

# Independent Clearance Before Consequence

Why accepted institutional control principles require an independent determination at the Execution Boundary

Mohammed Thameem

Founder and Director, OATHOR LTD, Abu Dhabi Global Market, United Arab Emirates

Version 1.0 | October 2026

---

## Abstract

Automated and AI systems increasingly initiate, route and carry out actions with operational, institutional, financial, legal or public consequence. This paper asks under what institutional conditions accountable consequential digital execution necessarily requires a determination, made independently of the acting system and of the function or provider that operates or supplies it, of whether each consequential action is cleared to proceed before its consequence exists.

The argument proceeds from principles that institutions, supervisors and engineers already apply: institutional accountability for outsourced and automated activity; segregation of duties; the self-review threat recognised in audit independence; the independence of protective functions from the systems they protect in functional safety; complete mediation in computer security; the time-of-check to time-of-use weakness; effective human oversight; and contemporaneous records. Securities regulation already applies several of these principles to automated actors through pre-order controls on algorithmic trading.

From these premises the paper derives seven requirements, each through the same five steps: the accepted condition, why existing controls may not satisfy it, why the property is required, what counts as a valid alternative, and where the argument does not apply. It then shows that, where the conditions of all seven are present, the seven are jointly necessary: a control that meets six but not the seventh leaves the institution's obligation unmet. Only after the derivation does the paper show that the seven requirements correspond to the published criteria ECI-1 to ECI-7 for Execution Control Infrastructure. The paper concerns requirements, not any product or implementation. It states where the requirements hold, where they weaken, and what evidence would refute them.

---

## 1. Introduction

Institutions have long separated the ability to act from the authority to let an action take effect. A payment is prepared by one officer and released by another. An order sent by a client's algorithm passes the broker's own checks before it reaches the market. A safety system in a process plant is kept independent of the control system it protects. In each case the separation exists because the consequence matters and cannot simply be undone.

Automation changes who, or what, prepares and carries out consequential actions. Software agents can now call payment functions, send communications, change records, commit organisations to terms and act through tools and other agents. When such an action is wrong, the question that matters is whether

anything independent of the acting system determined, before the consequence existed, that the specific action was cleared to proceed.

The controls relied upon for automated actions may sit inside the acting system, be configured and updated by the provider that supplies it, or depend on information and presentations that the acting system itself produces. This paper examines when that is sufficient and when it is not.

**What the paper claims.** Under the conditions stated in Section 7, principles that institutions and supervisors already accept require a determination that is made outside the acting system, beyond the actor's reach, free of control by the function or provider that operates or supplies the acting system, made before consequence for the specific action, applied as a matter of course, and recorded at the time. These seven requirements are jointly necessary where the conditions of all seven are present.

**What the paper does not claim.** It does not claim that every automated action requires such a determination. It does not claim that meeting the requirements guarantees that every determination is correct; the requirements describe the structure of independent control, not the quality of each decision, and adjacent controls remain necessary. It does not prescribe any technology, vendor, deployment arrangement, organisational structure or implementation method. And it does not claim that earlier work was unaware of separated or pre-execution control; Section 9 sets out that work.

---

## 2. Scope and working terms

The argument concerns consequential digital action across operational, institutional, financial, legal and public settings, including digital actions whose consequence is physical. It is not limited to AI agents or to any industry.

The following working terms are used in the derivation. Section 8 shows how they correspond to terms that OATHOR LTD has published and dated.

- **Consequential action:** a requested digital action whose effect, once it occurs, matters to the institution or to others and may not be fully reversible.
  - **Acting system:** every component that requests, prepares or carries out the action, including models, agents, workflows, orchestration, tools and the runtime in which they execute, and every system in a chain of delegation.
  - **Actor:** the acting system, anyone who requests the action or directs the acting system, and the instructions or content the acting system receives.
  - **Operating function or provider:** the internal function that operates the acting system, or the external provider that supplies it.
  - **Determination:** the decision whether a specific consequential action may take effect; only an action so cleared proceeds.
  - **Institutional requirement:** the institution's obligation, for an action for which it is accountable, to prevent rather than only detect consequence that cannot be fully reversed, with segregation of duties and evidence of control.
  - **Final preventable point:** the last point before an action's consequence exists at which the action can still be prevented.
-

### 3. Accepted premises

Each premise is stated in terms that do not depend on OATHOR's framework and is supported by sources from banking, market infrastructure, audit, information security, functional safety, pharmaceutical regulation or AI regulation.

**P1. Accountability remains with the institution.** An institution remains responsible for consequential actions carried out on its behalf, including through third parties and automated systems. Under the EU Digital Operational Resilience Act, financial entities using ICT third-party service providers "remain fully responsible for compliance with, and the discharge of, all obligations" (Regulation (EU) 2022/2554, Article 28(1)(a)). The EBA outsourcing guidelines state that "the outsourcing of functions cannot result in the delegation of the management body's responsibilities" (EBA/GL/2019/02, paragraph 35). The CPMI-IOSCO discussion paper on third-party reliance records that jurisdictions typically noted or implied that financial market infrastructures "retain ultimate responsibility for managing risks" from activities provided by third-party service providers (CPMI-IOSCO, 2026). The EU AI Act places obligations on deployers of high-risk AI systems as well as providers (Regulation (EU) 2024/1689, Article 26).

