

# ChatGPT vs. Copilot in Veeva: GxP Compliance Guide

By Adrien Laurent, CEO at IntuitionLabs • 2/19/2026 • 50 min read

veeva vault

microsoft copilot

gxp compliance

chatgpt enterprise

data integrity

21 cfr part 11

ai governance

life sciences ai



## Executive Summary

The adoption of generative AI tools such as OpenAI's ChatGPT and Microsoft's Copilot in life-sciences environments (particularly within Veeva's regulated content platforms) promises significant gains in productivity – from faster content generation to more efficient knowledge discovery <sup>(1)</sup> [rankstudio.net](http://rankstudio.net) <sup>(2)</sup> [www.veeva.com](http://www.veeva.com)). However, these benefits come with substantial **compliance and governance risks** under GxP regulations (Good Practices for clinical, manufacturing, quality, etc.). Generative models are “black box” systems that can produce unpredictable outputs (“hallucinations”) <sup>(3)</sup> [www.qualitestgroup.com](http://www.qualitestgroup.com) <sup>(4)</sup> [www.pharmalife.com](http://www.pharmalife.com), and their handling of sensitive data (patient information, proprietary research, etc.) must meet strict regulatory standards. Key concerns include data confidentiality (risk of leaking IP or personal data), data integrity (ensuring outputs are accurate and traceable), auditability (meeting Part 11-style audit trail requirements), and the need for validation in a GxP quality system.

This report examines these issues in depth. We compare ChatGPT and Microsoft Copilot from a compliance standpoint (see **Table 1**), analyze how each integrates with Veeva platforms (e.g. Veeva Vault, PromoMats, QualityDocs) and what controls they provide, and review regulatory expectations (FDA 21 CFR Part 11/EU GMP Annex 11, forthcoming Annex 22, GDPR/HIPAA, etc.). We draw on expert analyses, vendor documentation, industry surveys, and emerging guidelines. Existing case studies and pilot experiences – such as Veeva customers embedding AI into Vault workflows – are used to illustrate practical implications. In general, **Copilot's Microsoft-driven integrations** (using Azure/AAD, Graph Connectors, and built-in security certifications) can mitigate some risks (e.g. enforcing Vault permissions and audit logging <sup>(5)</sup> [rankstudio.net](http://rankstudio.net) <sup>(6)</sup> [learn.microsoft.com](http://learn.microsoft.com)), but do not eliminate the need for rigorous oversight. **ChatGPT**, when not used via such controlled connectors, entails greater exposure (external data flow, model training on prompts, lack of built-in compliance logging) <sup>(7)</sup> [www.qualitestgroup.com](http://www.qualitestgroup.com) <sup>(8)</sup> [www.fiercepharma.com](http://www.fiercepharma.com)). In all cases, life-sciences firms must apply risk-based governance (akin to computerized system validation) to incorporate these tools. Our analysis concludes with recommendations: pilot deployments with strict controls, user training to review AI outputs, clear policies, and alignment with evolving regulations (e.g. EU Annex 22, FDA AI guidance). By proactively addressing these governance needs, companies can harness generative AI safely in Veeva environments without compromising patient safety or data integrity.

## Introduction and Background

The healthcare and life-sciences industries are rapidly exploring generative AI (GenAI) to automate document creation, compliance checks, and data analysis <sup>(9)</sup> [www.pharmalife.com](http://www.pharmalife.com) <sup>(10)</sup> [www.veeva.com](http://www.veeva.com)). Two prominent AI assistants are OpenAI's **ChatGPT** (particularly the ChatGPT Enterprise edition) and **Microsoft 365 Copilot**, which leverages OpenAI's models within Microsoft's ecosystem. ChatGPT is a conversational large-language model (GPT-4 family) accessible via web or API, capable of generating text and code from prompts. Microsoft Copilot integrates AI into Office products (Word, Excel, Teams, etc.) and now into life-sciences apps via **Graph Connectors** that index external data (such as Veeva Vault content) for Copilot to query <sup>(11)</sup> [rankstudio.net](http://rankstudio.net) <sup>(5)</sup> [rankstudio.net](http://rankstudio.net).

**Veeva Systems** provides cloud-based content management and CRM solutions for regulated industries. Its flagship product, **Veeva Vault**, includes submodules such as PromoMats (marketing content management), QualityDocs (SOPs, batch records), and Vault RIM (regulatory information management) – all built to comply with regulations (GxP, 21 CFR 11, Annex 11, etc.). Employees use Veeva Vault to store and manage controlled documents, standard operating procedures, submissions data, and other regulated content. As such, Veeva environments are typically **validated/qualified computerized systems** within a Quality Management System (QMS) <sup>(12)</sup> [www.ey.com](http://www.ey.com) <sup>(13)</sup> [fdainspections.com](http://fdainspections.com)). Any integration or automation involving Vault or Vault data must preserve data integrity, access controls, audit trails, and other GxP requirements.

Generative AI in this context promises to **automate tedious tasks**: for example, Copilot can draft summaries of regulatory documents or suggest compliant marketing messages based on existing SOPs (<sup>[14]</sup> rankstudio.net) (<sup>[15]</sup> techcommunity.microsoft.com); Veeva AI Agents (announced in 2025) aim to embed AI assistance into Vault itself (<sup>[16]</sup> www.veeva.com). Past industry surveys show interest in AI for productivity: over 90% of life sciences companies have some AI initiatives, but concerns remain high – e.g. a 2024 survey found 65% of major pharma banned ChatGPT outright over data-leak fears (<sup>[8]</sup> www.fiercepharma.com), and 91% of respondents agreed that “data security/privacy” are major AI concerns (<sup>[17]</sup> www.fiercepharma.com). In response to these concerns, vendors and regulators are racing to define enclosure: Microsoft offers Copilot connectors with permission controls (<sup>[5]</sup> rankstudio.net), OpenAI publishes enterprise privacy commitments (<sup>[18]</sup> openai.com), and regulatory bodies (FDA, EMA) are drafting AI-specific guidance (e.g. EU GMP Annex 22, EU AI Act).

This report systematically analyzes **ChatGPT and Microsoft Copilot usage in Veeva environments**, focusing on **compliance risks and GxP governance**. We start by reviewing the technology and integration mechanisms, then delve into specific risk categories (data security, output validity, audit requirements, regulatory context, etc.). We include tables comparing the platforms and summarizing risk factors. Throughout, we cite industrial experience and expert commentary to ground our conclusions. Our goal is a thorough, evidence-based guide for companies considering or currently deploying these AI tools in regulated life sciences workflows.

## Foundations of ChatGPT and Microsoft Copilot

### ChatGPT Overview

**ChatGPT** is a family of large language models from OpenAI (e.g. GPT-4) made accessible via chat or API. It can parse natural-language prompts and generate fluent text, code, and analyses. In early 2023, OpenAI introduced **ChatGPT Enterprise**, a business-grade offering with additional privacy and admin controls. According to OpenAI, Enterprise users **own their data** and “we do not train our models on your data by default” (<sup>[18]</sup> openai.com). Transcripts and conversation logs can be captured by OpenAI’s Compliance API (<sup>[19]</sup> help.openai.com), and enterprise workspaces allow admins to set retention policies (<sup>[19]</sup> help.openai.com). The Enterprise version also supports SAML single sign-on, and OpenAI has achieved SOC 2 audit compliance (<sup>[20]</sup> openai.com) with data encryption (AES-256 at rest, TLS in transit) (<sup>[20]</sup> openai.com). Nonetheless, ChatGPT operates on OpenAI’s cloud; by default it stores user prompts/outputs for model improvement unless explicitly opt-out (<sup>[7]</sup> www.qualitestgroup.com) (<sup>[18]</sup> openai.com). In 2024, OpenAI updated its enterprise privacy policy to explicitly commit that “you own your inputs and outputs” and “you control how long your data is retained” (<sup>[18]</sup> openai.com). These controls reduce but do not eliminate underlying compliance risks (see below).

ChatGPT’s **capabilities** include answering questions, drafting documents, summarizing or translating text, and generating code. Compared to traditional automated systems, it can handle ambiguous or poorly-specified queries and produce human-like prose (<sup>[9]</sup> www.pharmalive.com). This generality makes ChatGPT attractive for many tasks (from research brainstorming to medical inquiry). However, the model is non-deterministic and can “hallucinate” incorrect or irrelevant information (<sup>[4]</sup> www.pharmalive.com) (<sup>[3]</sup> www.qualitestgroup.com). Its training on public data means it may inadvertently reveal copyrighted content (<sup>[21]</sup> www.qualitestgroup.com). Each ChatGPT session also lacks intrinsic audit trails: only enterprise log modes or APIs capture transcripts (<sup>[19]</sup> help.openai.com), and there is no built-in mechanism to verify which knowledge sources the model used to generate its answers.

