Same Model, Different Contract

An overview for teams choosing an AI assistant for content and coding work in industries that answer to SOC 2, GDPR, and HIPAA.

CTRL ALT Elite · Fortune 100-caliber infrastructure, operations, and security leadership

By René López, Founder and Principal · ORCID https://orcid.org/0009-0008-4993-0088

> **WHO THIS IS FOR**
>
> *Teams in healthcare, financial services, insurance, legal, and any business that sells to them, who are deciding which AI assistant to standardize on for writing, analysis, and software development. It covers ChatGPT, Claude, Microsoft Copilot, Gemini, Kimi, and the largest open-weight models you can run yourself, and it is written for the people who will have to answer an auditor, a regulator, or a customer's security questionnaire about the choice. If your organization produces a SOC 2 report, signs GDPR processing terms, or holds a HIPAA business associate agreement, this is written for you.*

Start with the contract, not the model

Most teams researching an AI assistant start with capability: which model writes the cleanest prose, which one fixes the failing test, which one handles a 200-page contract. Those are the right questions eventually. In a regulated business they are the second questions, because the consumer version of every major assistant fails the regulated-data test and the enterprise version of every major assistant can pass it, and the model underneath is the same in both.

ChatGPT Free and ChatGPT Enterprise run the same models. Consumer Claude and Claude Enterprise run the same models. Copilot with a personal account and Copilot with a work account are the same product at the same address. What changes between tiers is not the technology. It is the contract, the default settings for training and retention, and who is on the other end of the phone when something goes wrong. The right comparison is therefore not vendor against vendor. It is tier against tier, within whichever vendors your content and engineering teams actually want to use.

> *You are not choosing a model. You are choosing which agreement governs what happens to your data after you press Enter.*

This overview does three things. It explains what SOC 2, GDPR, and HIPAA actually require of an AI vendor, which is narrower and more specific than the marketing suggests. It walks each major product line tier by tier, including the hosted Kimi products and the largest open-weight models, using the vendors' own published terms as of October 2026. And it lays out the part no vendor's paperwork can cover for you: the controls that remain your responsibility after the contract is signed.

What the three labels actually mean

Each of the three frameworks answers a different question, and each puts the burden in a different place. Treating them as three flavors of the same checkbox is the first mistake most buyers make.

SOC 2 is a report about the vendor, not a certificate for you

A SOC 2 Type II report is an independent auditor's opinion that a service organization's controls operated effectively over a period, typically six to twelve months, against the AICPA Trust Services Criteria for security and, optionally, availability, confidentiality, processing integrity, and privacy. It is not a certification and it is not a pass or fail. It is a

document you read. Two things matter for AI buyers. First, the report has a scope: it names the products and infrastructure covered, and a vendor's report may cover the enterprise product and not the consumer one. Second, the report proves the vendor's controls, not yours. Your own SOC 2 auditor will still ask how you manage access to the tool, what data you allow into it, and how you monitor use. The vendor's report is one input to your vendor-management control, nothing more.

GDPR is about roles and a signed processing agreement

Under the General Data Protection Regulation, the organization that decides why and how personal data is processed is the controller, and a vendor that processes it on the controller's instructions is a processor. Article 28 requires a written contract between the two, the data processing agreement or addendum, with specific mandatory terms: processing only on documented instructions, confidentiality, security measures, sub-processor controls, assistance with data-subject rights, deletion or return at the end of service, and audit rights. If personal data leaves the European Economic Area, a transfer mechanism such as Standard Contractual Clauses or an adequacy decision is also required. The practical test is simple: if a vendor will not sign a DPA that names the product you are using, that product is not fit for EU personal data. Consumer terms of service are not a DPA, and the vendor in that arrangement is typically a controller in its own right, using your data for its own purposes.

HIPAA requires a business associate agreement before PHI moves

The HIPAA Privacy and Security Rules permit a covered entity to share protected health information with a vendor only if that vendor is a business associate under a written agreement that satisfies 45 CFR 164.504(e). The BAA obligates the vendor to safeguard PHI, restrict its use, report breaches, flow the terms down to subcontractors, and return or destroy PHI at the end of the relationship. There is no such thing as a HIPAA-certified product; the Department of Health and Human Services does not certify anything. What exists is a vendor that will sign a BAA covering a specific, named service, configured a specific way. Sending PHI to any AI product that is not covered by a signed BAA is a disclosure without authorization, whether or not the vendor's security is excellent.

