NTP Working Group
Internet-Draft
Intended status: Informational
Expires: 1 December 2026

L. Melegassi
Catellix

30 May 2026

Cross-Vantage Clock-Offset Coherence Bounds for NTP-Disciplined Measurement Vantages
draft-melegassi-ntp-mvps-clock-coherence-00

Abstract

The Network Time Protocol version 4 (NTPv4) [RFC5905] specifies how a host disciplines its clock to a reference time scale; Network Time Security [RFC8915] and the Message Authentication Code [RFC8573] authenticate the client-server exchange; and the NTP Best Current Practices [RFC8633] direct operators to "monitor their NTP instances to detect attacks" (Section 5.3) without specifying a quantitative, cross-host monitoring procedure. The Security Requirements document [RFC7384] establishes that an on-path adversary can impose a clock offset (Sections 3.2.2, 3.2.3, 3.2.6) and that a single client cannot always detect such an offset by itself.

This document makes ONE contribution and proves it: given two or more measurement vantages disciplined to a common reference and each declaring an NTP-tier offset bound, a deterministic cross-vantage detector exists that (a) NEVER fires on offsets that are legitimate under the [RFC5905] / [RFC8633] synchronization envelope, and (b) is GUARANTEED to fire on an injected single-clock offset above a closed-form threshold. Both properties are theorems with elementary proofs; no statistical assumption, no protocol change, and no claim about the NTP wire format are required. The detector is the cross-vantage clock-skew axis of the Multi-Vantage Path Snapshot (MVPS) framework [I-D.melegassi-ippm-mvps-bundle]; this document isolates and proves the part that is purely a consequence of [RFC5905]’s error envelope.

A second result governs what happens AFTER detection, when the environment itself collapses (vantages go dark, telemetry thins). We prove (i) that the false-positive-free property survives any telemetry collapse with two or more surviving vantages, and (ii) a data-processing ceiling: an AI/LLM analysis layer riding the gated signal cannot recover information the surviving vantages did not observe. The LLM’s role is therefore provably EXPLANATION of an already-detected collapse, never detection itself; its operating envelope (decision tiers, classification accuracy) is given honestly as a stated model, not a theorem.

A third result governs SPEED. Driving the same cross-vantage comparison at a Bidirectional Forwarding Detection (BFD, [RFC5880]) cadence instead of the legacy 60-second coherence tick gives a closed-form detection-latency window (the L_DL lemma) whose worst case equals the BFD detection time plus one signalling delay. Because the false-positive-free property (Theorem 1) is independent of the sampling rate, the gate may run at the fastest BFD cadence (multiplier 1) WITHOUT trading away its zero-false-alarm guarantee, detecting an injected offset in tens of milliseconds rather than tens of seconds.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Copyright Notice

Table of Contents

1. Introduction and Scope
2. Terminology and Model
3. Hypotheses (the conditions the theorems depend on)
4. Lemma 1: Legitimate Cross-Vantage Span Bound
5. Theorem 1: A False-Positive-Free Threshold Exists and 2*tau Is the Smallest One
6. Theorem 2: Guaranteed Detection and the Minimum Detectable Offset
7. Corollary: Why a Single Host Cannot Self-Detect (and Two Can)
8. Post-Detection AI Layer over the Collapsed Environment
    8.1. Theorem 3: Collapse-Robustness of the FP-Free Gate
    8.2. Theorem 4: The Data-Processing Ceiling on the AI Layer
    8.3. AI Decision Envelope (MODEL, not a theorem)
9. Detection Latency: Binding the Gate to BFD Timing (RFC 5880)
    9.1. Theorem 5: Closed-Form Detection-Latency Window (L_DL)
    9.2. Corollary: 1091x Dwell Reduction and M=1 Optimality
10. Mapping to RFC 8633 Section 5.3 and RFC 7384
11. What This Document Does NOT Claim
12. Empirical Confirmation (and the corner it does not exercise)
13. Refinements the Theorems Suggest (constructive, non-normative)
14. Security Considerations
15. IANA Considerations
16. References
Appendix A. Worked Numbers per NTP Tier
Appendix B. Detection-Latency Variants (L_DL receipt)

