IETF
Internet-Draft
Intended status: Standards Track
Expires: 27 October 2026

A. Subbiah
Independent

April 2026

IPv7: Identity-Centric Network Protocol for Security, Proxy Mitigation, and Operability
draft-subbiah-ipv7-00

Abstract

This document specifies a network-layer protocol, IPv7, that extends the Internet Protocol model with an identity-carrying address form and an origin-validation mechanism intended to mitigate abuse of residential proxy infrastructure. IPv7 replaces purely numerical source addressing with a hierarchical identity string and a Variable-Length Identity Block (VLIB) that carries an Ephemeral Identity Token (EIT), provider and tenant identifiers, role/policy signalling, and an Origin Signature verifiable by the originating provider. The protocol enables routers to apply policy and reputation signals at the network layer while limiting disclosure of a subscriber’s long-term identity to intermediate systems. This document addresses growing security challenges in Internet-connected devices (IoT), including smart TVs, appliances, and other residential endpoints that are vulnerable to residential proxy exploitation and botnet infection.

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 3 October 2026.

Table of Contents

1. Introduction

The current Internet architecture (IPv4 and IPv6) is built upon the principle of "reachability first, security second." IP addresses identify connection points or topological locations, not the identity or intent of the sender. This fundamental architectural gap has enabled the proliferation of residential proxy networks - a multi-billion-dollar market where malicious actors mask their identity behind legitimate consumer IP addresses to conduct fraud, credential stuffing, and distributed denial of service (DDoS) attacks.

With the explosive growth of Internet-connected devices (IoT) - including smart TVs, robotic vacuum cleaners, refrigerators, washing machines, and other consumer appliances - the security challenges have intensified. Most of these devices run Linux or Android-based systems with customizable network stacks that can be modified at the kernel level. However, current IPv4 and IPv6 lack the native mechanisms to authenticate the true origin of traffic or enforce network-layer policy based on device identity and trust.

IPv7 defines an address and header structure that allows an IPv7 node to (1) validate the binding between an asserted provider and the packet origin, and (2) convey coarse-grained policy and reputation signals to forwarding devices. In IPv7, the source address is treated as an authenticated assertion rather than a purely topological locator. This document describes the on-wire format, processing model, and security considerations for these mechanisms.

1.1. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

IPv7 Node: A host or router that originates, forwards, or terminates IPv7 packets.

Ephemeral Identity Token (EIT): A time-bound token carried in the User Identity component of the Variable-Length Identity Block (VLIB).

Provider ID: An identifier for the originating ISP or provider domain used for signature verification and policy decisions.

Origin Signature: A cryptographic signature that enables a verifier to validate the source/provider binding and detect unauthorised proxying.

Variable-Length Identity Block (VLIB): A structured identity block that carries EIT, service, location, provider, tenant, role/policy, and signature components.

Source-Provider Validation (SPV): A validation procedure that checks whether an IPv7 packet’s claimed Provider ID matches the provider that can validate the Origin Signature.

1.2. Key Design Pillars

Identity-Carrying Addressing: The address syntax includes explicit service, location, provider, tenant, and role/policy components to support operational attribution.

Origin Validation: Packets include an Origin Signature that enables Source-Provider Validation (SPV) at, or near, the first hop.

Trust and Reputation Signalling: Packets carry compact trust/ reputation indicators to support fast-path filtering and traffic management.

Privacy Preservation: The EIT mechanism is intended to limit disclosure of a subscriber’s long-term identity to intermediate routers while enabling accountability by the originating provider.

1.3. Goals and Capabilities

IPv7 is designed to achieve the following objectives:

Signal-Based Endpoint Integrity Monitoring: While IPv7 cannot perform malware analysis on endpoints, it SHOULD enable provider-edge routers to receive and act upon signals from machine-learning (ML) analysers. These signals allow routers to flag packets from potentially compromised devices based on behavioral analysis, traffic pattern anomalies, and reputation scoring. This approach shifts responsibility for anomaly detection to the provider's infrastructure rather than requiring it at the endpoint.