**P2. Segregation of duties.** No single party should both initiate a consequential action and approve it. The Basel Committee states that "an effective internal control system requires that there is appropriate segregation of duties and that personnel are not assigned conflicting responsibilities" (Basel Committee, 1998, Principle 6). Segregation of duties is also a control in information security management (ISO/IEC 27001:2022, Annex A, control 5.3) and a point of focus for control activities in internal control frameworks (COSO, 2013, Principle 10).

**P3. The self-review threat.** A party should not provide independent assurance over work it has produced or implemented. The IESBA Code defines the self-review threat as "the threat that a professional accountant will not appropriately evaluate the results of a previous judgment made, or an activity performed by the accountant or by another individual within the accountant's firm or employing organization" (IESBA, paragraph 120.6 A3(b)). Audit independence rules accordingly treat the design and implementation of an audit client's financial information systems or internal controls as incompatible with independent assurance (SEC Regulation S-X, Rule 2-01(c)(4)(ii); Regulation (EU) No 537/2014, Article 5(1), point (e)).

**P4. Prevention where consequence cannot be reversed.** Internal control distinguishes preventive control activities, which stop an undesired event, from detective ones, which identify it afterwards (COSO, 2013). Where a consequence cannot be fully reversed, detection leaves the consequence in place.

**P5. A check holds only for the conditions under which it was made.** Information security recognises the time-of-check to time-of-use weakness, in which a condition verified at one moment no longer holds when the resource is used (MITRE, CWE-367).

**P6. Complete mediation.** In computer security, "every access to every object must be checked for authority" (Saltzer and Schroeder, 1975). The reference monitor concept requires that the reference validation function "must always be invoked" and "must be tamper proof" (Anderson, 1972).

**P7. Independence of protective functions.** Functional safety requires that a safety instrumented system "be designed to be separate and independent from the BPCS [basic process control system] to the extent that the safety integrity of the SIS is not compromised", and that the design "take into consideration all aspects of independence and dependency between the SIS and BPCS" (IEC 61511-1:2016, clauses 11.2.4 and 11.2.9, as quoted by the UK Health and Safety Executive). Nuclear safety requires that "the levels of defence in depth shall be independent as far as is practicable" (IAEA, SSR-2/1 (Rev. 1), Requirement 7). The common

principle is that a protective function should be independent of the system it protects to the extent that a shared dependency could otherwise cause both to fail together.

**P8. Oversight must be able to intervene.** The EU AI Act requires that natural persons overseeing high-risk AI systems be enabled, as appropriate and proportionate, "to intervene in the operation of the high-risk AI system or interrupt the system through a 'stop' button or a similar procedure" (Regulation (EU) 2024/1689, Article 14(4)(e)).

**P9. Records are made at the time.** Data integrity guidance expects data to be recorded contemporaneously, at the time the activity is performed (FDA, 2018), and requires computer systems to be "designed to ensure that the execution of critical steps is recorded contemporaneously" (MHRA, 2018, section 6.12). The EU AI Act requires that high-risk AI systems "technically allow for the automatic recording of events" (Regulation (EU) 2024/1689, Article 12(1)).

**P10. Suppliers can change what they supply.** A provider or operating function can change a system's behaviour through updates, configuration and terms of service. Supervisory principles on third-party risk accordingly address prior notification of change by key service providers and their sub-contractors (BCBS principles, cited in CPMI-IOSCO, 2026, footnote 28).

**P11. Acting systems can be steered by the content they process.** Systems built on language models can be induced to take actions by instructions embedded in content they read, without any change to their code or permissions (Greshake et al., 2023).

**An existing application to automated actors.** Several of these principles are already applied by securities regulators to automated order flow. Under the US Market Access Rule, a broker-dealer must maintain controls that "prevent the entry of orders unless there has been compliance with all regulatory requirements that must be satisfied on a pre-order entry basis" (17 CFR 240.15c3-5(c)(2)(i)), including controls to prevent erroneous orders (paragraph (c)(1)(ii)), and these controls "shall be under the direct and exclusive control of the broker or dealer" (paragraph (d)), so that for customer orders the controls are under the broker's control rather than the customer's, subject to a limited allocation of specific regulatory controls to a customer that is itself a registered broker-dealer (paragraph (d)(1)). Under MiFID II, investment firms engaged in algorithmic trading "shall carry out the following pre-trade controls on order entry for all financial instruments" (Commission Delegated Regulation (EU) 2017/589, Article 15(1)) and must be able to cancel unexecuted orders immediately as an emergency measure (Article 12(1)). These rules do not require Execution Control Infrastructure. They show that regulators already accept, for one class of automated consequential actions, that checks must apply to each action before it takes effect, and, under the Market Access Rule, that for customer orders they must be under the broker's control rather than that of the customer whose system initiates the order.

---

## 4. The final preventable point

**Proposition.** For any consequential action with an identifiable point at which its consequence comes into existence, there is a last point before that consequence at which the action can still be prevented. After that point, control can detect, mitigate, reverse or compensate, but cannot prevent the original consequence.

This follows from the structure of the action. It requires no assumption about technology or institutions. Combined with P4, it implies that an institution that must prevent, rather than only detect, an irreversible consequence must have an effective control that operates no later than the final preventable point.

Two further observations complete the picture. First, systems that do not take the final step can still shape the action: they influence, initiate, route, approve, restrict, suspend, escalate or coordinate it. Their influence is relevant wherever it reaches the action that becomes consequence. Second, the relevant risk is that consequence proceeds without appropriate control before execution. Section 8 shows that these observations correspond to published terms.

---

## 5. Seven requirements

Each requirement below is derived through five steps: the accepted condition; why existing controls may not satisfy it; why the property is required; what counts as a valid alternative; and where the argument does not apply.

