BURAKOKTENLI.COM · GEORGETOWN UNIVERSITY · ORCID 0009-0001-8573-1667

BLADE-FINANCE Governance Simulator

Simulation User Guide

Version: v2.2 (High-Assurance Research Build)

Document: FINANCE-SIM-UG-001

Date: May 2026 · **Revision:** 1.0

Author: Burak Oktenli

Institution: Georgetown University, M.P.S. Applied Intelligence

ORCID: 0009-0001-8573-1667

Related Artifact: DOI 10.5281/zenodo.20374692

Document Control

Document ID	FINANCE-SIM-UG-001
Simulation Version	v2.2 (High-Assurance Research Build)
Author	Burak Oktenli, Georgetown University, ORCID 0009-0001-8573-1667
Date / Revision	May 2026 / Rev 1.0, Initial Release
Distribution	Public release, CC BY 4.0 (research artifact)
Related Artifact	DOI 10.5281/zenodo.20374692

Table of Contents

1. Purpose, Scope, and Assumptions

This guide provides operating procedures for the BLADE-FINANCE Governance Simulator, which implements authority governance for financial-sector AI decision systems. The simulator adjudicates a stream of synthetic transactions through the AUTHREX eight-stage pipeline and a four-tier HMAA authority model, appending every decision to a tamper-evident SHA-256 audit ledger. Six scenarios cover nominal traffic, deepfake authentication, AI-agent coordinated attack, high-value risk, a realistic mixed stream, and a low-and-slow coordinated ring, with a retrospective swarm-review module for ring recovery.

Intended audience: technical reviewers, academic peers, and collaborators seeking independent, reproducible verification of governance-pipeline behavior.

Scope: operation of the simulator. It does not cover the mathematical derivations (see the published paper) or the reference hardware specification (see the deposit files).

1.1 Assumptions and Constraints

- The user has a modern browser with JavaScript enabled.
- All transaction features and labels are synthetic research data, not captured financial-institution data.
- The simulator models governance decision-logic only; it does not model a production payment rail or a real fraud model.
- The SHA-256 audit chain uses a built-in JavaScript implementation verified byte-exact against FIPS 180-4 vectors; it is not a hardware HSM or TPM.
- Results demonstrate architectural and decision-logic behavior, not operational fraud-detection performance.

> **IMPORTANT.** This simulator is a research prototype at TRL 3-4. It is not for operational planning, transaction screening, or safety-critical decisions. All data are synthetic. The reported recall is an actionable-risk triage measure, not an empirical fraud-detection rate.

2. Quick Start

1. **Open** blade-finance-simulation.html in any modern browser (double-click works; no server needed).
2. **Select a scenario** from the scenario selector (start with Nominal, then try Realistic mixed).
3. **Click Run** to stream transactions, or **Step** to advance one record at a time.
4. **Watch the eight-stage pipeline** adjudicate each transaction and the HMAA authority tier route it to cleared, review, or blocked.
5. **Click Self-Test** (expect 9 of 9 pass) and **Export JSON** to save the session and audit trail for verification.

> **NOTE.** All computation runs client-side and offline. No data leaves the browser, there is no backend, and no network connection is required after the file is opened.

3. System Requirements and Security

Requirement	Specification
Browser	Chrome 90+, Firefox 88+, Safari 15+, Edge 90+
Protocol	None. Runs from a local file (file://) or any host; no HTTPS or secure context required.
Dependencies	None. Single self-contained HTML file, no CDN, no external scripts, no build step.
Display	Minimum 1280x720; 1920x1080 or larger recommended
CPU / Memory	Any modern processor; Monte Carlo (2,000 trials) completes in seconds on a single core
Network	Not required at any point; fully offline
Installation	None: no install, no login, no backend, no database, no cookies

3.1 Security Considerations

- **No data exfiltration:** all computation runs in the browser; nothing is sent to any server.
- **No third-party code:** the simulator is a single self-contained file with no CDN or external dependency, removing the supply-chain surface of remotely loaded scripts.
- **Audit integrity:** a SHA-256 hash chain links each ledger entry to the previous entry's hash over a canonical serialization; the built-in self-test re-walks and re-verifies the chain.
- **Ephemeral state:** there is no authentication and no persistence; all state is held in browser memory and is discarded when the tab closes.

