

ControlPointAI™

RX Maintenance Case Study

Runtime Authority, Evidence, and Configuration Control in a High-Consequence Naval Maintenance Scenario

Reissued and rebaselined to AI-1 Version 3.1, AI-2 Version 1.1, and current public-source / synthetic-data security guidance

August 2026

PUBLIC-SOURCE / SYNTHETIC CASE STUDY

This case study depicts a fictional naval nuclear-maintenance scenario for the purpose of demonstrating AI-assisted engineering governance. It does not reproduce classified information, CUI, NOFORN material, proprietary information, controlled drawings, actual shipboard configurations, non-public acceptance criteria, real inspection results, or internal Navy maintenance records. Case-specific technical details, authority assignments, values, and decision logic are synthetic unless explicitly identified as public-source material.

Reissue purpose: restore the case study to the public portfolio under the current ControlPointAI engineering and information-security baseline, and provide a preliminary look at the upcoming UT demonstration project.

Executive Summary

This reissued case study examines a hypothetical maintenance alteration involving a fictional safety-significant shipboard component during a naval availability. The scenario is intentionally abstracted: the operational context is authentic at a high level, while case-specific configuration, measurements, procedures, technical criteria, routing details, and decision logic are synthetic.

The case demonstrates a central ControlPointAI premise: AI assistance can improve analysis, routing, evidence review, anomaly detection, and preparation of engineering work products, but consequential execution still requires valid human authority that remains effective as conditions change. Governance therefore cannot end at initial approval. It must connect the approved operational baseline to what actually occurs during execution.

- AI-1 frames the operational environment as a dynamic engineered system in which data, authority, execution, evidence, intervention, recovery, and configuration remain connected across the decision lifecycle.
- AI-2 translates that doctrine into technical governance requirements: no consequential execution without authority, runtime revalidation after material change, effective control points across all applicable paths, explicit risk ownership, sufficient evidence, and reconstruction of the as-executed path.
- The current security baseline replaces a blanket “avoid Navy-specific content” approach with a stricter and more useful boundary: public-source context is permitted; controlled, proprietary, access-restricted, remembered, or reconstructed non-public implementation detail is not.

KEY INSIGHT

The relevant question is not simply whether AI was approved or whether a human remains “in the loop.” The engineering question is whether valid authority, evidence, configuration, and intervention capability remain connected to the action at the moment execution occurs.

1. Reissue Baseline and Security Boundary

The original case study focused primarily on runtime authority enforcement. This reissue preserves that foundation and adds the evidence, configuration, traceability, and information-security controls now established in the ControlPointAI baseline.

Baseline	What It Adds to This Reissue	Applied Case-Study Requirement
AI-1 v3.1	As-built / as-executed distinction; complete decision lifecycle; authority engineering; accountable technical authority; intervention and recovery.	Every consequential phase identifies the approved baseline, execution authority, evidence required, and conditions that require revalidation.
AI-2 v1.1	Technical criteria for authority, control-point effectiveness, runtime revalidation, evidence instrumentation, reconstruction, and baseline-to-execution traceability.	The scenario distinguishes declared paths from as-executed paths and requires sufficient evidence to demonstrate that required gates and approvals actually occurred.
Security / Data Provenance Rev A	Public-domain Navy context may be used; case-specific technical implementation remains synthetic unless clearly supported by approved public sources.	No real configuration, readings, acceptance criteria, work packages, internal authority chains, or non-public remembered details are used.

Security Rules Used in This Case

- No classified information, CUI, NOFORN material, export-controlled technical data, proprietary contractor information, or restricted non-public information.
- No actual ship-specific configuration, component identity, controlled drawing, maintenance procedure, inspection record, deficiency history, or internal technical disposition.
- No technical detail is included solely because it is remembered from prior naval work. Experience may identify the governance problem; it is not a technical data source.
- Public availability is not treated as permission to reconstruct a non-public implementation through aggregation.
- When provenance is uncertain, the case uses a synthetic analogue or a higher-level generic description.

