

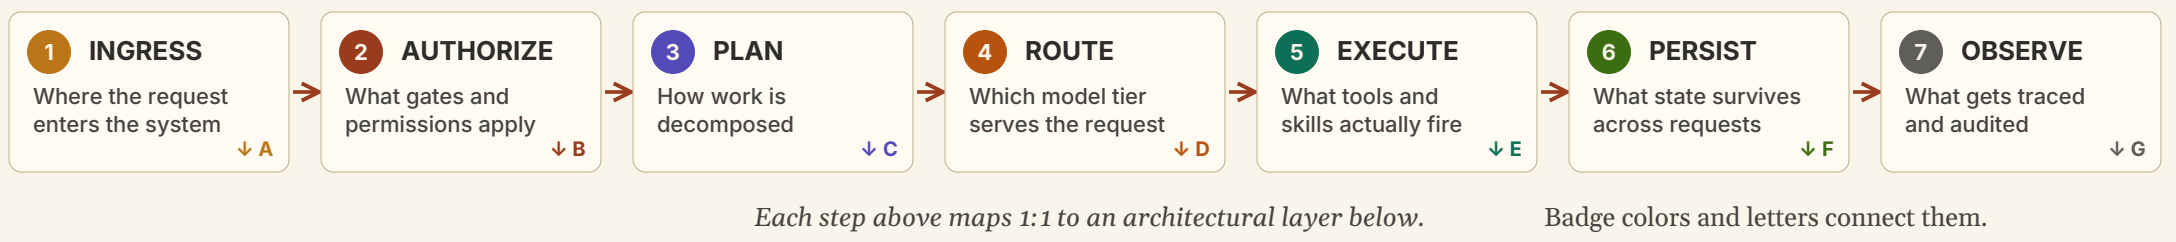
Open Source Agentic AI Architecture

A cost-effective alternative to platform AI agents

The disciplined open source path produces a defensible institutional posture at an order-of-magnitude lower cost than frontier-default managed alternatives (i.e. Claude), and at materially better auditability than any vendor attestation can replicate. The right answer is context-dependent — but the option set is wider than most institutions realize.