4. Interface Layout and Navigation

4.1 Panel Layout

The interface presents the transaction-governance workflow in linked panels. **Scenario panel:** scenario selector plus a Scenario Briefing describing the active threat. **Transaction feature panel:** the current record's features (device match, geo consistency, provenance, payee/device/IP concentration, account velocity, amount spike, deepfake and AI-agent indicators) with per-feature trust contributions. **Pipeline panel:** the eight governance stages with per-stage pass/hold/block status. **Authority panel:** the HMAA tier (T3/T2/T1/T0) and the resulting decision. **Metrics panel:** live triage counts and the precision, recall, and F1 with Wilson 95% confidence intervals. **Audit ledger:** the running SHA-256 chain head and entry count. **Swarm Review panel:** the retrospective ensemble result and escalation-delta.

4.2 Pipeline Stages

The eight-stage pipeline processes every transaction in order, with an audit append after the terminal decision:

```
VALIDATE → SATA → ADARA → MAIVA → HMAA → FLAME → ERAM → CARA  (→ AUDIT ledger append)
```

Stage	Function
VALIDATE	Schema and range validation; malformed records are isolated as invalid (DATA_FAULT) and never scored as benign
SATA	Source trust allocation from device match, geo consistency, and provenance
ADARA	Adversarial indicator scoring: deepfake artifacts and AI-agent patterns
MAIVA	Suspicion composition from adversarial, trust, coordination, and novelty signals
HMAA	Four-tier authority decision (T3 autonomous-clear, T2 supervised-review, T1 elevated-confirmation, T0 manual-hold)
FLAME	Cascade prevention and bounded deliberation window
ERAM	Independent risk escalation; can force elevated confirmation on high consequence
CARA	Graduated recovery and tamper-evident SHA-256 audit-chain entry

> **TIP.** Run scenario E (Low-and-slow ring) first as a normal stream, then click Swarm Review: the per-transaction path clears the ring legs, and the retrospective ensemble surfaces them as an escalation-delta. This is the clearest illustration of why the retrospective layer exists.

5. Operating Procedures

5.1 Startup

1. Open blade-finance-simulation.html locally, or navigate to it from burakoktenli.com/blade-finance-simulation.
2. Confirm all panels render. The status line should report the v2.2 build ready.
3. The default scenario initializes; the pipeline is idle until Run or Step is used.

5.2 Standard Operation

1. Select a scenario. Read the Scenario Briefing to understand the expected behavior.
2. Click Run to stream, or Step to advance one transaction and inspect each stage.
3. Observe the eight-stage pipeline and the HMAA tier routing each record to cleared, review, or blocked; invalid records isolate at VALIDATE.
4. Watch the metrics panel update precision, recall, F1, and the triage split with Wilson confidence intervals.
5. Click Monte Carlo for a multi-trial run with confidence intervals; click Self-Test for the deterministic suite.
6. For coordinated-ring scenarios, click Swarm Review to run the retrospective ensemble and read the escalation-delta.
7. Use the fault-injection toggles (Section 6) to stress the pipeline and observe fail-closed behavior.
8. Click Export JSON or Golden to save the session, metrics, and audit trail.

5.3 Shutdown

1. Export the session if you want to keep it.
2. Close the browser tab; all state is discarded.

> **IMPORTANT.** State is not persisted. Export before closing to preserve results.