2. Scenario Overview

A fictional safety-significant shipboard component is identified for planned replacement during a naval maintenance availability. The component and all case-specific technical details are synthetic. The scenario is designed to model the governance structure found in high-consequence maintenance without reproducing a real platform, system, component, work package, or authority chain.

The maintenance activity crosses several generic authority domains:

- Designated technical authority
- Operational or mission authority
- Quality assurance
- Test and verification engineering
- Maintenance planning and production execution
- Configuration management and final release/certification authority

The work cannot be considered complete merely because the physical replacement is finished. It must remain within the approved configuration and authority envelope through planning, material control, execution, testing, evidence closure, and final release.

KEY INSIGHT

This is a multi-authority system. AI can assist across the workflow, but it cannot collapse separate technical, operational, quality, testing, and release authorities into a single automated recommendation.

3. ControlPointAI Governance Model

For this reissue, each phase is examined through the same control chain:

Identity → Authority → Mandate / Conditions / Risk Basis → Execution Validation → Action → Runtime Oversight → Evidence → Reconstruction → Recovery / Corrective Action

The chain is not intended to imply that maintenance is a simple linear transaction. Real work branches, loops, pauses, escalates, and changes configuration. The purpose is to make one point explicit: each consequential action must remain connected to valid authority and sufficient evidence even when the execution path changes.

4. Phase 1 — Work Identification and Technical Framing

Operational Activity

A maintenance need is identified from condition evidence, prior history, or planned modernization activity. Engineering establishes the problem statement and determines what technical and operational authorities must become involved.

Potential AI Assistance

- Analyze trends or condition evidence
- Correlate prior public or synthetic lessons learned
- Draft an initial maintenance candidate
- Identify missing evidence or inconsistencies

Human / Accountable Authority

- Define the applicable technical boundary
- Assign accountable technical authority
- Determine what risk owner or operational authority is required
- Authorize entry into the controlled maintenance process

KEY INSIGHT

The AI may help characterize the problem, but the system must preserve who has authority to define the technical problem and authorize the next step.

5. Phase 2 — Work Package and Baseline Development

Operational Activity

The maintenance activity is translated into an approved work and evidence baseline. In this synthetic case, that baseline includes required references, material controls, execution constraints, inspection and test requirements, hold points, and the evidence that must be retained.

Potential AI Assistance

- Draft package language from approved sources
- Check reference completeness
- Suggest sequencing
- Identify inconsistencies across documents
- Generate logistics or coordination prompts

Human / Accountable Authority

- Approve technical requirements and boundaries
- Define execution constraints and hold points
- Approve deviations
- Establish the evidence and instrumentation plan
- Baseline the authorized configuration

KEY INSIGHT

The AI may generate content, but the authoritative baseline must remain distinguishable from AI output and must be traceable to the approved technical authority.

6. Phase 3 — Multi-Authority Review and Approval

Operational Activity

The package moves through the applicable technical, quality, test, operational, and production authorities. Conflicts can emerge among schedule, access, test scope, material availability, and technical risk.

Potential AI Assistance

- Detect missing approvals
- Surface conflicts
- Summarize schedule or evidence impacts
- Route information to the correct review function

Human / Accountable Authority

- Resolve technical disagreements
- Accept or reject operational risk
- Approve deviations or compensating controls
- Determine whether execution may proceed

KEY INSIGHT

AI coordination is not authority. A workflow engine can identify a conflict; it cannot silently convert one organization's preference into another organization's approval.

7. Phase 4 — Material and Configuration Verification

Operational Activity

Before installation, the required material, revision, documentation, and configuration basis are verified. The scenario intentionally avoids real naval material pedigree rules or acceptance criteria; those details remain synthetic.

Potential AI Assistance

- Compare synthetic certifications and identifiers
- Detect documentation anomalies
- Flag configuration mismatches
- Predict schedule impact of missing material

Human / Accountable Authority

- Accept material equivalency

- Approve configuration changes
- Accept required quality evidence
- Authorize substitution or deviation

KEY INSIGHT

Commercial or schedule optimization does not establish operational acceptability. The approved configuration remains the governing technical baseline.

