DOI https://doi.org/10.32689/maup.it.2026.2.6
UDC 004.056.5:004.72
JEL Classification: C63, L86, O33

A COMPROMISE-TOLERANT CROSS-LAYER RESILIENCE CONTROLLER ARCHITECTURE AND AN EQUAL-BUDGET BENCHMARK FOR PROGRAMMABLE NETWORKS

Dmytro Prokopovych-Tkachenko

АРХІТЕКТУРА МІЖРІВНЕВОГО КОНТРОЛЕРА КІБЕРСТІЙКОСТІ, СТІЙКОГО ДО КОМПРОМЕТАЦІЇ, ТА ПОРІВНЯЛЬНЕ ОЦІНЮВАННЯ ЗА РІВНОГО РЕСУРСНОГО БЮДЖЕТУ ДЛЯ ПРОГРАМОВАНИХ МЕРЕЖ

Д. І. Прокопович-Ткаченко

Abstract
This study defines a compromise-tolerant cross-layer resilience controller for programmable networks, combining a programmable data plane, federated control, provenance-aware state estimation, a digital twin, a formal policy shield, bounded local fallback, and deterministic reconciliation. B0-B8 were evaluated in a synthetic flow-level event-driven model under equal 16-vCPU, 32-GiB, and 1-Gbit/s budgets, nine scenarios, six scales, and 50 paired seeds. At 10^4, B7 produced mission loss of 0,061 versus 0,155 for B6, recovery of 21,6 versus 30,9 s, and goodput of 78,0% versus 69,4%. The mean paired relative estimator was 57,9% (95% bootstrap interval, 57,0-58,9%), whereas the ratio of displayed aggregate means was 60,6%; these answer different questions. Under controller compromise, B7 retained Qmin=0,668. The 250-ms p99 limit was first exceeded at 25,000. Conclusions are bounded by the synthetic model, parameterization, resource accounting, scenarios, and trust assumptions.
Key words: *programmable networks, cyber resilience, controller compromise, federated control, digital twin, equal-budget benchmark.*

Анотація
Запропоновано референтну архітектуру міжрівневого контролера кіберстійкості для програмованих мереж. Вона поєднує програмований рівень даних, федеративне керування, оцінювання стану з урахуванням походження доказів, цифровий двійник, формальний захисний контур, локальне резервування та детерміноване узгодження. B0-B8 оцінено у синтетичній потоковій подійній моделі за однакового бюджету 16 vCPU, 32 GiB і 1 Gbit/s, дев'яти сценаріїв, шести масштабів і 50 парних початкових значень. За 10^4 B7 отримала місійні втрати 0,061 проти 0,155 для B6, час відновлення 21,6 проти 30,9 с і goodput 78,0% проти 69,4%. Середній парний відносний ефект становив 57,9% (95% бутстреп-інтервал 57,0-58,9%), тоді як відношення агрегованих середніх дорівнює 60,6%; це різні оцінки. За компрометації контролера Qmin B7 становила 0,668. Межу p99=250 мс уперше перевищено при 25 000. Висновки обмежені синтетичною моделлю, параметризацією, ресурсним обліком, сценаріями та довірчими припущеннями.
Ключові слова: *програмовані мережі, кіберстійкість, компрометація контролера, федеративне керування, цифровий двійник, рівний ресурсний бюджет.*

1. Introduction. Software-defined networking (SDN) separates decision logic from packet forwarding, making behaviour programmable, observable, and rapidly reconfigurable [1]. P4, eBPF/XDP, and network function virtualization extend enforcement to switches, hosts, gateways, and service chains [2-4]. These mechanisms coordinate detection, routing, isolation, rekeying, and degraded operation, but they also concentrate authority in controllers, telemetry services, and policy stores.

Replication improves availability but does not ensure that replicated decisions are trustworthy or safe. Controller compromise, forged messages, stale state, and malicious applications remain systemic SDN risks [5]; fault-tolerant execution addresses consistency but not the provenance and authority of a decision [6]. Digital twins can forecast cross-domain effects [7,8], and time-dependent service curves capture resilience better than static availability [9]. Secure placement and integrity auditing add reusable mechanisms [10, 11], yet prior work does not specify how service should continue when the controller, evidence stream, or policy bundle is partly compromised.

RQ1 asks whether a reference controller can preserve bounded mission service during controller partition or compromise without producing unauthorized cross-layer actions. RQ2 asks whether provenance verification, prediction, formal shielding, local fallback, and reconciliation remain useful when every configuration receives the same compute, memory, bandwidth, workload, traces, seeds, and time. The contribution is a reference architecture with four operational planes plus two cross-cutting rails, a compromise-tolerant federation protocol, and a B0-B8 equal-budget benchmark. Claims are limited to the controlled synthetic model.

