CYBER-PHYSICS

Deterministic Operational Governance Architecture

Technical White Paper
Enterprise Evaluation Edition

Runtime:	Cyber-Physics Enterprise Runtime v2.0.0-ee
Validation Program:	Enterprise Engineering Validation Program
Document Status:	Technical Reference
Edition:	Enterprise Evaluation Edition
IP Status:	PCT/IB2026/053487

This document describes the architecture and observed engineering evidence supported by the supplied Evaluation Kit. It is not a certification, regulatory approval, or third-party endorsement.

Contents

1 Executive Summary

Cyber-Physics is a deterministic operational-governance architecture intended to provide a runtime governance layer for continuously processed telemetry and associated operational evidence.

The current enterprise evaluation artifact is the **Cyber-Physics Enterprise Runtime v2.0.0-ee**. The supplied runtime is evaluated as an immutable OCI runtime under controlled conditions. The evaluation scope focuses on deterministic governance processing, runtime stability, evidence integrity, replay behaviour, and recovery behaviour.

The associated **Enterprise Engineering Validation Program** was designed to provide evidence suitable for subsequent institutional technical evaluation. The supplied Evaluation Kit contains the runtime artifact, evaluation documentation, engineering reports, and raw evidence generated during the documented validation phases.

The current release is therefore an **evaluation baseline**, not a claim of universal production certification.

2 Architectural Position

2.1 Operational Role

The current EEVP defines Cyber-Physics as a deterministic governance layer whose runtime behaviour can be observed independently of a complete enterprise security-product deployment.

The Evaluation Kit separates:

1. validation of the Cyber-Physics runtime itself;
2. validation of a particular telemetry integration; and
3. validation of a complete enterprise deployment architecture.

These are distinct engineering questions. The baseline Evaluation Kit is primarily concerned with the first.

2.2 What the Evaluation Runtime Is Not

The Evaluation Kit does not, by itself, constitute a complete out-of-the-box integration with every SIEM, SOAR, EDR, firewall, cloud platform, or telemetry system. Deployment-specific integration may require adapters, normalization, APIs, message-bus integration, organizational access controls, monitoring, and other engineering.

The supplied evidence does not establish multi-node clustering, active-active or active-passive high availability, geographic redundancy, multi-region deployment, Kubernetes production orchestration, or universal enterprise-scale horizontal scaling.

3 Deterministic Runtime Model

3.1 Runtime Processing

The EEVP methodology defines a processed telemetry event as an individual **pulse**. During stress validation, latency is measured for each individual processed pulse.

For each processed telemetry event, the validation instrumentation records:

- processing latency;
- governance state;
- CFL metric;

- thermal pressure; and
- DGEC evidence.

3.2 Governance State Model

The governance model uses the following state progression:

$$\text{OBSERVE} \rightarrow \text{ALERT} \rightarrow \text{GOVERNANCE} \rightarrow \text{STABILIZED}.$$

3.3 Runtime Boundaries

The supplied Evaluation Kit defines a controlled runtime boundary including:

- OCI container deployment;
- CPU limit of 1 core for the evaluation runtime;
- memory limit of 512 MB for the evaluation runtime;
- read-only runtime filesystem as specified by the installation documentation;
- persistent evidence associated with the DGEC ledger.

These boundaries describe the supplied evaluation configuration and are not universal requirements for every future deployment.

4 Deterministic Governance Evidence Chain

The **Deterministic Governance Evidence Chain** is the cryptographic evidence mechanism used by the runtime to preserve continuity between governance events.

The EEVP methodology validates DGEC integrity by comparing the current `last_calculated_hash` with the subsequent event's `previous_hash` using SHA-256:

$$H_{n+1}^{\text{previous}} = H_n^{\text{calculated}}.$$

The acceptance criterion requires uninterrupted hash-chain continuity across the applicable validation cycle.

4.1 Sequence Observability

The field `DGEC_Last_Sequence` is treated by the current methodology as an instrumentation-observed sequence field. During the recorded endurance evidence, the external instrumentation probe reported the value 0. The methodology classifies this as the genesis/initial snapshot value exposed to the probe rather than an independent measurement of live DGEC progression.

Accordingly, sequence observability is not used as the primary DGEC integrity criterion. Cryptographic hash-chain continuity remains the primary integrity test.

5 Replay Verification

Replay validation provides an independent check of recorded governance evidence for the scenarios defined by the EEVP methodology.

Historical DGEC records are replayed through the CyberPhysicsKernel. For the validated replay scenarios, the following outputs are required to match:

- governance decision;
- CFL metric;
- governance state; and
- operational outcome.

The acceptance condition is:

$$\text{deterministic output equivalence} = 100\%$$

for the validated replay scenarios.