> **THE PHRASE TO WATCH FOR**
>
> *Every vendor in this paper uses some version of the words "supports your compliance." That phrasing is accurate and deliberate. A vendor can supply a SOC 2 report, sign a DPA, and sign a BAA. It cannot make your organization compliant, because most of the obligations under all three frameworks sit on your side of the contract.*

The one rule that predicts almost everything

Read the vendors' own terms side by side and one pattern explains almost every difference. Consumer tiers are governed by a privacy policy the vendor writes for its own benefit. Commercial tiers are governed by a customer agreement that the customer can negotiate, attach a DPA and BAA to, and hold the vendor to. Everything else follows from that.

On the consumer side, three things are true across the industry as of this writing. Conversations are used to improve the vendor's models unless the user finds and changes a setting. Some conversations are read by human reviewers, and reviewed conversations may be retained for years regardless of the user's deletion. And the vendor's own help pages tell users not to enter confidential information. That last point is worth pausing on: the vendor is telling you, in writing, that the product is not built for regulated data.

On the commercial side, the same three things invert. Customer content is excluded from training by default under the commercial terms. Human access is restricted to incident response, legal requirement, or the customer's explicit request. Retention is customer-controlled, with deleted content purged on a stated schedule, and zero-retention configurations are available for API traffic. A DPA is standard, and a BAA is available on named services.

> *The consumer product is the vendor's product. The commercial product is yours. The difference is written into the terms, and it is the whole difference.*

The product lines, tier by tier

The table below reflects each vendor's published terms and trust documentation as of October 2026. Vendors change these terms several times a year; treat this as a map of where to look, and verify against the primary source before signing. "Training" means whether prompts and outputs are used to train the vendor's models by default.

FIGURE 1: WHAT EACH TIER ACTUALLY PROVIDES

Product and tier	Training default	SOC 2 report	GDPR DPA	HIPAA BAA
ChatGPT Free, Plus, Pro	On, unless the user turns off model improvement or uses temporary chats	Not in scope	No; consumer terms	No
ChatGPT Business	Off by default under commercial terms	In scope (Type II)	Yes	No; BAA is offered on Healthcare, Enterprise with Regulated Workspace, Clinicians, and API with Modified Retention
ChatGPT Enterprise, Edu, Healthcare	Off by default	In scope (Type II)	Yes	Yes, on ChatGPT for Healthcare and Enterprise with Regulated Workspace; standard Enterprise is not on the eligible list
OpenAI API platform	Off by default; 30-day abuse retention, zero retention available	In scope (Type II)	Yes	Yes, on API with Modified Retention; Codex in the cloud excluded
Claude Free, Pro, Max	On by default since late 2025 unless the user opts out; opted-in data retained up to five years	Not in scope	No; consumer terms	No
Claude Team	Off under commercial terms	In scope (Type II)	Yes	No; BAA excludes Team
Claude Enterprise and first-party API	Off under commercial terms; zero retention available	In scope (Type II)	Yes	Yes, on Enterprise once the Primary Owner activates HIPAA compliance and accepts the BAA, and on the first-party API; Batch, Files, Code Execution, Computer Use, Web Fetch, connectors, and MCP servers excluded
Copilot, personal Microsoft account	Used for training unless turned off in settings	Not in scope	No; consumer terms	No
Copilot Chat, work or school account	Off; enterprise data protection applies on sign-in	In scope under Microsoft 365	Yes, Microsoft DPA	Yes, under the Microsoft 365 BAA on eligible plans; web search queries excluded
Microsoft 365 Copilot (licensed)	Off; processed inside the tenant boundary	In scope under Microsoft 365	Yes	Yes, on eligible plans and properly configured tenants
Gemini app, personal Google account	On by default; human reviewers see a sample; reviewed chats kept up to three years	Not in scope	No; consumer terms	No

