

— KNOWLEDGE, NOT A BLANK CANVAS

Turn how your team already 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 the limits you set — described in plain language, composed from familiar building blocks, and traceable from the very first run.

01 Describe

02 Compose

03 Govern

04 Promote

— ILLUSTRATIVE IMPACT

60%LESS MANUAL TRIAGE &
CLEANUP**45%**FASTER OBJECTIVE TO
FIRST RUN**3×**MORE AGENTS
PROMOTED PER
QUARTER**\$18**AVG. COST PER
AUTOMATED RUN

— WHAT'S INSIDE

Why teams build agents this way

Describe & Compose

Govern before it acts

Specialized agents

Observe & improve

Integrations

Enterprise readiness

Governed AI agents for enterprise IT operations — built from what your team already knows.

— WHY TEAMS BUILD THIS WAY

The fix already exists. It just lives in one person's head.

Most operational toil isn't a mystery — someone on the team already knows the fix. The problem is that the fix only runs when that person is awake, available, and willing to do it by hand again. Agent Builder takes the part they already know and turns it into something that runs the same way every time, inside limits the team decided on ahead of time.

Cut the repeat work

Investigation, enrichment, routing and cleanup that repeat every week, handled without someone re-deriving the fix from scratch each time.

Close things sooner

One flow carries a signal from first alert to an approved action — context, diagnosis, decision and execution, in the same place.

Build resilience that outlives one engineer

A standard response to a known failure mode doesn't need the one person who remembers how it was fixed last time.

Keep cost honest continuously

Waste, anomalies and idle spend get found and cleaned up on a schedule, not rediscovered once a quarter.

An agent that nobody can explain is not autonomy. It's a liability wearing a demo. Agent Builder is built around the opposite bet — that the fastest path to trustworthy automation runs through limits that were set on purpose, not discovered after something broke.

— WHERE THE EFFORT ACTUALLY GOES

Describing the objective

A sentence, not a spec document — the starting point most teams already have in a Slack thread.

Composing the logic

Wiring the triggers, skills, checkpoints and actions so the flow matches how the work really happens.

Setting the boundary

Deciding, once, what the agent may do alone and what needs a person — before anything runs.

— SIGNS A TASK IS READY TO BECOME AN AGENT

- Someone on the team can already explain the fix in a couple of sentences.
- The risky part of the fix is narrow and nameable — not the whole task.
- It comes up often enough that re-doing it by hand is a recognizable drain.
- A person would still want a say before the riskiest step executes.

— DESCRIBE

Start from an objective, not a blank canvas.

You don't design an agent from nothing. You describe what a good fix looks like, in the words you'd already use talking to a teammate, and the first draft — triggers, skills, decision logic, actions — comes back built around that objective.

```
describe "investigate a kubernetes crash loop"
# drafting workflow from objective...
  trigger: on-demand · pod CrashLoopBackOff event
  skills: fetch logs, describe pod, check recent deploys
  ai step: rank likely causes against last 90 days
  human step: confirm before restart in production
✓ draft ready · 4 steps · edit any block before you run it
```

— QUICK-START OBJECTIVES TEAMS ACTUALLY USE

Triage an incident and route it to the right team

Detect anomalous cloud spend

Diagnose infrastructure performance degradation

Investigate a Kubernetes crash loop

Find and clean up unused cloud disks

Rotate aging IAM keys with approval

— COMPOSE — THE BUILDING BLOCKS

01 Triggers

A schedule, a webhook, an ITSM event, or someone just asking.

02 Skills

The reusable moves an agent already knows how to make.

03 AI capabilities

Judgement grounded in the runbooks your org already trusts.

04 Flow controls

Handles the branch and the exception — the part that breaks a plain script.

05 Human steps

A person stays exactly where judgement matters, nowhere else.

06 Actions

The actual change, executed and logged — not just proposed.

