Network Working Group
Internet-Draft
Intended status: Informational
Expires: 28 November 2026

L. Melegassi
Catellix

28 May 2026

MVPS Vantage Localization Feasibility under MPLS Path Camouflage
draft-melegassi-ippm-mvps-vantage-mpls-00

Abstract

IP geolocation databases are known to be unreliable for network-path measurement purposes [Poese-2011]. Traceroute-based localization -- the practical alternative -- is further corrupted by invisible and opaque MPLS tunnels that suppress IP-TTL propagation, hiding intermediate hops and creating false direct links in the apparent network topology [Donnet-2012] [Vanaubel-2017] [Luttringer-2020].

This document formalises the interaction between MPLS path camouflage and the vantage-authentication problem of the Multi-Vantage Path Snapshot (MVPS) framework [I-D.melegassi-iab-mvps-architecture]. Three technical contributions are introduced.

First, Lemma L-GEO-1 (RTT Localization Bound) establishes the feasible location set for any MVPS vantage given RTT measurements to three or more anchor points, under the assumption that all traversed tunnels are explicit or implicit in the Donnet taxonomy (TTL propagation active).

Second, Lemma L-MPLS-1 (MPLS Camouflage Vulnerability) quantifies the correction term Delta_mpls that invisible and opaque tunnels introduce into the L-GEO-1 bound. For invisible tunnels this correction is unbounded without prior tunnel revelation; for opaque tunnels it is bounded by the hidden-hop count times the minimum per-hop propagation delay.

Third, Theorem T-CAM-1 (MPLS-Aware Camouflage Detection) proves that an MVPS bundle from three or more vantage-to-anchor paths, combined with DPR/BRPR tunnel-revelation probing [Vanaubel-2017] or its TNT implementation [Luttringer-2020], detects MPLS-camouflaged vantage impersonation with probability at least 1 - epsilon under the existing MVPS chi-squared coherence test (Theorem 2 of the v4.0 proof catalogue [v4-proof]). Three explicit caveats (T-CAM-1.A on the i.i.d. assumption of the DKW bound, T-CAM-1.B on the empirical FAR Hypothesis H3 of [v4-proof], and T-CAM-1.C on revelation soundness under adversarial operators) qualify the bound in operational deployment.

An auxiliary lemma L-GEO-1.1 (Anchor Geometry) characterises the necessary and sufficient angular distribution of anchors for L-GEO-1 to discriminate two candidate positions; this gives a deployable guideline for anchor selection.

A new phase label MPLS_CAMOUFLAGE_SUSPECTED is introduced and added to the MVPS phase taxonomy alongside LOCATION_CONSISTENT, LOCATION_MARGINAL, CAMOUFLAGE_SUSPECTED, and SPOOFED_VANTAGE.

Limitations explicitly disclosed include: the symmetric "RTT

inflation" attack (Section 10.2), the PHP tunnel coverage gap when the adversary controls the ingress LER (Section 10.3), and the alignment between the geometric vantage minimum (N >= 3) and the Byzantine vantage minimum (N >= 3f+1, Section 10.4).

All results are proved by discharging MVPS axioms A1..A5 against the structural assets of the Donnet MPLS taxonomy combined with the RTT-ellipsoid localization method. No new wire format is defined; no new codepoints are required. The document is informational.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The Multi-Vantage Path Snapshot (MVPS) framework [I-D.melegassi-iab-mvps-architecture] defines a formal structure for multi-point network coherence measurement. Its Architecture Invariance Theorem states that any instantiation satisfying five structural axioms (MVPS-A1 through MVPS-A5) mechanically inherits a catalogue of nine theorems and two lemmas from the v4.0 existence proof [v4-proof]. Instantiations demonstrated so far include classical Internet paths [I-D.melegassi-ippm-mvps-bundle], satellite orbital segments [I-D.melegassi-ippm-mvps-orbital], IXP meshes [I-D.melegassi-nic-ippm-mvps-ixp-vantage], and broadband-access CPE fleets [I-D.melegassi-ganascim-mvps-bbf-mesh].

All instantiations share a structural assumption that is easy to overlook: that the vantages claimed by the system actually occupy the locations they declare. This assumption is non-trivial. An adversary controlling an MPLS-capable infrastructure can place an MVPS vantage at a remote location while making it appear, to standard traceroute-based localization, as if the vantage is co-located with a legitimate network entry point. The mechanism is well known: invisible MPLS tunnels, as classified by Donnet et al. [Donnet-2012] and studied further in [Vanaubel-2017] and [Luttringer-2020], suppress IP-TTL propagation inside label-switched paths, eliminating intermediate hops from the traceroute output. The result is a false direct link between the ingress Label Edge Router (LER) and the egress LER. An adversary placing a vantage at the egress end of an invisible tunnel can cause the vantage to appear, from the outside, as if it were adjacent to the ingress LER -- potentially in a completely different geographic location.

This vulnerability has two compounding roots. First, IP geolocation databases -- the simplest localization tool -- are known to be unreliable [Poese-2011]: they mis-locate IP addresses frequently enough to render database-based vantage authentication impractical. Second, traceroute-based localization, the standard alternative, is blind to the hidden hops inside invisible tunnels and therefore reports the ingress-to-egress RTT without the ability to attribute it to intermediate topology.

This document addresses both roots simultaneously by formalising the problem using MVPS axiom A1 (the tick-lattice constraint, which encodes timing precision) and the Donnet MPLS taxonomy (which classifies how much of the path each tunnel type hides). The combination yields a three-part result:

(a) A localization lemma (L-GEO-1) valid for paths traversing only explicit or implicit tunnels, giving a closed-form feasible-location set from RTT measurements alone.

(b) A vulnerability lemma (L-MPLS-1) showing that invisible and opaque tunnels break L-GEO-1 unless prior revelation is performed, and quantifying the correction term Delta_mpls where revelation is partial.

(c) A detection theorem (T-CAM-1) showing that the MVPS chi-squared coherence test (inherited Theorem 2) detects MPLS-camouflaged vantage impersonation with probability at least 1 - epsilon, when combined with DPR or BRPR revelation [Vanaubel-2017] or their TNT implementation [Luttringer-2020].

The practical consequence is that an adversary attempting to camouflage a vantage via invisible MPLS tunneling faces two independent detection channels: the RTT-localization feasibility test (L-GEO-1 + L-MPLS-1) and the MVPS coherence residual (T-CAM-1). Neither channel alone is sufficient; both together close the gap.