1. Introduction and Scope

[RFC5905] disciplines one host's clock. [RFC8633] Section 5.3 asks operators to monitor for attack signatures and gives qualitative ones (bogus packet, zero-origin packet, bad MAC); a quantitative, cross-host agreement test is left, appropriately, to implementation. [RFC7384] Section 3.2 catalogues the offset-inducing attacks (spoofing, replay, delay manipulation) and observes that a delay attack in particular cannot be defeated by cryptography alone and

benefits from path redundancy. This document takes that observation as its starting point and supplies one such quantitative test, with proofs, as a complement to -- never a replacement for -- the existing NTP work.

The gap is therefore precise and acknowledged by the IETF: there is no standardized way to verify, from outside a host, that several NTP-disciplined hosts AGREE on the time to within what their declared stratum permits. This document does not propose to fill that gap by changing NTP. It proves that the agreement test is a one-line inequality on published offsets, and that with the correct threshold the test is provably free of false alarms against the [RFC5905] envelope while provably catching any large enough injected offset.

This is deliberately the SMALLEST provable statement. Everything that is not a theorem is moved to Section 9 ("What This Document Does NOT Claim").

2. Terminology and Model

The key words "MUST", "MUST NOT", "SHOULD", "MAY" are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals.

Vantage: an independent host running clock-disciplined measurement, synchronized to a common reference time scale per [RFC5905].

N: the number of vantages, N >= 2.

epsilon_i: the true offset of vantage i's clock relative to the common reference at the measurement instant (epsilon_i > 0 means the clock is ahead).

tau_i: a declared upper bound on |epsilon_i|, taken from the vantage's NTP tier (its stratum / synchronization class). tau_i is an operator-supplied calibration input, NOT a number normatively fixed by [RFC5905]; representative values are tabulated in Appendix A.

tau: the worst (largest) declared bound in the bundle, tau := max_i tau_i. Mixed-tier bundles are bound by their loosest clock.

o_i: the offset value vantage i PUBLISHES (its NTP-reported clock offset, e.g. the offset computed from the [RFC5905] Section 8 timestamp quadruple).

S: the cross-vantage span, S := max_i o_i - min_i o_i.

D_theta: the detector that raises a flag if and only if S > theta, for a fixed threshold theta >= 0.

3. Hypotheses (the conditions the theorems depend on)

The theorems in Sections 4-7 are CONDITIONAL on the following. Each is stated with a falsification path so a reviewer can break it.

H1 Common reference. All N vantages discipline to the same reference time scale (a common stratum-1/-2 source or an equivalent ensemble) for the duration of one measurement window. Falsification: per-vantage "refid" / source log diverges.

H2 Honest offset publication under H0. Under the null hypothesis

(no attack), each vantage publishes o_i = epsilon_i with |epsilon_i| <= tau_i. Estimator noise of the NTP offset computation is folded into tau_i (i.e. tau_i is taken large enough to also cover the measurement error of the offset itself). Falsification: a calibrated clock whose published offset exceeds its declared tier bound in the absence of attack.

H3 Cardinality. N >= 2. (N >= 3 is required by the surrounding MVPS axioms for Byzantine tolerance, but the two theorems of this document hold for N >= 2.) Falsification: trivial.

These three hypotheses are the ENTIRE basis. No distributional assumption on epsilon_i is made.

4. Lemma 1: Legitimate Cross-Vantage Span Bound

STATEMENT. Under H1-H2 (no attack), the cross-vantage span satisfies

$$S = \max_i o_i - \min_i o_i \le 2\tau .$$

PROOF. By H2, o_i = epsilon_i and |epsilon_i| <= tau_i <= tau for every i, hence o_i in [-tau, +tau]. The span of a finite set contained in an interval of width 2*tau is at most 2*tau:

$$\max_i o_i - \min_i o_i \le (+\tau) - (-\tau) = 2\tau . \quad \text{QED}$$

REMARK. This is exactly the "2*eps_NTP" term that already appears in the joint-skew bound of the MVPS BBF-mesh profile ([I-D.melegassi-ganascim-mvps-bbf-mesh] Theorem T-MESH-1), there written "maximum pairwise skew <= 2*eps_NTP + tau_RTT_max". Lemma 1 is the propagation-free specialization.

