

— THE BUILDING BLOCKS EVERY AGENT IS MADE OF

Write the method once. Let everyone run it.

Skills Builder turns one engineer's best procedure into a versioned, governed skill — with steps, tool calls, constraints and an output schema — so the right method runs first, for anyone, every time.

32 built-in skills 10 domains 4 routing checks

01

Match

Relevance to the question

02

Check

Connector, state, RBAC

03

Route

Pick the right few

04

Invoke

Only then, tools fire

— PUBLISHED BENCHMARK RANGES

40-60%

LOWER MTTR

~90%

LESS ALERT NOISE

1/3 week

BACK FROM FIREFIGHTING

Day 1

NEW-HIRE RAMP

db-pool-diagnostics.skill

v2.3 · enabled

```
domain: database
persona: [sre, dba]
connectors: [postgres, datadog]
steps:
- pull active sessions + wait events
- correlate pool saturation vs latency
- rank blocking queries by impact
```

— WHAT'S INSIDE

Why skills

Anatomy of a skill

The router

Autonomy & guardrails

Catalog & marketplace

Versioning

Outcomes

Versioned, routed, governed skills — one engineer's best method, runnable by the whole team.

— WHY SKILLS

Your best engineer's method shouldn't need them in the room.

Every team has someone who can walk into a messy incident and know which three things to check first. That knowledge is the most valuable thing on the rota, and right now it only runs when that person is awake. Skills write it down in a form the whole team — and every agent — can actually execute.

— FOR THE ENGINEER, DAY TO DAY

Stop starting from zero

Every investigation opens with an expert method already loaded, not a blank terminal at 3am.

Fewer escalations to you

When the skill encodes your runbook, on-call doesn't need to phone the one person who knows.

Do the interesting work

Hand the tenth identical CrashLoop to a skill and get the afternoon back for architecture.

Confidence to act

Constraints and RBAC are enforced by the skill, so moving fast is not the same as moving recklessly.

— FOR THE ORGANISATION, ON THE BALANCE SHEET

Downtime costs less

Most enterprises put an hour of outage above \$300K. Halving MTTR takes that straight off the top.

Capacity without headcount

The same team covers more estate, because the repeatable half of the work no longer needs a person.

Expertise stops walking out

When your best engineer moves on, their method stays behind as a versioned skill.

Consistent, auditable operations

Same procedure every time, with a record — which is what auditors actually ask for.

A junior engineer directing a skill follows the same steps as your best senior. That's the whole idea.

— THE DIFFERENCE ON A REAL NIGHT

AN INCIDENT WITHOUT SKILLS

- On-call opens a blank terminal and a stale wiki page
- Tries the obvious check, then the second-obvious one
- Phones the expert at step four
- Writes up notes in whatever format they like

THE SAME INCIDENT WITH SKILLS

- The router loads the expert method before anything runs
- The right three skills fire first, not four
- The expert reviews the result in the morning
- Output lands in one schema, ready to audit

— ANATOMY OF A SKILL

A recipe, not a prompt.

A prompt is a suggestion a model may or may not follow the same way twice. A skill is a structured recipe with five parts, each one doing a specific job — which is why it behaves the same at 3am as it did in the demo.

01

Steps

The order an expert would work in, written down once.

02

Tool calls

Exactly which systems it may query, with which parameters.

03

Constraints

What it must never do. The guardrail ships inside the skill.

04

Output schema

A predictable shape other skills and workflows can consume.

05

Connectors

The integrations it depends on, checked before routing.

— WHAT ONE LOOKS LIKE

db-pool-diagnostics.skill

v2.3 · enabled

```
domain: database
persona: [sre, dba]
connectors: [postgres, datadog]
autonomy: discover
steps:
  - pull active sessions + wait events
  - correlate pool saturation vs latency
  - rank blocking queries by impact
constraints:
  # never terminate a session
  - read_only: true
  - max_rows: 5000
output: root_cause, evidence[], confidence, next_actions[]
```

— WHY THE SHAPE MATTERS

- Steps make it repeatable — the order an expert works in doesn't drift.
- Constraints make it safe — read_only means read only, every run.
- Tool calls make it bounded — nothing outside the list gets touched.
- The schema makes it composable — its output feeds the next skill cleanly.