— THE DRAFT IS A STARTING POINT, NOT A COMMITMENT

Nothing runs from a draft

Compose only assembles the logic — Govern decides what may actually execute.

Every block is editable

Swap a skill, add a branch, tighten a trigger — before or after the first run.

The objective can change

Re-describe it, and the workflow re-drafts around the new objective, not from zero.

— GOVERN

Set the boundary before anything runs.

Every capability here exists to answer one question in advance, not after the fact: what may this agent do without asking, and where does it have to stop and check?

Approval gates

Sensitive changes stop and wait for a named human decision before anything executes.

Access scopes & permission boundaries

Each agent sees only the systems and data its objective requires — nothing wider.

Policy checks & guardrails

Policy is evaluated on every run, so behaviour can't drift outside the agreed envelope.

Error budgets & escalation paths

Past its budget, work escalates to the right team instead of retrying blind.

Safety controls for autonomous actions

What an agent may do unattended, set per environment, widened as confidence grows.

Audit trails for every step

Every decision, tool call, approval and outcome recorded and reviewable afterward.

— WHAT A BOUNDARY LOOKS LIKE IN PRACTICE

```
agent: iam-key-rotation    environment: production
detect_aging_key: auto    (daily scan, no side effects)
draft_rotation: auto      (prepared, not applied)
apply_rotation: approval  (role: security-lead, evidence attached)
on_error: escalate to on-call, do not retry unattended
```

— THE QUESTION EACH CONTROL ANSWERS

CONTROL	QUESTION IT ANSWERS
Approval gate	Who has to say yes before this specific action runs?
Access scope	What is this agent physically able to reach?
Policy check	Does this action still match the rules, right now?
Error budget	At what point does retrying stop being the safe choice?

— WHO TYPICALLY REVIEWS THIS BEFORE LAUNCH

Security

Enumerated action classes, not permissions inherited from a service account.

Engineering

The reversible steps automated first, so the team feels value quickly.

Risk & compliance

Approval routing and the audit trail shown as configuration, not intent.

— PROMOTE

Turn a working flow into a standing agent.

A workflow that worked once is a good sign. A workflow that keeps working, with its objective, thresholds and safety rules attached, is what actually scales past the person who built it.

```
workflow v4 · promote to agent · ready
objective: reclaim unused block storage · defined
threshold: success rate & latency SLO · set
cost control: per-run budget ceiling · set
safety: approval on tier-1 assets · required
observability: tracing & audit · on
# reusable across teams · versioned · reversible
```

- Defined objectives and success criteria
- Performance thresholds and cost controls
- Safety requirements carried into every run
- Observability from the very first execution

— A WORKFLOW VS. A PROMOTED AGENT

PROPERTY	AD-HOC WORKFLOW	PROMOTED AGENT
Ownership	Whoever built it, informally	A named objective and owner
Cost	Unmeasured	Tracked against a ceiling
Trust	Assumed after it works once	Earned against a threshold
Reuse	Copy-pasted, drifts over time	Versioned, shared as-is

— WHEN A WORKFLOW IS READY TO PROMOTE

- It has run enough times to show a stable success rate, not just a lucky first try.
- Its cost per run is known, not estimated after the fact.
- The approval points have been exercised for real, not just configured.
- Someone other than the builder could pick it up and understand it.

— SPECIALIZED AGENTS

Start faster from a foundation, not a blank agent.

Six starting points built around where enterprise IT actually loses time and money — each one composable, and each one governed the same way as anything you build from scratch.

Network Specialist

Investigates latency, connectivity and routing issues across hybrid and multicloud infrastructure.

Compute Specialist

Finds utilization gaps, right-sizing opportunities and container-level operational risk.

Storage Specialist

Surfaces disk health concerns, cleanup opportunities and performance drift before workloads feel it.

Database Specialist

Investigates performance, capacity and error patterns to catch avoidable risk early.

FinOps Agent

Detects spend anomalies, finds optimization opportunities and automates cost cleanup.