Section 8 additionally describes how AReST [Dekinder-2025], the 2025 tool for Advanced Revelation of Segment Routing Tunnels, extends the revelation corpus to SR-MPLS and SRv6 infrastructure, providing a forward-compatible path for T-CAM-1 as operators transition from classical MPLS to Segment Routing.

This document is informational. It defines no new wire format, no new codepoints, and no RFC 2119 MUST/SHOULD obligations. It proposes one addition to the MVPS phase taxonomy (MPLS_CAMOUFLAGE_SUSPECTED, Section 7) and one validator (Appendix B).

1.1. Motivation: Using the Taxonomy Against Itself

The central observation of this document is that Donnet’s MPLS taxonomy is both the attack surface and the defence toolkit. The four tunnel types (explicit, implicit, opaque, invisible) define exactly the surface area that an adversary can exploit; and the revelation techniques (DPR, BRPR, TNT, AReST) developed to MEASURE that surface are the same techniques that close the vantage-authentication gap. The MVPS coherence test then provides the statistical binding that makes detection mathematically precise.

1.2. Scope and Non-Goals

This document does not propose any modification to MPLS router behaviour, TTL-propagation defaults, or RFC 4950 [RFC4950] deployment. It does not request any allocation from IANA. Its sole technical contribution is the formal integration of the Donnet MPLS taxonomy with the MVPS vantage-authentication problem.

2. Terminology

The key terms used in this document are defined as follows.

MVPS	Multi-Vantage Path Snapshot framework, as defined in [I-D.melegassi-iab-mvps-architecture].
Vantage (v)	A measurement point participating in an MVPS bundle. Each vantage has a declared position p_v in geographic or topological space.
Anchor (a_i)	A reference node with a known, publicly verifiable position. Used as a fixed point for RTT-based localization. Suitable anchors include RIPE Atlas probes, CAIDA Ark nodes, and IXP route servers with published coordinates.
RTT(v, a_i)	The round-trip time measured from vantage v to anchor a_i. For localization purposes this is the minimum observed RTT over a calibration window.
c_fiber	Speed of light in standard single-mode optical fibre, approximately 2/3 * c_vacuum, i.e., approximately 2e8 m/s.
sigma_NTP	Per-vantage NTPv4 synchronisation error. Under MVPS axiom A1 the joint skew satisfies 2 * sigma_NTP + tau_RTT_max < T_tick.
Feasible Location Set (F_v)	The set of positions consistent with all RTT measurements to anchor set {a_i}, given timing precision sigma_geo = RTT_floor * c_fiber / 2.
Delta_mpls	Correction term introduced by an MPLS tunnel on the path from a vantage to an anchor. Zero for explicit/implicit tunnels; bounded for opaque; potentially unbounded for invisible.
LER	Label Edge Router: the ingress (iLER) or egress (eLER) router of an MPLS Label-Switched Path.
LSR	Label Switching Router: an intermediate router inside an MPLS LSP.
DPR	Direct Path Revelation: probing technique of [Vanaubel-2017] for revealing IP hops hidden inside invisible MPLS tunnels.
BRPR	Backward Recursive Path Revelation: recursive probing technique of [Vanaubel-2017].
TNT	Traceroute for Network Tunnels: implementation of DPR and BRPR in [Luttringer-2020].
AReST	Advanced Revelation of Segment Routing Tunnels: tool for SR-MPLS tunnel revelation [Dekinder-2025].
MVPS-A1..A5	The five structural axioms of the MVPS architecture [I-D.melegassi-iab-mvps-architecture].
T2	Theorem 2 of [v4-proof]: the Mahalanobis D^2 chi-squared coherence test with FAR control.

3. MPLS Tunnel Taxonomy

3.1. Four Canonical Types

The following classification is due to [Donnet-2012], with the opaque type subsequently revised and refined in [Vanaubel-2017] and further replicated at Internet scale in [Huddleston-2025].

Two binary features determine the visibility of an MPLS LSP to traceroute:

Feature F1. TTL-propagate. Whether the ingress LER copies the IP-TTL value into the MPLS LSE-TTL field (ttl-propagate ON), or instead sets LSE-TTL to 255 (ttl-propagate OFF, i.e., no-ttl-propagate).

Feature F2. RFC 4950 [RFC4950]. Whether LSRs include MPLS label-stack information in their ICMP time-exceeded messages.

The four types are:

Explicit (E). F1 = ON, F2 = yes.
All LSRs inside the LSP respond to traceroute and include MPLS labels in ICMP responses.
Full hop-by-hop visibility; semantic label information available.
Delta_mpls = 0.

Implicit (I). F1 = ON, F2 = no.
LSRs respond to traceroute but appear as ordinary IP routers (no label information).
RTT measurements are accurate.
Delta_mpls = 0.

Opaque (O). F1 = OFF, F2 = yes.
Ingress LER sets LSE-TTL = 255; LSRs do not respond to traceroute probes; only the exit hop (eLER) is visible. However, the LSE-TTL value returned by the eLER in its ICMP time-exceeded message reveals the tunnel length:
n_hidden = 255 - LSE-TTL - 1.
Delta_mpls is bounded: see Section 5.2.

Invisible (V). F1 = OFF, F2 = no (or RFC 4950 may be present but without PHP/UHP response).
All LSRs inside the tunnel are completely hidden. The ingress LER appears as a direct neighbour of the egress LER. No length information is available without revelation.
Delta_mpls is unbounded without DPR/BRPR.

3.2. Prevalence (2025 Measurements)

[Huddleston-2025] replicated the [Vanaubel-2017] large-scale MPLS study using 2025 vantage-point data. Key findings:

- At least 30% of Internet paths traverse at least one MPLS tunnel (consistent with [Donnet-2012]).

- Invisible (PHP) tunnels remain the most problematic type;

their fraction relative to total tunnels has remained consistent from 2019 to 2025 despite overall MPLS deployment declining.

- Each invisible tunnel hides an average of 5.7 routers per tunnel (2025 data).

- Explicit tunnels are partially replacing invisible UHP, implicit, and opaque tunnels, suggesting gradual improvement in traceroute transparency -- but not elimination.

These figures establish that invisible MPLS tunnels are not a legacy pathology; they are a current, persistent property of the Internet that any vantage-localization scheme must account for.