Fraud Mitigation through Network-Layer Attribution: IPv7 is intended to reduce the value and effectiveness of unauthorized residential proxying by providing cryptographic binding between a packet's asserted origin and the originating provider. This enables targets to distinguish legitimate residential traffic from proxy-relayed abuse, supporting fraud prevention at the network layer.

Provider-Operated Identity Domains: IPv7 assumes that providers operate local validation domains with independent policy. This architecture supports jurisdictional compliance and operational flexibility while maintaining cryptographic integrity across domain boundaries.

Legal Process Support: IPv7 provides a framework for identity disclosure under law enforcement requests. The exact legal and policy mechanisms for such disclosures remain jurisdictional and are outside the scope of this protocol specification, but the protocol is designed to support such processes at the provider level.

Network-Layer Policy Enforcement: By including role, policy, and trust-level information in the address itself, IPv7 enables routers to enforce intent-based policies without relying solely on application-layer controls or external reputation databases.

2. Non-Goals

IPv7 does not perform endpoint-side malware analysis or endpoint integrity verification.

IPv7 does not replace application-layer authentication or authorization mechanisms (e.g., MFA, session management, and access control).

IPv7 does not prevent social engineering or application-layer credential theft attacks.

IPv7 does not mandate a single global identity provider or global identity registry; the protocol assumes provider-operated validation domains.

IPv7 does not specify legal processes for identity disclosure; such processes remain jurisdiction-dependent and outside the scope of this protocol.

3. Problems with IPv4 and IPv6

3.1. Identity Masking via Residential Proxies

IPv4 and IPv6 addresses identify connection points, not people. Residential proxies exploit this by routing malicious traffic through legitimate home IPs, making it difficult for servers to distinguish legitimate residential activity from automated abuse without additional signals or application-layer analysis.

The market for residential proxies generates billions in annual revenue, while victims lack consistent protocol-level mechanisms to authenticate source/provider binding and often rely on application-layer detection and third-party mitigation.

3.2. Stateless Security (Dumb Headers)

IPv4/v6 headers contain routing information but no metadata about sender intent or reputation. Security is always an add-on (firewalls, WAFs) rather than built in.

This design creates high overhead for deep packet inspection and forces applications to implement their own reputation, trust, and policy enforcement, duplicating work across millions of servers.

3.3. Complexity and Human Error

IPv6 addresses (e.g., 2001:0db8:85a3::8a2e:0370:7334) can be difficult for operators to read and reason about directly, increasing the risk of configuration errors and reliance on tooling for interpretation.

This difficulty impairs auditing, troubleshooting, and incident response, making security logs cryptic and preventing operators from quickly understanding "who is attacking us."

3.4. Lack of Native Trust Tiers

In IPv4 and IPv6, packets are typically forwarded without an explicit, standardised trust tier at the network layer. Trust and prioritisation are commonly implemented using local policy and external systems after traffic is received.

DDoS attacks receive the same network resources as legitimate traffic. Routers cannot prioritise based on sender reputation, forcing targets to over-provision bandwidth or rely on external mitigation services.

4. IPv7 Solutions

4.1. Identity-Centric Addressing

IPv7 represents addresses as hierarchical identity strings. An IPv7 source address MUST include sufficient components to support provider validation and policy enforcement.

Format:
[EIT]/service.location.provider.tenant.role.trustlevel.reputationscope.[Origin_Signature]

Example: eit_7f3a9c2b/
web.nyc.exampleisp.home.guest.medium.local.ed25519sig_...

4.2. Eliminating Residential Proxies via Source-Provider Validation

IPv7 defines Source-Provider Validation (SPV) using the Provider ID and Origin Signature components. An IPv7 validating node (typically the first-hop router) MUST verify that the Origin Signature is valid under a public key associated with the asserted Provider ID. If validation fails, the packet MUST be dropped. Providers SHOULD rotate signing keys and deploy key protection mechanisms (e.g., HSMs) to reduce the impact of key compromise.

