CONVERGED SECURITY | RESEARCH COMPANIONS

PUBLIC SAMPLE EDITION

Governing AI-Driven Infrastructure

Telemetry, Autonomy and Systemic Assurance

Prof. Dr. Vladimir Bunic

Global Converged Security Board



Title and publication overview

Governing AI-Driven Infrastructure
Telemetry, Autonomy and Systemic Assurance
Public sample edition | Version 1.0 | 11 November 2026

Author: Prof. Dr. Vladimir Bunic
Owner and distributor: Global Converged Security Board

A secondary practical/reference companion to the primary article identified in the research bibliography, read alongside the issued GCSB standards, mandatory procedures and advisory playbooks.

Eight explanatory readings and three fictional reflection exercises develop questions of evidence, authority and practical judgement. Exact reading crosswalks connect the commentary to issued standards, mandatory procedures and advisory playbooks. This is not a course study guide, a complete training sequence or a research article.

Research method and limits. Qualitative multi-scenario reconstruction of risk propagation; not field validation of a deployed AI control system.

Publication chronology. The primary article's registered online publication date is 31 October 2026. The CSF standards cited here were issued 13 October 2026; procedures and playbooks followed on 19 October 2026. This companion was prepared for publication on 11 November 2026. Its crosswalk is a later educational comparison, not evidence that the article used these issued editions.

Copyright and reader information

Copyright © 2026 Global Converged Security Board.
All rights reserved.

Except as permitted by applicable law, reproduction, redistribution or adaptation of this companion requires prior written permission. Source articles retain their separate licences; the companion’s rights notice does not change them.

Applicable law and adopted issued standards govern requirements. Mandatory procedures implement parent requirements; playbooks are advisory. Implementation milestones use the starting event stated by the relevant parent standard. Record organisational adoption and certification commencement separately; the Main Standard uses documented adoption for a new adopter. Article publication, source issue and book publication dates do not start an organisation's implementation period.

All exercises are fictional. Use no personal or confidential workplace data and make no live operational changes. Full worked solutions, complete training scenarios and instructor guidance are excluded.

The companion itself has not been peer reviewed and is not a GCSB Journal article. Source articles retain their original authorship, institutional attribution and licences. No ISBN is assigned. Cover imagery is AI-generated illustration.

Suggested citation. Bunic, V. (2026). Governing AI-Driven Infrastructure: Telemetry, Autonomy and Systemic Assurance. Public sample edition (v1.0). Global Converged Security Board. https://doi.org/10.5281/zenodo.24129171

Contents

AI changes the decision pathway **4**

Telemetry integrity and model alignment **5**

Autonomy with bounded authority **6**

SOC and GSOC as evidence coordinators **7**

The CSF lens: integration does not require unbounded AI **8**

A model assurance record is a decision record **9**

Autonomy boundaries cross safety and privacy boundaries **10**

Systemic assurance needs defined measurements **11**

Sample exercise 1: A plausible but disputed reading **12**

Sample exercise 2: A wider autonomy boundary **13**

Sample exercise 3: Several small anomalies **14**

Research bibliography and evidence **15**

Reading crosswalk: explanations **16**

Reading crosswalk: sample exercises **17**

Issued-source bibliography **18**

Issued-source bibliography (continued) **19**

AI changes the decision pathway

The primary article examines critical infrastructure where AI, digital twins, IoT and autonomous OT participate in operational decisions. Risk assessment must therefore consider how information becomes a recommendation or action, not only whether a device is connected.

A useful distinction is between observation, prediction, recommendation and execution. An error in one stage may affect the next, but the effect depends on the authority and controls around each transition.

A sophisticated prediction can still rest on poor inputs. A connected sensor can produce plausible but incorrect telemetry. Governance asks who owns the data quality question and how contrary evidence reaches decision-makers.

Original companion example: a water-treatment model recommends an operating change while a physical inspection contradicts the inferred equipment state. The question is not simply which display is newer, but how evidence is reconciled before authority is exercised.

Reading bridge. GCSB-GDL-TIA-001, Play 1, Inventory and map the estate: identify ownership, dependencies and system purposes, including converged assets.

Telemetry integrity and model alignment

The article analyses telemetry corruption and digital twin drift as propagation pathways. A digital representation can gradually diverge from the real process, making a model appear coherent while its operational assumptions become unreliable.

Validation should distinguish a syntactically valid measurement from a credible account of reality. Time, provenance, calibration context and independent observations can affect whether a reading is suitable for a particular decision.

Model alignment asks whether the model’s assumptions and intended use still match the operating environment. A successful historical prediction does not establish continued suitability after maintenance, environmental change or a changed dependency.

