Identifying Nodes in ICMP Extended Errors

Bill Fenner
IETF 119

What's the problem?

More and more deployments are assigning only a single IPv4 address to routers (instead of one address per interface), and duplicating addresses used throughout the network due to scarcity.

An extreme example from draft-chroboczek-intarea-v4-via-v6:

```
$ traceroute -n 8.8.8.8
traceroute to 8.8.8.8 (8.8.8.8), 64 hops max, 52 byte packets
 1  192.168.0.1  1.894 ms  1.953 ms  1.463 ms
 2  192.0.0.8  9.012 ms  8.852 ms  12.211 ms
 3  192.0.0.8  8.445 ms  9.426 ms  9.781 ms
 4  192.0.0.8  9.984 ms  10.282 ms  10.763 ms
 5  192.0.0.8  13.994 ms  13.031 ms  12.948 ms
 6  192.0.0.8  27.502 ms  26.895 ms
 7  8.8.8.8  26.509 ms
```

Existing solutions

RFC5837 - Extending ICMP for Interface and Next-Hop Identification - allows inserting information about incoming and outgoing interfaces into an ICMP error.

Incoming Sub-IP	ifIndex	ifName	~~ifAddr~~	MTU
Incoming IP	ifIndex	ifName	ifAddr	MTU
Outgoing IP	ifIndex	ifName	ifAddr	MTU
Outgoing Nexthop	~~ifIndex~~	~~ifName~~	ifAddr	~~MTU~~

Limitations of RFC5837

- Addresses MUST be of the same family as the packet
  - E.g., RFC5549 deployments, with IPv6 nexthops for IPv4 routes, do not have an address assigned to the interface that can be used.
- Addresses require another lookup to determine node name

Proposal

https://datatracker.ietf.org/doc/draft-fenner-intarea-extended-icmp-hostid/

Add more information types like those specified in RFC5837:

- An IP address of appropriate scope (“globally-unique” is the most obviously appropriate)
- A node name

Workarounds

- Violate MUST in RFC5837, and include an IPv6 address when responding to an IPv4 packet
- Play fast and loose with “other human-meaningful name” wording from RFC5837; this interface is named “TwentyFiveGigE1/0/13@be2264.ccr22.sjc01.example.com”

  `The interface name MAY be some other human-meaningful name of the interface.`

Next steps

Gauge interest

Apply for code point