Internet Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: 19 October 2026

J. Thain
One Limited

17 April 2026

Internet Protocol Version 8 (IPv8)
draft-thain-ipv8-02

Abstract

Internet Protocol Version 8 (IPv8) is a managed network protocol suite that transforms how networks of every scale -- from home networks to the global internet -- are operated, secured, and monitored. Every manageable element in an IPv8 network is authorised via OAuth2 JWT tokens served from a local cache. Every service a device requires is delivered in a single DHCP8 lease response. Every packet transiting to the internet is validated at egress against a DNS8 lookup and a WHOIS8 registered active route. Network telemetry, authentication, name resolution, time synchronisation, access control, and translation are unified into a single coherent Zone Server platform.

IPv4 is a proper subset of IPv8. An IPv8 address with the routing prefix field set to zero is an IPv4 address. No existing device, application, or network requires modification. The suite is 100% backward compatible. There is no flag day and no forced migration at any layer.

IPv8 also resolves IPv4 address exhaustion. Each Autonomous System Number (ASN) holder receives 4,294,967,296 host addresses. The global BGP8 routing table is structurally bounded by ASN count rather than prefix count. WHOIS8 is a critical infrastructure service underpinning this model.

This document is one of the companion specifications:

* draft-thain-ipv8-02 Core protocol (this document)
* draft-thain-routing-protocols-00 BGP8, IBGP8, OSPF8, IS-IS8, CF
* draft-thain-rine-00 Regional Inter-Network Exchange
* draft-thain-zoneserver-00 Zone Server Architecture
* draft-thain-whois8-00 WHOIS8 Protocol
* draft-thain-netlog8-00 NetLog8 Protocol
* draft-thain-support8-00 ARP8, ICMPv8, Route8
* draft-thain-ipv8-mib-00 IPv8 MIB and SNMPv8
* draft-thain-wifi8-00 WiFi8 Protocol
* draft-thain-update8-00 Update8 and NIC Certification

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 19 October 2026.

Table of Contents

1. Introduction

1.1. 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.

1.2. The Network Management Problem

Modern network management is characterised by fragmentation. DHCP, DNS, NTP, logging, monitoring, and authentication are separate products, separately licensed, separately configured, and separately maintained with no shared awareness of network state. A device connecting to a network may require manual configuration of a dozen independent services before it is operational. Security is inconsistent -- some services are authenticated, others are not.

Failures require correlating data across systems that were never designed to work together.

This fragmentation scales with every device added. Small networks cannot afford the operational complexity. Large networks cannot afford the inconsistency. The global internet has no coherent mechanism to validate that advertised routes are legitimately held by their advertisers, or that packets transiting to external destinations have been validated against any registry of active routes.

IPv6 [RFC8200] addressed address exhaustion but did not address management fragmentation. After 25 years of deployment effort IPv6 carries a minority of global internet traffic. The operational cost of the dual-stack transition model, combined with the absence of management improvement, proved commercially unacceptable.

IPv8 addresses both problems simultaneously.

1.3. The IPv8 Management Philosophy

The central operational concept in IPv8 is the Zone Server -- a paired active/active platform that runs every service a network segment requires: address assignment (DHCP8), name resolution (DNS8), time synchronisation (NTP8), telemetry collection (NetLog8), authentication caching (OAuth8), route validation (WHOIS8 resolver), access control enforcement (ACL8), and IPv4/IPv8 translation (XLATE8).

A device connecting to an IPv8 network sends one DHCP8 Discover and receives one response containing every service endpoint it requires. No subsequent manual configuration is needed for any service. The device is fully operational -- authenticated, logged, time-synchronised, zone-policy-enforced -- before its first user interaction.

Every manageable element in an IPv8 network is authorised via OAuth2 JWT tokens [RFC7519]. Tokens are validated locally by the OAuth8 cache on the Zone Server without round trips to external identity providers. A device in a remote location with a temporarily unreachable cloud identity provider continues to authenticate normally -- the OAuth8 cache holds all public keys and validates signatures locally in sub- millisecond time. JWT tokens may be served by a local OAuth2 authority (home router operating in local authority mode) or by a cached enterprise OAuth2 provider. Authentication is universal, consistent, and requires no per-service credential management.

Firmware and software updates for L1-L4 stack components are managed via the Update8 protocol [UPDATE8]. Update8 defines a standard vendor feed format, Zone Server validated proxy, optional local caching, device criticality-based scheduling, and rollback prevention enforced in NIC hardware. Devices receive updates only from DNS-named sources validated by the Zone Server. Connection to an update source identified by IP address is blocked by default.

The 127.0.0.0/8 r.r.r.r range is permanently reserved as the IPv8 internal zone prefix space. Organisations assign internal zone prefixes (127.1.0.0, 127.2.0.0 etc) to network zones and regions. Internal zone addresses are never routed externally. No address conflict between zones is possible. An organisation may build a network of arbitrary geographic and organisational scale -- with dozens of regional zones, each containing thousands of devices -- using familiar routing protocols without any external address coordination.

1.4. Zone Server Redundancy and Even/Odd Addressing

Every IPv8 network segment operates with two Zone Servers, assigned the two highest addresses in the subnet: .254 (even) and .253 (odd). These are the two default gateways issued to every device in the DHCP8 lease response.

IPv8 traffic follows an even/odd affinity model: even-addressed hosts route via the even Zone Server (.254) and odd-addressed hosts route via the odd Zone Server (.253). Each Zone Server is the PVRST root for its respective VLANs. Under normal operation the load is distributed across both Zone Servers with no single point of failure. On Zone Server failure, traffic fails over to the surviving gateway automatically.

