AI inventory card

Fill in one card for your whole organization, before any AI strategy work, and keep it as a running list. Fill in what you know; every blank line gets an owner and a date. A blank line means: the count is not done, and that row is the next thing to fix.

QUESTION	ANSWER	OWNER / BY WHEN
What's on our approved list, and what do we already know runs outside it?		
By when is the first full count done? (Readiness and risks come after it.)		
Assistants staff bring themselves: who asks staff openly, and by when?		
Which regulated data, such as client or patient records, is most likely being pasted into unsanctioned assistants?		
Agents wired to team accounts (e.g. an assistant drafting shared-inbox replies): who reports what runs, team by team?		
AI a department built itself: who in each unit names what that unit runs?		
AI inside software we already bought: which vendors ship AI features, and who adds them to the inventory and the risk assessment?		
What is monitored and which sites, and is that published plainly to staff?		
How are the people being watched involved in how the watching runs?		
When watching catches accidental oversharing: who coaches in the moment, and when does discipline apply?		
Before any detection or disciplinary step: who makes the open ask about assistants and agents staff already run?		
Policy test: would admitting unapproved use make the next person likelier to speak up?	☐ Yes ☐ No	
Any tracking of AI use, such as developers' token use: protects privacy, limits the data collected, used only for its stated purpose? (all three)	☐ Yes ☐ No	
Which rows of the running list still lack a named finder or a date?		
Decision and date	☐ INVENTORY READY ☐ NOT READY: NEXT FIX _____ Date _____ Signed _____	

AI inventory card

Example (fictional): the count at a 400-person regional insurance broker

QUESTION	ANSWER	OWNER / BY WHEN
What's on our approved list, and what do we already know runs outside it?	6 apps chartered in March. Outside it so far: 2 browser assistants on the support desk and an AI note-taker add-in in the sales CRM	Head of IT operations · list refreshed monthly
By when is the first full count done? (Readiness and risks come after it.)	30 November, ahead of the January strategy offsite	Chief of staff
Assistants staff bring themselves: who asks staff openly, and by when?	Every team lead asks in the November one-to-ones; one open question, no names recorded	All team leads · by 30 Nov
Which regulated data, such as client or patient records, is most likely being pasted into unsanctioned assistants?	Client claims files and health questionnaires; 2 coached cases on the support desk so far	Compliance officer · by 15 Dec
Agents wired to team accounts (e.g. an assistant drafting shared-inbox replies): who reports what runs, team by team?	Claims wired an assistant to the shared inbox in August to draft replies; other teams report in the same open ask	Claims team lead · by 30 Nov
AI a department built itself: who in each unit names what that unit runs?	Underwriting's document summarizer, running since spring; every unit to name its own finder in the November ask	Head of underwriting · by 30 Nov
AI inside software we already bought: which vendors ship AI features, and who adds them to the inventory and the risk assessment?	9 of our 12 main SaaS vendors now ship AI features; 4 assessed so far	Procurement lead with IT risk · by 15 Dec
What is monitored and which sites, and is that published plainly to staff?	Browser monitoring covers 30 named AI sites; the list went to all staff in October's security note	IT security lead · reviewed quarterly
How are the people being watched involved in how the watching runs?	**Blank: not ready**	HR director · by 15 Dec
When watching catches accidental oversharing: who coaches in the moment, and when does discipline apply?	First-time oversharing is coached by the team lead in the moment; discipline only for repeat, deliberate sharing	HR director · policy note out 1 Dec
Before any detection or disciplinary step: who makes the open ask about assistants and agents staff already run?	The COO sends the open ask company-wide; detection rollout deferred until after the count	COO · by 30 Nov
Policy test: would admitting unapproved use make the next person likelier to speak up?	☒ Yes: the open ask in claims brought 5 more admissions the same week	[illegible]
Any tracking of AI use, such as developers' token use: protects privacy, limits the data collected, used only for its stated purpose? (all three)	Security's draft logs every developer prompt in full → **Fails the bar**: it collects more than the cost question needs	IT security lead · revised draft by 15 Dec
Which rows of the running list still lack a named finder or a date?	Marketing and operations have not named finders yet	Chief of staff · by 30 Nov
Decision and date	**☐ INVENTORY READY ☒ NOT READY. NEXT FIX: involve the watched staff, revise the token-tracking draft · Date 12 Nov · "Two rows not ready; the count continues."**	COO