BUYER FAQ AND ROLLOUT

Questions buyers ask about an AI workforce

Straight answers to the doubts that come up before a decision, and how a rollout works in stages.

Why a service

Why buy a managed service instead of a platform or an internal build?

Platforms are good products, and many can be customized. An internal build gives you full control. Both work when someone in your company has the time and skill to design the roles, connect the systems, check the output and keep it running. Ampity takes on that engineering and operating work. We design each role around its full ongoing responsibilities, build the connectors, scripts and checks it needs, and keep maintaining it under agreed coverage. Choose the option that matches who you want to carry the work.

What does Ampity stay responsible for after launch?

Monitoring, fixes to agreed connectors and checks, upkeep of the role's business context and examples, KRA and budget reviews, and agreed improvements. Support hours, response expectations and the change process are written in your proposal.

How does onboarding work?

Like a new hire. We agree the role charter, KRAs and authority with you, teach the role your processes with real examples, and give it only the access it needs. During supervised onboarding it prepares work for review, your team corrects it, and reviewed corrections become maintained examples that are retested.

What can an AI employee do on its own?

Everything inside its written authority, automatically: the agreed duties in its charter, such as record updates, reminders, reports and routine follow-ups. Specified decisions outside that authority go to a named person with the context to decide. Some actions are ruled out entirely. Authority grows only as reviews show reliable work, and you can narrow it at any time.

Is this a chatbot or an outbound messaging tool?

An AI employee owns ongoing operational responsibilities inside your systems, such as keeping records current, preparing briefs and packs, tracking commitments and coordinating handovers. Outreach can be included where it fits the agreed role and follows your audience, messaging, permission and sending rules. It is managed against role outcomes and KRAs, rather than just message volume.

Working with it

How do we give it new requests?

In plain language, through the channels configured for the role, such as email or team chat. Requests inside its charter are handled as normal work. Requests outside it are flagged instead of attempted, and we discuss whether they should become part of the role.

How do we teach it a correction?

Correct the work where you review it and say why. Ampity reviews the correction, adds it to the role's maintained examples or rules, and retests before relying on it. Corrections are not left to chance in a chat history.

Can it work in Microsoft Teams, Slack, email or on calls?

Where those channels are supported, configured and tested for your role. Email and team chat are common. Meeting and calling capabilities depend on the tools involved, their terms and any extra cost, and are scoped separately. We do not assume a channel works until it has been set up and tested in your environment.

Who owns the accounts and data? Is it secure?

AI provider and tool accounts are in your name, on plans that support the use, where the provider offers this. Your data stays in your systems unless the configuration sends it to a specific provider or tool, and we map which components process which data before go-live. Each role has least-privilege access and separate credentials, and secrets are kept in a secrets manager, not in prompts. We do not pool subscriptions or share credentials. Provider certifications belong to those providers and do not certify the AI Workforce service. The access and data plan for your engagement sets out the specific controls.

What if it gets something wrong?

Where checks are configured for a responsibility, they read your source records before work is marked complete, so many errors are caught before anyone relies on the result. Anything with consequences outside its authority waits for a named person. Activity records show what was done and why, which makes errors traceable. Each error is corrected, and the fix is added as a check or example and retested.

Does it replace people on our team?

It takes recurring operational work off your team so people spend more time on judgement, customers and decisions. Staffing decisions remain yours. We judge the role by its KRAs and its total operating cost per accepted result, not by headcount.

Limits, failures and coverage

Do our AI subscriptions cover all the usage?

Not always. Suitable subscriptions in your name can provide included capacity for supported work, and repeatable steps run as custom software rather than repeated model calls, which keeps usage down. But every subscription has limits and terms. We model usage against your volumes and current plans before the build.

When do we need API access or paid extra usage?

When a task falls outside what a subscription's terms support, when volumes exceed included limits, or when a step needs a capability only available through an API or another tool. These cases are identified and budgeted upfront, not discovered on an invoice.

What happens when a limit is reached or a system fails?

- At a provider limit, work queues or pauses under the agreed rule and you are told. Extra usage is used only within the agreed budget.
- When a system or connector fails, affected work is not marked complete. It is held as blocked, the issue is raised, and Ampity investigates within the agreed coverage.
- Temporary failures are retried within agreed limits. Repeated failures get a root-cause fix and a new check.

How is cost reported?

Each role has a budget. Usage is reviewed from the records each provider or tool supports, and cost is reported per accepted result in the regular KRA and budget reviews.

What support coverage do we get?

The coverage written in your proposal: support hours and channels, response expectations, what is included in ongoing operation and what is a change request. We do not offer round-the-clock coverage unless it is specifically scoped and agreed.

Can we leave? What do we keep?

Yes. The accounts are already yours. On exit you receive the documentation and configurations the contract sets out, ownership or licensing of custom components follows the proposal, and Ampity's access is removed and confirmed.

Evaluation and rollout

How should we evaluate it?

- Agree what an accepted output is for each responsibility before you start, and check it against your own records.
- Compare the same work over the same observation period, before and after.
- Count everything: provider plans, extra usage, software, implementation, management, and your team's review and rework time.
- Track KRAs, escalations and rework, not only volume. The KRA and economics workbook is built for this.

Rollout in stages

Stage	What happens	You decide	Move on when
1. Role design	Charters, KRAs, authority, data plan and cost model	Which roles, and the authority	Charter and plan approved
2. Build and connect	Agreed connectors, scripts and checks built and tested on your accounts	Access approvals	Tests pass on real data
3. Supervised onboarding	The role prepares work for review; corrections become maintained examples	Corrections and acceptance	Acceptance rules met consistently
4. Delegated operation	Agreed responsibilities run automatically within authority	What stays with a person	KRA reviews show reliable work
5. Review and expand	Regular KRA and budget reviews	Add, narrow or stop	Results justify the next step

Timing depends on the roles, the systems involved, access approvals and your team's review time. We agree a plan in the proposal rather than quote a fixed timeline. Commercial terms are set out in a written, scoped proposal.

Next step

Start with a role design session for one to three roles. You leave with draft charters, authority and a cost model, and decide from there.

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