4.3. Built-In Reputation and Hardware-Level Filtering

By including trust level and reputation scope in the address itself, routers can perform hardware-level filtering. Instead of a software firewall consulting a database, a router can detect an indicator such as "trustlevel.low" and apply local policy (e.g., throttling or quarantine). This can reduce per-packet CPU cost and avoid duplicating reputation logic across endpoints.

4.4. Granular Policy Enforcement

The role and policy fields enable intent-based networking. An IPv7 node MAY indicate that a packet is associated with a specific access intent (e.g., "read-only" or "admin"). If a packet asserts a role that is not authorised for a given path, a policy-enforcing device SHOULD reject the traffic before it reaches the destination service.

4.5. Human-Readable Auditing

IPv4 Example: 192.0.2.1

IPv7 Example: malicious_user/
bot.london.exampleisp.home.guest.0.global

This representation is intended to improve operator attribution during troubleshooting and incident response.

5. Darknet Diaries Case Studies: IPv7 as the Fix

5.1. Episode 172: SuperBox

Consumer network devices (e.g., illicit streaming boxes) may exhibit malicious behaviour, including unsolicited outbound connections, local network scanning, and traffic proxying. Such devices can be used to conceal an attacker’s origin behind a residential connection and to generate abuse traffic from "clean" residential IP space.

In IPv7, the first-hop validating node performs Source-Provider Validation (SPV) over the Provider ID and Origin Signature. Packets that traverse an unauthorised proxy path will fail validation and MUST be dropped. This allows abuse originating from compromised residential environments to be filtered near the source, reducing reliance on application-layer bot detection.

5.2. Episode 128: Gollumfun (Part 1) (Fraud and Identity Abuse)

Underground cybercrime ecosystems enable fraud at scale, including credential theft, account takeover, and carding. These activities are operationally supported by infrastructure that reduces attribution (e.g., use of compromised devices, proxies, and anonymisation techniques).

IPv7's Source-Provider Validation (SPV) and trust/reputation signalling are intended to improve attribution and early filtering of abuse traffic. Where fraud operations rely on unauthorised proxying through residential networks, first-hop validation can reduce the effectiveness of such infrastructure by requiring origin/provider binding for forwarded traffic.

5.3. Episode 110: Spam Botnets

Large-scale botnets can be formed from compromised endpoints and used to send spam or generate abuse traffic at high volume. These botnets commonly rely on weak authentication, unpatched software, and an inability for the network to attribute or rate-limit traffic close to its true source.

IPv7 supports first-hop origin validation (SPV) and explicit rate enforcement. A validating node MAY apply local policy to rate-limit or discard packets based on the Trust/Reputation field and the Rate-Limit Token. This enables abuse traffic to be constrained near its origin and reduces reliance on downstream filtering alone.

5.4. Episodes 45-46: Xbox Underground (Credential Reuse and Lateral Movement)

Attackers frequently use stolen credentials and credential reuse to gain initial access, then pivot laterally inside trusted networks to reach higher-value systems. Network-layer forwarding does not typically account for authenticated role assertions.

IPv7 allows forwarding devices to apply policy based on role/policy signalling carried in the VLIB. Role and policy assertions that affect forwarding decisions SHOULD be integrity-protected so that on-path modification causes validation failure. This supports network-layer containment by preventing packets asserting a low-privilege role (e.g., "guest") from being forwarded to high-privilege destinations (e.g., "admin-only").

6. Additional Motivation from Public Incident Reporting

6.1. Botnets Monetised as Residential Proxy Infrastructure

Recent public reporting describes large botnets composed of compromised consumer devices (e.g., routers, cameras, and low-cost streaming boxes) that are not only used for volumetric DDoS but are also rented out as residential proxy endpoints. Under IPv4 and IPv6, the network layer does not provide an origin-authentication primitive; downstream destinations observe traffic as originating from the compromised subscriber's IP address and cannot distinguish

legitimate use from unauthorised proxying without extensive application-layer detection.

IPv7 is intended to enable earlier filtering by requiring Source-Provider Validation (SPV) and by supporting explicit rate enforcement. A first-hop validating node can drop packets that fail provider binding and can apply local rate policy using the Trust/Reputation field and Rate-Limit Tokens.

