

IPv6 Formal Anycast and Functional Anycast Addresses

draft-smith-6man-form-func-anycast-addresses-01

IETF-106

Mark Smith
markzzzsmith@gmail.com

Current “Informal” IPv6 Anycast

RFC 4291:

“Anycast addresses are taken from the unicast address spaces (of any scope) and are not syntactically distinguishable from unicast addresses.”

IPv6 Anycast Configuration

'anycast' host interface manual address
configuration flag

Duplicate Address Detection (DAD) disabled per
RFC4862.

For reachability from off-link, injected into routing
protocol somehow.

Treated as Unicast

By:

routing domain/routing protocols

other hosts

Negative Sometimes

Duplicate “unicast” route (subnet/host) in routing domain:

Fault?

Yes - if not intended, needs to be fixed

No – it’s intended, it’s anycast, move on

Can’t tell from the address.

Negative Sometimes

Forwarding “Scopes/Domains” are coarse because Unicast:

- Global (GUA), Local Network (ULA), Link (Link-Local)

What about more discrete:

- e.g., “Realm”, “Administrative”, “Organization”?

Negative Sometimes

Applications or protocols can't distinguish between unicast and anycast without manual configuration.

E.g., Multipath TCP/MP-QUIC

1. establish with anycast
2. switch to unicast

Need to exclude anycast addresses during 2.

Well-Known, High Order Anycast Prefix?

ff00::/8 – IPv6 Multicast

aa00::/8 – IPv6 Anycast?

(aa - “anycast address”, fa - “formal anycast” an alternative)

Not My Idea

RFC 1546, “Host Anycasting Service”, November 1993

“There appear to be a number of ways to support anycast addresses, some of which use small pieces of the existing address space, others of which require that a special class of IP addresses be assigned.”

“The major advantages of using a separate class of addresses are that it is **easy to determine if an address is an anycast address** and **well-known anycast addresses are easier to support. ...**”

Existing Well-Known Anycast Addresses/Prefixes

Name/Purpose	Address/Prefix	Scope
Reserved Subnet Anycast Range	<subnet>,00-7f	Subnet Prefix's Scope (i.e. GUA, ULA, LL)
Mobile IPv6 Home-Agents anycast	<subnet>,7e	Subnet Prefix's Scope (i.e. GUA, ULA, LL)
IPv6 over GeoNetworking geographic anycast	<subnet>,7d	Subnet Prefix's Scope (i.e. GUA, ULA, LL)
Subnet-Router Anycast Address	<subnet>,0	Subnet Prefix's Scope (i.e. GUA, ULA, LL)

IANA: Internet Protocol Version 6 (IPv6) Anycast Addresses

Existing Well-Known Anycast Addresses/Prefixes

Name/Purpose	Address/Prefix	Scope	Globally Reachable
Discard-Only Address Block	100::/64	Network Local	No
Port Control Protocol Anycast	2001:1::1/128	Global	Yes
Traversal Using Relays around NAT Anycast	2001:1::2/128	Global	Yes
Direct Delegation AS112 Service	2620:4f:8000::/48	Global	Yes

Common Properties

Property	Comment
Encoding Services/Functions	These are not host/interface identifiers in unicast sense
Different Number Spaces	GUA Global (2000::/3), Non-GUA Global, Any subnet IID
Different Forwarding Domains/Scopes	Global, Local Network, Local Link (e.g. Link-Local Anycast address)

Common Properties

Property	Comment	IPv6 Multicast?
Encoding Services/Functions	These are not host/interface identifiers in unicast sense	✓
Different Number Spaces	GUA Global (2000::/3), Non-GUA Global, Any subnet IID	✓ (Global/unspecified, or within embedded unicast prefix domain)
Different Forwarding Domains/Scopes	Global, Local Network, Local Link (e.g. Link-Local Anycast address)	✓


Anycast is similar to unicast

(e.g., routing)

Anycast might also be similar to multicast?

