

Narrative Interfaces for Multi-Agent Software Development

AJ Barea • Rochester Institute of Technology • Center for Human-Aware AI (CHAI)

Abstract

Kourai Khryseai reimagines multi-agent development as a collaborative narrative experience. By translating psychological engagement models from game design into developer tooling, we shift the AI paradigm from silent automation to a transparent, human-on-the-loop "Forge Party."

From Static Tools to Narrative Teammates

The Problem: CLI interfaces flatten emotional context. AI acts as *tools*, not *teammates*.

The Solution: A Human-on-the-Loop (HOTL) narrative experience. Agents stream internal reasoning; developers guide critical decisions.

Psychological Interaction Models

Translating proven psychological engagement models from game design into developer tooling via two interlocking systems:

✦ **System Affinity** (*Mass Effect*): Per-agent trust score (0.0-1.0) across 4 tiers—Cold, Professional, Warm, Intimate. Rises when the developer accepts suggestions and engages with agent interjections; falls when advice is dismissed or quality checks are skipped.


✦ **Forge Virtues** (*Disco Elysium*): 8 psychological dimensions that act as inner agent voices. Each virtue has a vice (≤ 0.25) and apex (≥ 0.80) threshold. Crossing a threshold triggers the patron agent to interject with in-character dialogue.

The Forge Agents



The Personalization Ecology

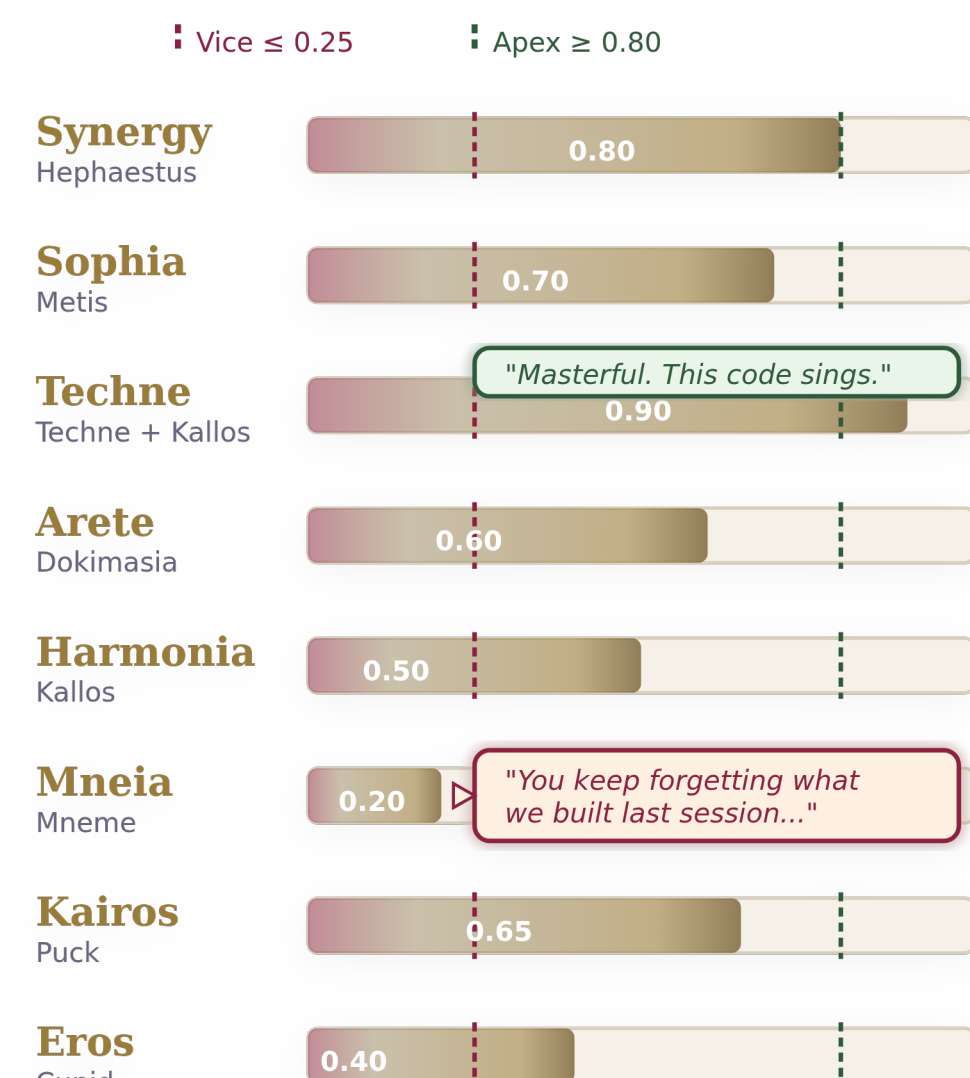
System Affinity
Inspired by Mass Effect's Paragon / Renegade




