Internet Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: 22 October 2026

N. Matsukami
Independent Researcher

20 April 2026

IPv4.5: A Locator/Identifier-Separated Extension to IPv4 with Post-Quantum Session Security

draft-matsukami-intarea-ipv45-00

Abstract

The global IPv4 address space was exhausted at the IANA level in 2011, and Carrier-Grade NAT (CGN) has since served as the primary operational workaround. CGN preserves connectivity but violates the end-to-end principle, complicates application development, and introduces substantial operational overhead.

This document specifies IPv4.5, a pragmatic extension of IPv4 that introduces a 96-bit address space organized as a Locator/Identifier separation: a 32-bit IPv4 Locator, a 16-bit Site Identifier, and a 48-bit Endpoint Identifier. IPv4.5 packets are encapsulated in UDP (port 4242) for transparent transit through existing IPv4 routers, NAT devices, and firewalls without requiring infrastructure changes.

Session security is established through a hybrid post-quantum key exchange combining ML-KEM-768 [FIPS203] and X25519 [RFC7748], performed once per session. Subsequent data protection uses symmetric AEAD ciphers (AES-256-GCM or ChaCha20-Poly1305).

The design enforces strict separation of concerns across four independent planes: data, control, identity, and cryptographic. Higher-level functions such as semantic routing and self-sovereign identity are explicitly out of scope for this specification.

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/.

Copyright Notice

Table of Contents

1. Introduction

1.1. Problem Statement

The IPv4 address pool managed by IANA was exhausted in February 2011. Subsequent exhaustion at each Regional Internet Registry (RIR) has forced network operators to rely on Carrier-Grade NAT (CGN) [RFC6598] as a large-scale address sharing mechanism. While CGN preserves IPv4 reachability, it introduces the following well-documented problems:

1. Violation of the end-to-end principle: applications cannot rely on a globally unique, reachable address for each endpoint.

2. Peer-to-peer impairment: NAT traversal for real-time communication requires additional infrastructure such as STUN, TURN, and ICE.

3. Logging complexity: multiple subscribers share a single public IPv4 address at any given time, complicating forensic and legal demands.

4. Single point of failure: CGN appliances represent concentrated failure domains in operator networks.

IPv6 [RFC8200] was designed to resolve address exhaustion definitively, but its adoption has progressed slowly after more than two decades of availability. As of 2026, IPv6 accounts for approximately 45-50% of Internet traffic by volume, and growth has decelerated. Key barriers include dual-stack operational cost, the absence of day-one benefit for early adopters, and a large installed base of IPv4-only applications and embedded devices.

Furthermore, neither IPv4 nor IPv6 was designed with post-quantum cryptography (PQC) in mind. Future quantum computers capable of running Shor’s algorithm would break classical elliptic-curve and RSA key exchange mechanisms, enabling "harvest now, decrypt later" attacks against recorded traffic.

1.2. Design Goals

IPv4.5 is designed to satisfy the following goals:

G1 (Address Exhaustion): Provide a globally unique address space sufficient for 6G/IoT requirements, without relying on address sharing.

G2 (Transparent Transit): Deploy immediately over existing IPv4

infrastructure, with no changes required to intermediate routers, NAT devices, or firewalls.

G3 (Day-One ROI): Provide immediate, measurable benefits to early adopters — including endpoint privacy and quantum-resistant session security — independent of the deployment fraction of IPv4.5 on the Internet.

G4 (Quantum Resistance): Resist passive traffic recording and future decryption by quantum computers, using NIST-standardized post-quantum key encapsulation [FIPS203].

G5 (Layered Design): Limit the network-layer specification to addressing and forwarding. Advanced functions are delegated to separate planes or companion specifications.

G6 (Privacy): Support per-session ephemeral Endpoint Identifiers to prevent cross-session endpoint tracking.

G7 (Political Acceptability): Require no central authority, mandatory validation service, or built-in surveillance mechanism.

1.3. Non-Goals

The following are explicitly OUT OF SCOPE for this document:

* Content-based or semantic routing at the network layer.

* Per-packet intent or policy fields in the data plane.

* Mandatory Decentralized Identifier (DID) or self-sovereign identity requirements.

* Per-packet post-quantum digital signatures.

* Traffic anonymization equivalent to Tor or mixnet overlays.

1.4. Relationship to Existing Work

IPv4.5 is inspired by, but distinct from, the Locator/Identifier Separation Protocol (LISP) [RFC9300]. LISP is an overlay tunneling protocol designed to scale BGP routing; IPv4.5 defines a new wire format with explicit addressing semantics and integrated session security. IPv4.5 borrows the Map-Server/Map-Resolver architecture of LISP [RFC9301] as one supported implementation of its EID-Locator Mapping Service.

The UDP encapsulation approach is analogous to QUIC [RFC9000], which demonstrated that a new transport protocol can achieve broad deployment by operating over existing UDP/IP infrastructure. IPv4.5 is complementary to IPv6 and does not compete with it.

1.5. Requirements Language

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Terminology

Locator
: A 32-bit IPv4 address that identifies the network attachment point of an IPv4.5 anchor. Locators are allocated from existing IPv4 address pools and are routed by unmodified IPv4 infrastructure using standard routing protocols (BGP, OSPF, IS-IS).

Site Identifier (Site ID)
: A 16-bit value allocated by the operator of an IPv4.5 anchor. Identifies a logical administrative domain (e.g., a data center, campus, or enterprise network) behind a single Locator.

Endpoint Identifier (EID)
: A 48-bit value that uniquely identifies an endpoint within a Site. EIDs may be statically assigned (persistent) or dynamically generated per session (ephemeral; see eNID).

