Shadow AI Ready

AI DILIGENCE PACK

Example SaaS Co.

Answers for the AI section of your security questionnaire

Prepared **25 September 2026** for **an enterprise retailer's vendor security questionnaire**
Product: **Example Product** · Pack contact: **Head of Security, security@example-saas.example**

SAMPLE. Example SaaS Co., Example Product, every feature named, and every date and figure in this pack are fictional and illustrative. No real people are named; owners are shown by role. This is not a real customer's pack.

AI diligence score

Self-reported inventory, scored 24 September 2026.

Low | **Med** | High

Score

66 /100

Medium procurement risk. Buyer likely asks follow-ups.

Top gaps from the score

1. SSO and DLP on consumer LLM accounts
2. Written acceptable-use policy for generative AI
3. Vendor risk review for AI subprocessors

Inventory

7 AI tools listed: 4 approved, 2 in review, 1 exception. Full inventory in section 01. Each gap is addressed in section 05.

Not a certification. Not legal advice. Not a GRC platform.

This pack drafts answers from information Example SaaS Co. reported to Shadow AI Ready on 24 September 2026. It has not been independently verified. Example SaaS Co. is responsible for reviewing every answer and confirming it is accurate before sending it to a customer. The score is a rule-based internal risk view, not an attestation of any kind.

CONTENTS

What is in this pack

Everything here is written to be pasted or forwarded. Adjust wording to match the exact question the buyer asked.

How to use it this cycle

1. **Answer the questionnaire.** Find the closest match in section 02 and paste the answer. Question IDs (A01, B03) are ours, not the buyer's.
2. **Attach the inventory and the FAQ** when the buyer asks for a list of AI tools or sub-processors, or a summary for their security team.
3. **Send the exceptions page** when the buyer asks about gaps or open items. Naming a gap with a date usually lands better than a vague yes.
4. **Adopt the policy stubs** only after your own review. They are a starting draft, not a finished policy.
5. **Keep the score card internal** if you prefer. A buyer copy without it is available on request.

Scope and basis

Answers describe Example SaaS Co. as of 25 September 2026, based on the intake submitted 24 September 2026 and any corrections sent since. Where the intake did not cover a point, the answer states what Example SaaS Co. needs to confirm. Nothing in this pack is a legal opinion or a statement of conformance with any law, standard, or framework.

01

AI tool inventory

All AI tools used in Example Product and by Example SaaS Co. staff, as of 25 September 2026. Owner is the person accountable for the tool's settings and contract.

Tool	Vendor	Where used	Data touched	Customer data	Training opt-out	DPA / sub-processor	Owner	Status
Azure OpenAI Service (GPT-4o)	Microsoft	Product	Forecast figures, SKU and location IDs, user questions	Yes	Not used for training (enterprise terms)	DPA signed; listed sub-processor	VP Engineering	Approved
ChatGPT Enterprise	OpenAI	Internal	Drafts, internal docs, research	Not permitted	Not used for training (enterprise terms)	DPA signed; not a sub-processor	IT Lead	Approved
GitHub Copilot Business	GitHub (Microsoft)	Internal (engineering)	Source code open in the editor	Not permitted	Not used for training (business terms); public code filter on	DPA signed; not a sub-processor	VP Engineering	Approved
Zoom AI Companion	Zoom	Internal	Internal meeting transcripts and summaries	No (off for external meetings)	Not used for training (vendor terms)	DPA signed; not a sub-processor	IT Lead	Approved
Gong	Gong	Internal (sales)	Recorded sales calls, prospect contact details	No product data	Training opt-out set in admin	DPA signed; review in progress (X3)	Sales Ops Lead	Review
Notion AI	Notion	Internal	Internal wiki pages	Not permitted	To confirm in review (X3)	DPA under review (X3)	IT Lead	Review
ChatGPT Plus (3 personal accounts)	OpenAI	Internal (marketing, sales)	Marketing copy and email drafts	Not permitted	Not controlled by company (personal plan)	None (personal terms) (X1)	IT Lead	Exception

Where used: Product means the tool processes customer data inside Example Product. Internal means staff productivity only. **Status:** Approved (in use under company terms), Review (being evaluated, restricted use), Exception (known gap, see section 05). **Sub-processor list:** example-saas.example/subprocessors.

