Verifiable runtime inputs

Policy graph
Rules and triggers

Domain graph
Agents, assets and roles

Episode graph
Pending effects

Normal admission

Check the exact action
Verify mandate and evidence
Pin source versions

Coordinated kill switch

Verify evidence and freshness
Authenticate the responder
Check scoped stop authority

Independent local stop

Local operator control or independent safety monitor

Direct protective path

Independent of the AI model, graphs, network and remote identity services.

Protected resource boundary

Stop overrides admission

Action decision

Allow, deny or escalate.
Commit only if permitted.

Contain effects

Block new effects and hold.
Cancel or hand off work safely.

Safe state or fallback

Reach a defined safe state or transfer to validated fallback.

Signed receipt to episode graph

Record decisions and enforced stops.
Log asynchronously and track unresolved work.

Safe resumption

Verify recovery and obtain separate authorization.
Release the hold only after both checks pass.