The aim of this study is to specify a compromise-tolerant cross-layer resilience controller architecture for programmable networks and evaluate its ability to preserve bounded mission service during controller partition or compromise while respecting authorization and safety constraints within an equal-budget synthetic benchmark.

The research objectives are to define the controller's operational planes, trust boundaries, federation, local fallback, and reconciliation mechanisms; formalize provenance-aware state estimation and constrained action selection supported by a digital twin and a policy shield; establish a common maximum resource envelope,

© Prokopovych-Tkachenko D. I., 2026
Стаття поширюється на умовах ліцензії відкритого доступу CC BY 4.0

threat scenarios, outcome metrics, and paired comparison procedure for configurations B0-B8; and assess mission outcomes, safety, resource costs, and scaling limits, including the contributions of individual mechanisms through ablation analysis. These objectives link the architectural specification and evaluation methods to the reported results within the stated scope of the synthetic model.

2. Materials and methods

2.1. Architecture and trust boundaries. The design follows the anticipate-withstand-recover-adapt logic of systems-security engineering and explicit authorization from Zero Trust [12, 13]. Ethernet, managed Wi-Fi, and 5G are treated as distinct profiles; 5G roles, QUIC and Multipath TCP continuity, signed build provenance, survivable updates, and software bills of materials are composed under common evidence and authority contracts [14-19]. Requirements include plane separation, explicit trust boundaries, tolerance of one compromised controller replica or collector class, bounded local survival, deterministic reconciliation, safe reversible action, clean deployment, and equal resource budgets.

Table 1

Plane responsibilities and authority boundaries

Plane/rail	State and role	Authorized output	Forbidden shortcut
Data	Packets, queues, links; P4/ eBPF	Forwarding and local enforcement	Identity or global-policy authorization
Federated control	Topology, intents, chains, epochs	SDN/NFV actions; local fallback	Unauthenticated evidence; shield bypass
Evidence/state	Collectors, clocks, provenance, estimator	State estimate and uncertainty	Direct enforcement
Decision	Twin, optimizer, policy, shield	Approved reversible action or abstention	Direct forwarding writes
Identity/session	Identity, keys, tokens, attestation	Authorize, rekey, revoke, sign	Treat authentication as health
Ops.	SIEM/SOAR, NDR, SOC/NOC, audit	Case, approval, override, export	Invoke unregistered actuators

TB-1 separates decision generation from actuation: data-plane agents accept only signed, time-bounded commands within their domain. TB-2 separates evidence construction from policy selection: urgency cannot relabel stale, dependent, or inconsistent observations as trustworthy. The federated plane uses monotonic epochs, attested voting weights, signed health claims, and deterministic action identifiers. Quorum for domain d is valid only when attested weight reaches its threshold and participating epochs differ by at most one:

$$q_d(t)=1 \Leftrightarrow \sum_{i\in C_d} a_i(t)v_i \geq Q_d \wedge \left[\max_i e_i - \min_i e_i\right] \leq 1. \quad (1)$$

Here $a_i(t)$ denotes an eligible attested replica, v_i its voting weight, Q_d the threshold, and ei its epoch. Loss of quorum freezes policy expansion and activates pre-authorized local actions with finite time-to-live, resource caps, and rollback points. Divergent members are quarantined; rejoining domains accept newer signed epochs, reject stale actions, merge conflict-free effects, and retain unexplained effects for review.

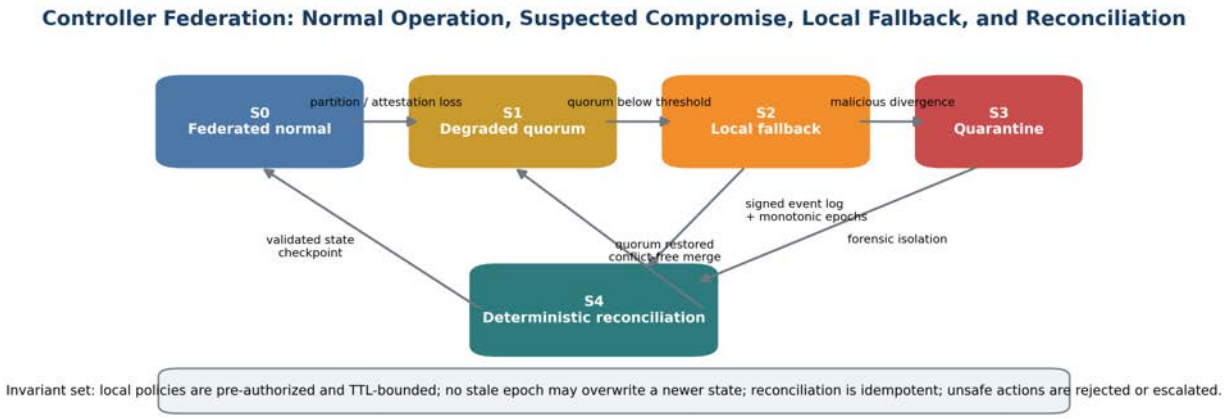


Fig. 1. Federation states and authority limits during partition, quarantine, local fallback, and reconciliation. Source: authors' model

