

AgentForge Operator Guide

Run workflows and audit spend on the localhost:8001 dashboard

Run workflows and read agent fleet spend at a glance.

Operator tasks on the localhost:8001 dashboard, not the report-studio CLI.

AUDIENCE AgentForge operators · platform engineers

Executive summary

GUI WALKTHROUGH Annotated screenshots	3 OPERATOR FLOWS Dashboard · workflows · transactions	Live SOURCE localhost:8001 console	Vol 5 DEEP DIVE AgentForge architecture
---	---	--	---

This guide walks the AgentForge dashboard at localhost:8001, the operator console for the agent fleet. Each section pairs operator steps with an annotated screenshot: text first, image second. AgentForge runs and audits workflows across projects, and the dashboard surfaces invocation counts, spend, approvals, and third-party transaction health. Captures refresh via `scripts/capture-platform-ui-screenshots.mjs`.

Key findings

- 1 Operator tasks on the AgentForge dashboard, not the report-studio CLI.
- 2 Hotspots call out the health, spend, run, and audit controls operators use.
- 3 Next steps link to Vol 5 architecture and Vol 10 operations.

Who should read this

Operators running and auditing the agent fleet on localhost:8001. Engineers continue to Vol 5 for the workflow engine after this how-to pass.

SECTION I

Overview

Who this guide is for and what you will produce.

This guide walks the AgentForge dashboard at localhost:8001, the operator console for the agent fleet. Each section pairs operator steps with an annotated screenshot: text first, image second. AgentForge runs and audits workflows across projects, and the dashboard surfaces invocation counts, spend, approvals, and third-party transaction health. Captures refresh via scripts/capture-platform-ui-screenshots.mjs.

Methodology

This brief is generated from the live NLT document suite (nlt-report-studio v0.2.0, brand nlt-v3.2). Diagrams and tables reflect build-time workspace discovery where noted.

Audience

- Investors & board — Vol 1-2
- Architects — Vol 3-4
- Engineering DD — Vol 5-8

Related proof

[portfolio.nltlabs.ai](#) · [briefs.nltlabs.ai](#)

SECTION II

Open the operator dashboard

Open the AgentForge dashboard and read fleet health and spend.

Follow the steps below, then confirm the dashboard state matches the screenshot.

- Open `http://localhost:8001/` in a desktop browser; the dashboard reads the running backend.
- Confirm the Connected badge is green, then set the window with the 7d, 30d, or 90d toggle.
- Read the KPI cards: Agents, Pending Approvals, Invocations, Total Cost, and System Health.
- Note the auto-approve banner; switch to Require Approval in Settings when manual review is needed.

