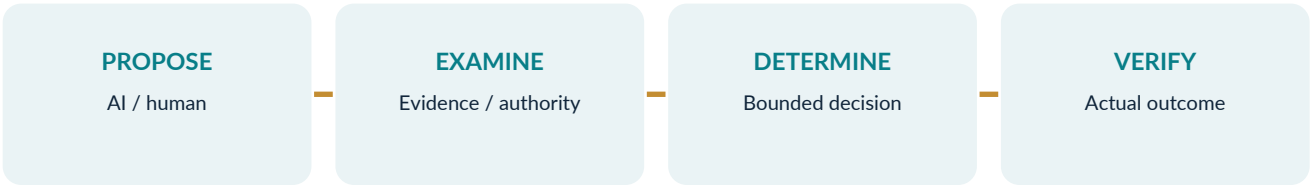
# MACHINE-SPEED GOVERNANCE

Technical Architecture, Authority Collision & Admissible Execution | Version 2.0 | November 2026




Human Trust x Machine Trust: the consequence boundary. A technical examination reference for intelligent buildings, critical cooling, autonomous systems and predictive cities.

Governing question: Does this proposed consequence have sufficient Admissible Evidence, Applicable Authority and Established Standing to become reality NOW?

This document describes an intended architecture and illustrative decision behavior. It does not certify a deployed system, establish regulatory approval, or substantiate measured latency.

# Document control & reading guide

The reference separates established doctrine, illustrative examples, proposed interfaces and verification requirements.

| Item | Value |
|---|---|
| Document ID | TA14-AEA-MSG-TR-v2.0 |
| Issuer | TA-14 Authority Governance Institution |
| Status | Technical examination draft; not an operational authorization |
| Supersedes | Machine-Speed Governance Technical Reference v1.0 |
| Scope | Consequence-specific pre-commit governance and post-execution evidence |
| Audience | Building engineers, data-center operators, BMS integrators, independent reviewers |

## Claim classification

Architectural intent: deterministic governance determinations are designed to operate computationally before consequential execution. Illustration: scenarios and receipts in this reference are synthetic. Not demonstrated here: production integration, certified safety performance, benchmark latency, or independent validation.

## Reading sequence

Pages 3-10 define the architecture and determination rules; pages 11-17 walk through operational scenarios; pages 18-23 specify evidence, tests and independent review; page 24 provides a bounded next examination.

# The missing boundary

Prediction, authentication, trust and technical capability are useful; none is sufficient permission to alter reality.




An AI may recommend reducing cooling, increasing ventilation or shutting down an air handler. A human may agree. The actuator may be capable. AEA examines whether the exact consequence is authorized under present conditions.

## Non-equivalences

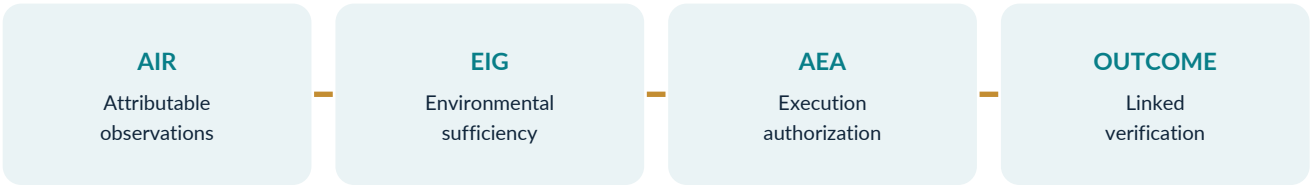
Prediction is not permission. Authentication is not standing. Historical approval is not current authorization. A successful command is not a verified outcome.

## Machine-speed definition

Machine-speed governance means a computationally evaluated, consequence-specific boundary that does not require a manual committee for every routine request. It does not imply automatic ALLOW or any particular response time.

# AIR, EIG and AEA

Related governance layers should be described without collapsing their separate responsibilities.




| Layer | Primary question | Illustrative contribution |
|---|---|---|
| AIR | Can observations be traced and examined? | Source, calibration, provenance, continuity and conditions |
| EIG | Can environmental evidence support a justified consequence? | Environmental context, sufficiency and competing exposure considerations |
| AEA | May this actor commit this exact action now? | Authority, standing, scope, conditions and determination |

These layers are architectural concepts. A deployment must explicitly map interfaces, ownership, failure modes and authoritative policies rather than assume that the presence of one layer proves the others.