6.2. Credential Abuse and the Limits of Source Addressing

Annual incident and breach analyses consistently identify stolen credentials and exploitation of exposed services as dominant initial access vectors. Once valid credentials are obtained, an attacker’s traffic is often indistinguishable at the IP layer from legitimate user traffic.

IPv7’s role/policy signalling and trust-aware forwarding are intended to provide additional containment and prioritisation options at the network layer; however, they are not a replacement for strong authentication and application-layer authorisation.

7. Technical Specification

7.1. IPv7 Packet Header Format

IPv7 packets begin with a fixed 40-byte header for rapid processing in routers, followed by a variable-length identity block (VLIB) containing hierarchical identity information.

7.1.1. Fixed Header Section (40 bytes)

The fixed header contains routing, flow identification, and trust signalling fields for fast-path processing.

7.1.2. Variable-Length Identity Block (VLIB)

The VLIB contains hierarchical identity information including EIT, provider, tenant, service, location, role, and Origin Signature components.

7.2. How Routers Process IPv7 Headers

7.2.1. Three-Stage Processing

Fast Path (Fixed Header): The router examines the Trust/Reputation octet. Packets below a locally configured threshold MAY be rate-limited or discarded, particularly under congestion.

Validation Path: The router performs SPV by verifying the Origin Signature against the asserted Provider ID. Packets that fail validation MUST be dropped.

Routing Path: The router uses provider, location, and role/policy information to select the next hop and apply local forwarding policy.

8. Operational Considerations

8.1. Key Management and Rollover

Operational deployments MUST define key lifecycle procedures for Origin Signature verification, including issuance, rotation, and revocation. Routers SHOULD support caching of provider public keys with explicit freshness bounds. Providers SHOULD support seamless rollover (e.g., overlapping validity windows) to avoid traffic loss during rotation events.

8.2. First-Hop Deployment Model

SPV is expected to be enforced primarily at, or near, the source attachment network (e.g., access edge, CPE gateway, or provider first-hop router). Core and transit networks SHOULD NOT be required to maintain subscriber-specific state to forward IPv7 traffic. Operators MAY deploy SPV selectively (e.g., only for certain roles, tenants, or destination prefixes) during incremental rollout.

8.3. Telemetry and Troubleshooting

To support incident response, implementations SHOULD provide counters and logs for SPV outcomes (e.g., signature failures, key-fetch failures, replay rejections) and SHOULD allow operators to correlate drops with Provider ID, tenant, and role/policy fields.

8.4. Policy and Misconfiguration Risk

Because role/policy signals may influence forwarding, operators MUST consider the risk of misconfiguration causing unintended denial of service. Implementations SHOULD support safe defaults (e.g., treat unknown roles as least privilege) and SHOULD support staged policy rollout with monitoring before enforcement.

8.5. Interconnection Considerations

Inter-domain deployment requires agreement on how Provider IDs are formed and how validation keys are distributed or discovered. Operators SHOULD support multiple authorised Provider IDs for multi-homed environments and SHOULD define local policy for traffic arriving from domains that do not support SPV.

8.6. Manageability Considerations

Implementations SHOULD expose operational configuration and telemetry for SPV and policy decisions via existing management frameworks (e.g., vendor CLIs, NETCONF/RESTCONF data models where available). At a minimum, operators SHOULD be able to configure trust thresholds, rate-limit behaviour, and fallback policy for key unavailability.

9. Deployment and Transition Considerations

9.1. Incremental Deployment Models

IPv7 is expected to be deployed incrementally. Early deployment models include: (1) dual-stack hosts and routers that support IPv7 alongside IPv4/IPv6; (2) IPv7-in-IPv6 (or IPv7-in-IPv4) encapsulation between IPv7-capable edges; and (3) translation gateways that terminate IPv7 and originate IPv4/IPv6 flows toward legacy destinations.

9.2. Coexistence and Negotiation

Where both peers support IPv7, an endpoint SHOULD prefer IPv7 for flows that benefit from origin validation and policy signalling. Name resolution mechanisms are expected to evolve to return both IPv7 and IPv6/IPv4 locators where applicable.