(e.g., used for services/functions)

Anycast in the Middle

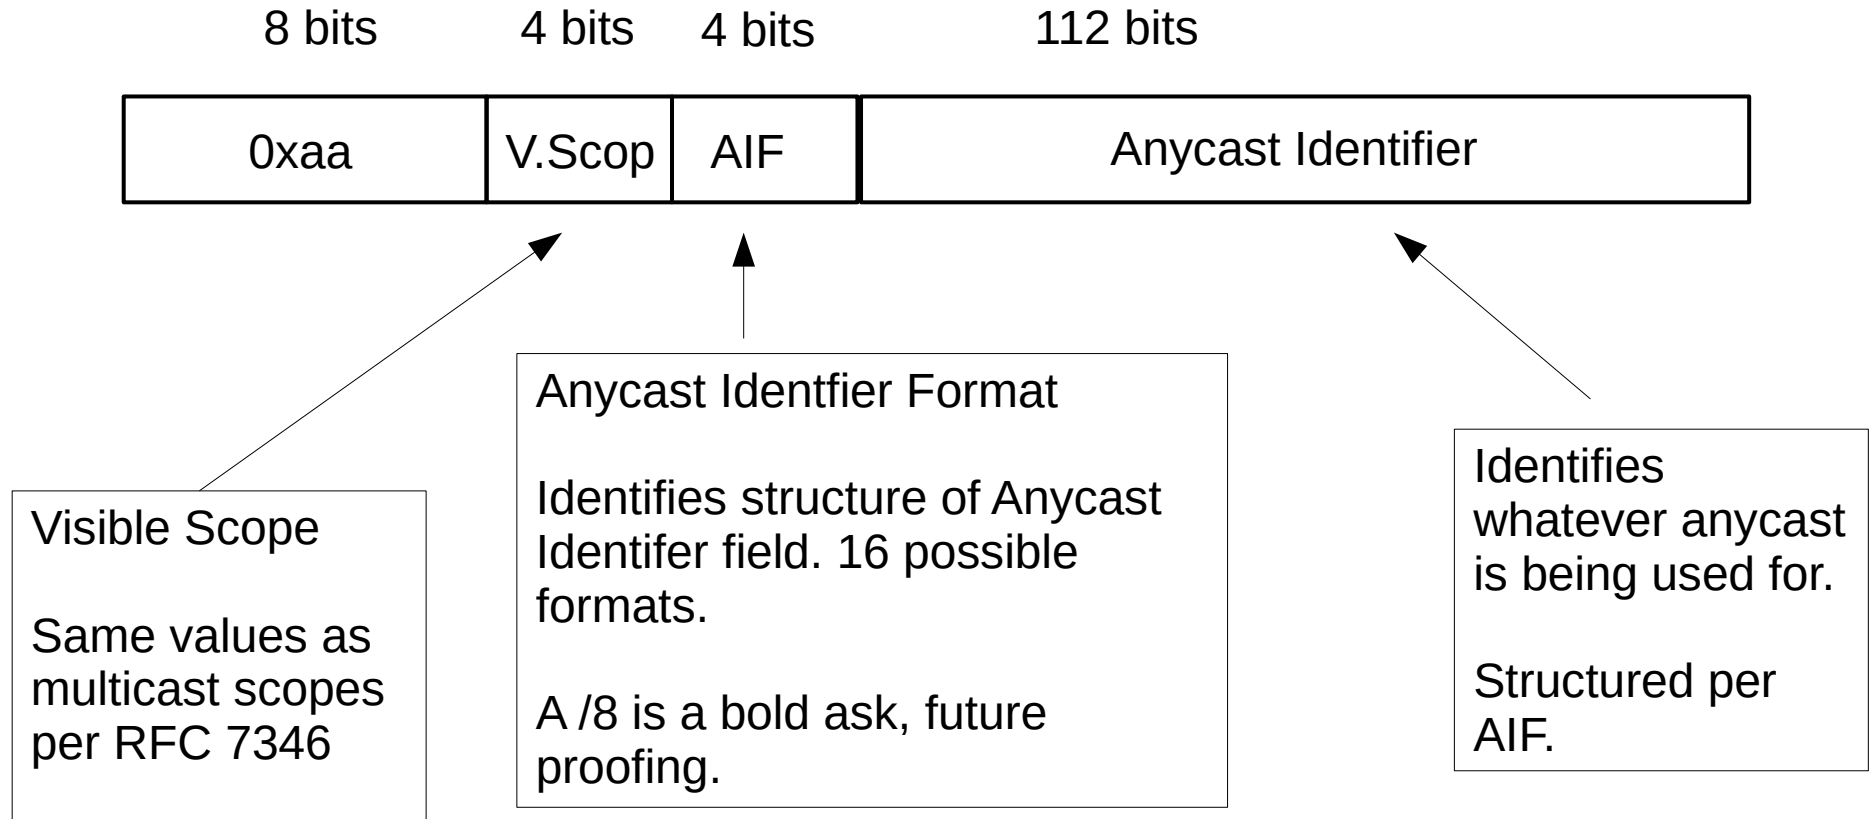


Formal Anycast Address Space

Seek IPv6 Multicast inspiration

- scopes
- embedded unicast prefix
- IANA well known function/service IDs
- Local network function/service IDs

IPv6 Formal Anycast Prefix



ICMPv6 Destination Unreachable

New Code:

“Edge of Visible Scope Reached”

Interim Equivalent:

Communication with destination administratively prohibited

IPv6 Destination Address Selection

Prefer Formal Anycast over Unicast by default

Anycast easily ignored if unicast preferred