Hosts SHOULD be assigned addresses matching their primary traffic path. Dual-NIC hosts SHOULD be issued one even and one odd address by DHCP8 -- one per NIC -- enabling full simultaneous use of both gateway paths and both physical links. Full redundancy is achieved with no additional configuration: loss of either NIC or either Zone Server results in automatic failover on the surviving path. Administrators MAY deviate from even/odd assignment where operational requirements justify it.

1.5. East-West and North-South Security

IPv8 addresses two distinct traffic security problems:

East-west security -- traffic between devices within a network -- is enforced by ACL8 zone isolation. Devices communicate only with their designated service gateway. The service gateway communicates only with the designated cloud service. Lateral movement between devices or zones is architecturally prevented by the absence of any permitted route to any other destination. Three independent enforcement layers provide defence in depth: NIC firmware ACL8, Zone Server gateway ACL8, and switch port OAuth2 hardware VLAN enforcement.

North-south security -- traffic from internal devices to the internet -- is enforced at the Zone Server egress by two mandatory validation steps. First, every outbound connection must have a corresponding DNS8 lookup -- no DNS lookup means no XLATE8 state table entry means the connection is blocked. Second, the destination ASN is validated against the WHOIS8 registry -- if the destination prefix is not registered as an active route by a legitimately registered ASN holder the packet is dropped. These two steps together eliminate the primary malware command-and-control channel: connection to hardcoded IP addresses without DNS resolution.

At the global routing level, BGP8 route advertisements are validated against WHOIS8 before installation in the routing table. A route that cannot be validated is not installed. Manual bogon filter list maintenance is eliminated. Prefix hijacking is architecturally difficult -- an attacker must compromise both an RIR registry entry and produce a validly signed WHOIS8 record.

1.6. Address Exhaustion

IANA completed allocation of the IPv4 unicast address space in February 2011. CGNAT has extended IPv4 life but introduces latency, breaks peer-to-peer protocols, and complicates troubleshooting. The address exhaustion problem is architectural and cannot be resolved within the 32-bit IPv4 address space.

IPv8 resolves address exhaustion as a consequence of its addressing architecture, not as a primary design goal. The 64-bit IPv8 address space provides 2^64 unique addresses. Each ASN holder receives 2^32 host addresses -- 4,294,967,296 addresses -- sufficient for any organisation at any scale without address exhaustion, CGNAT, or renumbering.

IPv4 is a proper subset of IPv8. An IPv8 address with r.r.r.r = 0.0.0.0 is an IPv4 address, processed by standard IPv4 rules. No existing device, application, or network requires modification to participate in an IPv8 network. The suite is 100% backward compatible. There is no flag day and no forced migration at any layer.

The global BGP8 routing table is structurally bounded by ASN count. The /16 minimum injectable prefix rule prevents deaggregation -- an ASN may advertise multiple /16 prefixes but never more specific. Route acceptance is validated against WHOIS8, which serves as a critical infrastructure service for the global routing system. The BGP4 routing table exceeded 900,000 prefixes with no architectural bound. The BGP8 routing table is bounded by the number of /16 prefixes held across all ASNs -- a structurally finite and manageable set.

1.7. Routing Protocol Improvements

IPv8 extends OSPF8 [RFC2328], BGP8 [RFC4271] (both iBGP8 and eBGP8), and IS-IS8 with a unified path quality metric -- the Cost Factor (CF).

CF is a 32-bit accumulated metric derived from seven components measured from TCP session telemetry: round trip time, packet loss, congestion window state, session stability, link capacity, economic policy, and geographic distance as a physics floor. CF accumulates across every BGP8 hop from source to destination. Every router independently selects the path with the lowest accumulated CF without coordination.

CF combines the dynamic composite path quality of EIGRP, the accumulated cost model of OSPF, and proportional load balancing across multiple paths -- in a single open versioned algorithm that operates end-to-end across AS boundaries. OSPF and EIGRP stop at the AS boundary. CF does not.

The geographic component of CF sets a physics floor -- no path can appear better than the speed of light over the great circle distance allows. A path that measures faster than physics permits is flagged immediately as a CF anomaly.

CF is an open versioned algorithm. CFv1 is the mandatory baseline. Future versions may add carbon cost, jitter, time of day, and application layer latency signals through the IETF process.

CF contribution for an IPv4-only hop is zero -- neither a positive nor a negative weight -- unless an operator explicitly assigns a CF policy weight to that hop. IPv4 hops are invisible to CF path optimisation by default. This ensures IPv4 transit paths are neither preferred nor penalised in CF path selection solely by virtue of their protocol version.

1.8. Backward Compatibility and Transition

IPv4 is a proper subset of IPv8:

IPv8 address with r.r.r.r = 0.0.0.0 = IPv4 address
Processed by standard IPv4 rules
No modification to IPv4 device required
No modification to IPv4 application required
No modification to IPv4 internal network required

IPv8 does not require dual-stack operation. There is no flag day. 8to4 tunnelling enables IPv8 islands separated by IPv4- only transit networks to communicate immediately. CF naturally incentivises IPv4 transit ASNs to upgrade by measuring higher latency on 8to4 paths -- an automatic economic signal without any mandate.

Transition phases are independent. Tier 1 ISPs, cloud providers, enterprises, and consumer ISPs may adopt IPv8 in any order and at any pace. 8to4 ensures interoperability throughout.

2. ARP8-Driven Version Selection

2.1. The Foundational Rule

An IPv8 host or router transmitting to a neighbor on a shared segment MUST use the IP protocol version matching that neighbor’s capability as recorded in the ARP8 cache. Capability is discovered at first contact via ARP8 dual probe and recorded permanently for the lifetime of the cache entry.

2.2. Neighbor Capability Discovery

On first contact with a neighbor, an IPv8 host issues two probes in parallel:

t=0ms ARP8 probe (IPv8 capable neighbor expected)
t=50ms ARP4 probe (IPv4 fallback if no ARP8 response)

The first response received determines the neighbor’s recorded capability. Once recorded, the capability entry persists for the lifetime of the ARP8 cache entry. No per-packet probing is required.