### 5.1 Requirement 1: made outside the acting system

**Condition.** The institution requires an independent check on the action (P2), and the check must not share a dependency that could cause it to fail together with the acting system (P7).

**Why existing controls may not satisfy it.** Guardrails, policy evaluations and safety instructions that run within the acting system share its runtime, its inputs and its exposure to manipulation (P11). They also share its change process: an update to the acting system can alter them at the same time (P10).

**Why the property is required.** A check that sits inside the acting system shares the acting system's failure modes. In the terms of functional safety, the protective function and the system it protects have a common cause of failure. Making the determination outside the acting system separates it from failure modes arising within that system. Requirements 2 and 3 address the separate questions of unilateral control by the actor and by its operating function or provider.

**Valid alternatives.** Any arrangement in which the determination is made by a component that does not share the acting system's runtime, inputs or change process meets the requirement, whatever it is called and however it is deployed. Naming a component as separate does not by itself place it outside; what matters is whether it shares the dependencies.

**Where it does not apply.** Where the action is fully and promptly reversible, or its consequence is negligible, a check within the acting system may be proportionate.

### 5.2 Requirement 2: beyond the actor's reach

**Condition.** Segregation of duties (P2) and complete mediation (P6) apply to the action.

**Why existing controls may not satisfy it.** A check located outside the acting system can still be within the actor's reach. The actor can defeat such a check in five ways, each the automated counterpart of a failure of segregation of duties:

- *Self-approval:* the actor determines the conditions or the outcome of its own clearance.
- *Override:* the actor alters the conditions that apply to it.
- *Self-attestation:* the actor satisfies a condition by its own assertion, or by information that only it supplies and that the determination does not establish independently.
- *Circumvention:* the actor bypasses the check, or the action can proceed without the check being applied or when no determination is made.

- *Substitution of the approver*: the actor puts another determination in place of the one relied upon.

A human approval can fall within the actor's reach where the person can establish the action only from what the actor presents, so that the actor controls a material condition of clearance.

**Why the property is required.** Where any of these is available to the actor by virtue of being the actor, the check provides no assurance that segregation of duties is meant to provide. A reference monitor that can be bypassed or tampered with is not a reference monitor (P6).

**Valid alternatives.** Any arrangement in which the actor cannot, by virtue of being the actor, determine, alter, satisfy, bypass or substitute the conditions of its own clearance meets the requirement. Submitting the action and its parameters for determination is not satisfying a condition by the actor's own assertion, where the action that proceeds is the action determined.

**Where it does not apply.** Where segregation of duties is not required for the action in question.

### 5.3 Requirement 3: free of control by the operating function or provider

**Condition.** The institution remains accountable for the action (P1), segregation of duties applies (P2), and the function or provider that operates or supplies the acting system can change it (P10).

**Why existing controls may not satisfy it.** A determination can be outside the acting system and beyond the actor's reach, yet remain under the control of the function or provider that operates or supplies the acting system: it may set the conditions, change them through updates or configuration, suspend the determination or override its outcome.

**Why the property is required.** Three accepted principles converge.

1. *Accountability (P1).* The institution remains fully responsible for the action. If the operating function or provider controls the conditions of clearance, the institution answers for outcomes whose clearance conditions another party can change unilaterally. Where the provider of the acting system can unilaterally set, change, suspend or override the conditions of clearance, the institution answers for outcomes it cannot itself control and cannot show that its retained responsibility is effective. This ground concerns external providers; for an internal operating function the requirement rests on grounds 2 and 3.
2. *Self-review (P3).* Assurance over a system is not independent when it is provided by the party that designed or implemented the system. A determination controlled by the party that operates or supplies the acting system is that party evaluating the actions of its own system.
3. *Shared dependency (P7, P10).* If one change process governs both the acting system and its clearance conditions, a single change, defect or compromise can alter both at once. That is the common-cause dependency functional safety requires protective functions to avoid.

The same reasoning applies to an internal operating function. Where an institution operates its own acting system, including a self-hosted or open-weight system, the function that operates it has the same ability to change it and the same relation to its outputs. Information security controls require segregation of conflicting duties, separation of development, test and production environments, and managed change (ISO/IEC 27001:2022, Annex A, controls 5.3, 8.31 and 8.32). Organisational ownership does not by itself make clearance independent of the function that operates the actor.

Two consequences follow, without being separate requirements. *Concentration*: where many institutions rely on controls governed by the same provider, a single change can alter clearance conditions across all of them at once, extending the concern about third-party concentration that supervisors already recognise for service availability (CPMI-IOSCO, 2026). *Jurisdictional exposure*: where the provider operates under another jurisdiction, provider control over clearance means the conditions for an institution's consequential actions can be changed from outside its legal and supervisory environment. This is a consequence of provider dependence, not a nationality test.

**Valid alternatives.** Independence here concerns authority over the determination, not organisational, technical or commercial isolation. A provider may host, integrate, support or supply technology used in the arrangement without controlling the determination. A relationship alone is not control. The question is who holds the practical ability to set, change, suspend or override the conditions or the outcome of the determination. A control operated abroad can meet the requirement; a locally operated control can fail it if the acting provider retains that ability.

**Where it does not apply.** Where the operating function or provider cannot change the acting system or its clearance conditions, or where the institution bears no accountability for the action.

#### 5.4 Requirement 4: before consequence

**Condition.** The institution must prevent the consequence, not only detect it (P4), and the consequence cannot be fully reversed.

