

LIVEWYER

AI AGENT CONTROL PLANE

Governed Agent Pilot

Executive summary

The background

Your engineers already use coding and operations agents on laptops and in CI. Your security team needs to know what each agent can reach, who authorised it and where the record is kept. A ban slows delivery; unmanaged access leaves your team exposed.

You need a governed way to run agents on your own estate, with evidence your platform and security teams can inspect.

The opportunity

The AI Agent Control Plane pilot gives two of your real workflows a place on your Kubernetes platform, with sandbox isolation, restricted access and scoped identities. Every pilot task produces a signed record of its authorisation, configuration and permitted access.

We select a deployment baseline with you, test 20 failure scenarios and hand over the findings with configuration and runbooks. Your production decision rests on the results from your own workflows.

£42,000 fixed price: a two-day Agent Readiness Workshop followed by a five-week pilot. Two workflows, one non-production cluster and one deployment baseline. Production implementation is separately scoped and quoted.

Pilot success criteria

Agree the criteria with your workflow owners, platform team and security lead before the build begins.

- **Useful work:** assess output quality, reviewer effort and operating cost against the tasks you selected.
- **Control behaviour:** test permitted access, policy enforcement and the operator's ability to stop tasks and revoke access.
- **Evidence:** verify signed permission records and document what activity logs can establish.

The five-week pilot

Agent Readiness Workshop · two days

Map existing agent use and data paths. Select two workflows, agree approved models and tools, and define acceptance criteria. Bring a platform owner, a security partner and someone able to accept each workflow's output.

The workshop precedes the schedule below. The build needs one non-production Kubernetes cluster, approved model access and representative test data.

Stage	Work and review
Week 1 Design and access	Confirm the selected baseline, isolation model, data paths and identity design. Prepare the runtime and agree signing-key ownership.
Week 2 First workflow	Implement the first workflow with restricted access, configuration checks and signed evidence export. Review its output with the workflow owner.
Week 3 Second workflow	Onboard the second workflow. Configure its scoped tool access and review the engineer and operator experience.
Week 4 Validation	Run the programme of 20 failure scenarios adapted to the baseline. Record observed controls, bypass paths and gaps with security.
Week 5 Handover	Review acceptance evidence, rehearse stop and revoke, and hand over configuration and runbooks. Present the production recommendation and quote.

Representative workflow choices

Dependency maintenance can end in a proposed change for human review. A diagnostic workflow can inspect a non-production fault using read-only access and return a verifiable report. The workshop selects your actual workflows.

The fixed pilot implements one baseline. A runtime comparison build or a later runtime change requires its own scope and validation.

The included readiness workshop is separate from LiveWyer's AI-Assisted Engineering training.

Make the controls inspectable

Included implementation

Configure sandbox isolation, restricted egress and scoped identities around the selected workflows. Place configuration checks at enforcement points the baseline supports, such as Kubernetes admission, CI or a runtime API, and document possible bypass paths.

Signed evidence records

Every pilot task produces a signed record linking task identity and authorisation to the configured image, approved model and tools, permitted destinations and policy version. Agree who holds signing keys and show how the client verifies a record.

A permission record describes authorised configuration. It does not by itself prove which tools were called, what data was read or what the agent did throughout its run. Assess runtime telemetry and retention separately and document collection gaps.

Data and model boundaries

Hosting agent execution on your Kubernetes does not mean model inference is local. Agree where prompts, tool data and outputs may be sent, who retains them and for how long.

Evaluate operating requirements

- Can an engineer submit work and interpret a rejection or failure?
- Which consequential actions need human approval?
- Which duration, retry, concurrency and cost limits can be enforced?
- Can an operator stop a task, revoke access and diagnose the outcome?

Treat these as acceptance questions. Record unsupported controls as gaps; selecting a runtime does not establish that it supplies them all.

Test failure and useful output

Week four runs a programme of 20 failure scenarios tailored to the baseline and selected workflows. Agree the scenario definitions and expected outcomes with security before execution.

Control behaviour

Include prompt injection, attempted data exfiltration and unapproved actions. Check enforcement and operator response, recording passed, failed and untested behaviour separately.

Workflow value

Have the owner judge the result against acceptance criteria. Record review time, retries and available model and infrastructure costs per accepted task.

Evidence for the decision

- Client-held configuration, version inventory and operating runbooks.
- Signed authorisation and configuration records with verification instructions.
- Test definitions, observations and a control-gap register.
- Workflow results, review effort and cost observations with their measurement limits.
- Knowledge transfer and a stop/revoke rehearsal with the client operator.

A passed scenario demonstrates the tested behaviour in the recorded environment. It does not certify the platform or prove protection against every attack. No savings or productivity result is assumed in advance.

Decide the next deployment step

Use the pilot evidence to proceed, narrow the use case, do further validation or stop. Production work is a separate decision with an explicit scope and quote.

Production questions

Review availability and recovery, tenant boundaries, data retention, capacity, support ownership and additional integrations. Moving to another runtime can require changes to policy checks, evidence collection and failure tests.

Your team retains the work

Hand over configuration and documented control requirements with the evidence definitions and runbooks. Agree the ownership of retained records and remove temporary access.

Relevant LiveWyer experience

Our [Siemens Digital Industries case study](#) demonstrates our Kubernetes platform delivery experience.

[AI-Assisted Engineering training](#) helps teams develop their use and review practices. Our [AI Engineering Principles](#) describe human responsibility for engineering output.

Discuss your agent workflows

Bring a representative task, its data and access needs, and the decision the pilot must support.

livewyer.io/contact/

£42,000 includes the two-day workshop and following five-week pilot. Availability, tax treatment, expenses and contractual terms are confirmed in the proposal.