4. Vantage Localization in MVPS

4.1. MVPS Axiom A1 and the Tick-Lattice Constraint

MVPS axiom A1 [I-D.melegassi-iab-mvps-architecture] requires that all vantages share a common tick lattice, i.e., that their clocks are synchronised to a common stratum with a joint skew bounded by:

$$2 * \sigma_{NTP} + \tau_{RTT,max} < T_{tick}$$

where σ_{NTP} is the per-vantage NTPv4 synchronisation error (typically < 1 ms on a well-peered stratum-2 source), $\tau_{RTT,max}$ is the maximum observed RTT from any vantage to the NTP server, and T_{tick} is the measurement cadence (typically 100 ms to 1 s in deployed MVPS bundles).

This constraint has a direct implication for localization. The minimum RTT from vantage v to any anchor a_i satisfies:

$$RTT_{min}(v, a_i) \geq 2 * D(p_v, p_{a_i}) / c_{fiber}$$

where $D(p, q)$ is the great-circle distance between positions p and q, and c_{fiber} is the speed of light in fibre (approx. 2e8 m/s). The inequality is tight for paths with no queuing delay and negligible processing delay. Equality does not hold in practice due to routing indirectness, but RTT_{min} provides a hard lower bound.

Note on A1 and timing precision:

$$\sigma_{geo} := RTT_{floor} * c_{fiber} / 2$$

is the localization uncertainty attributable to timing noise and routing indirectness. For σ_{NTP} < 1 ms, σ_{geo} < 100 km -- a resolution appropriate for inter-city localization but not intra-city. Sub-city localization requires additional probing (e.g., multi-anchor FRPLA [Vanaubel-2017]).

4.2. Lemma L-GEO-1: RTT Localization Bound

Scope note: L-GEO-1 is the idealised reference case in which no opaque or invisible MPLS tunnels intervene on any anchor path. In the public Internet of 2025, Section 3.2 establishes

that at least 30% of paths traverse some MPLS tunnel and invisible PHP tunnels remain prevalent; consequently L-GEO-1 alone is rarely applicable outside controlled environments (data-centre fabrics, intra-AS measurement, audited IXP meshes). In the public Internet, the operationally relevant form is L-MPLS-1 (Section 5), which extends L-GEO-1 to account for tunnel-induced corrections.

LEMMA L-GEO-1 (RTT Localization Bound under Transparent Paths).

Pre-conditions:
(P1) M >= 3 anchors {a_1, ..., a_M} with known positions.
(P2) All paths from vantage v to each a_i traverse only Explicit (E) or Implicit (I) tunnels (Delta_mpls = 0). See scope note above.
(P3) Minimum RTT r_i = RTT_min(v, a_i) is measured over a calibration window of at least n_calib samples.

Statement:
Under P1..P3, the feasible location set of v is:

F_v = INTERSECTION over i in {1..M} of
Ball(a_i, r_i * c_fiber / 2 + sigma_geo)

where Ball(c, r) denotes the set of positions within distance r of centre c.

A vantage claiming position p_c with p_c NOT in F_v is LOCATION_INFEASIBLE.

Proof sketch:
Under P2, RTT_min(v, a_i) >= 2 * D(p_v, p_{a_i}) / c_fiber (Section 4.1). Therefore D(p_v, p_{a_i}) <= r_i * c_fiber / 2. Adding sigma_geo for timing noise (bounded by A1) gives p_v in Ball(a_i, r_i * c_fiber / 2 + sigma_geo) for all i. The intersection over M >= 3 non-collinear anchors has bounded diameter (in R^3, three spheres in general position intersect in at most two points, and a fourth anchor resolves the ambiguity). If p_c lies outside this intersection, then D(p_c, p_v) > 0 for all physically feasible p_v, proving infeasibility.

Remark: P2 is the critical condition that Sections 5 and 6 relax. When invisible tunnels are present, r_i may undercount the true path length, inflating the apparent feasible set.

Remark on scope: L-GEO-1 provides a one-sided geometric constraint -- it can REJECT positions whose distance to some anchor exceeds the RTT-derived ball radius, but it CANNOT reject positions that happen to fall inside every ball even though they differ from the true location. Whether the intersection F_v actually discriminates p_c from p_r depends on the angular distribution of the anchors with respect to the line segment [p_r, p_c]; see Lemma L-GEO-1.1 below.

4.3. Lemma L-GEO-1.1: Anchor Geometry for Discrimination

LEMMA L-GEO-1.1 (Anchor Geometry).

Pre-conditions:
(P1') Same as L-GEO-1 pre-conditions P1..P3.
(P4) True position p_r and claimed position p_c with p_r != p_c.

Statement:

The feasible set F_v excludes p_c (i.e., L-GEO-1 detects the lie) if and only if there exists at least one anchor a_k such that:

$$D(p_c, a_k) > RTT_{min}(v, a_k) * c_{fiber} / 2 + \sigma_{geo} \geq D(p_r, a_k)$$

A sufficient geometric condition is that the anchor set {a_i} spans the sphere with enough angular diversity that for any two distinct candidate positions p, p′ on the surface of Earth, there exists at least one a_k satisfying $|D(p, a_k) - D(p', a_k)| > 2 * \sigma_{geo} + \Delta_{mpls,max}$.

Operational interpretation:
The lemma quantifies what "non-collinear anchors" means in L-GEO-1. Three anchors clustered in the same region (e.g., all in Western Europe) leave a large feasible set that may contain both p_r and p_c. Three anchors spanning continents (e.g., one each in North America, Europe, and East Asia) produce a smaller intersection that discriminates inter-continental displacement. Intra-continental claims (e.g., Miami vs. Newark) require either (a) anchors in multiple directions on the same continent, or (b) reliance on the MVPS coherence axes C_2 and C_3 (Theorem T-CAM-1) rather than L-GEO-1 alone.

Proof:
The biconditional is immediate from the definition of F_v (Lemma L-GEO-1, intersection over anchors of Ball(a_i, r_i * c_fiber/2 + sigma_geo)). p_c is in F_v iff for every a_k, $D(p_c, a_k) \leq r_k * c_{fiber}/2 + \sigma_{geo}$. The sufficient condition follows by triangle inequality applied to the pair (p_r, p_c). QED.

Recommendation: Operators SHOULD select anchors so that at least one pair (a_j, a_k) satisfies $D(a_j, a_k) > D(p_r, p_c)$ for the smallest geographic displacement the operator wishes to detect. For inter-city detection at city-pair scale (~1000 km), at least three anchors with mutual distances above 2000 km are required.

5. MPLS Camouflage Vulnerability

When one or more MPLS tunnels of type Opaque (O) or Invisible (V) lie on the path from vantage v to anchor a_i, the RTT measured at the probe source is the full end-to-end RTT from source to eLER. However, the TOPOLOGY inferred from the traceroute hop sequence is false: intermediate LSRs are absent, making the iLER appear as the direct neighbour of the eLER.