ISSN 2786-5460 (Print), ISSN 2786-5479 (Online)

2.2. State estimation and safe action selection. Each evidence record identifies source, role, measured and received time, clock uncertainty, epoch, software and policy version, integrity proof, provenance path, semantic relation, and action outcome. A valid signature is not equated with physical truth: a signed report may still contradict independent route, queue, transport, or application behaviour. The robust estimate is

$$\hat{\mathbf{s}}(t) = \arg\min_{\mathbf{s}} \{\mathcal{L}_\rho(\mathbf{s}) + \beta \parallel \mathbf{L}_x \mathbf{s} \parallel_1 + \gamma \parallel \mathbf{s} - \mathbf{s}^- \parallel_2^2\}, \quad (2)$$

where the bounded residual limits extreme observations, L_x penalizes cross-layer inconsistency, and s− is the prior or digital-twin prediction. The estimator exposes uncertainty and may abstain when alignment, provenance, or model coverage is insufficient. Mission quality $Q_m(t)$ is normalized relative to the required service level, and the primary outcome is the weighted area of degradation:

$$L_M = \frac{1}{T} \int_0^T \sum_{m \in \mathcal{M}} w_m \left[1 - Q_m(t)\right] dt. \quad (3)$$

Candidate policies minimize expected mission loss plus decision-time, reconfiguration, overhead, and safety terms, subject to explicit safety, authorization, stability, and budget constraints:

$$\pi^\star = \arg\min_\pi \mathbb{E}\left[L_M(\pi) + \lambda_t J_t + \lambda_r J_r + \lambda_o J_o + \lambda_s J_s \right], \quad (4)$$

$$g_{\text{safe}}(\pi) \le 0, \; g_{\text{auth}}(\pi) = 1, \; g_{\text{stab}}(\pi) \le 0,$$
$$g_{\text{budget}}(\pi) \le 0. \quad (5)$$

The digital twin forecasts candidate effects but cannot authorize them. Immediately before enforcement, the shield checks authorization, capacity, stability, ownership, session continuity, blast radius, and rollback. If a candidate fails, the previous policy may be reused only if it also satisfies the current invariants; otherwise the controller abstains. Autonomous authority is limited to observation and reversible actions. Mission-significant or irreversible change requires separate approval. Covered critical services must satisfy the bounded-degradation condition

$$Q_m(t) \ge q_{\text{safe}} \mid \wedge \mid T_{\text{recover}}\left(Q_m \ge q_{\text{target}}\right) \le \text{RTO}_m. \quad (6)$$

Identity is evaluated for each subject-resource-action tuple. Replicas, collectors, policy services, operators, workloads, and data-plane agents have distinct least-privilege roles; short-lived credentials, attestation where available, and separation of duties prevent one credential from both creating evidence and approving a high-impact response. Rekeying and revocation are modeled as service transitions because abrupt revocation can destroy session continuity. SIEM/SOAR and network detection exchange normalized evidence and action requests through a broker, but cannot send arbitrary switch or kernel commands. Deployable units carry signed digests, dependency graphs, policy bundles, build attestations, and SBOM fields linked to source, toolchain, base image, and version [17-19].

2.3. Threat model and benchmark scenarios. The adversary may compromise one controller replica, one collector class, selected hosts or gateways, a route-advertising component, or a software/configuration artifact; delay, replay, reorder, or forge telemetry; inject plausible false reports; partition controllers; and coordinate effects across layers. Simultaneous compromise of all trust roots or the required federation threshold, ASIC faults, and destruction of every communication path are outside the model.

Table 2

Attack and failure catalog

ID	Scenario	Primary property
RP	Route poisoning/blackhole	Path trust; convergence; safe reroute
DF	DDoS plus flash crowd	Admission fairness; goodput; recovery
CP	Controller partition	Local fallback; reconciliation
CC	Controller compromise	Quarantine; quorum; effect detection
CD	Clock/telemetry deception	Alignment; provenance; robust fusion
TS	TCP/QUIC stress	Transport control; path diversity
KR	Key/session compromise	Rekey/revoke without service collapse
PC	Policy conflict/oscillation	Shield; hysteresis; no loops
SC	Supply-chain/config defect	Attestation; canary; rollback; SBOM

Seven hard invariants govern every run: I1 authorization; I2 epoch monotonicity; I3 bounded blast radius; I4 finite-TTL reversibility; I5 absence of recursive control loops; I6 scenario-specific bounded service degradation; and I7 linkage of evidence, estimate, forecast, alternatives, constraints, action, acknowledgement, effect, and rollback by action ID. A hard-invariant failure is reported separately and cannot be offset by a favourable average score.

