

GLIMMER

Security and governance overview

How Glimmer protects your company's data, the controls your IT team can check, and the choices your company makes about AI.

VERSION	DATE	FOR	CLASSIFICATION	SECURITY CONTACT
1.2	16 October 2026	IT and security reviewers	Public. Share freely	Yong Jie Wong

Sharing this with your IT team? It answers the questions IT teams usually ask before they approve an AI tool. Sections 2 to 5 cover the architecture, data, controls and assurance.

1 Summary


Glimmer is an AI agent that works inside your company's chat app and business systems. That access makes it useful, and it is also the risk: an agent can read data, take actions, and be tricked by an instruction hidden in a file. An AI model can make mistakes, so Glimmer puts the protections in the setup around the model. Each one is visible to your IT team in the Glimmer console.

KEY FACTS			
WHERE IT RUNS Your cloud account, or a server in your office	REGION Your choice, including Singapore	CHAT ISOLATION One virtual machine per chat, kept apart by KVM	OUTGOING TRAFFIC Approved list only. Blocked attempts are logged
PASSWORDS Held by a separate program per system. Never sent to the AI	AI MODEL Claude, GPT, Gemini or open models, by the route you choose	ACTIVITY LOG Every request, action, approval and blocked connection	CERTIFICATION ISO 27001 in progress, expected January 2027

Contents

1	Summary	1	5	Assurance	3
2	Architecture	2	6	Seven questions to ask any AI vendor	4
3	AI models and data retention	2	7	Who does what	4
4	Security controls	3			

2 Architecture



3 AI models and data retention

Glimmer is not tied to one AI company. You choose the model, how it is reached and where it runs, and Glimmer works the same way with each. Teams can use different models on the same Glimmer, and you can switch at any time. Whichever route you choose, Glimmer sends only the question and the text it needs.

ROUTE	MODELS	WHERE IT RUNS	DATA RETENTION
Your own account with the AI company	Claude, GPT, Gemini	The AI company's regions	You choose. Standard: kept up to 30 days for misuse checks, then deleted. Zero retention: nothing kept; Anthropic and OpenAI approve it account by account, and we apply for you.
Through the cloud you already use	Amazon Bedrock, Google Vertex AI, Microsoft Azure	The region you pick. Claude and Gemini models can run in Singapore	Amazon Bedrock keeps nothing by default. Google Vertex AI keeps a 24-hour cache, which can be turned off.
In your own environment	Open models such as DeepSeek, Qwen and SEA-LION	Your own cloud or your own office	Questions never leave your environment.

Training and retention follow the terms of the route you choose. Terms checked against each company's own pages on 10 October 2026. Companies change their terms often, so check the current page before you decide.

TERMS USED IN THIS DOCUMENT			
Virtual machine A complete computer that exists only in software. Glimmer uses KVM, the software built into Linux, to keep each one apart.	Approved list The systems and websites Glimmer may connect to. Your company sets it. Also called an allow list.	Planted instruction A hidden line in a file or email, written to trick an AI into acting. Security teams call it prompt injection.	Zero retention An agreement under which the AI company keeps nothing once it has answered.

4 Security controls

The last column names where your IT team can check each control.

#	CONTROL	WHAT IT DOES	WHERE TO CHECK
4.1	Your own account	Glimmer runs in a cloud account in your company's name, on your own AI key. If you do not have one yet, we open it for you.	Your cloud account
4.2	Sealed chats	Each chat runs in its own virtual machine, kept apart by KVM, the technology Amazon and Google use to separate their customers.	Console: agents
4.3	Approved connections	Each chat reaches only the systems on your list. Any other connection is blocked, and the block is logged.	Console: connection log
4.4	Passwords kept from the AI	A separate program for each system holds its password and signs in for the AI. A tricked agent has no password to give away.	Console: connected systems
4.5	People approve key actions	A team lead picks which actions need a person's OK, and whose. When you connect a new system, every change it makes needs approval until someone clears it.	Console: approval rules
4.6	Every agent has an owner	A new rule for a whole team waits for the owner's approval. Every change is kept in a history.	Console: agents
4.7	Full activity log	Who asked, what the agent did, which AI model answered, and who approved it. Team leads see their own agents. Managers see all of them.	Console: connection log
4.8	Spend you can see	Cost in dollars per agent and per team, with no licence for every person.	Console: cost

5 Assurance

CERTIFICATION

ISO 27001

Our security controls follow the ISO 27001 framework. Glints is working toward ISO 27001 certification, expected January 2027.

TESTED AT GLINTS

40 of 40

Glints' own Glimmer passes everything it reads through Jev, a checker from TypeSafe, which receives each text to score it. In a test of 160 real texts, it flagged all 40 planted instructions.

SUPPORT ACCESS

Named staff

Our support team reaches your Glimmer through a private network. Each person signs in with their own work login, and you can switch the access off at any time.

6 Seven questions to ask any AI vendor

Ask these questions of every AI vendor, including us. Clear, specific answers are a good sign.

#	QUESTION	GLIMMER'S ANSWER
6.1	Where does our data sit, and whose account is it in?	Your account In a cloud account in your company's name. The region can be Singapore.
6.2	Does each conversation run apart from the others?	Yes Each chat has its own virtual machine.
6.3	Can the agent see our passwords?	No A separate program for each system holds its password.
6.4	Which websites can the agent reach?	Your list only Anything else is blocked and logged.
6.5	What checks what the agent reads for planted instructions?	Tested at Glints Jev, from TypeSafe. See section 5.
6.6	Which actions wait for a person, and who decides?	Your team leads They pick the actions and the approver.
6.7	Can we see a log of everything the agent did?	Yes Every request, connection and approval, in the Glimmer console.

From "Is it safe to let an AI agent into your company's systems?" by Yong Jie Wong: ai.glints.com/articles/is-an-ai-agent-safe

7 Who does what

YOUR COMPANY

- Owns the cloud account Glimmer runs in, and decides who can sign in.
- Approves the list of systems Glimmer may reach.
- Chooses the AI model, the route and the data retention.
- Names the team leads who set the approval rules.
- Can switch off our support access at any time.

GLINTS

- Sets up Glimmer in your account and keeps it running.
- Reaches it only through a private network, with named staff who each sign in with their own work login.
- Removes a person's access to every deployment when they leave our support team.
- Hands over on request: your own IT, or another company, can take over the running of Glimmer.

Bring your IT team to a call

Our engineers open the Glimmer console live and answer your IT team's questions.

ai.glints.com/security

Security questions: Yong Jie Wong, Co-founder, Glints · yongjie@glints.com

About Glints. A Southeast Asian hiring platform that runs its own business on Glimmer, in every department.

15M+
people use Glints

50,000+
companies hire through
Glints

Claude
Partner Network member