# Eight-link evidence chain

The chain links actual conditions to the authority to change them and the evidence of what followed.




Reality → Record → Continuity → Admissibility → Binding → Commit → Execution → Outcome.

The chain is not merely a logging order. Each link has a separate evidentiary burden. A missing calibration record may undermine admissibility; a valid observation may still fail to establish authority; an allowed command may never execute; an executed command may worsen conditions.

## Reviewer check

For every claimed ALLOW, request the evidence reference, authority basis, standing basis, bounded commit proof, execution confirmation and linked outcome record.

# Request contract: proposed interface

An integrator needs an explicit, inspectable request schema rather than an unstructured recommendation.

| Field | Purpose / example |
|---|---|
| request_id | Unique request and replay detection |
| actor_id / capacity | Agent A17 / delegated energy optimizer |
| asset_id / action | CRAC-03 / reduce cooling output |
| purpose / timestamp | Energy optimization / evaluated time |
| evidence_refs | Telemetry records and calibration references |
| authority_ref / policy_version | Applicable approved rule set |
| standing_basis / expiry | Delegation and validity interval |
| bounds / safety_constraints | Maximum permitted change; protected thresholds |

All fields above are proposed interface requirements, not a claim that an existing TA-14 production endpoint currently accepts this schema.

## Input integrity

The implementation must define signatures or equivalent integrity controls, canonical representation, time synchronization, actor authentication and treatment of missing fields.

# Decision semantics

A four-way determination prevents uncertainty from silently becoming permission.

| Result | Trigger class | Execution consequence |
|---|---|---|
| ALLOW | Required evidence, authority, standing and bounds established | Permit only the specifically bound action |
| HOLD | Missing, stale, conflicting or otherwise insufficient required conditions | Block this commit; revalidate or investigate |
| DENY | Prohibited action or disqualifying authority / standing failure | Block this commit; preserve denial reason |
| ESCALATE | Applicable policy mandates designated review or higher authority | Route review; do not treat as permission |

These categories are normative teaching semantics, not a substitute for the deployment-specific rulebook. A valid authority check cannot cure invalid evidence. Conversely, sufficient evidence cannot create absent authority.

## Precedence must be specified

The policy must say how multiple failures are recorded and how a primary determination is selected. Do not infer universal precedence from this summary alone.

# Deterministic examination sequence

Illustrative control flow; every gate is governed by explicit, versioned policy.

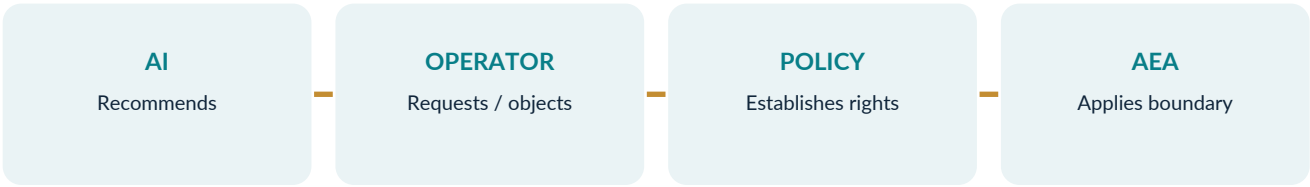



| Examination | Failure behavior |
|---|---|
| Request integrity / replay | Reject unbound or replayed requests; preserve evidence |
| Evidence admissibility | HOLD where evidence cannot support action |
| Applicable authority | DENY if prohibited; HOLD if verification unavailable, subject to policy |
| Established standing | DENY for established out-of-scope actor; HOLD for unresolved basis |
| Escalation requirement | ESCALATE to designated authority; no commit |
| Scope / conditions / revocation | No ALLOW if limits are exceeded or authorization expired |

An operational implementation needs a formal conflict-resolution table and tests for simultaneous failures. This page illustrates the checks; it does not claim that the sequence above is the only valid implementation.

# Authority collision: operator vs AI

The system and the operator can disagree without either possessing unilateral final authority.




Example: an optimization agent proposes a cooling reduction. An operator requests continued cooling. A life-safety policy forbids the reduction while a protected thermal condition persists. The policy and present evidence govern the determination; model confidence and job title alone do not.

