A CXO Briefing · Enterprise AI & Data Safety

Safely Using AI Within Your Organization

The four levels of data control, and which one your firm needs

A plain-English guide for leaders deciding how their teams should use AI without putting client data, intellectual property, or compliance at risk. Built to be read by you and passed to your technology lead.

Executive Summary

Your team is already using AI. The only real question is whether they are using it safely.

Most enterprise AI risk does not come from the technology. It comes from not knowing which level of control you are operating at. There are four, and the gaps between them are large: the difference between a setting and a contract, and between trusting a toggle and owning your infrastructure.

Level 1 is consumer tools used carelessly, where your data can train someone else's model. Level 2 is the same tools used responsibly, which is better but still rests on a switch you have to trust. Level 3 is enterprise agreements, where your data is not used for training by contract, and this carries most of what most firms do. Level 4 is self-hosting, where your data never leaves your walls, and almost every organization has some slice of data that deserves it. With India's data protection law now in force and full compliance due in 2027, the distance between "we use AI" and "we use AI safely" is becoming a board-level question.

This briefing lays out the four levels, what actually separates consumer terms from enterprise terms, and a simple guide for matching each kind of data to the right level of protection.

01 The Control Ladder

The four levels of AI data control

Read this as a ladder. As you climb, your control over your data rises, and so does the cost and effort required. The goal is not to reach the top. It is to match the level to the sensitivity of the data.

1
TIER 01

Consumer tools, used carelessly

Staff paste client data, contracts, or financials into free or personal AI accounts. This is the leak that almost certainly exists in your firm today, whether or not anyone has approved it.

DATA CONTROL — VERY LOW
COST & EFFORT — NONE

2
TIER 02

Consumer tools, used responsibly

The same paid apps, but with model-training turned off in privacy settings and basic rules in place. Safer, but you are trusting a toggle and a policy, not a contract, and some flagged conversations can still be reviewed.

DATA CONTROL — LIMITED
COST & EFFORT — LOW

3
TIER 03

Enterprise agreements

Business and enterprise plans from the major providers. Your data is not used for training by default, retention is limited, and you get admin controls and audit logs, all backed by a signed contract rather than a setting. This carries most of what most firms do.

DATA CONTROL — HIGH
COST & EFFORT — MODERATE

4
TIER 04

Self-hosted models

Open-weight models run on infrastructure you control, so prompts never leave your environment at all. Full sovereignty, the highest cost and skill requirement. For data that legally or competitively cannot touch a third party.

DATA CONTROL — TOTAL
COST & EFFORT — HIGH

02 The Immediate Fix

Before you climb: govern levels 1 and 2

Most firms cannot jump straight to enterprise tools across the board. The immediate priority is closing the leak that already exists, which is a policy problem before it is a technology one.

A one-page AI usage policy should cover

- **What never goes into a consumer AI tool:** client data, personal data, contracts, financials, passwords, or source code.
- **Training opt-out as a baseline:** if staff use paid personal accounts at all, training must be switched off, and someone must verify it. On consumer plans this defaults to on unless a person turns it off.
- **One sanctioned tool:** give people an approved enterprise option, so the safe path is also the easy path. Bans without an alternative just push usage underground.
- **Awareness, not fear:** most leaks are accidental. A short briefing on what is safe to paste prevents more incidents than any block list.

03 Consumer vs Enterprise

What actually changes when you pay for the business plan

The model can be identical. The terms governing your data are not. This is the single most misunderstood point among buyers, and the one your technology lead will want to see clearly.

	Consumer plans Free and personal paid accounts	Enterprise plans Team, Enterprise, and business accounts
Used to train their models	By default, unless you opt out	**No, by default**
What protects you	A setting you toggle	**A signed data-handling contract**
How long data is kept	Up to 5 years if training is on; 30 days if off	**A short window, often only days**
Admin and audit controls	Minimal	**Central admin, audit logs, single sign-on**
Option to store nothing at all	Not available	**Available for qualifying use cases**
Formal agreement for regulated or health data	No	**Available on eligible plans**

Exact terms and timelines differ by provider (OpenAI, Anthropic, Microsoft Azure, Google) and change often. Confirm current terms with each vendor before relying on them. One phrase to know is "zero data retention": even on enterprise plans, providers usually keep prompts for a short window to monitor abuse, unless you negotiate zero data retention, which removes that storage entirely.

04 Full Control

When even enterprise is not enough

Enterprise plans keep your data private, but it still travels to the provider's servers to be processed. For a small set of firms, and for a small slice of data inside almost every firm, that itself is the line they cannot cross.

Level 4 in practice

You run open-weight models, such as Llama, Mistral, Qwen, or DeepSeek, on infrastructure you control, either on your own servers or through managed cloud compute like Microsoft Azure AI Foundry, where the model runs inside your own private environment. Your prompts never leave your walls.

This is worth it when data volumes are high and steady, or when regulation or competitive sensitivity forbids any outside processing. For variable or low usage, a managed enterprise plan is usually cheaper, faster to stand up, and good enough. The build itself is a separate, deeper topic for a technical team.

05 Why Now

India's data protection law has arrived

This is no longer only good practice. It is becoming a legal obligation with real penalties.

From guidance to law

India's Digital Personal Data Protection rules were notified in November 2025 and are rolling out in phases. 2026 is the build year.

Nov 2025

Data protection rules notified; the Data Protection Board is established.

Nov 2026

The consent-manager framework becomes operational.

May 2027

Full compliance: notice, consent, security safeguards, breach reporting, and individual rights.

The stakes: penalties run up to ₹250 crore for failing to maintain reasonable security safeguards, and a breach must be reported within 72 hours. How your firm processes personal data through AI tools sits squarely inside this.

06 The Takeaway

Which level for which data

The practical takeaway. Match the data, not the hype. Most firms will use more than one level at once.

Public or low-stakes work Marketing copy, general research, brainstorming with no internal data	LEVEL 1–2
Internal, non-sensitive material Drafts, internal notes, operations that carry no personal or client data	LEVEL 2–3
Client data, financials, personal data, contracts Anything covered by confidentiality or data protection law	LEVEL 3
Highly regulated or sovereignty-bound data Where data legally or competitively cannot reach a third party at all	LEVEL 4

Drawn from the public terms and documentation of Anthropic, OpenAI, and Microsoft Azure, and India's Digital Personal Data Protection Rules, 2025. Current as of June 2026.