Product and tier	Training default	SOC 2 report	GDPR DPA	HIPAA BAA
Gemini in Workspace, business and enterprise editions	Off; not used to train without customer instruction	In scope under Workspace	Yes, Workspace DPA	Yes, on covered plans as HIPAA Included Functionality; some Gemini surfaces excluded
Gemini via Vertex AI on Google Cloud	Off under Cloud terms	In scope under Google Cloud	Yes, Cloud DPA	Yes, under the Google Cloud BAA
Kimi app (kimi.com), consumer	On by default after de-identification; opt-out only by request to support, account-wide, no per-conversation option; data stored in mainland China	None published	No; consumer terms	No
Kimi OpenPlatform API (Moonshot AI, Singapore)	On by default per the platform policy; no in-product opt-out; enterprise restrictions by contact only; data stored in Singapore	None published	Not offered publicly	No
Open-weight model, self-hosted (DeepSeek, Qwen, Kimi K3, GLM, Llama, gpt-oss, Mistral)	None; no vendor ever sees a prompt	Your own report must cover the platform	Not applicable to the model; DPA needed only with your hosting provider	Not applicable to the model; BAA needed only with your hosting provider
Open-weight model via a hyperscaler (Bedrock, Vertex AI, Azure AI Foundry)	Off under the cloud's terms	In scope under the cloud's report	Yes, the cloud DPA	Yes, under the cloud BAA where the model is a covered service

Sources: OpenAI enterprise privacy and business data pages; Anthropic consumer and commercial terms and privacy documentation; Microsoft Learn enterprise data protection documentation; Google Workspace generative AI privacy hub and HIPAA Included Functionality; Moonshot AI Kimi user privacy agreement and Kimi OpenPlatform privacy policy; AWS, Google Cloud, and Microsoft BAA scope lists. Verified October 2026. Vendor terms change frequently.

A few details in that table deserve their own line, because they are the ones that get organizations in trouble.

- **Sign-in decides the protection, not the URL.** Microsoft's Copilot Chat applies enterprise data protection when a user is signed in with a work account and consumer terms when signed in with a personal one, at the same address. The same is true of Gemini in a browser. An employee who opens the tool from a personal account has stepped outside every agreement your organization signed.
- **Consumer Claude changed sides in 2025.** Older comparisons still describe Claude's consumer tier as opt-in for training. That was true until late 2025. Under current consumer terms, conversations on Free, Pro, and Max plans are used for training unless the user opts out, and opted-in data can be retained for up to five years. Commercial plans were explicitly excluded from that change.
- **Team is not Enterprise.** Both OpenAI and Anthropic offer a mid-tier plan that excludes training and carries SOC 2 scope but does not come with a BAA. Anthropic's BAA covers HIPAA-ready Enterprise and the API, not Team. OpenAI's BAA is offered on ChatGPT for Healthcare, Enterprise with Regulated Workspace, ChatGPT for Clinicians, and the API with Modified Retention; standard Enterprise and Business are not on the eligible list. If you handle PHI, the mid-tier is not enough.
- **BAAs have carve-outs inside the product.** Microsoft's BAA does not cover web search queries sent from Copilot. OpenAI's excludes Codex in the cloud and local Codex execution on the developer's machine. Anthropic's excludes Batch, Files, Code Execution, Computer Use, Web Fetch, connectors, MCP servers, and beta products, and covers Claude Code only on specific paths, described in the coding note below. Google's HIPAA Included

Functionality does not extend to every Gemini surface, and Google's own help pages list Gemini in Chrome and Gemini Notebook as outside the BAA. A signed BAA is a list of named services, and the list is shorter than the product.

- **Kimi has the model and not the paperwork.** Moonshot's models are among the strongest available, which is why they show up in enterprises. But the consumer app's own data-usage page states that inputs and replies may be used for model training after de-identification, that opting out is by request to customer service, account-wide, with no per-conversation option, and that data is stored on servers in mainland China. The international API platform is run from Singapore and says plainly that submitted content is used to train and refine its models, with no in-product opt-out as of this writing; training restrictions are available only by contacting the company. No SOC 2 report is published and no BAA is offered. For regulated data, the hosted Kimi products do not clear the bar. The open-weight Kimi models, run on your own infrastructure, are a different matter, covered below.
- **Adjacent products do not inherit coverage.** GitHub Copilot is not covered by the Microsoft 365 BAA because Microsoft owns GitHub. Google AI Studio is not covered because Vertex AI is. Each product has its own scope, and the assumption of inheritance is a common audit finding.

The other column: open-weight models you run yourself

The option regulated industries ask about most is the one hosted-product comparisons usually leave out: downloading the model and running it inside your own perimeter. That option has changed character in the last eighteen months. The largest open-weight models are now within a few points of the closed frontier on most benchmarks, and several are released under MIT or Apache 2.0 terms.