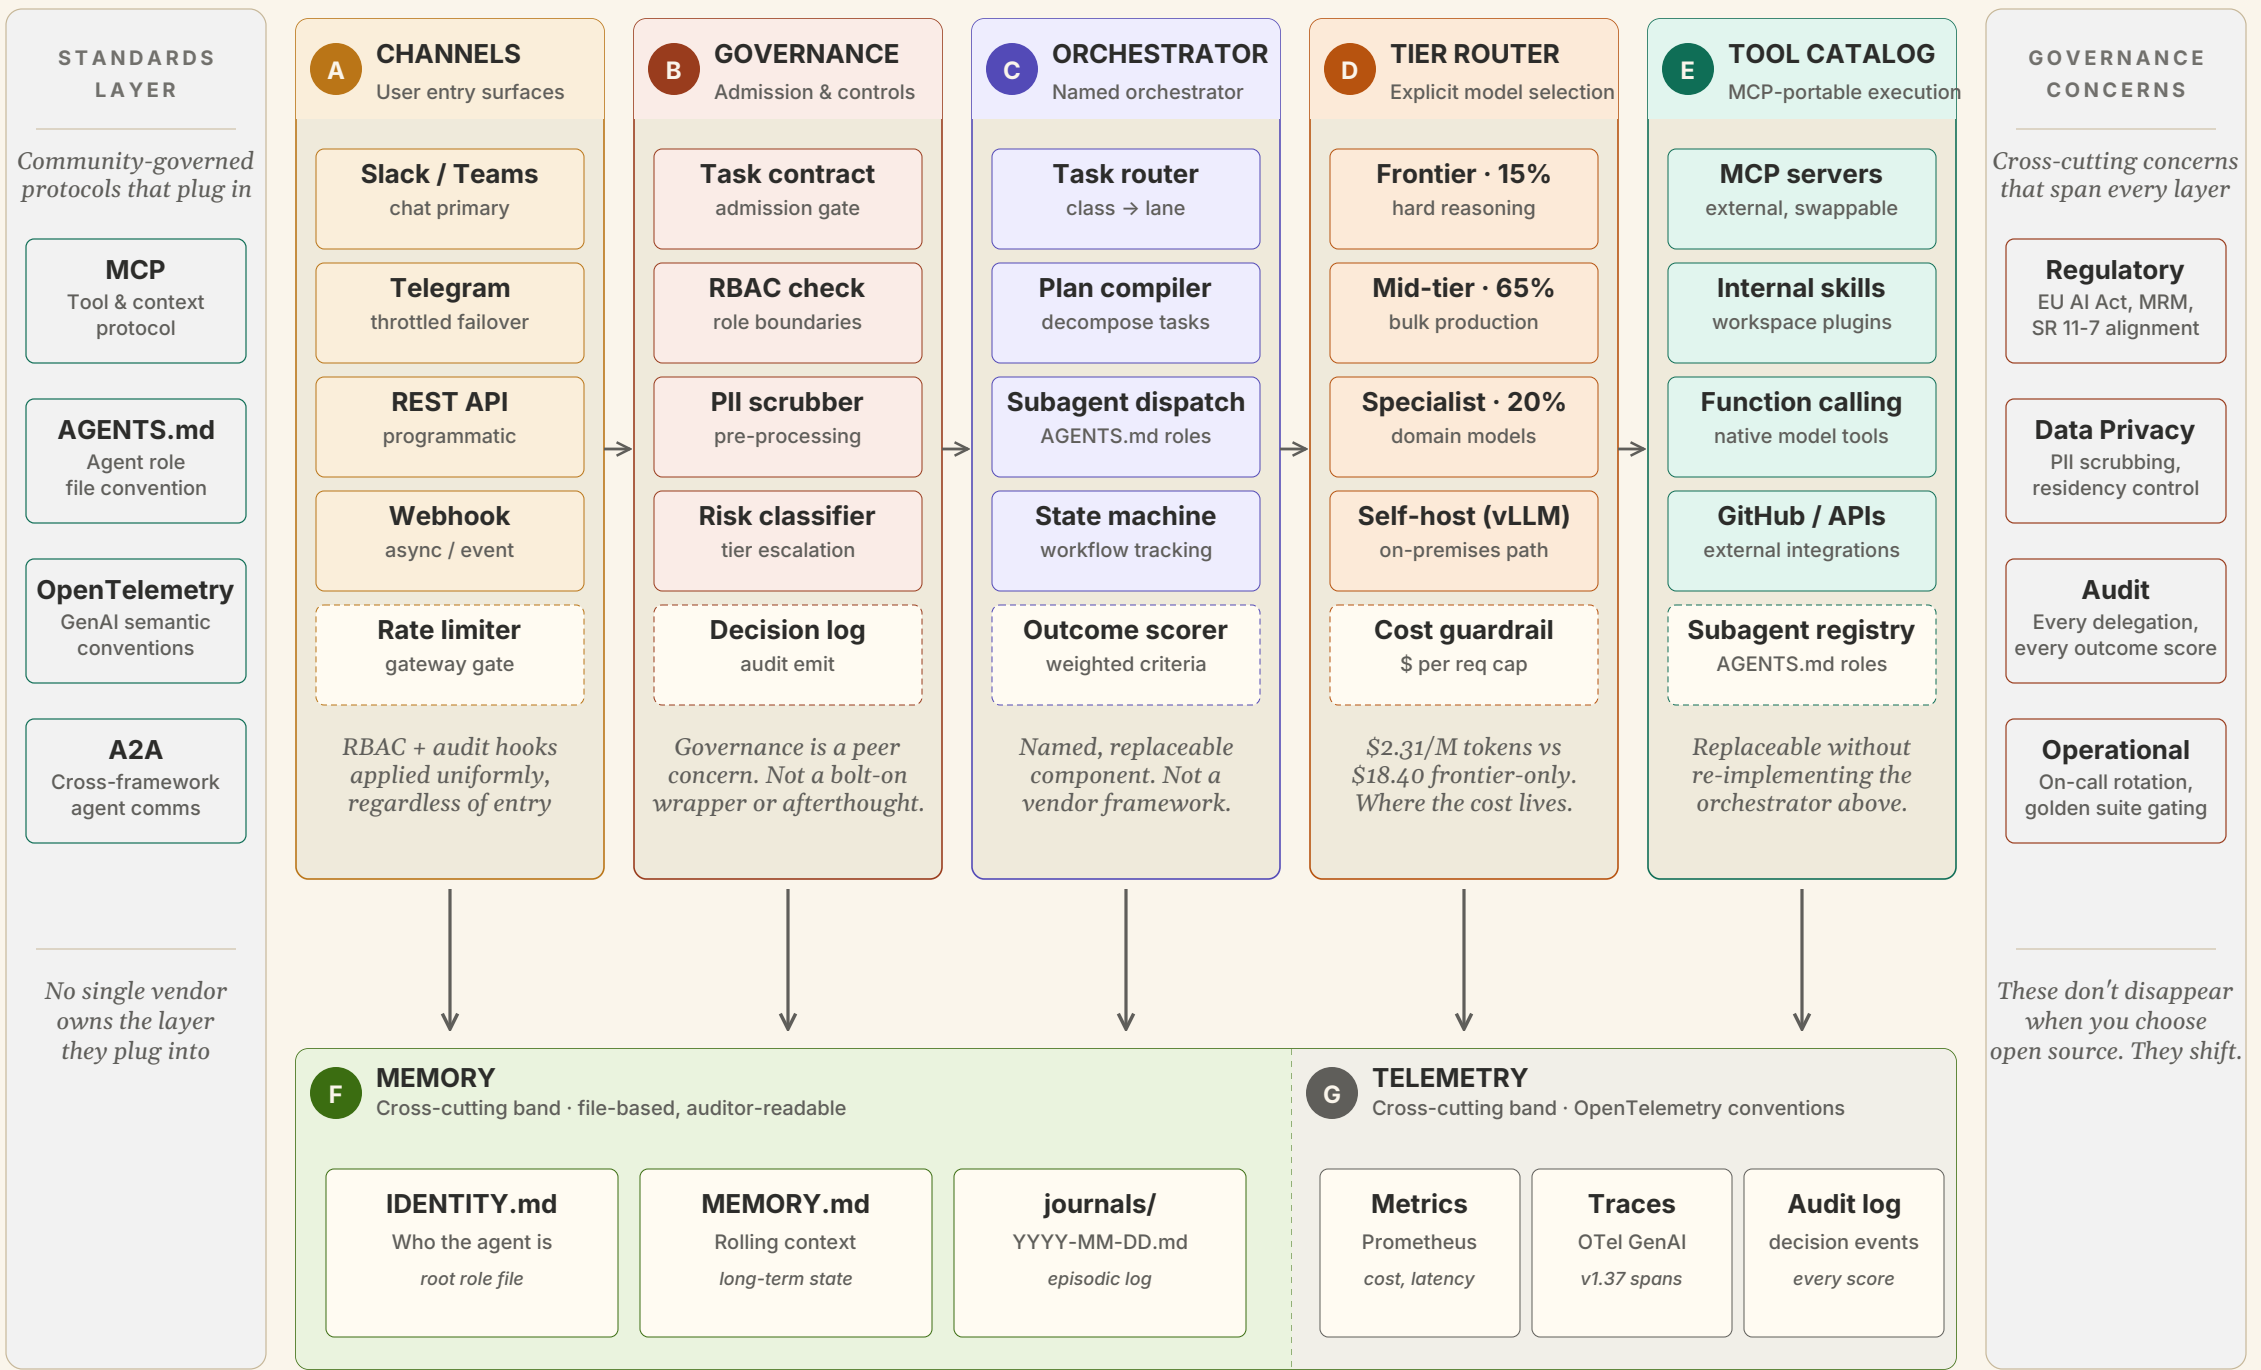
01 · THE REQUEST FLOW

Every request walks the same path. *No exceptions.*



