

Extending ICMP for Node Identification

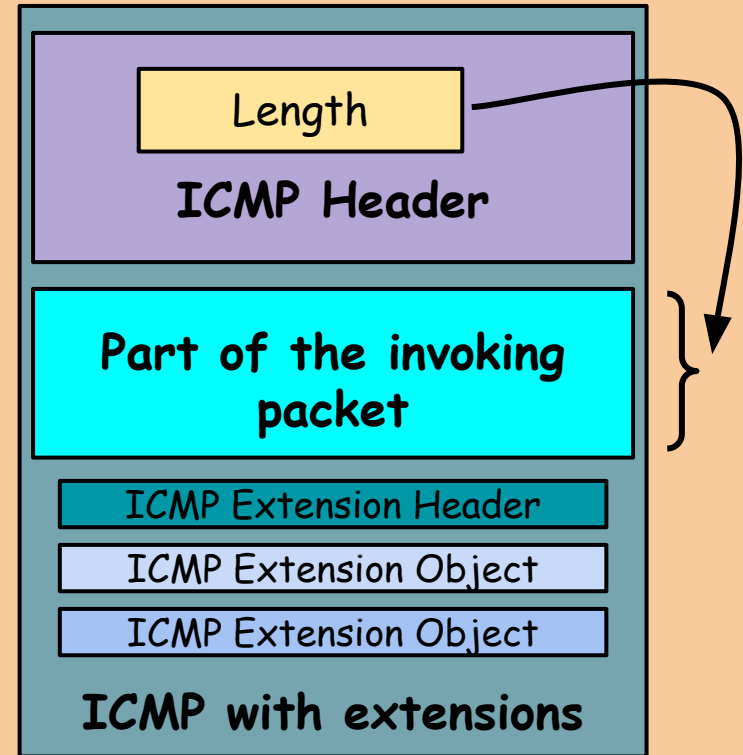
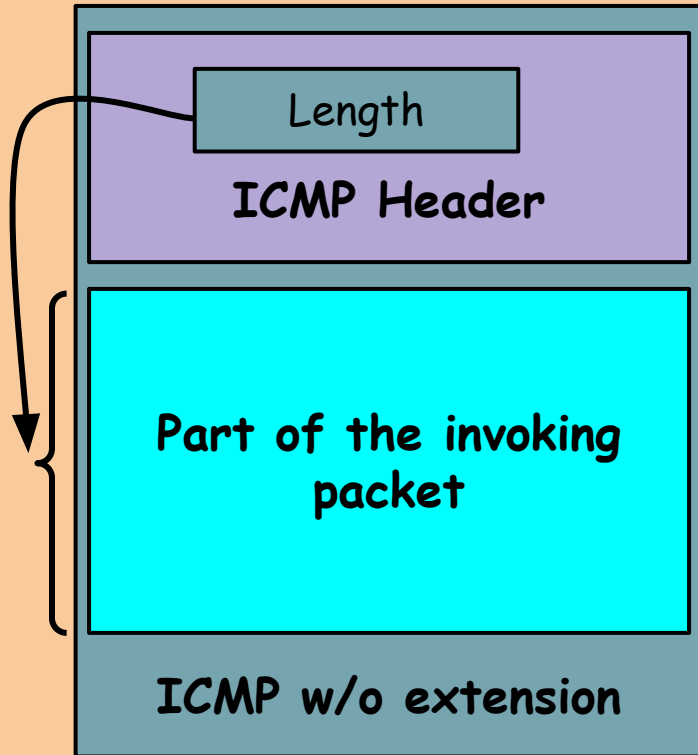
[draft-ietf-intarea-extended-icmp-nodeid](#)

Jen Linkova on behalf of Bill Fenner , Reji Thomas
IETF121, Nov 2024

IPv6-Only Networks

- Routers can be IPv6-only or have duplicated IPv4 addresses
- IPv4 traceroute can't represent IPv6-only hops
 - [draft-chroboczek-intarea-v4-via-v6](#)
 - [draft-equinox-v6ops-icmpext-xlat-v6only-source](#)

Quick Recap: ICMP Extensions (RFC4884)



Node Identification Object

- C-value: 5 (already allocated)
- Can contain two sub-objects (see next slides)
- Can be added to:
 - ICMPv4 Time Exceeded
 - ICMPv4 Destination Unreachable
 - ICMPv4 Parameter Problem
 - ICMPv6 Time Exceeded
 - ICMPv6 Destination Unreachable

3.2. Node IP Address Sub-Object

If the Node Identification Object identifies the node by address, the Object Payload contains an address sufficient to identify the node within the appropriate scope - global or as otherwise configured - as depicted in [Figure 2](#).




Figure 2: Node Identification Object - C-Type 2 Payload

Payload fields are defined as follows:

- **Address Family Identifier (AFI):** This 16-bit field identifies the type of address represented by the Address field. Values for this field represent a subset of values found in the IANA registry of Address Family Numbers (available from [IANA.address-family-numbers](#)). Valid values are 1 (representing a 32-bit IPv4 address) and 2 (representing a 128-bit IPv6 address).
- **Reserved:** This field MUST be set to 0 and ignored upon receipt.
- **Address:** This variable-length field represents an address of appropriate scope (global, if none other defined) that can be used to identify the node.

