

# How to Give a Surgical or Industrial Robot a Safe-Idle Fallback

If you build a robot that touches the physical world, you eventually face the same question: what should the machine do when it is no longer sure it is safe to act? This guide describes an architectural approach in which a governor computes execution readiness before every actuation and defaults to a safe, non-acting mode when readiness is insufficient. It is drawn from an architecture disclosed in U.S. Provisional Application No. 64/049,409, not a shipping library, and it centers on the Confidence-Governed Actuation inventive step.

---

## What You Are Building

You are building a control path that sits between a robot's planner and its actuators, so that no physical action reaches a motor, arm, cutter, valve, or infusion pump until a governor has decided the action is safe enough to take right now. When the governor cannot establish that confidence, the system defaults to a safe, non-acting state rather than executing the pending command anyway.

This is the problem faced by anyone shipping a surgical tool actuator, an industrial manipulator, an infusion pump, or a mobile robot that shares space with people. The planner is usually optimistic: it computes a trajectory and issues it. But sensing degrades, an operator enters the work envelope, a needed observation goes stale, or the

environment shifts in a way the plan did not anticipate. In those moments you do not want the last confident command to keep running. You want the machine to fall back to something safe, and you want a record of why.

The architecture described here comes from the Confidence-Governed Actuation inventive step, disclosed in U.S. Provisional Application No. 64/049,409. It is an approach you implement yourself, not a package you install.

## **Why the Obvious Approaches Fall Short**

The two conventional approaches are the interlock and the emergency stop. An interlock is a hard condition wired ahead of an actuator: if a guard door is open, the motor cannot run. An emergency stop is an operator-triggered or fault-triggered cut that halts motion. Both are real, useful, and standard functional-safety practice. Neither is what this guide replaces, and neither should be removed.

The structural gap is that these mechanisms are binary and uncredentialed. An interlock either permits or blocks; there is no middle behavior between full execution and complete cessation. An e-stop is a single signal with no notion of who issued it, whether that authority was legitimate, or how often it has been invoked. And critically, conventional sensor-to-actuator control terminates at the actuator command: the system issues the command and moves on, without comparing the observed effect against the expected effect, and without leaving a reconstructable record of why the action was permitted.

So when confidence quietly erodes, a binary interlock sees nothing to trip on, and the last confident command keeps executing. The disclosed architecture addresses this by making every actuation pass through a graduated, credentialed evaluation, with a safe non-acting mode as the natural low-confidence outcome rather than an exceptional one.

## The Architecture

The spec describes a confidence-governed execution primitive. Its core idea is that every physical actuation, of a braking system, a manipulator arm, a surgical tool actuator, a medical-dispensing actuator, a valve, and so on, is treated as a governed action that must be evaluated before it executes. The evaluation is done by a composite admissibility evaluator that considers the proposed actuation jointly with the observations the unit has consumed and other cognitive inputs. This is the physical-space extension of a confidence governor that, in the cited work, modulates an agent's readiness to act.

Four elements make the safe-idle fallback work.

First, **graduated actuation modes**. Instead of permit-or-deny, the disclosed mode selector maps the admissibility determination onto a range of modes. The spec enumerates, among others: a disabled mode where the actuator does not execute and the non-execution is recorded; a simulated mode that dry-runs without physical effect; an advisory mode that records what would have been done; a consultative mode that requests confirmation from a human operator or higher-authority endpoint and waits; a partial mode at reduced magnitude, rate, or scope; a constrained mode subject to added limits; a stage-gated mode that executes in stages with re-evaluation between them; and a full mode at nominal execution. As confidence rises the selector moves toward more autonomous modes; as it falls, the selector moves toward less autonomous modes. Your safe-idle fallback is the disabled mode: the low-confidence terminus of this ladder.

Second, **per-actuator confidence thresholds**. Each actuator carries thresholds that gate execution at each mode. A high-consequence actuator, the spec names propulsion in traffic, weapon-system terminal actuators, and medical-dispensing actuators, requires elevated admissibility across multiple contributing factors before it executes at full mode, while a low-consequence actuator such as an indicator light requires only

baseline admissibility. Thresholds are dynamically modulated: an actuator with a depleted capability envelope, or operating where forecast uncertainty is high, runs under elevated thresholds. This is how "unsure" becomes "stay idle" automatically.