02

Draft questionnaire answers

Written as Example SaaS Co., ready to paste. Where a buyer's question combines two of these, join the answers and trim.

A. AI IN THE PRODUCT

A01 **Does your product use artificial intelligence or machine learning? Describe the features.**

Yes. Example Product uses AI in three features: Example Feature A (plain-language summaries of why a demand forecast changed), Example Feature B (short descriptions of unusual shipment or inventory patterns), and Example Feature C (answers questions about the customer's own forecast data). These features are off until a customer admin enables them in workspace settings. AI output is presented to users as a suggestion that they can accept, edit, or discard. No AI feature makes decisions with legal or similarly significant effects on individuals.

A02 **Which third-party AI models or providers do you use, and how are they accessed?**

Example Product uses OpenAI models (GPT-4o) hosted in Microsoft's Azure OpenAI Service, accessed through Azure's enterprise API from our own Azure tenant, under our Microsoft agreement. We do not host or train our own foundation models. Requests are processed in the Azure East US 2 region. The full list of AI providers is in our AI tool inventory and sub-processor list (example-saas.example/subprocessors).

A03 **Is customer data used to train, fine-tune, or improve AI models, either yours or a third party's?**

No. Customer data is not used to train or fine-tune any model, ours or a provider's. Our agreement with Microsoft excludes customer inputs and outputs from model training, and we have not built any fine-tuned models on customer data. If this ever changes, we will give customers at least 30 days' notice and an opt-out before it takes effect.

A04 **What customer data is sent to AI models? Does it include personal or sensitive data?**

When a user invokes an AI feature, we send only the data needed for that request: aggregated forecast figures, SKU and location identifiers, and the text of the user's question. Forecast data is operational and does not normally contain personal data. We do not send user names, email addresses, or other account details with AI requests. We do not send authentication data, billing information, or raw uploaded files to any AI provider.

A05 **Do AI providers retain prompts or outputs? For how long?**

Microsoft processes each request to generate a response. Under our Azure OpenAI configuration, prompts and outputs may be stored for up to 30 days solely for abuse monitoring, then deleted. They are not used for training. Within Example Product, AI inputs and outputs are stored with the forecast they relate to in the customer's workspace and deleted with the customer's account data under our standard retention schedule.

A06 **Where is data processed by AI features stored and processed?**

AI requests are processed in the United States (Azure East US 2). Customer data at rest remains in our primary hosting region in the United States. We do not currently offer EU processing for AI features. Customers that require it can keep AI features disabled.

A07 **How is customer data kept separate between tenants in AI features?**

Each AI request is built only from data in the requesting customer's workspace, using the same tenant-scoped access checks as the rest of Example Product. We do not create embeddings or vector indexes of customer data. No customer's data is used to generate output for another customer.

A08 **Can customers disable AI features or opt out?**

Yes. AI features are off by default. A workspace admin turns them on or off in Settings and can limit them to named user groups. When AI features are disabled, no data from that workspace is sent to an AI provider. Disabling AI features does not affect the rest of the product.

A09 **Is AI-generated content identified to users?**

Yes. AI-generated content in Example Product is labelled with an "AI generated" tag, and each explanation links to the forecast figures it was based on, so users can tell it apart from system-calculated values and human-written content.

B. GOVERNANCE AND OVERSIGHT

B01 **Who is accountable for AI risk at your company?**

Our VP Engineering is accountable for AI use in the product and internally. Changes to AI features are reviewed in a monthly product security review with our Head of Security.

B02 **Do you maintain an inventory of AI systems and tools?**

Yes. We maintain an inventory of every AI tool used in the product and by staff, recording vendor, where it is used, data touched, whether customer data is involved, training opt-out status, DPA status, and owner. It was last reviewed on 25 September 2026 and is attached.

B03 **Is there human oversight of AI outputs? Are any decisions fully automated?**

AI features produce explanations and answers that a user reads and decides whether to act on. Example Product never changes a forecast, order, or setting based only on AI output. No AI output is used to make automated decisions about individuals, such as decisions on employment, credit, or access to services.