5. Theorem 1: A False-Positive-Free Threshold Exists and 2*tau Is the Smallest One

STATEMENT. Let D_theta flag iff S > theta.

(a) If theta >= 2*tau, then under H0 (H1-H2 hold, no attack) the detector NEVER flags: the false-positive probability against the [RFC5905]/[RFC8633] synchronization envelope is exactly zero, deterministically.

(b) theta = 2*tau is the SMALLEST such threshold: for any theta < 2*tau there exists a legitimate H0 configuration on which D_theta flags.

PROOF. (a) By Lemma 1, S <= 2*tau <= theta under H0, so the event S > theta cannot occur. No probability is involved; the bound is a set inclusion.

(b) Take N = 2 with epsilon_1 = +tau, epsilon_2 = -tau, both within tier (admissible under H2). Then S = 2*tau. For any theta < 2*tau, S = 2*tau > theta, so D_theta flags on a fully legitimate configuration. Hence no threshold below 2*tau is false-positive free. QED

COROLLARY 1.1. The unique smallest false-positive-free detector is D_{2*tau}. Operators SHOULD set theta = 2*tau, where tau is the worst declared tier bound in the bundle.

6. Theorem 2: Guaranteed Detection and the Minimum Detectable Offset

We now model an on-path adversary consistent with [RFC7384]

Section 3.2: by packet manipulation (3.2.1), spoofing (3.2.2), replay (3.2.3) or delay manipulation (3.2.6) the adversary imposes a single additive offset Delta > 0 on exactly one vantage k, so that vantage k publishes o_k = epsilon_k + Delta, while the other N-1 vantages remain within tier (|epsilon_i| <= tau, i != k).

STATEMENT.

(i) [Worst case over legitimate placements.] The infimum of the post-attack span over all admissible legitimate offsets is

$$\inf S' = \max(0, \Delta - 2\tau) .$$

Consequently D_theta is GUARANTEED to flag (for every admissible placement of the honest clocks) if and only if

$$\Delta > \theta + 2\tau .$$

(ii) [Well-synchronized baseline.] If the honest vantages are tightly synchronized (epsilon_i = 0 for i != k, and epsilon_k = 0 before the attack), then S' = Delta and D_theta flags iff Delta > theta.

Hence, writing Delta_min for the smallest reliably detected offset,

$$\theta < \Delta_{min} \leq \theta + 2\tau ,$$

collapsing to Delta_min = theta in the tightly-synchronized baseline and to the worst-case guarantee Delta_min = theta + 2*tau in general. With the recommended theta = 2*tau (Corollary 1.1):

$$\Delta_{min} = 2\tau \text{ (baseline) to } 4\tau \text{ (worst case).}$$

PROOF of (i). To minimize the post-attack span, the honest clocks and epsilon_k are chosen adversarially. Cluster the N-1 honest points at a single value a in [-tau, tau] and write the attacked point as b + Delta with b = epsilon_k in [-tau, tau]. The two distinct points are {a, b + Delta}, so

$$S'(a,b) = | a - (b + \Delta) | = | (a - b) - \Delta | .$$

Over a, b in [-tau, tau] the difference a - b ranges over [-2*tau, 2*tau]. Thus (a - b) - Delta ranges over [-2*tau - Delta, 2*tau - Delta], and

$$\inf_{a,b} |(a - b) - \Delta| = \begin{cases} 0, & \text{if } \Delta \leq 2\tau \text{ (0 is attainable),} \\ \Delta - 2\tau, & \text{if } \Delta > 2\tau \text{ (interval lies > 0).} \end{cases}$$

i.e. inf S' = max(0, Delta - 2*tau). The detector is guaranteed to flag for ALL placements iff this infimum exceeds theta, i.e. iff Delta - 2*tau > theta. QED

PROOF of (ii). With every honest offset 0 and epsilon_k = 0, the point set after attack is {0 (x (N-1)), Delta}, so S' = Delta and S' > theta iff Delta > theta. QED

REMARK (interpretation). The "2*tau gap" between the baseline and the worst case is not slack to be engineered away: it is exactly the region in which a genuine within-tier skew of the honest clocks is indistinguishable from a small injected offset. This indistinguishability is the same fact proven in Theorem 1(b); it is a property of the [RFC5905] envelope, not a deficiency of the detector.