because well known:

aa00::/8

Anycast Address Registration Protocol

“Host-based Anycast using MLD”

draft-haberman-ipngwg-host-anycast

Hop-by-Hop Using Anycast Address

Encode Hop-by-Hop function using anycast DA.

After local hop processing, egress anycast forwarding:

1. RPF check on SA, same as multicast
2. Exclude local node anycast DA instance
3. Forward based on remaining anycast routes

Hop-by-Hop Using Anycast Address

This is neither unicast or multicast forwarding.

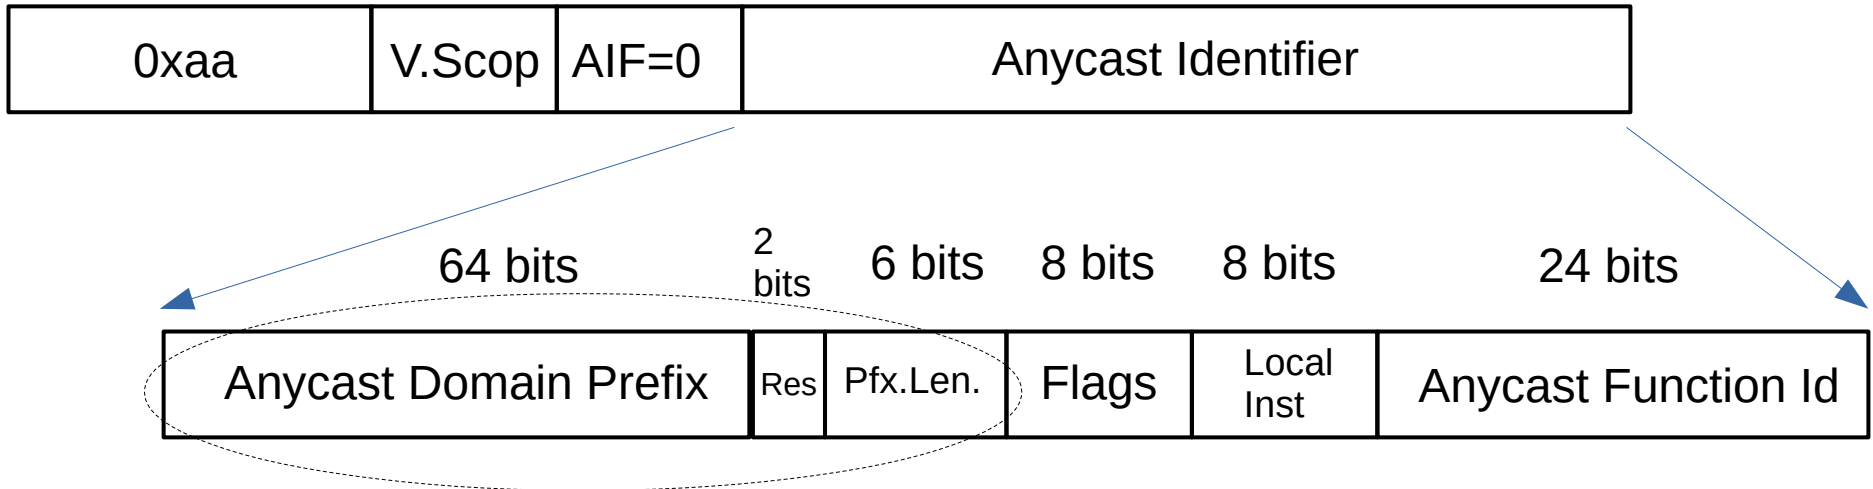
The series of hops is a series of IPv6 host hops.