6. Controls and Fault Injection

Control	Function
Run / Step / Pause / Reset	Stream, single-step, pause, or reset the transaction stream and metrics
Self-Test	Runs the deterministic test suite (expect 9 of 9 pass), including audit-chain re-walk
Monte Carlo	Multi-trial run reporting triage rates and detection metrics with Wilson 95% intervals
Swarm Review	Retrospective ensemble over the recent ledger window; reports rings and escalation-delta
Golden / Export JSON	Export the golden trace (fixed seed) or the full session and audit trail as JSON
Load Data / Sample Data / Replay	Ingest an external labeled benchmark, load the bundled sample, or replay an exported trace
Bypass Gov	Disables the governance gate to establish the no-governance baseline (governance lift)
Corrupt In	Injects malformed or out-of-range fields; expect VALIDATE to isolate them as invalid
Sensor Drop / Comms Lag / Op Error	Injects a dropped feature, communication latency, or an operator error
Byz Models	Enables Byzantine agent models in the retrospective swarm-review ensemble

7. Scenario Reference

ID	Name	Description	Expected Behavior
A	Nominal	Benign transaction stream	Cleared at T3; metrics show high clear rate, zero breaches
B	Deepfake authentication	Synthetic voice or video defeats step-up auth	ADARA deepfake artifacts plus SATA provenance gap; HMAA holds for review or block
C	AI-agent coordinated	Many individually-plausible transactions across shared payee, device, and IP	Coordination score rises; MAIVA elevates suspicion; review or block
D	High-value risk	Large transfer with partial anomaly	ERAM forces elevated confirmation even where HMAA alone would clear
MIX	Realistic mixed	Blended benign and adversarial stream (the reported metrics scenario)	Cleared 86.8%, review 11.9%, blocked or invalid 1.4%; zero FLAME breaches
E	Low-and-slow ring	Sub-threshold legs sharing a mule payee and botnet IP	Per-transaction path clears the legs; Swarm Review surfaces the ring as an escalation-delta

8. Metrics, Formulas, and Verification

8.1 Governance Scoring (reference equations)

The decision math is intentionally simple and auditable. All terms are clamped to [0,1].

```
sataTrust = clamp(0.40·deviceMatch + 0.30·geoConsistency + 0.30·provenance)
```

```
coordinationScore = clamp(0.28·payeeConc + 0.26·deviceConc + 0.24·ipBurst + 0.12·acctVelocity + 0.10·amountSpike)
```

```
pAttack = 1 - (1 - deepfakeArtifacts)(1 - agentPattern)(1 - coordinationScore)
```

```
adv = clamp(0.78·pAttack + 0.14·(1 - sataTrust) + 0.08·coordinationScore)
```

```
suspicion = clamp(0.42·adv + 0.25·(1 - sataTrust) + 0.23·coordinationScore + 0.10·novelPayeeFlag)
```

8.2 Key Metrics

- **Triage split:** cleared, review, blocked, invalid (per transaction).
- **Precision / recall / F1:** over the actionable-risk triage decision, reported with Wilson 95% confidence intervals. Recall here means actionable-risk triage coverage, not an empirical fraud-detection rate.
- **FLAME breaches:** deliberation-window violations; the governed pipeline target is zero.
- **Governance lift:** compare metrics with Bypass Gov off versus on.
- **Swarm Review:** ensemble agreement and ring detection over a look-back window, reported as an escalation-delta. This is ensemble agreement plus ring detection, not a Byzantine-fault-tolerance guarantee.

8.3 Verification Checklist

Action	Expected Result
Click Self-Test	9 of 9 pass, including audit-chain re-walk
Run scenario A (Nominal)	High clear rate at T3; zero FLAME breaches
Run scenario MIX, governance on	Cleared 86.8%, review 11.9%, blocked or invalid 1.4%; precision 0.917, recall 1.000, F1 0.957
Toggle Corrupt In	Malformed records isolate as invalid at VALIDATE, not scored as benign
Toggle Bypass Gov	Detection collapses relative to the governed run (governance lift baseline)
Run scenario E, then Swarm Review	Ring surfaced; escalation-delta lists actionable transactions among the cleared
Export and re-run with the same seed	Bit-identical outputs and audit head

9. Data Export and Reproducibility

Click Export JSON to save parameters, the metrics summary, the transaction history, and the audit trail. Click Golden to export the fixed-seed golden trace. To verify: export, note the seed, reload with the same seed and scenario, and confirm a bit-identical result and audit head.

