features generate more compliance documentation than the manual processes they replace, not less. You can trace the story of every screening decision, user response, and monitoring outcome in a system designed for inspection.

How does Veeva test the agents and keep them fit for purpose?

Veeva employs a multi-layered evaluation framework that combines traditional software testing with specialized AI validation to ensure reliability:

- **Deterministic qualification:** Standard QA testing validates the underlying software foundation, ensuring that prompt handling, agent orchestration, and integrations function correctly.
- **Probabilistic evaluation:** Because LLM outputs are non-deterministic, Veeva uses an automated methodology to measure performance against curated, expert-verified 'Ground Truth' datasets. These datasets cover expected use cases, edge cases, and adversarial scenarios specific to each agent.

- **Multi-method testing:** Agents are continuously measured using a combination of automated deterministic testing, 'LLM-as-a-judge' methods, and human evaluation.
- **Release-based cadence:** Performance and accuracy scores are re-evaluated during every major release cycle (three times per year) and whenever updates are made to an agent.

To maintain long-term quality and prevent degradation, Veeva integrates architectural safeguards and continuous monitoring as applicable to each agent:

- **Factual grounding (RAG):** Agentic Text Monitoring for End Users maintains accuracy through retrieval-augmented generation of PIs and Compliance Documents, which ensures responses are derived from a customer's specific, approved Vault data rather than the model's general training data, minimizing hallucinations and bias.
- **Predictable architecture and drift mitigation:** The system is designed for predictable performance to mitigate model drift over time. Models and real-world conditions are checked for drift and relevancy during every standard release testing process.
- **Human-in-the-loop:** A core design principle requires human experts to review, validate, and approve critical AI outputs — such as call reports populated by Agentic Voice or Agentic Text Monitoring-flagged notes.
- **Model stability:** To ensure a consistent user experience, model versions are typically held constant for four to eight months, with updates announced in the Vault CRM release documentation. As model providers update their model offerings or retire older versions, we continue to re-assess the best fit-for-purpose model and version.
- **Administrative observability:** Administrators can access Agent Instance and Agent Action Execution records to monitor usage patterns, success rates, and the quantity of AI interactions daily.
- **Object, field, and record security inheritance:** Agents automatically inherit the initiating user's existing data access permissions and sharing rules, ensuring AI-generated insights remain within established governance boundaries.

For legal and compliance stakeholders, the practical significance is this: Veeva evaluates the agents like regulated-industry software, measures them like AI systems, and governs them like any other Vault capability. Accuracy is not asserted — it is tested on a documented cadence against expert-verified datasets with results that can be discussed in your diligence process.

Summary

The strategy of restricting free text in pharmaceutical CRM systems was a reasonable response to the enforcement environment of the early 2000s. Today it has run its course. Documentation did not stop with restrictions on free text: It scattered into systems with no monitoring, no screening, and no analytical value. A contemporary risk assessment compares the Agentic Call Report not to silence, but to that scattered reality. The double standard that limits entry in CRM while tolerating invisible documentation elsewhere does not survive that comparison.

The Agentic Call Report offers a better path: voice-focused capture; real-time screening that catches PHI, adverse events, off-label information, and inappropriate content at the point of entry; retrospective agentic monitoring that replaces keyword noise with contextual review; and a system of record that produces Commercial Evidence for the business and documented oversight for the compliance program — at the same time, from the same free text.

You can adopt the features individually or together, with rules sourced appropriately for each audience: published guardrails for the field, independent criteria for the home office. The current direction of AI regulation — transparency, human oversight, testing, accountability — rewards this design.

Veeva is committed to providing life sciences organizations with tools that reflect the current regulatory and technology landscape. Please reach out to your Veeva account partner if you have questions or would like to discuss configuration options for your organization.

Appendix A:

Agentic Text Monitoring for End Users Screening Categories

The table below summarizes the out-of-the-box screening categories available in Vault CRM Agentic Text Monitoring for End Users, what each category detects, and how the agent responds. The agent assesses based on regulatory obligations most relevant to the category in the country in which the interaction takes place. The Corporate Compliance category is configurable via plain-language instructions and is additive to the other categories.