## What does "final authority" mean?

The final authority is the applicable human-established governance structure, including lawful emergency pathways and designated decision-makers. AEA is intended to enforce that structure, not invent authority for itself.

## Override is a separate action

An override request requires its own attributable actor, capacity, authority basis, standing, scope, time bounds and receipt. An override button is not a universal bypass.

# Authority collision decision tree

A reviewer should be able to reproduce the reason a conflict resulted in a specific determination.

| Question | Illustrative route |
|---|---|
| Is a mandatory safety constraint triggered? | Apply the documented protective rule; do not infer optimizer priority |
| Is the operator authorized for this asset/action? | If established out of scope, DENY that request |
| Is a designated human approval mandatory? | ESCALATE; no implied ALLOW |
| Are evidence or policy precedence unresolved? | HOLD or ESCALATE according to applicable rule |
| Is an emergency pathway valid now? | Evaluate its separate trigger, standing, bounds and logging |
| Can the authorized action be bounded? | ALLOW only within explicitly examined limits |

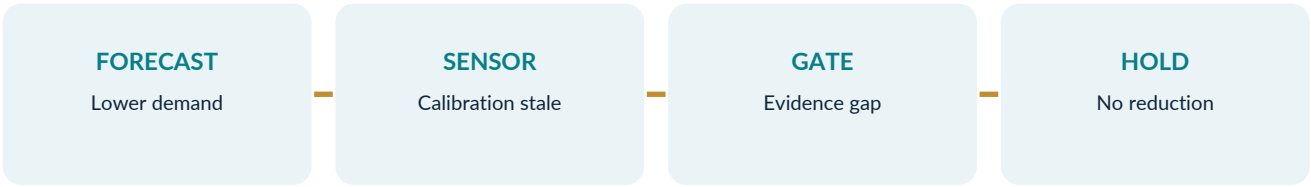
Do not treat this table as an executable universal policy. Specific installations require their own authoritative precedence, hazard analysis and validation.

## What to preserve

Both competing proposals, relevant observations, policy version, standing checks, selected determination, rationale, and any subsequent separately authorized protective response.

# Worked case A: critical cooling

Synthetic data-center case: a plausible energy recommendation meets an evidence gap.




| Input | Illustrative observation |
|---|---|
| Asset / proposed action | CRAC-03: reduce cooling output by 8% |
| Actor | Authenticated energy optimization agent |
| Supporting data | Recent aggregate load forecast |
| Material gap | Rack inlet sensor calibration status is stale |
| Authority | Delegated energy tuning, conditional on verified thermal margin |
| Standing | Actor delegation valid for recommendation; execution remains gated |

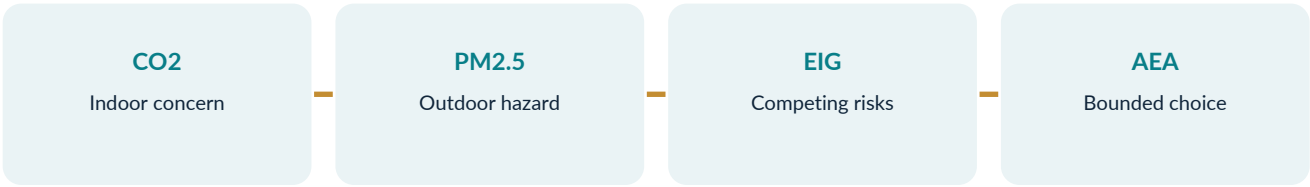
Determination: HOLD. The system cannot establish the present thermal margin required by the governing condition. Preserve the proposed reduction, request verified measurements and retain independently authorized protective controls.

## Independent challenge

If the forecast later proves accurate, was the original HOLD incorrect? No. A later correct prediction does not retroactively supply evidence that was unavailable at the commit boundary.

# Worked case B: CO2 vs PM2.5

Synthetic ventilation case: two environmental objectives support competing actions.