An adversary exploiting this property can:

(Attack-A) Place vantage v behind an invisible MPLS tunnel at geographic position p_r while declaring claimed position p_c = p_{iLER}, i.e., the position of the ingress LER. External probes will observe RTT(source, eLER) = RTT(source, v) without detecting the tunnel, and the false topology will show a direct link from iLER to v.

(Attack-B) Use PHP (Penultimate Hop Popping) within an invisible tunnel to cause the second-to-last LSR to decrement the MPLS TTL instead of the eLER, preventing the eLER from sending an RFC 4950 response. This defeats opaque-tunnel detection at the eLER, converting an opaque to a fully invisible tunnel from the measurement side.

5.1. Lemma L-MPLS-1: MPLS Camouflage Correction

LEMMA L-MPLS-1 (MPLS Camouflage Vulnerability).

Let P(v, a_i) be the set of MPLS tunnel segments on the path from vantage v to anchor a_i. For each tunnel segment t in P(v, a_i), let type(t) in {E, I, O, V} be its Donnet type, and n_h(t) be the number of hidden hops (zero for E and I).

Define the per-anchor correction:

```
Delta_mpls(v, a_i) :=
  SUM over t in P(v, a_i) where type(t) in {O, V} of
    n_h(t) * RTT_min_hop
```

where RTT_min_hop is the minimum propagation delay attributable to a single router hop. The choice of RTT_min_hop materially affects the magnitude of Delta_mpls and therefore the size of the corrected feasible set F_v^mpls. Operators MUST select RTT_min_hop using a defensible derivation; this document recommends the following calibration procedure:

(a) Estimate the per-hop propagation floor from the operator's own measurement infrastructure. For a representative sample of EXPLICIT (type-E) MPLS tunnels of known hop count n on the same anchor pool, compute per_hop_floor := median over tunnels of (RTT_explicit / n). Typical values observed in large-scale measurement (CAIDA Ark, RIPE Atlas) fall in the range 0.5-2 ms for co-located rack-to-rack hops and 2-5 ms for inter-PoP hops within the same metropolitan area.

(b) Choose RTT_min_hop as the 10th percentile of the per_hop_floor distribution. Choosing a low percentile is CONSERVATIVE for L-MPLS-1: it gives the adversary the maximum benefit of the doubt by subtracting the largest plausible Delta_mpls, shrinking F_v^mpls as little as possible.

(c) Re-calibrate RTT_min_hop quarterly or whenever the operator's anchor topology changes materially.

When operator-specific calibration is not available, this document specifies RTT_min_hop = 2 ms as a default. This default is justified as the approximate 10th percentile of the per-hop floor distribution reported in [Huddleston-2025] Table 4 for invisible-tunnel intra-tunnel hop counts in 2025 IPv4 measurements. Operators using the default SHOULD document this choice in their MVPS deployment notes; the default is NOT a normative constant of this specification.

Then the corrected feasible-location set under MPLS is:

```
F_v^mpls = INTERSECTION over i in {1..M} of
             Ball(a_i, (r_i + Delta_mpls(v, a_i))
                        * c_fiber / 2 + sigma_geo)
```

For type O tunnels:
n_h(t) is observable from the LSE-TTL value returned by the eLER (Section 3.1). Delta_mpls is bounded and computable. F_v^mpls is a superset of F_v but remains bounded.

For type V tunnels:
n_h(t) is unknown without DPR/BRPR/TNT revelation. In the worst case n_h(t) is unbounded (255 - 1 hops maximum in a single LSP label stack), so F_v^mpls degenerates to an unbounded set: L-GEO-1 cannot guarantee localisation.

5.2. Per-Type Analysis

Type	F1	F2	Delta_mpls	L-GEO-1 intact?
E	ON	yes	0	YES
I	ON	no	0	YES
O	OFF	yes	bounded	YES (superset, bounded)
V	OFF	any	unbounded*	NO (without revelation)

* Unless DPR/BRPR/TNT probing reveals n_h(t); see Section 8.

Corollary L-MPLS-1.1 (Implicit PHP Attack). The PHP variant of an invisible tunnel (Attack-B, Section 5) converts an operationally opaque tunnel into a type-V tunnel from the MEASUREMENT perspective, making n_h(t) unobservable via standard ICMP LSE-TTL inspection. DPR is required to recover n_h(t).

Operational impact: In the 2025 replication study [Huddleston-2025], invisible PHP tunnels hid an average of 5.7 hops per tunnel. At RTT_min_hop = 2 ms per hop, Delta_mpls = 11.4 ms per invisible tunnel, corresponding to a false position credit of approximately 1140 km per invisible tunnel. An adversary traversing three invisible tunnels in series could mask geographic displacement exceeding 3400 km -- effectively spanning a continent -- while appearing legitimate to any localization scheme that does not perform tunnel revelation.

6. MVPS-Aware Camouflage Detection

6.1. Theorem T-CAM-1: Detection via Coherence Test

THEOREM T-CAM-1 (MPLS-Aware Camouflage Detection).

Pre-conditions:
(Q1) MVPS bundle with N >= 3 vantages; vantage v claims position p_c.
(Q2) M >= 3 anchors with known positions and measured RTTs from v, calibrated over n_calib >= 18,500 observations (MVPS operational contract OC3).
(Q3) DPR/BRPR or TNT [Luttringer-2020] has been run on all paths from v to {a_i} and the revelation output has identified all type-V and type-O tunnel segments with their corrected n_h(t).
(Q4) The corrected feasible set F_v^mpls (Lemma L-MPLS-1) has been computed.

Statement:
Let p_r be the true position of vantage v, and let

p_c be its declared position. If p_c NOT in F_v^mpls, then the MVPS chi-squared coherence test (Theorem 2 of [v4-proof]) detects the vantage as LOCATION_INFEASIBLE with probability at least 1 - epsilon, where epsilon satisfies the DKW bound [I13]:

epsilon <= exp(- 2 * n_calib * gamma^2)

with gamma = (FAR_target / 2) and FAR_target the operator-chosen false alarm rate. For FAR_target = 0.01 and n_calib = 18,500, epsilon < 1e-9.

