

IPv6 multihoming status

The follow up to Manchester...



© 2006 - Netnod AB
<http://www.netnod.se/>

2006-01-10

Background

- This is more or less my personal take...
 - But I happen to be the co-chair of shim6 with Geoff Huston
- I made a similar presentation at RIPE51 and Geoff at Nanog

The problem



- Some end sites want multiple connections to different upstreams for
 - Resilience
 - Renumbering avoidance
- This does not *HAVE* to imply multiple upstream providers
 - But it can

The problem

- Routing system constraints
 - In order for the multiple upstreams to forward traffic to the end-site, a unique identifier is needed for the longest-prefix-match algorithm
- In IPv4 this is either of
 - PI address block
 - “more specific” PA
 - Multiple addresses on each node

Multihoming Today

- AS + PI
 - Get an AS
 - Get PI space
 - Advertise and use BGP
- More specific PA
 - Advertise the more specific route



Effects of Multihoming

- Leads to “uncontrolled” growth of the routing table
 - Can lead to problems in the future
- Would be better if each end-user/site could get a block from each provider
 - And be able to use both prefixes as source addresses in case of failures
 - Today this does not work due to inbound-filtering at the ISPs

The IETF effort

- The multi6 WG was tasked with inventorying possible solutions
 - And benchmarking/selecting a solution
 - Selected an architecture based on separating locator / identifier
- Work on protocol is moved to the shim6 WG

The SHIM6 Solution

- host-based solution (rather than host and router)
- network layer (rather than transport)
- discoverable negotiated capability
- no new identifier space



The SHIM6 Approach

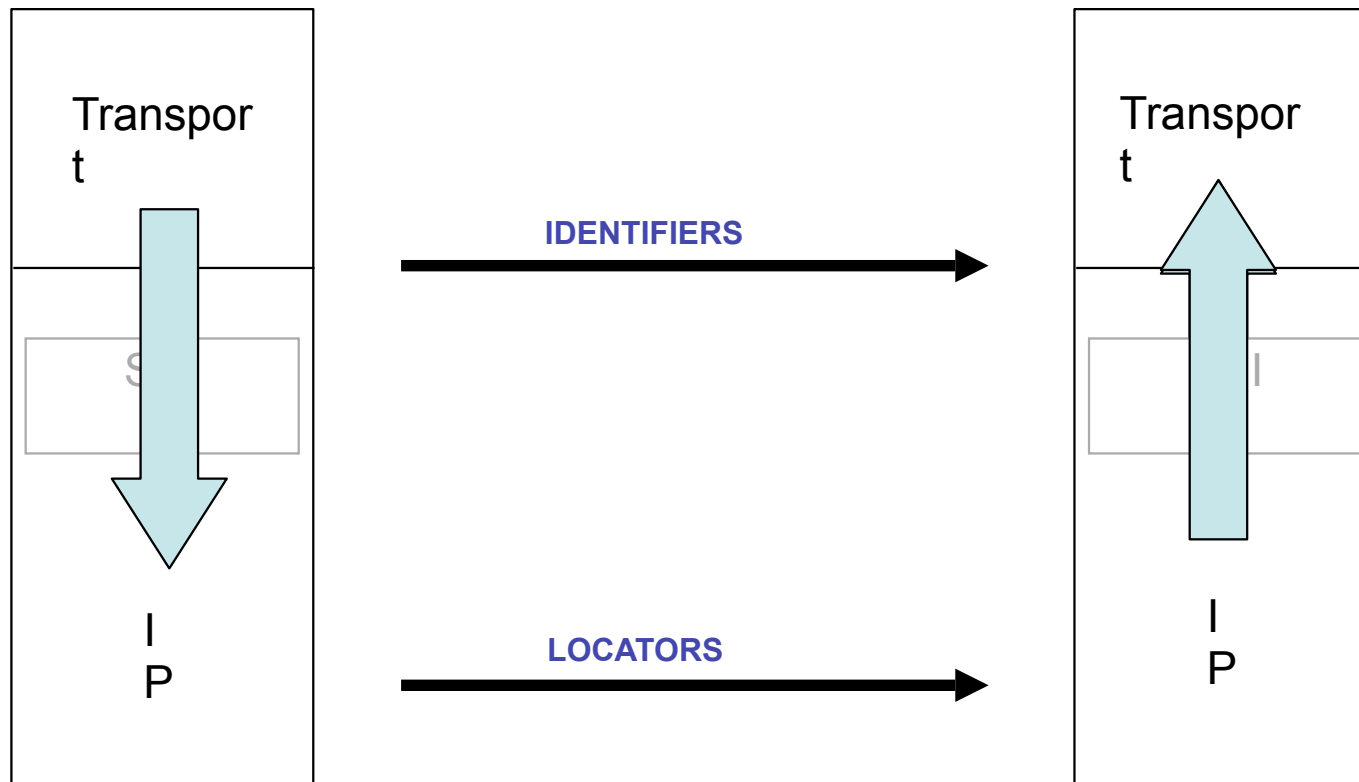
- a functional module at layer 3 (IP)
- the initial locator is the upper layer identifier (RFC3484 selection)
- subsequent negotiation to enable the Shim6 module for an upper layer identifier pair
- the Shim6 module translates upper layer identifiers into the currently active forwarding layer locators
- the upper layer identifier pair plus a context value forms the shared shim6 state identifier
- an IPv6 end-to-end header is used to signal SHIM6 context