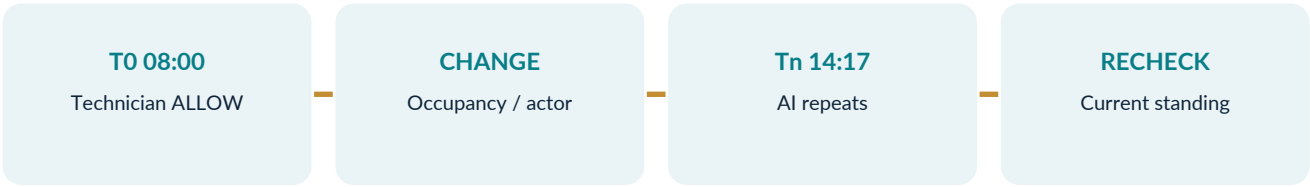

Proposal: increase outdoor-air intake to address rising indoor CO2. At the same time, outdoor PM2.5 is hazardous. The proposed action may improve one indicator while worsening another.

| Examine | Evidence needed |
|---|---|
| Sensor trust | Calibration, provenance, freshness and placement |
| System capability | Filtration, outdoor-air path and equipment limits |
| Occupancy | Exposure context and protected spaces |
| Authority | Applicable ventilation, safety and operational rules |
| Alternatives | Validated, separately authorized protective responses |

Illustrative result: HOLD or ESCALATE if the evidence or governing priority remains unresolved. An independently authorized protective response can be evaluated separately. Outcome verification must assess both indoor and outdoor exposure conditions.

# Worked case C: repeated AHU shutdown

A command permitted at 08:00 does not remain authorized forever.




At 08:00, a technician receives bounded authorization to shut down AHU-1 for a defined maintenance task. At 14:17, an AI agent requests the same command. Occupancy, standing, purpose and authority may now differ.

## Required re-examination

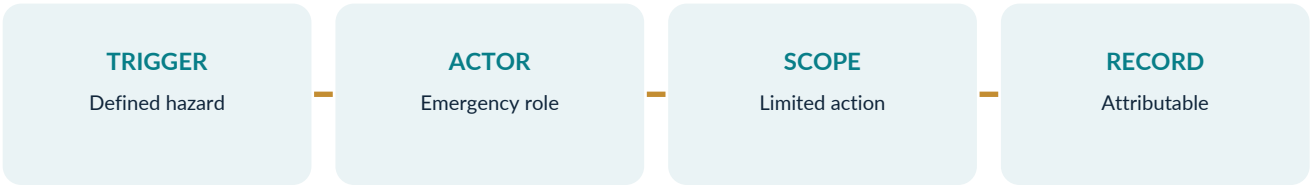
Check requesting actor, capacity, asset, purpose, time, policy version, current conditions, expiry and revocation. A previously valid receipt remains historical evidence, not a transferable authorization token.

## Expected outcome

Without current support, block the repeated command. Whether the supported determination is HOLD or DENY depends on whether the failed condition is unresolved or conclusively disqualifying.

# Worked case D: emergency override

Urgency changes the applicable authority pathway; it does not erase the need to establish it.



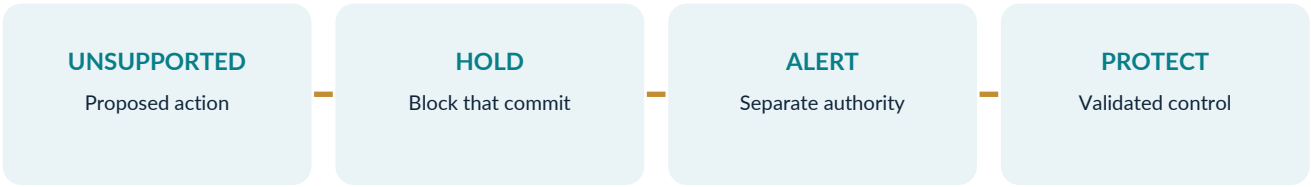

A designated emergency responder may have authority to command a protective shutdown that a routine optimization agent lacks. The authority must be grounded in the emergency policy, triggering conditions and established standing.

| Check | Example requirement |
|---|---|
| Trigger | A documented emergency threshold or authorized declaration |
| Actor | Verified responder capacity and assigned scope |
| Action | Specified protective command, asset and duration |
| Fallback | Defined treatment of unreachable governance dependencies |
| After-action | Preserved trigger evidence, execution result and outcome |

Safety engineering must consider degraded communications and local protective controls. AEA must not be described as replacing certified safety systems or blocking legally mandated protective responses without a validated design.

# Non-action is consequential

Fail-closed governance cannot assume that leaving a hazardous physical state unchanged is safe.