Third, **execution readiness that degrades and recovers**. The spec describes a confidence governor that reduces execution readiness proportionally as governance-credentialed inputs are lost at runtime, down to an infrastructure-denied mode operating on the unit's own sensors with governance-policy-defined conservatism, and restores readiness as corroborating observations accumulate again. In-progress actuations can be de-escalated mid-execution to a constrained, partial, simulated, or disabled mode when a newly arriving input indicates reduced admissibility.

Fourth, **reversibility awareness**. A commitment-point evaluator classifies actuations by reversibility and identifies, for a staged actuation, the stage beyond which continuation becomes irreversible. Before that commitment point, a drop in admissibility can interrupt the sequence without incurring the irreversible final stages. For a cutting tool or a dose, this is the difference between backing off and committing.

Throughout, the spec requires lineage recording: every evaluation, every mode selection, every de-escalation, every commitment-point transit, and the post-actuation verification outcome are recorded so the decision can be reconstructed afterward. A post-actuation verification step compares observed effects against expected effects and can raise thresholds for that actuator when discrepancies recur.

## **How to Approach the Build**

The following steps are illustrative and faithful to the disclosed structure. You are implementing them yourself.

1. **Interpose a governor between planner and driver.** No actuation command should reach a physical driver except through this stage. The proposed actuation identifies an actuator, a command, and parameters.
2. **Define your admissibility inputs.** Decide what evidence feeds the confidence determination: fresh sensor observations, staleness of required inputs, whether an operator is in the envelope, capability-envelope state, forecast uncertainty. The spec treats admissibility as composite; you choose the contributing factors.
3. **Enumerate modes for each actuator class and pick your idle mode.** At minimum, wire disabled, some reduced modes (partial, constrained, simulated), a consultative mode, and full. Your safe-idle fallback is the disabled mode. Decide what "idle" physically means for each actuator, holding position, retracting, or ceasing dispensing.

```
# Illustrative only; faithful to the disclosed mode-selection structure.
def select_mode(actuation, admissibility, thresholds):
    if admissibility < thresholds.disabled_floor:
        return "disabled"          # safe-idle fallback: do not execute, rec
    if admissibility < thresholds.constrained:
        return "consultative"      # ask a human / higher authority, then wa
    if admissibility < thresholds.full:
        return "partial"           # reduced magnitude / rate / scope
    return "full"
```

4. **Set per-actuator thresholds by consequence, and modulate them.** Give high-consequence actuators elevated floors. Raise thresholds when the capability envelope is depleted or forecast uncertainty is high, so degraded conditions push the selector toward idle without hand-written per-case rules.
5. **Make readiness degrade at runtime.** When credentialed inputs are lost, lower readiness and let the raised thresholds pull actions down the mode ladder. Support mid-execution de-escalation so an action already running drops to a reduced or

disabled mode on a new adverse input.

6. **Classify reversibility and detect the commitment point.** For staged, irreversible actuations, identify the last stage before continuation is irreversible, and allow interruption before it. Prefer reversible or late-commitment paths where admissible.
7. **Record lineage and verify effects.** Write each decision, mode selection, and de-escalation to a provenance record. After execution, compare observed against expected effect, and let recurring discrepancies raise that actuator's thresholds.

## **What This Does Not Give You**

This is an architecture, not a drop-in library, and there is no package to install. Every component above is something you build and integrate into your own control stack. The disclosure is a patent filing describing the approach; it is not benchmarked, not certified to any safety-integrity level, and not a production-proven product. Nothing here reports performance numbers, because the filing does not, and you should not assume any.

It does not replace your interlocks, your e-stops, or your functional-safety certification obligations. Treat it as a governance layer above those mechanisms, not a substitute for them. It gives you a graduated, recorded, confidence-driven path to a safe-idle default; it does not by itself guarantee any particular safety outcome, and the correctness of your thresholds, your admissibility inputs, and your definition of "idle" for each actuator is entirely your responsibility. Where an actuator has no meaningful non-acting state, the disabled mode is not automatically the safe choice, and you must reason about that case yourself.

## Disclosure Scope

The architecture in this guide, including the confidence-governed execution primitive, graduated actuation modes, per-actuator confidence thresholds, reversibility-aware commitment-point evaluation, and lineage-recorded actuation provenance, is disclosed in U.S. Provisional Application No. 64/049,409. This guide is educational. It describes an approach a developer can build; it is not a warranty, not a certification, and not an offer of software. Any implementation, and its fitness and safety for a given deployment, is the responsibility of the implementer.

