AIAI-EVD-012 | CHAPTER 12 | Governance guides

What can be done—and who may authorise it?

A practical distinction connecting Decision Intelligence, HumanGate, TrustMesh and multi-tool workflows.

Jessica Fogarty | AIAI.science
Public editorial edition 1.0 | 27 October 2026

Take away three things

1. Capability and permission need separate records.

2. Attach approval to the specific action and scope.

3. Verify the outcome and retain exceptions.

Begin with a precise action

Reading a record, preparing a recommendation and changing an external system have different consequences. The workflow needs to identify the actual action requested.

Decision Intelligence considers feasible options and trade-offs. The authority check then establishes which role can approve or perform the selected action in the relevant context.

Carry the conditions with the task

A useful approval has a purpose and scope. Provider controls should be able to represent the required boundary; an incompatible or weaker control becomes a review item.

The latest multi-MCP direction makes this compatibility question explicit. A proposed connection cannot assume that every provider enforces the same policy in the same way.

Record what happened

The result needs an observable check. A tool response or companion saying “done” should be connected to the evidence of the external state where that can be obtained.

Denials, returns for correction and partial completion are legitimate outcomes to preserve. They help distinguish a bounded action from an unsupported claim of complete automation.

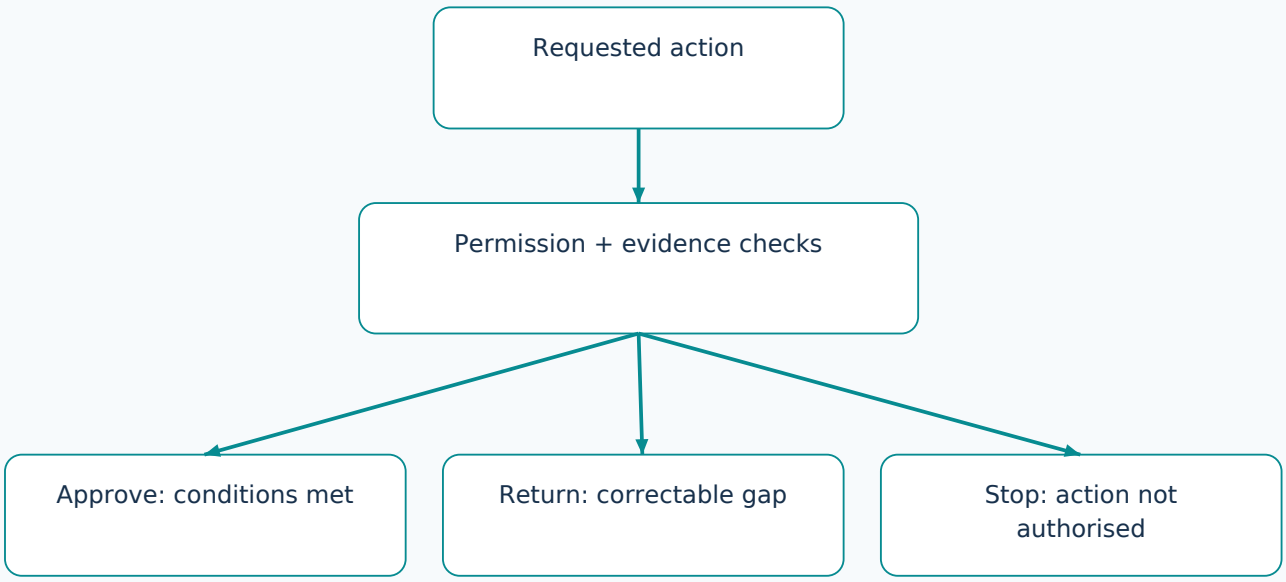
Make governance inspectable

For one consequential step, document the action, evidence, permitted role, approval conditions and expected result. Ask how the system behaves if any required part is missing.

The architecture is evaluated through those specific behaviours. It does not support a claim of universally perfect decisions or complete governance.

Figure 12. The idea in one view

FIGURE 12 / What can be done—and who may authorise it?



Concept illustration. Read the chapter for its scope and assumptions.

Conceptual explanation with illustrative inputs where shown. This figure is not an empirical result or a validation of a deployed system.

Basis and scope

Public architectural synthesis of HumanGate, TrustMesh, the AIAI.science blueprint and the latest recovered multi-MCP direction reviewed 27 October 2026. These concepts are not a universal security or governance guarantee.

Read online and explore related work

Open this chapter online

The online edition includes interactive representations where useful. The chapter remains complete as a static explanation.