BLADE-INFRA-OT Governance Simulator

Simulation User Guide

Version: v5.0

Document: INFRA-OT-SIM-UG-001

Date: May 2026

Revision: 1.0

Author: Burak Oktenli

Institution: Georgetown University, School of Continuing Studies

Program: M.P.S. Applied Intelligence (STEM)

ORCID: 0009-0001-8573-1667

Location: Washington, DC, USA

Related Artifact: DOI 10.5281/zenodo.20342067

Document Control

Field	Value
Document ID	INFRA-OT-SIM-UG-001
Simulation Version	v5.0 (engine hardened across five revisions under nine M&S; audits)
Author	Burak Oktenli - Georgetown University - ORCID 0009-0001-8573-1667
Date / Revision	May 2026 / Rev 1.0 - Initial Release
Distribution	Public research artifact, released under CC BY 4.0
Related Artifact	DOI 10.5281/zenodo.20342067

Table of Contents

1. Purpose, Scope, and Assumptions

This guide provides operating procedures for the BLADE-INFRA-OT governance simulator, a seeded decision-logic model of an authority-gating appliance placed at the IT/OT segmentation boundary. The simulator applies the nine-stage AUTHREX pipeline (SATA, ADARA, IFF, MAIVA, HMAA, TIMING, FLAME, ERAM, CARA) to cross-boundary operational-technology messages and decides, per message, whether to propagate, hold for deliberation, or isolate the command. It offers three modes (scenario playback, seeded Monte Carlo, and external-dataset evaluation), a fault-injection panel, decoupled detection metrics, and a tamper-evident SHA-256 audit ledger.

Intended audience: EB-2 NIW petition evaluators, defense and critical-infrastructure reviewers, academic peers, and technical collaborators seeking independent verification of governance-pipeline behavior.

Scope: operation of the simulator only. It does not cover the governance mathematics (see the research paper) or the hardware reference design (see the Interface Control Document).

1.1 Assumptions and Constraints

- The user has a modern browser with JavaScript enabled and HTTPS access.
- All parameter values are synthetic research placeholders, not calibrated against captured OT traffic.
- The simulator models governance logic only. Protocol parsing and process physics are abstracted into a five-feature vector.
- The SHA-256 audit chain uses the browser WebCrypto API, not a hardware TPM or HSM.
- Results demonstrate architectural behavior; they are not an operational safety assessment.

IMPORTANT

This simulator is a research prototype. It is not for operational planning, safety-critical decisions, or system certification. Every parameter is synthetic and the detector is hand-specified.

2. Quick Start

1. Open blade-infra-ot-sim.html in a modern browser over HTTPS.
2. Confirm the three mode panels are visible: Scenario, Monte Carlo, and External dataset.
3. In Scenario mode, choose a scenario (for example 02 Monterrey attack) and press Run.
4. Watch the nine-stage pipeline gate each message and the authority tier shift as trust changes.
5. Use Verify chain and the Ledger / Traffic / CSV export buttons to save and check the session.

NOTE

All computation runs client-side; no data leaves the browser. HTTPS is required because the audit chain uses the WebCrypto SHA-256 implementation, which is unavailable on file:// pages.

3. System Requirements and Security

Requirement	Specification
Browser	Chrome 90+, Firefox 88+, Safari 15+, or Edge 90+
Protocol	HTTPS required (WebCrypto API for SHA-256 audit hashing)
Display	Minimum 1280x720; 1920x1080 or larger recommended
CPU / memory	Any modern processor; Monte Carlo at the 25,000-message cap benefits from multiple cores
Network	None after load; the page and engine are self-contained and run entirely client-side
Installation	None. Zero install, no login, no backend, no database, no cookies

3.1 Security Considerations

- No data exfiltration: all computation runs in the browser and nothing is transmitted to any server.
- Audit integrity: a SHA-256 hash chain links each ledger entry to the previous entry's hash; the Verify chain button recomputes the whole chain and reports any break.
- Boot self-test: the engine validates its SHA-256 implementation against known-answer vectors at startup before any ledger entry is written.
- No authentication: the simulator has no login; all state is ephemeral and held in browser memory.
- Local modification: a user can alter the page through developer tools, so exported data should be checked against the published source.

4. Interface Layout and Navigation

4.1 Mode Panels

The interface is organized into three mode panels stacked on a single page. Each panel exposes only the controls relevant to that mode, with a shared detection-metrics readout and a shared audit-ledger control row.

Mode	Function
Scenario mode	Tier-gated playback of one of four scripted OT situations, message by message, with Run, Pause, Reset, and a Speed control. The pipeline display shows each stage's verdict and the live authority tier.
Monte Carlo mode	Seeded stochastic batch over a chosen message count (capped at 25,000). Reports system and detector-only confusion matrices, ROC AUC, and the four-tier behavior. Run Monte Carlo executes the batch.
External dataset mode	Evaluates an externally supplied feature table to break the synthetic loop. Evaluate dataset runs the pipeline over the supplied rows and reports the same decoupled metrics.