IPv4.5 Address
: The 96-bit triple [Locator | Site ID | EID] that globally and uniquely identifies an IPv4.5 endpoint.

Anchor
: An IPv4.5-aware node holding one or more Locators, responsible for encapsulating and decapsulating IPv4.5 packets to and from standard UDP/IPv4 packets for endpoints in its site.

Edge Translator
: A gateway that enables communication between IPv4.5 endpoints and IPv4-only endpoints, performing address and header conversion in both directions.

EID-Locator Mapping Service (ELMS)
: A distributed service that maps Endpoint Identifiers to their current Locators, enabling resolution of EIDs not directly resolvable via DNS.

Ephemeral Network Interface Designator (eNID)
: A per-session randomly generated EID used to prevent cross-session endpoint tracking. Modeled after IPv6 Privacy Extensions [RFC8981] and Semantically Opaque Interface Identifiers [RFC7217].

Intent Negotiation Protocol (INP)
: A control-plane protocol through which endpoints request session-level service characteristics (latency class, bandwidth class) from the network operator’s policy controller.

AX Record
: A new DNS resource record type conveying a 96-bit IPv4.5 address, analogous to the A record (IPv4) and AAAA record (IPv6).

3. Architecture Overview

3.1. Four-Plane Separation

IPv4.5 is organized into four independent planes, each with a well-defined scope and interface boundary:

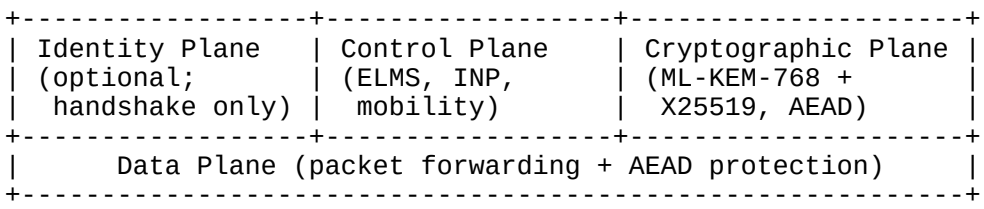

Figure 1: IPv4.5 Four-Plane Architecture

Data Plane (mandatory)
: Forwards IPv4.5 packets between anchors and endpoints. Applies AEAD encryption and decryption using symmetric session keys. Stateless except for per-session symmetric key material.

Control Plane (mandatory)
: Resolves EIDs to Locators, negotiates session policies, and handles endpoint mobility. Implemented as a separate daemon or process. MUST NOT be on the forwarding fast path.

Identity Plane (optional)
: Authenticates endpoints during the initial handshake using one of four supported identity models. MUST NOT be referenced by the data plane after handshake completion.

Cryptographic Plane (mandatory)
: Performs hybrid post-quantum key exchange to establish session keys. Provides derived key material to the data plane for symmetric encryption.

3.2. Operating Modes

IPv4.5 defines four operating modes for staged deployment:

Mode A (Pass-Through)
: Pure IPv4 operation. No IPv4.5 processing. All existing IPv4 traffic continues unmodified.

Mode B (Translated)
: An IPv4.5 endpoint communicates with a legacy IPv4-only endpoint through an Edge Translator. REQUIRED for deployments that must interoperate with IPv4-only peers.

Mode C (Encapsulated)
: An IPv4.5 endpoint communicates with another IPv4.5 endpoint using UDP encapsulation over standard IPv4 (protocol 17, port 4242). No changes are required to intermediate routers. This is the baseline operational mode for Phase 1 and Phase 2 deployment. All IPv4.5 implementations MUST support Mode C.

Mode D (Native, OPTIONAL)
: An IPv4.5 endpoint communicates with another IPv4.5 endpoint using IP Protocol 144 directly, without UDP encapsulation. Requires router firmware support. Intended for Phase 3 deployment.

4. Addressing

4.1. Address Format

An IPv4.5 address is 96 bits wide, consisting of three contiguous fields:

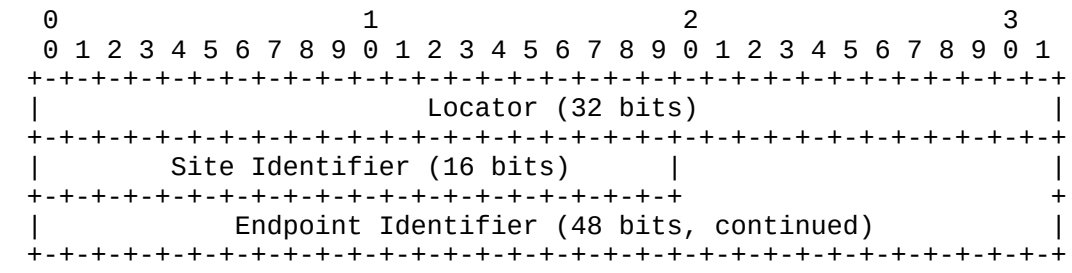

Figure 2: IPv4.5 96-bit Address Format

The total address space is 2^96 ≈ 7.9×10^28, sufficient for all projected 6G, IoT, and cloud workloads through at least 2050 by multiple orders of magnitude.

4.2. Address Allocation

Locator
: IPv4.5 reuses existing IPv4 address allocations from IANA and the Regional Internet Registries. No new global registry or allocation process is required. BGP announcements for Locators are identical to existing IPv4 BGP practice [RFC4760].

Site ID
: The anchor operator allocates Site IDs within its Locator's space. No registration with any external authority is required.

Endpoint ID
: EIDs may be assigned by an operator-managed allocation service (analogous to DHCP) or generated autonomously by the endpoint (analogous to IPv6 SLAAC). Implementations generating EIDs autonomously MUST use a cryptographically secure random number generator conforming to [RFC4086].