RFC 8200:

router a node that forwards IPv6 packets not explicitly addressed to itself.

host any node that is not a router.

IPv6 Functional Anycast



Anycast Domain Prefix: Up to 64 bit Unicast prefix identifying anycast domain, inspired by multicast RFC 3306

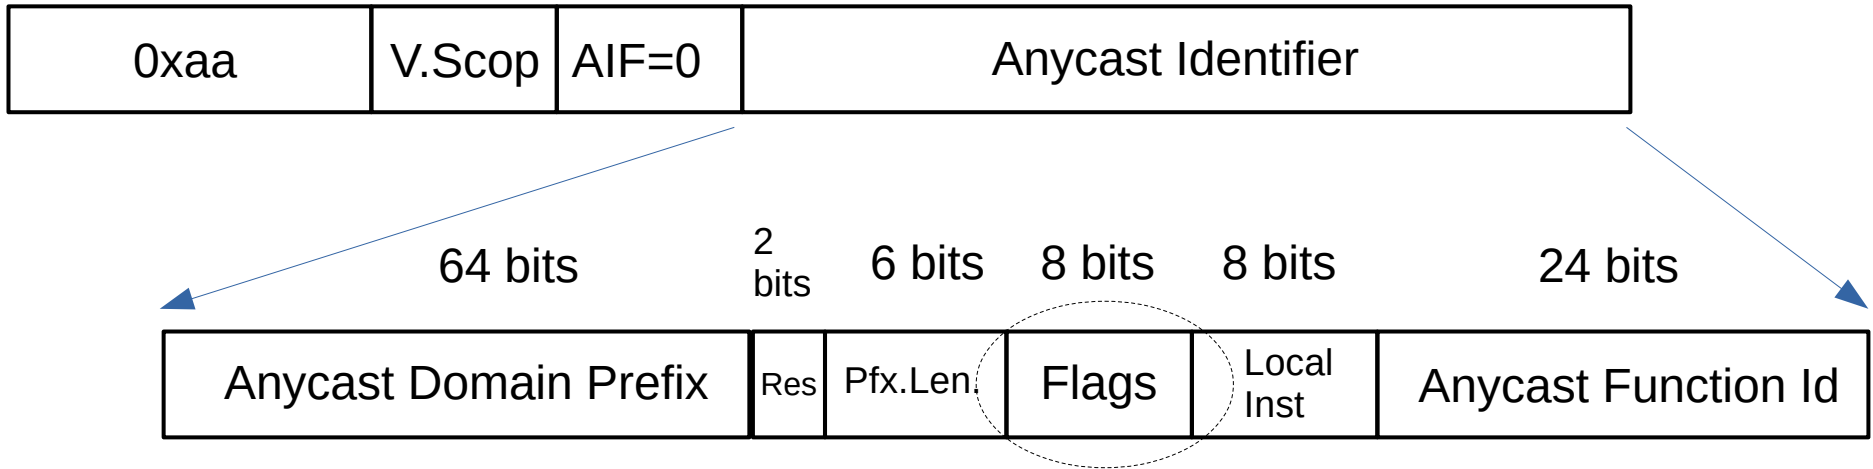
::/64 - "unspecified" or "this" domain

Anycast Domain routes can be aggregated here

Reserved: 2 bits reserved, zero upon send, ignored upon receipt.

Pfx.Len: 6 bit Anycast Domain Prefix length, 0 = 64. Informational.