Challenge mechanisms need an accountable owner. An operator who notices a contradiction should know how to escalate it without being required to prove an entire technical failure before the concern is heard.

Reading bridge. GCSB-GDL-TIA-001, Play 4, Operate and monitor the architecture: monitoring needs defined observations and follow-through, not merely collected logs.

Autonomy with bounded authority

The article considers autonomy escalation and automated misclassification. These raise a governance question: what may the system decide or change, under which conditions, and when must responsibility return to an authorised person?

Human involvement is not automatically meaningful. A person shown an opaque recommendation without context, time or authority may be unable to challenge it. Oversight should examine whether intervention is feasible and consequential.

Changes to an autonomy boundary deserve consideration as operational changes, not merely software features. Increasing an automated system’s reach can alter safety, identity, privacy and continuity assumptions across domains.

An exception or emergency mode should not silently become the normal operating model. Evidence of activation, authority, observed consequences and restoration of the intended boundary helps keep responsibility visible.

Reading bridge. GCSB-GDL-TIA-001, Play 3, Secure the build and change process. This companion gives no automated-control instructions or mandatory autonomy scale.

SOC and GSOC as evidence coordinators

The article proposes governance mechanisms connecting telemetry validation, autonomy monitoring, prediction-model alignment and systemic risk monitoring. A SOC or GSOC can help coordinate these perspectives without becoming the sole owner of every operational decision.

Systemic monitoring asks how apparently local anomalies combine across services or sectors. Similar alarms can have different consequences depending on dependencies, operating state and the ability to intervene.

A common operational picture should retain uncertainty. Conflicting telemetry, delayed observations and unexplained model behaviour should not be flattened into a single green status merely to simplify reporting.

The qualitative scenario method identifies plausible pathways. It does not establish incident frequencies, guaranteed detection rates or universal control effectiveness. These require appropriate empirical work and organisation-specific evidence.

Reading bridge. GCSB-GDL-RMA-001, Play 6, Integrate risk into decisions: use observations to inform authorised decisions and revisable risk explanations.

The CSF lens: integration does not require unbounded AI

GCSB-ST-CSF-001 v2.0 integrates cybersecurity, physical security and operational technology security. Personnel, information, identity, suppliers and privacy are cross-cutting concerns. AI is a capability within that scoped estate, not a fourth security domain or an independent authority.

Main Standard clauses 6.4.1 and 6.4.3 explicitly read monitoring and AI/ML deployment through GCSB-ST-CSF-TIA-001. The Technology Standard's clauses 6.1 and 6.6 require proportionate architecture and analytics assurance; a coordinated or federated architecture can evidence equivalence.

GCSB-PRC-CSF-001 Step 6 links technology, identity and data controls.
GCSB-PRC-TIA-001 Steps 1-2 require real scope, material dependencies and justified exclusions. A nonexistent OT estate should not be invented to make a diagram appear converged.

GCSB-GDL-CSF-001 Play 3 and GCSB-GDL-TIA-001 Play 2 provide advisory implementation perspectives. The research article's AI scenarios help question decision pathways; neither their presence nor this book establishes that an adopter must automate every operation.

A model assurance record is a decision record

GCSB-ST-CSF-TIA-001 clause 6.6 and GCSB-ST-CSF-RMA-001 clause 6.7 connect analytics to purpose, data provenance, assumptions, limitations, validation, uncertainty and accountable human authority. A plausible output is not proof that the input represents the physical estate.

GCSB-PRC-TIA-001 Step 7 makes model limitations, drift and error measures part of assurance. GCSB-PRC-RMA-001 Step 7 adds the risk-decision context. Validation should concern the intended task, operating conditions and consequences of error rather than an isolated accuracy headline.

GCSB-GDL-TIA-001 Play 4 supports monitoring interpretation. A digital twin discrepancy can prompt questions about timestamp alignment, source identity, changed equipment and contrary observation, without this companion prescribing process settings or a diagnostic sequence.

A useful assurance record distinguishes observed facts, model interpretation and decision authority. It states what evidence would trigger review and who can challenge or suspend reliance under authorised arrangements. No universal accuracy threshold or validated plant-control claim is introduced here.

Autonomy boundaries cross safety and privacy boundaries

GCSB-ST-CSF-TIA-001 clauses 6.7 and 6.9 distinguish authorised response from uncontrolled automation and require material changes to be assessed. An update that changes a recommendation into an executed action changes authority, not merely presentation.

GCSB-PRC-TIA-001 Steps 8 and 10 address documented authority, safety, privacy, continuity, logging, fallback and change evidence. Domain owners retain their responsibilities; coordinated governance does not let an analytics supplier assume the adopting organisation's decision rights.