9.3. Naming and Discovery Considerations

This document does not specify new DNS resource record types. A standards-track evolution of IPv7 would be expected to define how IPv7 locators and/or identity strings are represented for name resolution and service discovery.

9.4. Middleboxes, Firewalls, and Translation

Existing middleboxes that rewrite packet headers may interfere with IPv7 if they modify fields covered by the Origin Signature. Deployments SHOULD treat the VLIB as integrity-protected metadata and SHOULD avoid in-path modification.

10. Privacy Model: Hybrid Anonymity

IPv7 is designed to support accountability while limiting identity disclosure to on-path devices. The protocol uses Ephemeral Identity Tokens (EITs) such that intermediate routers can forward traffic and apply coarse-grained policy without requiring access to a subscriber’s long-term identifier.

10.1. Ephemeral Identity Tokens (EIT)

The User Identity field does not contain the user’s legal name or permanent identifier. Instead, it contains an Ephemeral Identity Token (EIT) - a time-bound opaque value renewed per session (typically every 24 hours). Intermediate routers see a "Verified" identity with a "High Trust" score but cannot determine the user’s identity across sessions.

10.2. ISP-Level Verification

The originating provider is assumed to maintain the ability to map an EIT to subscriber records using local operational data (e.g., provisioning, authentication, and accounting logs). The mechanisms and conditions under which such mappings are disclosed to third parties are out of scope for this document and depend on applicable policy and law.

10.3. Reputation without Identification

IPv7 separates conduct from identity. If an EIT is flagged for malicious activity, the reputation of the associated Tenant ID or Provider ID is affected - not the individual user. This creates a system of collective responsibility, incentivising ISPs and network operators to monitor tenant health without inspecting private user data.

10.4. Optional Selective Disclosure

Only the destination server (after mutual authentication) can request the true identity behind an EIT. Users can choose to publish their EIT to public key mapping via distributed PKI (DNS CERT records, blockchain), allowing applications to verify identity without revealing their legal name to network operators.

11. Advanced Security Features

11.1. Quantum-Resistant Cryptography

IPv7 incorporates a quantum-resistant signature option using CRYSTALS-Dilithium (NIST PQC standard). Devices negotiate the strongest mutually supported algorithm, enabling smooth migration from classical to post-quantum cryptography without protocol redesign.

11.2. Multi-Layer Authentication

Layer 1 (Origin Verification): Origin Signature validated at first-hop router.

Layer 2 (Identity Verification): EIT token validated against user’s public key certificate.

Layer 3 (Policy Enforcement): Role/Policy validated against destination ACL.

11.3. Built-In DDoS Mitigation

IPv7 packets include a Rate-Limit Token field. ISPs issue time-limited tokens allowing N packets per second. Routers enforce tokens at the hardware level, preventing botnet traffic from overwhelming targets. Malformed or forged tokens are silently dropped.

11.4. Reputation-Based Filtering

The Trust Level field is dynamically updated. Nodes observing malicious behaviour (failed authentication, port scans, botnet signatures) decrement the Trust Level. Once below threshold (e.g., 50/255), packets are rate-limited or quarantined. Trust can be restored through ISP intervention or automatic recovery.

12. Advanced Routing Features

12.1. Trust-Aware Path Selection

IPv7-BGP routing protocols compute paths based on topological distance and trust scores of intermediate ASes. A path with lower cost may be rejected if intermediate routers have low reputation, preventing traffic hijacking through compromised ASes.

12.2. Policy-Aware Forwarding

Routers make decisions based on Role/Policy, not just destination. A packet with a .guest role cannot reach .admin-only resources. Access control shifts from the application layer to the network layer, reducing server-side processing.

12.3. SLA Enforcement

Service Level Agreements can be enforced at the network layer using IPv7's explicit service identifiers and trust signals.