4.3. Textual Representation

The canonical textual representation of an IPv4.5 address is:

```
<dotted-decimal-ipv4>::<hex-site-id>::<hex-eid>
```

where the Endpoint Identifier is written as three 16-bit groups separated by colons. Leading zeros within each group MAY be omitted. The double-colon (::) serves as a field separator and MUST NOT be interpreted using IPv6 zero-compression rules.

Example:

```
192.0.2.1::1234::DEAD:BEEF:0001

Locator:      192.0.2.1
Site ID:      0x1234
Endpoint ID:  0xDEADBEEF0001
```

An implementation parsing an IPv4.5 address MUST reject strings that do not contain exactly two "::" separators.

4.4. Reserved Addresses

The following addresses are reserved and MUST NOT be assigned to operational endpoints:

Pattern	Purpose
0.0.0.0::0::0:0:0	Unspecified address
127.0.0.1::*::*	Loopback (all Site IDs and EIDs)
224.0.0.0/4::*::*	Multicast Locator range
LOCATOR::FFFF::*	Site-local scope
LOCATOR::FFFE::*	Link-local scope
LOCATOR::FFFF::FFFF:*	IPv4-Embedded EID (compatibility)

Table 1: IPv4.5 Reserved Addresses

For Edge Translator (Mode B) operation, a pure IPv4 endpoint with address A.B.C.D behind translator at address T is represented as:

```
T::FFFF::FFFF:<hex-A><hex-B>:<hex-C><hex-D>

Example: IPv4 198.51.100.42 behind translator 203.0.113.1:
203.0.113.1::FFFF::FFFF:C633:642A
```

4.5. Ephemeral Endpoint Identifiers (eNIDs)

To mitigate cross-session endpoint tracking, implementations SHOULD use a per-session randomly generated EID (Ephemeral Network Interface Designator, eNID). An eNID:

* MUST be generated using a cryptographically secure random number generator conforming to [RFC4086].

* SHOULD be regenerated at the start of each new session or after a configurable interval (recommended default: 24 hours).

* MUST NOT be reused across sessions to distinct destination endpoints.

* SHOULD be statistically indistinguishable from a static EID to on-path observers.

5. Packet Format

5.1. IPv4.5 Fixed Header

The IPv4.5 fixed header is exactly 20 bytes. When operating in Mode C, the Locator fields are carried in the outer IPv4 header and are NOT repeated in the IPv4.5 header.

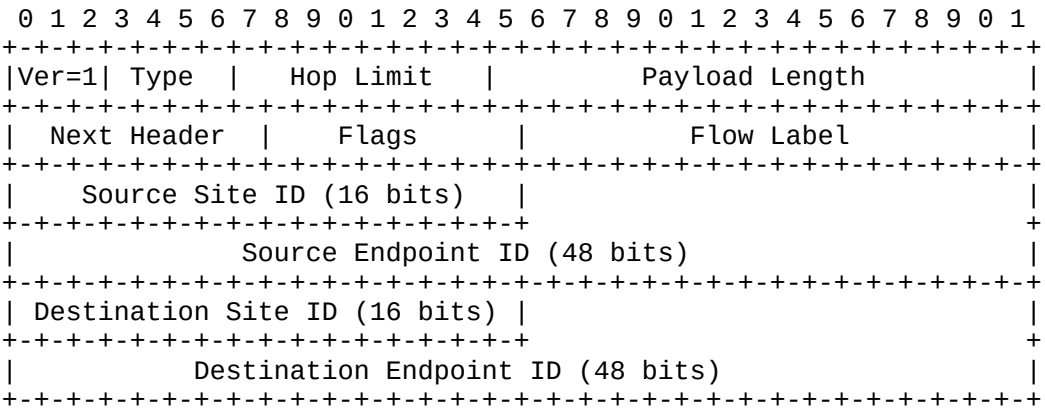

Figure 3: IPv4.5 Fixed Header (20 bytes)

5.2. Field Descriptions

Version (4 bits)
: MUST be set to 1 for this version of IPv4.5. Packets with Version ≠ 1 MUST be silently discarded.

Type (4 bits)
: Identifies the purpose of the packet:

Value	Name	Description
0	DATA	Carries upper-layer payload
1	HANDSHAKE	Session establishment and key exchange
2	CONTROL	Control-plane messages (ELMS, INP)
3	ECHO	Echo request/reply (diagnostic)
4	ERROR	Reports error conditions to the sender
5	KEEPALIVE	NAT/firewall session keepalive
6-14	—	Unassigned; see Section 14.5
15	—	Reserved; MUST NOT be used

Table 2

Hop Limit (8 bits)
: Initialized to 64 by sending implementations. Each IPv4.5-aware forwarding node MUST decrement this field by 1. See Section 6.3.

Payload Length (16 bits)
: Length in bytes of all content following the 20-byte fixed IPv4.5 header, including any extension headers.

Next Header (8 bits)
: Identifies the protocol of the immediately following header or upper-layer payload. Uses IANA-assigned Internet Protocol Number values.

Flags (8 bits)

Bit	Symbol	Meaning
0 (MSB)	E	Payload is AEAD-encrypted and authenticated
1	F	Fragmentation extension header is present
2	X	Additional extension headers are present
3	M	More fragments follow (used with F)
4	C	Compression (reserved for future use)
5	N	NAT traversal extension present
6-7	—	Reserved; MUST be sent as zero

Table 3

Flow Label (16 bits)
: Identifies a traffic flow for ECMP load distribution and QoS treatment. When assigned by INP (Section 8.2), reflects the negotiated SLA class. Otherwise, MUST be set to a pseudorandom value.