GCSB-GDL-TIA-001 Play 3 is an advisory build-and-change reference. The parent procedure is the source for mandatory response safeguards even when a playbook's traceability table points to a narrower set of steps.

The companion's software-update exercise asks which assumptions and authorities changed. It deliberately omits an approval decision or implementation plan. Human oversight should be examined for actual access to evidence and capacity to intervene, not simply the existence of a nominal approver.

Systemic assurance needs defined measurements

GCSB-ST-CSF-TIA-001 clause 6.10 and GCSB-ST-CSF-RMA-001 clause 6.6 require defined populations, calculations, evidence, measurement windows, exclusions, owners and escalation. An unexplained 'real-time' or coverage percentage can hide critical missing sources.

GCSB-PRC-TIA-001 Steps 5 and 10 separate integrated visibility from metric definition. Latency should be interpreted for a task and measured event. An aggregated dashboard should preserve contradictory evidence rather than flatten uncertainty into a single status.

GCSB-GDL-TIA-001 Plays 4-5 discuss operation and review. A SOC or GSOC can coordinate evidence while operational, privacy and executive owners retain their decisions. Several small anomalies sharing a supplier deserve a dependency question, not an unsupported causal conclusion.

The Main Standard clause 7 keeps its 12-month baseline distinct from the Technology Standard's 18-month domain track. The extended track is not a permission to postpone baseline risk controls. The article's reconstructed scenarios establish plausible pathways, not measured frequency or guaranteed detection.

Sample exercise 1: A plausible but disputed reading

Fictional situation. A plant model shows normal pressure. An authorised technician reports an inconsistent physical observation. The dashboard classifies the report as sensor noise.

Reflection prompts. Separate telemetry, interpretation and authority. What evidence could reconcile the disagreement? Who must be able to escalate it?

Your output. An evidence-question list. Do not alter any live process or prescribe a plant setting.

Further reading. GCSB-GDL-TIA-001, Plays 1 and 4. No solution is supplied.

Sample exercise 2: A wider autonomy boundary

Fictional situation. A software update changes an analytics tool from recommending maintenance priorities to automatically authorising selected changes. The asset owner was not consulted.

Reflection prompts. What governance assumptions have changed? Which specialists need to assess dependencies and intervention capability?

Your output. A change-review question set, not an implementation plan.

Further reading. GCSB-GDL-TIA-001, Play 3. No model approval decision is included.

Sample exercise 3: Several small anomalies

Fictional situation. Three connected service operators each report minor telemetry inconsistencies. Their shared supplier classifies each ticket separately and reports no material outage.

Reflection prompts. What shared dependency might deserve examination? How should a coordinator communicate a plausible systemic concern without claiming an established cause?

Your output. A concise uncertainty statement and evidence request.

Further reading. GCSB-GDL-RMA-001, Play 6. No complete incident scenario or answer key is included.

Research bibliography and evidence

Primary article. Bunic, V. (2026). GOVERNING CROSS-DOMAIN RISK IN AI-DRIVEN CRITICAL INFRASTRUCTURE. International Journal of Science and Research Archive, 21(1), 296-302. https://doi.org/10.30574/ijsra.2026.21.1.1877

Article history. Received 3 September 2026; revised 10 October 2026; accepted 12 October 2026. Online publication: 31 October 2026 (Crossref published-online field; not DOI creation or deposit date).

Source identity. Vladimir Bunic; Converged Security Institute (CSI), Barcelona, Spain.. Publisher article | Original article PDF. Sources checked 11 November 2026.

Declared journal review policy. Read the publisher policy. Public policy describes the journal’s stated process, not an independent audit of the article’s confidential reports.

Scope of this companion. Qualitative multi-scenario reconstruction of risk propagation; not field validation of a deployed AI control system. No quantitative benefit, causal effect or universal compliance outcome is inferred.

Source rights. Publisher page and original PDF state CC BY 4.0. No original article text, figures or tables are reproduced. New companion commentary retains its separate reserved-rights notice.

Related primary families. TIA, RMA, GLA, IRBC. The following bibliography links the full issued framework; only the indicated families are central to this volume.

Reading crosswalk: explanations

Standards v2.0 set adopted requirements; procedures v1.0 set mandatory steps; playbooks v1.0 offer advisory methods. The following is an interpretive reading crosswalk, not an operational checklist or a new requirement.

AI changes the decision pathway
GCSB-ST-CSF-TIA-001 §6.1 | GCSB-PRC-TIA-001 Step 2 | GCSB-GDL-TIA-001 Play 1
Trace observation, interpretation and authority through real interfaces.