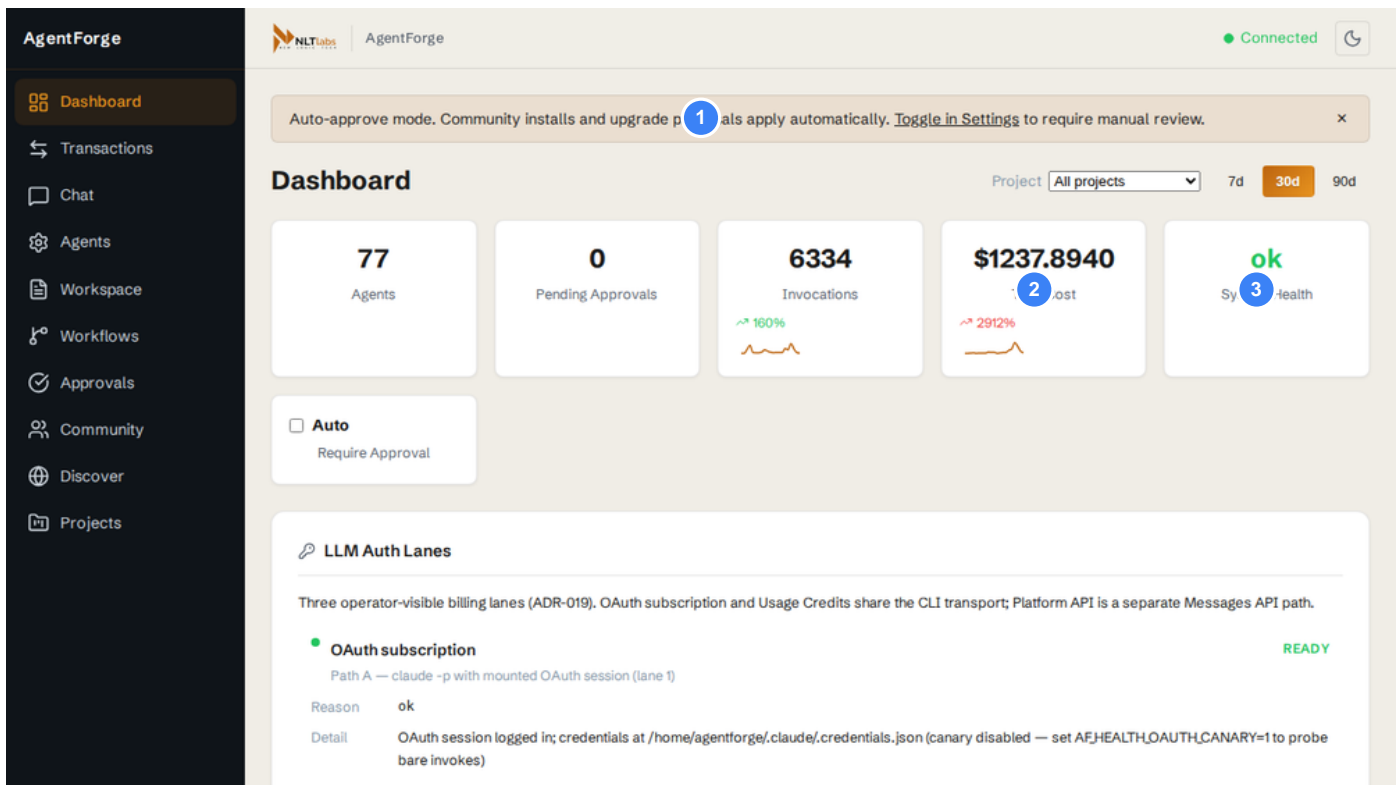


Figure 1 — AgentForge dashboard: fleet KPIs, auto-approve banner, and LLM auth lanes.

- 1 — Auto-approve banner; switch to Require Approval in Settings before trusting auto-installs.
- 2 — Total Cost is running spend for the selected window (7d/30d/90d toggle, top right).
- 3 — System Health: ok means the invoke path and auth lanes are live.

SECTION III

Run a workflow

Open the Workflows tab to browse and run a registered workflow.

Work the steps in order; each maps to a numbered hotspot on the screenshot.

- Select Workflows in the left navigation.
- Find a workflow card and read its node count and last-run status (COMPLETE, PARTIAL, CANCELED, DRY RUN).
- Press Run to launch the workflow; a workflow with no runs yet shows No runs yet.
- Return to Dashboard to confirm the Invocations and Total Cost KPIs moved after the run.