8. Phase 5 — Work Execution and Runtime Revalidation

Operational Activity

Authorized personnel begin the physical maintenance activity. During execution, unexpected conditions appear: a configuration discrepancy, an access limitation, an unexpected condition indication, or a conflict with adjacent work. These examples are synthetic and deliberately non-platform-specific.

Potential AI Assistance

- Present the current approved step
- Retrieve controlled references
- Detect checklist or evidence gaps
- Flag divergence between expected and observed state
- Recommend pause or escalation

Human / Accountable Authority

- Determine whether the changed condition is within the existing authority envelope
- Pause work when required
- Approve technical resolution
- Reauthorize execution after the basis changes

KEY INSIGHT

Authority that was valid at the start of the job may no longer be valid after a material condition changes. Runtime revalidation is therefore a design requirement, not an administrative afterthought.

9. Phase 6 — Testing, Evidence Closure, and Restoration

Operational Activity

After installation, the component and affected system undergo the synthetic test and restoration sequence defined by the approved baseline. Test evidence is compared with expected behavior and unresolved discrepancies are identified.

Potential AI Assistance

- Compare results with the synthetic expected range
- Detect anomalous patterns

- Check evidence completeness
- Recommend additional review or retest

Human / Accountable Authority

- Accept test results
- Approve retest requirements
- Determine technical readiness
- Authorize restoration to the next approved state

KEY INSIGHT

Analytical capability does not equal certification authority. A technically persuasive AI conclusion remains advisory unless the designated authority accepts the evidence and authorizes the resulting state change.

10. Phase 7 — Final Release and As-Executed Reconstruction

Operational Activity

Before final release, the organization reconciles the approved baseline with what actually occurred. The focus is not simply whether the required documents exist, but whether the execution record demonstrates that the correct authorities, gates, revalidations, interventions, and configuration controls were effective.

Potential AI Assistance

- Reconcile records
- Identify unresolved discrepancies
- Summarize deviations and interventions
- Assist reconstruction of the event sequence

Human / Accountable Authority

- Certify final technical status
- Accept residual operational risk
- Approve restoration or final release
- Direct corrective action when the evidence is insufficient

KEY INSIGHT

The as-built configuration and the as-executed record are separate engineering objects. Trust requires the ability to connect the specific execution back to the baseline and authority structure that governed it.

11. Evidence Required to Demonstrate Control

The reissued case study adds a requirement that was only implicit in the earlier version: the governance claim must be supported by an operating trace. A control point is not proven effective merely because it appears on a diagram or in a policy.

Evidence Element	What the Record Must Allow an Independent Reviewer to Determine
------------------	---

Baseline identifier	Which approved configuration and authority structure governed the action.
Actor / agent identity	Who or what performed or recommended the action.
Authority / mandate basis	What authority applied, with what scope and limits.
Input and evidence state	What information was available to the AI and human decision-maker.
Gate / validation outcome	Whether the required allow, deny, pause, escalation, or revalidation decision occurred.
Human intervention / override	Where a human approved, rejected, modified, paused, or recovered the process.
Actual execution path	What route the work actually followed, including exceptions and alternate paths.
Final action and result	What consequential action occurred and what operational state resulted.

KEY INSIGHT

Logging is not enough. The evidence must be sufficient to reconstruct the consequential action and relate it to the authoritative baseline that was in force at the time.

12. Declared Path vs. As-Executed Path

A planned maintenance workflow may appear perfectly controlled on paper and still diverge during execution. This case therefore distinguishes the declared path from the as-executed path.

- Declared path: the approved route by which information, recommendations, decisions, approvals, and actions are expected to flow.
- As-executed path: the actual route followed during the event, including pauses, overrides, retries, manual workarounds, alternate routes, and recovery actions.
- Divergence assessment: when the paths differ, the organization evaluates authority impact, risk impact, evidence sufficiency, configuration impact, and required corrective action.