---

## **Confidence-Governed Actuation** (</gover> [All 40 steps → \(/inventive-steps\)](#))

### **ned-actuation)**

Eleven graduated modes. Preemption budget. Reversibility-aware commitment. Harm-minimization.

Provisional application

### **PRIMARY TECHNICAL DISCLOSURE**

- [Confidence-Governed Actuation: Graduated Modes for Physical Systems \(/articles/confidence-governed-actuation-graduated-modes-for-physical-systems\)](/articles/confidence-governed-actuation-graduated-modes-for-physical-systems)

### **SECONDARY TECHNICAL**

- [The Graduated Actuation Mode Selector \(/articles/governed-actuation/graduated-mode-set\)](/articles/governed-actuation/graduated-mode-set)
- [Preemption Budget for Rate-Limited Override \(/articles/governed-actuation/preemption-budget\)](/articles/governed-actuation/preemption-budget)
- [Reversibility Classifier \(/articles/governed-actuation/reversibility-classifier\)](/articles/governed-actuation/reversibility-classifier)
- [Stage-Gated Commitment for Irreversible Actions \(/articles/governed-actuation/staged-commitment\)](/articles/governed-actuation/staged-commitment)
- [Harm-Minimization Deviation \(/articles/governed-actuation/harm-minimization-deviation\)](/articles/governed-actuation/harm-minimization-deviation)
- [Policy-Configurable Harm Ordering \(/articles/governed-actuation/policy-configurable-harm-ordering\)](/articles/governed-actuation/policy-configurable-harm-ordering)
- [Post-Actuation Verification \(/articles/governed-actuation/post-actuation-verification\)](/articles/governed-actuation/post-actuation-verification)

- [Discrepancy Classification Taxonomy \(/articles/governed-actuation/discrepancy-classification\)](/articles/governed-actuation/discrepancy-classification)
- [Actuation State as Mesh-Broadcast Observation \(/articles/governed-actuation/actuation-state-broadcast\)](/articles/governed-actuation/actuation-state-broadcast)
- [Cross-Unit Coordination Through Actuation-State Broadcast \(/articles/governed-actuation/cross-unit-coordination\)](/articles/governed-actuation/cross-unit-coordination)
- [Lineage-Recorded Provenance for Every Actuation \(/articles/governed-actuation/lineage-recorded-provenance\)](/articles/governed-actuation/lineage-recorded-provenance)

## APPLICATIONS · GENERAL

- [Autonomous Aviation Execution Under Governed Actuation \(/articles/governed-actuation/autonomous-aviation-execution\)](/articles/governed-actuation/autonomous-aviation-execution)
- [Compliant Autonomous Defense Engagement: Governing the Kill Chain With Admissibility-Gated Actuation \(/articles/governed-actuation/autonomous-defense-execution\)](/articles/governed-actuation/autonomous-defense-execution)
- [Autonomous Surgical Execution Under Governed Actuation \(/articles/governed-actuation/autonomous-surgical-execution\)](/articles/governed-actuation/autonomous-surgical-execution)
- [FDA PCCP for Surgical Robots: An Architecture for Change-Control Eligibility \(/articles/governed-actuation/fda-pccp-surgical-pathway\)](/articles/governed-actuation/fda-pccp-surgical-pathway)
- [Auditable Harm-Minimization for Autonomous Vehicles: Declared Harm Ordering, Not Emergent Heuristics \(/articles/governed-actuation/harm-minimization-autonomous-driving\)](/articles/governed-actuation/harm-minimization-autonomous-driving)
- [Certifying Autonomous Aircraft Under ICAO: Structural Evidence for Phase, Authority, and Sense-and-Avoid \(/articles/governed-actuation/icao-autonomous-aviation\)](/articles/governed-actuation/icao-autonomous-aviation)
- [Cobot Safety Compliance Under ISO 10218 and ISO/TS 15066: Architecting Auditable Collaborative-Robot Mode Decomposition \(/articles/governed-actuation/industrial-collab-robotics-iso\)](/articles/governed-actuation/industrial-collab-robotics-iso)
- [Preemption Budget for Multi-Vehicle Fleet Operations \(/articles/governed-actuation/preemption-budget-fleet-operations\)](/articles/governed-actuation/preemption-budget-fleet-operations)
- [14 CFR Part 107 Commercial Drone Operations \(/articles/governed-actuation/cfr-part-107-drones\)](/articles/governed-actuation/cfr-part-107-drones)
- [Enforcing 14 CFR Part 135 Operational Control in Software for Commuter, Charter, and eVTOL Air Carriers \(/articles/governed-actuation/cfr-part-135-commercial-air\)](/articles/governed-actuation/cfr-part-135-commercial-air)
- [Certifying Adaptive Airborne Software Under DO-178C with Governed Actuation \(/articles/governed-actuation/do-178c-airborne-software\)](/articles/governed-actuation/do-178c-airborne-software)
- [FDA 510\(k\) and De Novo Pathway for Autonomous Medical Devices \(/articles/governed-actuation/fda-510k-de-novo\)](/articles/governed-actuation/fda-510k-de-novo)
- [Meeting IEC 61508 Safe-State and Graduated-Response Requirements with Governed Actuation \(/articles/governed-actuation/iec-61508-functional-safety\)](/articles/governed-actuation/iec-61508-functional-safety)
- [ISO 14971 Risk Management for Autonomous Medical Devices \(/articles/governed-actuation/iso-14971-medical-risk\)](/articles/governed-actuation/iso-14971-medical-risk)

