BLADE-AGENT-HSM Emulator

Simulation User Guide

Version: v2.6
Document: AGENT-HSM-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
Related Artifact: DOI: 10.5281/zenodo.20299821

Document Control

Document ID	AGENT-HSM-SIM-UG-001
Simulation Version	v2.6
Author	Burak Oktenli — Georgetown University — ORCID 0009-0001-8573-1667
Date / Revision	May 2026 / Rev 1.0 — Initial Release
Classification	UNCONTROLLED — Research Artifact
Related Artifact	DOI: 10.5281/zenodo.20299821

Table of Contents

1. Purpose, Scope, and Assumptions

This guide provides operating procedures for the BLADE-AGENT-HSM emulator, an adversarial high-assurance browser emulator for the hardware root of trust that anchors autonomous AI agents. BLADE-AGENT-HSM is the hardware companion to the AUTHREX-AGENT software shim. The emulator reproduces the device behavior with real WebCrypto primitives: a five-opcode 64-byte ABI (audit_sign, pcr_extend, pcr_quote, tool_auth, spawn_quorum_sign), TPM-style PCR measurement chains, non-exportable-key audit signing, four-tier authority (T3 green / T2 amber / T1 red / T0 lock), HKDF per-tool tokens, spawn-quorum aggregation, and a multi-modal tamper cascade that zeroizes keys and latches the device to T0.

Intended Audience: EB2-NIW petition evaluators, defense/aerospace and AI-security reviewers, academic peers, and technical collaborators seeking independent verification of behavior.

Scope: Operation of the BLADE-AGENT-HSM emulator. Does not cover the reference hardware design (see the Zenodo deposit, the Interface Control Document ICD-AGENT-HSM-001, and the bill of materials) or the cryptographic theory (see the companion paper).

1.1 Assumptions and Constraints

- The user has a modern browser with JavaScript enabled and HTTPS access.
- The emulator runs real WebCrypto primitives (ECDSA P-256/P-384, SHA-256, HKDF); silicon timing is modeled, not measured.
- All parameter values are synthetic research placeholders, not calibrated against fabricated silicon.
- No certified hardware exists; the secure element, TPM, and tamper sensors are modeled in software.
- Post-quantum (ML-DSA) fields model interface shape only.

> **IMPORTANT**
>
> This emulator is a research demonstrator (TRL 2-3 silicon / 3-4 emulator). No certified hardware exists. No FIPS, Common Criteria, EAL, NSA, NASA, or DoD endorsement, validation, or certification of any kind is claimed. All parameters are synthetic.

2. Quick Start

Step 1. Open the emulator HTML file in Chrome over HTTPS.

Step 2. Review the panels: ABI console, PCR bank, authority-tier indicator, audit ledger, and tamper status.

Step 3. Select a failure-mode scenario (or normal flow) and start the emulator.

Step 4. Issue ABI opcodes (audit_sign, pcr_extend, pcr_quote, tool_auth, spawn_quorum_sign) and watch the PCRs and ledger.

Step 5. Run the self-test, then export the signed evidence bundle for verification.

> **NOTE**
>
> All computation runs client-side with the WebCrypto API. No data leaves your browser. Requires HTTPS (not file://).

3. System Requirements and Security Considerations

Requirement	Specification
Browser	Chrome 90+, Firefox 88+, Safari 15+, Edge 90+
Protocol	HTTPS required (WebCrypto API for ECDSA / SHA-256 / HKDF)
Display	Min 1280x720; recommended 1920x1080+
CPU/Memory	Any modern processor; ~200MB RAM peak for batch runs
Network	Internet for initial CDN load. All computation client-side after load.
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 with WebCrypto. No data is sent to any server.
- CDN dependencies load from cdnjs.cloudflare.com with Subresource Integrity hashes where available.
- Audit integrity: per-entry ECDSA P-256 signatures, a hash-chained ledger, and a P-384 signed golden-trace anchor.
- Trust-root caveat: the anchor public key travels inside the evidence file, so adversarial forgery resistance requires an out-of-band identity pin; absent that pin, the verifier reports integrity-and-same-session trust only.

