

- 1 GitHub Agentic Workflows Rollout Kit (Technical Preview)
  - 1.1 Who This Is For
  - 1.2 What Changed (Why This Matters)
  - 1.3 Operating Principles
  - 1.4 Recommended Rollout Path (4 Weeks)
  - 1.5 Phase 0 (Day 0-2): Secure the Baseline
  - 1.6 Phase 1 (Week 1): Issue Triage Pilot
  - 1.7 Phase 2 (Week 2-3): CI Failure Investigator
  - 1.8 Phase 3 (Week 4+): Reporting and Selective PR Creation
  - 1.9 Use Cases: Good Fit vs Keep Deterministic
    - 1.9.1 Good Fit (Agentic)
    - 1.9.2 Keep Deterministic (Standard GitHub Actions)
  - 1.10 Security Guardrails Checklist (Must-Haves)
  - 1.11 Cost Controls (Preview Reality)
  - 1.12 Quick Start (CLI)
  - 1.13 Starter Template 1: Issue Triage (Low Risk, High Value)
  - 1.14 Starter Template 2: CI Failure Investigator (Start with Issue Creation Only)
  - 1.15 Starter Template 3: Weekly Repository Health Report (Batch the Signal)
  - 1.16 Pilot Scorecard (4-Week Evaluation)
    - 1.16.1 Accuracy and Usefulness
    - 1.16.2 Operational Impact
    - 1.16.3 Risk and Control
    - 1.16.4 Cost
  - 1.17 Review Checklist for New Agentic Workflows
  - 1.18 Rollback / Kill Switch Procedure
  - 1.19 Common Failure Modes (and How to Avoid Them)
    - 1.19.1 1) Starting with PR creation
    - 1.19.2 2) Broad triggers
    - 1.19.3 3) No quality metrics
    - 1.19.4 4) Treating prompt injection as hypothetical
  - 1.20 Decision Rule: Scale, Hold, or Stop
  - 1.21 Source Verification (Checked on 2026-02-24)
  - 1.22 Final Notes

# 1 GitHub Agentic Workflows Rollout Kit (Technical Preview)

Version: 2026-02-24 Author: Mathieu Kessler Site: <https://www.talk-nerdy-to-me.com>  
Companion article: <https://www.talk-nerdy-to-me.com/blog/github-agentic-workflows-continuous-ai>

## 1.1 Who This Is For

This playbook is for platform engineers, DevOps teams, and engineering leaders who want to pilot GitHub Agentic Workflows safely and pragmatically.

It is designed for:

- Teams that already use GitHub Actions and want to add AI-assisted repository automation
- Teams that care about security guardrails and reviewability
- Teams that need a phased rollout plan (not a hype-driven “replace all YAML” migration)
- Teams that want measurable success criteria before scaling

It is not designed for:

- Replacing deterministic build/test/deploy jobs with agents
- Autonomous production changes
- Auto-merge pipelines controlled by AI

## 1.2 What Changed (Why This Matters)

GitHub Agentic Workflows (technical preview announced on February 13, 2026) introduces a new workflow authoring model:

- You author the workflow as Markdown with YAML frontmatter
- The `.md` file is the source of truth
- `gh aw compile` generates a hardened `.lock.yml` GitHub Actions workflow
- The lock file runs the selected coding agent (Copilot, Claude Code, Codex, or compatible custom engine)
- Agent writes are constrained through `safe-outputs`

This is useful for tasks that require judgment, not just deterministic execution:

- issue triage
- CI failure investigation
- documentation drift correction
- reporting and monitoring summaries
- bounded refactoring suggestions

Keep deterministic CI/CD jobs in normal GitHub Actions YAML.