2.3. Version Selection at Transmit Time

The key rule of IPv8 backward compatibility: an IPv8 device SHALL transmit only IPv4 packets to IPv4 devices. There are no exceptions. An IPv4 device on any segment with IPv8 devices will never receive a packet it cannot process.

For a neighbor recorded as IPv8-capable, the sender constructs a full IPv8 packet:

```
Version field:       8
Source address:      r.r.r.r.n.n.n.n  (64-bit, full prefix)
Destination address: r.r.r.r.n.n.n.n  (64-bit, full prefix)
```

For a neighbor recorded as IPv4-only, the sender constructs a standard IPv4 packet:

```
Version field:       4
Source address:      n.n.n.n  (32-bit, host portion only)
Destination address: n.n.n.n  (32-bit, host portion only)
```

The r.r.r.r prefix is not present on the wire for that hop. The IPv4 neighbor receives a packet indistinguishable from normal IPv4. An IPv4-only device on a shared segment with IPv8-capable devices never receives a packet with version 8 in the IP header. The mismatched-version drop case never arises because the IPv8 sender constructs the packet in the version the neighbor understands.

2.4. Router Forwarding

An IPv8 router forwarding to a next-hop IPv4-only neighbor applies the same rule. If a packet arrived as IPv8 on the inbound interface and the next hop is recorded as IPv4-only, the router MUST downgrade at the outgoing interface:

```
Inbound:   version 8, r.r.r.r.n.n.n.n source and destination
Outbound:  version 4, n.n.n.n source and destination
```

The r.r.r.r prefix is stripped from the source address. Session state is preserved via XLATE8 on the Zone Server for return traffic reconstruction. The IPv4 device on the outgoing segment receives a standard IPv4 packet and has no knowledge of IPv8 upstream.

A single IPv8 router MAY serve IPv8-capable and IPv4-only devices on different interfaces simultaneously, applying version selection independently per outgoing interface based on the ARP8-recorded capability of each next-hop neighbor.

2.5. Attribution at IPv4 Hops

When a packet is transmitted as IPv4 -- either host-to-IPv4- neighbor or router-downgraded at an IPv4 boundary -- the r.r.r.r prefix is not on the wire for that hop. ASN attribution and WHOIS8 validation apply only to hops where both endpoints are IPv8-capable. This is an accepted and correct property of the transition model: IPv4 communications do not carry IPv8 properties, in the same way that IPv4 communications today do not carry RPKI validation unless both endpoints support it.

2.6. Transition Properties

This model produces the following transition properties:

* IPv4-only endpoints never require modification.
* IPv4 devices on a shared segment with IPv8 devices continue to operate without configuration change.
* IPv4 devices behind IPv8 routers continue to operate because the router downgrades at the boundary.
* IPv4 applications on IPv8 hosts continue to operate because XLATE8 handles version translation on their behalf.
* No IPv4 device ever receives a packet with version 8 in the IP header.
* Transition is per-device and per-router, on each operator’s own schedule, with no flag day and no coordination requirement between operators.

Four lines of configuration enable IPv8 on a router. No existing IPv4 device, application, or network requires any modification.

3. Motivation and Problem Statement

3.1. Management Fragmentation

IPv4 network management has no coherent integrated model. The protocols that operate a network -- DHCP, DNS, NTP, syslog, SNMP, authentication -- were specified independently over four decades, share no common identity model, no common authentication mechanism, and no common telemetry format.

The consequences are operational: networks require specialist knowledge of each protocol independently. Security is inconsistent -- some services require authentication, others accept unauthenticated requests from any source. Failures require correlating logs across systems with different timestamp formats, different severity models, and different identity representations. Management scales with operational burden, not with network size.

IPv8 addresses this by defining a coherent management suite in which every service shares a common identity model (OAuth2 JWT), a common delivery mechanism (DHCP8), a common telemetry format (NetLog8), and a common authentication cache (OAuth8).

3.2. Address Exhaustion

IANA completed allocation of the IPv4 unicast address space in February 2011. Regional Internet Registries exhausted their allocations between 2011 and 2020. CGNAT extended IPv4 life at the cost of latency, peer-to-peer protocol breakage, and troubleshooting complexity.

IPv6 [RFC8200] was developed to address exhaustion. After 25 years of standardisation and deployment effort IPv6 carries a minority of global internet traffic. The dual-stack transition model -- requiring every device, application, and network to simultaneously support both protocols -- proved commercially unacceptable. The absence of a forcing function meant organisations could continue with CGNAT indefinitely.

IPv8 resolves address exhaustion without dual-stack operation. IPv4 is a proper subset of IPv8. The transition requires no flag day and creates no operational discontinuity.

3.3. Routing Table Growth

The BGP4 global routing table exceeded 900,000 prefixes in 2024 and grows without architectural bound. Prefix deaggregation -- advertising more specific prefixes to influence traffic engineering -- is the primary driver of growth. No protocol mechanism prevents it.

BGP4 has no binding relationship between what an ASN advertises and what it is authorised to advertise. Prefix hijacking, route leaks, and bogon injection are possible because there is no route ownership registry that border routers enforce as a condition of route acceptance.

IPv8 addresses both problems. The /16 minimum injectable prefix rule prevents deaggregation at inter-AS boundaries. WHOIS8 mandatory route validation creates a binding relationship between BGP8 advertisements and registered route ownership -- WHOIS8 is a critical infrastructure service for the global routing system. The global BGP8 routing table is bounded by the number of /16 prefixes held across all active ASNs, a structurally finite set.