### Microsoft Copilot Overview

**Microsoft Copilot** is Microsoft’s branding for AI assistants integrated into its products. As of 2024–2025, the relevant versions include Copilot for Microsoft 365 (Word, Teams, Outlook, etc.) and domain-specific Copilots (e.g. Copilot for

Security). Behind the scenes, Copilot leverages OpenAI's GPT models via Azure, but with Microsoft's infrastructure and governance. Crucially for regulated industries, Microsoft has built **Graph Connectors** that allow Copilot to access enterprise data sources. For example, Copilot Connectors have been released for Veeva Vault modules: Vault RIM, PromoMats, QualityDocs (and others) can be "crawled" into Microsoft Graph (<sup>[22]</sup> rankstudio.net). When a user queries Copilot, it can search this indexed content to produce context-aware answers. Because this pipeline runs inside the corporate Azure/Microsoft 365 tenant, Copilot queries over Entra ID (Azure AD) authentication, and search results respect the original Veeva permissions (<sup>[5]</sup> rankstudio.net).

In practice, Copilot integration works as follows: An administrator configures a Microsoft 365 Copilot connector for a Veeva Vault (e.g. Vault PromoMats). Microsoft's crawl engine logs in as a designated Vault user and indexes allowed documents (text/PDF) and metadata into the Graph index (<sup>[5]</sup> rankstudio.net) (<sup>[23]</sup> rankstudio.net). Users can then ask Copilot questions in Word, Teams, or a "Copilot in Teams" chat, and Copilot will search the Graph index. Copilot can cite matching document names or passages, thanks to Microsoft's "citation embedding" methodology (<sup>[23]</sup> rankstudio.net). Importantly, the connector is **read-only** (it only pulls data into Graph; it does not push anything back to Vault) (<sup>[24]</sup> rankstudio.net). Microsoft emphasizes that connectors "follow Microsoft's comprehensive security, compliance, and privacy approach" (<sup>[24]</sup> rankstudio.net), keeping data encrypted and subject to the tenant's compliance policies. For example, if "Respect Veeva Vault permissions" is enabled (the default), Copilot will only show each user the Vault documents they are already authorized to see (<sup>[5]</sup> rankstudio.net). Microsoft 365 auditing automatically logs Copilot-related activities (user queries, document hits) if audit logging is enabled for the tenant (<sup>[6]</sup> learn.microsoft.com).

Both ChatGPT and Copilot thus bring advanced natural-language AI to enterprise users, but they differ in their control models. ChatGPT (outside of tightly-managed setups) is a generic chatbot environment; it lacks any native enterprise permissions and typically would index data passed to it on the open internet. Copilot, by contrast, is "anchored" in the Microsoft ecosystem: it uses corporate identity and controlled connectors. This difference has major compliance implications (as detailed below). We analyze those implications in the Veeva-specific context of GxP-regulated data management.

## Integration with Veeva Environments

### Veeva Vault and AI Connectors

Veeva Vault is a GxP-compliant content management platform. Its modules hold regulated documents like regulatory submissions, SOPs, and promotional materials. In 2024–2026, Veeva has actively pursued AI integration. Besides its own Veeva AI Agents (announced April 2025) (<sup>[16]</sup> www.veeva.com), Veeva has worked with Microsoft to create Copilot connectors for Vault. Official documentation and partner guides (e.g. RankStudio, Microsoft TechCommunity) detail how to **deploy the Veeva Copilot connector** (<sup>[22]</sup> rankstudio.net) (<sup>[5]</sup> rankstudio.net). Currently available connectors include Veeva Vault PromoMats, QualityDocs, and Vault RIM (<sup>[22]</sup> rankstudio.net).

A summary of these connectors' capabilities is given by RankStudio: once the Vault is authorized via AAD OAuth, the Microsoft crawl service indexes Veeva content into Graph (<sup>[22]</sup> rankstudio.net) (<sup>[5]</sup> rankstudio.net). This "unstructured data" (text, PDF) then becomes searchable in Copilot. The executive summary of that guide notes users can "leverage Copilot to draft compliant promotional materials based on research documents in PromoMats" and "quality/regulatory teams can rapidly review SOPs, track deadlines, and prepare audit responses" (<sup>[14]</sup> rankstudio.net). Reported benefits include automated tagging, summarization, and faster content retrieval, *while preserving compliance by respecting Veeva's permission model* (<sup>[14]</sup> rankstudio.net). Indeed, key statistics underline Vault's prevalence: over 250 organizations (including 12 of the top 20 pharma) use Vault RIM, and 450+ use Vault PromoMats (<sup>[25]</sup> rankstudio.net). In short, the Copilot connectors are a strategic enabling technology to extend AI into Veeva workflows.

Third-party integrators have also built tools linking ChatGPT with Veeva. For example, OneTeg (an enterprise iPaaS vendor) advertises a “Veeva Vault – ChatGPT Integration” (“Veeva Vault Digital Asset Management (DAM) and ChatGPT Cognitive Computing/AI apps with ... just a few clicks” (<sup>[26]</sup>oneteg.com)). Such connectors typically pull Veeva data into an intermediate system that calls the ChatGPT API. They illustrate corporate demand for using ChatGPT on Vault content, but also highlight risks: unlike Microsoft’s approach, these connectors would send Veeva data through an external platform and into ChatGPT’s cloud, breaking the controlled boundary. No official Veeva–ChatGPT plugin exists to date.

## Use Cases in Veeva Workflows

Generative AI can assist Veeva users in multiple ways. For **promotional content (PromoMats)**: Copilot can suggest wording for HCP communications by referencing approved materials in Vault. It can perform “pre-submission checks” on ads, flagging issues like overstated claims (<sup>[27]</sup>www.veeva.com), thus reducing errors before Medical/Legal/Regulatory (MLR) review. In fact, Veeva’s own Quick Check Agent (Agentic AI) is designed to flag compliance issues such as efficacy exaggeration (<sup>[27]</sup>www.veeva.com). Similarly, Microsoft’s Copilot can automatically fetch the latest Prescription Information sheets when an employee writes a sales email, ensuring the content is “compliant and up-to-date” (<sup>[28]</sup>techcommunity.microsoft.com). In **regulatory affairs (Vault RIM)**: Copilot could pose queries about submission timelines, pulling from indexed regulatory calendars; or assist in drafting submission summaries by analyzing prior filings in Vault. The RankStudio report notes Copilot could help “draft compliant promotional materials (‘personalized content for customer interactions...based on research documents in PromoMats’)” (<sup>[14]</sup>rankstudio.net), and that Veeva clients have seen 50% faster speed-to-market and 40% more content reuse with PromoMats alone (<sup>[1]</sup>rankstudio.net), suggesting large productivity gains when AI aids these processes.

In **quality systems (QualityDocs Vault)**, Copilot might help employees find and interpret Standard Operating Procedures (SOPs) on-demand. For example, a QA manager could ask Copilot “What is our CAPA procedure for out-of-specification results?” and receive a summary citing the relevant SOP in Vault. In such cases, compliance requires that Copilot’s output is traceable to the actual SOP. Microsoft’s “citation embedding” (showing which Vault documents the answer draws from) is intended to address this need (<sup>[23]</sup>rankstudio.net). If ChatGPT were used instead, one would need to manually feed it SOP text or references (since ChatGPT has no direct Vault access), raising similar traceability challenges.

These use cases show the appeal of speed and contextual insight. However, because Vault data is highly validated and regulated, any AI assistance must preserve the chain of evidence: query inputs and AI outputs may become part of regulatory records if used in official processes. The next sections examine these governance and compliance considerations in detail.

## Compliance Risks and GxP Governance

Using generative AI on Veeva content introduces many compliance risk categories. We organize them below, with evidence from industry sources and regulations. We then summarize with comparative tables.

### Data Security and Confidentiality

- **Data Leakage and Privacy.** Generative AI tools often require sending data to third-party servers. With **ChatGPT**, employees might be tempted to copy proprietary text or business queries into the ChatGPT interface. By default, **OpenAI retains all prompts and outputs** to improve its models (<sup>[7]</sup>www.qualitestgroup.com). This has already led to high-profile breaches: for example, a Samsung engineer accidentally pasted proprietary code into ChatGPT, which was then seen by other users (<sup>[29]</sup>www.qualitestgroup.com). Compliance experts warn that “all data fed into ChatGPT will remain on their servers forever” unless users opt out (<sup>[7]</sup>www.qualitestgroup.com). In a regulated context, this is

unacceptable: confidential chemistry formulas, clinical data, or manufacturing SOPs should never leave the company's secure systems. In fact, a 2024 survey found **65% of top 20 pharma companies banned ChatGPT**, explicitly citing fear that "sensitive internal data could be leaked" <sup>(8)</sup> [www.fiercepharma.com](http://www.fiercepharma.com)). Half of life-science respondents said their companies have prohibited ChatGPT use (with similar concern for Google Bard) <sup>(8)</sup> [www.fiercepharma.com](http://www.fiercepharma.com)).