9.1 Reproducibility Guarantee

Property	Value
PRNG	mulberry32 (32-bit seeded); zero Math.random() calls on computation paths
Audit chain	SHA-256 over a canonical serialization, built-in JS implementation, byte-exact against FIPS 180-4
Golden trace	Scenario MIX, seed 42, 1,000 records
Deterministic head hash	`7d5aaab4...`
Running checksum	`0x06acd6be`
Confidence intervals	Wilson score intervals over the trial count (2,000 for the reported run)
Cross-browser / platform	Deterministic across Chrome, Firefox, Safari, Edge on Windows, macOS, Linux

> **TIP.** Any divergence from head hash `7d5aaab4...` or checksum `0x06acd6be` at seed 42 / MIX / 1,000 records indicates a modified engine or a non-conforming SHA-256 implementation.

10. Limitations and Threat Considerations

Limitation	Description
Simulation-only evidence	Browser-based decision-logic only; no captured transaction data and no hardware measurements (TRL 3-4)
Synthetic data	All features and labels are synthetic research parameters, not empirically calibrated
Triage, not fraud detection	Reported recall is an actionable-risk triage measure; absolute rates illustrate architecture, not field performance
Swarm review scope	Ensemble agreement and ring detection; no quorum-intersection or Byzantine-fault-tolerance safety bound is claimed
Software cryptography	SHA-256 is a JS implementation; per-record ECDSA and HSM custody are design-specified, not exercised here
Ephemeral state	All state is in memory; closing the tab discards it

10.1 Threat Considerations

- **Local modification:** a user can alter the code via developer tools; verify exported results against the published source and the golden-trace head hash.
- **Browser extensions:** a malicious extension could modify the DOM or state; run in a clean or private window for trustworthy results.
- **No supply-chain surface:** because the file is self-contained with no CDN or external script, there is no remote-code-injection vector to mitigate.

11. Troubleshooting

Problem	Likely Cause	Solution
Blank screen after open	JavaScript disabled or a render error	Enable JavaScript; open the console (F12) for details; try a clean browser profile
Self-Test does not show 9 of 9	Modified engine or non-conforming environment	Re-download the published source; confirm the golden-trace head hash
Runs slowly on large Monte Carlo	Many concurrent tabs	Close other tabs; reduce the trial count
Export button does nothing	Downloads blocked	Allow downloads for the page in browser settings
Controls unresponsive	Tab lost focus	Click inside the simulator window and ensure the tab is active

12. Glossary and References

12.1 Glossary

Term	Definition
AUTHREX	Authority Regulation and Execution, the governance framework root
VALIDATE	Schema and range validation; isolates malformed input as invalid (DATA_FAULT)
SATA	Source trust allocation from device, geo, and provenance evidence
ADARA	Adversarial indicator scoring (deepfake artifacts, AI-agent patterns)
MAIVA	Multi-attribute suspicion composition gating the authority decision
HMAA	Four-tier authority model: T3 autonomous-clear, T2 supervised-review, T1 elevated-confirmation, T0 manual-hold
FLAME	Cascade prevention with a bounded deliberation window
ERAM	Independent risk escalation gate
CARA	Graduated recovery and tamper-evident audit-chain entry
Escalation-delta	Actionable transactions surfaced by the retrospective swarm review among those the per-transaction path cleared
mulberry32	Seeded 32-bit pseudo-random generator used for reproducibility

12.2 References

1. Oktenli, B. (2026). BLADE-FINANCE Governance Node (v1.0). Zenodo. DOI 10.5281/zenodo.20374692.
2. Simulation artifact: burakoktenli.com/blade-finance-simulation
3. Project page: burakoktenli.com/blade-finance
4. U.S. Department of the Treasury (2026). Financial Services AI Risk Management Framework.
5. NIST (2023). AI Risk Management Framework (AI RMF 1.0).

12.3 Contact

Burak Oktenli, Georgetown University, M.P.S. Applied Intelligence. Website burakoktenli.com, ORCID 0009-0001-8573-1667. For questions about the simulator or the governance-architecture research program, use the contact form at burakoktenli.com.

End of Document