

shim6: what now?

IETF-71, March 10, 2008
Φιλαδέλφεια, PA, US

Hjitsch van Beijnum

Are we done?

- Close to publication:
 - HBA: ✓
 - Proto: ✓
 - REAP: ✓
- That's it?

Not quite

- Still unaddressed:
 - ingress filtering
 - "initial failures"
 - traffic engineering
 - proxy implementation

Ingress filtering

- We have:
 - an ICMPv6 "wrong source addr" message
 - Marcelo Bagnulo's/Christian Huitema's old drafts about tunnels to egress router
- We don't have:
 - source address dependent routing
 - address rewriting in routers

Ingress filtering (2)

- Without something for this, shim6 doesn't work in practice!
- So:
 - write down what we know
 - or do more?

"Initial failures"

- When a site has multiple locators
- We pick the wrong one initially
- Solve this:
 - we don't: apps try different addresses?
 - mapping mechanism?
 - reachability detection?

Traffic engineering

- Two issues:
 - do it in the first place
 - centralized control of TE
- Do we need extra protocols?
- Use the DNS?

Proxy shim6

- Makes for easy shim6 deployment
- Enterprises like central control
- Easier to do traffic engineering

Rewriting in routers

- Attractive:
 - solves ingress filtering
 - get hint for free
- Quantum behavior of addresses: can change it as long as nobody has looked at it
- But: some packets may be rewritten, not others
- Unless.... we do NAT

NAT

- NAT is bad
- Bad NAT is much worse than less bad NAT
 - In IPv6 no need to overload ports
 - I-to-I NAT: incoming sessions work

NAT (2)

- Can compensate for checksum difference when rewriting address:
 - works with TCP, UDP, ICMPv6, DCCP
 - RTP: no checksum (but on UDP usually)
 - SCTP: CRC32, but not over address
- So all (?) protocols work **statelessly**
- Referrals and path changes still break, though

NAT (3)

- NAT allows for easy proxy shim6
- Do I-to-I NAT grudgingly rather than end up with port overloading NAT?

Flame away!