- [ISO/SAE 21434 Compliance: Binding Automotive Cybersecurity to Runtime Control \(/articles/governed-actuation/iso-21434-automotive-cyber\)](/articles/governed-actuation/iso-21434-automotive-cyber).
- [Meeting ISO 21448 SOTIF for Autonomous Driving with Governed Actuation \(/articles/governed-actuation/iso-21448-sotif\)](/articles/governed-actuation/iso-21448-sotif).
- [ISO 26262 Functional Safety for Autonomous Driving \(/articles/governed-actuation/iso-26262-functional-safety\)](/articles/governed-actuation/iso-26262-functional-safety).
- [NHTSA ADS-DV Requirements and AV STEP \(/articles/governed-actuation/nhtsa-ads-dv-requirements\)](/articles/governed-actuation/nhtsa-ads-dv-requirements).

## APPLICATIONS · SPECIFIC

- [Aurora Driver vs Governed Actuation for Reversibility-Aware Execution \(/articles/governed-actuation/aurora-self-driving\)](/articles/governed-actuation/aurora-self-driving)
- [Governed Surgical Autonomy Beyond Intuitive da Vinci \(/articles/governed-actuation/intuitive-davinci\)](/articles/governed-actuation/intuitive-davinci).
- [Waymo Driver vs Governed Actuation: The Externally Verifiable Commitment Gate \(/articles/governed-actuation/waymo-driver\)](/articles/governed-actuation/waymo-driver)
- [CMR Versius and Governed Actuation for Partial Surgical Autonomy \(/articles/governed-actuation/cmr-versius\)](/articles/governed-actuation/cmr-versius).
- [Cruise Robotaxi vs Governed Actuation: A Stage-Gated, Reversibility-Aware Substrate \(/articles/governed-actuation/cruise-self-driving\)](/articles/governed-actuation/cruise-self-driving).
- [Kodiak Robotics vs Governed Actuation for Autonomous Trucking \(/articles/governed-actuation/kodiak-self-driving-trucking\)](/articles/governed-actuation/kodiak-self-driving-trucking)
- [Mobileye SuperVision vs Governed Actuation: The Tier-Transition Gap \(/articles/governed-actuation/mobileye-supervision\)](/articles/governed-actuation/mobileye-supervision).
- [Tesla FSD Supervised vs Governed Actuation: The Stage-Architecture Gap \(/articles/governed-actuation/tesla-fsd-supervised\)](/articles/governed-actuation/tesla-fsd-supervised)
- [Zoox Alternative: Governed Actuation Beyond the Purpose-Built Robotaxi \(/articles/governed-actuation/zoox-amazon\)](/articles/governed-actuation/zoox-amazon).
- [Governed Actuation Beyond Archer Aviation Midnight eVTOL \(/articles/governed-actuation/archer-maker\)](/articles/governed-actuation/archer-maker).
- [Asensus Senhance vs Governed Surgical Actuation \(/articles/governed-actuation/asensus-senhance\)](/articles/governed-actuation/asensus-senhance)
- [BMW Personal Pilot vs Governed Actuation \(/articles/governed-actuation/bmw-personal-pilot\)](/articles/governed-actuation/bmw-personal-pilot)
- [Boston Dynamics Atlas vs Governed Humanoid Actuation \(/articles/governed-actuation/boston-dynamics-atlas\)](/articles/governed-actuation/boston-dynamics-atlas)
- [Governed Actuation Beyond Boston Dynamics Spot \(/articles/governed-actuation/boston-dynamics-spot\)](/articles/governed-actuation/boston-dynamics-spot)