shim6 - protocol

- Current thinking is that the base header will look remarkably like a HIP header
 - but it is *NOT!*
- Some issues are still TBD but we have come a far way....


Initial Contact

No SHIM state active
Locator Selection using RFC3484
Locators and Identifiers are Equivalent



SHIM6 Activation

SHIM active
Current Locator Sets exchanged
Locators and Identifiers are Equivalent




IDENTIFIERS



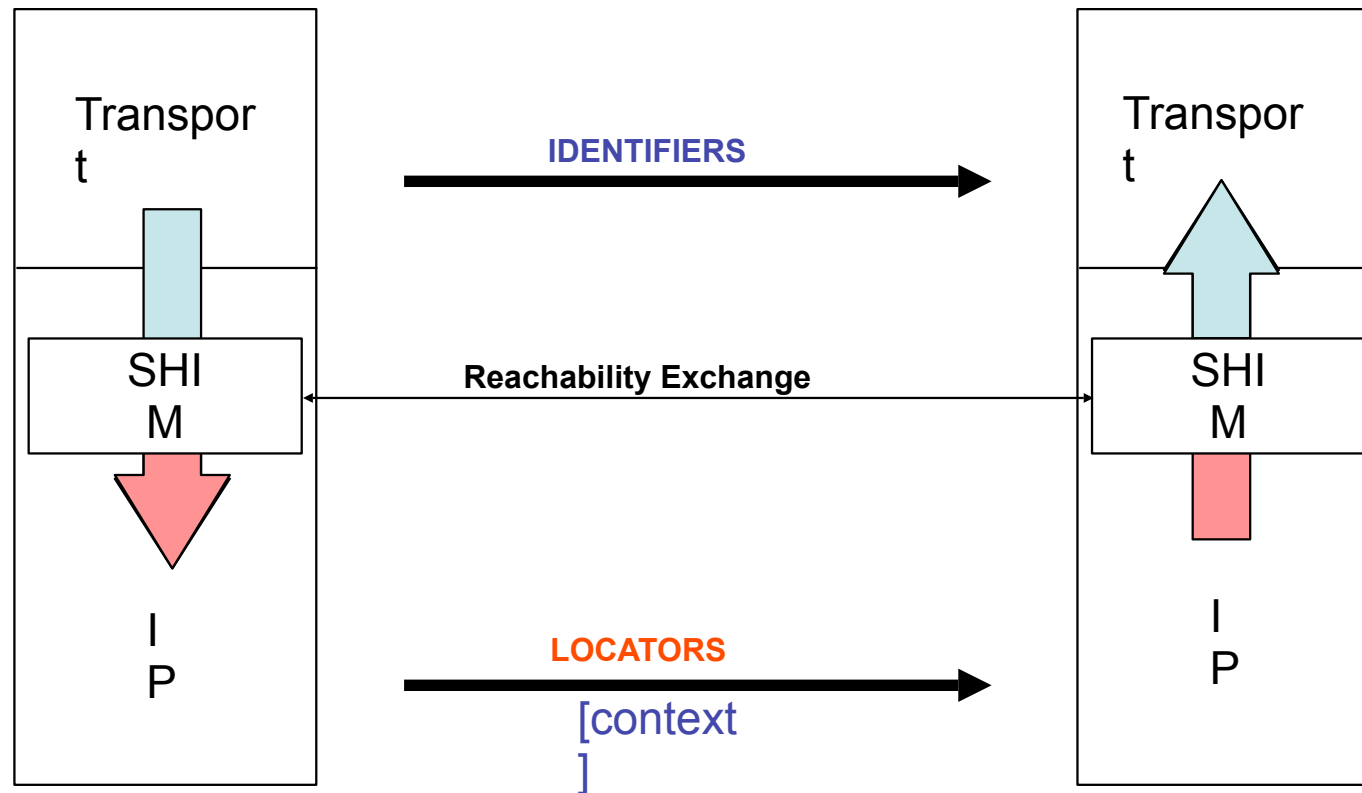
LOCATORS

[context
]