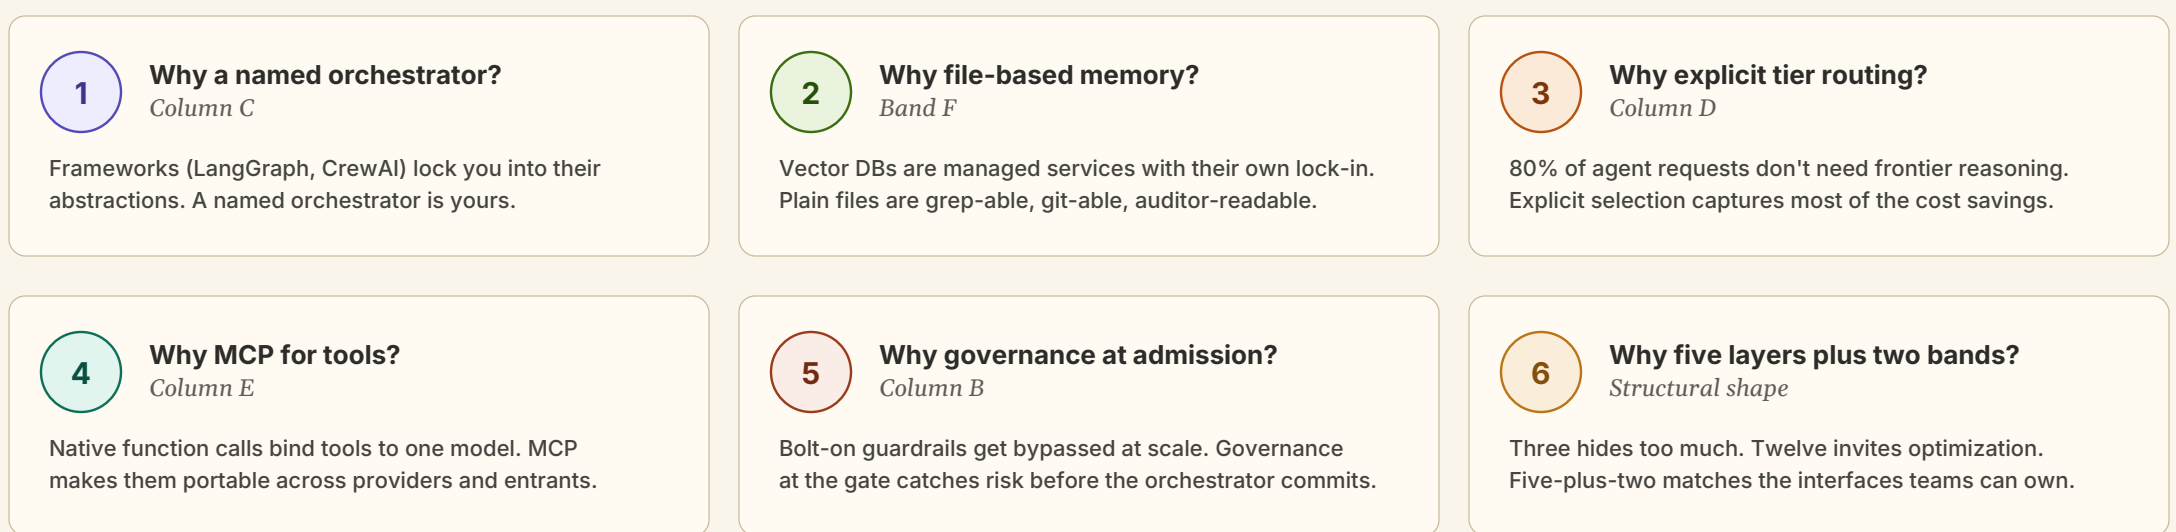
02 · THE ARCHITECTURE

Five layers, two cross-cutting bands. *Every component is replaceable at the protocol layer.*



03 · DESIGN DECISIONS


Why each layer is structured the way it is. *The rationale behind the architecture.*