- [Boston Dynamics Stretch Alternative: Governed Actuation for Warehouse Robotics \(/articles/governed-actuation/boston-dynamics-stretch\)](#)
- [EHang EH216-S vs Governed Actuation: Audit-Time Evidence for Autonomous eVTOL \(/articles/governed-actuation/ehang-evtol\)](#)
- [Eve Air Mobility eVTOL Alternative: Governed Actuation Beyond the Flight Control Law \(/articles/governed-actuation/eve-air-mobility\)](#)
- [Ford BlueCruise vs Governed Actuation: Beyond Per-Feature Gating \(/articles/governed-actuation/ford-bluecruise\)](#)
- [GM Super Cruise vs Governed Actuation: Stage-Gated L2+/L3 Commitment \(/articles/governed-actuation/gm-super-cruise\)](#)
- [Honda SENSING 360 alternative: governed actuation beyond mode-switch ADAS \(/articles/governed-actuation/honda-sensing-360\)](#)
- [Hyundai HDA and Highway Driving Pilot vs Governed Actuation \(/articles/governed-actuation/hyundai-hda\)](#)
- [Ottava vs Governed Surgical Autonomy: Actuation-Authority Decomposition Beyond a Single-Snapshot Review \(/articles/governed-actuation/jnj-ottava\)](#)
- [Joby Aviation and Governed eVTOL Actuation: A Beyond-Flight-Control Alternative \(/articles/governed-actuation/joby-aviation\)](#)
- [Lilium Jet alternative: governed actuation beyond fly-by-wire \(/articles/governed-actuation/lilium-jet\)](#)
- [Medtronic Hugo Alternative: Governed Actuation for Bounded Surgical Autonomy \(/articles/governed-actuation/medtronic-hugo\)](#)
- [Mercedes Drive Pilot vs Governed Actuation: the Certified L3 Hand-Back \(/articles/governed-actuation/mercedes-drive-pilot\)](#)
- [Rivian Driver+ Alternative: Governed Actuation Beyond L2 Handoff \(/articles/governed-actuation/rivian-driver-plus\)](#)
- [Smith Nephew CORI vs Governed Surgical Actuation \(/articles/governed-actuation/smith-nephew-cori\)](#)
- [Stellantis STLA AutoDrive vs Governed Actuation: A Multi-Brand, Multi-Region Level 3 Alternative \(/articles/governed-actuation/stellantis-stla-autodrive\)](#)
- [Stryker Mako vs Governed Surgical Actuation: A Reversibility-Aware Alternative \(/articles/governed-actuation/stryker-mako\)](#)
- [Toyota Teammate Alternative: Governed Actuation Beyond Binary Driver Handoff \(/articles/governed-actuation/toyota-teammate\)](#)
- [Vertical Aerospace VX4 vs governed actuation \(/articles/governed-actuation/vertical-aerospace\)](#)
- [Volocopter VoloCity Air Taxi vs Governed Actuation \(/articles/governed-actuation/volocopter-air-taxi\)](#)

