

Enterprise Agentic AI Architecture Guide

A companion worksheet for CMOs, AI directors and product leaders

Integrated.Social

Companion resource to: *Agent Washing: True Agentic AI vs Deterministic Workflows*

What this guide is for

Use this worksheet to separate a fixed automation from an AI system with meaningful decision-making autonomy. It is designed for the point after the demo: when a team needs to decide what a system can do, what a person must approve, and how success will be assessed.

The aim is not maximum autonomy. It is the right level of autonomy for the job, with clear guardrails where reliability, safety and accountability matter.

The four-level autonomy spectrum

Level 1: Deterministic workflow

A trigger starts a predefined sequence. The same input is intended to follow the same path. This can be reliable and useful, but it is not independently planning a route to an objective.

Typical fit: repeatable routing, notifications, data hand-offs and controlled production steps.

Level 2: Knowledge assistant

The system can retrieve approved business context and propose an answer or next step, while its access and outputs remain bounded.

Typical fit: internal knowledge retrieval, research support and draft preparation.

Level 3: Action agent

The system can select from approved tools to progress a defined objective. Its actions should be constrained by permissions, logs, spending limits and human approval gates.

Typical fit: monitored operational tasks, campaign support and structured cross-system work.

Level 4: Multi-agent coordination

Multiple specialist roles contribute to a complex, open-ended objective. This level requires explicit roles, hand-offs, escalation rules and governance rather than a collection of unobserved prompts.

Typical fit: genuinely complex work where the path cannot be fully scripted in advance.

Four readiness questions

1. **What outcome is the system responsible for progressing?** State the business objective before naming a model or vendor.
 2. **Which decisions may it make alone?** Separate recommendations, reversible actions and consequential actions.
 3. **Which tools and data can it access?** Limit access to what the task requires and record how access is used.
 4. **What happens when it is uncertain or outside scope?** Define a human escalation route before launch.
-

Hybrid architecture model

Most enterprise use cases benefit from combining predictable infrastructure with bounded AI reasoning:

- **Knowledge foundation:** approved content, product and customer context.
- **Decision layer:** a clear objective, constraints and evaluation criteria.
- **Action layer:** narrow tools with permissions appropriate to the risk.
- **Human sign-off:** checkpoints for sensitive, irreversible or customer-facing actions.
- **Evidence layer:** logs, outcomes and review notes that allow the team to improve the system.

A workflow can be an important part of this model. The test is whether the system can make an appropriate, governed decision when the input does not match a fixed script.

Decision rights matrix

Activity	Can the system propose?	Can it act without approval?	Required control
Summarise approved internal material	Yes	Usually	Source boundaries and quality review
Draft a campaign or response	Yes	Usually not	Human editorial approval
Update a reversible internal record	Yes	Sometimes	Role-based permissions and audit log
Spend budget, send external communications or change customer data	Yes	No by default	Named human approval and clear escalation
Make a high-impact or irreversible decision	Yes	No	Senior accountable owner

Eight-point readiness assessment

Mark each item **Ready**, **Needs work** or **Not defined**.

1. The intended business outcome is written in plain language.
2. The system's level of autonomy is explicit.
3. Approved data sources and prohibited sources are documented.
4. Tool permissions are limited to the task.
5. Human approval gates are defined for consequential actions.
6. Failure, uncertainty and escalation paths are tested.
7. Logs and outcome measures support review.
8. A named owner is accountable for operating the system.

Go / no-go prompt

Proceed only when the team can explain what the system may decide, what it may do, what it must never do, and who takes responsibility when the expected path changes.

Next step

Read the accompanying article for the full discussion of agentic AI, deterministic workflows and practical enterprise governance:

Agent Washing: True Agentic AI vs Deterministic Workflows

<https://integrated.social/blog/agent-washing-true-agentic-ai-vs-deterministic-workflows-enterprise-guide-2026/>

For a structured view of your current stack, start with the **Free AI Answer Readiness Score**:

<https://integrated.social/free-ai-growth-audit/>

This guide is educational material, not legal, security or procurement advice. Organisations should apply their own governance, security and compliance requirements to each use case.