By contrast, **Microsoft Copilot** connectors do *not* send Vault data to an external AI service. Instead, Veeva documents are *indexed* into the organization's own Microsoft Graph tenant. As one analysis explains, the connector "only indexes content that the designated Vault user can see... Copilot results only show items viewable by the querying user" <sup>(15)</sup> [rankstudio.net](http://rankstudio.net)). There is no write-back: Copilot operates on a *read-only* knowledge base pulled into Graph <sup>(24)</sup> [rankstudio.net](http://rankstudio.net)). Thus, Veeva content never transits to OpenAI servers or beyond the corporate cloud. Microsoft further emphasizes comprehensive tenant-controlled security (encryption, retention policies, etc.) <sup>(24)</sup> [rankstudio.net](http://rankstudio.net)). In practice, this means Copilot queries stay within corporate data governance and do not inherently expose new sensitive data to the outside. However, companies must still control how users phrase queries to avoid inadvertently retrieving disallowed information (e.g. someone might use Copilot to extract data they are not supposed to see if permissions were misconfigured, or might still quote Vault content into ChatGPT from outside).

- **Data Residency and Governance.** If ChatGPT is used at all, data typically flows to US-based servers (OpenAI's cloud). For non-US regulated data (e.g. EU personal data under GDPR), this is legally problematic: recent rulings (Schrems II) have barred transfer of EU personal data to the US without strong safeguards. A compliance analyst notes that "all content [ChatGPT] is processed and stored in the US... making the transfer of personal data to the US unlawful" for many EU companies <sup>(30)</sup> [www.qualitestgroup.com](http://www.qualitestgroup.com)). By contrast, Microsoft Copilot (Azure/Microsoft 365) allows data to remain within chosen regions/tenants subject to local compliance controls. Azure's data centers can be selected to meet GDPR or HIPAA requirements. This means that, under proper configuration, Veeva data need not leave its jurisdiction with Copilot. Nonetheless, any data used by either AI tool must still comply with HIPAA (for PHI) or GDPR (for personal data) if applicable. In sum, using ChatGPT on regulated data carries severe privacy risk without enterprise-only deployment, whereas Copilot's model is more amenable to corporate governance (though still requiring due care).

## Data Integrity and Output Reliability

- **Hallucinations and Accuracy.** Generative AI can produce plausible but incorrect answers. OpenAI itself warns that "ChatGPT will occasionally make up facts or 'hallucinate' outputs" <sup>(3)</sup> [www.qualitestgroup.com](http://www.qualitestgroup.com)). In a compliance setting, an AI hallucination could have serious consequences. Imagine GPT drafting a batch-production report with invented test results, or answering a regulatory query with fictitious references. As one industry expert notes, **regulatory and legal teams demand accuracy** and "what assurances can companies have that a bot's output will never be violative... or become subject to those hallucinations?" <sup>(4)</sup> [www.pharmalive.com](http://www.pharmalive.com)). The answer given by that expert was clear: *restrict the data set*. In biotech, Hallucinations are mitigated by feeding the model only pre-approved, controlled content: "They must pull from a restrictive, proprietary, and closed data set" <sup>(4)</sup> [www.pharmalive.com](http://www.pharmalive.com)).

Copilot's approach partially embodies this advice: since it searches the indexed Vault content, its answers (in principle) can be traced back to real documents. Microsoft's Graph integration also includes "citation embedding," which annotates Copilot outputs with links to the source documents <sup>(23)</sup> [rankstudio.net](http://rankstudio.net)). This makes outputs "verifiable" by design. Nevertheless, Copilot can still output incorrect summaries if underlying docs are flawed or if the AI misinterprets. Both OpenAI and Microsoft caution that human review remains essential: employees *must* check any AI-generated content for accuracy. A Qualitest guide specifically recommends that "Legal and compliance leaders should require employees to review any output generated by ChatGPT for accuracy, appropriateness, and actual usefulness" <sup>(31)</sup> [www.qualitestgroup.com](http://www.qualitestgroup.com)).

- **Reproducibility and Traceability.** GxP guidelines emphasize that results must be **reproducible and attributable** in a quality system. Traditional CSV systems have fixed algorithms; the same input always yields the same output. ChatGPT, however, can give different answers to the same prompt on different days (due to model updates and sampling randomness). This non-determinism poses a challenge for reproducibility. Copilot (with OpenAI 4) is similarly a probabilistic model, though Microsoft may freeze updates for enterprise connectors. In either case, the lack of deterministic logs is a concern.

Regulatory frameworks expect a chain-of-custody for data. For an AI-assisted document, one should be able to document *which version of the AI (which model checkpoint) was used*, what exact prompt and context was supplied, and what output it produced. Consultation with AI should arguably generate an audit trail. Copilot's logging (via Microsoft Purview) provides better support here than ChatGPT's default (which retains no such logs outside of Record Mode). Additionally, organizations should include the AI step in their validation and change-control processes: e.g., if ChatGPT is used to draft an SOP, the resulting SOP must go through formal approval with records of those review steps, just as if it were written by a person.

## Audit Trails, Records, and Part 11/GxP Validation

- **Audit Trails and Electronic Records.** Under regulations like FDA's 21 CFR Part 11, any **electronic record** or electronic signature must have secure, time-stamped audit trails that record creation, modification, or deletion of records (<sup>[13]</sup> [fdainspections.com](https://www.fda.gov/inspections-compliance-enforcement-and-civil-investigation)). In an AI context, the question is: if ChatGPT or Copilot contributes to a regulated document, how is that contribution recorded? Copilot, being part of M365, benefits from Microsoft's built-in audit infrastructure: any query and result passed through Office can be logged in the Unified Audit Log (if enabled) (<sup>[6]</sup> [learn.microsoft.com](https://learn.microsoft.com)). The Microsoft Learn documentation confirms that **user interactions with Copilot are automatically logged** as part of the audit (for example, "AIAppInteraction" events) (<sup>[32]</sup> [learn.microsoft.com](https://learn.microsoft.com)). Thus, an organization could retrieve who asked Copilot what question and when, satisfying part of the audit trail requirement. (However, note Microsoft's pay-as-you-go model for non-Microsoft AI logs (<sup>[32]</sup> [learn.microsoft.com](https://learn.microsoft.com)) – producers of integration need to ensure auditing is turned on.)

ChatGPT does not inherently log conversations in a way accessible to downstream compliance systems. Chat history is visible to the user but not exported by default. ChatGPT Enterprise can capture transcripts via its Compliance API (<sup>[19]</sup> [help.openai.com](https://help.openai.com)), but these must be explicitly configured. Without such logging, an inspection would have no record of what prompts were used or what advice ChatGPT gave. Therefore, any use of ChatGPT in a GxP process would either have to be very limited (with documented manual recording) or require a policy that all ChatGPT sessions go through an enterprise logging mechanism. The **FDAMAG (AI/Part11) guide** notes that ignoring these controls "can lead to systems that are not defensible during a regulatory inspection" (<sup>[33]</sup> [fdainspections.com](https://www.fda.gov/inspections-compliance-enforcement-and-civil-investigation)). For both tools, it is critical that companies treat AI interactions as part of the computerized system and extend their audit trails to cover them (e.g. by instrumenting API calls or by manual note capture).

- **Computerized System Validation (CSV)/Good Practice.** In GxP-regulated industries, software tools used for regulated tasks must generally be validated. This means documenting requirements, verifying they are met, and controlling changes. Traditional systems follow GAMP-5 life cycle or equivalents. The big question is: **Should ChatGPT or Copilot be "validated" like any other tool?**

Experts feel that, in the context of quality systems, the use of any AI tool in core processes likely triggers validation expectations. Christian Johner, an authority in med-device quality, explains that "in many cases, there is a regulatory obligation to validate tools such as ChatGPT, particularly when they are used in processes within the QMS" (<sup>[34]</sup> [blog.johner-institute.com](https://blog.johner-institute.com)). In other words, if a ChatGPT output directly influences a regulated decision or document, that use case should be risk-assessed and, if high-risk, validated accordingly. This might involve demonstrating that the AI tool consistently meets requirements (perhaps via sampling outputs, testing edge cases, etc.) and documenting any limitations. Johner suggests a risk-based approach: validate GPT usage in high-impact processes, but no need to validate every casual use. He warns against over-cautious bans: "Not using GPTs or LLMs due to CSV requirements

jeopardizes competitiveness”<sup>(35)</sup> [blog.johner-institute.com](https://blog.johner-institute.com)) – implying firms must adapt validation strategies rather than outright forbid the technology.

For **Copilot connectors**, the toolchain is a combination of Microsoft’s system and the Veeva connector. While Microsoft’s infrastructure is already validated in broad terms (Azure, M365, etc. are certified cloud platforms), the organization should assess the Copilot connector as part of its environment: e.g., test that it indexes correctly, respects permissions (as a configuration test), and that the Graph search works as intended. Because the connector is in “preview” (per Microsoft blogs<sup>(36)</sup> [techcommunity.microsoft.com](https://techcommunity.microsoft.com))), cautious deployment (starting with a small user group) is advised<sup>(37)</sup> [rankstudio.net](https://rankstudio.net)). Any unexpected behavior should be resolved before full rollout. Ultimately, the entire AI-assisted workflow (e.g. “Copilot drafts MLR review email”) should be validated end-to-end as a change to the existing content process. Both FDA’s CSA guidance and GAMP 5 recommend documented risk analysis and continuous monitoring for non-traditional tools like AI.

