ceph days LONDON 2026

The Rise of the AI Programmer

Implications for Ceph

Ernesto Puerta (IBM)
epuerat@ibm.com

Heavily-assisted-by: **ChatGPT, GitHub Copilot, Gemini**

ceph days
LONDON 2026

A Short Story

ceph days LONDON 2026

Back to 2011...

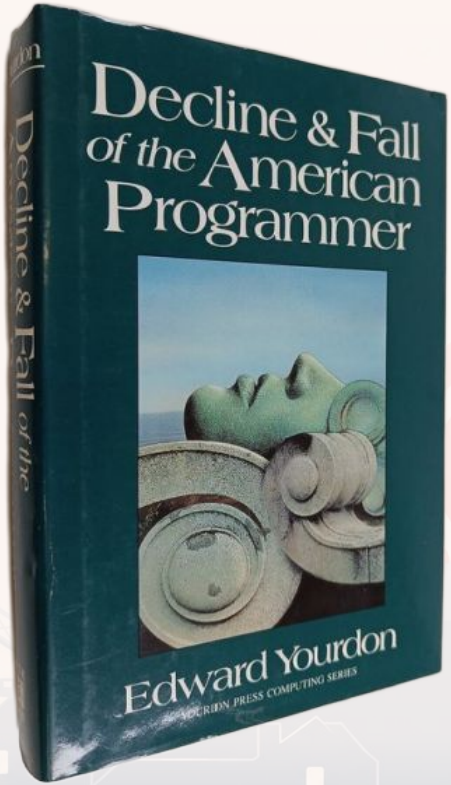

ceph days LONDON 2026

Decline & Fall of the American Programmer (E. Yourdon, 1990)

American *Programmers are ...*

- *More* ***expensive***
- *Less* ***productive***
- ***Lazier***
- *More* ***individualistic***
- *Make more* ***mistakes***

... than ***Foreign*** *Programmers*

ceph days LONDON 2026

Decline & Fall of the Human Programmer (2026)

Human *Programmers are ...*

More ***expensive***
Less ***productive***
Lazier
More ***individualistic***
Make more ***mistakes***

... than ***AI*** *Programmers*

ceph days LONDON 2026

“AI is here to help you”

ceph days
LONDON 2026

Using AI in Ceph

ceph days LONDON 2026

LLMs & Agents

ceph days LONDON 2026

Using AI in Ceph: Agentic SDLC

Problem: We (Ceph community) only have visibility/traceability over this. PRs are public.

Planning & Design
Requirements, architecture, design

Implementation
Initial coding and feature creation

Review
Analysis, summarization & suggestions

QA
Automated and Manual Testing

Documentation
Technical and User Docs

Build & Release
Final release to production

Maintenance
Support, troubleshooting, bug-fixing, ...

ceph days LONDON 2026

AI in Ceph: Policy

Still **under discussion**: https://github.com/ceph/ceph/pull/66788

Key points:

- **LLM-assisted code should not be banned** from Ceph.
- However, each component may set **stricter requirements**.
- **Disclosure is expected** (`Generated-by:` or `Assisted-by:` git trailers).
- Diverse views about providing the prompt.
- **Human contributors remain responsible/accountable** for correctness (e.g.: PR Review) and the DCO/legal implications (ownership).