Proof sketch:
Under Q1..Q4, the corrected L-GEO-1 test (L-MPLS-1) maps p_c to LOCATION_INFEASIBLE. This infeasibility manifests as a systematic offset in at least one coherence axis C_j (j in {1,2,3}): the RTT-based C_1 axis reports the ingress-to-egress RTT (which is physically consistent with p_{eLER}), while the path-topology inferred C_3 axis (Jaccard similarity on touched-object sets) is inconsistent with p_c, since the hidden hops DO influence the routing table of p_{eLER} even though they are invisible to traceroute.

The joint Mahalanobis D^2 on (C_1, C_2, C_3) detects this inconsistency at threshold q_J (MVPS Theorem 2 + Theorem 4, [v4-proof]). Calibration over n_calib samples bounds the FAR by the DKW inequality (Imported Result I13 of [v4-proof]). Both theorems inherit from the Architecture Invariance Theorem [I-D.melegassi-iab-mvps-architecture] since axioms A1..A5 are satisfied by any vantage that participates in a valid MVPS bundle. QED.

Remark: The detection relies on C_3 (topological axis) being inconsistent. If the adversary also spoofs the routing-table content of the eLER (not just its location), detection requires the Byzantine-robust extension described in Theorem 9 of [v4-proof], which bounds the adversarial bias on the centroid.

Caveat T-CAM-1.A (Independence assumption). The DKW bound I13 of [v4-proof] requires that the n_calib calibration observations be independent and identically distributed (i.i.d.). An adversary aware of the calibration window can degrade the i.i.d. assumption by correlating tunnel activation with diurnal traffic patterns, BGP convergence events, or peering reconfigurations. Operators SHOULD partition the calibration window across multiple non-overlapping epochs (recommendation: four 90-minute windows separated by at least 24 hours) and verify per-epoch FAR stability before asserting the global bound.

Caveat T-CAM-1.B (Empirical FAR hypothesis). Theorem 2 of [v4-proof] guarantees the chi-squared distribution of D^2 under the conditions of axioms MVPS-A1..A3. The realized false-alarm rate within +/- 25% of nominal is Hypothesis H3 of [v4-proof], which is empirically supported but NOT formally proven for non-Gaussian C(t) distributions (see [v4-proof] Section "Hypotheses for empirical validation"). The epsilon < 1e-9 figure quoted above inherits this empirical conditioning. Operators with strict FAR requirements SHOULD validate Hypothesis H3 on a per-deployment basis using the DKW-bound test specified in [v4-proof] OC3, and tighten n_calib if observed FAR departs from nominal by more than 25%.

Caveat T-CAM-1.C (Revelation soundness). Pre-condition Q3

assumes that TNT or AReST revelation produces accurate n_h(t). As detailed in Section 10.1, this assumption fails under adversarial MPLS operators (Attacks C and D). In such environments the effective epsilon is bounded by the minimum of the DKW bound and the revelation success probability, which must be characterized operationally.

6.2. Corollary T-CAM-1.1: CWT Cross-Binding

COROLLARY T-CAM-1.1 (CWT Cross-Binding).

Under the CWT trust model [I-D.melegassi-santos-ippm-mvps-cwt], a vantage v that:

(a) presents a valid CWT token (T-AUTH-CWT-1 is satisfied), AND
(b) is flagged LOCATION_INFEASIBLE by the L-MPLS-1 test

is classified as MPLS_CAMOUFLAGE_SUSPECTED.

Rationale: CWT authentication establishes cryptographic identity of the measuring process; it does not authenticate the physical location. A valid CWT token from a vantage at p_r, presented under a claimed location p_c NOT in F_v^mpls, is a combination that the CWT model cannot rule out but that the L-MPLS-1 localization can. The MPLS_CAMOUFLAGE_SUSPECTED label precisely captures this disjunction: "we cannot deny the identity, but physics denies the location."

7. Phase Taxonomy Extension

The MVPS phase taxonomy, introduced in the base bundle [I-D.melegassi-ippm-mvps-bundle], is extended by this document with one new phase label and a new detection pathway.

Revised MVPS phase taxonomy (phase labels ordered by severity):

COHERENT
Normal operation. $D^2 < q_J$. L-GEO-1 passes.

DRIFTING
Coherence degrading. D^2 approaching q_J.
L-GEO-1 passes.

LOCATION_CONSISTENT
L-GEO-1: claimed position p_c IN F_v.
Revelation: no invisible tunnels detected on anchor paths.

LOCATION_MARGINAL
L-GEO-1: p_c within sigma_geo of boundary of F_v.
Revelation: no invisible tunnels, but path is MPLS-rich.
Operator should increase anchor count to M >= 5.

MPLS_CAMOUFLAGE_SUSPECTED [NEW -- this document]
L-MPLS-1: invisible or opaque tunnels found on anchor paths.
Corrected F_v^mpls excludes p_c.
CWT: authentication status may be valid or invalid.
ACTION: Run DPR/BRPR on all anchor paths; if revelation confirms n_h(t) and p_c remains outside F_v^mpls, escalate to CAMOUFLAGE_CONFIRMED.

CAMOUFLAGE_SUSPECTED
L-GEO-1: p_c NOT in F_v (no MPLS tunnels involved).
CWT: authentication valid (identity present but location physically infeasible without tunneling).

SPOOFED_VANTAGE
L-GEO-1/L-MPLS-1: p_c outside feasible set.
CWT: authentication INVALID.
Full rejection; remove vantage from bundle.

The ordering is informational; operators may choose their own escalation policy. The MVPS phase state machine [I-D.melegassi-ippm-mvps-bundle] treats any phase from MPLS_CAMOUFLAGE_SUSPECTED upward as requiring operator intervention.

8. Tunnel Revelation Integration

8.1. DPR and BRPR (Classical MPLS)

[Vanaubel-2017] introduced two techniques for revealing IP hops hidden inside invisible MPLS tunnels.

Direct Path Revelation (DPR):
DPR sends probes with systematically decremented MPLS TTL values toward the tunnel, causing individual LSRs to issue ICMP time-exceeded messages. The source IP addresses of these messages reveal the hidden hops in forward order. DPR requires that the measurement vantage be co-located with (or close to) the iLER, so that it can manipulate the MPLS label stack.

Backward Recursive Path Revelation (BRPR):
BRPR uses traceroute probes from the OUTSIDE toward the tunnel egress, with TTL values set to exactly reach each hidden LSR from the probe source. This does not require access to the iLER. BRPR is iterative: it discovers hops one by one from the eLER backward.

For the MVPS vantage-localization use case, BRPR is the preferred technique because:

(a) The measurement point (probe source) is NOT inside the tunnel (the adversary's tunnel is between the claimed vantage and the anchor).
(b) BRPR requires only standard traceroute probing capability from the probe source, with no access to the iLER.