5.3. Mode C Encapsulation

In Mode C, an IPv4.5 packet is encapsulated in UDP [RFC0768] over IPv4 [RFC0791]:

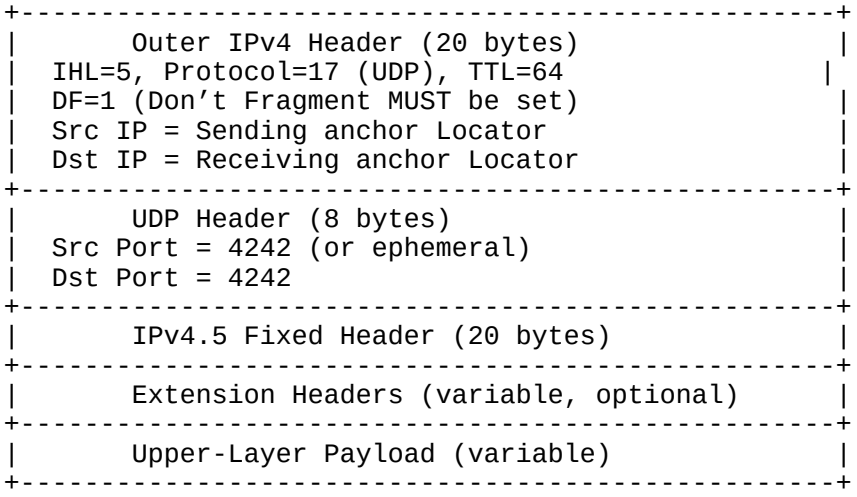

Figure 4: IPv4.5 Mode C Packet Structure

The outer IPv4 DF bit MUST be set to enable Path MTU Discovery (Section 6.1). The outer IP TTL SHOULD be set to 64.

5.4. Extension Headers

Extension headers provide optional functionality. They are chained using the Next Header field and MUST be processed in the order they appear.

Next Header	Extension Type	Description
0x00	Hop-by-Hop Options	Options processed by every forwarding node
0x01	Routing	Explicit source routing via designated anchors
0x02	Mobility	Locator change notification for mobile endpoints
0x03	Fragmentation	Source-originated fragmentation and reassembly
0x04	Authentication	Handshake signature and certificate chain
0x05	ESP-X	AEAD-encrypted and authenticated payload container
0x06	Destination Options	Options processed only by the destination
0xFF	No Next Header	No content follows

Table 4: IPv4.5 Extension Header Types

5.5. Fragmentation Extension Header

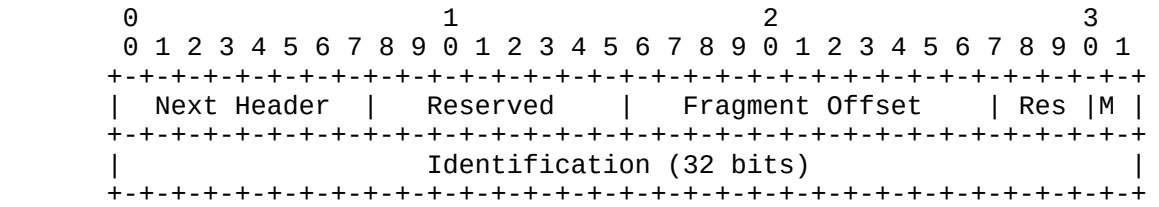

Figure 5: IPv4.5 Fragmentation Extension Header (8 bytes)

Fragment reassembly MUST time out after 60 seconds. Incomplete fragment sets MUST be silently discarded.

6. Forwarding

6.1. Path MTU Discovery

All IPv4.5 implementations MUST support Path MTU Discovery (PMTUD) as defined in [RFC1191] and [RFC8201]. Implementations SHOULD also implement Packetization Layer Path MTU Discovery (PLPMTUD) [RFC4821] as a fallback when ICMP messages are filtered.

The minimum IPv4.5 Path MTU is 1280 bytes. In Mode C, the minimum effective payload MTU is:

1280 - 20 (outer IPv4) - 8 (UDP) - 20 (IPv4.5) = 1232 bytes

The Don’t Fragment (DF) bit MUST be set in the outer IPv4 header. When a router drops an oversized packet, it MUST return an ICMP Fragmentation Needed (Type 3, Code 4) message to the sending anchor. The anchor MUST propagate this information to the originating endpoint via an IPv4.5 ERROR message (Type 4, Error Code MTU_EXCEEDED).

6.2. Fragmentation

Source endpoints MUST fragment packets that exceed the discovered Path MTU using the IPv4.5 Fragmentation Extension Header (Next Header 0x03). Intermediate nodes MUST NOT fragment IPv4.5 packets.

6.3. Hop Limit Processing

When an IPv4.5-aware node forwards a packet, it MUST decrement the Hop Limit by 1. If the resulting value is zero:

1. The node MUST NOT forward the packet.
2. The node SHOULD send an IPv4.5 ERROR message (Type 4, Error Code HOP_LIMIT_EXCEEDED) to the source Locator.
3. The packet MUST be discarded.

7. Address Resolution

7.1. DNS AX Records

A new DNS resource record type, AX, is defined to carry IPv4.5 addresses. An AX record contains a 96-bit IPv4.5 address in network byte order and supports standard TTL semantics identical to A and AAAA records.

```
host.example.com.  3600  IN  AX    192.0.2.1::0100::DEAD:BEEF:0001
host.example.com.  3600  IN  A     192.0.2.1
host.example.com.  3600  IN  AAAA  2001:db8::1
```

When A, AAAA, and AX records are all present, endpoints SHOULD use a Happy Eyeballs v3 algorithm (extending [RFC8305]) with AX as the highest-preference address family.