ceph days LONDON 2026

Reviewing Code: Adoption

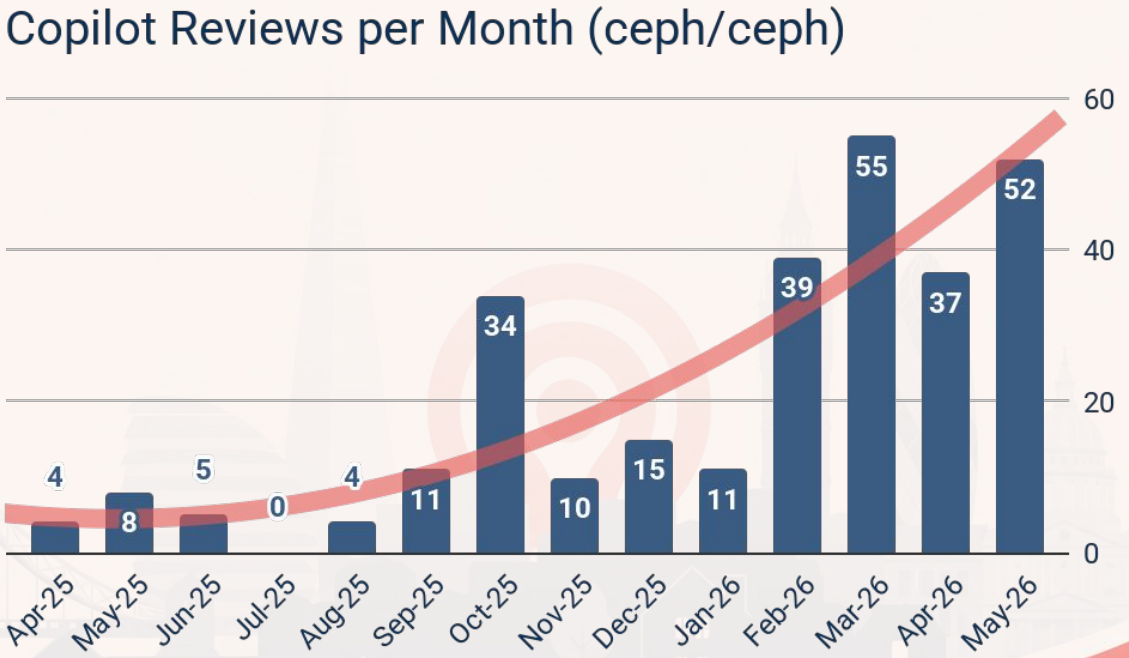

ceph days LONDON 2026

Reviewing Code: How to Enable

On-Demand Review

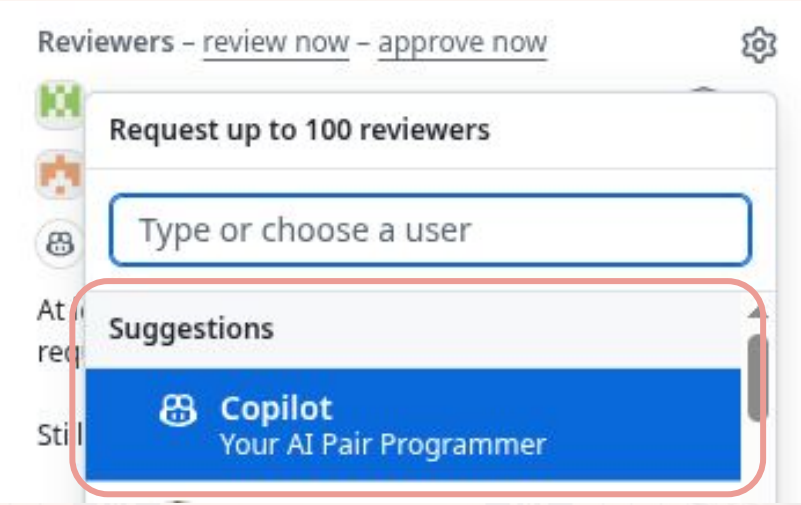

Automatic Review

Per-contributor (GitHub profile settings):

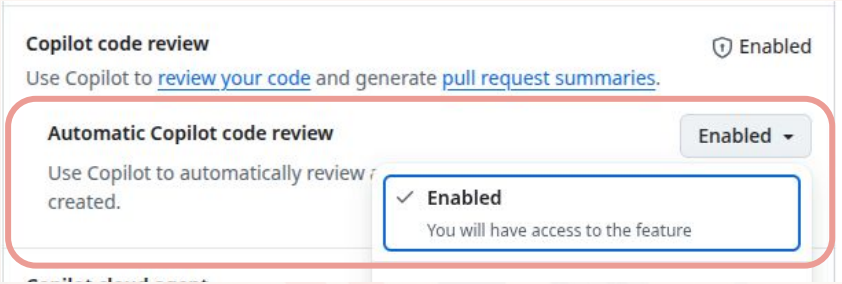

Per-repository (branch protection rules):