- **Regulatory Frameworks (Annex 11/Part 11, EU AI Act, etc.).** Generative AI sits at the intersection of several regulatory schemes. Under FDA 21 CFR Part 11 (USA) and EU GMP Annex 11 (EU), any automated system producing electronic records for regulated activities must be validated and auditable<sup>(13)</sup> [fdainspections.com](https://fdainspections.com))<sup>(38)</sup> [www.ey.com](https://www.ey.com)). Annex 11 requires “formal validation” of computerized systems, typically via IQ/OQ/PQ protocols. The FDA’s recent push for Computer Software Assurance (CSA) emphasizes a risk-based approach to validation, particularly for modern software. In effect, if ChatGPT or Copilot output becomes part of a regulated submission or quality record, the system that generated it falls under these rules.

Specific to AI, regulators are beginning to issue new guidance. The FDA (in Jan 2025) released draft guidance on AI models for drug submissions (proposing to require transparency about models used)<sup>(39)</sup> [www.fda.gov](https://www.fda.gov)) and device AI. Meanwhile, the EU has published a draft **Annex 22 to EudraLex vol 4 (GMP)** dedicated to AI, currently under consultation (expected finalization ~2026)<sup>(40)</sup> [www.ey.com](https://www.ey.com)). Annex 22 aims to “bridge traditional pharmaceutical manufacturing standards with modern AI-driven processes”<sup>(40)</sup> [www.ey.com](https://www.ey.com)). An EY analysis notes that compliance with **Annex 11/22, the upcoming EU AI Act, GDPR, and CSV/CSA principles is essential** to ensure regulatory trust<sup>(38)</sup> [www.ey.com](https://www.ey.com)). The EU AI Act (targeting foundation models and high-risk AI) could also apply, though generative tools used internally might be considered lower-risk if no safety-critical decisions are based solely on them.

In the U.S., any use of AI for tasks that influence drug approval documentation or safety reporting would have to meet Part 11 and current FDA guidance on electronic records. The FDAMAG guide summarizes: Part 11’s bedrock requirements (audit trails, unique signatures, access controls) still apply<sup>(13)</sup> [fdainspections.com](https://fdainspections.com)), and the novel “black box” nature of AI makes it hard to satisfy the traditional CSV mindset<sup>(41)</sup> [fdainspections.com](https://fdainspections.com)). To comply, companies may need to implement compensating controls (detailed below). In sum, while a definitive “AI regulation” for GxP is still evolving, existing frameworks clearly cover any AI-enabled system used in regulated operations.

## Intellectual Property and Copyright

Generative AI models are trained on vast text corpora which include copyrighted and proprietary material. There is a real risk that an AI-generated output might inadvertently infringe someone else’s IP, or reveal confidential information that the model “remembered” from training. OpenAI’s policy notes that ChatGPT “may collect personal information (this might also include your intellectual property)” and use it to improve services<sup>(21)</sup> [www.qualitestgroup.com](https://www.qualitestgroup.com)). Qualitest emphasizes that employees “should not share any confidential or sensitive information with ChatGPT and other LLM models”<sup>(42)</sup> [www.qualitestgroup.com](https://www.qualitestgroup.com)). Moreover, since some training datasets contain published drug labels and medical texts, an AI might regurgitate phrases verbatim, leading to copyright issues if reused without attribution.

By contrast, Microsoft’s Copilot connectors index *only* the organization’s own content (no external proprietary text is ingested). Therefore, Copilot’s outputs, when limited to that index, pose less risk of infringing external IP. The company still must ensure that outputs do not breach the copyright terms of the user’s own data sources (e.g. licensors of content). In practice, this means Copilot would not generate new external text beyond what is in Vault, whereas ChatGPT might do

so. Also, OpenAI's "user policy" suggests that outputs technically belong to the user (<sup>[18]</sup> [openai.com](https://openai.com)), but corporations must implement their own IP governance (e.g., lawyers should review AI-generated disclaimers). In regulated docs, it remains advisable for humans to verify that any generated text (e.g. assists in writing a procedure) is original or properly cited, to avoid plagiarism or IP conflict.

## Security and Threats

Both tools must be considered under a security lens. **Cybersecurity attacks** targeting AI systems are an emerging concern: malicious actors might attempt to poison training data, reverse-engineer models, or deceive AI into generating harmful outputs. ChatGPT has already experienced security incidents; for example, in early 2023 OpenAI temporarily shut down the service due to a bug that exposed user chat titles (<sup>[43]</sup> [www.qualitestgroup.com](https://www.qualitestgroup.com)). A compliance advisor warns that "the system can also generate phishing frauds; ChatGPT itself can be hacked and its behaviour could be altered" (<sup>[44]</sup> [www.qualitestgroup.com](https://www.qualitestgroup.com)). Any compromise of the AI service used in a regulated workflow (e.g. a rogue AI suggestion in a batch record) could lead to serious harm.

Copilot runs on Microsoft's secure Azure infrastructure, which benefits from industry-standard protections (SOC/FedRAMP, encryption, threat monitoring). However, it too can have vulnerabilities: if the connector indexed incorrect documents or if Azure services were down, Copilot answers would be affected. Importantly, both systems should be included in enterprise security assessments. Organizations should use tools to profile AI source authenticity when necessary (especially if Copilot is extended via custom "agents"). Also, as with any web-based tool, **phishing risk** arises: e.g. if employees are tricked into giving credentials on a fake ChatGPT or Copilot UI.

## Regulatory and Compliance Perspectives

Applying the above risks to Veeva/GxP yields the following insights:

- **Restricted Access and Policy:** Because of the high stakes, many companies start by banning ChatGPT on corporate networks. The FiercePharma survey found that a majority of large pharma companies have prohibited ChatGPT use for precisely this reason (<sup>[8]</sup> [www.fiercepharma.com](https://www.fiercepharma.com)). By contrast, Copilot (being integrated into sanctioned enterprise IT) is often allowed under IT policies, albeit with guidance. The key is formal policy: companies should **define where and how each AI tool may be used**. For example, ChatGPT might be barred from use on any Vault content, whereas Copilot via approved connectors might be permitted for specific tasks. Policies must also forbid uploading PHI or export-control data into ChatGPT (this is a compliance requirement akin to any third-party SaaS).
- **Audit and Documentation:** Auditing measures must treat AI activities as professional acts. For ChatGPT, this may mean capturing screen footage or transcripts of any conversation used to create a regulated document. For Copilot, IT can centrally retain logs. In either case, the goal is to preserve a compliance-grade record of how an AI contributed.
- **Quality Systems and Training:** GenAI use should be incorporated into the Quality System. This includes training employees on *safe AI usage*. Surprisingly, the FiercePharma survey found that **fewer than 60%** of life-sciences companies had provided any training or guidelines about safe ChatGPT use (<sup>[45]</sup> [www.fiercepharma.com](https://www.fiercepharma.com)). It is crucial to fix that: all users of ChatGPT/Copilot should receive guidance on data sensitivity (e.g. "never share confidential data") and output verification. The notion of AI as an "assistant, not author" should be enforced: the output must be reviewed by a qualified person.
- **Vendor Qualification:** Generative AI services may be considered "suppliers" or "cloud services" in a GMP vendor management sense. While OpenAI and Microsoft maintain security certifications, companies should include their AI services in supplier qualification assessments. For example, verify that Microsoft holds relevant certifications (ISO 27001, SOC 2/3, FedRAMP (<sup>[46]</sup> [www.genieai.com](https://www.genieai.com))) and that OpenAI meets ANSI SOC 2 (which it has) (<sup>[20]</sup> [openai.com](https://openai.com)). Companies must document reliance on these providers as part of their quality records, and have contingency plans in case of service interruption.
- **Ethical and Bias Considerations:** Life-science companies must ensure AI tools do not produce biased or unethical recommendations. Even if not a direct regulatory issue, biased output (e.g. a Copilot suggestion that implicitly favors a patient subgroup) could expose the company to compliance risk. No source covered bias extensively, but best practice (per international guidance) is that output is checked for fairness, and critical decisions are never automated without oversight.

## Summary of Risk Factors

The following table contrasts ChatGPT vs Copilot on key compliance aspects in the Veeva context. Citations to sources are provided in each row for evidence of the claims.