2.4. Equal-budget design and statistics. Evaluation uses a controlled flow-level event-driven model, not vendor hardware or a production deployment. Each run lasts 120 s; the disturbance begins at 10 s. Nine scenarios use 50 paired seeds at 10^2, 10^3, 10^4, 2.5×10^4, 5×10^4, and 10^5 flows/devices. All configurations share the same maximum resource envelope:

$$\mathbf{b}\left(B_j\right)=\mathbf{b}_0=\left[16,32,1,120,D,T,S\right],\ \forall j\,. \tag{7}$$

The vector fixes 16 vCPU, 32 GiB RAM, a shared 1-Gbit/s management link, 120 s, identical demand D, traces T, and seeds S. Unused capacity is recorded; B7 cannot exceed the envelope. Every run starts from an empty state store with pinned manifests, fixed thresholds, a demand file, scenario file, and seed file; no learned state is reused.

Table 3

Equal-budget configurations

ID	Configuration	Capability and deliberate omission
B0	Standard stack	Native timeouts; no resilience controller
B1	Strong single layer	Scenario-specific IDS/ACL; no cross-layer state
B2	Zero Trust + SIEM/SOAR	Identity/playbooks; no twin or adaptive control
B3	Rule-based SDN/NFV	Static failover; centralized deterministic rules
B4	Routing/multipath	Path diversity; no session/app coordination
B5	Central optimizer	Complete telemetry; weak local fallback
B6	Distributed policies	Fast local survival; weak global coordination
B7	Full controller	Federation, provenance, twin, shield, fallback, reconcile
B8	Ablations	Separate variants omit provenance, twin, shield, or fallback

Clean deployment is part of the comparison. The harness verifies the signed manifest and SBOM, starts management services, instantiates the data-plane profile, waits for readiness, injects the fixed disturbance, archives raw logs and action records, and then destroys the environment. Empty state stores and pinned component digests prevent warm-state reuse or baseline-specific tuning. Public artifacts exclude real addresses, keys, production topologies, and sensitive detection rules.

Primary outcomes are normalized mission loss and bounded degradation; secondary outcomes include minimum quality, recovery, weighted goodput, tail latency, CPU, RAM, management traffic, rollback, and modeled unsafe actions. The predeclared scale knee is

$$N^{\star}=\min\left\{N:d_N>D_{\max}\left|\vee\right|\ c_N>U_{\max}\left|\vee\right|m_N>M_{\max}\right\}, \tag{8}$$

with p99 decision latency d_N, CPU c_N, RAM m_N, D_{max}=250 ms, U_{max}=85%, and M_{max}=85% of 32 GiB. Common random numbers pair workload and disturbance variation. Relative mission-loss effects are computed within each scenario-seed pair:

$$\delta_j=\frac{1}{|\mathcal{S}||\mathcal{K}|}\sum_{s,k}\frac{L_M\left(B_j,s,k\right)-L_M\left(B7,s,k\right)}{L_M\left(B_j,s,k\right)}. \tag{9}$$

The reported 95% interval uses 10,000 paired bootstrap resamples. Deadline failures, rollback failures, unsafe actions, and negative comparisons remain in the data. The interval quantifies repeated variability within the synthetic model, not uncertainty over all real networks; run-level reproduction should also examine scenario-stratified or hierarchical resampling.

3. Results and discussion.

3.1. Conformance and aggregate performance. The model instantiated four operational planes and the identity/session and operations rails. Each B7 actuation traversed an evidence gate, estimator, constrained decision, and shield. B7 passed the implemented requirements, and no modeled unsafe action occurred in 2,700 runs. This is internal conformance to the implemented specification, conditional on its invariants, roots of trust, event equations, and scenario coverage; it is not an independent security proof. The audit chain linked evidence, uncertainty, predictions, rejected alternatives, TTL, observed effect, and rollback.