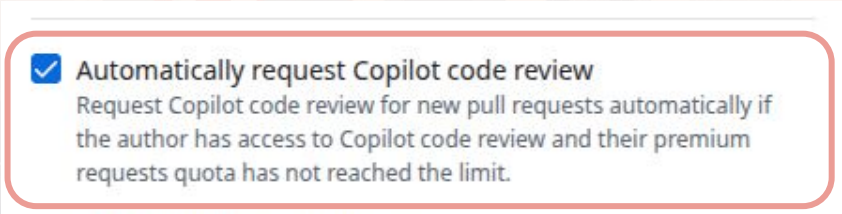

ceph days LONDON 2026

Reviewing Code: Evolution

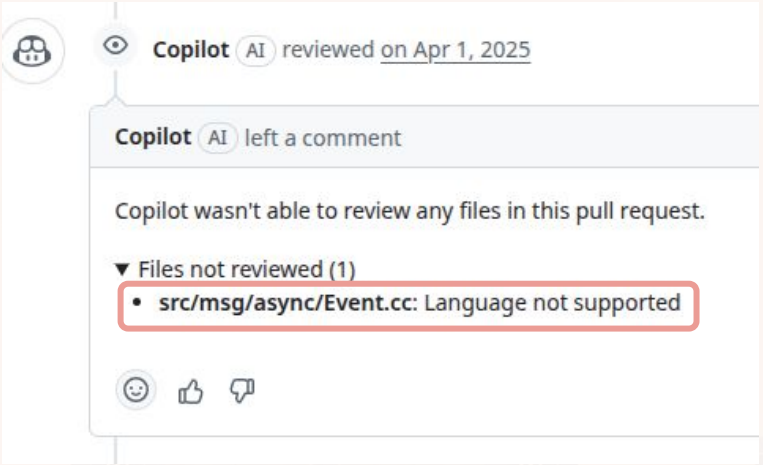

Apr 2025

Jun 2026

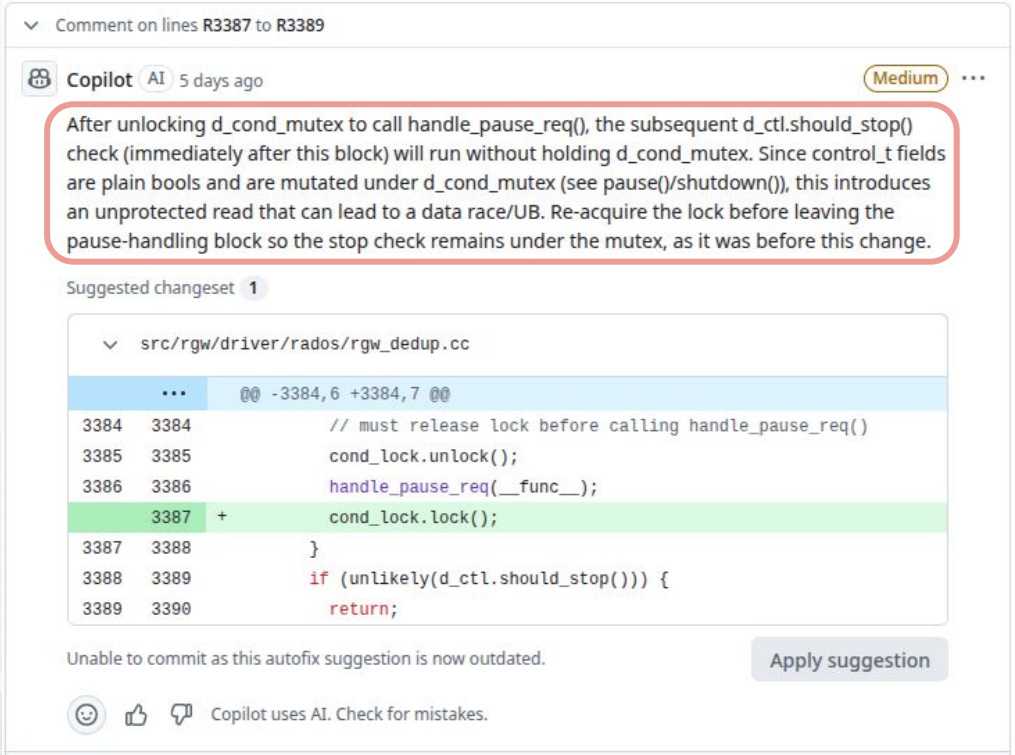

ceph days LONDON 2026

Coding: Customizing AI

Custom Instructions

.github/copilot-instructions.md

Skills

```
---
name:
ceph-commit-message-review
description: Reviews git
commit messages against Ceph
community conventions and
Linux-style best practices.
license: LGPL-3.0
metadata:
  author: ceph-community
  version: "1.0"
```

Custom Agents

.github/agents/

MCP Servers:

.mcp.json

- GitHub [default]
- Redmine [unofficial]
- Sentry
- Ceph [unofficial]
- Teuthology [💡]

Hooks

.github/hooks/

```
"userPromptSubmitted": [
  {
    "type": "command",
    "bash": "log-prompt.sh",
  }
]
```

ceph days LONDON 2026

Coding: Prompt Engineering

1. English Only

Other languages consume 20%-600% more tokens.
Cost is inversely correlated to HDI.