**Why existing controls may not satisfy it.** Monitoring, post-execution review and audit operate after the final preventable point.

**Why the property is required.** By the proposition in Section 4, a determination made after the final preventable point cannot prevent the consequence. Where an action is carried out in several steps, the determination must precede the first step that is consequence. A preparatory step that has an effect outside the acting system that stopping the action cannot undo, such as disclosure of data to a third party or communication of a commitment, is a consequential action in its own right. The Market Access Rule applies the same logic to orders on a pre-order entry basis.

**Valid alternatives.** Any determination effective no later than the final preventable point.

**Where it does not apply.** Where the consequence is fully and promptly reversible.

#### 5.5 Requirement 5: the specific action

**Condition.** The institution's control objective concerns individual actions and the conditions that exist when they proceed (P5).

**Why existing controls may not satisfy it.** Permissions, mandates, policies and sessions apply to classes of actions and are established before the action. They show that actions of a class may be permitted, not that this action, with its actual parameters, is acceptable under present conditions.

**Why the property is required.** A check made earlier, or for a class, may no longer hold when the action proceeds (P5). The determination must concern the specific action, with its actual parameters, at the time it is to proceed. The MiFID II pre-trade controls apply on order entry, to each order. The requirement is that the action that proceeds is the action that was determined; this paper does not address how any implementation secures that correspondence, and a determination at the boundary does not by itself remove every change that might occur afterwards.

**Valid alternatives.** A determination made for each occurrence against a standing mandate meets the requirement. A determination over a group of actions meets it only where it decides each action in the group.

**Where it does not apply.** Where the institution's obligation concerns only the class of actions and not their individual instances.

### 5.6 Requirement 6: as a matter of course

**Condition.** The institution relies on the control for a class of consequential actions (P6).

**Why existing controls may not satisfy it.** Sampling, exception-based review and review triggered by the actor leave actions within the class without a determination.

**Why the property is required.** Complete mediation requires that every access be checked, not a sample. The Market Access Rule applies its controls to orders on a pre-order entry basis, not by sampling. Where the institution relies on a control for a class of consequential actions, every action in that class must reach the determination. A label such as low-risk or routine does not by itself take an action outside the class where the action otherwise becomes consequence. Practices such as maker-checker are often applied above thresholds; the requirement does not forbid defining the class. It requires that the class be stated, that exclusions be recorded, and that no consequential action within the class relied upon proceeds without a determination.

**Valid alternatives.** Any arrangement in which every action within the class relied upon reaches the determination, however the class is defined.

**Where it does not apply.** Outside the class of actions for which the control is relied upon.

### 5.7 Requirement 7: recorded at the time

**Condition.** The institution must be able to show that each determination occurred, when, and for which action (P1, P9).

**Why existing controls may not satisfy it.** Records assembled afterwards from logs, summaries or recollection can be incomplete, and can be shaped by parties with an interest in the outcome, including the actor and its operating function or provider.

**Why the property is required.** Data integrity principles require records of critical steps to be made at the time of performance (P9). Accountability for a determination depends on evidence that it existed before the consequence. The record must be created contemporaneously with the determination and must identify it sufficiently to show that it existed at that time.

**Valid alternatives.** Any form of record that is created contemporaneously and identifies the determination meets the requirement. The paper does not prescribe format, storage or technology.

**Where it does not apply.** Where the institution has no need to evidence individual determinations.

---

## 6. Why the seven belong together

The seven requirements are not independent preferences. Each closes a way in which the institutional requirement fails even when the other six are met. The table describes, for each, a control that meets the other six but not that one.

| Missing requirement                               | A control meeting the other six                                                                                       | What fails                                                                                                                                                                                                                                                        |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. Outside the acting system                      | Meets Requirements 2 to 7, but the determination is performed as one of the acting system's own components            | The determination remains part of the system it decides on; naming the component separately does not take it outside. Any failure, manipulation or change of that system, whatever its source, can therefore reach the determination and the action together (P7) |
| 2. Beyond the actor's reach                       | Separate component applied to every action, but the actor can satisfy its conditions by its own assertion             | Segregation of duties (P2)                                                                                                                                                                                                                                        |
| 3. Free of provider or operating-function control | Separate and beyond the actor, but its conditions are set and updated by the acting system's provider                 | Accountability (external provider), self-review, shared dependency (P1, P3, P7)                                                                                                                                                                                   |
| 4. Before consequence                             | Independent, action-specific and recorded at the time, but its outcome arrives only after the final preventable point | Prevention of irreversible consequence (P4)                                                                                                                                                                                                                       |
| 5. Specific action                                | Independent and timely, but decides only whether the actor is permitted                                               | Time-of-check to time-of-use (P5)                                                                                                                                                                                                                                 |
| 6. As a matter of course                          | Independent, but applied by sampling or on exception                                                                  | Complete mediation (P6)                                                                                                                                                                                                                                           |
| 7. Recorded at the time                           | Independent and complete, but evidenced by reconstruction                                                             | Contemporaneous records and accountability (P1, P9)                                                                                                                                                                                                               |

The claim is joint necessity under the stated conditions. It is not a claim of sufficiency: a control meeting all seven can still make a wrong determination because of defective conditions, incomplete information or a governance failure. The seven requirements describe the structure of independent control; they do not guarantee the correctness of each decision, and they do not make any control layer the institution's ultimate authority. The determination decides whether the action takes effect: only a cleared action proceeds. A determination that does not decide this is advisory and cannot be the preventive control that P4 requires.