— SKILL ROUTER

Four checks before any tool fires.

Routing isn't a similarity search that grabs the closest match and hopes. Every candidate skill has to pass four tests, and the ones that fail are dropped — visibly — before anything runs.

Relevance
embedding

Does the question actually match what this skill says it is for?

Availability
connector

Is the system this skill depends on connected right now?


Skill state
state

Enabled, drafted or deprecated — only live skills are routed.

Permission
RBAC

Your role decides what runs. Denied skills are dropped, not hidden.

— ONE QUESTION, NARROWED

32 Skills in the catalog

27 Available to this persona

3 Routed for the question

✓ Tools invoked

3 of 32

skills actually ran — the right three, first

— WHEN A SKILL IS FILTERED OUT

IF A CHECK FAILS...	WHAT THE ROUTER DOES
Relevance is low	The skill is never a candidate for this question.
Its connector is down	It is routed around; the next best live skill takes its place.
It is drafted or deprecated	It is skipped. Only enabled versions ever run.
Your role is not allowed	It is dropped and shown as denied — not silently hidden.

You can test which skills a query would select, and why, before it ever runs against production.

— WHAT A ROUTER TEST SHOWS YOU

```
router test "why is checkout p99 spiking in eu-west?"  
k8s-pod-diagnostics relevance 0.41 skipped — below threshold  
aws-investigation relevance 0.72 connector down — routed around  
latency-correlation relevance 0.91 routed  
db-pool-diagnostics relevance 0.88 routed  
deploy-timeline relevance 0.84 routed  
session-terminator relevance 0.66 denied by RBAC — shown, not hidden
```

— GOVERNANCE & GUARDRAILS

Every skill declares its own autonomy.

Classification lives on the skill itself, not on the agent that calls it. So an agent can never do more than its skills allow — the ceiling travels with the work, wherever it gets used.

1 discover. Read only. Investigates and reports; cannot change a thing, safe on any site.

Earned by: the default for every new skill

2 manage. Scoped change on tagged targets, within a blast radius, approval where required.

Earned by: a clean record in discover, plus tagged targets

3 automate. Runs unattended on its own.

Earned by: an acceptance rate that has actually earned the promotion

— CONSTRAINTS, WRITTEN INTO THE SKILL

```
constraints:
# travels with the skill, on every agent that calls it
- read_only: true
- max_rows: 5000
- never: [terminate_session, drop_index]
- requires_approval_on: [prod-tier-1]
```

Guardrails bolted onto an agent can be forgotten on the next agent. Guardrails written into the skill can't.

- New skills start in discover — nothing changes until it has a track record.
- Promotion to automate is based on acceptance rate, not on confidence.

— HOW REAL SKILLS END UP CLASSIFIED

EXAMPLE SKILL	CLASS	WHY
db-pool-diagnostics	discover	Reads sessions and metrics; never touches a connection
cache-flush (tagged services)	manage	Changes tagged targets only, approval on checkout
state-snapshot-cleanup	automate	Earned it after months of accepted runs in manage
memory-leak-detection	discover	New from the marketplace — starts at the bottom

— CATALOG, WORKSPACE, MARKETPLACE

Three places a skill comes from.

You rarely need to write a skill from nothing. Start from the built-in catalog, clone and tune it in your workspace, or install a specialist skill when Lumi notices a gap.

Company catalog**Built in by UnityOne**

32 production-tested skills across every domain, versioned and maintained for you.

Workspace**Yours and your team's**

Clone a built-in skill and tune it, or write a new one straight from your own SOP.

Marketplace**Installed on demand**

Specialist skills added when Lumi spots a gap — recommended, not guessed at.

— ORGANISED BY DOMAIN, FILTERED BY PERSONA

DOMAIN	THE KIND OF WORK SKILLS COVER
Compute & Kubernetes	Pod diagnostics, crash loops, memory pressure, right-sizing
Network	Latency correlation, routing and connectivity checks
Storage & database	Pool saturation, blocking queries, disk and volume health
Cloud	Account investigation, service health, SKU and cost context
Identity & security	Access review, policy drift, suspicious activity triage
Observability & services	Deploy timelines, alert correlation, dependency tracing

— WHAT AN AI RECOMMENDATION LOOKS LIKE

Add Memory Leak Detection

Memory-leak patterns fired 14x last week and the closest skill only managed 76% acceptance. Installing the specialist is expected to add ~12 points.