2. Don’t Be Polite

Use imperative commands.
Phrases like "Please" waste tokens and invite sycophancy.

3. Ground the Model

Prioritize examples and concretions.
Specify inputs, outputs, and enumerations.

4. Pass Compact Data

Use TOON, CVS, YAML or Markdown.
Use delimiters and lists for clear structure.

5. Adopt Spec Driven Development (SDD)

Transforming a version-controlled spec.md file directly into production-ready code.

ceph days LONDON 2026

Coding: Spec Driven Development (SDD)

ceph days LONDON 2026

The future of Coding: Human IN the Loop

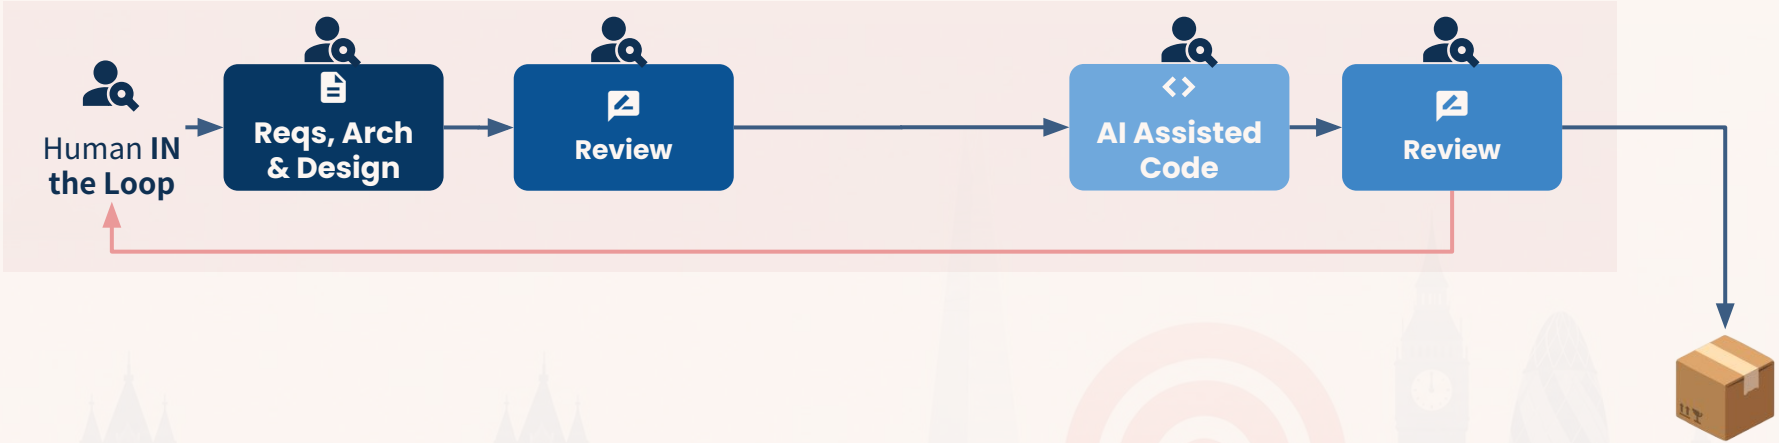

"Human in the Loop does not scale"