**Table 1.** Comparison of ChatGPT and Microsoft Copilot in Veeva/GxP context.

Aspect	ChatGPT (OpenAI)	Microsoft Copilot (M365)
<b>Data Integration</b>	No native Veeva connector; requires custom integration (e.g., via external API or iPaaS) ([26] oneteg.com). Often means copying text into ChatGPT UI or connecting via intermediary service.	Official Graph connectors for Veeva Vault exist (PromoMats, QualityDocs, RIM) ([22] rankstudio.net). These index Vault docs into Microsoft Graph for Copilot.
<b>Data Location &amp; Flow</b>	Data goes to OpenAI's cloud (by default US-based). OpenAI retains all prompts/outputs for training/improvement (unless opted-out) ([7] www.qualitestgroup.com) ([18] openai.com). Risk of sensitive data leaving control.	Data remains in corporate Azure tenant. Copilot connectors "only index content the Vault user can see" ([5] rankstudio.net). No Vault data is sent to external AI; all knowledge stays in Graph under tenant control.
<b>Permissions &amp; Access</b>	None. ChatGPT has no concept of Vault permissions; any user query could expose unauthorized info (if the user pastes it). No built-in filtering.	Honors Veeva access. If "Respects Vault permissions" is enabled (default), Copilot only returns documents the querying user is authorized to see ([5] rankstudio.net). (Admins watch out not to loosen this setting as it raises risk.)
<b>Output Validation</b>	Prone to "hallucinations" – nothing guarantees accuracy ([3] www.qualitestgroup.com). No citations unless specially requested (but ChatGPT's citation ability is unreliable). Must always be fact-checked by users.	Primarily searches validated content. Copilot supports "citation embedding" linking answers to source docs ([23] rankstudio.net). Outputs can still be incorrect, but they are grounded in actual Veeva documents for verification.
<b>Security &amp; Compliance</b>	ChatGPT Enterprise offers SOC 2, encryption, admin controls ([20] openai.com). ChatGPT Business version is similar. (Public ChatGPT has no enterprise guarantees.) Data default to US servers – GDPR/HIPAA issues ([30] www.qualitestgroup.com). OpenAI's new commitments say they will not train on your data by default ([18] openai.com), but.	
<b>COPILOT</b>		Microsoft cloud (Azure/M365) has extensive certifications (ISO 27001, SOC 1/2, FedRAMP, etc.) ([46] www.genieai.co). Data obeys tenant's retention/encryption policies ([24] rankstudio.net). Compliance APIs in Microsoft handle audit logging ([6] learn.microsoft.com).
<b>Audit Logging</b>	No built-in audit log of queries in standard ChatGPT. ChatGPT Enterprise can capture transcripts via Compliance API ([19] help.openai.com). Otherwise, record-keeping is manual.	Copilot interactions are logged in Microsoft 365 Audit (or Purview) automatically (when enabled) ([6] learn.microsoft.com). Organizations can view "AI application" logs to trace user queries.
<b>Compliance Cert.</b>	OpenAI publishes SOC 2 for enterprise offerings ([20] openai.com). No FedRAMP; unclear ISO. Data residency remains US (problematic for non-US).	Microsoft Copilot (as part of Microsoft Cloud) inherits Azure/AAD certifications (FedRAMP High, HIPAA-BAA eligible, etc.) ([46] www.genieai.co). Strong data sovereignty controls ([47] www.genieai.co) (regional datacenters).
<b>GxP Validation</b>	If ChatGPT output is used in regulated work, that use must be risk-assessed/validated (CSV) ([34] blog.johner-institute.com). Many companies default to banning it due to validation challenges ([8] www.fiercepharma.com).	Copilot connectors are enterprise-supported features; they should be tested (e.g. that indexing works correctly) but building on an accredited platform. Still require documentation of use-case validation and QMS review.

## Case Studies and Examples

While large-scale deployment of generative AI in GxP workflows is nascent, early pilots and user experiences provide insight:

- Modernizing PromoMats Workflows (Moderna).** Veeva's own case study reported that Moderna became the first customer to use Veeva AI's PromoMats Quick Check Agent (an integrated GenAI feature). Moderna's Global Marketing Ops Director noted that implementation was "very easy" because the AI agent "runs alongside our regular workflow without causing any disruption" <sup>(48)</sup> [www.veeva.com](http://www.veeva.com)). The agent flags compliance issues (like overstated claims) before MLR review, "rapidly sav[ing] countless hours" by summarizing review context <sup>(27)</sup> [www.veeva.com](http://www.veeva.com)). Moderna saw up to a 50% reduction in cycle times for content approvals **without compromising compliance** <sup>(48)</sup> [www.veeva.com](http://www.veeva.com)). This case highlights how a **validated, context-aware** AI embedded in Veeva can improve speed while maintaining quality. Notably, Moderna's usage was within Veeva's official framework (Veeva AI Agents), not a generic ChatGPT.
- Commercial CRM Enhancement (Bayer, GSK, Otsuka).** At the 2026 Veeva Summit, executives from Bayer, GSK, Boehringer, and Otsuka described how they architected "AI-ready" data platforms for sales and marketing. Key focus was unifying global data to power AI-driven sales plans (e.g. pre-call engagement planning, content recommendation) <sup>(49)</sup> [www.veeva.com](http://www.veeva.com)) <sup>(50)</sup> [www.veeva.com](http://www.veeva.com)). For instance, Veeva demonstrated a Vault CRM "Voice Agent" that auto-transcribes rep calls into CRM and a "Free Text Agent" that flags any compliance issues in call notes <sup>(51)</sup> [www.veeva.com](http://www.veeva.com)) <sup>(52)</sup> [www.veeva.com](http://www.veeva.com)). Otsuka Europe's representative explicitly mentioned working "closely with Veeva to address early compliance resistance" and to "confidently capture rich insights" from these agents <sup>(53)</sup> [www.veeva.com](http://www.veeva.com)). In other words, the company recognized stakeholders' compliance concerns and planned governance before full rollout. Bayer commented on building a master data architecture to feed AI accuracy <sup>(54)</sup> [www.veeva.com](http://www.veeva.com)), recognizing that poor data leads to AI errors. These examples illustrate that leading firms treat AI as a **data quality and governance challenge**: they prioritize harmonized, approved data feeds and stakeholder readiness to accept AI assistance.
- Integration Pilots (Copilot Connectors).** Microsoft's healthcare blog and documentation also include solution examples. For instance, one example shows a marketer in Word asking Copilot for the latest HCP-approved diagram, and Copilot fetching it from Vault PromoMats. Microsoft's blog claims Copilot can "monitor and ensure that all promotional materials... adhere to industry regulations," making compliance checks "more robust" <sup>(55)</sup> [techcommunity.microsoft.com](http://techcommunity.microsoft.com)). While this is marketing language, it suggests that Copilot is envisioned as a compliance assistant. Another example: a field rep using Copilot during a sales call to suggest content from Vault, thereby avoiding the risk of using outdated or incorrect information mid-call <sup>(28)</sup> [techcommunity.microsoft.com](http://techcommunity.microsoft.com)). No public study has quantified these pilots, but Microsoft reports thousands of Copilot customers industry-wide, and life-sciences adoption is growing.
- Internal Guidelines (Implementation Strategies).** Industry analysts recommend careful rollout strategies. For example, a LinkedIn article advising life-sciences companies on Copilot stresses **data centralization first**: store controlled documents (SOPs, protocols, agreements) in an accessible repository (e.g. SharePoint) to leverage Copilot's search <sup>(56)</sup> [www.linkedin.com](http://www.linkedin.com)) <sup>(57)</sup> [www.linkedin.com](http://www.linkedin.com)). One case study scenario described copying a Quality Agreement from the original DMS into SharePoint for a pilot, with a label: "Please refer to the Veeva document for official purposes. These copies are for training and information only as part of our pilot" <sup>(57)</sup> [www.linkedin.com](http://www.linkedin.com)). Such disclaimers ensure that the AI-trained version is not mistaken for the validated source. This approach – mirroring validated data into a sandbox environment with clear audit scope – exemplifies how companies can safely test AI. Other experts advise rolling out to a limited group, verifying that indexed content and permissions behave correctly before general use <sup>(58)</sup> [rankstudio.net](http://rankstudio.net)).

These examples underline that **successful use of AI in GxP settings requires blending technology with process changes**: deploying connectors or agents is just the first step. Equally important are policies (bans vs. guidelines), training (on AI "dos and don'ts"), and incremental pilots (with metrics like time saved, error reduction). The early adopters in life sciences are moving cautiously but clearly see a future in which AI becomes "baseline" for content lifecycle management <sup>(59)</sup> [rankstudio.net](http://rankstudio.net)) <sup>(60)</sup> [www.veeva.com](http://www.veeva.com)).

## Data Analysis and Discussion

Based on the above evidence, several themes emerge:

- Permission Boundaries Matter.** The Copilot connectors demonstrate that it is **technically feasible** to grant an AI assistant access to Vault data while preserving ACLs and audit logs <sup>(5)</sup> [rankstudio.net](http://rankstudio.net)). In a study by RankStudio, they found that with "Respect Veeva Vault permissions" enabled, Copilot only returns content the user is allowed to see <sup>(5)</sup> [rankstudio.net](http://rankstudio.net)). Conversely, any uncontrolled ChatGPT use trivially breaches these boundaries. This means that from a compliance perspective, Microsoft's approach is inherently safer. However, it places trust in Microsoft/Azure. Companies should still validate that the connector was set up securely (correct Azure app, federated IDs mapped properly, minimal scope) <sup>(5)</sup> [rankstudio.net](http://rankstudio.net)).