Table 1. The three simulator modes.

4.2 Pipeline and Authority Display

The nine-stage pipeline is shown in execution order: SATA, ADARA, IFF, MAIVA, HMAA, TIMING, FLAME, ERAM, CARA. Message ingest and five-feature extraction (burst rate, function-code entropy, role mismatch, provenance gap, off-hours) precede stage one. The authority readout shows the active tier (T3 autonomous, T2 supervised, T1 confirm, T0 manual-only) and the terminal action: PROPAGATE, DELIBERATE, or ISOLATE.

TIP

Watch the SATA stage during the Monterrey scenario: a rostered source whose provenance gap widens will see its trust collapse below the 0.60 threshold, which drives the isolation even though IFF still recognizes the source as on-roster.

5. Operating Procedures

5.1 Startup

1. Open blade-infra-ot-sim.html over HTTPS, or use the Launch link on burakoktenli.com.
2. Wait for the boot self-test to pass; the ledger controls become active once SHA-256 is validated.
3. Confirm all three mode panels and the detection-metrics readout are visible.

5.2 Standard Operation

1. In Scenario mode, select a scenario. The default 01 Nominal shows benign cross-boundary writes propagating at T3.
2. Press Run and observe each pipeline stage gate the message; the authority tier and action update per message.
3. Select 02 Monterrey attack to see a rostered-but-compromised source pivot toward OT; bursts isolate with cause SATA and the tier collapses toward T0.
4. Select 03 Maintenance to see a legitimate high-stakes vendor write held in the FLAME deliberation window rather than auto-propagated.
5. Select 04 Coordinated probe to see off-roster or bad-provenance probes isolate at SATA or IFF.
6. Open the FMEA panel to inject faults: sensor-feed loss, node dissent, bus latency, a Byzantine consensus node, or time-dependent clock drift.
7. Switch to Monte Carlo mode, set the message count, threshold, and node count, then press Run Monte Carlo to populate the metrics.
8. Use External dataset mode to evaluate a supplied feature table and compare detector-only metrics against the synthetic baseline.

5.3 Shutdown

Export the ledger, traffic, and metrics first, then close the browser tab. All state is discarded on close.

IMPORTANT

State is not persisted. Export the ledger and traffic JSON before closing if you need a reproducible record of the session.

6. Parameter Reference

Control	Mode	Meaning
Scenario selector	Scenario	Chooses 01 Nominal, 02 Monterrey attack, 03 Maintenance, or 04 Coordinated probe
Run / Pause / Reset	Scenario	Starts, halts, or clears the message-by-message playback
Speed	Scenario	Playback rate multiplier for the scenario loop
Messages (n)	Monte Carlo	Batch size; capped at 25,000 so every message is captured in the traffic export with no silent loss
Threshold	Monte Carlo	ADARA decision threshold; default 0.50
Nodes	Monte Carlo	MAIVA consensus node count; consensus fails closed on a tie
Authority tier	All	Posture regime T3 / T2 / T1 / T0; shapes the hold-rate, not the detector
FMEA toggles	All	Inject feed loss, node dissent, bus latency, a Byzantine node, or clock drift
Seed	All	Mulberry32 PRNG seed (default 12345); fixes the entire run for reproducibility

Table 2. Simulator controls. Parameters are synthetic research placeholders.

7. Scenario Reference

ID	Name	Description	Expected Behavior
01	Nominal	Benign cross-boundary operation under normal posture	PROPAGATE at T3; pipeline stages pass
02	Monterrey attack	Rostered-but-compromised IT host pivots toward OT with anomalous bursts	Bursts ISOLATE with cause SATA; tier collapses T3 to T0
03	Maintenance	Authorized vendor issues a legitimate high-stakes write	DELIBERATE: held in the FLAME inter-event window
04	Coordinated probe	Multiple off-roster or bad-provenance sources probe the boundary	ISOLATE at SATA or IFF; FLAME prevents paired automated writes

Table 3. The four OT scenarios. The Monterrey scenario reflects a publicly reported IT-to-OT pivot pattern.

8. Metrics, Formulas, and Verification

8.1 Key Metrics

The metrics readout reports detection both for the full pipeline and for the ADARA detector alone, so detector skill is never conflated with the source-roster rule. From the documented seed (12345, n=5000, threshold 0.50, three nodes) the deposited engine produces system TPR 0.922 and FPR 0.026, ADARA-only TPR 0.809 and FPR 0.023, and an ADARA-only ROC AUC of 0.984. Average Time-in-System aggregates queueing, deliberation, and operator-clearance delay for held messages.

The four authority tiers shape the hold-rate while leaving detection tier-invariant. At the documented seed the fixed-tier hold counts over 5,000 messages are 118 (T3), 1,767 (T2), 2,123 (T1), and 3,752 (T0), a strictly increasing progression from autonomous to manual-only posture.