ceph days LONDON 2026

The future of Coding: Human ~~IN~~ ON the Loop

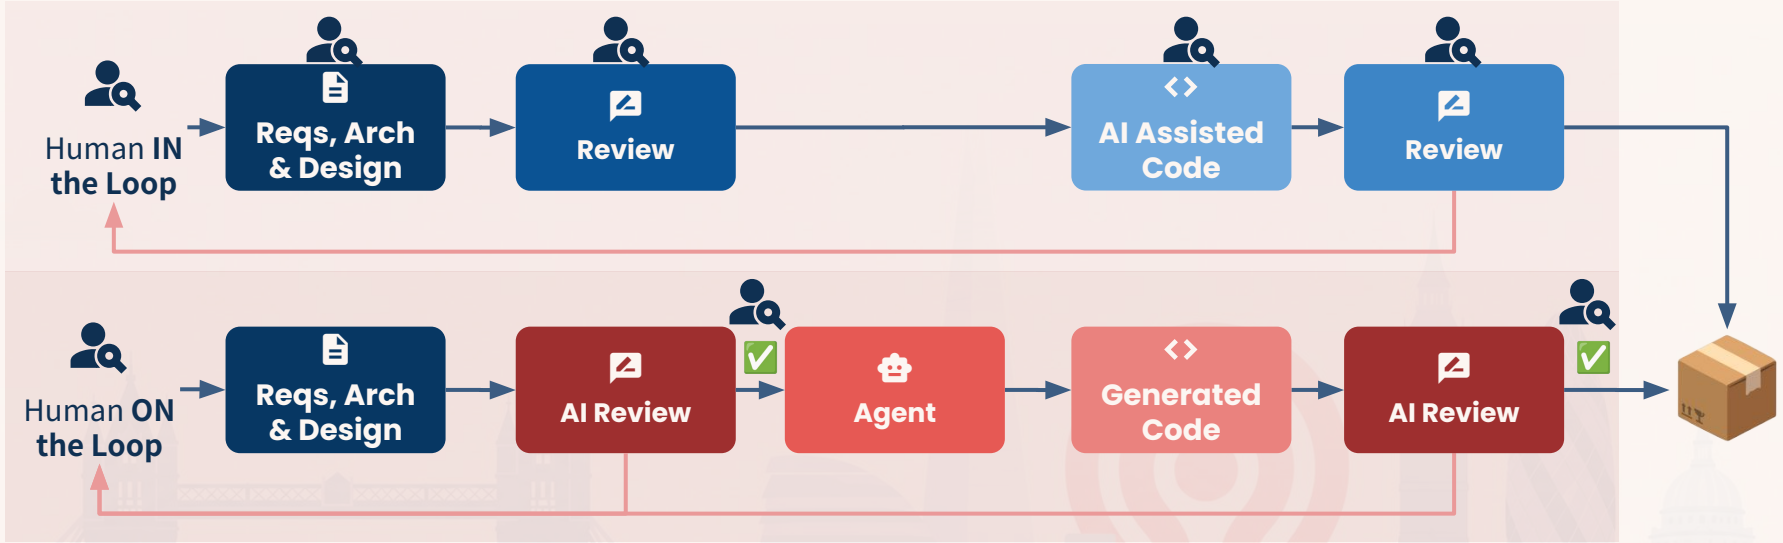

ceph days LONDON 2026

The future of Coding: Human AT the Edge

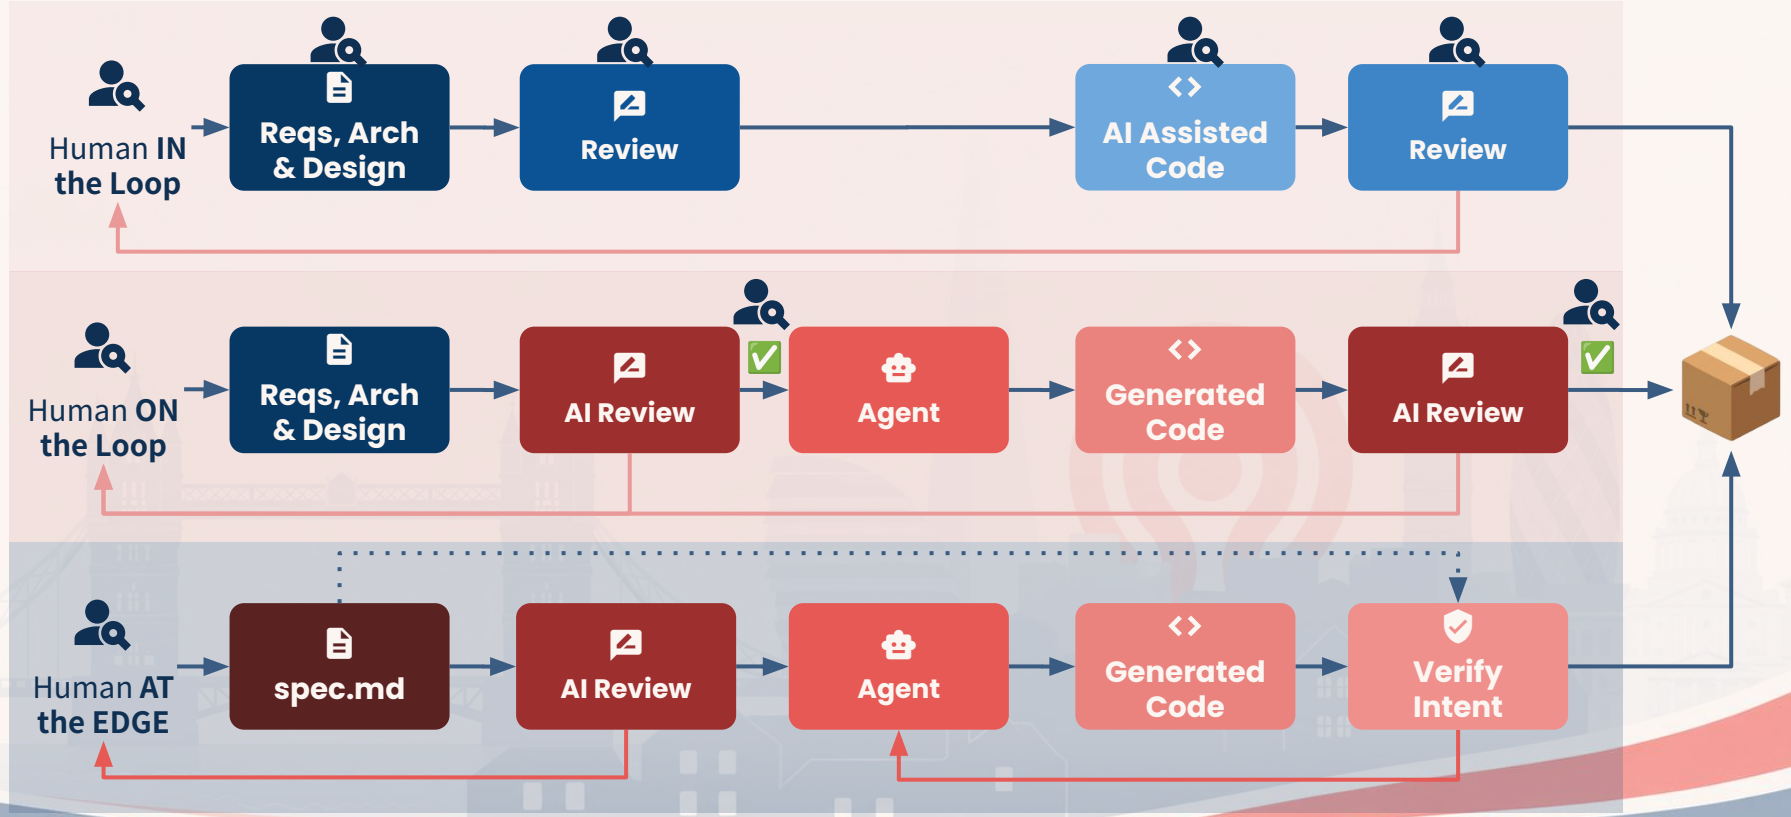

ceph days LONDON 2026

AI in Ceph: Security

Claude Mythos

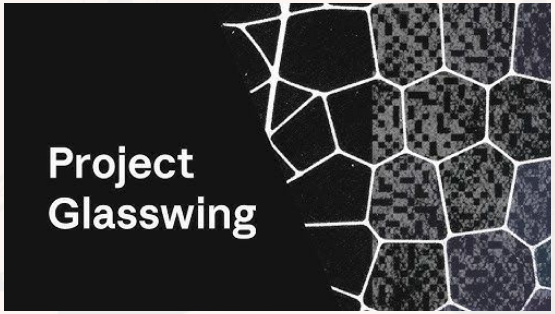

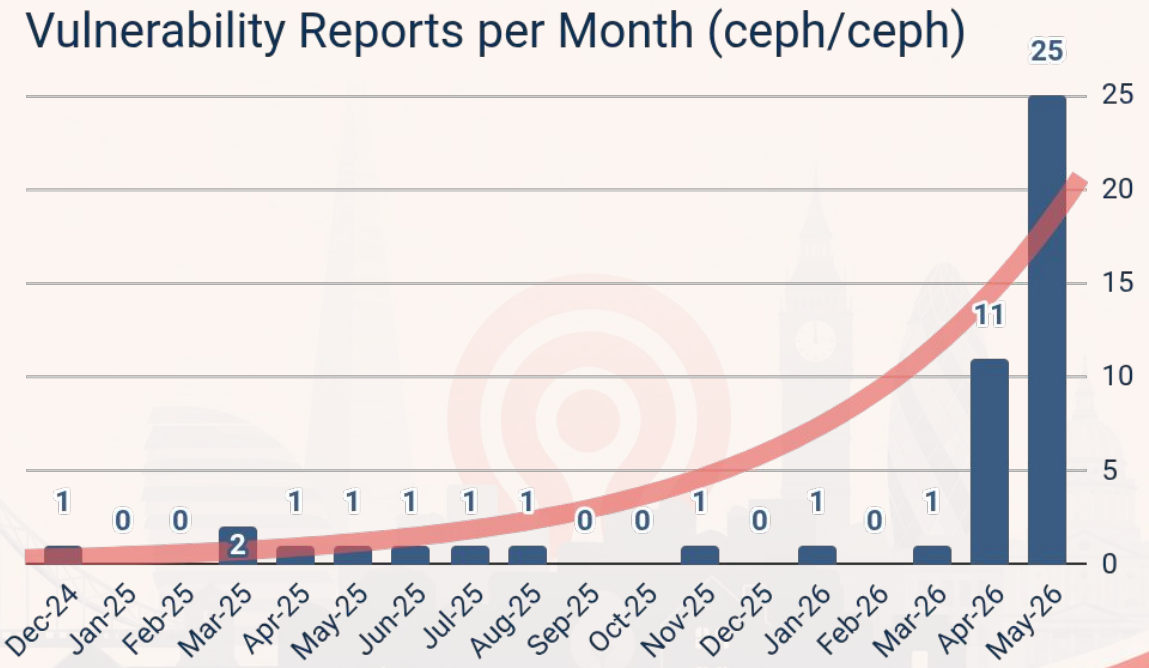

ceph days LONDON 2026

AI in Ceph: Troubleshooting

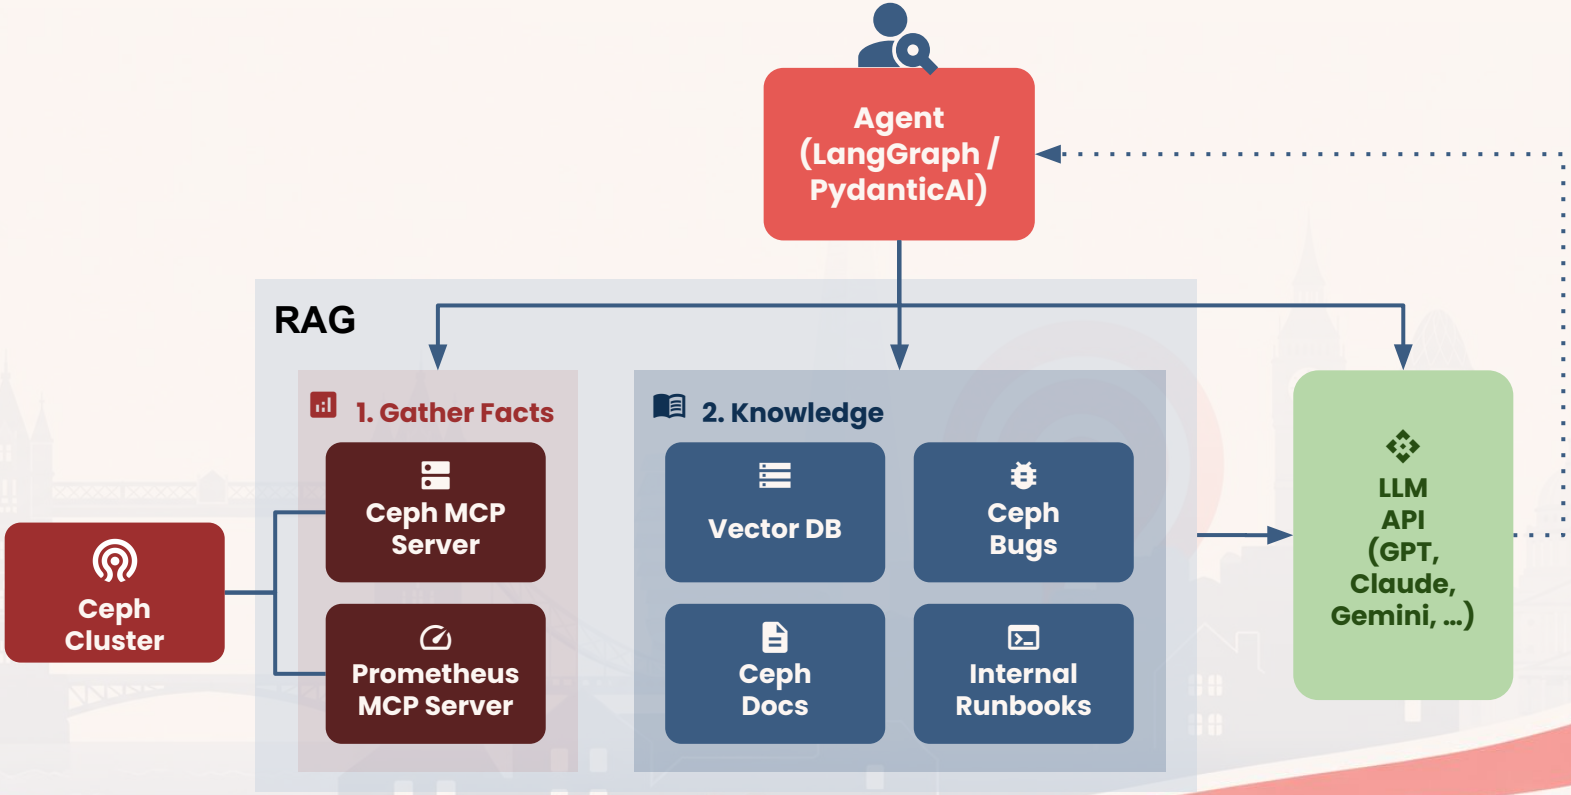

ceph days
LONDON 2026

The Dark Side of the AI

ceph days LONDON 2026

Current Challenges

AI Slop

Poor quality, hallucinations, and unverified AI outputs that clutter the codebase.

Mitigations:

- Treat AI output as untrustworthy.
- Define hard contracts (APIs, schemas).
- Establish project-wise AI rules (e.g.: readable code over redundant docs).
- Beware of AI bias (over-engineering, over-abstractions, etc).