## 1.3 Operating Principles

1. Guardrails before convenience
2. Start with low-blast-radius workflows
3. Measure quality before scaling
4. Keep humans in the review loop
5. Treat agent-readable content as untrusted input

## 1.4 Recommended Rollout Path (4 Weeks)

### 1.5 Phase 0 (Day 0-2): Secure the Baseline

Checklist:

- Confirm branch protections and required reviews are enabled
- Identify pilot repository owners and on-call maintainers
- Limit pilot scope to one repo with active maintainers
- Document rollback steps (disable workflow + revert `.md` / `.lock.yml`)
- Define a weekly pilot review cadence (30 minutes is enough)
- Require `gh aw audit` before enabling any workflow

Do not start with:

- PR creation workflows
- org-wide rollout
- broad triggers (push on many branches, all workflow runs, etc.)

### 1.6 Phase 1 (Week 1): Issue Triage Pilot

Why first:

- High maintainer value
- Low blast radius
- Easy to review output quality (labels/comments)
- Good environment for prompt tuning

Success criteria (week 1):

- Triage precision (human agrees with classification/label)  $\geq 85\%$
- Clarification comments useful  $\geq 70\%$
- False-positive duplicate suggestions low enough that maintainers still trust outputs

### 1.7 Phase 2 (Week 2-3): CI Failure Investigator

Why second:

- High leverage when CI failures are noisy
- Clear evidence sources (logs, recent commits, workflow metadata)
- Easy to evaluate actionability of output

Guardrails:

- Start with `create-issue` safe-output only (not PR creation)
- Require confidence ratings (high / medium / low)
- Require evidence references (log snippet, failing job, commit hash)
- Keep the trigger scoped to one CI workflow on `main`

## 1.8 Phase 3 (Week 4+): Reporting and Selective PR Creation

Add only after quality is proven:

- weekly repository health reports
- documentation drift proposals
- narrow-scope refactoring PRs

Before enabling agent-created PRs:

- Verify CODEOWNERS coverage
- Verify required review rules
- Verify reviewer response SLAs
- Verify rollback procedure is exercised

## 1.9 Use Cases: Good Fit vs Keep Deterministic

### 1.9.1 Good Fit (Agentic)

- Issue triage and support routing
- Duplicate detection and clarification comments
- CI failure investigation and evidence summarization
- Weekly project/repo health reporting
- Documentation maintenance and drift checks
- Bounded refactoring suggestions for human review

### 1.9.2 Keep Deterministic (Standard GitHub Actions)

- Build pipelines
- Unit/integration test execution
- Packaging and release automation
- Deployments to staging/production
- Secret rotation or credential management
- Infrastructure mutation without human approval

## 1.10 Security Guardrails Checklist (Must-Haves)

Use this checklist for every new Agentic Workflow.

- Minimal frontmatter permissions (avoid broad writes)
- Agent runs read-only unless a safe-output is declared
- safe-outputs restricted to exactly what is needed (labels/comments/issues first)
- Narrow allowlists (labels, title prefixes, targets)
- Review generated `.lock.yml` in every PR (it is executable)
- Run `gh aw audit` before enabling and after changes
- Keep branch protections and required reviews enabled
- Treat issues/PR bodies/comments/commit messages as untrusted input
- Review unexpected output behavior immediately and disable if needed

- Re-verify docs and CLI behavior after upgrades (technical preview changes happen)

## 1.11 Cost Controls (Preview Reality)

GitHub's launch materials document two cost layers:

- GitHub Actions usage (minutes/storage)
- Model usage (GitHub Copilot plan consumption on GitHub.com, or BYO provider billing for self-hosted Claude/Codex setups)

Control cost with these tactics:

- Start with narrow triggers (`issues.opened`, one `workflow_run`)
- Prefer scheduled summaries when batching is acceptable
- Delay PR creation until value is proven
- Benchmark engines by workflow type (triage vs CI diagnosis may need different models)
- Track "cost per useful output," not just total spend
- Set quotas/limits for self-hosted API keys