7. Corollary: Why a Single Host Cannot Self-Detect (and Two Can)

STATEMENT. A single client with one time source cannot, from its own observations alone, detect a consistent symmetric offset attack of magnitude Delta. Two vantages on distinct paths can, per Theorem 2.

ARGUMENT. [RFC5905] Section 8 computes a client's offset from the timestamp quadruple (T1, T2, T3, T4) as theta_hat = ((T2 - T1) + (T3 - T4)) / 2. A symmetric delay attack that adds d to both directions, or a server-time shift of Delta, leaves the client's internal consistency checks satisfied: there is no second, independent observable against which T-quadruple can be contradicted. [RFC7384] Section 3.2.6 states this directly -- delay attacks "cannot be prevented by cryptographic means" and mitigation requires redundant, diverse paths. The offset is therefore UNOBSERVABLE to a lone client.

With N >= 2 vantages on distinct paths and a common reference, the attacked vantage's published offset diverges from the others, and Theorem 2 converts that divergence into a deterministic detection guarantee. This is the precise, provable sense in which a multi-vantage construction adds detection power that no single [RFC5905] client possesses.

This document claims nothing stronger: it does not prevent the attack, does not authenticate the exchange (that is [RFC8915] / [RFC8573]), and does not identify WHICH vantage is wrong without the N >= 3 Byzantine machinery of the surrounding MVPS axioms.

8. Post-Detection AI Layer over the Collapsed Environment

The deterministic detector of Sections 4-7 answers one question: "do the vantages still agree on time within their envelope?" When the answer is no, the environment is, in the operational sense, COLLAPSING: clocks diverge, vantages may be going dark, telemetry thins. The MVPS framework places an AI/LLM analysis layer on top of that signal [I-D.melegassi-mvps-ai-coherence]. This section states precisely -- and proves where it can -- what that layer can and cannot do. The scope of every claim is tagged ANALYTICAL (a theorem), MODEL (a stated operating model), or CONJECTURE (open).

Collapse model. Let each of the N vantages independently still report in the current window with probability h in (0, 1] (the "environment health"; 1-h is the fraction gone dark through link failure, blackhole, or noise). Let M be the number of survivors.

8.1. Theorem 3: Collapse-Robustness of the FP-Free Gate [ANALYTICAL]

STATEMENT.

(a) For EVERY realization with M >= 2 survivors, the false-positive-free property of Theorem 1 holds verbatim with theta = 2*tau. No degree of telemetry collapse can manufacture a false alarm.

(b) If an offset Delta > theta is injected on one vantage k (baseline corner, Theorem 2(ii)), the probability the surviving bundle still catches it is, exactly,

$$P_{detect}(h) = h * (1 - (1 - h)^{(N-1)}) .$$

PROOF. (a) Theorem 1(a) is a statement about the M surviving published offsets only; its proof (Lemma 1) used N nowhere. Restrict the index set to the survivors: S_survivors <= 2*tau <= theta still holds. Hence no false positive, for any M >= 2.

(b) The injected offset is observable only if vantage k itself survives (probability h) AND at least one other vantage survives to

form a span (probability 1 - (1-h)^(N-1), by independence). The two events are independent, giving the product. QED

READING. Detection POWER degrades as the environment collapses, but the ZERO-false-alarm guarantee does not: the gate fails safe. An aggregate (rather than pairwise) coherent statistic degrades on the smooth A*sqrt(h) slope rather than a cliff [I-D.melegassi-mvps-ai-coherence]; that aggregate refinement is tagged MODEL there and is not needed for (a)-(b) here.