ISSN 2786-5460 (Print), ISSN 2786-5479 (Online)

Table 4

Aggregate results at 10^4 flows/devices (nine scenarios; 50 paired seeds)

Cfg.	LM	Recovery, s	Goodput, %	Q_{min}	p99, ms	Rollback, %
B0	0.376	62.3	50.9	0.165	2374	87.1
B1	0.302	49.1	57.0	0.221	700	88.7
B2	0.280	50.8	59.1	0.274	1044	92.2
B3	0.256	46.9	61.2	0.323	233	93.8
B4	0.223	38.2	64.0	0.352	188	92.0
B5	0.233	49.7	63.6	0.416	339	96.0
B6	0.155	30.9	69.4	0.464	162	92.0
B7	0.061	21.6	78.0	0.686	228	98.9

At 10^4, B7 improved mission outcomes relative to B6 but used 63.1% CPU and 14.3 GiB RAM versus 41.1% and 9.2 GiB. The mean paired relative loss reduction was 57.9% (reported 95% bootstrap interval, 57.0–58.9%). In contrast, the ratio of displayed aggregate means, (0.155−0.061)/0.155, is 60.6%. The paired estimator gives equal weight to within-pair relative effects; the ratio compares two aggregate means. They answer different questions, and the interval is conditional on the implemented resampling and model.

3.2. Compromise, scale, and operational trade-offs. The largest modeled advantage occurred under controller compromise. B7 produced mission loss 0.073, Qmin=0.668, and recovery 27.0 s, versus 0.215, 0.328, and 37.5 s for B6 and 0.679, 0.106, and 127.3 s for B5. Under controller partition, B7 achieved loss 0.048, Qmin=0.717, recovery 19.2 s, and 100% bounded degradation; B6 achieved 0.093, 0.563, 22.9 s, and 98%. These differences match the modeled roles of quarantine, bounded local authority, and reconciliation. The advantage was smaller for transport stress, where established path mechanisms were already effective. B4 and B6 also used fewer resources and often decided faster, so the full stack is not optimal for every single-domain network.

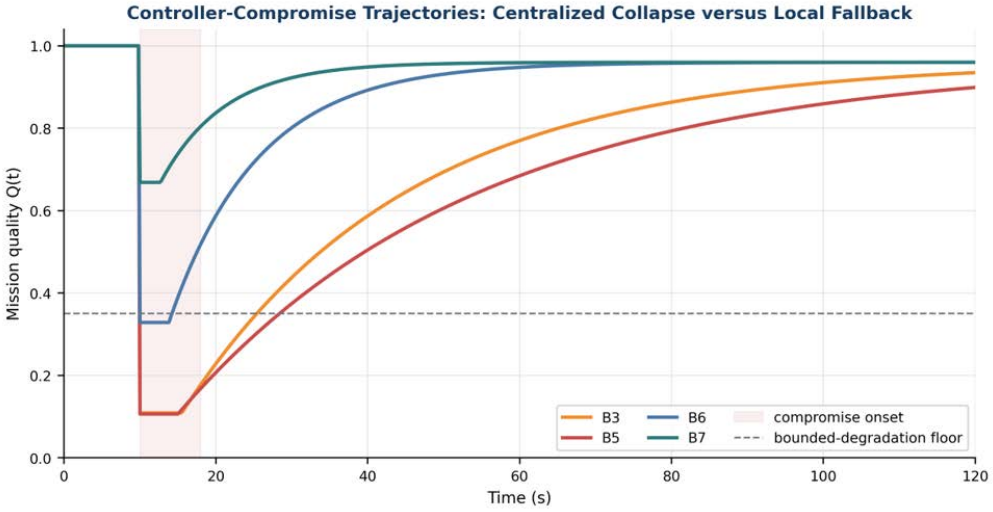


Fig. 2. Representative mission-quality trajectories under controller compromise at 10^4 flows/devices. Source: authors' model

Table 5

Controller partition and compromise at 10^4 flows/devices

Scenario	Cfg.	LM	Q_{min}	Recovery, s	Bounded, %
CP	B3	0.323	0.259	68.1	0
CP	B5	0.426	0.214	92.9	0
CP	B6	0.093	0.563	22.9	98
CP	B7	0.048	0.717	19.2	100
CC	B3	0.530	0.109	95.0	0
CC	B5	0.679	0.106	127.3	0
CC	B6	0.215	0.328	37.5	2
CC	B7	0.073	0.668	27.0	100

B7 stayed below the 250-ms p99 decision limit through 10^4. The first evaluated violation occurred at 25,000 (252 ms), followed by 278 ms at 50,000 and 358 ms at 100,000. CPU reached 85.7% at 100,000; RAM was 22.0 GiB. Combined control and telemetry traffic reached 205.8 Mbit/s, below the nominal 1-Gbit/s link. Tail decision latency, not bandwidth exhaustion, was the first binding constraint.