- Traceability Through Citations.** Verifiable AI output is a key compliance goal. Microsoft's "semantic index and citation embedding" helps: it can link answers back to Vault documents ([23] rankstudio.net). Qualitative evidence suggests that such traceability is valuable, but not foolproof – ultimately a human must verify the link is correct. No similar citation feature exists in ChatGPT (aside from unreliable "source" outputs). Thus if ChatGPT is used, organizations may require employees to manually footnote or cross-ref outputs. We strongly recommend any AI-generated content be treated as a *draft to be reviewed*, not as final documentation (even if it cites sources).
- Audit and Governance Are Essential.** The audit trails in Copilot (via Microsoft) are a huge enabler. For example, all user interactions with Copilot can flow into the company's Purview audit log with negligible extra setup ([32] learn.microsoft.com). In contrast, the default ChatGPT experience provides no such enterprise log. At best, one could enable "conversation history" retention or use the Compliance API, but this is an add-on and not as seamless. We note that Microsoft clearly signals to auditors that nothing about the connector violates the local system controls: the data is still "subject to the organization's compliance policies (e.g. data retention, encryption)" ([24] rankstudio.net). Life-science quality systems will need to capture Copilot logs (and retrain verification logs for ChatGPT) as part of the 21 CFR 11 audit trail.
- Content Quality and Validation.** The inherent variability of AI outputs contradicts the GxP principle of "consistency." In FDA terms, AI is more like a soft-coded expert system than a fixed calculation. The FDA compliance guide explicitly warns that the "black box" nature of AI conflicts with CSV, since regulators expect predictable, testable software behavior ([41] fdainspections.com). In practical terms, this means life-science firms cannot deploy ChatGPT or any AI "as is" for generation of controlled content without putting validation layers around it. For example, if Copilot is used in drafting an SOP revision, the revision must still go through all normal change-control steps (review by SMEs, QMS approvals, etc.). Some suggest applying GAMP 5's risk framework to the AI itself: for low-level tasks it may be used with less oversight, but for high-impact outputs (e.g. drafting a regulatory submission), stringent validation is needed.
- Regulatory Outlook.** Regulators are aware of these issues. The EMA and FDA have made preliminary moves: the EU is developing Annex 22 (drafted by mid-2025) to explicitly cover AI in GMP ([40] www.ey.com); the FDA has issued draft guidances on AI/ML in medical devices and is expected to focus on model credibility in submissions. Companies should watch these developments. For now, the safe assumption is that any deployment of generative AI in GxP must satisfy existing rules (Annex 11/Part 11, ALCOA data integrity, etc.). Notably, compliance frameworks like ICH Q9 (risk management) and FDA's CSA already encourage flexible approaches for novel tech – implying that a risk-based approach to validating LLM usage is acceptable if documented. In practice, that means getting ahead by crafting AI SOPs: e.g. how prompts are formulated, how often outputs are tested, how models are updated. As one expert put it, meeting Part 11 with AI "requires acknowledging and addressing these specific hurdles head-on" ([33] fdainspections.com).

Given these findings, the safest strategy is a **controlled pilot** approach. The evidence suggests starting small, measuring impact, and iterating on controls. Key performance indicators might include reduced labor hours (as in Veeva's 50% efficiency gains ([1] rankstudio.net)), fewer errors (e.g. compliance issues caught by AI vs. by humans), and user satisfaction. Meanwhile, compliance metrics (audit log completeness, validation status) should be tracked. Over time, more robust use cases can be unlocked, potentially extending agents to QC Vault, safety databases, clinical records, etc. ([61] rankstudio.net). Importantly, early adopters should document their experiences (anonymized case studies) to build industry best practices. (For example, Veeva encourages joining its AI community for knowledge sharing ([62] www.veeva.com).)

## Tables

**Table 1.** Comparing ChatGPT vs. Microsoft Copilot (M365) in Veeva/GxP environments (see discussion above).

Aspect	ChatGPT (OpenAI)	Microsoft Copilot (M365)
<b>Data Integration</b>	No native Veeva connector; requires custom integration (e.g. via external API/PaaS) ([26] oneteg.com). Often means copying Vault content into chat.	Official Graph connectors for Veeva Vault exist (PromoMats, QualityDocs, RIM) ([22] rankstudio.net). These index Vault docs into Microsoft Graph for Copilot.
<b>Data Location &amp; Flow</b>	Data is sent to OpenAI's cloud by default (US servers); OpenAI retains all prompts/outputs for model training/improvement unless opted out ([7] www.qualitestgroup.com) ([18] openai.com). Risk of sensitive data leaving control.	Data remains in corporate Azure/M365 tenant. Content is pulled into Graph; Copilot does not send data back out. All data stays under the organization's compliance policies (retention, encryption) ([24] rankstudio.net).
<b>Permissions &amp; Access</b>	ChatGPT has no built-in knowledge of Vault permissions. If a user inputs data they can see, they can query it; but ChatGPT could inadvertently expose logic or patterns to others ([7] www.qualitestgroup.com).	Copilot respects Vault ACLs. If "Respect Vault permissions" is enabled (default), Copilot results only show items the user is allowed to see ([5] rankstudio.net). (Administrators should not override this setting, as it would breach compliance.)

Aspect	ChatGPT (OpenAI)	Microsoft Copilot (M365)
Output Validity	May "hallucinate" incorrect information or blend unrelated sources ([3] www.qualitestgroup.com) ([4] www.pharmalive.com). No guaranteed citations (unless prompt carefully, which ChatGPT may ignore). Users must fact-check all outputs.	Uses actual indexed documents as the basis. Copilot embeds citations so answers can be traced to Vault content ([23] rankstudio.net). ChatGPT-like inaccuracies can still occur (if the question's answer isn't in any doc), but outputs are generally grounded, verifiable content.
Security & Compliance	ChatGPT Enterprise enforces encryption (AES-256/TLS) and SOC 2 compliance ([20] openai.com). No FedRAMP/ISO by default (dataset is global). Data in US may violate EU/GDPR rules ([30] www.qualitestgroup.com). OpenAI (as of 2026) commits "we do not train our models on your data by default" ([18] openai.com), and customers own inputs/outputs and retention.	Copilot runs on Azure (ISO 27001, FedRAMP High, HIPAA BAA available) ([46] www.genieai.co). and Microsoft applies "strong data sovereignty controls" ([47] www.genieai.co). Veeva data in Graph is encrypted and subject to corporate retention policies ([24] rankstudio.net).
Audit Logging	No native audit trail of conversations in standard ChatGPT. ChatGPT Enterprise can log transcripts via a Compliance API ([19] help.openai.com), but this must be configured. Otherwise, capturing chat history is manual.	User Copilot interactions (queries, results) are automatically logged in the Microsoft 365 Audit Log if enabled ([6] learn.microsoft.com), providing an easily accessible trace of Copilot usage. (Non-Microsoft AI logs filter to pay-as-you-go billing, but Microsoft Copilot is covered under standard audit.)
GxP Validation & Governance	Any regulated output from ChatGPT use must be validated with a risk-based approach ([34] blog.johner-institute.com). Many firms ban it to avoid validation burden ([8] www.fiercepharma.com). If used, it should be within a controlled pilot with QA oversight.	The Copilot connector is a Microsoft-sanctioned feature; organizations should test it (e.g. verify indexing and ACL behavior ([5] rankstudio.net)). It then works within the validated M365 platform. Still, usage cases should be documented, and outputs should go through normal review/approval steps.
IP/Copyright	Trained on public internet; output can inadvertently copy copyrighted content or internal proprietary patterns ([21] www.qualitestgroup.com). Risk of IP exposure.	Indexes only customer's own (licensed) content. Does not generate new content beyond what's in Vault. Copyright/IP risk is limited to verifying Vault content usage.
Provider Dependency	OpenAI as an external service. Service downtime or changes (e.g. ChatGPT went offline for bug fixes ([63] www.fiercepharma.com)) are beyond company control.	Microsoft's SLA and Azure infrastructure apply. Copilot availability ties into corporate cloud; outages are managed like any Azure service.

Table 2. Risk profile and mitigation in Veeva/GxP use of generative AI.