8.2. Theorem 4: The Data-Processing Ceiling on the AI Layer [ANALYTICAL]

Let x be the (hidden) true state of the collapsed environment, y the published multi-vantage observations (offsets, spans, hop data) on which the gate fired, and g(y) any AI/LLM analysis of y -- a classification, an explanation, a remediation hint. Because the LLM sees only y, the variables form a Markov chain

$$x \rightarrow y \rightarrow g(y) .$$

STATEMENT. $I(x ; g(y)) \leq I(x ; y)$.

PROOF. Direct application of the data-processing inequality [Cover-Thomas] to the chain x -> y -> g(y). QED

CONSEQUENCE (the honest division of labour). No AI or LLM post-processing can recover information about the collapsed environment that the surviving vantages did not capture. Therefore:

- DETECTION is the job of the deterministic gate (Theorems 1-3), whose guarantees are exact and adversary-independent.
- The AI/LLM layer's job is EXPLANATION within the information y already contains: naming the likely failure cause, ranking hypotheses, drafting an operator-readable account of the collapse. It provably cannot substitute for a missing vantage.

This is why adding the LLM does not weaken any guarantee in this document: it operates strictly downstream of, and bounded by, the gated signal. It also tells operators where the real lever is -- not a better model, but more/better-placed vantages (the Layer-3 program of [I-D.melegassi-mvps-ai-coherence], out of scope here).

8.3. AI Decision Envelope (MODEL, not a theorem)

On top of the proven gate, the framework reports an AI decision tier as a function of the detection power p [MVPS-AI-ENVELOPE]:

tier	condition	meaning
PERFECT	p >= 0.90	full-confidence decision
OPTIMAL	0.70 <= p<0.90	AI compensates; high confidence
GOOD	0.55 <= p<0.70	AI still decides above legacy floor
COLLAPSE	p < 0.55	degraded toward chance

For a coherent effect spread across N = 32 vantages with the A*sqrt(h) model, the tier holds at PERFECT/OPTIMAL/GOOD down to h = 0.5 (half the telemetry lost) while a single-vantage monitor of the same spread effect never leaves COLLAPSE -- the "AI prevails where the environment collapses" envelope.

SCOPE. The tier thresholds (0.90 / 0.70 / 0.55) are DESIGN CHOICES, not theorems; the A*sqrt(h) degradation is a stated MODEL; embedding this envelope in a captured real attack is a CONJECTURE pending the

live lab. Any classification accuracy figure (e.g. the diagonal-Gaussian failure-cause classifier reported elsewhere at macro-F1 ~ 0.72) is EMPIRICAL with a declared methodological semi-circularity and is explicitly NOT claimed as a guarantee here.

9. Detection Latency: Binding the Gate to BFD Timing (RFC 5880)

A detector is only as useful as it is fast: an attacker's dwell time is exactly the detection latency. NTP's own disciplining is deliberately slow (poll intervals of seconds to thousands of seconds, [RFC5905] Section 13), and the legacy MVPS coherence tick is 60 s. This section binds the cross-vantage gate to the timing discipline of Bidirectional Forwarding Detection [RFC5880], whose detection time is itself a published closed form, and shows the gate inherits a tens-of-milliseconds latency WITHOUT weakening Theorem 1.

Onset-phase model. The gate samples on a tick lattice $t_k = k*T_{tick}$. An injected offset Delta (large enough to be detectable by Theorem 2) appears at onset t_0 with phase $\phi := t_0 - \text{floor}(t_0/T_{tick})*T_{tick}$ in $[0, T_{tick})$. An alarm requires M consecutive above-threshold ticks (the detection multiplier). $\tau_{RTT} \geq 0$ is the one-way latency for the alarm to reach the acting subscriber.