4. Interface Layout and Navigation

4.1 Panel Layout

The interface presents the five-opcode ABI console, the TPM PCR bank (PCR0 tier, PCR1 ledger, PCR2 tool policy, PCR3 spawn, PCR4 tamper cause, PCR5-7 reserved), the four-tier authority indicator (T3 green / T2 amber / T1 red / T0 blinking red + alarm), the hash-chained audit ledger, the per-tool HKDF token panel, the spawn-quorum aggregator, and the tamper-status block (active mesh, voltage-glitch, thermal). Controls include scenario selection, opcode issue, self-test, and evidence-bundle export.

> **TIP**
>
> Hover over interface elements for tooltips. Gauges include ARIA labels for screen-reader accessibility.

5. Operating Procedures

5.1 Startup

1. Navigate to the simulation URL or click Launch Simulation from burakoktenli.com.
2. Wait for loading (2-5 seconds). CDN scripts load from cdnjs.cloudflare.com.
3. Verify the interface loads completely. All panels should be visible.

5.2 Standard Operation

1. Open the emulator via Launch Simulation or navigate to blade-agent-hsm-simulation.html.

2. Issue audit_sign to sign a 32-byte ledger hash with the modeled SE051 ECDSA P-256 key; the entry chains into PCR1.

3. Issue pcr_extend / pcr_quote to extend and attest the PCR bank; PCR0 holds the active authority tier.

4. Issue tool_auth to HKDF-derive a per-tool token bound to the active tier; above-tier requests are refused.

5. Issue spawn_quorum_sign to verify N-of-M sub-agent signatures and aggregate them via the secure element.

6. Run a failure-mode scenario: voltage glitch, electronic-warfare hazard with recovery, split-brain ledger reconciliation, or Byzantine fault isolation; observe the tamper cascade zeroize keys and latch T0.

7. Run the software-only-versus-HSM baseline to quantify what the hardware anchor adds.

8. Run the in-browser self-test (275 deterministic checks across seven batteries).

9. Export the signed evidence bundle (event count, final PCR digests, trace SHA-256, P-384 anchor).

5.3 Shutdown

1. Export session data. 2. Close the browser tab (all state discarded).

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

6. Parameter Reference

Parameters are pre-configured per scenario (tier ceilings, tool policy, quorum N-of-M, tamper thresholds). See Section 4 for interface controls and the companion paper and ICD-AGENT-HSM-001 for the full ABI frame layout and PCR allocation.

7. Scenario Reference

Normal flow plus adversarial failure-mode scenarios. Select a scenario before starting.

Scenario	Description
Normal flow	Baseline lifecycle: tiered actions, signed ledger entries, and PCR attestation.
Voltage glitch	Out-of-window rail excursion trips the cascade; keys zeroize and the device latches T0.
EW hazard with recovery	Electronic-warfare disturbance and the modeled recovery path.
Split-brain ledger	Divergent ledger states reconciled against the hash chain and anchor.
Byzantine fault isolation	A misbehaving sub-agent isolated during spawn-quorum aggregation.
Anchor re-key attack	Attempted forgery via anchor re-key, rejected when an out-of-band identity pin is set.

8. Metrics, Formulas, and Verification

8.1 Key Metrics

Metric	Definition
ABI Opcodes	audit_sign, pcr_extend, pcr_quote, tool_auth, spawn_quorum_sign (64-byte frame).
PCR Bank	PCR0 tier, PCR1 ledger, PCR2 tool policy, PCR3 spawn, PCR4 tamper cause, PCR5-7 reserved.
Authority Tier	T3 autonomous / T2 supervised / T1 confirmed / T0 halt-and-lock (keys zeroized).
Audit Signature	Per-entry ECDSA P-256; hash-chained; P-384 signed golden-trace anchor.
Validation	275 deterministic checks across seven batteries (275/275); software-only-vs-HSM baseline.

8.2 Verification Checklist

Action	Expected Result
Start emulator / self-test	Self-test passes 275/275 deterministic checks.
Issue audit_sign	Entry signed with a non-exportable key; PCR1 extends; signature verifies.
Request above-tier tool_auth	Token derivation is refused; the action is blocked.
Trigger a tamper scenario	Keys zeroize; PCR4 records the cause; the device latches T0; only pcr_quote is served.
Export evidence bundle	Bundle downloads; the verifier re-derives PCRs and validates every signature.
Truncate the trace	Verification fails (the P-384 anchor detects truncation/substitution).

