Multi6 WG and DT status

Presenter: Erik Nordmark

Multi6 status

- Done Goals for a multihoming solution as RFC - RFC 3582
- Done Final solicitation of proposals
- Done Begin architectural evaluation of proposals
- Done First draft of architectural evaluation
- Oct 04 Submit informational I-D to IESG on how multihoming is done today
  - just finished WG Last Call - revision needed
- Oct 04 Submit informational I-D to IESG on security threats
  - in AD's hands, sort of - revision needed
- Nov 04 Submit informational I-D to IESG on architectural evaluation
  - just finished WG Last Call - revision needed
- Dec 04 Identify proposal(s) for further development, recharter
  - we are here
- Jan 05 Submit informational I-D to IESG on practical questions
  - just finished WG Last Call - revision needed

Design Team status

- DT formed at the San Diego IETF
  - Look at L3 shim approach
- Members: J. Arkko, I. van Beijnum, M. Bagnulo, G. Houston, E. Nordmark, M. Wasserman, J. Ylitalo
- Delivered 5 I-Ds with name **-multi6dt-**
- Was discussed in multi6 WG this week
  - Will form basis for architecture draft
  - A separate WG in internet area likely to work on specifying the protocol

What did we try to accomplish?

- Minimal or no additional dependency on DNS
  - Work for hosts without FQDNs
- An approach which allows application referrals to work
- Good enough security
  - Avoid time-shifting attacks if possible
- Think about privacy concerns
- Supports or extensible to handle mobility
- Think about avoiding hard /64 bit boundary

Design Team approach (1)

- A L3 shim between IP endpoint and routing sub-layers
  - Below fragmentation, IPsec
  - Provide “service” to all transport protocols
- No new ID name space
  - AAAA records contain same thing as today
  - Applications/transports use “upper-layer ID”
    - Any one of the locators from the AAAA RRset
    - Doesn't change during the connection
  - Shim switches locators when a failure

Design Team approach (2)

- Using Hash-based addresses (or CGA) to prevent redirection attacks
  - When host has a fixed set of addresses, the verification is just a hash computation
  - Changing set of addresses require using CGA i.e., verification using public-key crypto
- Testing/probing to find a working locator pair after a failure
  - Due to interaction between ingress filtering and routing the locator pairs might need to be different in the two directions
  - Can handle this without much additional complexity

Issues from the DT

- Need to handle ingress filtering
  - Exit router selection based on source address for small sites?
  - Non-DT draft addresses this
    - draft-huitema-multi6-ingress-filtering-00
- Actual packet formats
  - Overloading flow label vs. adding 8 byte extension header after rehoming
- Interaction with applications and transport protocols
  - Started discussion at open apps area meeting

Interesting things we haven't explored in depth (1)

- State management
  - What exactly identifies the multi6 context state?
  - Do the peers coordinate when they discard the state?
- Using non-reachable locators as ULIDs
  - Example: ULAs
  - Nothing in the approach and drafts prevents this
- Apps using DNS reverse and forward for non-reachable locators?
  - There might be issues about DNS and interaction with IPv6 nodes that are not multi6 aware

Interesting things we haven't explored in depth (2)

- Handle subnet prefixes with more or less than 64 bits
  - No text about this yet
  - Unclear whether broader community is concerned about hard-coding the /64 boundary forever

Other things needed

- Need some understanding of what policy controls should (and can) be provided when using multiple, provider-allocated address prefixes
  - In IPv4 with provider independent address BGP provides tools to do this
  - With multiple, aggregated PA prefixes things are different
  - If you are interested in this please get involved