As of October 2026 the field at the top is led by Moonshot's Kimi K3, at roughly 2.8 trillion parameters; Alibaba's Qwen3.8-Max, a 2.4 trillion parameter sparse model; DeepSeek V4-Pro at about 1.6 trillion; and Z.ai's GLM-5.3 at roughly 750 billion. Meta's Llama 4 family, OpenAI's gpt-oss 120B, and Mistral Large 3 round out the credible enterprise options, and smaller checkpoints such as Qwen3.8-27B, gpt-oss 20B, and Mistral Small 4 bring near-frontier quality to one or two GPUs. "Open weight" means the weights are downloadable. It does not mean the training data is published, and it does not mean the license is unconditional.

Self-hosting inverts every row of the compliance table. Nothing leaves your network, so training on your data is not a question, a DPA with the model's author is unnecessary because they never process anything, and no BAA is required for the model itself because there is no business associate. What you have instead is total responsibility. Your own SOC 2 report must now cover the inference platform. Your GDPR obligations run to the infrastructure provider, if any, not the lab. Your HIPAA obligations are the same as for any internal system holding PHI: access control, audit, encryption, minimization, and a risk analysis. The vendor's checkmarks are gone because the vendor is gone.

FIGURE 2: WHO CARRIES WHAT UNDER EACH DEPLOYMENT MODEL

Obligation	Closed model, enterprise tier	Open-weight, hyperscaler-hosted	Open-weight, self-hosted
Keeps prompts out of training	Vendor, by contract	Cloud provider, by contract	You, by architecture
SOC 2 evidence for the AI platform	Vendor's report, scoped to the product	Cloud provider's report, scoped to the service	Your own report; the platform is in your scope
GDPR processing terms	Vendor DPA	Cloud DPA	Only with whoever hosts the hardware; none if on-premises
HIPAA business associate	Vendor BAA on named services	Cloud BAA where the model service is covered	None required; you are the covered entity's own system
Patching, model updates, incident response	Vendor	Shared with the cloud provider	You, entirely

Obligation	Closed model, enterprise tier	Open-weight, hyperscaler-hosted	Open-weight, self-hosted
License exposure	Commercial subscription	Passed through by the cloud	Yours to read: Llama's user cap and EU multimodal exclusion, Qwen3.8-Max revenue-sharing terms, Kimi K3's conditions for large commercial products, GLM-5.3's conditional license; DeepSeek, gpt-oss, Mistral, and Qwen3.8-27B are MIT or Apache 2.0

The middle column is where most regulated organizations land: an open-weight model served inside AWS Bedrock, Google Vertex AI, or Azure AI Foundry inherits the hyperscaler's SOC 2 scope, DPA, and BAA, and the model's author is never a party to your data.

Three practical warnings for anyone taking the self-hosted path.

- **Trillion-parameter models need a cluster.** The models at the top of the leaderboards require multi-GPU H200 or B200 class hardware to serve at any useful speed. The realistic on-premises choices for most organizations are the distilled and mid-size checkpoints, or the largest models rented through a hyperscaler under that hyperscaler's contracts.
- **The lab's hosted API is not the same as the lab's weights.** DeepSeek's API stores data in China under laws that compel cooperation with state intelligence work; DeepSeek's weights on your servers store nothing anywhere. Kimi's API trains on your content by default; Kimi K3 on your hardware cannot. When someone in your organization says "we use DeepSeek," the first question is which one.
- **Read the license the way you read a contract.** Several of the strongest open-weight models carry conditions on user counts, geography, revenue, or resale of inference. The licenses change per checkpoint, not per vendor. The model that is fine for an internal tool may not be fine for a product you sell, and your customers' security teams will ask.

Which tier your data actually requires

The tables above answer "does this product have the paperwork." The question a team has to answer is "which paperwork does our data need." Working backward from the data classification gives a cleaner rule than working forward from the product list, and it settles most content-versus-coding debates quickly: source code, customer records, and draft contracts are all confidential business data at minimum, and the tier they require is the same.

FIGURE 3: MINIMUM TIER BY DATA CLASS

If the tool will see	Minimum acceptable arrangement	Why
Public information, generic drafting, code with no secrets	Any tier, with a policy that nothing confidential is entered	No regulated obligation attaches; the risk is habit, not the data.
Internal confidential business data (plans, pricing, contracts, source code)	Commercial tier with training excluded and a DPA; SSO-enforced accounts	Your NDAs, trade-secret position, and customer contracts assume the data is not being used to train a public model or read by a vendor's reviewers.
Personal data of EU or UK residents (customers, employees, prospects)	Commercial tier with a signed DPA naming the product, plus a transfer mechanism and data residency where offered	GDPR Article 28 requires the contract; without it you cannot lawfully engage the processor.