7.2. EID-Locator Mapping Service (ELMS)

For EIDs not directly resolvable via DNS, ELMS provides Locator resolution. ELMS is a distributed, decentralized service. Three implementation architectures are defined:

Architecture 1 (LISP-Compatible)
: Follows the LISP Map-Server/Map-Resolver architecture [RFC9300] [RFC9301].

Architecture 2 (DHT-Based)
: Implemented as a Kademlia-based [KADEMLIA] distributed hash table. No central server is required.

Architecture 3 (BGP SAFI)
: EID-to-Locator mappings distributed via BGP using SAFI 144 (see Section 14.4). Suitable for large ISP deployments.

ELMS resolution MUST be performed by the control-plane daemon. Results MUST be cached locally. ELMS queries MUST NOT block the forwarding fast path.

8. Control Plane

8.1. Session Establishment Overview

IPv4.5 session establishment follows these steps:

1. The initiating endpoint resolves the target’s Locator via DNS (AX record) or ELMS.

2. A HANDSHAKE packet (Type 1, Subtype 0x01 Client Hello) is sent to the target’s Locator.

3. The responder processes the Client Hello and sends a HANDSHAKE reply (Subtype 0x02 Server Hello).

4. Both endpoints independently derive identical symmetric session keys using HKDF (see Section 10.3).

5. Subsequent DATA packets (Type 0) are encrypted using the derived keys.

8.2. Intent Negotiation Protocol (INP)

Endpoints MAY request specific session service characteristics through INP. INP operates at the session level only. Per-packet intent fields in the data plane are explicitly prohibited by this specification.

Class	Description	Authorization
BEST_EFFORT	Default forwarding treatment	None required
LOW_LATENCY	Expedited Forwarding treatment	Operator policy
HIGH_BANDWIDTH	Guaranteed bandwidth reservation	Operator policy
DETERMINISTIC	Hard real-time guarantee	Pre-arranged contract
BACKGROUND	Lower-than-default priority	None required

Table 5: INP SLA Classes

8.3. Endpoint Mobility

When an endpoint changes its network attachment point (Locator), it MUST:

1. Generate a new IPv4.5 address using the same EID under the new Locator.

2. Send a CONTROL message (Mobility Notification) to active peers, conveying the new Locator.

3. Update its ELMS registration with the new Locator.

Existing session keys remain valid across a Locator change. Re-authentication is NOT required.

9. Identity Plane

Identity authentication in IPv4.5 is OPTIONAL. When used, it occurs exclusively during the HANDSHAKE phase. The data plane MUST NOT carry identity assertions after handshake completion.

9.1. Anonymous Mode (Default)

When no identity model is selected, the session uses only ephemeral EIDs (eNIDs) and ephemeral cryptographic key pairs. Anonymous mode is the DEFAULT and MUST be supported by all implementations.

9.2. X.509 PKI Mode

Endpoints authenticate using X.509 certificates [RFC5280] compatible with the existing Web PKI. This mode is RECOMMENDED for server-to-client authentication in public-facing services.

9.3. DID Mode

Endpoints authenticate using W3C Decentralized Identifiers [W3C-DID]. DID documents are resolved from the decentralized identifier infrastructure; no certificate authority is required. This mode is OPTIONAL.

9.4. Pre-Shared Key (PSK) Mode

Endpoints sharing a pre-established symmetric secret authenticate using PSK mode. PSK MUST be combined with the hybrid PQC key exchange (Section 10) to preserve forward secrecy. RECOMMENDED for constrained IoT devices and closed enterprise deployments.

10. Cryptographic Plane

10.1. Hybrid Post-Quantum Handshake

IPv4.5 session establishment uses a hybrid key exchange combining:

* *X25519* [RFC7748]: Classical elliptic-curve Diffie-Hellman, providing security against classical computers.

* *ML-KEM-768* [FIPS203]: NIST-standardized lattice-based Key Encapsulation Mechanism, providing security against quantum computers capable of running Shor’s algorithm.

The hybrid construction ensures session key security as long as at least one of the two component algorithms remains unbroken.

10.2. Handshake Message Formats

10.2.1. Client Hello (Type 1, Subtype 0x01)

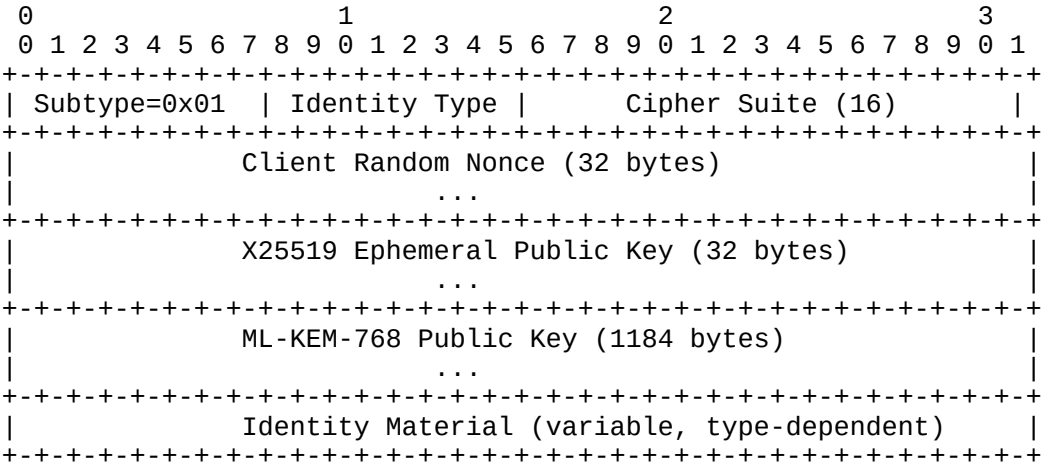

Figure 6: IPv4.5 Client Hello Message

The total Client Hello size is approximately 1252 bytes before identity material, which exceeds the minimum Mode C MTU of 1232 bytes. Sending implementations MUST fragment the Client Hello using the Fragmentation Extension Header when the discovered path MTU is insufficient.

