KIZATA

Account architecture and IP

A method for determining what a tier permits
Foundation Session 0, AI Operating Core · September 2026

This reference accompanies question three of Foundation Session 0. It sets out a method rather than a set of answers, because provider terms change and a method outlasts any particular version of them.

Nothing here is legal advice. It is the sequence a practitioner uses to establish what is actually permitted before deciding what to enter into a tool, and where the question stops being answerable by a provider at all.

1. Three questions, routinely treated as one

Most discussion of AI and confidentiality collapses three separate questions into a single one, usually phrased as whether a tool is safe. Collapsed that way, the question has no actionable answer. Separated, each has a different source of authority.

QUESTION	ANSWERED BY
Do inputs train the model?	The account tier and the governing agreement
Who owns the output?	The governing agreement, and it is not always the user
May this material be entered at all?	Obligations to clients, employers and regulators. **Not by the provider**

The third question carries the professional risk and no provider can answer it. A permissive contract does not release an obligation of confidentiality owed to a client. A firm may be contractually free to enter material that it is professionally prohibited from disclosing to any third party, and the provider's terms are silent on that point by design.

Practical consequence. Questions one and two are research tasks with findable answers. Question three is a judgement that belongs to whoever holds the obligation, and it is answered by reading engagement letters and professional rules rather than terms of service.

2. Where consumer and business tiers commonly differ

The pattern below is stable across major providers. The specific position of any one provider at any one time is not, which is why section 3 exists.

	CONSUMER TIER	BUSINESS OR ENTERPRISE TIER
Training on inputs	Commonly permitted by default, sometimes with an opt-out	Commonly excluded by contract rather than by a setting
Retention and deletion	Set by the provider, changed at the provider's discretion	Set by the administrator, within contractual limits
Output ownership	Granted to the user in most terms, with carve-outs	Defined explicitly in the agreement
Indemnity for output	Generally absent	Sometimes offered, usually conditioned on using the product as directed
Administrative control	None. The user is the account	A tenant administrator can enable and disable capabilities
Audit and logging	Not available to the user	Commonly available, and commonly required by the buyer

Two observations that matter more than the table.

The two tiers frequently look identical in use. The same interface, the same model name, the same features. Tier is a billing fact, not a visible one, which is why section 3 step 1 reads the billing record rather than the screen.

An opt-out and a contractual exclusion are not equivalent. An opt-out is a setting that a person can change, that a new device or a re-install can reset, and that carries no remedy if it fails. A contractual exclusion is an obligation with a counterparty. Where the distinction matters, only the second is worth relying on.

3. The method. Four steps, in order

Step 1. Identify the account type from the billing record

Not from the interface, and not from the model name. Open the billing or subscription record and read what is being paid for, by whom, and under which entity. A personal card paying for a personal subscription used for company work is a consumer tier regardless of how the work is described.

Common finding. An organisation believes it holds a business tier because some staff do. Individuals who signed up independently are on consumer tiers inside the same organisation, and nobody has an inventory.

Step 2. Locate the terms that govern that tier

A business or enterprise tier is usually governed by a negotiated or click-through agreement that is not the public terms page. The public page frequently states that it does not apply where a separate agreement exists.

Where to look. The order form or subscription agreement, any data processing addendum, and any product-specific terms incorporated by reference. Incorporated documents are part of the agreement and are commonly where the operative language sits.

Step 3. Find the training and retention clauses by name

Read what the agreement says about three things specifically:

- Whether inputs, outputs, or both may be used to train or improve models
- How long either is retained, and what triggers deletion
- Whether human review of inputs is permitted, and under what conditions

Human review is the clause most often missed. A provider may exclude training on inputs while permitting human review for abuse detection or quality. Those are different exposures and a confidentiality analysis that only asks about training will not surface the second.

Step 4. Confirm who administers the tenant, and what they have enabled

An administrator setting can override a contractual permission in either direction. A capability the agreement permits may be disabled centrally, and a capability assumed to be off may be on.

This step decides whether certain work is possible at all. Scheduled or automated runs, connections to file stores and mail, and data-retention controls are commonly administrator-gated. A person using the tool has no reliable way to see which of these apply to them.

Where there is no administrator, there is no tenant. That is itself a finding, and it means steps 2 and 3 are the only controls in place.

4. What the method produces, and what to do with it

The output is a written statement of fact about a contract, not an opinion about a tool. Something close to: this account is a business tier under agreement X; inputs are excluded from training under clause Y; retention is thirty days, administrator-configurable; human review is permitted for abuse detection; scheduled runs are disabled at the tenant.

Where the answer cannot be located, that absence is the finding. An account whose governing terms cannot be identified should be treated as a consumer tier until proven otherwise. This is the most common outcome of a first pass and it is not a failure of the method.

A short list is more useful than a long policy. Three or four categories of material that do not enter any tool, written down, is a control that a person can actually apply. It is also the beginning of a policy, written before anyone asked for one.

Where this stops being a training question. An organisation that cannot answer question three from section 1 has a governance question rather than a skills question, and the two are addressed differently. Establishing which one is in play is the useful result of this exercise.

5. Where this sits in the six weeks

Account architecture is foundation rather than curriculum. It determines what is possible before any fundamental is applied, which is why it is taught in the free hour rather than in a paid week.

The program builds on it at three points.

Week one records the account tier and the administrative constraints during setup, because they decide which of two tool layers a participant works in and whether week six can schedule a job or must use a calendar trigger instead.

Week four connects mail, calendar, files and project systems, and sets a standard for which material stays with a person. That standard is derived from the answer to section 1, question three.

Week six teaches two fundamentals directly downstream of this reference. *Found or figured* marks what came from a document and what was concluded. *Back to the source* checks a claim against an original rather than a summary. Together they are the verification practice that a confidentiality analysis assumes but does not supply.

One discipline the program applies to itself. Nothing a participant works on enters a system Kizata controls. Work is shown in session or shared through a link the participant holds. That is the same practice taught in week six, applied to the program first.

What follows

AI Operating Core. Six weeks, one hour a week, live and online. Ten participants. Thirteen named fundamentals across five families, installed on participants' own accounts using their own work, with two workflows measured in week one and again in week six.

Six hours in session. Nine hours applied. Fifteen in total.

No cohort has completed, so the program promises a measurement and not a result. The first measured result is dated October 27, 2026.

Current cohort dates, terms and booking: kizata.ai/course