

Your Agent Architecture Has a Name. **Actually, Eight of Them.**

Understanding these eight recurring patterns is critical for designing reliable, scalable, governable AI systems.

SWIPE →

Eight patterns, four questions

01 Hierarchical & Coordinating — who's in charge

02 Distributed & Context-Aware — who's on the edge

03 RAG-Based & Planning — how work gets grounded and sequenced

04 Human-in-the-Loop & Self-Learning — where humans stay in control

01

Hierarchical Agents

Planners decompose goals, orchestrators sequence subtasks, executors call the tools and APIs — mirroring an organisation chart, top to bottom.

Used in: end-to-end customer journeys, incident management, multi-domain AI operating systems

04

Human-in-the-Loop Agents

Operate semi-autonomously, then pause at defined decision points — approvals, plan selection, escalation — for human confirmation or override.

Used in: regulated decision-making, operations centres, expert workflows

06

RAG-Based Agents. Retrieve documents, data, or facts at run-time and ground reasoning in that context — cutting hallucination and keeping outputs traceable.

Used in: enterprise knowledge assistants, decision support, hybrid analytics



The important question is not which label fits a given product demo, but which of these patterns are necessary, how they're composed, and how they're governed.

Terence Kok

Enterprise AI Strategist

Centralised control vs. edge resilience

HIERARCHICAL / COORDINATING

Single planner or coordinator

Clear control, one point of failure

Best for: structured, multi-step workflows

DISTRIBUTED

Autonomy at each node


Resilient to partial failure

Best for: IoT, edge, federated systems

Production systems don't pick **one**.

A representative architecture: a Planning Agent inside a Hierarchical structure, Coordinating Agents managing cross-domain handoffs, RAG-Based Agents grounding decisions in enterprise knowledge, Context-Aware Agents running at the edge, Human-in-the-Loop gates at key decision points, and Self-Learning Agents iterating within defined guardrails.

The eight, **layered**.



All eight patterns, with where each is used in production:
terencekok.com/blog/eight-practical-types-ai-agents-emerging-real-systems

Eight patterns. **One governable system.**

The question isn't which label fits your demo — it's which patterns your use case actually needs, and how you instrument them.



Terence Kok · Enterprise AI Strategist

SAVE · COMMENT · FOLLOW