Refresh AWS few-shot examples

Calibration drifted 14 points in 30 days — newer instance families were missing. Promoting the drafted v4 is expected to lift it from 74% to 86%.

— A TYPICAL WORKSPACE, AT A GLANCE

32

COMPANY CATALOG

2

IN THIS WORKSPACE

6

MARKETPLACE, READY

27 / 32

FIT THE SRE PERSONA

— GOOD TO KNOW

- Built-in skills are maintained for you as platforms change.
- Clones keep a link to the original, so you can see what you changed.
- Marketplace installs arrive with an expected-gain estimate.
- Workspace skills stay private until you share them. not a guess.

— VERSIONING & LIFECYCLE

Draft it, promote it, roll it back.

Skills change as your estate changes. Every version carries its own calibration, so you can see whether an edit made a skill better before it replaces the one everyone is using.

```
aws-investigation version history
v2 retired calibration 71%
v3 enabled calibration 74% # drifting — new SKUs missing
v4 drafted calibration 86% on replay

action: promote v4 → enabled (v3 kept for one-click rollback)
```

— EVERYTHING SKILLS GIVE YOU

Built-in catalog

32 composable skills across ten operational domains.

Persona-based

Filter to SRE, DBA, NOC or platform — see only what fits.

AI recommendations

Spots missing skills and drifting calibration, proposes the fix.

Cloning & customisation

Fork any skill, edit steps and constraints, keep the original.

Versioning

Draft, promote and roll back, with calibration tracked per version.

Skill Router

Test which skills a query selects, and why, before it runs.

Connectors

Each skill declares its integrations; unavailable ones are routed around.


Guardrails

Constraints live in the skill, so it only does authorised work.

Custom from SOPs

Turn your own runbooks into skills the whole team can direct.

— CALIBRATION, TRACKED PER VERSION



Illustrative history for one skill. Drift from 88% to 74% is what an estate changing underneath a skill looks like.

— HOW VERSIONS STAY HONEST

- Drafts replay against real past incidents before promotion.
- Every version keeps its own calibration score.
- Rollback is one step — the previous version is never deleted.
- Drift raises a recommendation, not a surprise.

— WHY IT MATTERS

The same work, done faster — and cheaper.

Skills turn one engineer's best method into something the whole team can run, every time, without waiting for the person who knows.

— BEFORE AND AFTER, BY THE NUMBERS

MEASURE	TYPICAL STARTING POINT	WITH SKILLS
Mean time to resolve	~4 hours	40-60% lower — the right method runs first, not fourth
Alert noise	~900 alerts a day	~90% less — correlation skills collapse floods
Time on firefighting	About a third of the week	Back to building the roadmap
New-hire ramp	3-6 months	Day one — the skill carries the method

40-60%

LOWER MTTR

~90%

LESS ALERT NOISE

1/3 week

BACK FROM FIREFIGHTING

Day 1

NEW-HIRE RAMP

Ranges reflect published AIOps and ITOps benchmarks — ITIC, Forrester, New Relic, 2024-2026. Your results depend on your estate and starting point.

— FILTERED TO THE PEOPLE WHO USE IT

SRE / DevOps

Latency, deploys, pods

DBA

Pools, locks, queries

NOC

Network and service health

Platform

Cloud, capacity, cost

— GETTING STARTED, IN THREE MOVES

1**Start from the catalog**

Pick the built-in skill closest to the work you repeat most.

2**Clone and constrain**

Tune the steps to your estate and write the limits in.

3**Test the router, then promote**

Check what a query selects, run it in discover, earn the next class.

Over 90% of enterprises put an hour of downtime above \$300K. Cutting MTTR in half is not an efficiency story — it's a budget line.

**Turn your best runbook
into a skill.**

Start from a built-in skill, clone it, and make it match how your team actually works.

GET IN TOUCH

contact@lumioops.ai

CATALOG

32 built-in skills

LINKEDIN

linkedin.com/lumioops