How Affinity Changes

- ▲ Accept agent suggestions
- ▲ Complete tasks with high virtue
- ▲ Engage with agent interjections
- ▼ Dismiss or ignore agent advice
- ▼ Repeated low virtue scores
- ▼ Skip quality checks (anti-patterns)

Forge Virtues
Inspired by Disco Elysium's Skill System




Crossing a threshold triggers the patron agent to interject with in-character dialogue.



SQLITE MEMORY CORE
A2A PROTOCOL RUNTIME

- ✦ **System Affinity:** Per-agent trust tiers that unlock deeper collaboration and narrative branches.
- ✦ **Memory Core:** SQLite-backed episodic (threads) & semantic (hierarchies) memory.



- ✦ **Forge Virtues:** 8 psychological axes with vice/apex interjections from patron agents.
- ✦ **Virtue Web:** Radar chart of the developer's psychological profile across all 8 dimensions.

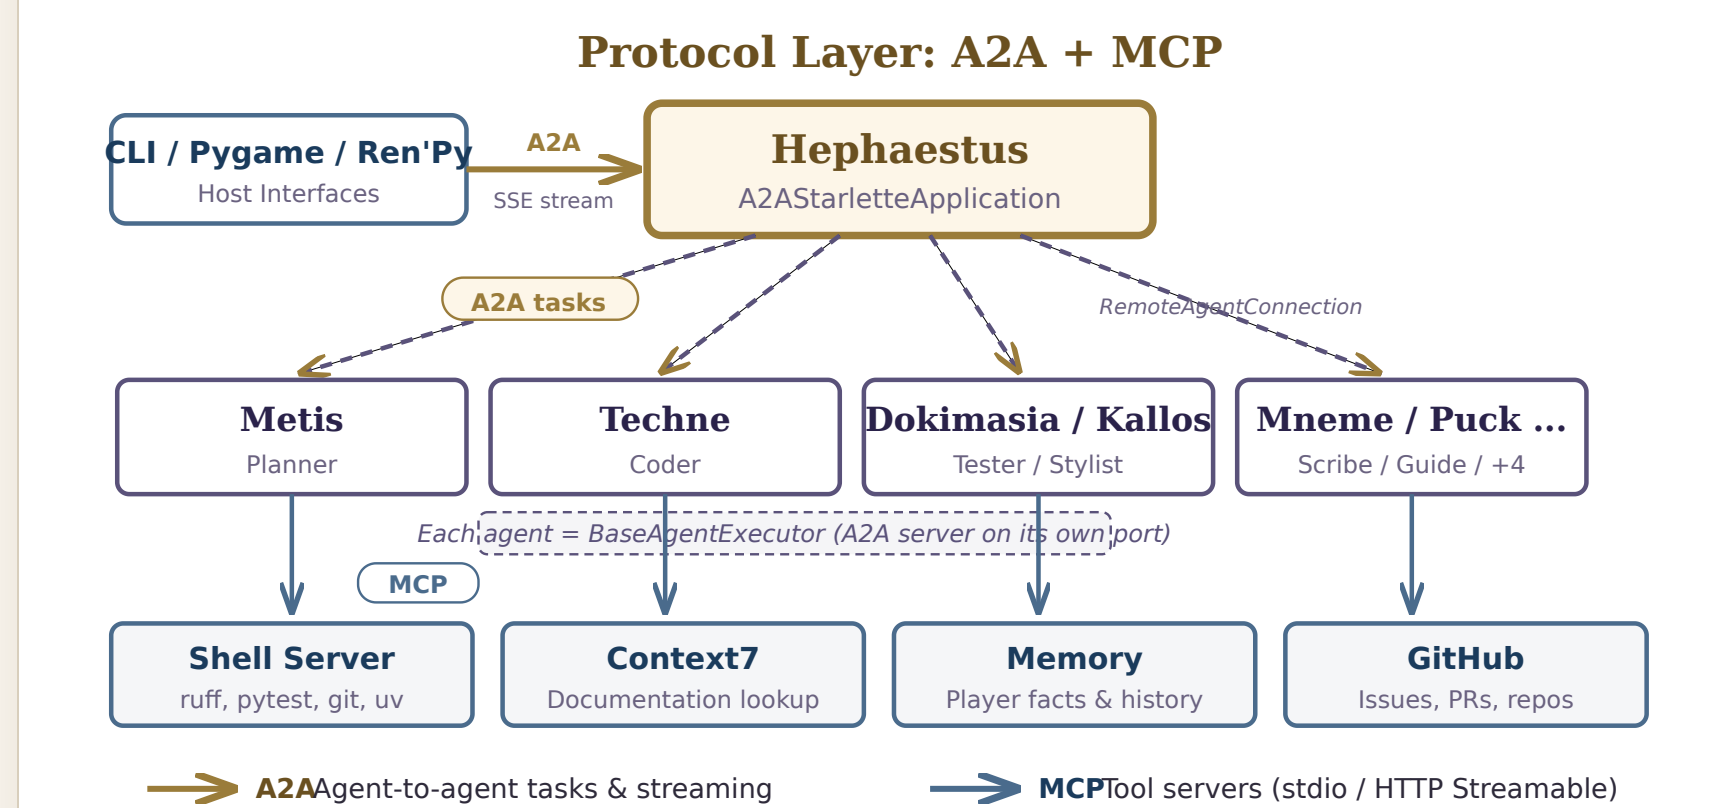
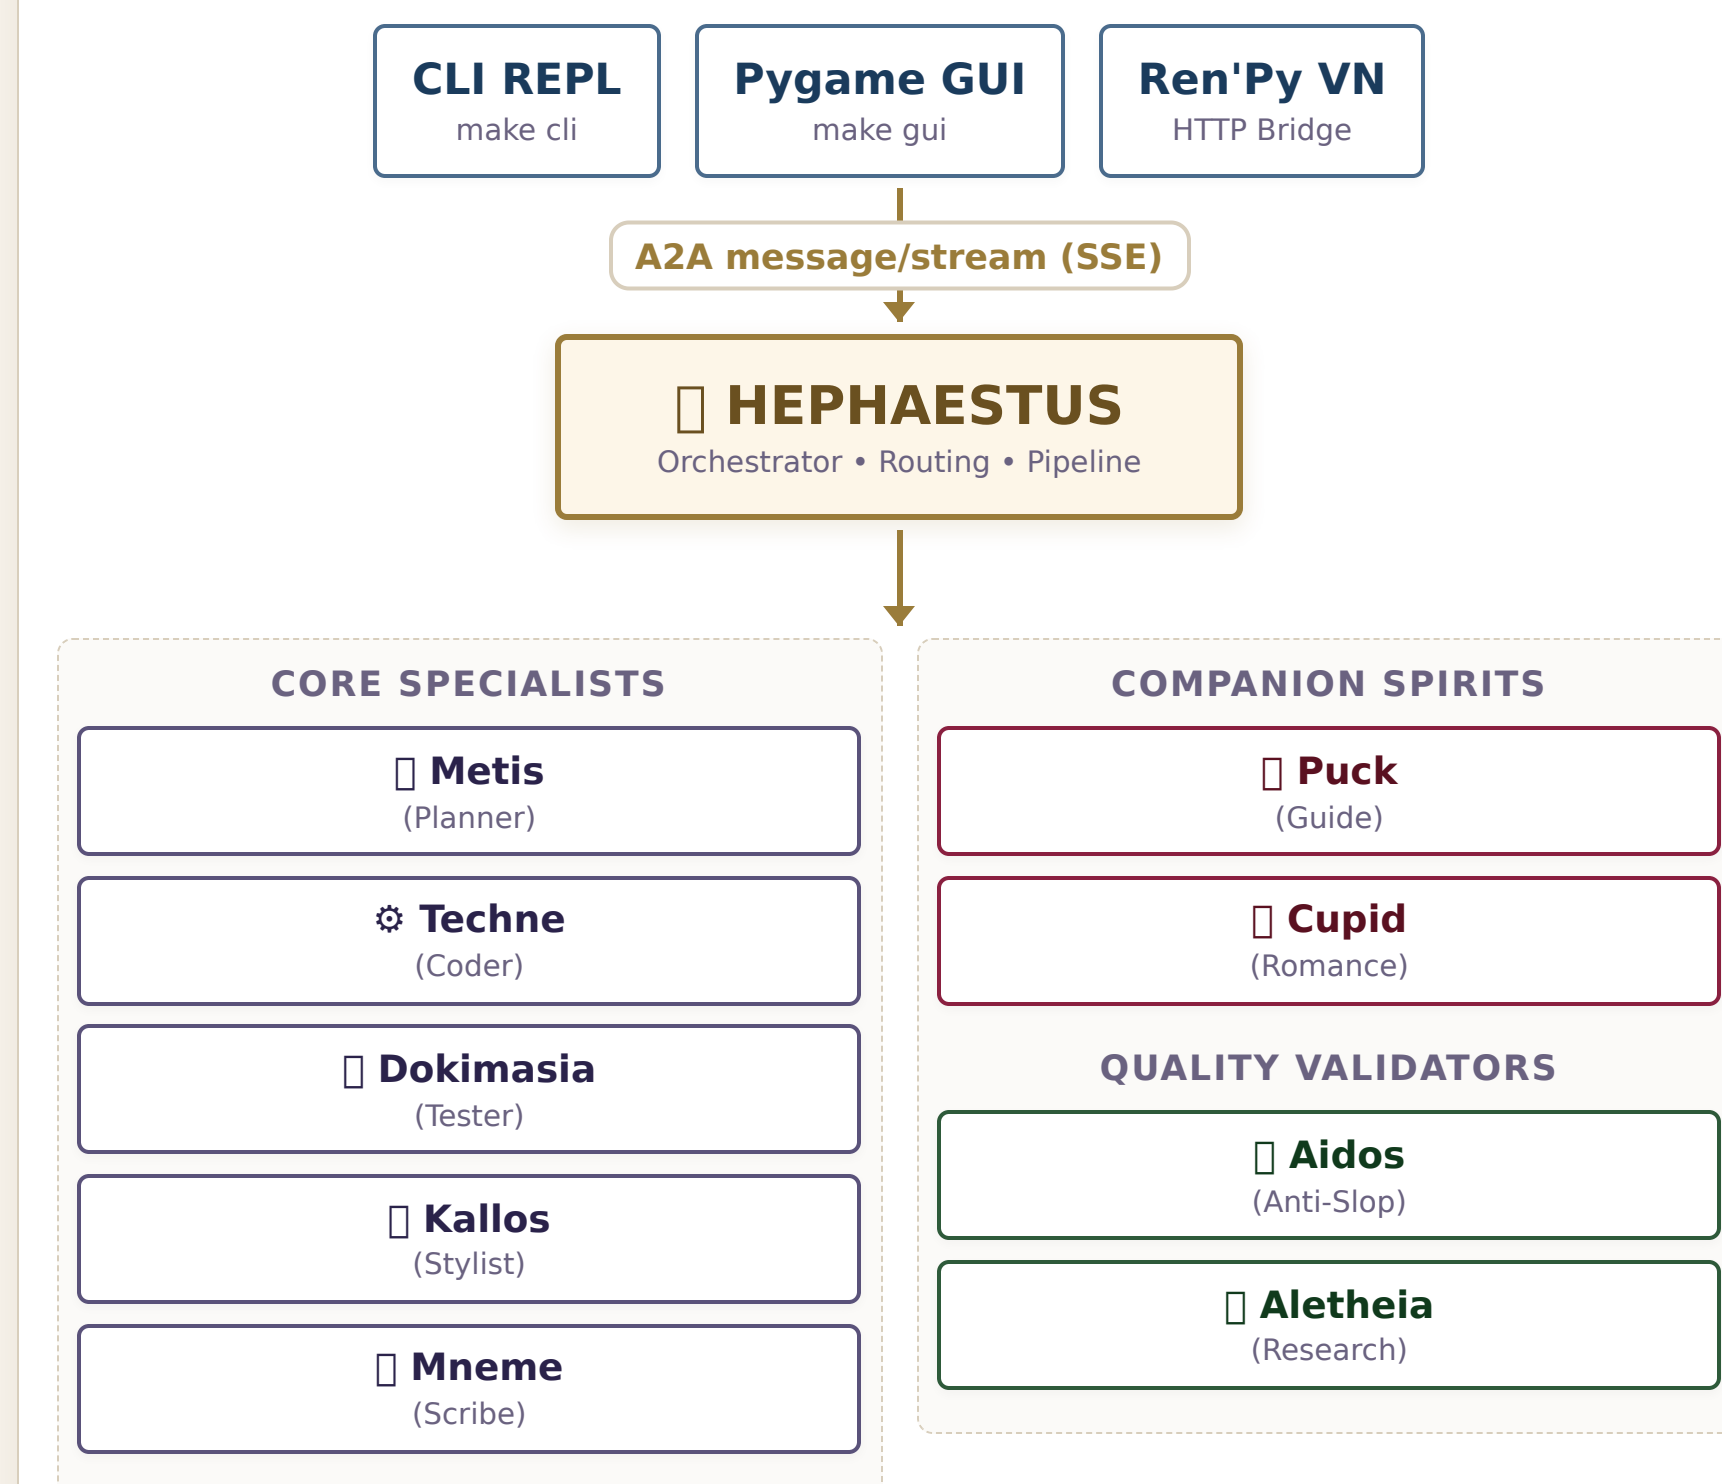
CHAI Focus: Human-Centered AI

This ecology negotiates user trust and evolves interactions over time. Specialized validation agents (*Aidos* for anti-slop, *Aletheia* for research) observe without owning virtues—ensuring quality without bias.

Host-Agnostic Narrative Delivery

Key Insight: UI is an independent design dimension, fully decoupled from agent capabilities. Three equal-citizen hosts connect to the exact same 10-agent backend.

System Architecture



- ✦ **A2A:** Agent-to-Agent communication.
- ✦ **Observability:** OTel & Jaeger.
- ✦ **Core Engine:** LiteLLM & Starlette.

Future Horizons

Beyond reactive assistance, we are exploring **Deep Narrative Branches** where agent personalities evolve permanently based on long-term virtue profiles. Our goal is a development environment that acts as a mentor, turning engineering into a journey of shared growth and divine craftsmanship.