B04 **How do you test AI features for accuracy and reliability?**

Before release, each AI feature is tested against a fixed set of 200 historical forecasts with known outcomes, and our data science team checks explanations for consistency with the underlying figures. Numbers shown to users always come from the forecast engine, not from the model. We monitor quality after release through weekly sampling of outputs and user feedback ratings. Users can flag an incorrect output from within the product, and flagged outputs are reviewed by the data science team within five business days.

B05 **How do you assess and reduce bias in AI outputs?**

Our AI features describe demand, shipment, and inventory data, not people, which limits exposure to bias about individuals. Before release, we review prompts and sample outputs for misleading or one-sided explanations, such as attributing a change to a single cause when the data shows several. We do not use protected characteristics as inputs to AI features.

B06 **Do you follow a recognised AI risk framework?**

We use the NIST AI Risk Management Framework (AI RMF 1.0) as a reference for how we identify and manage AI risk. We do not claim formal conformance with, or third-party attestation against, any AI standard. We track obligations under laws that apply to our customers, including the EU AI Act, GDPR, and US state privacy laws, and will update our practices as they take effect.

B07 **Will you notify customers before adding an AI sub-processor or materially changing AI features?**

Yes. We give at least 30 days' notice before adding a new sub-processor that will process customer data, through email to account admins and an update to our sub-processor page. Customers can object under the terms of our DPA. Material changes to how AI features use customer data are announced the same way.

C. SECURITY, LOGGING, AND INCIDENTS

C01 **How do you protect AI features against prompt injection, data leakage, and misuse?**

User input is passed to the model as data inside a fixed system prompt, never as instructions that can change what the model can access. The model has no tools that write data or call external systems, and its output is displayed as text only. AI features run with the permissions of the requesting user and cannot read data that user cannot access. Provider API keys are stored in Azure Key Vault and rotated every 90 days.

C02 **Are AI features included in your security testing?**

Yes. AI features go through the same secure development process as the rest of Example Product: code review, dependency scanning, and an annual third-party penetration test. The June 2026 test included the AI features and prompt injection cases.

C03 **Do you log AI feature usage? Who can access the logs, and how long are they kept?**

We log metadata for every AI request: workspace, user, feature, time, and token count. Prompt and output text is stored only as part of the customer's workspace data (see A05), not in operational logs. Access to these logs is limited to the engineering on-call team and security staff, through SSO with MFA. Logs are kept for 13 months.

C04 **How are AI-related incidents handled and communicated?**

AI-related incidents follow our incident response plan, which covers incorrect or harmful output at scale, unintended data exposure through an AI feature, and provider security incidents. We can disable AI features across all workspaces within minutes using a global feature flag. We notify affected customers of a security incident involving their data within 72 hours of confirming it.

C05 **How do you handle deletion requests for data processed by AI features?**

Data processed by AI features is deleted with the rest of the customer's data under our standard deletion process, within 30 days of a verified request or contract end. We do not create embeddings or fine-tuned models from customer data, so there is no derived AI data to delete separately. Copies kept by the provider for abuse monitoring expire within 30 days.

D. VENDORS AND SUB-PROCESSORS

D01 **Are your AI providers listed as sub-processors? Do you have DPAs in place?**

Yes. Microsoft (Azure OpenAI Service) processes customer data for AI features in Example Product and is listed as a sub-processor, under a signed DPA. It is the only AI sub-processor. AI tools used internally by staff do not receive customer data from Example Product, so they are not sub-processors. Our current sub-processor list is at example-saas.example/subprocessors.

D02 **How do you evaluate an AI vendor before adopting it?**

AI vendors that will process customer data go through our security review before contract: data use and training terms, retention, sub-processors, hosting region, security documentation, and a DPA. Microsoft (Azure OpenAI Service) was reviewed this way in March 2026. Two internal AI tools, Gong and Notion AI, were adopted by teams without a documented review. Those reviews are in progress, and we are applying one written checklist to every AI tool by 30 November 2026 (exceptions page, X3).