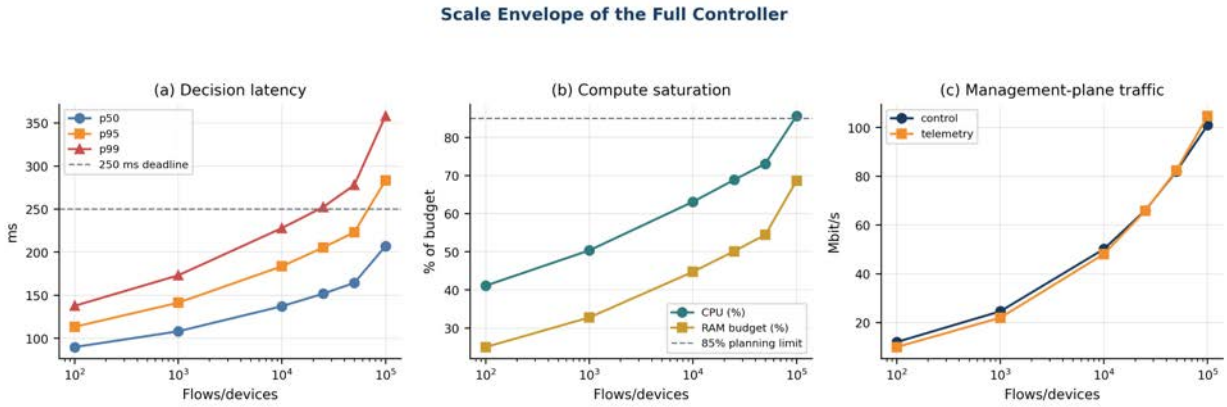


Fig. 3. B7 scale envelope; the p99 decision deadline is crossed before the common compute and link budgets are exhausted. Source: authors' model

Table 6

B7 scale results aggregated over nine scenarios

Flows/devices	p99, ms	CPU, %	RAM, GiB	Ctrl+tel, Mbit/s	LM
100	138	41.1	8.0	22.0	0.051
1,000	173	50.4	10.5	46.6	0.054
10,000	228	63.1	14.3	98.3	0.061
25,000	252	68.9	16.1	131.8	0.065
50,000	278	73.1	17.4	164.5	0.068
100,000	358	85.7	22.0	205.8	0.076

The knee is a reference envelope for the stated model and budget, not a universal capacity limit. A large deployment would need domain sharding, lower-cardinality telemetry, cached local policies, and global coordination reserved for cross-domain events. Faster hardware may move the knee upward; constrained gateways or dynamic radio networks may move it downward. At small and medium scale, B4 is rational for path-dominated risk and B6 for partition tolerance with safely pre-authorized local actions.

At 10^4, control and telemetry contributed 50.2 and 48.1 Mbit/s; at 10^5 they reached 101.0 and 104.8 Mbit/s. Although the combined load remained below 1 Gbit/s, consuming more than one fifth of a shared management path may compete with recovery traffic. Storage is limited through aggregation, retention classes, and signed manifests: public audit records retain hashes, summaries, decisions, and non-sensitive metrics, whereas restricted evidence follows the minimum approved retention period. Energy values are accounting-model outputs rather than physical measurements.

3.3. Ablations and validity. Ablations separated the mechanisms' roles. Removing provenance gates raised mission loss to 0.111. Removing the twin raised loss to 0.109 and recovery to 27.4 s while reducing p99 to 192 ms and CPU to 52.5%. Removing the shield admitted eight modeled unsafe actions in 450 runs and lowered rollback success to 87.8%. Removing local fallback raised loss to 0.167 and recovery to 45.1 s and reduced bounded degradation to 76.2%.

Table 7

B8 ablation results at 10^4 flows/devices

Variant	LM	Rec., s	p99, ms	Unsf.	Rollb., %	Bound., %
B7	0.061	21.6	228	0	98.9	100.0
B8-E	0.111	24.4	210	0	97.6	100.0
B8-T	0.109	27.4	192	0	99.1	100.0
B8-S	0.088	25.5	215	8	87.8	98.0
B8-F	0.167	45.1	241	0	98.9	76.2

ISSN 2786-5460 (Print), ISSN 2786-5479 (Online)

The components therefore trade distinct assurances: provenance supports evidence trust, the twin improves prediction at a latency cost, the shield protects action safety, and fallback preserves service under partition. The eight B8-S violations also show why average mission loss cannot absorb a hard safety outcome.

All quantitative findings are model-bounded. The study used synthetic traces, flow-level event equations, normalized service weights, and model-based resource and energy accounting; it did not execute hardware-in-the-loop or validate vendor P4/eBPF, SDN/NFV, Wi-Fi, or 5G systems. Results depend on weights, thresholds, recovery functions, and the assumption that independent trust roots and the federation threshold are not jointly compromised. The clean build demonstrates reproducibility only within the controlled bundle. Zero modeled unsafe actions for B7 does not prove the absence of unknown states. Independent code review, release of run-level data and seeds, hardware-in-the-loop validation, sector calibration, and external replication are required before operational claims.