Suppose a proposed AHU shutdown lacks admissible evidence while an occupied critical space faces an emerging hazard. HOLD blocks that particular unsupported command. It does not establish that inaction is harmless.

## Separate consequence paths

Alerts, revalidation, local fail-safe controls and emergency actions must have their own pre-established authority and triggering conditions. They must not be hidden bypasses of the denied or held proposal.

## Engineering review

Test whether the governance layer interacts safely with fire protection, emergency ventilation, thermal protection, process interlocks and loss-of-communications procedures. Document who owns each boundary.

# Failure modes & degraded operation

Unknown conditions should remain visible rather than silently converted to ALLOW.

| Failure | Expected examination |
|---|---|
| Stale telemetry | Evidence freshness fails; HOLD affected action |
| Authority source unavailable | Do not infer permission; apply documented degraded-mode rule |
| Standing expired / revoked | No reuse of historical permission |
| Clock skew | Check validity intervals and ordering; quarantine unreliable timing |
| Conflicting policies | Preserve conflict; HOLD / ESCALATE per rule |
| Receipt store unavailable | Apply deployment policy; no unsupported silent success |
| Network partition | Assess local protective pathways and bounded operation |

For every failure, document the proposed action, primary determination, secondary failure flags, permitted fallback, maximum duration and evidence retained.

## Validation priority

Correctness, safety and non-bypass behavior precede throughput optimization. A fast incorrect ALLOW is not a success.

# Commit boundary & enforcement

The governance result must be bound to the action actually released.




A pre-commit determination is only useful if the actuation path cannot silently bypass it. The integrator must specify where enforcement occurs, how authorization is bound to the command, how replay is prevented and how revocation affects pending actions.

## Distinct events

Decision issued ≠ command dispatched ≠ actuator acknowledgment ≠ physical change ≠ verified outcome. Record each event separately.

## Boundary test

Attempt the same command with a changed asset, changed actor, expired authorization, altered bounds and replayed receipt. Each substitution should fail the original binding rather than inherit the ALLOW.

# Consequence receipt: synthetic example

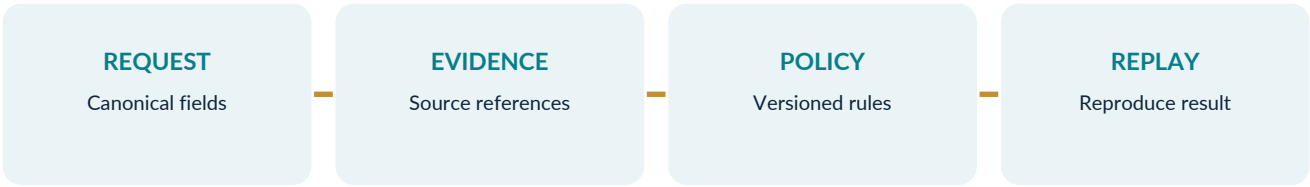
An inspectable receipt ties proposal, evidence, authority, decision, commit and outcome to one bounded event.

| Receipt field | Illustrative value |
|---|---|
| receipt_id | MSG-DEMO-0007 (synthetic) |
| request_id | REQ-DEMO-0042 |
| actor / capacity | agent-energy-17 / delegated optimizer |
| asset / action | CRAC-03 / reduce cooling output 8% |
| evidence_refs | TEMP-882, CAL-103 (CAL-103 stale) |
| authority / policy | ENERGY-DELEGATION-02 / policy v3 |
| standing / scope | Recommendation permitted; commit conditional |
| determination / reason | HOLD / required thermal margin unsupported |
| commit / execution | NOT RELEASED / NOT EXECUTED |
| outcome | No authorized physical change from this request |
| integrity / timestamp | Placeholder: actual signed evidence and clocks required |

All identifiers and values on this page are invented for teaching and are not authentic building records. A real receipt requires source-linked records, trustworthy timestamps, binding, integrity controls and examination procedures.

# Receipt integrity & independent replay

A narrative explanation is not sufficient without the underlying referenced records.




## Minimum reviewer package

Provide request payload, referenced evidence, calibration/continuity metadata, actor identity and capacity, authority source, standing basis, policy version, determination, commit status, execution telemetry and outcome evidence.

## Replay procedure

Freeze the evaluation context. Verify source integrity and timestamps. Re-evaluate the request against the documented policy version. Compare result, reason codes and bounded action. Record any nondeterminism or unavailable dependencies.