12.4. Real-Time Media (Voice and Video) Quality Signalling

Interactive voice and video conferencing is sensitive to latency, jitter, and packet loss. IPv7's explicit QoS Class and role/policy fields enable routers to apply local policy (e.g., low-latency forwarding, congestion handling, or admission control) based on authenticated context. Where deployed, trust/reputation and provider binding can also assist operators in prioritising interactive media flows from validated sources during congestion.

12.5. On-Demand Streaming Video Delivery Optimisation

On-demand streaming video remains throughput-driven but is sensitive to rebuffering and sudden congestion. IPv7's explicit service identifier, location, provider/tenant context, and QoS class can enable policy-aware routing decisions such as preferring local caches, applying per-tenant traffic engineering, and providing differentiated treatment for premium or latency-sensitive segments. SPV and trust signalling can help reduce abuse that leverages residential proxy infrastructure to scrape or attack streaming services.

13. IPv4 vs IPv6 vs IPv7 Comparison

A comprehensive comparison table comparing IPv4, IPv6, and IPv7 across key attributes (addressing, security, policy, scalability, and operability) would be included in a full implementation.

14. Use Case: Botnet Attack Mitigation

14.1. Scenario

A malicious actor in Region A wants to perform credential stuffing on a bank in Region B. To bypass the bank's firewall, the attacker rents access to a residential proxy service routing traffic through a legitimate home router in Region C (compromised via weak credentials).

14.2. IPv4/IPv6 Failure

The Packet Path: Attacker (192.0.2.5) - Proxy (198.51.100.14) - Bank

The Mechanism: The proxy server strips the attacker IP and replaces it with a residential IP.

Bank Result: The bank sees a legitimate residential customer.

Detection: Not generally feasible at the network layer using IP addressing alone. The bank typically relies on bot-detection and fraud controls at higher layers (with false positives and false negatives).

14.3. IPv7 Solution

The Packet Header:
[eit_7f3a9c]/web.london.exampleisp.home.guest.low.[sig_ed25519]

Step 1 (Identity Bind): EIT is cryptographically tied to the home router's hardware.

Step 2 (Signature Verification): The attacker cannot generate a valid Origin Signature without the provider's signing key.

Step 3 (The Drop): The first ISP router notices a Provider ID/Origin Signature mismatch. The packet is dropped at the edge. The attacker never reaches international cables or the bank.

14.4. Cost Shift

In IPv4/v6, many defensive controls against proxy abuse are applied at the destination (or by downstream mitigation providers). IPv7 shifts some enforcement capability toward the source attachment network by enabling first-hop validation and policy-based dropping. This can increase attacker cost in scenarios that depend on unauthorised proxying through residential networks.

15. Use Case: Interactive Conferencing (Voice and Video)

15.1. Scenario

An enterprise user participates in an interactive voice/video conference (e.g., Microsoft Teams or Zoom) from a residential access network while other household devices concurrently generate bulk traffic (e.g., software updates or large downloads). The conferencing application requires low latency and low jitter to maintain conversational quality, but the access uplink and upstream aggregation links are intermittently congested.

15.2. IPv4/IPv6 Limitations

Under IPv4 and IPv6, routers generally forward packets without an authenticated statement of application intent. QoS treatment is often inferred from DSCP markings or local policy, but DSCP is not consistently preserved or honoured across networks. As a result, conferencing traffic competes with bulk flows during congestion.

15.3. IPv7 Approach

With IPv7, an endpoint MAY indicate interactive media intent using the QoS Class field and MAY include a role/policy appropriate for real-time collaboration. Where local policy permits, the first-hop validating node can apply Source-Provider Validation (SPV) and then provide preferential queueing or congestion handling for validated real-time flows (e.g., lower latency treatment) while rate-limiting or deprioritising bulk traffic.

15.4. Operational Notes

Operationally, real-time flows benefit from stable queueing and bounded latency. Implementations SHOULD expose counters for drops and queueing outcomes by QoS Class and role/policy to support troubleshooting. Deployments MUST consider privacy implications of any per-application classification.