Deployment should progress from observe-only monitoring to decision support, guarded reversible actions, federated fallback, and only then conditional cross-domain autonomy. Advancement requires measured evidence quality, stable invariants, rollback testing, and explicit authority. Shared-mode compromise of a trust root, a systematically wrong twin, or a mis-specified invariant can still defeat several planes; a schema migration that loses monotonic epochs or action IDs can also make reconciliation non-deterministic. These residual risks belong in the deployment profile rather than being hidden by the aggregate score.

4. Conclusions. A compromise-tolerant controller was specified around explicit trust boundaries, provenance-aware evidence, constrained decisions, bounded local fallback, and deterministic reconciliation. Under equal modeled budgets at 10^4, B7 improved mission loss, recovery, goodput, and minimum service quality relative to B6 while using more compute, memory, and decision time. The paired relative mission-loss estimator was 57.9%, whereas the ratio of displayed means was 60.6%; both are reported because they summarize the data differently.

The benchmark also identified limits: p99 first exceeded 250 ms at 25,000, and simpler baselines remained preferable in some single-domain or latency-sensitive conditions. Ablations exposed distinct safety, prediction, provenance, and local-survivability trade-offs within the synthetic model. The contribution is a bounded reference architecture and evaluation method, not proof of universal field effectiveness. Future work should publish a frozen reproducibility package, use hierarchical inference, validate real controller and data-plane implementations, perform hardware-in-the-loop tests, and obtain independent replication.

Conflict of Interest. The author declares no conflict of interest.

Funding. This research received no external funding. Institutional Review Board Statement: not applicable; the synthetic benchmark involved neither human participants nor animals. Informed Consent Statement: not applicable.

Data Availability. Run-level data, scripts, configurations, seeds, parameter files, and manifests are maintained in a controlled working package. No public repository or DOI is claimed. Public release requires security review and removal of sensitive deployment information; until release, independent numerical reproduction remains limited.

Use of Artificial Intelligence. OpenAI GPT-5.6 Pro was used for language drafting, document structuring, and generation of controlled modeling and visualization scripts. The author reviewed and edited the outputs and accepts full responsibility. AI assistance does not replace independent code review, invariant testing, or external validation. Conflicts of Interest: the author declares no conflicts of interest.

Author Contributions. Conceptualization, methodology, software, validation, formal analysis, investigation, data curation, writing-original draft, writing-review and editing, visualization, and project administration: Dmytro Prokopovych-Tkachenko. The author approved the manuscript and accepts responsibility for its content.

References:

1. Kreutz, D., Ramos, F.M.V., Verissimo, P.E., Rothenberg, C.E., Azodolmolky, S., Uhlig, S. (2015). Software-Defined Networking: A Comprehensive Survey. Proc. *IEEE*. 103, 14–76. https://doi.org/10.1109/JPROC.2014.2371999.

2. Bosshart, P., Daly, D., Gibb, G., Izzard, M., McKeown, N., Rexford, J., Schlesinger, C., Talayco, D., Vahdat, A., Varghese, G., Walker, D. (2014). P4: Programming Protocol-Independent Packet Processors. ACM SIGCOMM Comput. Commun. Rev. 44, 87–95. https://doi.org/10.1145/2656877.2656890.

3. Vieira, M.A.M., Castanho, M.S., Pacifico, R.D.G., Santos, E.R.S., Camargos, E.P.M., Jr., Vieira, L.F.M. (2020). Fast Packet Processing with eBPF and XDP: Concepts, Code, Challenges, and Applications. *ACM Comput. Surv*. 53, 1–36. https://doi.org/10.1145/3371038.

4. Mijumbi, R., Serrat, J., Gorricho, J.-L., Bouten, N., De Turck, F., Boutaba, R. (2016). Network Function Virtualization: State-of-the-Art and Research Challenges. *IEEE Commun. Surv. Tutor*. 18, 236–262. https://doi.org/10.1109/COMST.2015.2477041.

5. Scott-Hayward, S., Natarajan, S., Sezer, S. (2016). A Survey of Security in Software Defined Networks. *IEEE Commun. Surv. Tutor*. 18, 623–654. https://doi.org/10.1109/COMST.2015.2453114.

6. Katta, N.P., Zhang, H., Freedman, M.J., Rexford, J. Ravana: Controller Fault-Tolerance in Software-Defined Networking. In Proceedings of the 1st ACM SIGCOMM Symposium on Software Defined Networking Research, Santa Clara, CA, USA, 17–18 June 2015; Article 4, pp. 1–12. https://doi.org/10.1145/2774993.2774996.