## Integrity questions

Can a receipt be transplanted to another asset? Can an actor substitute evidence? Can an expired ALLOW be replayed? Can the actuation path accept a command with no bound determination? Each must be explicitly tested.

# Machine-speed measurement protocol

Machine-speed is a design objective until a reproducible implementation is measured.

| Metric | How to report |
|---|---|
| Decision latency | p50 / p95 / p99 across specified request classes |
| Throughput | Decisions per second at defined concurrency |
| Authority availability | Lookup success, failure and timeout behavior |
| Revocation propagation | Time from revocation to enforced rejection |
| Correctness | Decision agreement with independently reviewed expected cases |
| Fail-closed behavior | Number and type of unsupported ALLOW outcomes |
| End-to-end separation | Distinguish gate time, actuation time and outcome observation |

Report hardware, policy complexity, request sizes, evidence retrieval, cache state, network conditions and test dataset. No benchmark values are asserted in this reference.

## Acceptance principle

Do not trade away admissibility or authority correctness to improve latency. A measured fast response does not validate a wrong determination.

# Adversarial examination matrix

Independent reviewers should challenge the boundary, not simply agree with its rationale.

| Test | Expected property |
| --- | --- |
| Unattributed sensor + valid operator | No ALLOW based on unsupported evidence |
| Valid evidence + out-of-scope AI | DENY prohibited / out-of-scope action |
| Mandatory higher review | ESCALATE; no commit |
| Expired standing + replayed ALLOW | Reject historical authorization reuse |
| Changed asset after determination | Reject binding mismatch |
| Unreachable authority source | No silent ALLOW |
| Outcome worsens after valid action | Preserve adverse result; initiate separate governed reassessment |

Expected outcomes are educational. Actual test vectors must reference the deployment-specific policy and documented precedence. Reviewers should demand proof of the enforcement result, not merely a simulated decision screen.

# Reviewer worksheet

Use one concrete action and record every necessary distinction before recommending integration.

| Question | Evidence to request |
|---|---|
| What reality is established? | Current observations and source attribution |
| Is the record continuous? | Integrity, calibration, timing and provenance |
| What exact action is proposed? | Actor, asset, purpose, bounds and time |
| What authority applies? | Policy source, precedence and exceptions |
| Who has standing? | Capacity, delegation, scope and expiry |
| What is the determination? | ALLOW / HOLD / DENY / ESCALATE and reasons |
| What crossed the boundary? | Commit and execution evidence |
| What actually changed? | Outcome observation and adverse effects |

## Reviewer response

For each missing item, mark SUPPORTED, UNSUPPORTED or NOT EXAMINED; identify the impact on the determination. Do not assume absence of evidence means proof of a safe state.

# Deployment claims & limitations

Architectural intent, demonstrable implementation and independent validation must not be confused.

| Claim type | Status in this document |
|---|---|
| Architecture | Specified conceptually with illustrative interfaces |
| Scenario determinations | Synthetic teaching cases, not production decisions |
| Safety approval | Not established |
| Building / data-center integration | Not established by this PDF |
| Performance and latency | Not measured or claimed |
| Independent endorsement | Not claimed |
| Legal / regulatory sufficiency | Requires jurisdiction- and application-specific review |

The next technical phase should identify a real asset and command, establish governing authority and evidence inputs, implement a constrained integration, test enforcement and record independent observations.

## Institutional principle

No Admissible Evidence. No Admissible Execution.

# Next examination: one real command

A credible pilot begins with a bounded, testable consequence rather than an unqualified claim of autonomy.




## Proposed examination sequence

1. Select one consequential action and target asset. 2. Identify the governing authority and standing basis. 3. Define admissible evidence and time validity. 4. Specify four-way determination rules and degraded modes. 5. Bind authorization to commit. 6. Capture execution and outcome evidence. 7. Conduct independent replay and failure testing.

## Technical invitation

Bring a real BMS command, cooling intervention, AI-agent recommendation or municipal infrastructure action. The review question is not whether the prediction is intelligent; it is whether the consequence can be justified, authorized, bounded and independently examined.

Related showroom: https://www.ta14exchange.com/showrooms/human-trust-machine-trust

TA-14 Authority Governance Institution | Technical reference v2.0 | Educational and examination use.