10.2.2. Server Hello (Type 1, Subtype 0x02)

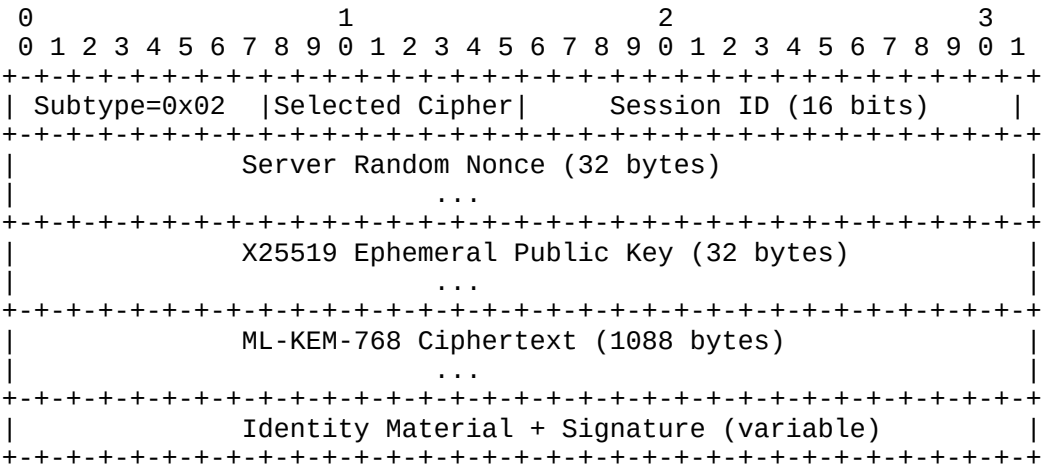

Figure 7: IPv4.5 Server Hello Message

10.3. Key Derivation

After the handshake exchange, both parties independently compute:

```
shared_x25519 = X25519(local_eph_priv, remote_eph_pub)     [32 bytes]
shared_kem    = ML-KEM-768.Decapsulate(ct, local_mlkem_priv) [32 bytes]

ikm  = shared_x25519 || shared_kem                          [64 bytes]
salt = client_nonce  || server_nonce                        [64 bytes]

tx_key = HKDF-SHA-256(ikm, salt, "IPv4.5-C2S-KEY-v1", 32)  [32 bytes]
rx_key = HKDF-SHA-256(ikm, salt, "IPv4.5-S2C-KEY-v1", 32)  [32 bytes]
```

The HKDF construction follows [RFC5116] and [RFC8446] Section 7.1. Labels are fixed ASCII strings. Implementations MUST verify that both X25519 and ML-KEM operations succeed before proceeding to data exchange.

10.4. Supported Cipher Suites

Code	Cipher Suite	Recommendation
0x01	AES-256-GCM	RECOMMENDED for hardware with AES-NI
0x02	ChaCha20-Poly1305	RECOMMENDED for software-only implementations
0x03-0xFF	—	Reserved

Table 6: IPv4.5 Cipher Suites

Both suites provide Authenticated Encryption with Associated Data (AEAD) as specified in [RFC5116].

10.5. Nonce Management

The AEAD nonce (96 bits) for DATA packet protection is constructed as:

```
nonce = session_id (16) || reserved (16) || sequence_no (64)
```

The sequence number MUST start at 1 and MUST be incremented by 1 for each DATA packet transmitted. Implementations MUST initiate key rotation before the sequence number reaches 2^32. Nonces MUST NEVER be reused under the same session key.

10.6. Forward Secrecy and Key Rotation

All session key material is derived from ephemeral key pairs generated fresh for each session. Long-term identity keys are used only to authenticate the handshake and are never used for encryption.

Active sessions MUST be re-keyed upon whichever of the following thresholds occurs first:

* 3600 seconds (1 hour) of elapsed session time; or

* 2^32 DATA packets sent under the current session key.

Key material from the previous epoch MUST be explicitly zeroized using compiler-barrier-protected routines before being overwritten.

11. Backward Compatibility Considerations

IPv4.5 provides *transparent IPv4 transit*, not full backward compatibility.

Mode A traffic (pure IPv4) is entirely unaffected. Mode C traffic traverses unmodified IPv4 routers as standard UDP/4242 traffic. Firewalls blocking non-standard UDP ports may prevent IPv4.5 sessions. Implementations SHOULD attempt UDP/4242 first, and if blocked, retry using UDP/443.

Operators deploying IPv4.5 SHOULD configure firewalls to permit inbound and outbound UDP/4242 during the transition period.

12. Privacy Considerations

A static EID enables correlation of an endpoint’s sessions across time and network positions, analogous to MAC address tracking. Implementations SHOULD use ephemeral EIDs (eNIDs) for user-facing applications. See Section 4.5.

The Locator is carried in cleartext in the outer IPv4 header. The Site ID is carried in cleartext in the IPv4.5 fixed header. Both fields are observable to all on-path entities. No protection of Locator or Site ID confidentiality is provided by this specification.

IPv4.5 does not protect against traffic analysis based on timing, packet sizes, or communication volumes. Applications requiring traffic analysis resistance SHOULD use Tor or mixnet overlays operating over IPv4.5 sessions.

13. Security Considerations

13.1. Threat Model

This specification is designed to protect against: passive eavesdropping (including by future quantum computers via "harvest now, decrypt later"); active injection, modification, or replay of IPv4.5 packets; endpoint impersonation during session establishment; and session hijacking.

