CONTROL AND DATA

Control and data in an Ampity AI workforce

How accounts, access, authority and data are handled. Specific controls are agreed and documented for each engagement. This guide explains the model and the questions to settle before go-live.

Accounts in your name

The AI provider and tool accounts your workforce uses are set up in your name, on plans that support the intended use, where the provider offers this. Ampity manages them with your authorization.

- You see the provider's own billing and usage records.
- Ampity does not pool subscriptions, share logins, or reuse one customer's subscription access for anyone else.
- Where a subscription's terms do not support a use, that work runs on API access or another supported option, scoped and budgeted upfront.
- If you stop working with Ampity, the accounts remain yours.

Least-privilege access

- Each role gets only the systems, records and actions its charter needs.
- Read access is kept separate from write access where your systems allow it.
- Roles that handle different kinds of data use separate credentials.
- Access is reviewed when responsibilities change, and removed at exit.

Delegated authority, written down

Automatic work inside the written authority is the point of an AI employee. Approvals are kept for the specified decisions outside it. Authority changes only with your approval.

Does automatically	Brings to a named person	Never does
Agreed duties in the charter, such as record updates, reminders, reports and routine follow-ups	Specified decisions, such as a first message to a new contact or a change above an agreed limit	Actions you rule out, such as changing contract terms, deleting records or moving money

Examples are illustrative. Each role's authority is agreed during role design.

Secrets, data boundaries and processing

Secrets and credentials

- Stored in a secrets manager or vault in the agreed environment, not in prompts, documents or code repositories.
- Rotated under your policy, and when people or vendors change.
- Never shared between customers and never pasted into chats.

Data boundaries

- For each role, agree which data it may read, which it may change and which it must never see.
- Sensitive categories, such as payroll, health or government identifiers, are excluded unless explicitly agreed.
- Each step sends a model only what that step needs.
- Repeatable steps run as code, which can keep some data from being sent to a model at all.

Who processes your data

Depending on the configuration, data may be processed by your own systems, the AI providers chosen for each role, channel tools, and the infrastructure the operating software runs on. We map this for each role before go-live.

Component	What it may process	Whose terms apply
Your systems, such as CRM, email and files	Your records	Your existing agreements
AI model or assistant provider	The text a step sends and the response	The provider’s terms on your account
Channel or calling tools, where used	Messages, transcripts or call audio	That tool’s terms
Ampity operating software and hosting	Task records, logs and configurations	Your proposal and the hosting provider’s terms

Local-only processing is not assumed. Where it is required, it is scoped as a specific requirement. Provider certifications belong to those providers and do not certify the AI Workforce service.

Records, retention and exit

Activity records and logs

- Actions, decisions, approvals and checks are recorded for each task, as configured for the role.
- What is logged, and where, is agreed during design.
- Logs can contain business data, so they follow the same access rules as the data itself.
- Retention follows your policy where the systems involved allow it. Providers' own retention applies to the data they process.

Questions to settle before go-live

- Which accounts will be in our name, and on which plans?
- Which systems and actions does each role need?
- What exactly is in the written authority?
- Which data categories are excluded?
- Which providers and tools process data, and in which regions, where that matters to us?
- What is logged, where, and for how long?
- Who is notified, and how, if something goes wrong?

Exit and handover questions

- Which documentation and configurations do we receive?
- How are custom components owned or licensed?
- How is Ampity's access removed, and how is that confirmed?
- What happens to logs and stored context held in Ampity-managed systems?

Next step

Ask for an access and data plan for the roles you are considering. It maps accounts, access, authority and data processing before anything is built.

hello@ampity.io · https://ampity.io/company/contact