## 7. Where the requirements hold

The requirements depend on the conditions stated in Section 5 being present. The dimensions below inform proportionality. They are not a scoring formula and are not to be added together. The institution's applicable obligations, whether legal, supervisory, contractual or to affected parties, determine what assurance is required; the dimensions indicate how strongly the conditions are engaged.

| Dimension                 | Conditions strongly engaged                                                             | Conditions weakly engaged                   |
|---------------------------|-----------------------------------------------------------------------------------------|---------------------------------------------|
| Reversibility             | Consequence cannot be reversed or fully compensated                                     | Consequence fully and promptly reversible   |
| Magnitude                 | Material financial, legal, operational or public effect                                 | Negligible effect                           |
| Velocity and aggregation  | Many actions whose effects accumulate                                                   | Few, isolated actions                       |
| Detection delay           | Errors found late or not at all                                                         | Errors found immediately and reliably       |
| Delegation depth          | Action shaped by several systems or agents                                              | Single, well-understood step                |
| Influence over the action | System routes, prepares or executes the action                                          | System provides only general information    |
| Obligations               | Institution answerable to supervisors, clients, affected parties or under contract      | No obligation concerning the outcome        |
| Manipulability and change | Actor processes untrusted content; operating function or provider can change the system | Closed, fixed system without external input |

A rare action can still be grave, and a reversible action can still be subject to a specific requirement for independent approval. Where the conditions are not engaged, lighter controls may meet the institution's actual obligations.

#### Boundary cases.

- *Composite actions*: the determination precedes the first step that is consequence.
- *Preparatory steps with external effect*: consequential actions in their own right.
- *Delegation chains*: the requirements apply to the action that becomes consequence, and the acting system includes every system in the chain.
- *Actions spanning institutions*: each institution's accountability attaches to the consequence for which it answers.
- *Gradual or aggregated consequence*: the relevant class is the one for which the control is relied upon; exclusions by value or type are recorded and justified against that class.
- *Physical effects*: where a digital command produces a physical effect, the final preventable point precedes the command that produces it.

## 8. Correspondence with the published terms

The arguments in Sections 4 to 7 rest on P1 to P11 and do not rely on OATHOR's definitions. Where a published term already stated a property, its wording has been used; this section shows the correspondence with the terms and criteria OATHOR LTD has published, with their dates.

**Execution Boundary.** The Execution Boundary is the final control location before a requested digital action becomes operational, institutional, financial, legal or public consequence. *Corresponds to the final preventable point (Section 4).*

**Execution Control Infrastructure.** Execution Control Infrastructure is the independent control layer that governs the Execution Boundary and determines whether a consequential action is cleared to proceed.

**Independent Clearance Determination.** The independent determination, made at the Execution Boundary and before consequence, of whether a consequential action is cleared to proceed. Only a cleared action proceeds. *Corresponds to the determination meeting Requirements 1 to 7.*

**Independence.** Independence is the structural separation between the system seeking to create a consequence and the control layer that determines whether that consequence may proceed. A clearance determination is independent when the actor seeking execution cannot, by virtue of being the actor, unilaterally determine, alter, satisfy, bypass or substitute the conditions under which its own consequential action is cleared to proceed. Independence also requires that the determination is not controlled by the function or provider that operates or supplies the acting system. *Corresponds to Requirements 1 to 3.*

**Human Authority Layer.** The Human Authority Layer is the requirement that accountable human authority remains effective as systems act with increasing autonomy. Execution Control Infrastructure is the independent control layer at the Execution Boundary.

The Human Authority Layer is a requirement, not a reviewer or a decision-making function. As automated actions grow in number and speed, review of each action by a person becomes impracticable. Accountable authority then remains effective only if the conditions set under that authority are applied to every action before consequence, independently of the acting system and its provider, and people remain able to set those conditions and to intervene (P8). The seven requirements describe that condition.

**Execution Proximity.** The degree to which a system may influence, initiate, route, approve, restrict, suspend, escalate, coordinate or execute a consequential action, independent of whether it makes the final decision. *Corresponds to the second observation in Section 4.*

**Execution-Boundary Risk.** The risk that consequence proceeds without appropriate control before execution.

| Requirement                                       | Published criterion |
|---------------------------------------------------|---------------------|
| 1. Outside the acting system                      | ECI-1               |
| 2. Beyond the actor's reach                       | ECI-2               |
| 3. Free of operating-function or provider control | ECI-3               |
| 4. Before consequence                             | ECI-4               |
| 5. The specific action                            | ECI-5               |
| 6. As a matter of course                          | ECI-6               |
| 7. Recorded at the time                           | ECI-7               |

The published criteria state: "A control that answers yes to all seven meets the criteria for Execution Control Infrastructure. A control that answers no to any of the first three remains inside the authority of the system it governs." The criteria are reproduced in Appendix A.

**Chronology.** OATHOR's Foundation Paper, Version 1.0, May 2026, set out the category, the Execution Boundary and the independent control layer, and stated that "a verifiable control record is generated at the execution boundary" and is "produced at execution time, not retrospectively". The Human Authority Layer,

Execution Proximity and Execution-Boundary Risk were set out on 24 June 2026, in consultation submissions to the DIFC Authority and the European Commission AI Office, as recorded on OATHOR's Category Reference (OATHOR LTD, 2026b). The full definition of Independence and the criteria ECI-1 to ECI-7 were published on 24 September 2026, and application notes in October 2026. Later developments are not attributed to earlier documents.