3.4. Requirements for a Viable Successor

* R1. Integrated management -- common identity, authentication, telemetry, and service delivery across all network services.
* R2. Single stack operation -- no dual-stack requirement.
* R3. Full backward compatibility -- existing IPv4 applications unchanged. IPv4 is a proper subset of IPv8.
* R4. Full backward compatibility -- RFC 1918 internal networks unchanged.
* R5. Full backward compatibility -- CGNAT deployments unchanged.
* R6. Vastly expanded address space.
* R7. Implementable as a software update without hardware replacement.
* R8. Human readable addressing consistent with IPv4 operator familiarity.
* R9. East-west and north-south traffic security enforced by protocol, not by manual configuration.
* R10. Structurally bounded global routing table.

IPv8 satisfies all ten requirements.

4. IPv8 Address Format

4.1. Structure

An IPv8 address is a 64-bit value:

r.r.r.r.n.n.n.n

* *r.r.r.r* -- 32-bit ASN Routing Prefix
* *n.n.n.n* -- 32-bit Host Address (identical semantics to IPv4)

4.2. Address Space

2^64 = 18,446,744,073,709,551,616 unique addresses.
2^32 ASN prefixes x 2^32 host addresses per ASN.

4.3. IPv4 Representation in IPv8

0.0.0.0.n.n.n.n

Packets with r.r.r.r = 0.0.0.0 MUST be routed using standard IPv4 rules applied to the n.n.n.n field. IPv4 is a proper subset of IPv8. No modification to IPv4 devices, applications, or networks is required.

4.4. ASN Encoding in r.r.r.r

The 32-bit ASN is encoded directly into r.r.r.r as a 32-bit unsigned integer in network byte order:

```
ASN 64496 (Example-A)   = 0.0.251.240
ASN 64497 (Example-B)   = 0.0.251.241
ASN 64498 (Example-C)   = 0.0.251.242
```

4.5. Internal Zone Prefix (127.0.0.0/8)

The 127.0.0.0/8 range of the r.r.r.r field is permanently reserved for internal IPv8 zone prefixes. Internal zone prefixes identify network zones within an organisation’s private addressing space.

127.x.x.x.n.n.n.n

Where x.x.x identifies the internal zone. Examples:

```
127.1.0.0.n.n.n.n   Internal zone 1 (e.g. Americas)
127.2.0.0.n.n.n.n   Internal zone 2 (e.g. Europe)
127.3.0.0.n.n.n.n   Internal zone 3 (e.g. Asia Pacific)
```

Internal zone prefix rules:

* MUST NOT be routed externally beyond the organisation’s AS boundary.
* MUST NOT appear on WAN interfaces or public internet links.
* MUST NOT be used in eBGP8 advertisements.
* MAY be used freely within an organisation’s internal routing infrastructure via OSPF8, IS-IS8, and IBGP8.
* Provides 2^56 effective internal addresses across all zone prefixes. No internal address conflict is possible between zones.
* Enables organisations to build geographically distributed, region-routed private networks of arbitrary scale without external address coordination.

ASN numbers that encode to the 127.0.0.0/8 range in the r.r.r.r field (ASN 2130706432 through ASN 2147483647) are reserved for internal zone use and MUST NOT be allocated by IANA for public internet routing.

4.6. Inter-Company Interop Prefix (127.127.0.0)

The 127.127.0.0 prefix is reserved as the standard inter- company interoperability DMZ. When two organisations need to interconnect without exposing their internal zone addressing, both deploy XLATE8 engines facing a shared 127.127.0.0 address space. Full specification in [ZONESERVER] Section 16.9.

4.7. Two-XLATE8 Interop Model

```
Company A                          Company B
---------                          ---------
127.1.0.0.x   XLATE8-A  127.127.0.0  XLATE8-B  127.2.0.0.x
```

Properties:

* Company A never sees Company B’s 127.2.0.0 addresses.
* Company B never sees Company A’s 127.1.0.0 addresses.
* Each company controls exactly what it exposes.
* No address overlap possible. No NAT complexity.
* Setup time: minutes per service exposed.

4.8. Private Interop ASN (ASN 65534)

ASN 65534 is reserved for private inter-company BGP8 peering consistent with [RFC6996]:

0.0.255.254.x.x.x.x

ASN 65533 (0.0.255.253.x.x.x.x) is reserved for documentation and testing purposes.

4.9. RINE Peering Prefix (100.0.0.0/8)

The 100.0.0.0/8 range of the r.r.r.r field is permanently reserved for the Regional Inter-Network Exchange (RINE) peering fabric. RINE addresses are used exclusively for AS-to-AS peering link addressing at IXPs and private interconnect facilities. Full specification in [RINE].

RINE addresses:

* MUST NOT be advertised in the global BGP8 routing table.
* MUST NOT be assigned to end devices.
* MUST be filtered at all eBGP8 border routers.

4.10. Interior Link Convention (222.0.0.0/8)

The n.n.n.n range 222.0.0.0/8 is the well-known IPv8 interior link address convention. Every AS MAY use <own-asn>.222.x.x.x for router-to-router interior link addressing within their AS.

This convention is analogous to RFC 1918 [RFC1918] for IPv4 -- universally recognised, universally filtered, never routed externally, never an endpoint.

4.11. Address Usage Model

Address Space	Usage	Routable
127.x.x.x.n.n.n.n	Internal devices (all zones)	Never
127.127.0.0.n.n.n.n	Inter-company interop DMZ	Private
100.x.x.x.n.n.n.n	RINE peering links only	Never
<asn>.222.x.x.x	Interior router links	Never
0.0.255.254.n.n.n.n	Private BGP8 peering	Private
<own-asn>.n.n.n.n	Explicit public services only	Global
0.0.0.0.n.n.n.n	IPv4 compatible (r.r.r.r = 0)	IPv4 only

Table 1

Most devices on most networks use 127.x.x.x internal addressing. Public ASN addresses are used only for explicitly public-facing services.