Legal

Copyright concerns, compliance with evolving regulations, and liability/accountability for AI-generated code.

Mitigations:

- Release a Community AI policy.
- Treat AI as a tool, not as a legal entity.
- Mandatory disclosure of AI usage.
- Strong provenance and code tracking.
- Run license compliance tools to detect code similarity.
- Extra reviews for AI-assisted code.
- Limit AI for specific parts of the project.

Security

Risks involving data leaks, prompt injection, and the need for robust sandboxing environments.

Mitigations:

- Always sandbox AI agents.
- Isolate CI from AI.
- Only allow signed MCP servers.
- Never inject prompts from code.
- Sanitize context.
- Every Agent can become an attacker if fed with a rogue prompt.

ceph days LONDON 2026

Longer Term Risks

Vendor Lock-in

Dependency on proprietary AI platforms, making it difficult to switch providers or regain independence.

Mitigations:

- Multi-model by design,
- Open Standards
- Keep prompts/spec/knowledge in git.
- Maintain a human-only fallback.
- Invest on local/LF-supported LLMs.

Developer Detachment

Loss of core engineering skills and motivation, as developers rely more and more on AI for higher-level tasks.

Mitigations:

- Keep silos where the human craft is highly valuable and rewarding.
- Define “AI forbidden” areas.
- Maintain craft rituals.
- Proactively share AI practices with less AI-skilled engineers.

AI-Shoring

The shift of work to automated AI systems, threatening labor markets and global economics.

Mitigations:

- Enforce junior hiring quotas.
- Large orgs/projects should comply with human quotas (% human written code).