- [Wing Delivery Drones vs Governed Actuation: A Governed Alternative for BVLOS Flight \(/articles/governed-actuation/wing-delivery-drones\)](#)
- [Wisk Aero vs Governed Actuation: A Certification Substrate for Autonomous eVTOL \(/articles/governed-actuation/wisk-aero\)](#)
- [Zimmer Biomet ROSA vs Governed Surgical Actuation \(/articles/governed-actuation/zimmer-rosa\)](#)
- [Zipline Drone Delivery vs Governed Actuation for Autonomous Aircraft \(/articles/governed-actuation/zipline-drone-delivery\)](#)
- [Does AGCO Fendt IDEAL and PTx OutRun Have Governed Actuation for Autonomous Farming? \(/articles/governed-actuation/agco-fendt-ideal\)](#)
- [Bear Robotics Servi vs Governed Actuation: A Serving Robot That Can Defer, Constrain, and Roll Back \(/articles/governed-actuation/bear-robotics-servi\)](#)
- [Built Robotics Alternative: Governed Actuation for Autonomous Construction Equipment \(/articles/governed-actuation/built-robotics-construction\)](#)
- [Governed Actuation Beyond Cat MineStar Autonomous Mining \(/articles/governed-actuation/caterpillar-minestar\)](#)
- [CNH Industrial Autonomous Agriculture vs Governed Actuation \(/articles/governed-actuation/cnh-industrial-autonomous\)](#)
- [Governed Actuation Beyond Embark and Locomotion: An Autonomous-Trucking Alternative \(/articles/governed-actuation/embark-locomotion-legacy\)](#)
- [Epiroc Mobilaris vs Governed Actuation for Mining Fleets \(/articles/governed-actuation/epiroc-mobilaris\)](#)
- [Gatik Alternative: Governed Actuation for Middle-Mile Trucking \(/articles/governed-actuation/gatik-trucking\)](#)
- [John Deere See and Spray alternative: governed actuation beyond selective spraying \(/articles/governed-actuation/john-deere-see-spray\)](#)
- [Komatsu FrontRunner Alternative: Governed Actuation for Autonomous Haulage \(/articles/governed-actuation/komatsu-frontrunner\)](#)
- [Plus autonomous trucking vs governed actuation: a Class 8 alternative frame \(/articles/governed-actuation/plus-ai-trucking\)](#)
- [Sandvik AutoMine vs Governed Actuation: A Commitment-Layer Alternative for Underground Autonomy \(/articles/governed-actuation/sandvik-automine\)](#)
- [Governed Autonomous Trucking Beyond Torc Robotics \(/articles/governed-actuation/torc-robotics-trucking\)](#)
- [TuSimple Autonomous Trucking vs Governed Actuation \(/articles/governed-actuation/tusimple-trucking\)](#)
- [Waabi Foundation Model Trucking vs Governed Actuation \(/articles/governed-actuation/waabi-trucking\)](#)

- [ABB Robotics vs Governed Actuation \(/articles/governed-actuation/abb-robotics\)](/articles/governed-actuation/abb-robotics)
- [FANUC Robotics vs Governed Actuation \(/articles/governed-actuation/fanuc-robotics\)](/articles/governed-actuation/fanuc-robotics)
- [Kiwibot Delivery Robots vs Governed Actuation: A Sidewalk-Robot Alternative \(/articles/governed-actuation/kiwibot-delivery\)](/articles/governed-actuation/kiwibot-delivery)
- [KUKA Robotics vs Governed Actuation: The Commitment Gate Above the Controller \(/articles/governed-actuation/kuka-robotics\)](/articles/governed-actuation/kuka-robotics)
- [Lyft Autonomous vs Governed Actuation: A Marketplace Dispatch Gate Beyond Vendor Availability \(/articles/governed-actuation/lyft-autonomous-magna\)](/articles/governed-actuation/lyft-autonomous-magna)
- [Governed Actuation vs Nuro Autonomous Delivery \(/articles/governed-actuation/nuro-delivery\)](/articles/governed-actuation/nuro-delivery)
- [Starship Technologies Alternative: Governed Sidewalk-Delivery Actuation \(/articles/governed-actuation/starship-technologies\)](/articles/governed-actuation/starship-technologies)
- [Governed Dispatch Beyond Uber Freight Autonomous Trucking \(/articles/governed-actuation/uber-freight-autonomous\)](/articles/governed-actuation/uber-freight-autonomous)
- [Universal Robots Cobots vs Governed Actuation: Admissibility-Gated Cobot Execution \(/articles/governed-actuation/universal-robots-cobots\)](/articles/governed-actuation/universal-robots-cobots)
- [Yaskawa Motoman Alternative: Governed Actuation Beyond Certified Motion \(/articles/governed-actuation/yaskawa-robotics\)](/articles/governed-actuation/yaskawa-robotics)

---

[Confidence-Governed Actuation overview → \(/governed-actuation\)](/governed-actuation)