5. Address Classes

r.r.r.r Value	Class	Description
0.0.0.0	IPv4 Compatible	Route on n.n.n.n using IPv4 rules
0.0.0.1 through 99.255.255.255	ASN Unicast	Route to ASN, deliver to n.n.n.n
		Public internet

		routing via eBGP8
100.0.0.0 through	RINE Peering	AS-to-AS peering link addressing
100.255.255.255		MUST NOT be globally routed
101.0.0.0 through	ASN Unicast	Route to ASN, deliver to n.n.n.n
126.255.255.255		Public internet routing via eBGP8
127.0.0.0 through	Internal Zone Prefix	Internal zone identifier
127.255.255.255		MUST NOT be routed externally
128.0.0.0 through	ASN Unicast	Route to ASN, deliver to n.n.n.n
ff.fe.ff.ff		Public internet routing via eBGP8
ff.ff.00.00	Cross-ASN Multicast	General cross-ASN multicast
ff.ff.00.01	OSPF8 Reserved	OSPF8 protocol multicast traffic
ff.ff.00.02	BGP8 Reserved	BGP8 peer discovery multicast
ff.ff.00.03	EIGRP Reserved	Reserved. Deprecated. Vendor ext.
ff.ff.00.04	RIP Reserved	Reserved. Deprecated.
ff.ff.00.05	IS-IS8 Reserved	IS-IS8. Vendor extensible.
ff.ff.00.06 through	Cross-ASN Multicast	Available for future IANA
ff.ff.ef.ff	(available)	assignment.

ff.ff.f0.00 through	Reserved	Future use.
ff.ff.fe.ff		
ff.ff.ff.ff	Broadcast	Maps to L2 broadcast. MUST NOT be routed.

Table 2

The n.n.n.n range 222.0.0.0/8 is reserved by convention for interior link addressing per Section 3.10.

5.1. Anycast

Anycast is not a separate address class in IPv8. Anycast is a routing property implemented via eBGP8. The Cost Factor (CF) metric defined in [ROUTING-PROTOCOLS] routes each packet to the nearest BGP8 instance by measured cost automatically.

6. IPv8 Packet Header

6.1. Header Format

IPv8 uses IP version number 8 in the Version field. The header extends IPv4 by replacing the 32-bit src/dst address fields with 64-bit equivalents.

<CODE BEGINS>

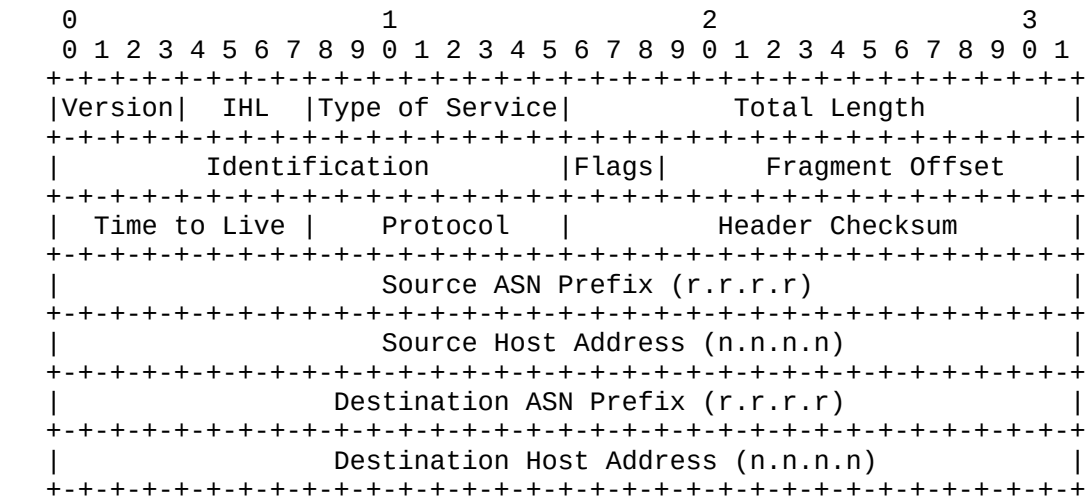

<CODE ENDS>

The IPv8 header is 8 octets longer than the IPv4 header.

6.2. Socket API Compatibility

Existing IPv4 applications use the standard BSD socket API with AF_INET and sockaddr_in. The IPv8 compatibility layer intercepts socket calls transparently -- the application has zero IPv8 awareness. New applications MAY use AF_INET8 with sockaddr_in8:

```
struct sockaddr_in8 {
    sa_family_t    sin8_family;   /* AF_INET8 */
    in_port_t      sin8_port;     /* port number */
    uint32_t       sin8_asn;      /* r.r.r.r ASN prefix */
    struct in_addr sin8_addr;     /* n.n.n.n host address */
};
```

7. ASN Dot Notation

Format: <ASN>.<n>.<n>.<n>.<n>

Where ASN is the autonomous system number and n.n.n.n is the host address. Example ASNs 64496-64511 are reserved for documentation per [RFC5398].

```
64496.192.0.2.1  = 0.0.251.240.192.0.2.1  (Example-A)
64497.192.0.2.1  = 0.0.251.241.192.0.2.1  (Example-B)
```

All IPv8-compliant implementations MUST accept ASN Dot Notation in all contexts where an IPv8 address is accepted.