ITSM Agent

Classifies, enriches and routes requests, and closes the loop with the service desk.

— FROM FOUNDATION TO FIT

1 Pick the closest specialist

The foundation that already matches the shape of the problem.

2 Compose in your own logic

Add the triggers, checkpoints and exceptions specific to your estate.

3 Govern and promote as usual

Same approval gates, same audit trail, same path to a standing agent.

— WHICH TEAM USUALLY STARTS WHERE

TEAM	USUAL STARTING SPECIALIST	TYPICAL FIRST OBJECTIVE
Platform / SRE	Network, Compute	Investigate a recurring latency or crash-loop pattern
Cloud & FinOps	FinOps Agent, Storage	Find and clean up idle disks or orphaned instances
Security	Network, custom skill	Rotate aging credentials with an approval gate
IT operations	ITSM Agent	Classify and route incoming service requests

— OBSERVE, GOVERN, IMPROVE

Autonomy you can explain after the fact.

Autonomy without accountability just moves the risk somewhere less visible. Agent Builder makes agent behaviour observable from the first run — what it did, how well, what it cost, and where it created value.

Measure performance

Execution speed, throughput, acceptance, completion quality and SLO attainment, per agent and across the fleet.

Enforce safety & governance

Approval requests, policy checks, guardrail events and every action executed, deferred or reverted.

Understand cost & efficiency

Spend across models, tokens, tools and infrastructure, compared against the outcomes it produced.

Trace every decision

Trigger through investigation, decision, tool call, approval, execution and outcome — the complete record.

— CONNECTS TO WHAT YOU ALREADY RUN

Observability & incident intelligence

Metrics, logs, monitors, alerts and entity context, so agents investigate with the same picture your team has.

Cloud & infrastructure

Query and act, with limits, across supported cloud, Kubernetes, infrastructure and database environments.

Engineering & IT operations

Source control, ticketing, collaboration, email, APIs and webhooks — the systems work already moves through.

Knowledge & operational context

Runbooks, playbooks, SOPs and internal docs, with private, team or org-wide visibility.

— UNDER THE HOOD

6

SPECIALIST FOUNDATIONS

4STAGES, START TO
STANDING AGENT**100%**RUNS LOGGED &
ATTRIBUTED**0**ACTIONS OUTSIDE THE
AGREED SCOPE

Model orchestration picks the right model per step, continuously evaluated, with reasoning that can stay entirely in-boundary where that matters.

— ENTERPRISE READINESS

Five things that have to hold together at once.

Moving past disconnected automation and one-off AI experiments takes more than a good demo. These five capabilities are what make governed agents something an enterprise can actually rely on.

01 Governed execution

Policies, approval gates, access scopes and controlled actions — enforced on every run, not just documented.

02 Operational transparency

End-to-end tracing and outcome visibility, so any action an agent took can be explained afterward.

03 Reusable operational intelligence

Workflows, skills and organizational knowledge that compound across teams instead of expiring with the incident.

04 Continuous optimization

Observability, cost measurement and safety tuning — autonomy extended on evidence, not assumption.

05 Enterprise-wide scale

One governed model across ITOps, SRE, cloud, FinOps and infrastructure teams — not five separate tools.

— WHAT EVERY PLAN INCLUDES

- Describe-to-draft workflow creation
- Six specialized agent foundations
- Full governance controls from day one
- A free playground — no setup required

Capabilities and integrations reflect LumiOps published materials for Agent Builder. Outcome figures referenced elsewhere in this brochure are illustrative and vary by estate size, objective scope and guardrail configuration. Third-party product names are trademarks of their respective owners.

— GETTING STARTED, IN THREE MOVES

1 Open the playground

No setup — describe an objective and see the first draft in minutes.

2 Compose with your team

Add the branches and checkpoints specific to how you actually work.

3 Govern, then promote

Set the boundary once, run it for real, promote it once it earns trust.

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