Telemetry integrity and model alignment
GCSB-ST-CSF-TIA-001 §6.6 | GCSB-PRC-TIA-001 Step 7 | GCSB-GDL-TIA-001 Play 4
Examine model drift, provenance and task-specific validation.

Autonomy with bounded authority
GCSB-ST-CSF-TIA-001 §6.7 and 6.9 | GCSB-PRC-TIA-001 Step 8 and 10 | GCSB-GDL-TIA-001 Play 3
Recognise autonomy changes as authority and operational changes.

SOC and GSOC as evidence coordinators
GCSB-ST-CSF-RMA-001 §6.6-6.7 | GCSB-PRC-RMA-001 Step 7 | GCSB-GDL-RMA-001 Play 5
Preserve uncertainty in systemic risk reporting.

Reading crosswalk: sample exercises

Standards v2.0 set adopted requirements; procedures v1.0 set mandatory steps; playbooks v1.0 offer advisory methods. The following is an interpretive reading crosswalk, not an operational checklist or a new requirement.

Sample exercise 1: A plausible but disputed reading
GCSB-ST-CSF-TIA-001 §6.6 | GCSB-PRC-TIA-001 Step 7 | GCSB-GDL-TIA-001 Play 4
Question conflicting readings without issuing plant instructions.

Sample exercise 2: A wider autonomy boundary
GCSB-ST-CSF-TIA-001 §6.7 and 6.9 | GCSB-PRC-TIA-001 Step 8 and 10 | GCSB-GDL-TIA-001 Play 3
Request authorised change evidence without approving deployment.

Sample exercise 3: Several small anomalies
GCSB-ST-CSF-RMA-001 §6.2 and 6.6 | GCSB-PRC-RMA-001 Step 2 and 7 | GCSB-GDL-RMA-001 Play 2
Explore shared dependencies without claiming established cause.

Issued-source bibliography

Primary standards and CCSO Policy: v2.0, issued 13 October 2026. Mandatory procedures and advisory playbooks: v1.0, issued 19 October 2026. Verify the current controlled edition and applicable provisions before use.

GCSB-CSF-STRAT | Converged Security Framework: Strategic Standards for Unified Risk Management and Organisational Resilience
Read issued primary source | GCSB-PRC-STRAT-001 | GCSB-GDL-STRAT-001

GCSB-ST-CSF-001 | Converged Security Framework Main Standard
Read issued primary source | GCSB-PRC-CSF-001 | GCSB-GDL-CSF-001

GCSB-ST-CSF-TRA-001 | Training and Awareness
Read issued primary source | GCSB-PRC-TRA-001 | GCSB-GDL-TRA-001

GCSB-ST-CSF-GLA-001 | Governance and Leadership Framework
Read issued primary source | GCSB-PRC-GLA-001 | GCSB-GDL-GLA-001

GCSB-ST-CSF-TIA-001 | Technology Integration and Architecture
Read issued primary source | GCSB-PRC-TIA-001 | GCSB-GDL-TIA-001

Issued-source bibliography (continued)

Primary standards and CCSO Policy: v2.0, issued 13 October 2026. Mandatory procedures and advisory playbooks: v1.0, issued 19 October 2026. Verify the current controlled edition and applicable provisions before use.

GCSB-ST-CSF-IRBC-001 | Incident Response and Business Continuity Framework
Read issued primary source | GCSB-PRC-IRBC-001 | GCSB-GDL-IRBC-001

GCSB-ST-CSF-IAM-001 | Identity and Access Management Standard
Read issued primary source | GCSB-PRC-IAM-001 | GCSB-GDL-IAM-001

GCSB-ST-CSF-DPP-001 | Data Protection and Privacy Standard
Read issued primary source | GCSB-PRC-DPP-001 | GCSB-GDL-DPP-001

GCSB-ST-CSF-RMA-001 | Converged Security Risk Management and Assessment Standard
Read issued primary source | GCSB-PRC-RMA-001 | GCSB-GDL-RMA-001

GCSB-POL-CSF-CCSO-001 | Chief Converged Security Officer (CCSO) Policy
Read issued primary source | GCSB-PRC-CCSO-001 | GCSB-GDL-CCSO-001

RESEARCH COMPANIONS | PUBLIC SAMPLE

Governing AI-Driven Infrastructure

Telemetry, Autonomy and Systemic Assurance

From research ideas to better questions: eight explanatory readings, three fictional exercises and exact issued-source crosswalks. A separate educational companion, not a standard or peer-reviewed research article.

Prof. Dr. Vladimir Bunic
Global Converged Security Board
Owner and distributor | All rights reserved