9.1. Theorem 5: Closed-Form Detection-Latency Window (L_{DL}) [ANALYTICAL]

STATEMENT. Under the onset-phase model, the detection latency of the clock-skew gate is

$$\tau_{detect}(\phi) = M*T_{tick} - \phi + \tau_{RTT},$$

and therefore

$$\tau_{min} = (M - 1)*T_{tick} + \tau_{RTT} \quad \text{(best, } \phi \to T_{tick}^-\text{)}$$
$$\tau_{E} = (M - 1/2)*T_{tick} + \tau_{RTT} \quad \text{(expected, } \phi \text{ uniform)}$$
$$\tau_{max} = M*T_{tick} + \tau_{RTT} \quad \text{(worst, } \phi = 0\text{)}.$$

All three are linear in M with slope T_{tick}; the spread $\tau_{max} - \tau_{min} = T_{tick}$ is exactly one tick.

PROOF. The alarm fires at tick index $k_0 + M$, i.e. at $t_{alarm} = (k_0 + M)*T_{tick}$, so $t_{alarm} - t_0 = M*T_{tick} - \phi$; adding τ_{RTT} gives $\tau_{detect}(\phi)$. The three corners follow by substituting $\phi \to T_{tick}^-$, integrating uniformly, and substituting $\phi = 0$. This is Lemma L_{DL}, proved in full in [MVPS-L-DL] Section 2 and validated to the millisecond against the Coherence-BFD benchmark in its Section 4. QED

RFC 5880 binding. Identify M with the BFD Detection Multiplier and T_{tick} with the negotiated BFD transmit interval ([RFC5880] Section 6.8.4, Detection Time = Detection Multiplier x transmit interval). Then τ_{max} is exactly the BFD detection time plus one signalling latency τ_{RTT}. The gate thus rides BFD's own, already-standardized liveness clock.

9.2. Corollary: 1091x Dwell Reduction and M=1 Optimality [ANALYTICAL]

COROLLARY 5.1 (dwell reduction). For the legacy tick (T_{tick} = 60000 ms, M = 1, τ_{RTT} = 5 ms), τ_{max} = 60005 ms. For a BFD-echo cadence (T_{tick} = 50 ms, M = 1, τ_{RTT} = 5 ms), τ_{max} = 55 ms. The attacker's offset-injection dwell window shrinks by a factor 60005/55 ~ 1091.

COROLLARY 5.2 (M = 1 optimality for a deterministic gate). In a

statistical detector the multiplier M > 1 exists to suppress false alarms. Here it is unnecessary: by Theorem 1 the false-positive rate is identically zero at EVERY tick, independent of T_tick and M. Hence the latency-minimizing configuration M = 1 (fire on the first above-threshold tick) loses NOTHING in false alarms while achieving the smallest possible tau_max = T_tick + tau_RTT. Acceleration to BFD cadence is, for this gate, free of any precision/false-alarm trade-off -- a property a statistical NTP-skew monitor does not have.

READING. An on-path adversary who injects a detectable clock offset ([RFC7384] Sections 3.2.2/3.2.3/3.2.6) is caught within T_tick + tau_RTT ~ tens of milliseconds at BFD cadence, versus tens of seconds at the legacy tick and versus the seconds-to-kiloseconds of NTP's own poll discipline -- with zero false alarms either way.

10. Mapping to RFC 8633 Section 5.3 and RFC 7384

[RFC8633] Section 5.3 ("Detection of Attacks through Monitoring") asks operators to monitor for attack signatures. This document supplies one quantitative, host-external signature with proven error behavior:

RFC 8633 Sec 5.3 request	This document
"monitor ... to detect" signature: bogus/zero/MAC (no false-alarm guarantee)	D_{2*tau} over published offsets additive offset Delta (Theorem 2) zero false alarms vs RFC 5905 envelope (Theorem 1)

RFC 7384 threat	Detected when ... (Theorem 2)
3.2.2 Spoofing 3.2.3 Replay 3.2.6 Delay manipulation	imposed offset Delta > theta+2*tau (worst case) / Delta > theta (baseline); single-host blind by the Corollary of Section 7