Once revelation is complete, n_h(t) is known for each tunnel segment, Delta_mpls is computable, and the corrected feasibility test F_v^mpls can be evaluated (Theorem T-CAM-1, pre-condition Q3).

TNT [Luttringer-2020] implements DPR and BRPR in a single tool (forked from scamper [scamper]). It is the recommended implementation for integrating tunnel revelation into an MVPS measurement pipeline.

8.2. AReST Integration (Segment Routing)

As operators migrate from classical MPLS LSPs to Segment Routing (SR-MPLS and SRv6), the tunnel-camouflage threat surface migrates with them. SR-MPLS tunnels can exhibit the same visibility categories as classical MPLS tunnels, depending on SID type and TTL propagation configuration.

[Dekinder-2025] (AReST -- Advanced Revelation of Segment Routing Tunnels) extends the revelation corpus to SR-MPLS infrastructure. For MVPS vantage-localization purposes, AReST provides the same output as TNT: the revealed list of hidden hops for each SR tunnel segment on the anchor paths.

The integration is mechanical: replace the TNT revelation step in pre-condition Q3 of Theorem T-CAM-1 with AReST for SR-MPLS paths. All other steps, including the L-MPLS-1 correction and the T-CAM-1 coherence test, are unchanged.

Forward compatibility note: this document recommends that MVPS implementations maintain a revelation backend abstraction that can be satisfied by either TNT (classical MPLS) or AReST (SR-MPLS/SRv6), with the backend selected based on the MPLS label type observed in the Explicit-tunnel responses on anchor paths.

9. Deployment Considerations

9.1. Anchor Selection

For the L-GEO-1 and L-MPLS-1 bounds to be useful, anchors must satisfy three properties:

(i) Known geographic position, independently verifiable (e.g., RIPE Atlas site data, published IXP location).
(ii) Paths from the vantage under test to each anchor are route-stable over the calibration window.
(iii) Anchors are geographically distributed, not co-located in the same facility.

Suitable anchor pools: RIPE Atlas anchors (globally distributed, publicly queryable, route-stable by design), CAIDA Ark vantage points, or IXP route-server addresses published in the PeeringDB database.

9.2. Calibration Window

n_calib >= 18,500 observations (MVPS OC3) yields epsilon < 1e-9 in Theorem T-CAM-1 at FAR_target = 0.01. At a 1-second measurement cadence, this requires approximately 5.1 hours of continuous measurement. Operators SHOULD run an initial calibration phase of at least 6 hours before treating MVPS_CAMOUFLAGE_SUSPECTED labels as actionable.

9.3. Revelation Frequency

TNT/BRPR probing is heavier than normal MVPS path probing. Operators SHOULD run revelation on:

(a) Initial vantage enrollment.
(b) After any BGP route change that affects anchor paths.
(c) At a low-frequency periodic interval (e.g., weekly) to detect newly deployed tunnels.

10. Security Considerations

An adversary with access to MPLS-capable infrastructure can deploy invisible tunnels to camouflage the true geographic position of an MVPS vantage. The techniques in this document detect such camouflage but do NOT prevent it. Prevention requires either:

(a) Operator-side enforcement of ttl-propagate on all MPLS edges (eliminating invisible tunnels), or
(b) Cryptographic location attestation (e.g., hardware-anchored GPS or eLoran timestamps) that is independent of network-layer path measurement.

Neither (a) nor (b) is proposed in this document; they are operational choices outside the MVPS framework.

The Byzantine-robust extension (Theorem 9 of [v4-proof]) provides additional protection when the adversary also manipulates the routing-table content of the compromised vantage. Operators facing sophisticated adversaries SHOULD also deploy the geometric-median centroid estimator (MVPS Design D9(ii)) and enforce N >= 3f + 1 vantages, where f is the maximum number of Byzantine vantages the operator wishes to tolerate. Pre-condition Q1 of Theorem T-CAM-1 (N >= 3) is the geometric minimum for localization, NOT the Byzantine minimum; see Section 10.4 below.

10.1. Limitations of DPR/BRPR/TNT under Adversarial Operators

Theorem T-CAM-1 pre-condition Q3 assumes that DPR, BRPR, or TNT [Vanaubel-2017] [Luttringer-2020] revelation correctly identifies hidden hops n_h(t) for every tunnel segment. This assumption is sound when the MPLS operator merely CONCEALS tunnel topology (the threat model under which DPR/BRPR were originally analysed) but is NOT sound when the operator is itself the adversary.

Specifically, an adversary who controls the MPLS infrastructure can:

(Attack-C) Forge ICMP time-exceeded responses. DPR and BRPR rely on receiving ICMP responses from intermediate LSRs. Source IP, TTL value, and timestamps in these responses are not authenticated. An adversarial LSR can fabricate responses consistent with a benign tunnel topology, causing TNT to report a smaller n_h(t) than the true value.

(Attack-D) Suppress revelation probes. An adversary can rate-limit or drop probes whose pattern matches known DPR/BRPR signatures, leaving the defender with no observation at all (which under Q3 must be treated as "no tunnel found", a soft failure of the test).

This document does NOT solve Attack-C or Attack-D. Operators deploying T-CAM-1 in environments where the on-path MPLS operator may be adversarial SHOULD:

(i) Run revelation from multiple geographically and administratively independent probe sources, accepting n_h(t) only when at least two independent sources report values within tolerance.
(ii) Augment T-CAM-1 with the CWT trust binding (Corollary T-CAM-1.1) so that a forged revelation report cannot in itself validate a forged vantage.
(iii) Treat absence of revelation response (Attack-D) as equivalent to "invisible tunnel suspected" rather than "no tunnel".

Limitation note: even with mitigations (i)-(iii), an adversary who controls the entire forwarding path between vantage and anchor remains outside the protection envelope of this document. Such adversaries require physical-layer attestation (Security Considerations (b)) which is out of scope here.

10.2. RTT Inflation Attack (Dual of Camouflage)

The MPLS camouflage analysed in Sections 5 and 6 is the case where invisible tunnels HIDE distance, causing the apparent feasible set F_v to undercount the true path length. The symmetric "RTT inflation" attack is the case where the adversary INFLATES the measured RTT to claim a distant location.

Mechanisms for RTT inflation include:

- Kernel-side deterministic delay injection in the vantage's TCP/UDP probe response path.
- BGP path prepending to force a longer AS path.
- Routing through a deliberately distant intermediate hop under operator control.

