KubeCon | CloudNativeCon
North America 2026

Cadence Workflow ContribFest: Build Durable AI Agent Examples

Salt Lake City, Utah, United States





Diana Zawadzki Shijie Sheng

Start here

Event guide

Setup instructions

Warm-up exercise

Issue pool



Star the repository

cadence-workflow/cadence-ai-samples

How the repository is organized

recipes/	Runnable examples and their runbooks
contribfest/	Session setup, warm-up and issue pool
CONTRIBUTING.md	Human contribution and claim process
AGENTS.md	Coding-agent contract and context map

What you will learn

Temporary rate limit	A bounded Activity retry policy
Permanent 404	A clear stop without repeated attempts
Recorded results	Progress that survives Worker recovery
Human decisions	Signals for review and approval

How to claim an available task

1. Choose an open, unassigned status:ready issue.

2. Comment with your handle and proposed branch name.

3. Wait for recorded ownership and status:in-progress.

4. Then create issue-N-short-slug in your fork.

While you wait: get set up

Python 3.12 or 3.13. Cadence and Ollama running.

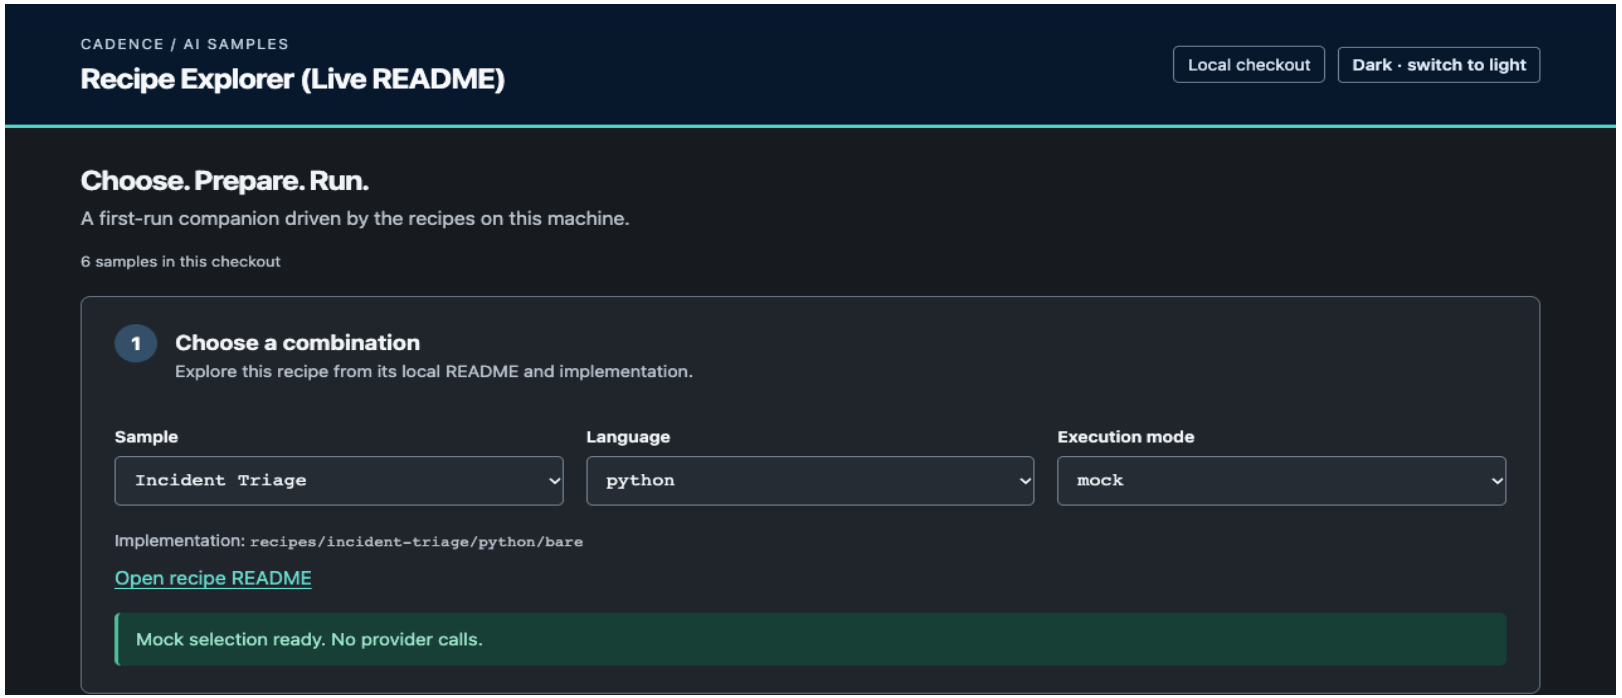
```
ollama pull qwen2.5:7b
```

```
cd recipes/incident-triage/python/bare
./run.sh setup
.venv/bin/python -m unittest discover -s tests
```