9. Data Export and Reproducibility

Click export to save the signed evidence bundle: the deterministic golden trace, the event count, the final PCR digests, the trace SHA-256, and the P-384 signed anchor. The bundled verifier re-derives the PCR chain, checks the ECDSA signature on every entry, and validates the anchor; truncation or substitution fails.

9.1 Reproducibility Guarantee

Property	Value
Crypto	Real WebCrypto: ECDSA P-256/P-384, SHA-256, HKDF (SubtleCrypto API)
Determinism	Deterministic golden trace; seeded generator; zero Math.random() in computation paths
Anchor	P-384 signed over event count + final PCR digests + trace SHA-256
Test batteries	Seven batteries, 275 checks (275/275), confirmed over three reruns
Baseline	Software-only-versus-HSM comparison quantifies the hardware-anchor contribution

10. Limitations and Threat Considerations

Limitation	Description
Emulator-Only Evidence	Browser-based emulation. No fabricated silicon and no hardware measurements.
Modeled Silicon Timing	Timing is modeled, not measured; performance figures are design targets.
Modeled Secure Hardware	The secure element, TPM, and tamper sensors are modeled in software.
Interface-Only PQC	Post-quantum (ML-DSA) fields model interface shape only.
Single-Session State	All state is held in memory; closing the tab discards all data.

10.1 Threat Considerations

• CDN compromise: dependencies load from cdnjs.cloudflare.com. Mitigation: Subresource Integrity hashes where available.

• Browser extensions could modify the emulator DOM/state. Mitigation: test in Incognito mode for clean results.

• Anchor-key trust: the evidence anchor public key travels in the bundle; pin the device attestation identity out-of-band, otherwise the verifier reports integrity-and-same-session trust only.

ASSURANCE BOUNDARY

This is a research demonstrator (TRL 2-3 silicon / 3-4 emulator). No certified hardware exists. No FIPS, Common Criteria, EAL, NSA, NASA, or DoD endorsement, validation, or certification of any kind is claimed. Silicon timing is modeled, not measured; post-quantum fields model interface shape only; the audit anchor requires an out-of-band identity pin for adversarial forgery resistance.

11. Troubleshooting

Problem	Likely Cause	Solution
Black screen after loading	Render error or CSP violation	Open F12 Console for details. Try Chrome Incognito.
Self-test fails	WebCrypto unavailable (file://)	Serve over HTTPS; WebCrypto requires a secure context.
Controls not responding	Tab lost focus	Click inside the emulator window. Ensure the tab is active.
Export not working	Download blocked	Allow downloads from the simulation domain in browser settings.
Loading never completes	CDN scripts blocked	Disable ad blockers. Allow cdnjs.cloudflare.com.

12. Glossary and References

12.1 Glossary

Term	Definition
BLADE-AGENT-HSM	Beam-Layer Authority for Directed Engagements, Agent Hardware Security Module
AUTHREX-AGENT	The software authority-lifecycle shim that this device anchors in hardware
ABI	Application Binary Interface — the fixed five-opcode 64-byte command frame
PCR	Platform Configuration Register — TPM measurement register (PCR0-7)
HKDF	HMAC-based Key Derivation Function — per-tool, tier-bound token derivation
Secure Element	CC EAL6+ device (modeled on NXP EdgeLock SE051) holding non-exportable keys
Tamper Cascade	Mesh/voltage/thermal event -> key zeroize -> PCR4 cause -> T0 lock

12.2 References

[1] Oktenli, B. (2026). BLADE-AGENT-HSM: A Reference Hardware-Root-of-Trust Design and Verified Emulator for Agentic-AI Authority Governance. DOI: 10.5281/zenodo.20299821.

[2] Oktenli, B. (2026). Emulator artifact. https://burakoktenli.com/blade-agent-hsm-simulation.html

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

[4] CISA, NSA AI Security Center, and Five Eyes (2026). Careful Adoption of Agentic AI Services (1 May 2026).

[5] FY26 NDAA Sections 1513 and 6601; NIST SP 800-53 Rev. 5; FIPS 140-2/140-3; TCG TPM 2.0 Library Specification.

12.3 Contact

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

End of Document