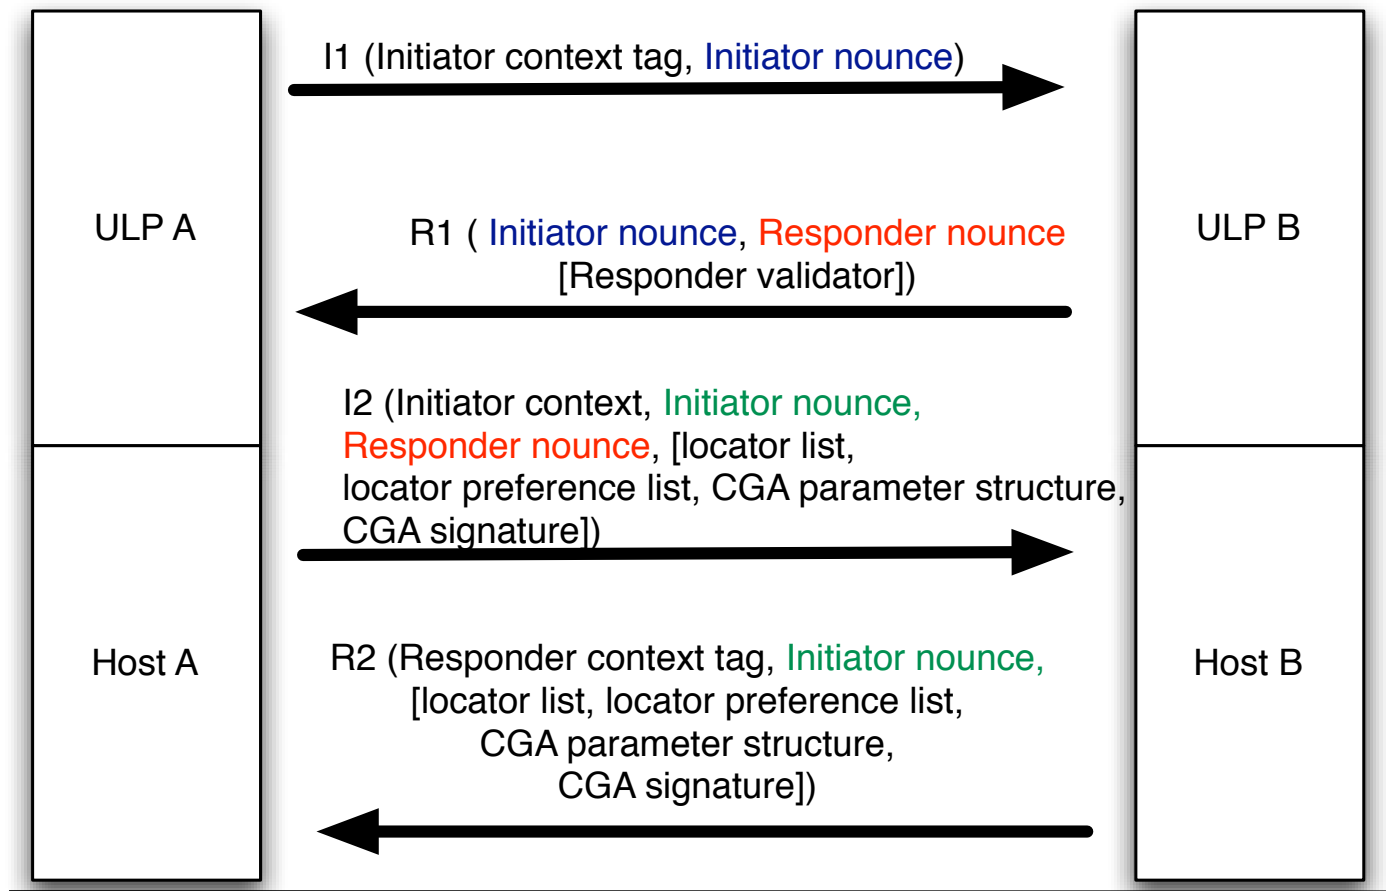


SHIM6 Locator Failure and Recovery

Detect locator failure
Explore for functioning locator pair
Use new locator pair – preserve identifier pair



Shim6 - protocol



2005-10-11

SHIM6 Control Elements

- initial handshake (4-way) and locator set exchange
- locator list updates
- explicit locator switch request
- keepalive
- reachability probe exchange
- No-Context error exchange

SHIM6 WG Approach

- base protocol specification
 - protocol exchange and packet formats
 - address specification: CGA and HBA
 - functional decomposition
- refinements
 - upper layer signalling
 - traffic engineering hooks
 - contactless shim6
 - failure detection refinements
 - ingress filtering / source address path selection

?