Setup instructions

Optional: the Samples Explorer

```
./scripts/explore.sh
```



Recipe docs
Service checks
Run commands

Terminal or
browser.
Your choice.

Help: #cadence-contributors on CNCF Slack

Today

1 Welcome and setup check

2 Warm-up together, about 20 minutes

3 **Your workflow or a contribution**

4 Share results and hand off

Meet Incident Triage

A fictional checkout failure

Synthetic incident. Real model calls.

Activities make the model and tool calls.

The Workflow coordinates them and waits for review.

Step 1: watch it investigate

```
.venv/bin/python exercises/retry_policy_live/run_exercise.py
--scenario baseline
```

Open the printed Cadence Web Run link.
Read the model action after each SPEND line.
The draft waits for your review.

Type reject to finish.

Step 2: see it fail

```
.venv/bin/python exercises/retry_policy_live/run_exercise.py
--scenario rate-limit
```

The gateway returns HTTP 429.
The starter makes no retry.

gateway_logs_unavailable
Zero model calls.

Step 3: your fix

Edit exercise_options.py: SCAN_ACTIVITY

Wait at least 1 second.
Use a finite number of attempts.
Do not retry permanent failures.

Cadence Activity retry documentation

Step 4: run both scenarios

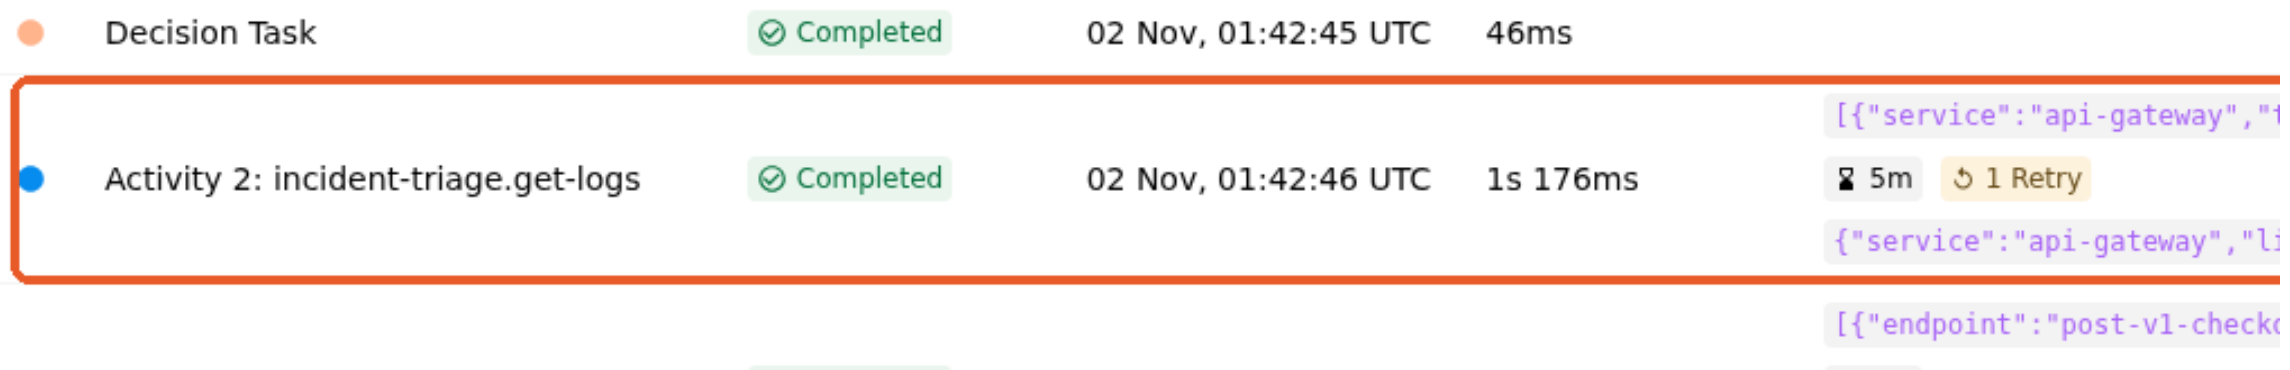
```
.venv/bin/python exercises/retry_policy_live/run_exercise.py
--scenario rate-limit
```

```
.venv/bin/python exercises/retry_policy_live/run_exercise.py
--scenario missing-resource
```