This is deliberately bounded: replay equivalence is evidence about the documented scenarios and methodology, not a universal guarantee covering every possible external dependency.

6 Enterprise Engineering Validation Program

The Enterprise Engineering Validation Program exercises the supplied OCI runtime through isolated external validation probes.

The documented coverage includes:

1. validation readiness;
2. baseline benchmarking;
3. runtime continuous operation;
4. stress testing;
5. engineering chaos and recovery;
6. endurance and operational stability;
7. final benchmarking; and
8. enterprise-readiness assessment.

The program evaluates:

- deterministic governance decisions;
- runtime continuity;
- memory, thread, and file-descriptor behaviour;
- DGEC integrity;
- replay determinism;
- recovery after defined forced failures; and
- resource behaviour under the documented workload.

6.1 Acceptance Criteria

Metric	Success Standard
RSS memory growth	Less than 1% growth over the observed endurance plateau
Thread leakage	Zero orphaned or leaked runtime threads
File descriptor leakage	Zero leaked file descriptors after continuous operation
DGEC hash integrity	100% uninterrupted hash-chain continuity
Replay recovery matching	100% deterministic replay equivalence for validated scenarios
Runtime availability	At least 99.9% over the 105-hour endurance phase, excluding planned chaos injection
Self-healing	Recovery of governance sequence and DGEC hash chain after defined forced failures
Kernel latency	Average latency below 2.0 ms under the defined normal enterprise workload

The acceptance documentation distinguishes the defined long-term target window of 30 days from the currently supplied evidence, which includes a completed 105-hour continuous endurance run.

7 Measurement Methodology

7.1 Endurance Sampling

During endurance validation, resource metrics are sampled every 60 seconds: RSS memory, CPU utilization, thread count, file descriptor count, and socket availability.

7.2 CPU Measurement

CPU utilization is measured for the Cyber-Physics Runtime process using the defined process-level instrumentation:

`cpu=process.cpu_percent(interval=0.1).`

The methodology distinguishes:

- measurement interval: 0.1 seconds;
- evidence sampling cadence: 60 seconds;
- measurement target: the Cyber-Physics Runtime process.

The 60-second cadence determines when observations are recorded; it does not mean CPU is averaged continuously over the entire 60-second period.

7.3 RSS Measurement

RSS is measured through container metrics or process instrumentation. The endurance evidence reports a baseline of 172.41 MB and a final recorded value of 172.73 MB:

$$\frac{172.73 - 172.41}{172.41} \times 100 = 0.1856\% \approx 0.19\%.$$

This is endpoint-to-baseline growth, not peak RSS excursion.

7.4 Kernel Latency

Kernel execution latency is measured with `time.perf_counter()` from immediately before processing `x_features` until governance completion. The documented measurement excludes external I/O operations and the final DGEC SHA-256 append.

8 Observed Engineering Evidence

The final engineering scorecard records:

Indicator	Recorded Observation
Memory stability	PASS; peak 173.03 MB
CPU stability	PASS; average 2.1%
Observed memory drift indicator	0.00% in the scorecard
Replay integrity	100% for the recorded scorecard set
DGEC integrity	100% for the recorded scorecard set
Hash continuity	0 recorded breaks in the scorecard
Chaos recovery	PASS; 1533.54 ms recovery value in the scorecard
Average latency	11.42 ms in the scorecard
Worst latency under stress	362.44 ms (P99) in the scorecard
Operational availability	100% continuous in the scorecard
Replay executions	105 cycles
Replay failures	0
Hash verification	105 / 105

Interpretation boundary. These are recorded engineering observations from the supplied evidence and are not universal guarantees.

The acceptance criteria define an average kernel-latency threshold below 2.0 ms for the defined normal enterprise workload, while the final scorecard separately reports an overall average latency of 11.42 ms and a stress P99 of 362.44 ms. This White Paper therefore does not claim that those scorecard figures establish universal compliance with every latency acceptance condition; the applicable workload, measurement method, and raw evidence must be examined by an evaluator.

9 Chaos Engineering and Recovery

The chaos phase evaluates runtime behaviour under controlled failure scenarios, including forced process termination, memory-related failure conditions, socket disruption, restart and state loading, power-failure simulation, replay recovery, and mixed-chaos scenarios.

The final Phase 4 assessment reports:

- observed logical recovery rate: 100% for validated recovery scenarios;
- observed system panic/deadlock rate: 0% within validated scenarios;
- observed time-to-first-decision after crash: less than 1500 ms.

These findings are limited to the documented scenarios and evidence set. Sequence-observability findings are distinguished from cryptographic DGEC integrity. The chaos results are not exhaustive coverage of all possible failure modes.