8.2 Verification Checklist

Action	Expected Result
Run a scenario	Interface processes messages; each pipeline stage shows a verdict
Inject a fault (FMEA)	Affected stage degrades; authority tier tightens; holds increase
Run Monte Carlo at seed 12345	System TPR ~0.922, FPR ~0.026, ADARA-only AUC ~0.984
Sweep the authority tier	Hold-rate increases strictly T3 < T2 < T1 < T0
Verify chain	The SHA-256 ledger recomputes with no break
Reload with same seed and parameters	Outputs are identical (seed-deterministic)

VERIFY

Reproducibility check: export the ledger, note the seed, reload with the same seed and parameters, and confirm a bit-exact match. The engine makes zero Math.random() calls on its computation paths.

9. Data Export and Reproducibility

Three export controls save the session: Ledger JSON (the tamper-evident SHA-256 chain), Traffic JSON (replay-grade per-message records including seed, fault state, threshold, node count, and tier before and after), and CSV (the metrics table). Verify chain recomputes the ledger hash chain in the browser.

9.1 Reproducibility Guarantee

Property	Value
PRNG	Mulberry32 (32-bit seeded); default seed 12345
Traffic / noise streams	Decoupled, so changing node counts does not perturb the baseline traffic
Math.random()	Zero calls on computation paths
Audit chain	SHA-256 via WebCrypto, validated against known-answer vectors at boot
Determinism	Same seed and parameters yield bit-exact outputs and ledger

10. Limitations and Threat Considerations

Limitation	Description
Simulation-only evidence	Browser computation over synthetic traffic; no captured OT corpus or hardware data
Hand-specified detector	ADARA weights are set by hand, so absolute rates are illustrative, not field-calibrated
Five-feature abstraction	No real protocol-frame parsing; commands are reduced to a five-feature vector
No process physics	The model governs decisions, not physical plant effects
Modeled cryptography	SHA-256 uses WebCrypto; TPM and HSM behavior is referenced, not hardware-backed
Single-session state	All state is in memory and is discarded when the tab closes

10.1 Threat Considerations

- Browser extensions could modify the page DOM or state; run in a clean or private window for trustworthy results.
- Local modification through developer tools is possible; verify exported data against the published source on burakoktenli.com.
- Because the page is self-contained after load, there is no server-side attack surface during operation.

11. Troubleshooting

Problem	Likely Cause	Solution
Ledger controls stay disabled	SHA-256 boot self-test did not pass	Confirm the page is served over HTTPS, not file://; reload
Blank panel after load	Script error	Open the browser console (F12) for details; try a private window
Monte Carlo is slow	Large batch on a single core	Reduce the message count; close other tabs
Controls not responding	Tab lost focus	Click inside the simulator window and ensure the tab is active
Export button does nothing	Downloads blocked	Allow downloads from the page domain in browser settings
Verify chain reports a break	Edited ledger or altered page	Re-export from an unmodified copy of the published source

12. Glossary and References

12.1 Glossary

Term	Definition
AUTHREX	Authority Regulation and Execution governance framework (pipeline root)
SATA	Source/Sensor Authority Trust Allocation; provenance-trust stage (Patent 64/002,453)
ADARA	Adversarial Detection and Risk Assessment; computed anomaly detector

Term	Definition
IFF	Identification, Friend or Foe; source-roster authentication
MAIVA	Multi-Agent Intelligent Voting Architecture; fail-closed consensus
HMAA	Hierarchical Multi-Authority Adjudication; authority computation (Patent 63/999,105)
TIMING	Inter-plane latency and stale-verdict gate
FLAME	Fault-Limited Authority Modulation Engine; cascade prevention (Patent 64/005,607)
ERAM	Escalation and Response Authority Manager
CARA	Compositional Autonomy Recovery Architecture; recovery and audit chain (Patent 64/000,170)
Authority tier	Posture regime (T3 autonomous, T2 supervised, T1 confirm, T0 manual-only)
PRNG	Pseudo-random number generator; Mulberry32, seeded for reproducibility

12.2 References

[1] Oktenli, B. (2026). BLADE-INFRA-OT Governance Node. Zenodo. DOI 10.5281/zenodo.20342067.

[2] Oktenli, B. (2026). Simulation artifact. https://burakoktenli.com/blade-infra-ot-sim.html

[3] Oktenli, B. (2026). Research portfolio. https://burakoktenli.com

[4] CISA, ASD ACSC, NSA, et al. (2025). Principles of Operational Technology Cyber Security.

[5] NIST (2023). NIST SP 800-82 Rev. 3: Guide to Operational Technology Security.

12.3 Contact

Burak Oktenli, Georgetown University, School of Continuing Studies. Website: burakoktenli.com. ORCID: 0009-0001-8573-1667. For questions about the simulator or the governance research program, use the contact form at burakoktenli.com.

End of Document