8. DNS A8 Record Type

* *Type:* A8 (IANA assignment pending)
* *Format:* 64-bit IPv8 address in network byte order
* RFC 1918 addresses MUST NOT be published as A8 records in public DNS.
* An IPv8 resolver SHOULD request both A and A8 records.
* For IPv4 applications on IPv8 hosts, the resolver returns the n.n.n.n portion; the stack prepends r.r.r.r transparently.
* A8 responses SHOULD be an even/odd pair -- one even address and one odd address. Even addresses are served via the even Zone Server gateway (.254) and odd addresses via the odd Zone Server gateway (.253). This enables clients to open parallel streams to the same host across independent gateway paths, providing load distribution and redundancy without stateful load balancer infrastructure. Administrators MAY deviate from this convention where operational requirements justify it.

Example records:

```
ns1.example.com.  IN  A8  0.0.59.65.192.0.2.1
ns1.example.com.  IN  A8  0.0.59.65.192.0.2.2
```

9. Routing Protocol Behaviour

9.1. Mandatory Routing Protocols

Protocol	Scope	Function	Status
eBGP8	Inter-AS	Mandatory EGP for public internet	MANDATORY
IBGP8	Inter-zone	Mandatory for internal zone routing	MANDATORY
OSPF8	Intra-zone	Mandatory for intra-zone routing	MANDATORY
IS-IS8	Intra-AS	Available in all L3 stacks	MUST BE AVAIL.
Static	All scopes	Mandatory for legacy and VRF routing	MANDATORY

BGP4	Transition	IPv4 AS compatibility	TRANSITION

Table 3

9.2. Deprecated Routing Protocols

Protocol	Status in IPv8	Notes
RIP/RIPv2	DEPRECATED	Replaced by OSPF8
EIGRP	DEPRECATED	Vendor extensible

Table 4

9.3. eBGP8 - Mandatory Exterior Gateway Protocol

eBGP8 is the mandatory exterior gateway protocol. All L3 devices MUST implement eBGP8. eBGP8 is 100% backward compatible with BGP4 [RFC4271]. Full specification in [ROUTING-PROTOCOLS].

The minimum injectable prefix at inter-AS boundaries is /16. Prefixes more specific than /16 MUST NOT be advertised across AS boundaries.

9.4. IBGP8 - Inter-Zone Routing

IBGP8 distributes WHOIS8-validated external routes throughout an autonomous system with full CF metric awareness.

CF_total = CF_external + CF_intrazone

9.5. OSPF8 - Intra-Zone Routing

OSPF8 is OSPFv2 [RFC2328] extended with a CF export interface. All L3 devices MUST implement OSPF8. Full specification in [ROUTING-PROTOCOLS] Section 10.3.

9.6. IS-IS8 - Optional Interior Gateway Protocol

IS-IS8 MUST be available in all IPv8 L3 routing stacks. Carriers and operators MAY deploy IS-IS8 at their discretion. IPv8 makes no recommendation regarding IGP selection. Full specification in [ROUTING-PROTOCOLS] Section 10.4.

9.7. Two-Tier Routing Table

Tier	Scope	Index	Description
1	Global	r.r.r.r	Routes to correct AS border router
2	Local	n.n.n.n	Identical to existing IPv4 routing table

Table 5

When r.r.r.r = 0.0.0.0 the Tier 1 lookup is bypassed.

9.8. VRF - Virtual Routing and Forwarding

VRF is mandatory for all IPv8 L3 devices. The management VRF (VLAN 4090) and OOB VRF (VLAN 4091) MUST be implemented on all IPv8-compliant devices. VRF isolation is a routing table property that cannot be bypassed by software misconfiguration.

10. ICMPv8

ICMPv8 extends ICMP [RFC792] to support 64-bit IPv8 addresses. ICMPv8 is backward compatible with ICMPv4. Both versions MUST be supported simultaneously. ICMPv8 carries full 64-bit IPv8 addresses in Echo, Destination Unreachable, Time Exceeded, Redirect, and Parameter Problem messages. Path MTU Discovery is extended for the larger IPv8 header.

ICMPv8 version selection follows the ARP8 rule defined in Section 2: ICMPv8 messages transmitted to an IPv4-only neighbor MUST be constructed as ICMPv4 with 32-bit addresses. An IPv4-only device never receives an ICMPv8 message. Full specification in [SUPPORT8].

11. Multicast

11.1. Intra-ASN Multicast

```
0.0.0.0.224.0.0.0/4   All intra-ASN multicast
0.0.0.0.239.0.0.0/8   Administratively scoped intra-ASN
```

Packets with r.r.r.r = 0.0.0.0 and n.n.n.n in the multicast range MUST NOT be forwarded beyond the local AS boundary.

11.2. Cross-ASN Multicast

ff.ff.00.00.n.n.n.n	General cross-ASN multicast
ff.ff.00.01.n.n.n.n	OSPF8 protocol traffic
ff.ff.00.02.n.n.n.n	BGP8 peer discovery
ff.ff.00.03.n.n.n.n	EIGRP (reserved, deprecated)
ff.ff.00.04.n.n.n.n	RIP (reserved, deprecated)
ff.ff.00.05.n.n.n.n	IS-IS8 (reserved, vendor ext.)

11.3. Cross-ASN Multicast Group Assignments

ff.ff.00.00.224.0.0.1	All IPv8 routers
ff.ff.00.00.224.0.0.2	All IPv8 Zone Servers
ff.ff.00.00.224.0.0.5	OSPF8 all routers
ff.ff.00.00.224.0.0.6	OSPF8 designated routers
ff.ff.00.00.224.0.0.10	IBGP8 peer discovery
ff.ff.00.00.239.0.0.0/8	Administratively scoped

12. Anycast

Anycast in IPv8 is a routing property implemented via eBGP8 and the Cost Factor (CF) metric. No special r.r.r.r prefix is required. CF routes each packet to the nearest instance by measured path cost.

13. Broadcast

The r.r.r.r value ff.ff.ff.ff is permanently reserved for broadcast and maps to the Layer 2 broadcast address. Packets with r.r.r.r = ff.ff.ff.ff MUST NOT be routed beyond the local network segment.