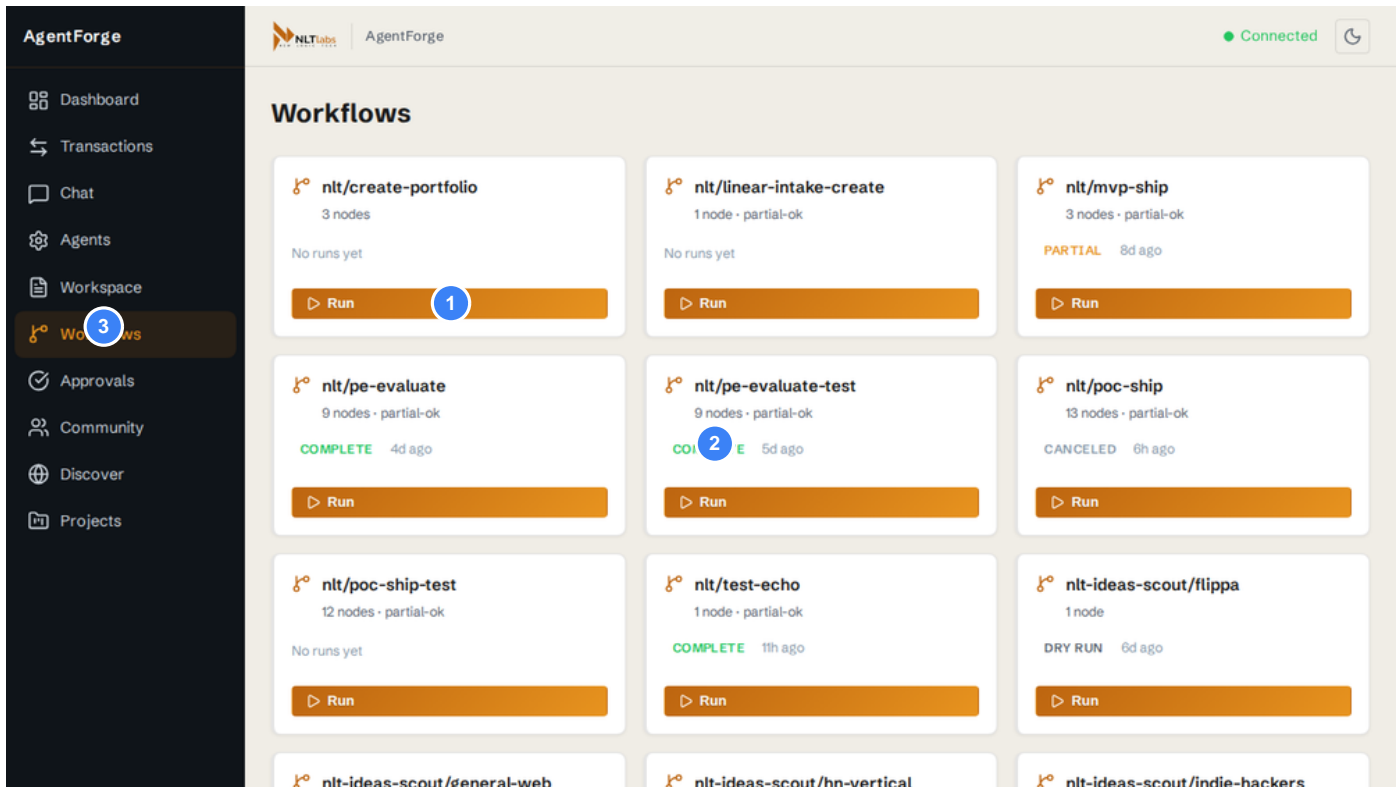


Figure 2 — Workflows tab: the registered workflow catalog with run controls.

- 1 — Run launches the workflow; node count and last-run status sit above the button.
- 2 — Status chip shows COMPLETE, PARTIAL, CANCELED, or DRY RUN with the last-run age.
- 3 — Workflows tab lists every registered graph across projects (nlt/*, nlt-ideas-scout/*).

SECTION IV

Audit third-party spend

Open the Transactions tab to audit third-party calls and spend.

Read each signal below and investigate any non-zero error count before trusting the run.

- Select Transactions in the left navigation.
- Read HTTP Calls, Search / Fetch, and MCP Web Tools for outbound vendor volume; taps-brain traffic is excluded.
- Check HTTP Errors and Avg HTTP Latency, and investigate any non-zero error count first.
- Confirm Provider Health shows each vendor (for example anthropic: ok) before trusting the run.