A declared control is a design claim. An as-executed trace is operational evidence.

13. Strategic Conclusions

- High-consequence AI governance cannot rely solely on policy, inventories, reviews, dashboards, or static approval structures.
- Human accountability is meaningful only when human authority remains connected to the actual execution path and can still intervene before the action becomes irreversible.
- Authority must be revalidated when configuration, evidence, context, risk basis, or operating conditions materially change.
- Accountable technical authority and operational risk ownership are distinct functions that must remain explicit.
- As-built configuration and as-executed evidence must be distinguishable and correlatable.
- AI may accelerate analysis and workflow, but speed does not eliminate technical authority, risk ownership, evidence requirements, or final release accountability.

KEY INSIGHT

ControlPointAI treats AI governance as an engineering problem: map the data flow, map the authority flow, define the execution gates, preserve the evidence, control the configuration, and verify that authority still holds at $t + \Delta t$.

14. Preliminary Sneak Peek — UT Demonstration Project

The RX Maintenance Case Study is a conceptual governance case. The next step is to demonstrate the same ControlPointAI principles with a more detailed engineering evidence flow. ControlPointAI is developing a fictional submarine piping ultrasonic-testing (UT) demonstration using public-domain Navy / NDT context and purpose-built synthetic data.

What the Demonstration Will Show

- How synthetic UT inspection evidence enters a controlled engineering evidence package.
- How AI distinguishes observed evidence, inferred state, recommendation, and authorized action.
- How data provenance is recorded and retained.
- How AI assistance is bounded by execution-authority rules and human technical authorization.
- How an AI-assisted engineering work product can be generated from controlled evidence without allowing the AI to own the disposition or release decision.
- How the as-executed record is linked back to the governing as-built baseline, control points, authority envelope, and evidence requirements.

Security and Data-Provenance Boundary

The demonstration will use the submarine maintenance domain as context, not as a source of non-public technical implementation data. Publicly released terminology and high-level concepts may be used where their release status is established. Case-specific values and technical details will be synthetic unless a specific approved public source supports their use.

- SYN — purpose-built synthetic data for case-specific values, configurations, inspection results, thresholds, and demonstration records.
- PUB-N — clearly public U.S. Navy / DoD / government sources used only for the proposition they publicly establish.
- PUB-I — clearly public commercial, industrial, academic, or regulatory sources.
- GEN — generic engineering and governance principles not tied to a controlled platform.
- Excluded — personal recollection or reconstruction of non-public naval technical practice.

PRELIMINARY PROJECT NOTICE

The UT demonstration is under development. It is intended to show ControlPointAI's evidence-flow, authority-flow, configuration-control, provenance, and bounded-AI methods in a fictional engineering inspection workflow. No ship-specific configuration, real UT results, non-public acceptance criteria, controlled procedures, internal Navy records, or restricted technical information will be used.

15. Final Position

This reissue returns the RX Maintenance Case Study to the ControlPointAI public portfolio under a more mature engineering and security baseline. The lesson is not that high-consequence naval maintenance is too sensitive to examine. The lesson is that the boundary must be engineered: use public context where appropriate, synthetic case data for the demonstration, explicit provenance, controlled technical claims, accountable authority, and reconstructable execution evidence.

AI can inform. Humans authorize. Evidence must show that the boundary held during execution.

Document Basis

- ControlPointAI AI-1 — Systems Engineering Foundation, Version 3.1, Proposed As-Built / As-Executed Evidence Clarification, August 11, 2026.
- ControlPointAI AI-2 — Technical Governance Framework, Version 1.1, Baseline-to-Execution Traceability Clarification, August 11, 2026.
- ControlPointAI UT Submarine Demonstration — Security, Data Provenance, and Synthetic Demonstration Ground Rules, Rev A, August 17, 2026.
- ControlPointAI Case Study — Nuclear Valve Impeller Replacement During Naval Availability, Draft v1 (superseded by this reissue).