This specification does NOT protect against: compromised endpoints; physical hardware access; legal compulsion of key material; traffic analysis; or vulnerabilities in cryptographic implementations.

13.2. Protections Provided

Confidentiality: All DATA packets are protected by AEAD using session keys derived from the hybrid PQC handshake.

Integrity: AEAD authentication detects and causes rejection of any modified, truncated, or corrupted DATA packet.

Replay Protection: Monotonically increasing sequence numbers combined with AEAD authentication prevent replay.

Forward Secrecy: Ephemeral key pairs ensure compromise of long-term key material does not expose past sessions.

Quantum Resistance: The hybrid ML-KEM-768 + X25519 key exchange is secure against quantum computers running Shor’s algorithm. Symmetric keys of 256 bits retain approximately 128 bits of security against Grover’s algorithm.

Endpoint Authentication (optional): When identity modes are used, signature algorithms SHOULD use ML-DSA-65 [FIPS204] for quantum resistance.

13.3. Implementation Requirements

Implementations MUST:

* Use constant-time implementations for all cryptographic operations to prevent timing side-channel attacks.

* Explicitly zeroize all key material after use using compiler-barrier-protected routines.

* Validate all header fields in received packets before processing.

* Implement PMTUD and respond to ICMP Fragmentation Needed.

Implementations MUST NOT:

* Reuse AEAD nonces under any circumstances.

* Continue data transmission after sequence number exhaustion without successful re-keying.

* Log session key material in plaintext.

13.4. Denial-of-Service Considerations

HANDSHAKE processing involves asymmetric cryptographic operations. Implementations SHOULD implement HANDSHAKE rate limiting (recommended: no more than 1000 new sessions per second per anchor) and SHOULD employ a lightweight proof-of-work or cookie mechanism to deter amplified HANDSHAKE flooding.

ELMS deployments SHOULD use Anycast addressing, multi-tier caching with signed responses, and rate limiting of Map-Request messages per source address.

14. IANA Considerations

14.1. UDP Port Number

IANA is requested to assign UDP port 4242 in the "Service Name and Transport Protocol Port Number Registry":

* Service Name: ipv45
* Transport Protocol: UDP
* Port Number: 4242
* Description: IPv4.5 Data Transport (Mode C encapsulation)
* Assignee: N. Matsukami <info@insane-projects.com>
* Reference: This document

14.2. IP Protocol Number

IANA is requested to assign IP Protocol Number 144 in the "Assigned Internet Protocol Numbers" registry:

* Decimal: 144

* Keyword: IPv4.5
* Protocol: Internet Protocol version 4.5
* Reference: This document

14.3. DNS Resource Record Type

IANA is requested to assign a new DNS RR type in the "Domain Name System (DNS) Parameters" registry:

* Type: AX
* Value: [to be assigned by IANA]
* Meaning: IPv4.5 Address Record
* Reference: This document

The wire format of the AX record carries a 96-bit IPv4.5 address in network byte order (12-byte RDATA field).

14.4. BGP Subsequent Address Family Identifier (SAFI)

IANA is requested to assign SAFI value 144 in the "Subsequent Address Family Identifiers (SAFI) Parameters" registry:

* Value: 144
* Name: IPv4.5 EID-Locator Mapping
* Reference: This document

14.5. IPv4.5 Packet Type Registry

IANA is requested to create a new registry "IPv4.5 Packet Types" within a new "IPv4.5 Parameters" registry group. Future assignments in values 6-14 require Standards Action [RFC8126].

Value	Name	Reference
0	DATA	This document
1	HANDSHAKE	This document
2	CONTROL	This document
3	ECHO	This document
4	ERROR	This document
5	KEEPALIVE	This document
6-14	Unassigned	Standards Action
15	Reserved	—

Table 7: IPv4.5 Packet Types (Initial Values)

14.6. IPv4.5 Error Code Registry

IANA is requested to create a new registry "IPv4.5 Error Codes" within the "IPv4.5 Parameters" registry group:

Value	Name	Description
0	UNSPECIFIED	Unspecified error
1	HOP_LIMIT_EXCEEDED	Hop Limit reached zero
2	MTU_EXCEEDED	Packet exceeds path MTU
3	UNREACHABLE	Destination unreachable
4	AUTH_FAILED	Handshake authentication failure
5	POLICY_DENIED	INP policy rejection
6-255	Unassigned	Standards Action

Table 8: IPv4.5 Error Codes (Initial Values)

15. References

15.1. Normative References

[FIPS203] National Institute of Standards and Technology (NIST), "Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM)", NIST FIPS 203, August 2024, <https://doi.org/10.6028/NIST.FIPS.203>.

[FIPS204] National Institute of Standards and Technology (NIST), "Module-Lattice-Based Digital Signature Standard (ML-DSA)", NIST FIPS 204, August 2024, <https://doi.org/10.6028/NIST.FIPS.204>.

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <https://www.rfc-editor.org/info/rfc768>.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/info/rfc791>.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990, <https://www.rfc-editor.org/info/rfc1191>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <https://www.rfc-editor.org/info/rfc4086>.

[RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, DOI 10.17487/RFC4821, March 2007, <https://www.rfc-editor.org/info/rfc4821>.

[RFC5116] McGrew, D., "An Interface and Algorithms for Authenticated Encryption", RFC 5116, DOI 10.17487/RFC5116, January 2008, <https://www.rfc-editor.org/info/rfc5116>.

[RFC7748] Langley, A., Hamburg, M., and S. Turner, "Elliptic Curves for Security", RFC 7748, DOI 10.17487/RFC7748, January 2016, <https://www.rfc-editor.org/info/rfc7748>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <https://www.rfc-editor.org/info/rfc8201>.

