

A SCOPING CHECKLIST FOR MSPS: BUILD, EXTEND, OR BUY

Before you build it: **know which layer is yours**

What decides whether an agent survives a hundred accounts is settled before the first line of code.

Every CloudOps platform has three layers. The **foundation** is the same for every MSP. **Specialist capability** moves fast and needs constant upkeep. **Your knowledge**, including the agents your customers talk to, is your IP. Each area below is tagged by layer.

Who owns each layer on each path

	Build	Extend	Buy
Customer experience	You	You	MontyCloud
Layer 3 · Your knowledge	You	You	You
Layer 2 · Specialist capability	You	MontyCloud	MontyCloud
Layer 1 · Foundation	You	MontyCloud	MontyCloud
What you answer	Every section, with a funded owner.	05, 08, and 09. Ask MontyCloud to answer the rest.	08 and 09. Ask MontyCloud to answer 01 to 07.

01 TENANT MODEL AND AUTHORITY

Layer 1 · Foundation

- ☐ Who or what holds credentials in each customer account, and at what privilege level?
- ☐ Can context or memory from one customer reach another? Is that prevented by design, or by intent?
- ☐ How is authority delegated per customer, who can change it, and where is that change recorded?

02 POLICY AND APPROVALS

Layer 1 · Foundation

- ☐ Which action classes may execute without a human, and does that answer differ per customer?
- ☐ Is policy evaluated before execution, or checked after the fact?
- ☐ Where does an approval request go, who is named on it, and what evidence arrives attached to it?
- ☐ Is agent authority earned level by level, against evidence, or granted after a good demo?

03 FAILURE AND CONTAINMENT

Layer 1 · Foundation

- ☐ When a run fails halfway, what halts the steps that depend on it?
- ☐ What state is captured before a change so it can be reversed, and what happens to actions that cannot be?
- ☐ Who is notified, through which system, and what do they see first: the record, or a mystery?
- ☐ Does a contained failure become a test case, or does the same surprise return?

04 EVIDENCE

Layer 1 · Foundation

- ☐ Can you produce, per customer and per period, what the agent did, under what authority, with whose approval, and whether the intended result occurred?
- ☐ Is that record produced by the system that did the work, or reconstructed afterward from logs written for another purpose?
- ☐ Does one record satisfy both your customer and your program reviews?

Only 21% of organizations have a mature governance model for agentic AI (Deloitte, State of AI in the Enterprise 2026, n=3,235). AWS MSP Program Validation Checklist (VCL) 8.0 carries 61 controls across six sections, with AI and agentic controls growing from 2 to 24.

05 CUSTOMER-FACING ACCESS

Layer 1 · Foundation

- ☐ Which dashboards, reports, and self-service automations will your end customers see?
- ☐ How is each customer's access scoped to its own accounts and data?
- ☐ Can you trace what a customer viewed or ran, and revoke that access at once?

06 COST

Layer 1 · Foundation

- ☐ What does a single run cost, by workflow and by customer?
- ☐ Do you know that before you price the service, or after the quarter closes?
- ☐ Where are the caps, who is alerted before one is hit, and what happens when it is?
- ☐ Which runs are doing deterministic work that a documented automation could do with no model in the path?

Agentic workflows can consume hundreds of thousands of tokens in a single session (EY, June 2026). Only 26% of large organizations have real-time visibility into AI operating costs (KPMG AI Quarterly Pulse Survey Q2 2026).

07 SCALE PAST THE FIRST ACCOUNT

Layer 1 · Foundation

- ☐ What changes between the first customer and the hundredth?
- ☐ How is one workflow configured to each customer's risk appetite without forking it?
- ☐ Can you see across every customer without mingling their data underneath?
- ☐ When one customer's run fails, what guarantees the other ninety-nine are unaffected?

08 UPKEEP AND OWNERSHIP

All layers

- ☐ Who owns this twelve months from now, and what is the budget line called?
- ☐ What happens when a model version, MCP server, or AWS API changes underneath you?
- ☐ How do you know the agent still performs as it did at review, and what test tells you?
- ☐ How closely does your provider work with AWS and model providers, and how fast does it absorb change?
- ☐ If the engineer who built it leaves, what is handed over?

09 YOUR KNOWLEDGE

Layer 3 · Your knowledge

- ☐ Which of your services, SOPs, and playbooks would a customer notice if they disappeared?
- ☐ How much engineering time goes to layer 3 today, and how much to layers 1 and 2?
- ☐ Which agents will your customers talk to, and who owns them?
- ☐ What would your team build first if layers 1 and 2 were already handled?

A note on what this list is not. None of these is a reasoning problem, so a better model next year will not solve them. They are operating problems. In the AI era, the decision is not build versus buy. It is build, extend, or buy, and it comes down to where your best people spend the next two years. You remain the trusted advisor to your customers. We are glad to walk through any line with your team.