3.3. Node Name Sub-Object

Figure 3 depicts the Node Name Sub-Object:




Figure 3: Node Identification Object Node Name Sub-Object

The Node Name Sub-Object MUST have a length that is a multiple of 4 octets and MUST NOT exceed 64 octets.

The Length field represents the length of the Node Name Sub- Object, including the length and the node name in octets. The maximum valid length is 64 octets. The length is constrained to ensure there is space for the start of the original packet and additional information.

The second field contains the human-readable node name. The node name SHOULD be the sys:hostname [RFC7317], if less than 64 octets, or the first 63 octets of the sys:hostname, if the sys:hostname is longer. The node name MAY be some other human-meaningful name of the node. The node name MUST be padded with ASCII NUL characters if the object would not otherwise terminate on a 4-octet boundary.

The node name MUST be represented in the UTF-8 charset [RFC3629] using the Default Language [RFC2277].

(Planned) Changes Since Adoption

<https://github.com/fenner/icmp-node-id/>

- Security considerations:
 - Including the Node Identification
 - Configurable
 - Off by default
 -except for IP/ICMP translators (SHOULD be enabled)

Next Steps?

- Feedback?
- Anything else to change/add?

Using Dummy IPv4 Address and Node Identification Extensions for IP/ICMP translators

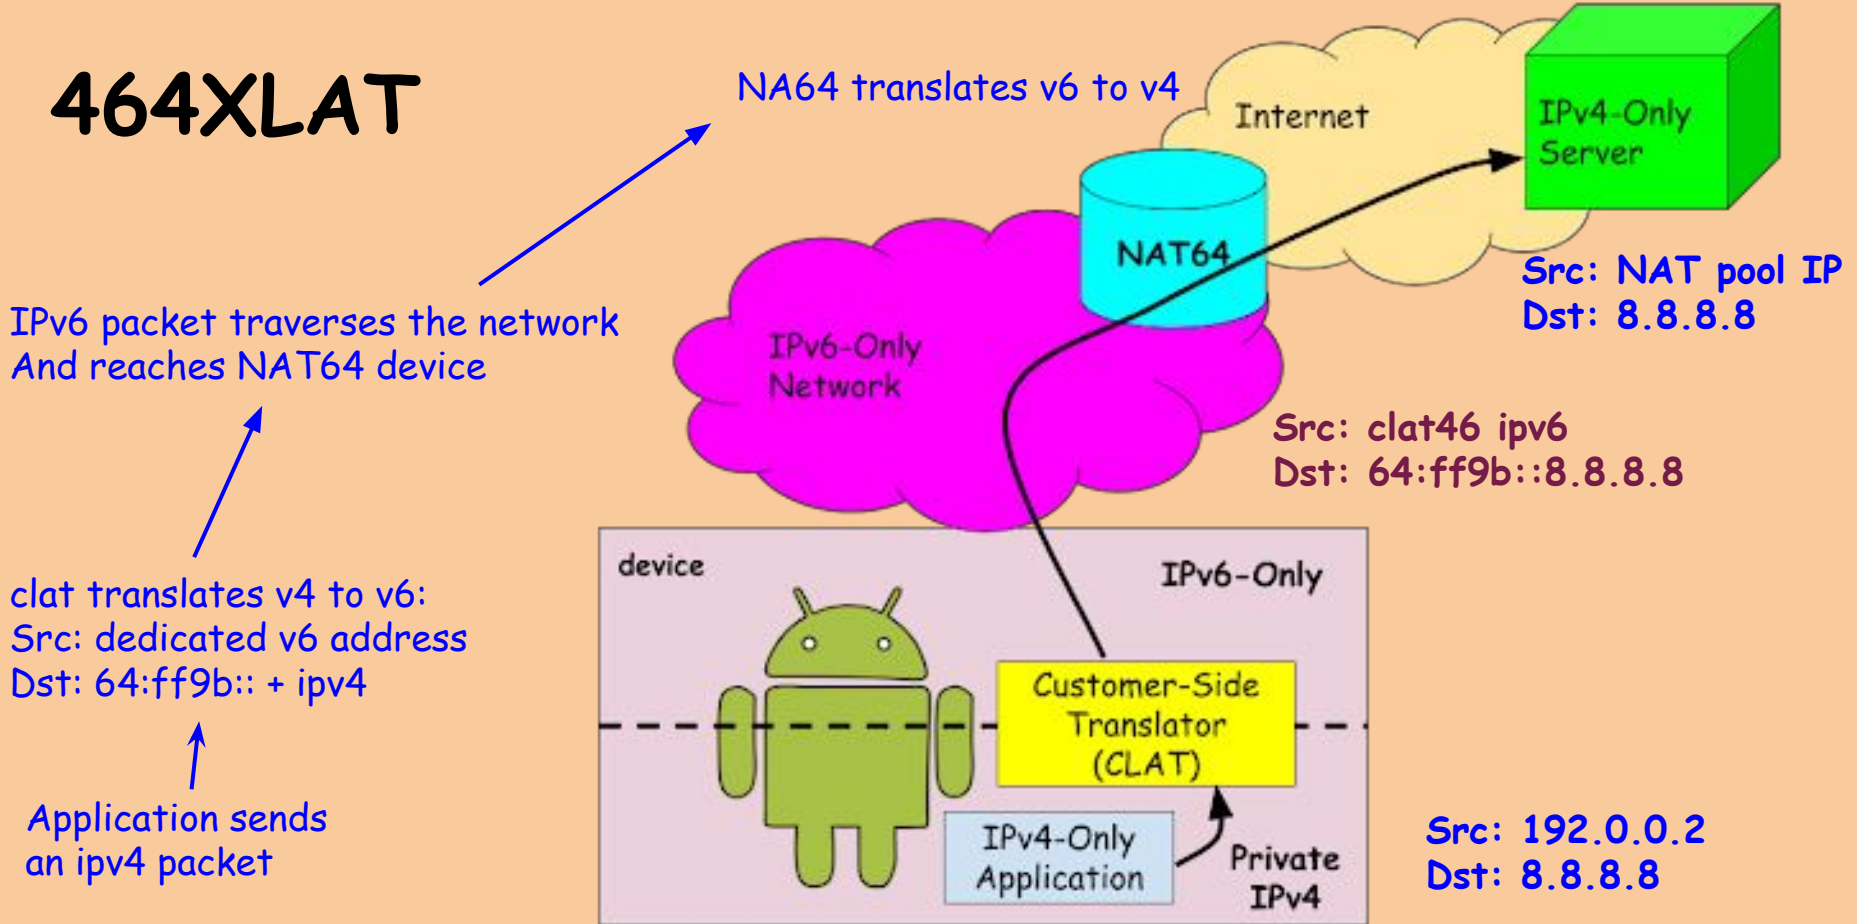
[draft-equinox-v6ops-icmpext-xlat-v6only-source](#)