[RFC8305] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, December 2017, <https://www.rfc-editor.org/info/rfc8305>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

[RFC8981] Gont, F., Krishnan, S., Narten, T., and R. Draves, "Temporary Address Extensions for Stateless Address Autoconfiguration in IPv6", RFC 8981, DOI 10.17487/RFC8981, February 2021, <https://www.rfc-editor.org/info/rfc8981>.

[RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021, <https://www.rfc-editor.org/info/rfc9000>.

[RFC9300] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, Ed., "The Locator/ID Separation Protocol (LISP)", RFC 9300, DOI 10.17487/RFC9300, October 2022, <https://www.rfc-editor.org/info/rfc9300>.

[RFC9301] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, Ed., "Locator/ID Separation Protocol (LISP) Control Plane", RFC 9301, DOI 10.17487/RFC9301, October 2022, <https://www.rfc-editor.org/info/rfc9301>.

15.2. Informative References

[KADEMLIA] Maymounkov, P. and D. Mazieres, "Kademlia: A Peer-to-peer Information System Based on the XOR Metric", Proceedings of IPTPS 2002, 2002.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <https://www.rfc-editor.org/info/rfc1034>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <https://www.rfc-editor.org/info/rfc1918>.

[RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <https://www.rfc-editor.org/info/rfc2474>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, DOI 10.17487/RFC3022, January 2001, <https://www.rfc-editor.org/info/rfc3022>.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <https://www.rfc-editor.org/info/rfc4760>.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <https://www.rfc-editor.org/info/rfc5280>.

[RFC6598] Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space", BCP 153, RFC 6598, DOI 10.17487/RFC6598, April 2012, <https://www.rfc-editor.org/info/rfc6598>.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <https://www.rfc-editor.org/info/rfc6830>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <https://www.rfc-editor.org/info/rfc7217>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC8926] Gross, J., Ed., Ganga, I., Ed., and T. Sridhar, Ed., "Geneve: Generic Network Virtualization Encapsulation", RFC 8926, DOI 10.17487/RFC8926, November 2020, <https://www.rfc-editor.org/info/rfc8926>.

[RFC9299] Cabellos, A. and D. Saucez, Ed., "An Architectural Introduction to the Locator/ID Separation Protocol (LISP)", RFC 9299, DOI 10.17487/RFC9299, October 2022, <https://www.rfc-editor.org/info/rfc9299>.

[W3C-DID] Sporny, M., Guy, A., Sabadello, M., and D. Reed, "Decentralized Identifiers (DIDs) v1.0", W3C Recommendation, July 2022, <https://www.w3.org/TR/did-core/>.

Appendix A. Deployment Considerations

A.1. Phase 1 (2027-2030): Early Adoption

IPv4.5 is deployed as a user-space overlay using Mode C encapsulation. No kernel changes are required. Day-one benefits include quantum-resistant session encryption, ephemeral EIDs for endpoint privacy, and global reachability without CGN for endpoints behind cooperative ISPs.

A.2. Phase 2 (2030-2035): Infrastructure Integration

OS kernel integration, router support for ELMS BGP SAFI, and deployment of Edge Translators for IPv4 interoperability. ISPs offering native IPv4.5 addressing can eliminate CGN for IPv4.5-capable subscribers.

A.3. Phase 3 (2035+): Native Protocol

IP Protocol 144 (Mode D) is assigned and native IPv4.5 forwarding is enabled in router hardware. IPv4 continues as the backward-compatible locator plane.

Appendix B. Open Issues

1. DNS AX Record Wire Format: A companion specification is needed for the complete AX record format, including DNSSEC signing.

2. TLS WG Coordination: Alignment with draft-ietf-tls-hybrid-design should be evaluated.

3. UDP/443 Fallback: The mechanism for detecting UDP/4242 blockage and falling back to UDP/443 needs specification.

4. Edge Translator HA: High-availability state synchronization for Edge Translators requires a separate specification.

5. ELMS DHT Specification: The Kademlia-based architecture requires a companion document.

6. Handshake Fragmentation: The Client Hello fragmentation strategy requires detailed specification for interoperability.

7. BGP-X Extensions: The BGP SAFI 144 specification requires a separate document.

Appendix C. Change Log

C.1. draft-matsukami-intarea-ipv45-00

Initial submission to the IETF Datatracker (April 2026).

* Address space redesigned to globally unique 96-bit addresses with explicit Locator/Identifier separation.

* Semantic communication, per-packet intent fields, mandatory DID/SSI, and per-packet PQC signatures removed from scope.

* Backward compatibility claim refined to "transparent IPv4 transit" with four explicit operating modes (A through D).

* Four-plane separation (data, control, identity, cryptographic) introduced.

* Hybrid PQC (ML-KEM-768 [FIPS203] + X25519 [RFC7748]) specified, referencing NIST FIPS 203 finalized August 2024.

Acknowledgements

The author thanks the IETF int-area mailing list participants for their ongoing discussion of IPv4 transition mechanisms and addressing architecture.

The Locator/Identifier separation concept builds directly on foundations laid by LISP [RFC6830] and its successors [RFC9300] [RFC9301]. The QUIC protocol [RFC9000] demonstrated that a new transport protocol can achieve wide deployment via UDP encapsulation and strongly influenced the Mode C design. The TLS 1.3 security model [RFC8446] informed the session security architecture.

The eNID privacy mechanism is modeled after IPv6 Privacy Extensions [RFC8981] and Semantically Opaque Interface Identifiers [RFC7217].

Author’s Address

Nao Matsukami
Independent Researcher
Email: info@insane-projects.com