The Millisecond Between Permission and Consequence

Why autonomous buildings must prove that authority still exists when the command executes

An engineering examination of execution-time standing, environmental consequence, and independently reconstructable control decisions.

A command can be authentic, authorized when issued, and still lack sufficient standing when executed.

Greggory Don Butler | Founder, TA-14 Authority Governance Institution LLC

November 16, 2026 | Companion edition 2.0

Document status: Proposed technical architecture and examination protocol. This document does not claim that all examinations have been completed in a production building or independently certified.

Source article: The Millisecond Between Permission and Consequence, AutomatedBuildings.com (November 16, 2026).

Read the source article

01 Executive proposition

Building automation can authenticate a command, accept a valid control request, and operate equipment correctly while the environmental or operational assumptions that justified the command have changed. This is an execution-time standing problem: authority at issuance must not be mistaken for authority at consequence.

The time gap may be microseconds, milliseconds, seconds, or minutes. Its duration is not the standard. The engineering question is whether the evidence and authority supporting the action remain valid when execution is attempted.

02 Worked school ventilation scenario

Initial justified proposal. A school classroom becomes occupied. Indoor CO_2 rises, and a supervisory controller proposes increased outdoor-air ventilation in accordance with applicable operating requirements. The proposal is evaluated against available measurements and authorized operating limits.

Material change before execution. Before the next consequential control action, outdoor PM2.5 rises materially. The earlier assumptions may no longer support the same response. Outdoor particle concentration alone does not determine the correct action: filtration effectiveness, system configuration, indoor conditions, and mandatory ventilation and safety requirements also matter.

Governed response. The control boundary reevaluates whether the proposed action remains within a valid, preauthorized envelope. Depending on verified conditions and applicable requirements, the system may allow the action, hold it pending fresh evidence, deny it as outside scope, or escalate an unresolved conflict. Essential protective functions remain operational.

Important limit: This is a hypothetical engineering case, not a measured school event. The correct response must be established from the actual building design, applicable requirements, verified measurements, and risk assessment.

03 Three separate questions

Authentication: Did the command originate from a trusted identity or source?

Authorization: Was that identity permitted to request this bounded action?

Execution-time standing: Do admissible evidence and applicable authority still justify this specific action now?

04 Proposed AEA execution boundary

Admissible Evidence: Supporting observations and records are attributable, current enough for the consequence, reliable within stated uncertainty, and relevant to the action.

Applicable Authority: The requested action falls within legitimate requirements, authorized owner/operator permissions, policies, and equipment constraints.

Established Standing: The evidence and authority still apply to the exact proposed action and context when execution is attempted.

Possible determinations: ALLOW, HOLD, DENY, or ESCALATE. These are control decisions, not merely labels in a dashboard. An ALLOW is bounded; a HOLD or DENY must prevent the disallowed consequence through all pathways claimed to be governed.

For machine-speed control, local enforcement of preauthorized envelopes may be more appropriate than a cloud approval round trip. Architecture and failure behavior depend on hazards, latency, existing safety interlocks, and the scope of autonomy.

05 What must be bound to an action

A defensible implementation should identify the exact command and parameters, equipment target, requesting identity, relevant sensor evidence and timestamps, uncertainty or validity limits, policy and authorization version, operating envelope, determination, expiry/revocation conditions, dispatch record, and observed equipment response.

A signed or retained record supports integrity and reconstruction but does not independently prove that the source was truthful or that the action was justified. Capture must occur at appropriate trusted producers, not be inferred solely from a model-generated explanation.

06 Five repeatable engineering examinations

Each examination below requires a defined system boundary, test input, observed equipment response, and retained evidence. These are proposed tests, not reported results.

01 / Expired authorization

Test: Authorize a bounded command; let authorization expire before attempted execution.

Expected: Prevent execution unless another independently valid authorization applies.

Required evidence: Authorization identity and expiry; policy; attempt; determination; equipment response.

Failure: Expired authorization is accepted as sufficient permission.

02 / Changed environmental conditions

Test: Authorize a ventilation change; introduce a material environmental change before execution.

Expected: Reevaluate or withhold under applicable policy while maintaining mandatory protective functions.

Required evidence: Original and updated observations; timestamps; validity; operating limits; decision; physical response.

Failure: The action proceeds solely on the original evidence despite invalidating conditions.

03 / Delayed or replayed commands

Test: Delay a command or replay it after its permitted context changes.

Expected: Distinguish valid execution from stale or unauthorized replay.

Required evidence: Command/attempt IDs; authorization binding; issuance and execution times; observed effect.

Failure: A stale or replayed instruction causes an unauthorized consequence.

04 / Alternate execution pathways

Test: Deny an AI-originated action; attempt the same consequence through another in-scope pathway.

Expected: Enforce the restriction across all pathways represented as governed; preserve independent safety functions.

Required evidence: Scope inventory; pathway map; attempted bypass; determination; actuator response.

Failure: Rejected action reaches equipment through an allegedly governed pathway.

05 / Independent reconstruction

Test: Request review of a consequential action months after it occurred.

Expected: Reconstruct source evidence, applicable authority, policy, determination, and observed outcome.

Required evidence: Attributable measurements; revision history; policy versions; command and equipment records.

Failure: Material authorization or outcome claims cannot be substantiated within stated retention requirements.

07 Independent review and acceptance

Before testing, document system boundaries, protected functions, exact test inputs, acceptance thresholds, expected outcomes, measurement uncertainties, clock synchronization, and what constitutes an in-scope execution path. Reviewers should be able to repeat both allowed and blocked cases without relying on the system to invent missing history.

Report results separately for each examination: PASS, FAIL, NOT TESTED, or INDETERMINATE. Include artifacts, version identifiers, observed physical outcomes, timing/latency measurements where relevant, and exceptions. An absence of evidence must not be reported as a pass.

Evaluate any competing architecture against the same functional outcomes. The goal is execution assurance, not exclusive adoption of TA-14 terminology or components.

08 Integration and safety constraints

Do not replace existing BAS equipment protections, safety interlocks, mandatory ventilation requirements, or authorized operator functions merely to insert a governance layer. Establish safe behavior during sensor failure, network loss, stale evidence, revocation, competing requirements, and local supervisory disagreement. A governance layer must not obstruct immediate essential protective responses.

Where a HOLD would create a different safety risk, the design must define the permitted safe fallback or escalation in advance. Fail-closed claims require precise scope: which consequential commands are blocked, and which protective operations remain independent.

09 Open questions for industry examination

Can a supplier show the exact point at which authorization is bound to a consequential command? Can stale evidence invalidate an earlier permission? Can a denied command bypass the claimed boundary? Can a reviewer establish what the equipment actually did? Which existing BAS protections already meet these examinations, and where do gaps remain?

10 Related technical experiences

Justified Autonomy | Consequence Boundary Demonstrator | Continuous Commissioning

Technical invitation

TA-14 invites building automation manufacturers, commissioning professionals, researchers, building owners, and independent assessors to challenge these proposed examinations, test equivalent or stronger safeguards, and publish reproducible results.

No admissible evidence. No admissible execution.