Lemma L-MPLS-1 does not detect inflation because Delta_mpls only SUBTRACTS hidden-hop time from the measured RTT; it never challenges measured RTT as anomalously large. An adversary at true position p_r who inflates RTT(v, a_i) by tau_inflate appears to occupy a Ball(a_i, (r_i + tau_inflate) * c_fiber / 2) which can extend to a falsely distant p_c.

Detection of inflation requires two complementary techniques not formalized in this document:

- Multi-anchor RTT consistency: comparing measured RTT to the minimum RTT predicted by the speed-of-light floor 2 * D(p_c, a_i) / c_fiber. Excessive ratio measured/ floor across multiple anchors is suspect.
- Cross-stratum NTP/PTP timing audit: an adversary inflating RTT generally also inflates timestamps at the wire, which can be detected by comparing to an external time reference (e.g., GPS PPS or NIST stratum-1).

Operators SHOULD treat both camouflage (this document) and inflation (this section) as a coupled threat surface and deploy detection for both. A "Lemma L-INFL-1" formalizing inflation detection is left to a future document.

10.3. PHP Tunnel Coverage Gap

Corollary L-MPLS-1.1 (Section 5) notes that a PHP (Penultimate

Hop Popping) configuration converts an operationally opaque tunnel into a type-V tunnel from the measurement perspective, requiring DPR to recover n_h(t). However, DPR requires that the measurement source be CO-LOCATED with (or have privileged access to) the ingress LER of the tunnel under inspection [Vanaubel-2017].

In an adversarial scenario the defender does NOT have access to the adversary's iLER, by construction. Hence:

- BRPR can be attempted from the defender's side; it partially recovers n_h(t) for type-V tunnels but its success rate degrades when PHP is combined with selective label-stack popping.
- TNT, which implements both DPR and BRPR, is constrained to its BRPR mode in this case.

Operators facing PHP-rich adversarial environments SHOULD:

(i) Increase n_calib and tighten FAR_target to compensate for the increased revelation uncertainty.
(ii) Treat any anchor path showing PHP-suspect ICMP response patterns as inflating Delta_mpls to its worst-case bound (n_h(t) = 17 hops, the 99th-percentile observed in [Huddleston-2025]) rather than the average.
(iii) Prefer anchors connected via SR-MPLS or SRv6 infrastructure where AReST [Dekinder-2025] applies, since AReST's revelation primitives operate on the segment list rather than relying on LER-side label manipulation.

10.4. Pre-condition Alignment with Byzantine Bound

Pre-condition Q1 of Theorem T-CAM-1 requires N >= 3 vantages. This is the GEOMETRIC minimum for trilateration. The MVPS architecture also imposes a BYZANTINE minimum of N >= 3f + 1 for resilience against f compromised vantages (axiom MVPS-A5, Theorem 9 of [v4-proof]). These two minima are independent:

- For pure localization with f = 0 (trusted vantages, MPLS infrastructure may be hostile), N = 3 suffices.
- For localization with f = 1 (one vantage may be compromised in addition to MPLS hostility), N >= 4.
- For localization with f = 2, N >= 7.

Operators MUST select N as the maximum of the geometric and Byzantine minima for their threat model. This document's probability bound (epsilon < 1e-9 with n_calib = 18,500, FAR_target = 0.01) assumes N satisfies BOTH minima.

11. IANA Considerations

This document has no IANA actions.

12. Acknowledgments

This document would not exist without the 14-year corpus of work on MPLS tunnel revelation by Benoit Donnet (Universite de Liege) and his collaborators. The four-type taxonomy of MPLS tunnels (explicit, implicit, opaque, invisible), the DPR and BRPR revelation primitives, the TNT implementation, and the AReST extension to Segment Routing form the structural foundation on which Lemma L-MPLS-1 and Theorem T-CAM-1 are built. In particular, [Donnet-2012], [Vanaubel-2017],

[Luttringer-2020], and [Dekinder-2025] provide the measurement-theoretic vocabulary that makes the MVPS vantage-authentication problem tractable. Any errors of formalisation or attribution in the present document are the author's own.

The author also thanks the IPPM, INTAREA, and DISPATCH working groups for the discussions that shaped the MVPS architecture series referenced herein, and the Catellix engineering team for the validator scaffolding referenced in Appendix B.

13. References

13.1. Normative References

[I-D.melegassi-iab-mvps-architecture]
Melegassi, L., "Multi-Vantage Path Snapshot: Architecture Invariance Theorem", draft-melegassi-iab-mvps-architecture-00, May 2026.

[I-D.melegassi-ippm-mvps-bundle]
Melegassi, L., "Multi-Vantage Path Snapshot: Bundle Envelope and Coherence Algebra", draft-melegassi-ippm-mvps-bundle-00, May 2026.

[I-D.melegassi-santos-ippm-mvps-cwt]
Melegassi, L. and R. Santos, "Coherent-Witness Trust for MVPS Vantage Authentication", draft-melegassi-santos-ippm-mvps-cwt-00, May 2026.

[RFC4950] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "ICMP Extensions for Multiprotocol Label Switching", RFC 4950, DOI 10.17487/RFC4950, August 2007, <https://www.rfc-editor.org/rfc/rfc4950>.

[v4-proof] Melegassi, L., "MVPS Mathematical Existence Proof -- Version 4.0", May 2026, <https://www.catellix.com/static/download/MVPS_MATHEMATICAL_EXISTENCE_PROOF_V4.txt>.

13.2. Informative References

[Donnet-2012]
Donnet, B., Luckie, M., Merindol, P., and J.-J. Pansiot, "Revealing MPLS Tunnels Obscured from Traceroute", ACM Computer Communication Review, vol. 42, no. 2, pp. 87-93, DOI 10.1145/2185376.2185388, April 2012.

[Vanaubel-2017]
Vanaubel, Y., Merindol, P., Pansiot, J.-J., and B. Donnet, "Through the Wormhole: Tracking Invisible MPLS Tunnels", ACM Internet Measurement Conference (IMC 2017), DOI 10.1145/3131365.3131378, November 2017.

[Luttringer-2020]
Luttringer, J.-R., Vanaubel, Y., Merindol, P., Pansiot, J.-J., and B. Donnet, "Let There Be Light: Revealing Hidden MPLS Tunnels with TNT",

IEEE Transactions on Network and Service
Management, vol. 17, no. 2, pp. 1239-1253,
DOI 10.1109/TNSM.2019.2962278, June 2020.