11. What This Document Does NOT Claim

- No change to the NTP wire protocol, packet format, modes, or algorithms. NTPv4 [RFC5905] and NTPv5 (work in progress) are untouched.
- No replacement for authentication. Integrity and server authentication remain [RFC8915] (NTS) and [RFC8573] (MAC).
- No relativity. The only physical inequality used elsewhere in MVPS is the propagation lower bound RTT >= 2*d/v_g, a triangle inequality at the medium signal speed v_g. Terrestrial IP propagation is sub-relativistic; this document makes NO appeal to special relativity, and the "Einstein"/"Lorentz" labels used in some companion material are editorial only and are avoided here.
- No detection below the envelope. A stationary offset Delta <= theta (and, in the worst placement, Delta <= theta+2*tau) is provably indistinguishable from legitimate within-tier skew (Theorems 1(b), 2). This is a hard limit, stated openly, not a deficiency to be tuned away.
- No single-host capability (Section 7).
- No AI/LLM detection. By Theorem 4 the AI layer cannot detect what the vantages did not observe; it explains an already-gated collapse. Its decision tiers and any accuracy figure are a stated MODEL / EMPIRICAL result (Section 8.3), never a guarantee.

- The tier bounds tau_i (Appendix A) are operator calibration inputs, not numbers fixed by any RFC.

12. Empirical Confirmation (and the corner it does not exercise)

The reference implementation (MVPS axis C10) was run against a graded single-vantage offset sweep Delta in {0, 1, 3, 10, 50, 200, 500} ms with the other vantages published at offset 0 and a stratum-1 tier (tau-class threshold 50 ms). The detector did not fire at Delta <= 50 ms and fired at Delta = 200 ms and Delta = 500 ms. This is exactly Theorem 2(ii) (baseline corner, flag iff Delta > theta) for the implementation's configured threshold.

HONEST GAP. Because every honest vantage was published at EXACTLY 0, this sweep exercises only the baseline corner. It does NOT exercise (a) the worst-case placement of Theorem 2(i) (honest clocks spread to +/- tau), nor (b) the false-positive boundary of Theorem 1(b) (legitimate span = 2*tau). A complete validation MUST add both corners; see the receipt produced alongside this document.

13. Refinements the Theorems Suggest (constructive, non-normative)

This document builds ON the NTP architecture [RFC5905], its security analysis [RFC7384], and its operational guidance [RFC8633]; nothing here is a criticism of that body of work. The two items below are refinements the theorems let us make to OUR OWN reference implementation, recorded so other implementers can reproduce them and so the working group can question the conventions differently if it prefers.

R1 Threshold convention (theta = tau vs theta = 2*tau). Our reference code currently uses the per-vantage tier bound directly as the span threshold (theta = tau). Theorem 1(b) shows the smallest FALSE-POSITIVE-FREE convention is theta = 2*tau, so we adopt 2*tau and state the choice openly. This is a definitional matter, not an error: an operator who declares tau as ALREADY a pairwise (max-minus-min) envelope would correctly keep theta = tau. The recommendation is therefore to state, per deployment, whether tau is a per-vantage or a pairwise bound, and to derive theta accordingly. Reviewers are invited to challenge this convention.

R2 Explicit "unverified" status when telemetry is absent. When a bundle carries no clock-offset telemetry, the theorems say nothing can be concluded about agreement. For transparency we recommend reporting an explicit "clock_unverified" status in that case rather than a bare "ok", so that a consumer is never led to infer agreement that was never measured. This is a reporting improvement, fully backward compatible.

14. Security Considerations

The detector is a monitoring aid, not a control: it raises a flag, it does not correct or authenticate time. An adversary who keeps an injected offset below the envelope (Delta <= theta in the baseline) is provably invisible to it; operators MUST treat the detector as a lower bound on detectable manipulation, layered beneath [RFC8915] authentication and [RFC8633] operational practice. An adversary able to corrupt all vantages' published offsets identically defeats the span test; path and reference diversity (per [RFC7384] Section 3.2.6) is therefore a precondition, captured by Hypothesis H1.

15. IANA Considerations

This document has no IANA actions.

16. References

16.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010, <https://www.rfc-editor.org/info/rfc5905>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017.