14. Compatibility and Transition

14.1. Single Stack Operation

IPv8 does not require dual-stack operation. IPv4 is a proper subset of IPv8 with r.r.r.r = 0.0.0.0. There is no flag day and no forced migration.

14.2. IPv4 Network Compatibility

Networks that have not deployed IPv8 continue to operate unchanged. IPv8 border routers apply ARP8-driven version selection at each outgoing interface: the r.r.r.r prefix is stripped and the packet is downgraded to IPv4 for any next-hop neighbor recorded as IPv4-only in the ARP8 cache. No configuration is required on the IPv4 side.

14.3. 8to4 - IPv8 Across IPv4-Only Networks

IPv8 traverses IPv4-only networks without pre-configured tunnels. A large service provider MAY leave their entire IPv4 core unchanged for 20 years. IPv8 traffic traverses it automatically using per-hop anycast encapsulation.

Every ASN publishes an IPv4 anycast address in its WHOIS8 record. All BGP8 edge routers in that ASN advertise this address into the surrounding BGP4 fabric as a normal IPv4 route. The global BGP4 table carries it with zero modification. When a BGP8 router needs to reach a destination ASN across IPv4-only hops, it looks up the destination ASN in its WHOIS8 cache, gets the anycast address, encapsulates the IPv8 packet in a standard IPv4 envelope addressed to that anycast address, and forwards it. IPv4-only routers in the core forward the outer IPv4 packet using normal routing. The nearest BGP8 edge router for the destination ASN decapsulates and delivers.

This model requires zero tunnel configuration on any device. All available paths through the IPv4 core are used by normal BGP4 path selection. A network that would have required 100,000 pre-configured tunnels under the tunnel model requires zero under the anycast model.

Full specification in [ROUTING-PROTOCOLS] Section 11.

14.4. Transition Sequence

Phase 1: Tier 1/2 ISP routers deploy IPv8 via software update.
Phase 2: Cloud providers deploy IPv8 internally.
Phase 3: Enterprise networks optionally adopt ASN prefixes.
Phase 4: Consumer ISPs deploy IPv8.

8to4 tunnelling enables each phase to interoperate with phases not yet completed. There is no dependency between phases.

15. CGNAT Behaviour

CGNAT devices require no modification. IPv8-aware CGNAT MUST NOT modify the r.r.r.r field during translation. Only the n.n.n.n field is subject to NAT translation. CGNAT operators without an ASN MUST use r.r.r.r = 0.0.0.0.

15.1. XLATE8 Even/Odd Load Balancing

When an IPv4 client connects to an IPv8 host via an XLATE8 gateway, the destination host may have both an even and an odd A8 address. The XLATE8 gateway SHOULD pass both addresses through to the IPv4 client where the client is capable of using both. Where the IPv4 client is not capable, the XLATE8 gateway MAY perform load balancing internally, distributing connections across the even and odd addresses of the destination host. This provides IPv4 clients with the benefit of IPv8 even/odd load distribution transparently.

16. Application Compatibility

Existing IPv4 applications require no modification. The IPv8 socket compatibility layer transparently manages r.r.r.r via DNS8 interception and XLATE8. New applications MAY use AF_INET8 and sockaddr_in8 as defined in Section 5.2.

17. Cloud Provider Applicability

IPv8 resolves VPC address overlap, VPC peering complexity, and multi-cloud routing through ASN-based disambiguation. The 127.x.x.x internal zone prefix enables cloud providers to assign unique zone prefixes to customer VPCs without address renumbering. Each customer VPC receives a unique 127.x.x.x zone prefix -- no two customer networks can overlap regardless of RFC 1918 address reuse within each VPC.

18. Device Compliance Tiers

18.1. Tier 1 - End Device

End devices MUST implement: Route8 unified routing table, static routes, VRF (management plane), two default gateways (even Zone Server .254, odd Zone Server .253), DHCP8 client, ARP8, ICMPv8, TCP/443 persistent connection to Zone Server, NetLog8 client, ACL8 client-side enforcement, management VRF (VLAN 4090), OOB VRF (VLAN 4091), gratuitous ARP8 on boot.

End devices SHOULD use the even gateway for even-addressed traffic and the odd gateway for odd-addressed traffic, failing over to the surviving gateway on failure. Dual-NIC end devices SHOULD request one even and one odd address from DHCP8.

End devices MAY implement: OSPF8, IS-IS8, eBGP8, IBGP8.

End devices DO NOT require any routing protocol to reach their default gateway.

18.2. Tier 2 - L2 Network Device

L2 devices MUST implement: 802.1Q trunking, VLAN auto- creation on tagged traffic, management VRF (VLAN 4090), OOB VRF (VLAN 4091), switch port OAuth2 binding, LLDP, NetLog8 client, ARP8 (management plane only), ICMPv8 (management plane only), PVRST, Zone Server as PVRST root capability, sticky MAC binding, Zone Server MAC notification.

18.3. Tier 3 - L3 Network Device

L3 devices MUST implement: All Tier 1 requirements, eBGP8, IBGP8, OSPF8, IS-IS8 (available), static routes, VRF (full), XLATE8 (mandatory on edge devices), WHOIS8 resolver, ACL8 gateway enforcement, Zone Server services (if Zone Server role), PVRST root capability, switch port OAuth2 binding support.

18.4. Spanning Tree - PVRST Mandatory

PVRST is mandatory for all IPv8 L2 and L3 devices. MST is not recommended. Zone Servers are PVRST roots by default:

* Primary Zone Server (.254): PVRST root for even VLANs, priority 4096.
* Secondary Zone Server (.253): PVRST root for odd VLANs, priority 4096.

