GovRAMP

Version 1.0

Asking Your Providers About AI

August 2026

A practical guide for government members evaluating generative AI capabilities in cloud products.

Asking Your Providers About AI

Why This Matters

AI is increasingly built into the cloud products government already uses—often added through routine updates, and not always with clear visibility. GovRAMP's AI Disclosure flags on the Authorized and Progressing Product Lists when AI is part of a product.

This guide is the companion conversation: the questions to ask once AI is in the picture. The aim is transparency and a shared understanding of who is responsible for what—not a checklist, and not a guarantee of security. Use the questions that fit your data, your use case, and your risk tolerance.

When to Ask

- **Lead with data.** Settle how your data is handled before evaluating any feature.
- **Early and often.** Raise the questions in procurement (RFI / RFP / RFQ) and revisit them across the relationship—don't wait for a major release.
- **Make it a partnership.** A direct question gives a busy provider the prompt to disclose newly added AI.

Who to Bring

An AI lead. Name a "quarterback" to own the conversation, supported by:

- The business / data owner closest to the system
- Security (CISO), Privacy (CPO), Legal, and Procurement
- The provider's AI product owner—an engineer who knows how it works, not a sales contact

Questions to Ask

System & Provenance

- Is AI part of this product today, and which features use it?
  - If not, is it in your product development roadmap?
- Is the model built in-house or supplied by a third party—and which one?
- Where is the model hosted and operated, and where does our data flow?
- What is the model's geographic origin and supply chain?
  - An air-gapped open-source model may still derive from a foreign base model.

Asking Your Providers About AI

Questions to Ask

Customer Data Handling

- Is our data processed by AI?
  - Which types—public, internal, confidential, or regulated?
- Is our data used to train, tune, or improve the model?
- Is opting out the default—and can we opt out of AI use and of data-for-training separately?
- Beyond training, how long is our data retained as context, and how is it protected?
  - Context data is often the most exposed if a model is misused.
- Who is responsible for protecting our data at each point: our agency, your organization, or a third-party AI or model provider?
- Which controls are you responsible for operating, and which controls must our agency configure or manage?
- What protections, such as encryption, access control, or redaction, apply, and how can we independently validate that they are enforced?
- Can we request deletion or expungement of our data?

Feature Activation & Controls

- Is AI on by default, or opt-in? (Leading practice: opt-in, after administrators are trained.)
- Can administrators enable AI for only certain users, or is it on or off for all?
- Can administrators disable AI features without losing core functionality?
  - Is any AI capability mandatory?
- Do end users receive advance notice when they are interacting with AI?

Agency & Permissions

- What can the AI actually do—read-only summarization, or execute actions, change configurations, and make decisions?
- Does it require human confirmation before consequential actions?
  - Broader permissions invite broader risk.

Transparency, Logging and Change Control

- Do administrators get audit logs and visibility when AI is used?
- Can we stay on a model version or delay updates?
  - Are rollback and version-control mechanisms available?
- How and when will you notify us before a model changes?
  - What is the threshold and the notice period?

Asking Your Providers About AI

Questions to Ask

Risks, Guidance and Disclosure

- What known risks or limitations should we weigh before enabling or relying on this functionality?
- What documentation supports our risk / benefit assessment?
- Has this product's AI been noted via GovRAMP's AI Disclosure?
  - What AI documentation have you uploaded to the PMO portal for us to review?

What Strong Answers Look Like

- AI is opt-in rather than on by default; both "use AI" and "use our data to train" default to off.
- Administrators can disable AI, manage it by use case, and stay on or roll back a model version.
- Claims about data use are verifiable—not only stated.
- Clear documentation exists, and the provider can name the product's risks and limitations candidly.

Documentation to Request *(in place of, or alongside, a questionnaire)*

Rather than ask providers to complete another form,
request the artifacts they may already maintain:

- ISO/IEC 42001 certificate
- NIST AI Risk Management Framework (AI RMF) alignment documentation
- Vendor-published model card
- GovAI Vendor Template (Templates & Resources | City of San José)

With Thanks to the AI Security Task Force

This guide was developed with input from the **GovRAMP AI Security Task Force**. We thank Task Force members for contributing their time, expertise, and perspectives to help shape practical guidance for more informed conversations about AI in cloud products.

GOVRAMP.ORG | INFO@GOVRAMP.ORG | LINKEDIN.COM/COMPANY/GOVRAMP