16.2. Informative References

[RFC7384] Mizrahi, T., "Security Requirements of Time Protocols in Packet Switched Networks", RFC 7384, October 2014, <https://www.rfc-editor.org/info/rfc7384>.

[RFC8573] Malhotra, A. and S. Goldberg, "Message Authentication Code for the Network Time Protocol", RFC 8573, June 2019.

[RFC8633] Reilly, D., Ed., Stenn, H., and D. Sibold, "Network Time Protocol Best Current Practices", BCP 223, RFC 8633, July 2019, <https://www.rfc-editor.org/info/rfc8633>.

[RFC8915] Franke, D., Sibold, D., Teichel, K., Dansarie, M., and R. Sundblad, "Network Time Security for the Network Time Protocol", RFC 8915, September 2020.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010, <https://www.rfc-editor.org/info/rfc5880>.

[MVPS-L-DL] Melegassi, L., "MVPS Detection-Latency Unified Lemma (L_DL)", Catellix technical note (docs/MVPS_DETECTION_LATENCY_LEMMA.txt and scripts/validate_detection_latency_lemma.py), 2026.

[I-D.melegassi-coherence-bfd] Melegassi, L., "Coherence-BFD: Sub-Second Multi-Vantage Coherence Liveness", Work in Progress, Internet-Draft.

[I-D.melegassi-ippm-mvps-bundle] Melegassi, L., "Multi-Vantage Path Snapshot (MVPS)", Work in Progress, Internet-Draft.

[I-D.melegassi-ganascim-mvps-bbf-mesh] Ganascim, R., Melegassi, L., and G. Ganascim, "MVPS over the Broadband Forum CPE Stack", Work in Progress, Internet-Draft.

[I-D.melegassi-mvps-ai-coherence] Melegassi, L., "AI/LLM Coherence Layer over MVPS", Work in Progress, Internet-Draft.

[MVPS-AI-ENVELOPE]

Melegassi, L., "Stealth-vs-Detection Envelope and AI
Decision Tiers for MVPS", Catellix technical note
(docs/MVPS_STEALTH_DETECTION_AI_ENVELOPE.txt and
scripts/analyze_stealth_detection_ai_envelope.py), 2026.

[Cover-Thomas]
Cover, T. and J. Thomas, "Elements of Information Theory",
2nd ed., Wiley, 2006 (data-processing inequality,
Theorem 2.8.1).

Appendix A. Worked Numbers per NTP Tier

The tau values below are calibration inputs, not RFC-normative. They reflect typical accuracy classes; operators MUST substitute their own.

tier	tau (ms)	theta = 2*tau	Delta_min (baseline..worst)
atomic	1	2	2 ms .. 4 ms
gps_ptp	5	10	10 ms .. 20 ms
ntp_s1	50	100	100 ms .. 200 ms
ntp_s2	200	400	400 ms .. 800 ms
ntp_s3_plus	500	1000	1 s .. 2 s

Reading: a bundle whose loosest clock is stratum-3+ cannot, by Theorem 1, detect any injected offset smaller than ~1 s without additional (e.g. longitudinal) information; tightening the worst clock is the only way to lower Delta_min.

Appendix B. Detection-Latency Variants (L_DL receipt)

The worst-case latency tau_max = M*T_tick + tau_RTT of Theorem 5, evaluated for the five Coherence-BFD benchmark variants ([MVPS-L-DL] Section 4; tau_RTT = 5 ms throughout). The p95 column is the measured benchmark; it matches tau_max to the millisecond.

variant	T_tick(ms)	M	tau_max(ms)	p95(ms)
V0 legacy tick	60000	1	60005	60005
V1 BFD fast	50	3	155	155
V2 BFD demand	1000	1	1005	1005
V3 BFD echo	50	1	55	55
V4 BFD hybrid	50	3	155	155

The latency-minimizing, false-alarm-free configuration is V3 (M = 1, fastest tick): tau_max = T_tick + tau_RTT = 55 ms, a 1091x reduction from the legacy tick (Corollary 5.1) at zero false-alarm cost (Corollary 5.2).

Author’s Address

Leonardo Melegassi
Catellix
Email: melegassi@catellix.com