Step 4: check both outcomes

429: (0, 429), then (1, 200)
Graph once. Review reached.

404: (0, 404), then stop. No planning.
Three 404 attempts = wrong.



Optional: guardrail and review

With the same run_exercise.py runner:

`--scenario repeat-guard`

The Workflow blocks the third identical call.

`--scenario review`

revise: rewrite from collected evidence
guidance: investigate again

Now build: choose a path

Bring your own workflow, #118

Explore and file findings, #119

Take an available recipe task

Propose an exercise extension

Bring your own workflow

Wrap	Run your agent in an Activity started by a Workflow.
Break it	Make the riskiest call fail on purpose.
Make it durable	Separate the call. Add a retry policy and timeout.
Steer, stretch	Use a Signal for human approval.

Stop at any depth.

Recipe tasks

#110 Test that revision reuses evidence and guidance collects more

Signals, Queries and Workflow tests. Python only.

#112 Fix the Incident Triage clean-copy check

Self-contained samples. Python only.

Available when open, unassigned and status:ready.

Recipe tasks

#121 Test that a lost reply does not save the conclusion twice

Idempotent writes. Python only.

#111 Demonstrate recovery from a rate-limited model call

Activity retries. Cadence, Ollama and qwen2.5:7b.

Available when open, unassigned and status:ready.

Recipe tasks

#120 Demonstrate expiration of an unanswered triage approval

Durable timers. Cadence, Ollama and qwen2.5:7b.

#109 Help the triage model identify the deepest observed error

Human review. Cadence, Ollama, llama3.2 and qwen2.5:7b.

Available when open, unassigned and status:ready.

Stuck: setup or model call

SETUP FAILED

Follow the printed fix. Ask a helper.

MODEL CALL FAILED

Keep the output and show a helper.
Alternative: --model-id llama3.2-local

Stuck: local connections

Optional Explorer: port busy

```
./scripts/explore.sh --port 0
```

Optional Explorer: Ollama warm-up fails

Check ollama list, then pull the missing model.

Query times out: start the matching Worker.

Finish or hand off

A draft PR

git commit -s and validation evidence

Or an issue handoff comment

Fork and branch. Changed paths. Commands and results.
Untested work, blockers and the exact next step.

This week with Cadence

Thu Nov 12, 2:40–3:10 PM
Agents Fail. Workflows Recover.
Salt Palace, Ballroom ACE

Earlier: Tue Nov 10, 3–7 PM
Maintainer kiosk P-15B

#cadence-contributors

Samples



Cadence Workflow ContribFest: Build Durable AI Agent Examples

Session feedback

KubeCon CloudNativeCon North America 2026

Laptop can’t run the model

Pair with a neighbor.

Choose a ready Python-only task:
#110, #112 or #121.

Or browse in the optional Explorer’s mock mode.