(was: [draft-equinox-intarea-icmpext-xlat-source](#))

David 'equinox' Lamparter, Jen Linkova
IETF121, Nov 2024

Problem Statement

464XLAT



Dual-stack INTERNET

ICMPv4
src: 198.51.100.1

NAT64 (PLAT), prefix 64:ff9b::/96

IPV6-Only Network

ICMPv6
src: 64:ff9b::198.51.100.1

ICMPv6
src: 2001:db8:a::a

IPv6-only host, running CLAT, prefix 64:ff9b::/96

ICMPv4, src: 198.51.100.1

ICMPv4, src: ????????

problem!

Operational Issues

- Traceroute can not represent IPv6-only hops
 - Confused users
 - Unhappy NOC
- ICMPv6 Packet Too Big are not translated
 - Broken PMTUD

Proposed Solution

Terminology: Untranslatable

- “Untranslatable” IPv6 address:
 - doesn't belong to NAT64 prefix
 - there is no explicit address mapping

ICMPv6 -> ICMPv4 Translation

1. Update RFC7915 (IP/ICMP Translation Algorithm):
 - a. Define how to translate "untranslatable" IPv6 addresses
 - b. Preserve information about the original IPv6 address

Translating the Untranslatable


Whenever a translator translates an ICMPv6 Destination Unreachable, ICMPv6 Time Exceeded or ICMPv6 Packet Too Big ([RFC4443]) to the corresponding ICMPv4 ([RFC0792]) message, and the IPv6 source address in the outermost IPv6 header is an untranslatable one, the translator SHOULD use the dummy IPv4 address (192.0.0.8) as IPv4 source address for the translated packet.

Address Block	Name	RFC	Allocation Date	Termination Date	Source Destination Forwardable			Globally Reachable	Reserved-by-Protocol
192.0.0.0/24 [2]	IETF Protocol Assignments	[RFC6890], Section 2.1	2010-01	N/A	False	False	False	False	False
192.0.0.0/29	IPv4 Service Continuity Prefix	[RFC7335]	2011-06	N/A	True	True	True	False	False
192.0.0.8/32	IPv4 dummy address	[RFC7600]	2015-03	N/A	True	False	False	False	False

Preserving the Original IPv6 Source Address

Node Identification Object

Node IP Address Sub-Object



Please Read the Draft

- Details on the translator behaviour:
 - adding new Extension Structure
 - appending extension object to existing structure
- Updates to RFC7915 (translation algorithm):
 - "When translating ICMPv4 -> ICMPv6, follow recommendations in draft-equinox-v6ops-icmpext-xlat-v6only-source"

Open Issues

- Current text contains many implementation details:
 - When to add an extension object, before or after the translation?
 - Before: easier to translate
 - After: can be added when ICMPv6 PTB is translated to ICMPv4 "Fragmentation Needed"
- Are we over specifying?
 - Maybe just define how the packet must look like?

Comments?