E. EMPLOYEE USE OF AI

E01 **Do you have a policy governing employee use of generative AI tools?**

Yes, in draft. Our AI acceptable-use policy was circulated to all staff as interim guidance on 23 September 2026. Formal adoption and staff acknowledgement are due by 15 October 2026 (exceptions page, X2). The policy lists approved tools, sets rules for what data may be entered into them, requires company-managed accounts, and makes each employee responsible for reviewing AI output they use.

E02 **Which AI tools are approved for staff, and how is unapproved use handled?**

Approved tools today: ChatGPT Enterprise, GitHub Copilot Business, and Zoom AI Companion, plus Azure OpenAI Service inside Example Product. Gong and Notion AI are in use with restrictions while their review completes. New tools are requested through an IT request ticket and reviewed before use. In our September 2026 inventory, three staff in marketing and sales reported using personal ChatGPT Plus accounts for drafting. They have been told not to enter any non-public data and are moving to ChatGPT Enterprise by 31 October 2026 (exceptions page, X1).

E03 **Can employees enter customer data into AI tools? What prevents it?**

No. Customer data from Example Product is not permitted in any AI tool used by staff. It is processed by AI only inside the product, through Azure OpenAI Service. Production data access is limited to on-call engineers through SSO with MFA. Approved AI tools are provisioned through Okta with company settings enforced. We do not yet run DLP on AI tools, and the personal accounts above sit outside SSO; closing that is exceptions item X1.

E04 **Do engineers use AI coding assistants? How do you protect source code and handle generated code?**

Yes. Engineers use GitHub Copilot Business, provisioned through Okta with company settings. Under the business terms our code is not used to train models, and the public code filter is enabled. No other coding assistant is approved. All code, including AI-assisted code, goes through the same pull request review and automated checks before merge. Engineers remain responsible for code they commit.

E05 **Do employees receive training on safe AI use?**

AI use is covered briefly in security onboarding today. A short module on the AI policy and its data rules will be added when the policy is formally adopted in October 2026, with an annual refresher after that.

03

Acceptable use of AI tools

Policy owner: VP Engineering · Version: 0.9 (interim) · Effective: 15 October 2026 (planned)

> Policy stubs. Starting text only. Review with your own counsel and leadership before adopting. Delete this note once adopted.

1. Purpose and scope

This policy sets how Example SaaS Co. employees, contractors, and interns may use AI tools for work. It covers chat assistants, coding assistants, meeting and writing assistants, and AI features inside other software. It applies on any device used for Example SaaS Co. work.

2. Approved tools

Only tools marked Approved in the AI tool inventory may be used with Example SaaS Co. data. Tools in Review may be used only within the limits noted on the exceptions page. Today these are: ChatGPT Enterprise, GitHub Copilot Business, and Zoom AI Companion, plus Azure OpenAI Service inside Example Product. Gong and Notion AI are in use with restrictions while their review completes. The inventory is maintained by the VP Engineering and is the source of truth.

3. What data may go into AI tools

Data class	Examples	Approved tools	Personal or unapproved tools
Public	Published docs, website copy	Allowed	Allowed
Internal	Plans, internal docs, non-sensitive code	Allowed	Not allowed
Confidential	Customer data, contracts, source code, financials	Customer data: only inside Example Product (Azure OpenAI). Other Confidential data: approved tools only	Never
Restricted	Credentials, keys, secrets, payment card data	Never	Never

4. Accounts

Use only company-managed accounts, signed in through Okta, for work. Do not use personal AI accounts for Example SaaS Co. work, even for Public data, so that company settings (training opt-out, retention) apply.

5. Coding assistants

Use only GitHub Copilot Business with the company's settings. Do not paste secrets, credentials, or customer data into prompts. Review AI-suggested code as you would a colleague's and do not commit code you do not understand.

6. You own the output

AI output can be wrong. You are responsible for checking anything AI produced before it goes to a customer, into the product, or into a decision. Do not present AI output as verified fact without checking it.

7. Customer-facing use

Do not use AI to make decisions about individuals (hiring, performance, eligibility). Label AI-generated content that customers see in the product. Customer-facing AI features follow the product review process, not this policy alone.