Risk Category	ChatGPT (OpenAI)	Microsoft Copilot (M365)
Data Confidentiality	High risk if used freely. Default logging by OpenAI ([7] www.qualitestgroup.com) risks leaking OPD/IP. Survey: 65% of top pharmas explicitly banned it over data leak concerns ([8] www.fiercepharma.com).	Data stays in company's cloud. Connector respects Vault ACLs ([5] rankstudio.net), so confidentiality is maintained across systems. (Admin must ensure connector user has minimal rights.)
Data Integrity (Hallucination)	High risk of incorrect output ("hallucinate") ([3] www.qualitestgroup.com). Any content provided by ChatGPT must be reviewed; direct use in compliance docs is unsafe.	Lower risk, since output is based on indexed Vault documents and includes citations ([23] rankstudio.net). Still, responses need human verification, but errors are easier to trace back to source.
Auditability	Poor. No built-in logging of queries in standard ChatGPT. Must use compliance API or manual recording ([19] help.openai.com) ([6] learn.microsoft.com).	Good. Copilot use is logged in M365 audit trails ([6] learn.microsoft.com). Questions asked and documents retrieved can be reviewed by IT/Audit.
Validation (21 CFR/Annex11)	Not inherently validated; if used for regulated tasks, must be validated per CSV. Johner Institute warns auditors will expect GPT validation in QMS processes ([34] blog.johner-institute.com). Many companies avoid it.	The connector is a new component but uses existing validated platforms (Azure, M365). It should be covered by CSV authorizations, though connectors may need test-qualification. Outputs still follow normal document validation cycles.
Copyright/IP	ChatGPT's training data includes copyrighted text; employees must be careful not to infringe when using outputs ([21] www.qualitestgroup.com). Sharing proprietary Vault data with ChatGPT risks exposing company IP.	Outputs come from company's Vault content. No external copyrighted data is involved. Must ensure Vault document rights (e.g. licensed materials) cover AI use.
User Policy/Training	Requires strong governance. Without guidelines, users may upload PHI or trade secrets. Survey says <60% of companies provided any ChatGPT policy ([45] www.fiercepharma.com) – a known shortfall.	Still needs policy, but many companies already allow Copilot under IT. Key is to train users that Copilot is an assistant and outputs require review.
Regulatory Enforcement	If abused, risk of regulatory action for data breach or non-compliance. Given bans and scrutiny ([8] www.fiercepharma.com) ([41] fdinspections.com), ChatGPT usage in GxP is seen as high hazard.	As part of corporate systems, Copilot is more expected. Regulators will still inspect usage (e.g. audit logs, change control). Proper controls can satisfy Part 11/GMP examiners more easily than free-form ChatGPT.

## Implications and Future Directions

Generative AI is rapidly moving from experimental to expected capability. For the life sciences, the question is not if these tools will be used, but how to **use them responsibly**. Several **implications and future trends** can be anticipated:

- **AI Governance Frameworks:** Companies should extend their Quality and IT governance to cover AI. This means formal SOPs for AI tool usage, training programs, performance monitoring, and incident response. For example, a standard might say: "All ChatGPT/Copilot interactions involving Vault data must be documented in the system's audit trail, and any output used in a GxP deliverable must be reviewed by a qualified person." Some organizations are already building "AI Governance Committees" to oversee compliance (drawing from ISPE and FDA guidance). GAMP 5 and quality risk management principles should be applied: e.g. classify AI as either a configurable tool or as an "AI model service," and define validation scope accordingly.
- **Evolving Regulations:** Companies must stay attuned to AI regulations. The EU's draft Annex 22 will likely define requirements for AI in GMP processes: early drafts emphasize risk-based validation similar to current Annex 11 practices, but updated for AI specifics. The EU AI Act (likely finalized in 2026) may impose transparency obligations (e.g. documenting AI model providers, bias mitigations). In the US, the FDA may update Part 11 or issue AI-specific guidance (for example, the FDAMAG notes that enforcing Part 11's requirements on opaque AI could be difficult without new rules <sup>(41]</sup> [fdainspections.com](#)). Staying involved in regulatory discussions or industry consortia will help Veeva users prepare when new standards arrive.
- **Tool Maturation:** The capabilities of these AI tools will continue to improve. Copilot connectors will likely graduate from preview to GA, with better support for large files, OCR for scanned docs, and maybe even real-time syncing <sup>(64]</sup> [rankstudio.net](#)). ChatGPT (or its competitors) may develop enterprise features such as on-premises hosting or better markdown citations. Meanwhile, Veeva's own AI agenda ("Agentic AI") promises built-in assistants that natively understand Vault context <sup>(16]</sup> [www.veeva.com](#) <sup>(65]</sup> [www.veeva.com](#)). As the technology matures, the distinction between ChatGPT, Copilot, and Veeva's internal AI may blur: Microsoft and Veeva are partners, and Veeva AI was said to be LLM-agnostic <sup>(66]</sup> [www.veeva.com](#)). Life-sciences firms should track these developments, as future regulatory integration may rely more on Veeva's or Azure's offerings than on standalone ChatGPT.
- **Insurance and Liability:** An emerging topic is who is liable if AI leads to a compliance breach or patient harm. If a Pharmacopeial error slips through because an AI approved it, that could trigger 483 observations or worse. While not addressed directly in the literature cited here, companies should consider this risk. Potential mitigations include AI "black box" insurance (some insurers offer policies for AI-related errors) and contractual protections (e.g., Microsoft's and OpenAI's terms of service disclaimers).
- **Ethical and Societal Factors:** Finally, the use of AI in regulated domains raises ethical questions (e.g. transparency with patients when AI is used). For instance, if ChatGPT drafts part of a patient communication, should that be disclosed? GxP culture emphasizes accountability, so "disguised" AI outputs might be frowned upon. Industry bodies may eventually issue codes of practice. For now, the focus is on internal governance, but companies should monitor broader debates (e.g. NIH AI guidelines, healthcare ethics opinions).

## Conclusion

ChatGPT and Microsoft Copilot offer transformative capabilities for life-sciences organizations, including those using Veeva's platforms. They can dramatically accelerate document creation, analysis, and compliance tasks. However, **the compliance risks are real and multi-faceted**. Uncontrolled use of ChatGPT can easily violate data privacy, integrity, and audit requirements; many pharma firms have already banned it to avoid these dangers <sup>(8]</sup> [www.fiercepharma.com](#)). Conversely, Copilot – when used via Veeva connectors – provides a more controlled path, with enterprise-level security and permission enforcement <sup>(5]</sup> [rankstudio.net](#) <sup>(6]</sup> [learn.microsoft.com](#)). Even so, neither tool obviates the need for human oversight.

The governance strategy is clear: **apply risk-based validation and controls** to AI use. This means (a) categorizing use cases (e.g. brainstorming vs. final report writing), (b) ensuring AI is hosted or accessed in a compliant manner (prefer EMS-based connectors over open ChatGPT), (c) training staff to treat AI outputs critically, and (d) documenting everything in the QMS (audit trails, SOPs, change controls). Regulators expect such precaution. As Johner's guidance emphasizes, auditors will ask about validation of AI tools <sup>(34]</sup> [blog.johner-institute.com](#)), and companies must be ready with evidence that actions and outcomes are controlled.

Looking ahead, we anticipate that integration of AI into regulated workflows will only increase. Early adopters have demonstrated feasibility (e.g. Moderna's experience <sup>(48]</sup> [www.veeva.com](#)). Standards are evolving to catch up (e.g. Annex 22 for AI in GMP), and regulators encourage safe innovation. By proactively addressing the compliance requirements – rather than banning the technology outright – life-sciences companies can derive the productivity gains noted (50% faster processes <sup>(1]</sup> [rankstudio.net](#)) while maintaining the trust of regulators and the public.

**Key Takeaway:** GenAI in Veeva environments is **no longer hypothetical**. It demands a comprehensive governance approach: enterprise-grade integration (Copilot connectors), risk-based validation, auditability, and trained oversight. When implemented thoughtfully, it can become a powerful assistant; when overlooked, it can become a compliance liability. Organizations should pilot these tools with full awareness of the considerations outlined here and adapt their quality systems accordingly.