- Learning should remain a core human skill.
- Deep-expertise Human Reserves.

ceph days LONDON 2026

15 years later…
The Rise of the Human Programmer

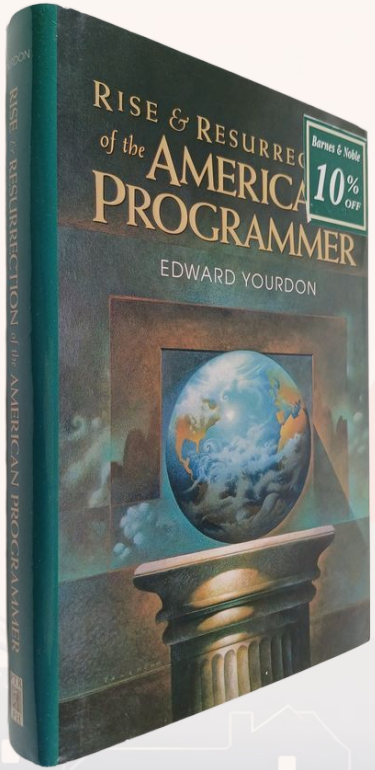

ceph days LONDON 2026

Any Hope?

Humans

Intent
Focus on the "Why" and purpose

System understanding
Holistic view of architecture

Tradeoffs
Complex decision balancing

Accountability
Ownership of outcomes

Long-term ownership
Sustainable lifecycle management

LLMs

Pattern synthesis
Matching and combining data

Token prediction
Probabilistic next-step generation

Local optimization
Specific task performance

Plausible output
Realistic but not verified content

Fast iteration
Rapid drafting and implementation

ceph days LONDON 2026

AI may **not replace** programmers (yet)…

But **projects** with **AI-skilled** programmers will **outpace** those without.