8. Requesting a new tool

Request new AI tools through an IT request ticket. The VP Engineering reviews the vendor's data use, training, retention, and contract terms and responds within 10 business days. Until approved, the tool may be used only with Public data.

9. Mistakes and reporting

If you enter data into an AI tool that this policy does not allow, report it to security@example-saas.example within 24 hours. Reporting quickly matters more than the mistake. Good-faith reports are not grounds for discipline.

10. Acknowledgement and review

Everyone covered by this policy acknowledges it at onboarding and once a year. The policy owner reviews it every six months and when a new AI tool is approved.

04

Buyer FAQ

AI at Example SaaS Co.: short answers for your security team. Details are in the attached questionnaire answers and AI tool inventory.

Does Example Product use AI?

Yes, in three optional features (Example Features A, B, and C): explanations of forecast changes, notes on unusual patterns, and answers to questions about your own forecast data. They are off until your admin turns them on.

Do you train models on our data?

No. Customer data is not used to train or fine-tune any model, ours or a provider's.

Which AI providers process our data?

One: Microsoft, through Azure OpenAI Service in the US East 2 region, under a DPA. Listed on our sub-processor page: example-saas.example/subprocessors.

How long are prompts and outputs kept?

Microsoft may keep prompts and outputs for up to 30 days for abuse monitoring only, then deletes them. In Example Product, AI output is stored with the forecast it explains and deleted with your data.

Can we turn AI features off?

Yes. AI features are off by default. Your admin controls them per workspace and per user group.

Is a person in the loop?

Yes. AI output is a suggestion a user reviews. Example Product does not make automated decisions about individuals.

How do your employees use AI?

Through company accounts with SSO: ChatGPT Enterprise, GitHub Copilot Business, and Zoom AI Companion, with Gong and Notion AI under review. Customer data from Example Product is not permitted in staff AI tools.

What happens if something goes wrong?

AI incidents follow our incident response plan. We can switch AI features off across all workspaces immediately and notify affected customers within 72 hours of confirming an incident involving their data.

Are there known gaps?

Yes, and we list them with target dates on the exceptions page. Currently three: retiring a few personal AI accounts, formally adopting our AI policy, and finishing documented reviews of two internal AI tools. None involves how Example Product processes your data.

Who do we contact with more questions?

Head of Security: security@example-saas.example. Written questions are answered within 3 business days.

05

Exceptions

> Where our AI controls are not yet where we want them, what is in place meanwhile, and when we expect to close each item. We would rather you read it here than find it later.

#	Gap	Current state	Interim control	Planned remediation	Owner	Target
X1	Personal AI accounts outside SSO (SSO and DLP on consumer LLM accounts)	Three staff in marketing and sales use personal ChatGPT Plus accounts for drafting. Company settings, SSO, and DLP do not apply to these accounts.	Written instruction (22 Sep 2026) not to enter any non-public data. ChatGPT Enterprise seats issued to all three.	Move all three users to ChatGPT Enterprise through Okta. Block consumer ChatGPT sign-in on managed browsers. Evaluate DLP for AI tools.	IT Lead	31 Oct 2026
X2	AI acceptable-use policy not formally adopted	Policy is written (section 03) but not yet approved by leadership or acknowledged by staff.	Data rules from section 03 sent to all staff as interim guidance (23 Sep 2026).	Leadership approval, then staff acknowledgement tracked in the HR system. Add to onboarding.	VP Engineering	15 Oct 2026
X3	No documented review for two internal AI tools (Vendor risk review for AI subprocessors)	Gong and Notion AI were adopted by teams without a documented security review. DPAs are signed (Gong) or under review (Notion).	Notion AI limited to Internal data. Gong limited to recorded sales calls. Neither receives Example Product customer data.	One written AI vendor checklist (data use, training, retention, sub-processors, DPA) applied to every tool in the inventory, and repeated at renewal.	Head of Security	30 Nov 2026

Reviewed 25 September 2026 by the VP Engineering. Items are removed from this page only when the remediation is complete, not when it starts. Questions: security@example-saas.example.