18.5. NIC Broadcast Rate Limits

IPv8 certified NIC firmware enforces broadcast and control packet rate limits that cannot be overridden by software. These limits apply to broadcast and unauthenticated traffic only and do not constrain unicast data throughput:

```
Broadcasts:            10 per second maximum
User unauthenticated:  10 per second, max 30 per minute
User authenticated:    100 per second, max 300 per minute
```

The DHCP8 Zone Server is the sole authority for rate limit elevation. Full NIC certification specification in [UPDATE8].

19. Security Considerations

19.1. ASN Prefix Spoofing

IPv8 border routers MUST implement ingress filtering validating that the source r.r.r.r of received packets matches the ASN of the BGP8 peer. Consistent with BCP 38 [RFC2827].

19.2. Internal Zone Prefix Protection

The 127.x.x.x internal zone prefix MUST NOT appear on WAN interfaces. Border routers MUST drop packets with 127.x.x.x source or destination r.r.r.r on external interfaces. NetLog8 SEC-ALERT MUST be generated for each violation.

19.3. RINE Prefix Protection

The 100.x.x.x RINE prefix MUST NOT appear in eBGP8 advertisements or on non-peering interfaces. NetLog8 SEC-ALERT MUST be generated for each violation.

19.4. Interior Link Convention Protection

Border routers MUST filter received BGP8 advertisements containing n.n.n.n addresses in the 222.0.0.0/8 range. NetLog8 E3 trap MUST be generated for each violation.

19.5. RFC 1918 Address Privacy

RFC 1918 private addresses in n.n.n.n remain non-routable on the public internet consistent with IPv4 behaviour.

19.6. Cross-ASN Multicast Filtering

Routing protocol reserved prefixes ff.ff.00.01 through ff.ff.00.05 MUST be filtered at all border routers.

19.7. /16 Minimum Prefix Enforcement

Prefixes more specific than /16 MUST NOT be accepted from external BGP8 peers. Such advertisements MUST be rejected and logged via NetLog8 as SEC-ALERT.

20. IANA Considerations

20.1. IP Version Number

IANA is requested to assign version number 8 in the IP Version Number registry to Internet Protocol Version 8.

20.2. Internal Zone Prefix Reservation

IANA is requested to reserve the r.r.r.r range 127.0.0.0 through 127.255.255.255 as the IPv8 internal zone prefix space. ASN numbers 2130706432 through 2147483647 MUST NOT be allocated for public internet routing.

20.3. RINE Prefix Reservation

IANA is requested to reserve the r.r.r.r range 100.0.0.0 through 100.255.255.255 as the IPv8 RINE peering fabric range. ASN numbers 1677721600 through 1694498815 MUST NOT be allocated for public internet routing.

20.4. Interior Link Convention

IANA is requested to document the n.n.n.n range 222.0.0.0/8 as the IPv8 interior link address convention.

20.5. Cross-ASN Multicast Range

IANA is requested to establish a registry for IPv8 cross-ASN multicast prefixes within ff.ff.00.00 through ff.ff.ef.ff.

20.6. Broadcast Reservation

IANA is requested to reserve r.r.r.r value ff.ff.ff.ff as the IPv8 broadcast address.

20.7. DNS A8 Record Type

IANA is requested to assign a DNS resource record type number for the A8 record type defined in Section 7.

20.8. Multicast Group Assignments

IANA is requested to assign multicast groups within ff.ff.00.00.224.0.0.0/24 as defined in Section 10.3.

20.9. Private ASN Reservations

IANA is requested to reserve ASN 65534 for private inter- company BGP8 peering and ASN 65533 for documentation and testing purposes consistent with [RFC6996].

21. Informative References

[RFC1918] Rekhter, Y., "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996, <https://www.rfc-editor.org/info/rfc1918>.

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

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998, <https://www.rfc-editor.org/info/rfc2328>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering", BCP 38, RFC 2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[RFC4271] Rekhter, Y., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC5398] Huston, G., "Autonomous System (AS) Numbers Reserved for Documentation Use", RFC 5398, December 2008, <https://www.rfc-editor.org/info/rfc5398>.

[RFC6996] Mitchell, J., "Autonomous System (AS) Reservation for Private Use", BCP 6, RFC 6996, July 2013, <https://www.rfc-editor.org/info/rfc6996>.

[RFC7519] Jones, M., "JSON Web Token (JWT)", RFC 7519, May 2015, <https://www.rfc-editor.org/info/rfc7519>.

[RFC792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981, <https://www.rfc-editor.org/info/rfc792>.

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

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

[RINE] Thain, J., "Regional Inter-Network Exchange", Work in Progress, Internet-Draft, draft-thain-rine-00, April 2026, <https://datatracker.ietf.org/doc/html/draft-thain-rine-00>.

[ROUTING-PROTOCOLS] Thain, J., "IPv8 Routing Protocols", Work in Progress, Internet-Draft, draft-thain-routing-protocols-00, April 2026, <https://datatracker.ietf.org/doc/html/draft-thain-routing-protocols-00>.

[SUPPORT8] Thain, J., "IPv8 Support Protocols -- ARP8, ICMPv8, and Route8", Work in Progress, Internet-Draft, draft-thain-support8-00, April 2026, <https://datatracker.ietf.org/doc/html/draft-thain-support8-00>.

[UPDATE8] Thain, J., "Update8 and NIC Certification", Work in Progress, Internet-Draft, draft-thain-update8-00, April 2026, <https://datatracker.ietf.org/doc/html/draft-thain-update8-00>.

[ZONESERVER] Thain, J., "IPv8 Zone Server Architecture", Work in Progress, Internet-Draft, draft-thain-zoneserver-00, April 2026, <https://datatracker.ietf.org/doc/html/draft-thain-zoneserver-00>.

Author’s Address

Jamie Thain
One Limited
Hamilton
Bermuda
Email: jamie@one.bm