[Dekinder-2025]
Dekinder, F., Vermeulen, K., and B. Donnet,
"Autonomous Systems under AReST: Advanced
Revelation of Segment Routing Tunnels",
ACM Internet Measurement Conference (IMC 2025),
DOI 10.1145/3730567.3764436, October 2025.

[Huddleston-2025]
Huddleston, J., Luckie, M., and A. Marder,
"Replication: Characterizing MPLS Tunnels over
Internet Paths", ACM Internet Measurement
Conference (IMC 2025), 2025.

[Poese-2011]
Poese, I., Uhlig, S., Kaafar, M. A., Donnet, B.,
and B. Gueye, "IP Geolocation Databases:
Unreliable?", ACM Computer Communication Review,
vol. 41, no. 2, pp. 53-56,
DOI 10.1145/1971162.1971171, April 2011.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding
Detection (BFD)", RFC 5880,
DOI 10.17487/RFC5880, June 2010.

[scamper] Luckie, M., "Scamper: A Scalable and Extensible
Packet Prober for Active Measurement of the
Internet", ACM Internet Measurement Conference
(IMC 2010), 2010.

[I13] Massart, P., "The Tight Constant in the
Dvoretzky-Kiefer-Wolfowitz Inequality",
Annals of Probability, vol. 18, no. 3, 1990.

Appendix A. Worked Example: Invisible-Tunnel Attack

This appendix illustrates Attack-A (Section 5) with concrete numbers, using the 2025 prevalence data from [Huddleston-2025]. The scenario is constructed so that the attack defeats L-GEO-1 in isolation but is caught by L-MPLS-1 after tunnel revelation; this is the regime where the contribution of this document is operationally significant.

Note A.1 (Physics constraint). Because RTT_min(v, a_i) >= 2 * D(p_r, p_{a_i}) / c_fiber by Section 4.1, an adversary cannot REDUCE the measured RTT below the great-circle floor. Examples in which the adversary "appears closer" than physically possible are inadmissible. The example below respects this floor: the adversary's true location p_r is closer to the anchor than the claimed location p_c, so MPLS inflation of the RTT is what creates room for the false claim. See also the dual "RTT inflation" attack discussed in Section 10.2.

Scenario:
- Adversary places vantage v at p_r = Newark, NJ, USA.
- Adversary claims vantage position p_c = Miami, FL, USA (e.g., to satisfy a regional SLA or geo-licensing

requirement that demands a Southeast U.S. presence).
- D(p_r, p_c) approximately 1750 km.
- Anchor a_1 located in Chicago, IL.
- D(Newark, Chicago) approximately 1170 km (physical RTT floor ~= 11.7 ms).
- D(Miami, Chicago) approximately 2090 km (physical RTT floor ~= 20.9 ms).
- One invisible MPLS tunnel on the path from v to a_1, hiding 5.7 hops (2025 average per tunnel, see Section 3.2).
- RTT_min_hop = 2 ms (per-hop conservative estimate, see Section 5.1).

Without revelation:
- RTT(v, a_1) measured = 25 ms. This value is physically admissible: 25 ms > 11.7 ms (Newark-Chicago floor) AND 25 ms > 20.9 ms (Miami-Chicago floor), so neither location is geometrically rejected by RTT alone. The adversary's true path (Newark -> MPLS LSP hiding 5.7 routers in a detour through Dallas -> Chicago) produces a higher RTT than the direct Newark-Chicago path would, plausibly attributable to BGP indirectness.
- L-GEO-1 (uncorrected): F_v = Ball(Chicago, 25 ms * c_fiber / 2 + sigma_geo) approx Ball(Chicago, 2500 km). Miami (at 2090 km from Chicago) is INSIDE F_v. Attack succeeds against L-GEO-1; vantage appears legitimate.

After BRPR/TNT revelation:
- One invisible tunnel detected on the Newark-Chicago path; 5.7 hidden hops recovered by BRPR probing.
- Delta_mpls(v, a_1) = 5.7 * 2 ms = 11.4 ms.
- Corrected bound: (25 - 11.4) ms * c_fiber / 2 + sigma_geo approx Ball(Chicago, 1360 km).
- Miami-Chicago distance = 2090 km > 1360 km.
- p_c = Miami is NOT in F_v^mpls.
- Vantage is flagged MPLS_CAMOUFLAGE_SUSPECTED.

After MVPS coherence test:
- C_3 (topological axis): Jaccard similarity on touched objects between the actual Newark-originated AS path and the AS path expected from a Miami-originated probe to Chicago. Illustrative values (calibration-dependent; see Note A.2): typical co-located coherence approximately 0.85 +/- 0.05; observed value approximately 0.30 +/- 0.10. D^2 > q_J at FAR_target = 0.01.
- Phase escalates to CAMOUFLAGE_CONFIRMED.

Note A.2 (Illustrative Jaccard values). The values 0.85 and 0.30 above are operationally typical for the BGP-AS topology of the U.S. East Coast as observed in CAIDA Ark and RIPE Atlas datasets (2024-2025). They are NOT theoretical constants and MUST be re-calibrated per anchor pool and per measurement epoch before being used as decision thresholds. See Section 9.2 and the MVPS calibration contract OC3 [I-D.melegassi-ippm-mvps-bundle].

The attack is defeated by the combination of:
(1) TNT/BRPR tunnel revelation (Donnet's techniques), and
(2) MVPS coherence test (Theorem 2 + Theorem 4 of [v4-proof]).

The example is deliberately conservative (one tunnel, one anchor, modest geographic displacement). Attacks involving multiple chained invisible tunnels or larger displacements produce proportionally larger Delta_mpls corrections and are easier to detect once revelation is performed.

Appendix B. Validator Notes

A companion validator is being developed at:

scripts/validate_vantage_localization.py

The validator takes as input:
- Anchor positions {p_{a_i}} (lat/lon)
- Measured RTTs {r_i} from vantage to each anchor
- TNT/AReST revelation output (hidden hop counts per tunnel)
- Claimed vantage position p_c

It outputs:
- F_v (L-GEO-1 feasible set, assuming transparent paths)
- F_v^mpls (L-MPLS-1 corrected feasible set)
- Membership of p_c in F_v and F_v^mpls
- Phase label recommendation (from Section 7)
- JSON receipt for SHA-256 verification

The validator follows the same structure as scripts/validate_ixp_vantage.py (D-18) and is designed for exit-code 0 on PASS, 1 on FAIL.

Author’s Address

Leonardo Melegassi
Catellix
Andradina, SP
Brazil
Email: melegassi@catellix.com
URI: https://www.catellix.com