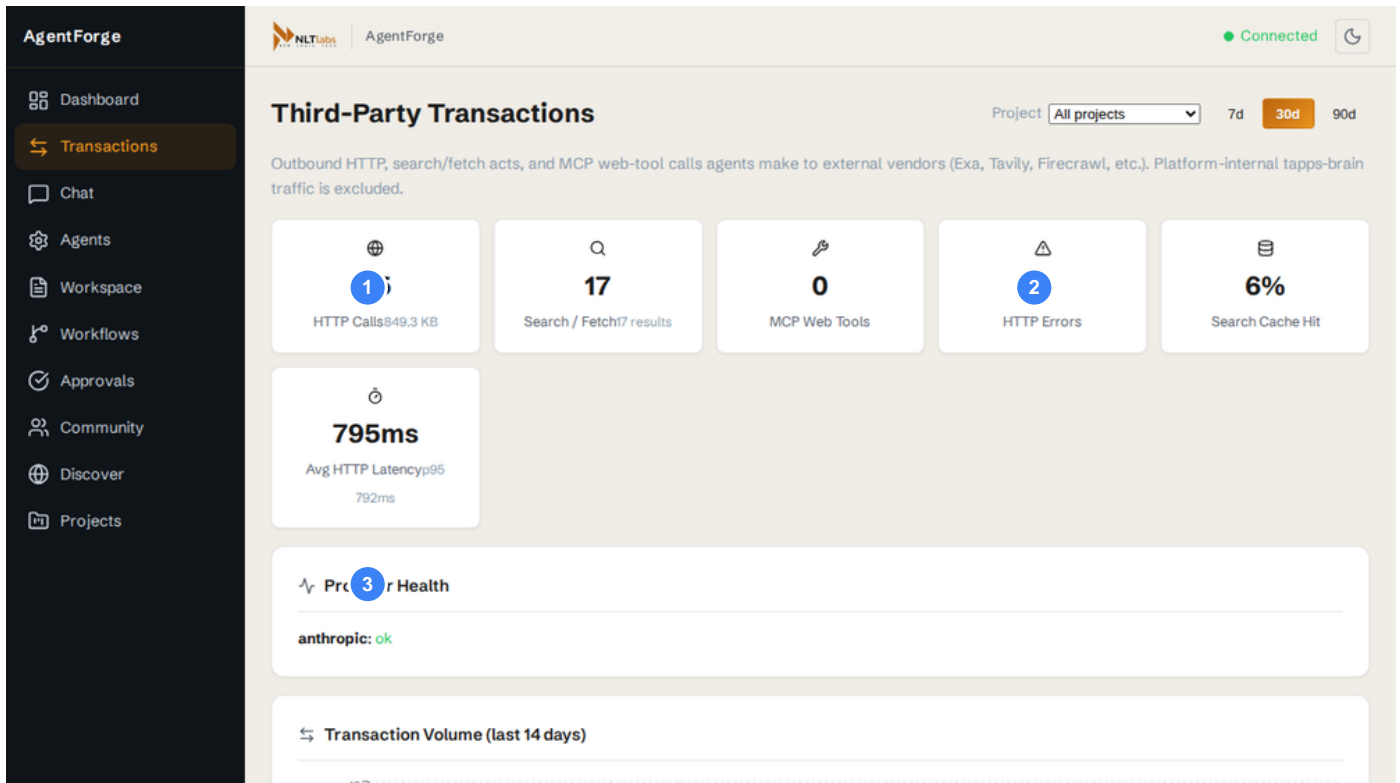


Figure 3 — Third-Party Transactions: outbound call volume, error rate, latency, and provider health.

- 1 — HTTP Calls and payload size: outbound vendor traffic for the window (taps-brain excluded).
- 2 — HTTP Errors: a non-zero count here is the first signal to investigate a failing provider.
- 3 — Provider Health: per-vendor status (anthropic: ok) for the LLM and search providers.

SECTION V

Next steps

Where to go after your first workflow.

- Deep dive: AgentForge architecture is Vol 5 of the suite (docs/NLT_DOCS_SUITE.md).
- Operations: Vol 10 covers fleet runbooks, approvals, and cost controls.
- Other tabs: Agents, Approvals, Projects, and Community extend the operator console.
- Refresh UI shots: node scripts/capture-platform-ui-screenshots.mjs (needs localhost:8001 live).

Engineer path

Engineers extending AgentForge should read Vol 5 for the workflow engine and audit internals, and use the report-studio CLI to regenerate these operator guides.

Confidential — prepared for NLT Labs stakeholders. Outward names follow docs/NLT_BRAND.md. Internal checkout names appear only in builder ops volumes (9-10).