


When an AI agent crosses the line: detect, contain and restore operations

What has to be in place when an AI agent does something it isn't allowed to do — in IT systems, IoT fleets, OT and on machines — and how is it detected, stopped and released again?



Where do you start — and which route is yours?

An agent did something nobody approved

Start at: Triage: what crossed which boundary

We are about to give an agent write access

Start at: The boundary in three parts

Our agent works with machines or robots

Start at: Robot cells and cobots

We want to restart after an incident

Start at: Release before return to service

FOUR ROUTES, STOP BY STOP · DARK SQUARES = SHARED TRUNK

1 · Technical writers

Technical writers and information developers

- 1 The boundary in three parts
- 2 CCMS, DMS and knowledge bases
- 3 Advisory agents on historian and MES
- 4 The containment ladder
- 5 Robot cells and cobots
- 6 CNC machines and production lines
- 7 Release before return to service
- 8 Review and update the documentation

2 · Technical marketers

Technical marketers and product communicators

- 1 The boundary in three parts
- 2 Scope, allowlist and own identity
- 3 Connected products and their apps
- 4 Device fleets and OTA updates
- 5 AMR and AGV fleets
- 6 The containment ladder
- 7 Check the reporting duties
- 8 Release before return to service

3 · Technical project managers

Technical project managers, product owners and operations leads

- 1 The boundary in three parts
- 2 Approval gates for irreversible actions
- 3 ERP, CRM and ticketing
- 4 Agents with write access to SCADA, DCS or recipes
- 5 Triage: what crossed which boundary
- 6 The containment ladder
- 7 Preserve the evidence
- 8 Release before return to service

4 · Documentation systems developers

Developers of documentation and information systems

- 1 The boundary in three parts
- 2 Agent trace and tool-call log
- 3 Signals of a boundary breach
- 4 Code repositories, CI/CD and cloud
- 5 Agents in PLC engineering
- 6 Building automation and energy management
- 7 The containment ladder
- 8 Release before return to service

WHAT DOES NOT CHANGE ON ANY ROUTE

- Responsibility for what an agent does stays with people and the operating company; an agent is never the accountable party.
- Containment stops further harm; it doesn't undo actions already taken.
- An agent stop is not an emergency stop: safety functions, emergency stops and safe states of machines stay in their control and safety systems.
- Reporting deadlines run from awareness, whether or not containment is finished.
- Naming a control, a standard or a law on this map never establishes compliance.

Limits: The map shows typical patterns per system type, not the design of a specific plant, product or IT landscape. Maturity labels are the editors' assessment as of October 21, 2026: "product" means some platforms offer the control, not that it is established practice. Standards for agent identity and delegation are still being drafted. The map is no substitute for the risk assessment of a machine, the security risk assessment of a product or plant, or legal advice on reporting duties.

GET IN TOUCH



Do you have questions?
Get in touch!

contact@aitechdoc.world

The full map with every station, its evidence and the mini course:
knowledge.aitechdoc.world/context-cards/ai-agent-boundary-breach-detect-contain-restore