---

## 9. Convergent traditions and related work

### 9.1 Separated control across domains

In several domains where consequences are high, institutions have developed separated control before the irreversible step: maker-checker in banking (Basel Committee, 1998); the two-person rule in high-consequence operations; multi-signature authorisation, under which a transaction requires signatures from several independent keys (Andresen, 2011); segregation of duties and managed change in information technology (ISO/IEC 27001:2022, controls 5.3, 8.31 and 8.32); the reference monitor and complete mediation in computer security (Anderson, 1972; Saltzer and Schroeder, 1975); the independence of safety instrumented systems from basic process control (IEC 61511-1:2016); independent levels of defence in depth in nuclear safety (IAEA, SSR-2/1); the self-review threat in audit independence (IESBA); and pre-order controls on automated market access (17 CFR 240.15c3-5; Regulation (EU) 2017/589).

These traditions do not implement the same architecture and none establishes all seven requirements. Their value as evidence is that, independently and across domains, institutions facing irreversible consequence have repeatedly concluded that the party that acts, or that builds and operates the means of acting, should not be the only party that decides whether the action takes effect. The requirements derived here are consistent with that recurring conclusion and carry it to consequential digital action.

### 9.2 Current work on pre-execution control of automated actions

Frameworks for AI agent security describe risks including manipulation through content, excessive agency and misuse of tools (OWASP, 2025; MITRE ATLAS; NIST, 2023).

Three recent works address control of actions before execution.

- *Faramesh* (Fatmi, 2026) proposes a non-bypassable boundary at which an autonomous agent's intended action is authorised against policy before execution, with provenance records for audit. Its scope is autonomous agent systems. It does not state a requirement that the decision be independent of the provider or operator of the agent system.
- *A SCITT profile for pre-execution AI action authorization records* (Munoz, 2026) specifies signed authorisation records bound to a specific request and recorded in a transparency service. Its scope is AI actions. It states that the issuer of authorisations and the transparency service may be the same operator; this concerns the record, not the provider of the acting system.
- *External Admission Boundary for Pre-Execution AI Action Control* (B., 2026) requires that an allow decision be issued outside the executor's trust domain, defined to include the system, workflow, agent, runtime, platform or tool environment, and that the executor cannot approve, bypass or self-issue its admission. Its scope extends beyond AI to automated workflows and other execution-capable systems. It does not name the function or provider that operates or supplies the executor.

These works establish that external, non-bypassable, action-specific control before execution, with supporting evidence, was being proposed in early and mid-2026. This paper does not claim otherwise.

### 9.3 What the complete requirement adds

Relative to the traditions and works above, the requirement derived here combines:

1. **A location defined by consequence**, across operational, institutional, financial, legal and public settings, rather than by resource access or a particular class of system, together with explicit treatment of systems that influence an action without taking the final step.
2. **Independence from the operating function or provider**, as a requirement in its own right, grounded in accountability, the self-review threat and shared dependency, and applying to internal functions as well as external providers.
3. **Seven requirements shown to be jointly necessary**, each derived from accepted premises and each with stated conditions, valid alternatives and limits.
4. **Institutional accountability as the frame**: the requirements are what keeps an institution's responsibility effective as its actions become automated.
5. **A test applicable to any control**, whatever it is called and whoever supplies it.

Overlap with earlier work on individual elements is acknowledged. The contribution is the complete requirement and its derivation.

---

## 10. Counterarguments

**"A human in the loop is enough."** A person can take part in an independent determination. The approval meets the requirements only if it does: it must be beyond the actor's reach, including the actor's ability to shape what the person is shown, and it must apply as a matter of course. Where volume makes review of each action by a person impracticable, the requirement does not disappear; it must be met in another way, with people setting conditions and able to intervene.

**"Monitoring and rollback are enough."** For fully reversible actions they may be, as Section 7 states. For irreversible actions they identify the consequence after it exists.

**"Guardrails inside the agent are enough."** They are useful adjacent controls. They are within the acting system and share its exposure (Requirement 1).

**"The provider is trustworthy."** The argument assumes no bad faith. Segregation of duties, self-review and functional safety are structural safeguards against error, manipulation, change and compromise, not accusations.

**"An independent check can still be wrong."** Yes. The claim is necessity under stated conditions, not sufficiency.

**"Who controls the controller?"** Independence is relative to the actor and to the function or provider that operates or supplies the acting system. It does not require an infinite chain. The control's own integrity is a security matter in its own right, and accountability for the outcome remains with the institution.

**"This moves liability to the control layer."** It does not. Responsibility cannot be delegated (P1). The independent determination keeps the institution's own conditions effective.

**"An independent provider creates a new concentration."** This is a real consideration. The requirements are neutral as to who meets them and allow many independent providers. Concentration of the control layer should be assessed like any critical service.

**"What if the control is unavailable?"** The requirements apply to the control as it operates, including when it or the acting system fails, is unavailable or is being updated. An action that proceeds without the control being applied, or when no determination is made, does not meet Requirements 2 and 6.

**"It is too slow or costly."** The requirements are proportionate to the conditions in Section 7. Where those conditions hold, the cost is weighed against a consequence that cannot be undone.

**"A different product or name avoids the requirement."** Where the full requirement applies, any arrangement must satisfy its substantive conditions. Different terminology or implementation does not remove them.