## 1.12 Quick Start (CLI)

```
# Install the Agentic Workflows CLI extension
gh extension install github/gh-aw

# Authenticate (GitHub and/or model providers depending on engine)
gh aw auth

# Add a starter workflow from the official examples repo
gh aw init githubnext/agentics#issue-triage

# Compile Markdown source to a hardened .lock.yml workflow
gh aw compile .github/workflows/issue-triage.md

# Audit before enabling
gh aw audit .github/workflows/issue-triage.md

# Dry run locally against a real event (replace 123)
gh aw run .github/workflows/issue-triage.md "issues.123" --watch

# Inspect execution logs/artifacts
gh aw logs .github/workflows/issue-triage.md
```

Notes:

- Re-check `gh aw --help` and the CLI reference before rollout; preview commands/options can change.
- Review both the `.md` source and generated `.lock.yml` in code review.

## 1.13 Starter Template 1: Issue Triage (Low Risk, High Value)

```

---
on:
  issues:
    types: [opened]
permissions: read-all
safe-outputs:
  add-labels:
    labels:
      - bug
      - enhancement
      - question
      - needs-triage
  add-comment:

```

### # Issue Triage Agent

Classify each newly opened issue and help maintainers reduce intake noise.

#### ## Goals

1. Determine if this is a bug, enhancement request, or question.
2. Add one primary label from the allowlist.
3. If key details are missing (repro steps, version, logs), add a short comment asking for the missing information.
4. If a likely duplicate exists, mention the matching issue number in the comment.

#### ## Rules

- Be concise and neutral.
- Do not speculate about root cause unless evidence is explicit.
- If confidence is low, label only as ``needs-triage`` and ask a clarifying question.

## 1.14 Starter Template 2: CI Failure Investigator (Start with Issue Creation Only)

```

---
on:
  workflow_run:
    workflows: ["CI"]
    types: [completed]
    branches: [main]
permissions:
  actions: read
  contents: read
  issues: read

```

```

pull-requests: read
safe-outputs:
  create-issue:
    title-prefix: "[ci-failure] "
    labels:
      - ci-failure
      - needs-investigation
---
```

### # CI Failure Investigator

Investigate failed CI runs on main and create an issue with a human-reviewable summary.

#### ## Tasks

1. Read failing job logs and identify the first actionable error.
2. Review recent commits likely related to the failing job.
3. Classify the failure as one of:
  - regression
  - flaky test
  - infrastructure/tooling issue
4. Create an issue with:
  - affected workflow/job
  - suspected root cause
  - confidence (high/medium/low)
  - suggested next action (fix, rollback, rerun, or escalate)

#### ## Safety

- Never propose secrets handling changes.
- Do not invent file paths or commit hashes not present in the evidence.

## 1.15 Starter Template 3: Weekly Repository Health Report (Batch the Signal)

```

---
on:
  schedule:
    - cron: "0 14 * * 1"
  workflow_dispatch:
permissions: read-all
safe-outputs:
  create-issue:
    title-prefix: "[weekly-repo-health] "
    labels:
      - reporting
      - repo-health
---
```

### # Weekly Repository Health Report

Produce a concise weekly report for maintainers.

**Include:**

- New issues opened/closed this week
- PR throughput (opened, merged)
- Top recurring CI failure themes
- Documentation drift signals (if visible)
- Recommended follow-ups (max 5)

Keep the report factual. Separate observations from recommendations.

## 1.16 Pilot Scorecard (4-Week Evaluation)

Track these during the pilot and review weekly.