16. Use Case: On-Demand Streaming Video

16.1. Scenario

A user streams on-demand video (e.g., Netflix or Amazon Prime Video) from a nearby cache while other traffic on the access network is bursty. The streaming service adapts bitrate based on measured throughput and rebuffering events. During peak periods, some paths to caches exhibit congestion or intermittent loss, reducing quality of experience (QoE).

16.2. IPv4/IPv6 Limitations

In IPv4 and IPv6 deployments, cache selection and traffic engineering commonly rely on DNS-based steering, anycast, and proprietary telemetry. The IP layer itself provides limited authenticated context about service intent, tenant boundaries, or delivery class.

16.3. IPv7 Approach

IPv7 can provide explicit, coarse-grained delivery context via the Service ID, location, provider/tenant components, and QoS Class. Operators MAY use this context to implement policy-aware routing such as preferring local caches, applying per-tenant traffic engineering, or selecting congestion-avoidant paths during peak demand. Where streaming abuse relies on unauthorised proxying through residential networks, SPV and trust/reputation signalling can provide additional inputs for early filtering and rate enforcement near the source attachment network.

16.4. Operational Notes

QoE (quality of experience) is impacted by sustained throughput, start-up delay, and rebuffering frequency. Operators SHOULD consider how QoS Class and service identifiers interact with congestion management to avoid starving other traffic.

17. Implementation Considerations

17.1. Signature Verification Performance

Ed25519 signature verification at line rate is feasible on modern hardware. CPUs with SHA-3 and Ed25519 instructions (AVX-512, ARM SVE) verify approximately 100,000 signatures per second per core. This is sufficient for backbone routers processing millions of packets per second using dedicated crypto accelerators (e.g., Intel QuickAssist, ARM TrustZone).

17.2. Trust Level Updates via Gossip Protocol

Rather than updating Trust Levels on every malicious packet (massive state overhead), IPv7 uses a gossip protocol where nodes periodically exchange reputation data. Low-trust addresses are marked in Bloom filters for fast negative lookups, reducing memory footprint.

17.3. EIT Generation and Rotation

EITs are generated using HASH(user's_long_term_public_key + session_nonce + ISP_signing_key). ISPs can verify EIT validity without learning which specific user generated it. EITs are valid for 24 hours; users request new tokens via an authenticated channel.

17.4. Backward Compatibility and Incremental Deployment

Detailed deployment and transition considerations are discussed in Section 10. Implementers SHOULD ensure that signed fields are treated as integrity-protected metadata and that any translation or tunnelling preserves the deployment model's security assumptions.

17.5. Scalability Considerations

IPv7 is intended to support Internet-scale routing and forwarding by minimising per-flow state in the network core. SPV is expected to occur primarily at, or near, the first hop; intermediate domains are not required to perform subscriber-specific validation. Trust/ reputation processing is designed for fast-path evaluation using compact fields and locally configured policy thresholds.

17.6. Resiliency and High Availability

Deployments MUST consider failure modes for validation and key management. If SPV cannot be performed (e.g., key retrieval failure), an implementation SHOULD support a locally configurable policy that defines whether traffic is dropped, rate-limited, or forwarded with reduced trust.

17.7. AI/ML-Assisted Policy and Telemetry (Non-Normative)

This document does not require artificial intelligence (AI) or machine learning (ML). However, operators MAY use ML models to improve operational outcomes, for example: predicting QoE (quality of experience) degradation for interactive media, detecting anomalous traffic patterns that correlate with proxy abuse, and tuning local policy thresholds for trust/reputation and rate enforcement.

17.8. Futuristic Extensions (Non-Normative)

Intent negotiation: A future extension could allow endpoints to negotiate acceptable QoS classes and policy constraints without relying on out-of-band configuration.

In-network telemetry signals: Standardised metadata for congestion exposure or path health could improve troubleshooting and closed-loop traffic engineering without payload inspection.

Privacy-preserving feature export for ML-based operations: A future extension could define standardised, coarse-grained telemetry features for ML-assisted anomaly detection and QoE prediction without requiring payload inspection or stable user identifiers.

Federated learning for cross-domain models: A future extension could enable federated learning approaches in which multiple operators train shared models for anomaly detection or QoE prediction while keeping raw telemetry local.