---

## 11. Illustrative patterns

The patterns below illustrate how the institutional requirement fails when one requirement is absent. They are illustrations, not a survey of incidents, and no organisation or product is identified.

| Pattern                                                                                                                               | Requirement absent |
|---------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| An agent reads content containing embedded instructions and carries out an action its operator did not intend (Greshake et al., 2023) | 1 and 2            |
| A broad permission granted for one task is used for a different consequential action                                                  | 5                  |
| An update by the provider changes how actions are evaluated, without the institution's agreement                                      | 3                  |
| A person approves an action on the basis of a summary composed by the acting system                                                   | 2                  |
| Small actions excluded from review accumulate into a material loss                                                                    | 6                  |
| After an event, the record of what was approved is assembled from partial logs                                                        | 7                  |
| An action is reviewed after it has taken effect                                                                                       | 4                  |

---

## 12. What would refute the argument

The argument would be weakened or refuted by evidence of any of the following.

1. **The obligation does not require independence.** That, in a setting where the conditions of Section 5 are present, the institution's underlying obligation is fully met by an arrangement in which the actor or its operating function or provider controls the determination. The challenge must address the obligation itself, not the performance record of a particular dependent control.
2. **Complete recovery.** That for a class of actions believed to be irreversible, full and timely recovery is available, placing the class outside Requirement 4.

3. **Accountability without authority.** That legal and supervisory frameworks accept that an institution discharges its responsibility for actions whose clearance conditions another party can change unilaterally.
4. **Inapplicability to automated actors.** That segregation of duties, self-review and the independence of protective functions do not apply where the acting party is an automated system. The market access rules in Section 3 are evidence against this.

Each is a claim about evidence or about the content of accepted principles, and each can be examined independently of OATHOR.

---

## 13. Implications

**For institutions.** For each type of consequential action that an automated system can take, an institution can identify the control it relies upon to determine whether the action proceeds, and test that control against the seven requirements. Section 7 indicates where the full requirements are engaged.

**For supervisors.** Principles of accountability, segregation of duties, third-party risk, self-review and human oversight already point toward these requirements. For example, the Bermuda Monetary Authority has proposed in consultation that, for higher-impact actions, controls should not rely solely on prompts or model instructions (Bermuda Monetary Authority, 2026, proposed Guidance Note, paragraph 74). Where expectations for automated systems are developed, the distinction between controls governed by the acting system or its provider and controls independent of both may be relevant to whether existing principles are met in practice.

**For providers.** Providers can supply technology used in an independent arrangement without controlling the determination. Designing systems so that clearance of their consequential actions can be determined independently is compatible with the requirements and with the provider's commercial role.

---

## 14. Conclusion

Every consequential action has a last point at which it can be prevented. What must happen at that point is a matter of principles that institutions, supervisors and engineers already apply: responsibility cannot be delegated; no party should approve its own consequential action or assure its own work; a protective function should be independent of what it protects to the extent that a shared dependency could cause both to fail; every relevant action must be checked; a check holds only for the conditions under which it was made; oversight must be able to intervene; and records must be made at the time.

Applied to automated actors and to the functions and providers that operate and supply them, these principles require, under the conditions stated, a determination that is outside the acting system, beyond the actor's reach, free of control by the operating function or provider, made before consequence for the specific action, applied as a matter of course and recorded at the time. The seven requirements are jointly necessary, and they correspond to the published criteria ECI-1 to ECI-7.

As consequential digital execution becomes increasingly autonomous, accountable human authority cannot depend on direct human intervention at every action. Three roles must then remain distinct: the institution retains accountable authority; the Independent Clearance Determination decides whether the specific consequential action is cleared to proceed; and the acting system carries out the action. Execution Control

Infrastructure neither assumes the institution's authority nor becomes the executing system. A technically separate component is not thereby independent: what matters is who holds effective authority over the conditions of clearance. Where an institution requires independent control before consequence, that control must remain effective at the Execution Boundary, beyond unilateral control by the acting system and by the function or provider that operates or supplies it. The seven requirements state the conditions that this determination must meet, and together they keep institutional authority, independent clearance and execution separate, without transferring the institution's responsibility to the control layer.

Different implementations are possible. The required properties are not optional where the institutional conditions make them necessary.

The requirement does not depend on any name or supplier. It depends on the premises. Where they hold, it follows. Only a cleared action proceeds.

---

## Declarations

**Competing interests.** The author is the founder and director of OATHOR LTD, which developed and published the Execution Control Infrastructure architecture and qualification criteria discussed in this paper. OATHOR LTD has commercial interests relating to this field and has filed patent applications in it.

**Generative AI assistance.** During preparation of this manuscript, the author used Claude (Anthropic) to assist with drafting, language editing, manuscript review and the identification and checking of references. The architectural concepts and ECI qualification criteria discussed in this paper were previously published by OATHOR LTD. The author directed the work, reviewed and revised the manuscript, verified the supporting sources and takes full responsibility for its arguments, accuracy and conclusions.

**Version and licence.** Version 1.0, October 2026, published by OATHOR LTD. Later versions, including any version submitted for peer review, will be identified by version number. Copyright 2026 OATHOR LTD. All rights reserved. Quotation with attribution to the author and OATHOR LTD is permitted. The canonical definitions and the criteria ECI-1 to ECI-7 reproduced in this paper may be reproduced word for word under the permission published at <https://oathor.com/terms#licence>.

**Citation.** Thameem, M. (2026). Independent Clearance Before Consequence: Why accepted institutional control principles require an independent determination at the Execution Boundary. Version 1.0, October 2026. OATHOR LTD.