If the tool will see	Minimum acceptable arrangement	Why
Protected health information	A tier and service explicitly named in a signed BAA, configured per the vendor's HIPAA guidance, with excluded features disabled; or an open-weight model inside your own HIPAA-compliant environment	HIPAA prohibits the disclosure without the BAA regardless of the vendor's security posture. A self-hosted model removes the business associate but not the Security Rule.
Guest profiles, loyalty and reservation records, and casino or gaming data (player tracking, cashless wagering, surveillance)	Commercial tier with training excluded and a DPA for guest records; cardholder data never enters a prompt (tokenized or redacted upstream); gaming system data stays inside the regulated environment unless the state gaming regulator's rules and your internal controls expressly allow the tool	Hospitality and gaming operators carry PCI DSS for cards, state privacy law for guest records, and gaming commission rules for player and surveillance data. None of the three frameworks in this paper covers those, and a vendor's SOC 2 report does not satisfy a gaming regulator.
Cardholder data, government CUI, or other data under a specific regime	Do not use general-purpose assistants; use the vendor's government or dedicated cloud offering, or a tokenized pipeline that never exposes the raw data	PCI DSS, CMMC, and similar regimes are not addressed by any of the three frameworks in this paper.

A NOTE ON CODING ASSISTANTS

Engineering teams often evaluate coding tools on a separate track from content tools, and the paperwork frequently gets missed on that track. Three points. Source code is confidential business data and usually contains secrets, customer identifiers, and sometimes PHI in test fixtures, so a coding assistant needs at least the commercial tier with training excluded and a DPA. BAA scope is narrower for coding surfaces than for chat: GitHub Copilot is not covered by the Microsoft 365 BAA; Anthropic's BAA covers Claude Code only in the terminal or Claude Desktop with the HIPAA configuration applied and a direct connection to Anthropic, or with a Console key under zero data retention, and excludes the VS Code and JetBrains extensions, gateway and cloud-provider paths, and Claude Code on web and mobile; OpenAI's BAA excludes Codex in the cloud and local Codex execution on the developer's machine; and agentic coding tools that run commands on a developer's machine sit outside most vendors' HIPAA-ready configurations. And an API key is not a contract: a developer with a personal API key on the consumer platform is outside every agreement the company signed. The cleanest path for regulated engineering teams is a coding tool that authenticates through the organization's commercial tenant or a hyperscaler endpoint, with the same SSO, logging, and secret-scanning controls applied as anywhere else.

What the contract does not do for you

A signed DPA and BAA cover the vendor's processing. They do not cover the path the data took to get there, who was allowed to send it, or what your own systems did with the answer. Under every one of these frameworks, the controls below remain the customer's responsibility, and they are where audits actually fail.

- **Identity.** Enforce single sign-on into the commercial tenant and block consumer endpoints for the same products at the network or identity layer. The most common real-world failure is not a vendor breach; it is an employee with a personal account and a work problem.
- **Permissions.** Assistants that read your environment inherit its sharing mistakes. Microsoft's own documentation says Copilot supports HIPAA compliance for properly configured implementations, and configuration mostly means fixing over-permissive SharePoint and OneDrive sharing that has been wrong for years. Run the permissions review before the rollout, not after the first surprising answer.
- **Data minimization.** A BAA authorizes disclosure of PHI the workforce member was already permitted to disclose for that task, not everything on the screen. Redaction and DLP in front of the assistant is the one control that does not depend on any vendor's policy staying the same.

- **Retention and audit.** Set retention to the shortest period your records obligations allow, enable the vendor's audit logs and compliance APIs, and keep the evidence somewhere your auditor can query. The Security Rule's audit-control requirement is yours, not the vendor's.
- **Workforce training.** The single sentence that prevents most incidents: if you would not put it in an email to an outside party, do not put it in an assistant that is not on the approved list. Make the approved list short and visible.
- **Vendor monitoring.** Terms in this space have changed materially at least once a year. Assign an owner to re-read the training policy, the BAA scope, and the SOC 2 report period annually, and to catch the day a vendor moves a product from one column to the other.