## References

- Johner Institute, “*Validating ChatGPT: What medical device manufacturers need to consider*,” Oct 29, 2025, (Johner notes auditors will ask about GPT validation and that it must be used with risk-based validation in QMS processes (<sup>[34]</sup> [blog.johner-institute.com](https://blog.johner-institute.com))).
- RankStudio, “*Veeva Copilot Connectors: Guide to M365 Graph Integration*,” Nov 3, 2025 (describes Copilot-Vault connectors, highlighting semantic search and compliance benefits (<sup>[22]</sup> [rankstudio.net](https://rankstudio.net)) (<sup>[14]</sup> [rankstudio.net](https://rankstudio.net))).
- Microsoft Tech Community, Michael Gannotti, “*Copilot Agents Solutions Series – Connecting Veeva Vault PromoMats*,” Apr 16, 2025 (discusses integrating Copilot with PromoMats for compliance-enhanced content access (<sup>[55]</sup> [techcommunity.microsoft.com](https://techcommunity.microsoft.com)) (<sup>[67]</sup> [techcommunity.microsoft.com](https://techcommunity.microsoft.com))).
- Veeva Systems, *Press release “Announcing Veeva AI*,” Apr 29, 2025 (introduces Veeva AI Agents/Shortcuts; emphasizes AI that is “secure and compliant for life sciences” (<sup>[16]</sup> [www.veeva.com](https://www.veeva.com)) (<sup>[68]</sup> [www.veeva.com](https://www.veeva.com))).
- Microsoft Learn, “*Audit logs for Copilot and AI applications*,” Aug 20, 2025 (explains that Copilot interactions generate logs under Microsoft’s audit framework (<sup>[6]</sup> [learn.microsoft.com](https://learn.microsoft.com))).
- EY Switzerland, “*AI validation in pharma: maintaining compliance and trust*,” Oct 21, 2025 (explains that compliance with Annex 11/22, EU AI Act, GDPR, and CSV/CSA is essential for AI in GxP (<sup>[38]</sup> [www.ey.com](https://www.ey.com)) (<sup>[40]</sup> [www.ey.com](https://www.ey.com))).
- Gregg Fisher & Mike Spitz, “*Practical applications for ChatGPT... in biopharma*,” PharmaLive (Feb 2024) (industry expert guidance on restricting ChatGPT’s training data and requiring citations in outputs (<sup>[4]</sup> [www.pharmalife.com](https://www.pharmalife.com)) (<sup>[69]</sup> [www.pharmalife.com](https://www.pharmalife.com))).
- Gourav Pandey, “*Unlocking the Power of Microsoft 365 Copilot in the Life Sciences Industry*,” LinkedIn, Nov 29, 2024 (advice on data centralization and pilot implementation; sample disclaimer linking back to Veeva documents (<sup>[57]</sup> [www.linkedin.com](https://www.linkedin.com)) (<sup>[56]</sup> [www.linkedin.com](https://www.linkedin.com))).
- Veeva Systems (Europe) blog, Philipp Luik, “*How Top Biopharmas Transform Commercial Execution with Agentic AI*,” Jan 9, 2026 (describes AI use cases in CRM and PromoMats; notes first Veeva AI customer improved speed/compliance, and reports that scalable AI hinges on a unified data foundation (<sup>[2]</sup> [www.veeva.com](https://www.veeva.com)) (<sup>[70]</sup> [www.veeva.com](https://www.veeva.com))).
- Qualitest, Somel Pal, “*Demystifying Compliance Risks: A Guide to Navigating ChatGPT Safely*,” Jun 30, 2023 (summarizes ChatGPT risks: data leaks, OpenAI’s default data collection, hallucination, IP/copyright, etc. (<sup>[29]</sup> [www.qualitestgroup.com](https://www.qualitestgroup.com)) (<sup>[3]</sup> [www.qualitestgroup.com](https://www.qualitestgroup.com))).
- FiercePharma, Andrea Park, “*Two-thirds of top 20 pharmas have banned ChatGPT... survey finds*,” Apr 19, 2024 (survey data on pharma AI policies: 65% banned ChatGPT citing data leak risks (<sup>[8]</sup> [www.fiercepharma.com](https://www.fiercepharma.com)); >80% worried about data security (<sup>[17]</sup> [www.fiercepharma.com](https://www.fiercepharma.com)); most hadn’t trained employees on AI use (<sup>[45]</sup> [www.fiercepharma.com](https://www.fiercepharma.com))).
- ISPE *Pharmaceutical Engineering*, Heitmann et al., “*ChatGPT, BARD, and Other Large Language Models Meet Regulated Pharma*,” Jul/Aug 2023 (peer-reviewed article identifying LLM risks and use cases; notes some companies already block ChatGPT in pharma networks (<sup>[71]</sup> [ispe.org](https://ispe.org)) and emphasizes that LLM outputs must be critically reviewed and documented (<sup>[72]</sup> [ispe.org](https://ispe.org)) (<sup>[73]</sup> [ispe.org](https://ispe.org))).
- FDA Inspections Blog, “*AI and FDA Part 11 Compliance: A Complete Guide for 2025*,” Sept 27, 2025 (guidance on AI compliance: reviews Part 11 basics; highlights the “black box” validation challenge for AI (<sup>[41]</sup> [fdainspections.com](https://fdainspections.com));

recommends strong AI governance and SOPs for model lifecycle (<sup>[74]</sup> fdainspections.com)).

- OpenAI Enterprise Privacy page, “Our commitments (Jan 2026),” OpenAI Help (outlines enterprise data commitments: not training on customer data by default, customer owns inputs/outputs, SOC 2 audit, encryption (<sup>[18]</sup> openai.com) (<sup>[20]</sup> openai.com)).

## External Sources

- [1] <https://rankstudio.net/articles/veeva-copilot-connector-guide#:~:prepa...>
- [2] [https://www.veeva.com/eu/wp-content/uploads/2023/12/%3AUsers%3Acaitinrothert-mbpr16%3ALibrary%3ACaches%3AAdobe%20InDesign%3AVersion%2019.0%3Aen\\_US%3AInDesign%20ClipboardScrap1.pdf#:~:Desig...](https://www.veeva.com/eu/wp-content/uploads/2023/12/%3AUsers%3Acaitinrothert-mbpr16%3ALibrary%3ACaches%3AAdobe%20InDesign%3AVersion%2019.0%3Aen_US%3AInDesign%20ClipboardScrap1.pdf#:~:Desig...)
- [3] <https://www.qualitestgroup.com/insights/blog/demystifying-compliance-risks-a-guide-to-navigating-chatgpt-safely/#:~:The%2...>
- [4] <https://www.pharmalive.com/practical-applications-for-chatgpt-and-other-large-language-models-in-biopharma/#:~:lies%...>
- [5] <https://rankstudio.net/articles/veeva-copilot-connector-guide#:~:Throu...>
- [6] <https://learn.microsoft.com/en-us/purview/audit-copilot#:~:This%...>
- [7] <https://www.qualitestgroup.com/insights/blog/demystifying-compliance-risks-a-guide-to-navigating-chatgpt-safely/#:~:With%...>
- [8] <https://www.fiercepharma.com/marketing/two-thirds-top-20-pharmas-have-banned-chatgpt-and-many-life-sci-call-ai-overrated-survey#:~:In%20...>
- [9] <https://www.pharmalive.com/practical-applications-for-chatgpt-and-other-large-language-models-in-biopharma/#:~:Numer...>
- [10] [https://www.veeva.com/eu/wp-content/uploads/2023/12/%3AUsers%3Acaitinrothert-mbpr16%3ALibrary%3ACaches%3AAdobe%20InDesign%3AVersion%2019.0%3Aen\\_US%3AInDesign%20ClipboardScrap1.pdf#:~:The%2...](https://www.veeva.com/eu/wp-content/uploads/2023/12/%3AUsers%3Acaitinrothert-mbpr16%3ALibrary%3ACaches%3AAdobe%20InDesign%3AVersion%2019.0%3Aen_US%3AInDesign%20ClipboardScrap1.pdf#:~:The%2...)
- [11] <https://rankstudio.net/articles/veeva-copilot-connector-guide#:~:produ...>
- [12] [https://www.ey.com/en\\_ch/insights/life-sciences/gxp-and-ai-tools-compliance-validation-and-trust-in-pharma#:~:AI%20...](https://www.ey.com/en_ch/insights/life-sciences/gxp-and-ai-tools-compliance-validation-and-trust-in-pharma#:~:AI%20...)
- [13] <https://fdainspections.com/ai-fda-part-11-compliance-guide/#:~:The%2...>
- [14] <https://rankstudio.net/articles/veeva-copilot-connector-guide#:~:Key%2...>
- [15] <https://techcommunity.microsoft.com/blog/healthcareandlifesciencesblog/copilot-agents-solutions-series-%E2%80%93-connecting-veeva-vault-promomats/4404882#:~:Once%...>
- [16] <https://www.veeva.com/resources/announcing-veeva-ai/#:~:Veeva...>
- [17] <https://www.fiercepharma.com/marketing/two-thirds-top-20-pharmas-have-banned-chatgpt-and-many-life-sci-call-ai-overrated-survey#:~:And%2...>
- [18] [https://openai.com/enterprise-privacy/?industry=data\\_processing\\_hosting#:~:You%2...](https://openai.com/enterprise-privacy/?industry=data_processing_hosting#:~:You%2...)
- [19] <https://help.openai.com/en/articles/11664471#:~:imme...>
- [20] [https://openai.com/enterprise-privacy/?industry=data\\_processing\\_hosting#:~:Compr...](https://openai.com/enterprise-privacy/?industry=data_processing_hosting#:~:Compr...)
- [21] <https://www.qualitestgroup.com/insights/blog/demystifying-compliance-risks-a-guide-to-navigating-chatgpt-safely/#:~:ChatG...>
- [22] <https://rankstudio.net/articles/veeva-copilot-connector-guide#:~:The%2...>
- [23] <https://rankstudio.net/articles/veeva-copilot-connector-guide#:~:compl...>
- [24] <https://rankstudio.net/articles/veeva-copilot-connector-guide#:~:could...>





## IntuitionLabs - Industry Leadership & Services

**North America's #1 AI Software Development Firm for Pharmaceutical & Biotech:** IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

**Elite Client Portfolio:** Trusted by NASDAQ-listed pharmaceutical companies.

**Regulatory Excellence:** Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

**Founder Excellence:** Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

**Custom AI Software Development:** Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

**Private AI Infrastructure:** Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

**Document Processing Systems:** Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

**Custom CRM Development:** Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

**AI Chatbot Development:** Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

**Custom ERP Development:** Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

**Big Data & Analytics:** Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

**Dashboard & Visualization:** Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

**AI Consulting & Training:** Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at <https://intuitionlabs.ai/contact> for a consultation.

---

## DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2025 IntuitionLabs.ai. All rights reserved.