---

## Appendix A. Criteria for Execution Control Infrastructure

Source: OATHOR LTD, Criteria for Execution Control Infrastructure, published 24 September 2026, <https://oathor.com/execution-control-infrastructure>

- **ECI-1:** Is the determination made outside the system that takes the action?
- **ECI-2:** Is it beyond the reach of the actor, so that the actor cannot determine, alter, satisfy, bypass or substitute the conditions of its own clearance?
- **ECI-3:** Is it free of control by the function or provider that operates or supplies the acting system?
- **ECI-4:** Is it made before consequence exists, while the action is still preventable?

- **ECI-5:** Does it decide whether this specific action is cleared to become consequence, or only whether the actor is permitted?
- **ECI-6:** Is it applied at the Execution Boundary as a matter of course, not only as an occasional review or exception?
- **ECI-7:** Is the record created at the moment of determination, rather than reconstructed afterwards?

A control that answers yes to all seven meets the criteria for Execution Control Infrastructure. A control that answers no to any of the first three remains inside the authority of the system it governs.

Application notes: <https://oathor.com/execution-control#application-notes>

---

## References

Anderson, J. P. (1972). Computer Security Technology Planning Study, Vol. I. ESD-TR-73-51. Electronic Systems Division, United States Air Force. Section 3.2.2.

Andresen, G. (2011). BIP 11: M-of-N Standard Transactions. Bitcoin Improvement Proposals.

B., F. (2026). External Admission Boundary for Pre-Execution AI Action Control. IETF Internet-Draft draft-b-external-admission-boundary-00, 16 May 2026. Individual submission.

Basel Committee on Banking Supervision (1998). Framework for Internal Control Systems in Banking Organisations. Bank for International Settlements, September 1998.

Bermuda Monetary Authority (2026). Consultation Paper: The Responsible Use of Artificial Intelligence in Financial Services in Bermuda, 14 August 2026.

COSO (2013). Internal Control: Integrated Framework. Committee of Sponsoring Organizations of the Treadway Commission.

CPMI-IOSCO (2026). FMIs' reliance on third-party service providers: challenges and risks. Discussion paper, September 2026.

European Banking Authority (2019). Guidelines on outsourcing arrangements. EBA/GL/2019/02, 25 February 2019.

European Union (2014). Regulation (EU) No 537/2014 on specific requirements regarding statutory audit of public-interest entities. Article 5.

European Union (2017). Commission Delegated Regulation (EU) 2017/589 (regulatory technical standards on algorithmic trading). Articles 12 and 15.

European Union (2022). Regulation (EU) 2022/2554 on digital operational resilience for the financial sector. Article 28.

European Union (2024). Regulation (EU) 2024/1689 (Artificial Intelligence Act). Articles 12, 14 and 26.

Fatmi, A. (2026). Famesh: A Protocol-Agnostic Execution Control Plane for Autonomous Agent Systems. arXiv:2601.17744, 25 January 2026.

FDA (2018). Data Integrity and Compliance With Drug CGMP: Questions and Answers. Guidance for Industry, December 2018.

Greshake, K., Abdelnabi, S., Mishra, S., Endres, C., Holz, T. and Fritz, M. (2023). Not what you've signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. arXiv:2302.12173.

IAEA (2016). Safety of Nuclear Power Plants: Design. Specific Safety Requirements SSR-2/1 (Rev. 1). Requirement 7.

IEC 61511-1:2016. Functional safety: Safety instrumented systems for the process industry sector, Part 1. Clauses 11.2.4 and 11.2.9, as quoted in UK Health and Safety Executive, Safety Bulletin CEMHD02-2023.

IESBA. International Code of Ethics for Professional Accountants (including International Independence Standards). Paragraph 120.6 A3(b).

ISO/IEC 27001:2022. Information security, cybersecurity and privacy protection: Information security management systems: Requirements. Annex A, controls 5.3, 8.31 and 8.32.

MHRA (2018). 'GXP' Data Integrity Guidance and Definitions, Revision 1, March 2018. Sections 6.1 and 6.12.

MITRE. CWE-367: Time-of-check Time-of-use (TOCTOU) Race Condition. Common Weakness Enumeration.

MITRE. ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems.

Munoz, C. (2026). A SCITT Profile for Pre-Execution AI Action Authorization Records. IETF Internet-Draft draft-munoz-scitt-permit-profile-00, 14 May 2026. Individual submission.

NIST (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. <https://doi.org/10.6028/NIST.AI.100-1>

OATHOR LTD (2026a). Execution Control Infrastructure for High-Consequence Digital Systems. Foundation Paper, Version 1.0, May 2026. <https://oathor.com/whitepaper>

OATHOR LTD (2026b). Category Reference: Human Authority and Execution Control Infrastructure. <https://oathor.com/category-reference>

OATHOR LTD (2026c). Criteria for Execution Control Infrastructure (ECI-1 to ECI-7), 24 September 2026. <https://oathor.com/execution-control-infrastructure>

OATHOR LTD (2026d). Application notes for ECI-1 to ECI-7. <https://oathor.com/execution-control#application-notes>

OWASP GenAI Security Project (2025). OWASP Top 10 for Agentic Applications for 2026. Released 9 December 2025.

Saltzer, J. H. and Schroeder, M. D. (1975). The protection of information in computer systems. Proceedings of the IEEE, 63(9), 1278-1308.

United States (2010). Risk Management Controls for Brokers or Dealers with Market Access. 17 CFR 240.15c3-5.

United States. Regulation S-X, Qualifications of accountants. 17 CFR 210.2-01(c)(4)(ii).