> **THE PATTERN**
>
> *Every enforcement action and audit finding we have seen in this area has the same shape. The vendor's paperwork was fine. The organization's controls around the tool were not: a personal account, a folder shared to everyone, PHI in a prompt nobody was supposed to write, and no log to show what happened. The contract was never the weak point.*

A thirty-day path for a regulated organization

This does not require a transformation program. It requires four decisions and a few weeks of configuration.

Week 1: Inventory and classify. Find every AI assistant in use, including personal accounts on work devices, and map each to the data classes above. Most organizations discover three to five times more tools than IT knew about.

Week 2: Pick one commercial tier per vendor and sign the paperwork. Standardize on the tier your highest data class requires. Execute the DPA. If PHI is in scope, execute the BAA and enable the vendor's HIPAA-ready configuration, which typically disables excluded features automatically.

Week 3: Enforce the path. SSO into the commercial tenant, consumer endpoints blocked, retention set, audit logging on, DLP or redaction in front of the tools that will see regulated data, and a permissions review of anything the assistant can read.

Week 4: Publish the list and train to it. One page: which tools are approved, for which data, from which account. Then a quarterly review with a named owner, because the vendors will move again.

The point

The question teams start with, which AI tool protects our data, has an answer that is less about the vendor than the buyer expects. The four major Western providers will exclude your data from training, supply a SOC 2 report, sign a DPA, and sign a BAA, if you are on the right tier, using the named service, signed in the right way. None of them will do any of that for the free product your staff are already using. Kimi's hosted products offer none of it yet, however good the model. And the open-weight route offers all of it and none of it at once: no vendor to trust, and no vendor to share the burden.

The contract is the easy part. It takes a procurement cycle and a signature. The controls around it, identity, permissions, minimization, logging, and training, are the part that determines whether the contract means anything when an auditor asks. That is the same work regulated industries have always done for every new system that touches their data. AI did not change the work. It changed how fast the tool arrived, and how many people adopted it before anyone asked the question.

Work with CTRL ALT Elite

CTRL ALT Elite delivers Fortune 100-caliber technology and security leadership on a fractional basis, as operators, not advisors. We have run infrastructure and audit-response functions for regulated data at airline, cruise line, and healthcare-platform scale, and we have sat on the buyer's side of the SOC 2, GDPR, and HIPAA conversations with every vendor named here. We help boards and operators choose the right tier, execute the right paperwork, and build the controls around it that actually pass an audit.

If your staff are already using AI tools and you cannot yet say which ones, with which data, under which agreement, that is the conversation we have.

CTRL ALT Elite

reboot@ctrlaltelite.io · 954-516-6909 · ctrlaltelite.io · Pompano Beach, FL
René López, Founder and Principal · ORCID https://orcid.org/0009-0008-4993-0088

Disrupt · Innovate · Elevate

© CTRL ALT Elite. Shareable in full with attribution. Vendor-neutral; no product is endorsed. Vendor terms, scopes, and defaults described here are as published in October 2026 and change frequently; verify against each vendor's current terms, trust portal, and BAA before relying on them. This paper is not legal advice.

Sources: OpenAI, Enterprise privacy (openai.com/enterprise-privacy) and Business data privacy, security, and compliance (openai.com/business-data); OpenAI Trust Portal (trust.openai.com); OpenAI Help Center, HIPAA-eligible products and functionality. Anthropic, Consumer Terms, Commercial Terms, Privacy Policy, and HIPAA-ready services documentation, including the Privacy Center article on protected health information (anthropic.com and privacy.claude.com). Microsoft Learn, Enterprise data protection in Microsoft Copilot and Microsoft Copilot Chat, and Microsoft Copilot Chat privacy and protections; Microsoft Online Services BAA in-scope services list. Google Workspace Help, Generative AI in Google Workspace Privacy Hub, HIPAA Included Functionality, and Gemini Apps Privacy Hub. Moonshot AI, Kimi User Privacy Agreement (kimi.com) and Kimi OpenPlatform Privacy Policy (platform.kimi.ai), and Kimi Help, Data Usage and Sharing (kimi.com). Model cards and licenses for Kimi K3, Qwen3.8, DeepSeek V4, GLM-5.3, Llama 4, gpt-oss, and Mistral on Hugging Face and the labs' repositories; AWS, Google Cloud, and Microsoft HIPAA eligible-services lists. AICPA Trust Services Criteria; Regulation (EU) 2016/679 Article 28 and Chapter V; 45 CFR 164.502(e) and 164.504(e).