### 1.16.1 Accuracy and Usefulness

- Triage precision (human agrees with label/classification): target  $\geq 85\%$
- Clarification comments judged useful: target  $\geq 70\%$
- CI investigation issues rated “actionable”: target  $\geq 75\%$

### 1.16.2 Operational Impact

- Maintainer time saved in intake triage (weekly estimate)
- Time to identify likely CI root cause (before vs after)
- Percentage of agent outputs requiring manual correction

### 1.16.3 Risk and Control

- Workflows passing `gh aw audit`: target 100%
- Incidents caused by unsafe agent behavior: target 0
- Runs with over-broad permissions/triggers: target 0

### 1.16.4 Cost

- Actions minutes per workflow type
- Model/API spend (or Copilot consumption signal)
- Cost per useful output (comment, issue, report, PR proposal)

## 1.17 Review Checklist for New Agentic Workflows

Use this in PR review before enabling a workflow.

- What exact human pain point does this workflow reduce?
- Is the trigger scoped to the minimum useful event set?
- Are permissions minimal?
- Are safe-outputs minimal and allowlisted?
- Is the output format easy for humans to review quickly?
- Is rollback documented?

- Are success metrics defined?
- Has `gh aw audit` been run and attached to the PR?

## 1.18 Rollback / Kill Switch Procedure

If output quality degrades or behavior is unexpected, disable first and investigate second.

```
# Disable a workflow quickly  
gh aw disable .github/workflows/issue-triage.md  
  
# Re-audit after changes  
gh aw audit .github/workflows/issue-triage.md  
  
# Re-enable when ready  
gh aw enable .github/workflows/issue-triage.md
```

Operational note:

- You can also disable the generated GitHub Actions workflow in the repository UI if needed.

## 1.19 Common Failure Modes (and How to Avoid Them)

### 1.19.1 1) Starting with PR creation

Problem: - Too much trust too early

Fix: - Start with comments/labels/issues only

### 1.19.2 2) Broad triggers

Problem: - Cost spikes - Noise - Poor maintainability

Fix: - Start narrow, measure, then expand

### 1.19.3 3) No quality metrics

Problem: - Team “feels” progress but cannot prove value

Fix: - Define precision/actionability/cost metrics before the pilot starts

### 1.19.4 4) Treating prompt injection as hypothetical

Problem: - Unsafe assumptions about untrusted content

Fix: - Assume hostile input is possible and rely on guardrails (permissions, safe-outputs, sanitization, review)

## 1.20 Decision Rule: Scale, Hold, or Stop

At the end of the pilot:

Scale if:

- Accuracy is strong
- Maintainers report real time savings
- Cost is acceptable
- No control failures

Hold and tune if:

- Value is real but precision/actionability is inconsistent

Stop if:

- Maintainer overhead increases
- Costs exceed value
- Workflow outputs are noisy or untrustworthy
- Guardrails are repeatedly bypassed or misconfigured

## 1.21 Source Verification (Checked on 2026-02-24)

Official sources used for this playbook:

- GitHub Changelog (technical preview launch): <https://github.blog/changelog/2026-02-13-agentic-workflows-a-new-way-to-automate-repository-tasks-in-public-preview/>
- GitHub Launch Blog (billing, guardrails, examples): <https://github.blog/ai-and-ml/github-copilot/automate-repository-tasks-with-github-agentic-workflows/>
- Agentic Workflows docs home: <https://github.github.io/gh-aw/>
- Workflow structure docs: <https://github.github.io/gh-aw/concepts/workflow-structure/>
- Security model docs: <https://github.github.io/gh-aw/concepts/security-model/>
- CLI reference: <https://github.github.io/gh-aw/reference/cli/>
- AI engines docs: <https://github.github.io/gh-aw/concepts/ai-engines/>
- Safe outputs reference: <https://github.github.io/gh-aw/reference/frontmatter/safe-outputs/>
- gh-aw repository (MIT): <https://github.com/github/gh-aw>
- AWF repository: <https://github.com/githubnext/awf>

## 1.22 Final Notes

- Agentic Workflows is technical preview software. Re-verify commands and defaults before production rollout.
- The `.md` file is your intent; the `.lock.yml` file is executable infrastructure. Review both.

- Keep the human review loop. GitHub states PRs created by Agentic Workflows are not auto-merged.