IPv6 Functional Anycast



Flags: 8 bit flags field, inspired by multicast address Flags field

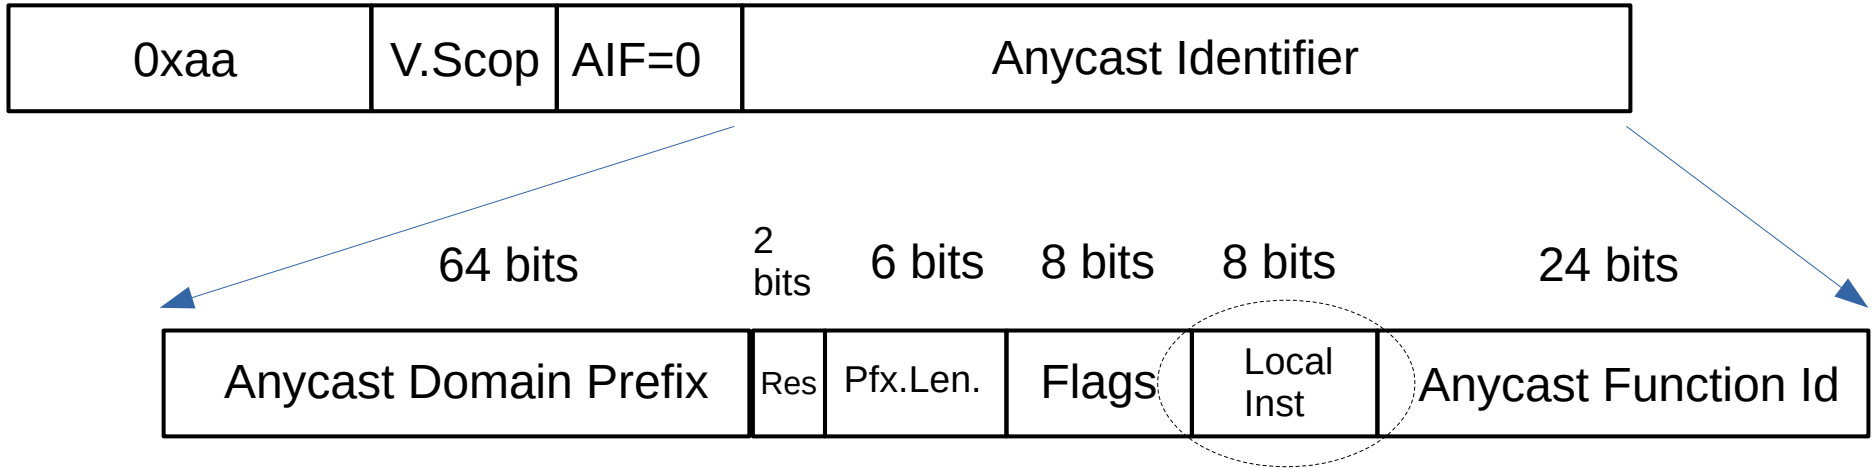
High Order T bit = Transient Anycast Function Id

T = 0, Well Known, IANA Assigned

T = 1, Transient, local anycast domain assigned

(i.e. same meaning as T bit in multicast Flags)

IPv6 Functional Anycast



Local Instance:

8 bit local instance, version or revision of following Anyc. Func. Id

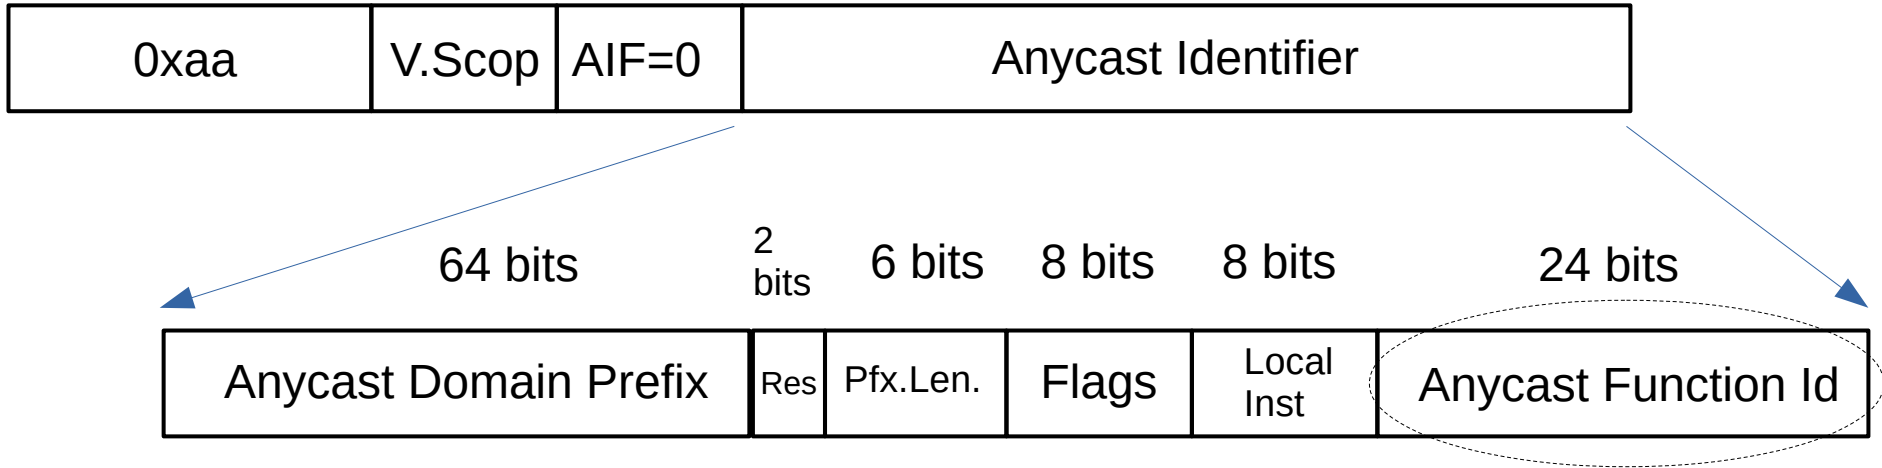
e.g. Dev function, Local Inst =1; Prod function, Local Inst = 0

Move Dev to Prod, change Local Inst from 1 to 0.

Default value of zero

Can be used to increase Anyc. Func. Id to 32 bits if desired.

IPv6 Functional Anycast



Anycast Function Identifier:

24 bit Function or Service Identifier

T=0 flag, IANA well known

T=1 flag, Local anycast domain assigned

Could be simpler?

Value of Prefix Length field a bit debatable to me.

With only 1 flag value, perhaps encode it somehow somewhere else?

Example Use Case 1 – Internet DDoS Impervious ISP Anycast DNS Resolvers

Requirements:

Reachable to all ISP's customers

Not reachable from the Internet

Attributable to the ISP

Example Use Case 1 – Internet DDoS Impervious ISP Anycast DNS Resolvers

Formal Anycast: aa00::/8

Visible Scope: 0xb - Network Service Provider (New “Multicast” Scope, >
Organization, < Global)

Anycast Identifier Format: 0x0 – Functional Anycast

Example Use Case 1 – Internet DDoS Impervious ISP Anycast DNS Resolvers

Anycast Domain Prefix: 2001:db8::(/64) – really an ISP's GUA

Prefix Length: 0x0 (/64)

Flags: 0x0 (T = 0, IANA assigned well known AFI)

Local Instance: 0x00 (Production)

Anycast Function Identifier: 0x000053, 0x000054, 0x000055

(IANA Well Known for DNS1, DNS2, DNS3 resolvers)

Example Use Case 1 –
Internet DDoS Impervious ISP Anycast DNS Resolvers

aab0:2001:db8::0053

aab0:2001:db8::0054

aab0:2001:db8::0055

Example Use Case 2 - Organization's Thingo service in development

Formal Anycast: aa00::/8

Visible Scope: 0x8 - Organization

Anycast Identifier Format: 0x0 – Functional Anycast

Example Use Case 2 - Organization's Thingo service in development

Anycast Domain Prefix: fdxx:yyyy:zzzz:1234:(/64) – org's ULA

Prefix Length: 0x0 (/64)

Flags: 0x80 (T = 1, Local organizations AFI)

Local Instance: 0x6a (IETF 106 revision)

Anycast Function Identifier: 0x544e47 (Thingo service)

Example Use Case 2 - Organization's Thingo service in development

aa80:fdxx:yyyy:zzzz:1234:0080:6a54:4e47

Status

Idea mostly baked I think.

Been thinking about it and working on draft for around 3 years to see if it would pan out – a /8 is a bold ask.

Discovering “Host-based Anycast using MLD” suggests others have thought of similar.

I’m at IETF-106, so present!

Thoughts?

Questions?