Multipath-aware identity binding: A future extension could support binding a single identity to multiple simultaneous paths (e.g., multi-access) while preserving provider validation properties.

Programmable policy objects: A constrained policy object format could enable richer, interoperable policy enforcement while remaining bounded for router fast paths.

Cryptographic agility: Explicit algorithm agility beyond the initial signature set to support post-quantum transitions and future primitives.

Encrypted metadata: Selective encryption of portions of the VLIB could reduce metadata leakage while retaining verifiability where required.

18. Security Considerations

18.1. Signature Key Compromise

If a provider signing key is compromised, an attacker could forge Origin Signatures for that provider. Providers SHOULD protect signing keys using appropriate controls (e.g., hardware security modules) and SHOULD define rapid revocation and rollover procedures. Validating nodes SHOULD support key freshness bounds and SHOULD be able to reject signatures from revoked keys based on locally configured policy.

18.2. Replay Attacks

Attackers could capture and replay valid packets. Mitigation: IPv7 packets include a timestamp and a nonce. Routers maintain a short-lived cache of (packet_hash, nonce) pairs and reject duplicates.

18.3. Role Escalation

An attacker could attempt to assert a higher-privilege role (e.g., "admin"). Role/policy fields that affect forwarding or access decisions SHOULD be integrity-protected (e.g., covered by the Origin Signature) so that unauthorised modification causes validation failure. Destination systems SHOULD continue to perform application-layer authorisation independent of network-layer signalling.

18.4. Privacy Leakage via Traffic Analysis

An adversary could correlate behavioural patterns (timing, volume) to de-anonymise users. Mitigation: Users employ packet padding and random inter-packet delays. ISPs provide "privacy mode" where multiple users share a single EIT during a session.

18.5. Trust Level Depletion Attack

An attacker might attempt to influence trust signals to cause legitimate traffic to be deprioritised. Implementations that accept trust updates from the network MUST authenticate and authorise such updates. Nodes SHOULD apply rate-limits and sanity checks to reputation inputs and SHOULD support administrative override procedures for recovery from erroneous or malicious downgrades.

19. IANA Considerations

This document makes no request of IANA at this stage. If the mechanisms described here are progressed on the standards track, IANA actions would be required to allocate protocol code points and establish registries to support interoperability, including allocation of an IP version number for IPv7, creation of registries for Next Header types, Role/Policy encodings, and Signature Algorithms.

20. References

20.1. Normative References

RFC 2119 - Key words for use in RFCs to Indicate Requirement Levels

RFC 8200 - Internet Protocol, Version 6 (IPv6) Specification

RFC 8032 - Edwards-Curve Digital Signature Algorithm (EdDSA)

FIPS 204 - Module-Lattice-Based Digital Signature Standard (CRYSTALS-Dilithium)

20.2. Informative References

Darknet Diaries Podcast - Episodes 172, 128, 110, 45, 46, 13

NIST SP 800-207 - Zero Trust Architecture

NIST Post-Quantum Cryptography Standardization Project

RFC 9330 - Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture

EU Regulation (EU) 2016/679 - General Data Protection Regulation (GDPR)

21. Conclusion

This document describes IPv7, an identity-carrying, origin-validating network-layer protocol intended to reduce abuse enabled by unauthenticated source addressing and residential proxy infrastructure. IPv7 introduces an explicit identity structure (VLIB), an Origin Signature for Source-Provider Validation, and fields that support policy and reputation signalling.

With the rapid proliferation of IoT devices running Linux and Android-based systems, the ability to modify network stacks at the kernel level creates both opportunities and challenges. IPv7 provides a framework for these devices to authenticate their origin, signal their intent and policy requirements, and participate in a trustworthy network ecosystem.

Further work is required to evaluate deployability, incremental transition mechanisms, and interoperability with existing IP networks, and to refine the protocol into a form suitable for standards-track consideration and eventual RFC publication.

Author’s Address

Arunkumar Subbiah
Independent
Email: arunkumar.subbiah@apexadversary.com