

AI Usage

1 user attributed to a person, 1 unattributed client

AI CLIENTS

3

distinct AI tools observed:
apps, IDEs, CLIs

USER SESSIONS

0

AGENT SESSIONS

1,315

1,315 total

UNIQUE USERS

2

people or devices detected

AI PROVIDERS

1

anthropic

100%

Unprotected AI Usage

seen, not inspectable

1,309 nonbusiness AI connections seen. Unprotectable.

Business: 2 users, 1,843 requests. Inspected.

Nonbusiness: 0 users, 0 requests. Not securable.

Seen, not inspectable: 1,309 browser AI connections to claude.ai, chatgpt.com, chat.openai.com. Content cannot be inspected (bot-managed or pinned TLS).

4 endpoints enrolled. 0 bypassing (lower bound).

Risk Metrics

posture grade A

A

928/950

RISK TYPE COUNT

1 critical

0 high

0 medium

2 low

2 would block

40 would warn

1,843 inspected

DRIFT

0

behavioral deviations

SPRAWL

3

AI tools, 1.5 per user

RISK / INTERACTION

0.03

critical-equivalent events

Top Risks

Volume

1 Data Exfiltration

CRITICAL

2

avg severity 39.4

2 Adversarial Prompt Chaining

HIGH

1

avg severity 24.6

3 Prompt Injection

CRITICAL

1

avg severity 46.2

Compliance and Accountability

6 required roles with no owner

COMPLIANCE IN SCOPE

4

GDPR, ISO/IEC 42001, ISO/IEC 27001, HIPAA

RISKS IN SCOPE

5/21

from the register your regimes reach

ROLES UNOWNED

6

of 6 required; most exposed Compliance Officer (3)

Known Issues and How They Close

appliance is in **Training** mode; these steps convert visibility into protection


ISSUE	TIMELINE	WHAT IT TAKES
Turn on blocking Nothing is being stopped today. 2 actions would have been blocked had enforcement been on.	1-2 weeks READY NOW	Staff time only
Assign owners to 6 unowned roles 6 of 6 roles your selected compliance requires have nobody named against them, starting with Compliance Officer, AI Ethics Officer.	This week READY NOW	One decision per role, in Govern
Disable chat.openai.com + chatgpt.com access 1,309 browser AI connections seen and none of them inspectable.	30 days NEEDS SETUP	Policy decision + network change
Enrol every device Route all AI traffic through the appliance.	30 days NEEDS SETUP	IT rollout
Attribute every agent 1 AI client resolves to a fingerprint, not a person. Assign them by identity or email domain in Govern.	30-60 days READY NOW	IT rollout + Govern assignments
Review new AI tools before use 3 tools appeared this period without a sign-off step.	This quarter NEEDS SETUP	One owner, ~2 hrs/week

Usage Versus Risk

risk per interaction improving, 90% vs prior period

AI IDENTITIES SEEN 45 none traceable to a named person	AI TOOLS IN USE 1 3 consumer surfaces we cannot inspect	RISK PER 100 INTERACTIONS 2.2 severity-weighted, the did-we-improve number ▼ 89.5% vs prior period
--	---	--

45 AI identities are using 1 AI tools, producing 2.2 risk per 100 interactions. More actors reach for more tools, and each tool is another place company data can leave. The third number strips out volume, so it answers whether the organisation improved rather than whether it simply did more. 33,646 interactions inspected this window.



All three lines are indexed so an average day equals 100, which is what makes them readable together. Usage climbing while risk per 100 stays flat is a good outcome even though nothing fell. Showing the latest unbroken run of 9 daily buckets; earlier data is separated by a capture gap.

Main Areas of Concern

3 measured this window

- You cannot revoke what you cannot name.** 45 of 45 AI identities (100%) resolve to a rotating fingerprint rather than a person. When someone leaves, their agents keep running; during an incident there is no account to suspend and nobody to ask.
- Every verdict this period was advisory.** The platform decided to block 2 times and every request was forwarded anyway. Each event was recorded after it completed, not prevented. Detection without enforcement changes the record, not the outcome.
- A data-protection question about consumer AI has no answer today.** claude.ai, chatgpt.com, chat.openai.com are consumer surfaces: usage is visible as 1,309 connections, but what was sent is not inspectable (bot-managed or pinned TLS).

Tool Sprawl

4 surfaces, 3 not inspectable

AI TOOL	FIRST SEEN	INSPECTABLE
api.anthropic.com	2026-08-27	INSPECTED
claude.ai 1,131 connections	seen this window	NOT INSPECTABLE
chatgpt.com 141 connections	seen this window	NOT INSPECTABLE
chat.openai.com 37 connections	seen this window	NOT INSPECTABLE

Sprawl is measured as distinct destination surfaces, not vendors. Two tools from the same vendor are two places data can go and are counted separately.

Impact on AI Governance

4 selected regimes

WHAT THE RULES EXPECT	FRAMEWORK	WHERE WE STAND
Know what personal data goes into AI tools	GDPR	Answerable for inspected traffic; 1,309 consumer connections are not inspectable
An AI management system with assigned responsibilities	ISO/IEC 42001	6 required roles have no owner
AI use covered by the ISMS and its asset inventory	ISO/IEC 27001	Records exist for inspected traffic; none for 1,309 consumer connections
Control PHI disclosure to AI tools	HIPAA	Answerable for inspected traffic; 1,309 consumer connections are not inspectable

How the Risk Number Is Calculated

critical-equivalent events

Risk is not a count of events. Each event converts to critical-equivalent events: a figure of 47 means the equivalent of 47 serious incidents.

BAND	SEVERITY	COUNTS AS	READ IT AS
Critical	80-100	1.00	one serious incident
High	60-79	0.50	two of these equal one critical
Medium	40-59	0.20	five of these equal one critical
Low	1-39	0.05	twenty of these equal one critical
None	0	0.00	not counted

Bands step by two to four times because harm is not linear: even weighting would let a flood of trivia bury one genuine breach. Band edges match the console, the period in progress is excluded, and comparisons use trailing averages.

All figures measured on monitored posture as of 2026-09-23. Risk is severity-weighted: critical 1.0, high 0.5, medium 0.2, low 0.05.