

ERROR-HANDLING -13

John Scudder

July 21, 2014



DRAFT STATUS

WGGLC for draft -07 initiated May 7, 2014

Led to much discussion with many productive comments

Six drafts later, we have -13... and it's still not quite ready

Two remaining open issues:

- What's a "valid IP host address"?
- What to do with the Traffic Engineering path attribute?

WHAT'S A "VALID IP HOST ADDRESS"?

RFC 4271 requires that the NEXT_HOP be a "valid IP host address". (RFC 4760 inherits this.)

- But, it doesn't say what that means.

Nailing this down has proven surprisingly slippery.

- Clearly unreasonable to define from scratch
- There must be a reference to cite... right?

IANA REGISTRIES

IANA maintains an “IPv4 Special-Purpose Address Registry” and an “IPv6 Special-Purpose Address Registry”

- These have various attributes, including whether an address can be a “destination” and whether that address is “forwardable”
- If an address can’t be a destination or isn’t forwardable, that means it’s not a “valid IP host address”... right?

A TWISTY MAZE OF PASSAGES

But what about IPv4-mapped IPv6 addresses? (Used in VPNv6, 6PE)

- OK, make an exception for them.
- And punt on other AFI/SAFI.
- And exceptions should be configurable, of course.

Are we done yet?

- Nope!

... ALL DIFFERENT

Robert Raszuk points out that RFC 5575 (Flow-Spec) makes the next hop optional (sort of)

- So, do we continue to dot i's, cross t's, and add codicils to cover every permutation of next hop?

Chris Hall (June 14) does a nice analysis of errors as “venial”, “mortal” and “trivial”

- “all next-hop errors are venial errors. So the draft doesn't, AFAICS, need to get sucked into defining “valid host address”.

ARE ALL NEXT-HOP ERRORS “VENIAL”?

In MP_REACH, next-hop length defines the beginning of the NLRI section

- So if the next-hop length is invalid, maybe we can't properly find the NLRI, so can't do treat-as-withdraw. That's a “mortal” error (session reset).

This comes back to an underlying assumption of the draft: generally confine ourselves to syntax errors, let purely semantic ones go.

WHEN IN DOUBT, PUNT

Jeff Haas proposes (June 16), radical simplification, rather than trying to dot every i and cross every t:

- “If the next hop field contains a semantically incorrect address within the context of deployed features and address family, treat as withdraw behavior should be used. ”

And I would add, if syntactically incorrect in the same context (wrong length), session reset.

PROPOSAL

Adopt something like the language on the preceding slide

Accept that the document will not provide a detailed prescription for every case

- We somehow have muddled through 25 years of BGP anyway

Remove “martians” discussion entirely

TE PATH ATTRIBUTE

The draft tries to cover every extant BGP path attribute (that doesn't already have compliant spec language and isn't deprecated)

For most attributes, this was straightforward.

Not so for the TE path attribute (RFC 5543).

When in doubt, punt!

PROPOSAL FOR TE PATH ATTRIBUTE

“an implementation that determines (for whatever reason) that an UPDATE message contains a malformed Traffic Engineering path attribute MUST handle it using the approach of “treat-as-withdraw””

- Possibly, attribute discard would be OK (if the TE attribute is strictly an optimization)

This language is already in -13

NEXT STEPS

Update -14 with new, reduced “valid IP host address” section

New WGLC