Screening Category	What the Agent Detects	How the Agent is Grounded (U.S. Example)	How the Agent Responds
Protected Health Information (PHI) / Personally Identifiable Information (PII)	Personal health or identifying information entered into unencrypted free text fields — including names, birthdates, identifying case details, and clinical information about identifiable individuals as defined by local regulations	No PII per Staff Manual Guide 3297.4 and OMB Circular No. A-130; no PHI per HIPAA Safe Harbor §164.514(b)	Alerts user in real time; configurable to prevent save until identifiable data is de-identified or removed from the note
Adverse Events (AEs) and Product Quality Complaints (PQCs)	Language indicating a patient or HCP experienced an unexpected or undesirable reaction, clinical event, or health outcome potentially associated with a company product; language indicating a product defect, contamination concern, labeling error, packaging issue, or other quality event	No AEs as defined in 21 CFR 314.80 and 21 CFR 310.305 for drugs and 21 CFR Part 803 for devices; no SAEs as defined in 21 CFR 314.80 for drugs and 21 CFR Part 803 for devices; no PQCs as defined in 21 CFR 211.198 for drugs and 21 CFR Part 820 for devices	Alerts user in real time to report through the pharmacovigilance system; configurable to provide a link to the appropriate reporting system and prevent save until the AE or PQC is removed from the note
Off-Label Information	Language indicating an HCP has asked about an unapproved use, indication, dosing, or patient population outside the product's approved labeling	Claims must be consistent with the date-versioned Prescribing Information (RAG-assisted); reference Title 21 CFR and current FDA guidance documents; special considerations for combination products, comparative claims, clinical trial discussions, etc.	Alerts user in real time; configurable to prevent save until the off-label information is removed; Medical Inquiries can be created directly from the call report's action menu

Screening Category	What the Agent Detects	How the Agent is Grounded (U.S. Example)	How the Agent Responds
Inappropriate / Profane Content	Profane, harassing, discriminatory, or otherwise inappropriate language	Baseline content filters from the LLM; rules listed in the company-specific Corporate Compliance Document(s) (RAG-assisted)	Alerts user in real time; configurable to prevent save until content is revised or removed
Restricted and Not-Allowed Products	Company detail products that are restricted and/or not allowed for accounts on the call	End user's My Setup Products and the accounts' restricted/ allowed product data	Alerts the user in real time if company products mentioned in the note violate existing business logic for Restricted and Allowed Product features; configurable to prevent save until content is revised or removed
Corporate Compliance	Company-specific rules listed in plain-language documents configured in the Vault	Rules listed in the country- and/or product-specific Corporate Compliance Document(s) (RAG-assisted)	Alerts the user in real time; configurable to prevent save until the text is revised or removed

Appendix B:

Comparing Approved Notes, Agentic Text Monitoring for End Users, and Agentic Text Monitoring for Home Office

	Approved Notes	Agentic Text Monitoring for End Users	Agentic Text Monitoring for Home Office
Technology	Traditional software; keyword and phrase matching against defined monitoring rules	Large language model; contextual analysis of free text	Large language model; contextual analysis of free text
Timing	Retrospective, scheduled background job after records are submitted	Real time, at the point of entry, before the record is saved	Retrospective, near real-time after records are submitted
Primary audience	Compliance and monitoring teams (home office)	End users (field)	Compliance and monitoring teams (home office)
Source of rules	Phrase lists authored by compliance users	Out-of-the-box categories grounded in country-relevant regulatory obligations, plus the customer's corporate compliance documents (RAG)	Plain-language monitoring criteria authored and maintained by the compliance organization
End user experience	Invisible to end users	Interactive; alerts highlight the concern and direct users to the right channel	Invisible to end users
False positive profile	High; phrase matching cannot read negation, context, or intent	Low; contextual analysis respects legitimate use cases and regulatory safe harbors	Low; contextual analysis of what the note actually says
Compliance role	Detection	Prevention	Detection
Human review	Compliance users review flagged phrases	End user resolves alerts before save	Compliance users review flagged phrases