

— KNOWLEDGE, NOT A BLANK CANVAS

Turn how your team fixes things into governed agents.

Your engineers already know the fix for the crash loop, the orphaned disk, the stale IAM key. Agent Builder turns that know-how into agents that act inside limits you set — scoped and traceable from the first run.

60%

LESS
MANUAL
TRIAGE &
CLEANUP

45%

FASTER
OBJECTIVE
TO FIRST
RUN

3×

MORE
AGENTS
PROMOTED
PER
QUARTER

\$18

AVG. COST
PER
AUTOMATED
RUN

— DESCRIBE. COMPOSE. GOVERN. PROMOTE.

01 Describe**Start from an objective**

Say what needs fixing in plain language — it drafts the triggers, skills and actions.

02 Compose**Assemble the logic**

Wire up triggers, skills, reasoning steps, checkpoints and actions to match how the work actually runs.

03 Govern**Set the limits first**

Approval gates, access scope, policy checks and safety ceilings — fixed before anything executes.

04 Promote**Make it a standing agent**

Objectives, cost ceilings and safety rules travel with it, so it runs the same way every time.

— GOVERNED BEFORE IT ACTS

- **Approval gates.** Sensitive changes wait on a named human.
- **Access scopes.** Each agent sees only what its job needs.
- **Policy checks.** Checked on every run, not just the first.
- **Error budgets.** Past budget, work escalates — it doesn't retry blind.
- **Safety ceilings.** What runs unattended, set per environment.
- **Full audit trail.** Every call and outcome, reviewable after.

— START FROM A SPECIALIST, NOT A BLANK AGENT

- **Network.** Latency, routing and connectivity issues.
- **Compute.** Right-sizing and container risk.
- **Storage.** Disk health and cleanup targets.
- **Database.** Capacity and error patterns.
- **FinOps.** Spend anomalies and idle waste.
- **ITSM.** Ticket triage, enrichment, routing.

Put Agent Builder to work.

Describe an objective, watch it draft the workflow, and promote what works into a standing agent.

GET IN TOUCH
contact@lumiops.ai

PLAYGROUND
Try it, no setup

LINKEDIN
linkedin.com/lumiops