10 Endurance and Operational Stability

The endurance assessment covers 105 hours of continuous runtime. The reported RSS endpoint growth is approximately 0.19%, calculated from the documented baseline and final endpoint values. The assessment also reports DGEC integrity, replay determinism, and governance stability for the evaluated run.

These results are evidence from the validated environment and workload. They do not establish identical resource consumption or performance for every deployment.

11 Evaluation Package and Reproducibility

The Enterprise Evaluation Kit is organized into:

```
01_Runtime/
02_Documentation/
03_Reports/
04_Evidence/
05_Support/
```

The runtime is delivered as an immutable OCI artifact for evaluation. The release documentation instructs evaluators not to modify, rebuild, patch, or otherwise alter the supplied runtime during the initial evaluation, because doing so creates a separate evaluation configuration.

The evidence package contains raw validation evidence, engineering assessments, endurance records, stress records, chaos records, and validation artifacts required to examine the documented results.

Before execution, package integrity is verified against the supplied SHA-256 manifest.

12 Security and Enterprise Boundaries

The Evaluation Kit is intentionally bounded.

It is not an identity and access-management platform and does not replace organizational IAM, SSO, MFA, RBAC, LDAP, SAML, or OIDC controls.

It is not a complete production operations plan. Enterprise deployment may require monitoring, alerting, backup, disaster recovery, incident response, change management, configuration management, hardening, patch management, and lifecycle management.

It is not a universal security certification. The evaluator remains responsible for applying its own security, compliance, risk, and deployment acceptance criteria.

13 Known Limitations

The current evidence establishes behaviour only within the documented evaluation scope. Important limitations include:

- no complete validation of every enterprise telemetry integration;
- no complete validation of multi-node or geographic high availability;
- no complete validation of Kubernetes production orchestration;
- no universal enterprise-scale horizontal-scaling claim;
- performance remains dependent on host architecture, contention, telemetry characteristics, event rate, topology, and interfaces;

- chaos scenarios are controlled and not exhaustive of all failure modes;
- replay determinism is bounded by validated replay scenarios and appropriate capture/normalization of external dependencies; and
- the Evaluation Kit is not a substitute for an organization's independent security assessment.

14 Independent Institutional Evaluation

The purpose of the current release is to provide a bounded, evidence-backed artifact suitable for independent institutional technical evaluation.

The intended next step is to enable qualified engineering, research, and security-architecture organizations to:

1. inspect the documented architecture and methodology;
2. verify the supplied evidence;
3. reproduce selected evaluation procedures;
4. challenge the stated acceptance boundaries;
5. run additional organization-specific tests; and
6. determine whether deeper evaluation is warranted.

The central distinction is:

$$\boxed{\text{Engineering evidence} \neq \text{Independent certification}}$$

15 Research and Intellectual Property Status

The project documentation identifies the intellectual-property status as:

PCT/IB2026/053487

The existence of an intellectual-property filing does not constitute independent validation of the technical claims in this document. Technical claims remain bounded by the supplied methodology, evidence, and evaluation scope.

16 Conclusion

Cyber-Physics Enterprise Runtime v2.0.0-ee is presented in this edition as a controlled deterministic operational-governance runtime accompanied by an engineering validation and evidence package.

The current release combines:

- a defined runtime governance model;
- a cryptographically linked evidence mechanism;
- replay validation for documented scenarios;
- controlled endurance, stress, and chaos evaluation;
- explicit measurement methodology; and
- explicit limitations and evaluation boundaries.

The resulting artifact is intended to support the next stage: **independent institutional technical evaluation**.

Any organization conducting such an evaluation may apply additional tests, alternative workloads, or stricter acceptance criteria. Such scrutiny is not outside the purpose of this document; it is the intended next step.

A Source Alignment

This White Paper was drafted against the current Enterprise Evaluation documentation and evidence set, including:

- `02_Documentation/Acceptance_Criteria.md`
- `02_Documentation/Benchmark_Methodology.md`
- `02_Documentation/EEVP_Validation_Plan.md`
- `02_Documentation/Known_Limitations.md`
- `02_Documentation/Release_Notes_v2.0.0-ee.md`
- `03_Reports/EEVP_Readiness_Report.md`
- `03_Reports/Endurance_Assessment.md`
- `03_Reports/Final_Benchmark_Assessment.md`
- `03_Reports/Phase4_Final_Assessment.md`
- `03_Reports/Cyber-Physics_Engineering_Scorecard.pdf`
- `04_Evidence/` raw validation evidence.

Document Control Note

This is a new technical reference document for the current Enterprise Evaluation Edition. It is intentionally not a revision of the previous v3.6 document.