7. Fett, M., Kraft, M., Wilking, F., Goetz, S., Wartzack, S., Kirchner, E. (2024). Medium-Level Architectures for Digital Twins: Bridging Conceptual Reference Architectures to Practical Implementation in Cloud, Edge and Cloud-Edge Deployments. *Electronics*, 13, 1373. https://doi.org/10.3390/electronics13071373.

8. El-Hajj, M. (2024). Leveraging Digital Twins and Intrusion Detection Systems for Enhanced Security in IoT-Based Smart City Infrastructures. *Electronics,* 13, 3941. https://doi.org/10.3390/electronics13193941.

9. Cho, H., Sung, J.-H., Kang, H.-J., Jang, J., Shin, D. (2025). Quantifying Cyber Resilience: A Framework Based on Availability Metrics and AUC-Based Normalization. *Electronics*, 14, 2465. https://doi.org/10.3390/electronics14122465.

10. Bocci, A., Forti, S., Ferrari, G.-L., Brogi, A. (2023). Declarative Secure Placement of FaaS Orchestrations in the Cloud-Edge Continuum. *Electronics*, 12, 1332. https://doi.org/10.3390/electronics12061332.

11. Yang, F., Sun, Y., Gao, Q., Chen, X. (2024). EDI-C: Reputation-Model-Based Collaborative Audit Scheme for Edge Data Integrity. *Electronics*, 13, 75. https://doi.org/10.3390/electronics13010075.

12. Ross, R., Pillitteri, V., Graubart, R., Bodeau, D., McQuaid, R. (2021). NIST SP 800-160 Vol. 2 Rev. 1: Developing Cyber-Resilient Systems: A Systems Security Engineering Approach. National Institute of Standards and Technology: Gaithersburg, MD, USA, https://doi.org/10.6028/NIST.SP.800-160v2r1.

13. Rose, S., Borchert, O., Mitchell, S., Connelly, S. (2020). NIST SP 800-207: Zero Trust Architecture. National Institute of Standards and Technology: Gaithersburg, MD, USA, https://doi.org/10.6028/NIST.SP.800-207.

14. 3GPP TS 33.501; Security Architecture and Procedures for 5G System. 3rd Generation Partnership Project: Sophia Antipolis, France, Release 18.

15. Iyengar, J., Thomson, M. (2021). RFC 9000: QUIC: A UDP-Based Multiplexed and Secure Transport. Internet Engineering Task Force: Fremont, CA, USA, https://doi.org/10.17487/RFC9000.

16. Ford, A., Raiciu, C., Handley, M., Bonaventure, O., Paasch, C. (2020). RFC 8684: TCP Extensions for Multipath Operation with Multiple Addresses. Internet Engineering Task Force: Fremont, CA, USA, https://doi.org/10.17487/RFC8684.

17. Torres-Arias, S., Afzali, H., Kuppusamy, T.K., Curtmola, R., Cappos, J. in-toto: Providing Farm-to-Table Guarantees for Bits and Bytes. In Proceedings of the 28th USENIX Security Symposium, Santa Clara, CA, USA, 14-16 August 2019; pp. 1393–1410. https://doi.org/10.5555/3361338.3361435.

18. Samuel, J., Mathewson, N., Cappos, J., Dingledine, R. Survivable Key Compromise in Software Update Systems. In Proceedings of the 17th ACM Conference on Computer and Communications Security, Chicago, IL, USA, 4-8 October 2010; pp. 61–72. https://doi.org/10.1145/1866307.1866315.

19. National Telecommunications and Information Administration. The Minimum Elements for a Software Bill of Materials (SBOM). U.S. Department of Commerce: Washington, DC, USA, 2021. Available online: https://www.ntia.gov/report/2021/minimum-elements-software-bill-materials-sbom (accessed on 17 August 2026).

Відомості про авторів

Англ.	Укр.
Dmytro Prokopovych-Tkachenko Candidate of Technical Sciences, Associate Professor, Head of the Department of Cybersecurity and Information Technologies, University of Customs and Finance, Doctoral Candidate, State University of Information and Communication Technologies omega2417@umsf.edu.ua ORCID: 0000-0002-6590-3898 Scopus Author ID: 57201777932 Web of Science ResearcherID: Q-2918-2018	Прокопович-Ткаченко Дмитро Ігорович кандидат технічних наук, доцент, завідувач кафедри кібербезпеки та інформаційних технологій, Університет митної справи та фінансів, докторант, Державний університет інформаційно-комунікаційних технологій omega2417@umsf.edu.ua ORCID: 0000-0002-6590-3898 Scopus Author ID: 57201777932 Web of Science ResearcherID: Q-2918-2018

Дата надходження статті: 15.05.2026
Дата надходження виправленої версії статті: 30.06.2026
Дата прийняття статті: 07.07.2026
Дата публікації статті: 26.10.2026