04 · THE COST EVIDENCE

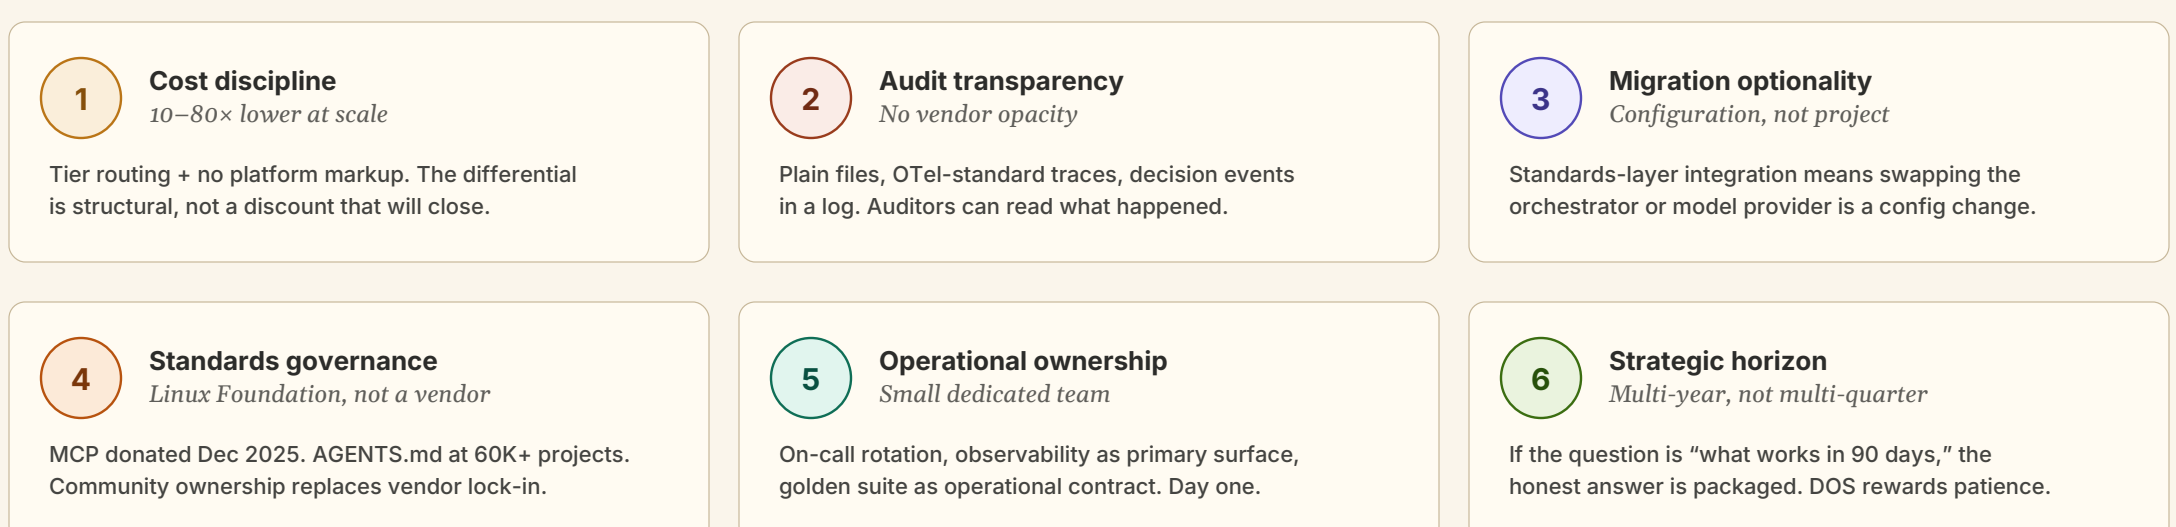
10–80× lower cost at scale. *Structural, not transient.*

Monthly operating cost at 10,000 agent requests · 2026 reference data · institutional workload mix



05 · WHAT MAKES IT INSTITUTIONALLY VIABLE

Six **discipline** conditions separate *viable* open source from *hopeful* open source.



You don't need a packaged platform to build production agentic AI. The disciplined alternative is cheaper, more transparent, and yours to own.