

An Update on Multihoming in IPv6 Report on IETF Activity

IPv6 Technical SIG
1 Sept 2004
APNIC18, Nadi, Fiji
Geoff Huston


Resiliency in IP

- How do you create a service that's available 100% of the time?
 - Use a server architecture and location environment that uses sufficient resiliency to provide 100% availability
 - Connect to the Internet using a service provider than can provide 100% guaranteed availability
- 100% network availability?
 - Multiple connections to a single provider?
 - No – there's a single routing state that is vulnerable to failure
 - Multiple Connections to multiple providers
 - More attractive, potentially allowing for failover from one provider to another in the event of various forms of network failure

Current approach


Either:

- Obtain a local AS
- Obtain PI space
- Advertise the PI space to all upstream providers
- Follow routing



Or:


- Use PA space fragment from one provider
- Advertise the fragment to all other upstream providers
- Follow routing



The cost of routing

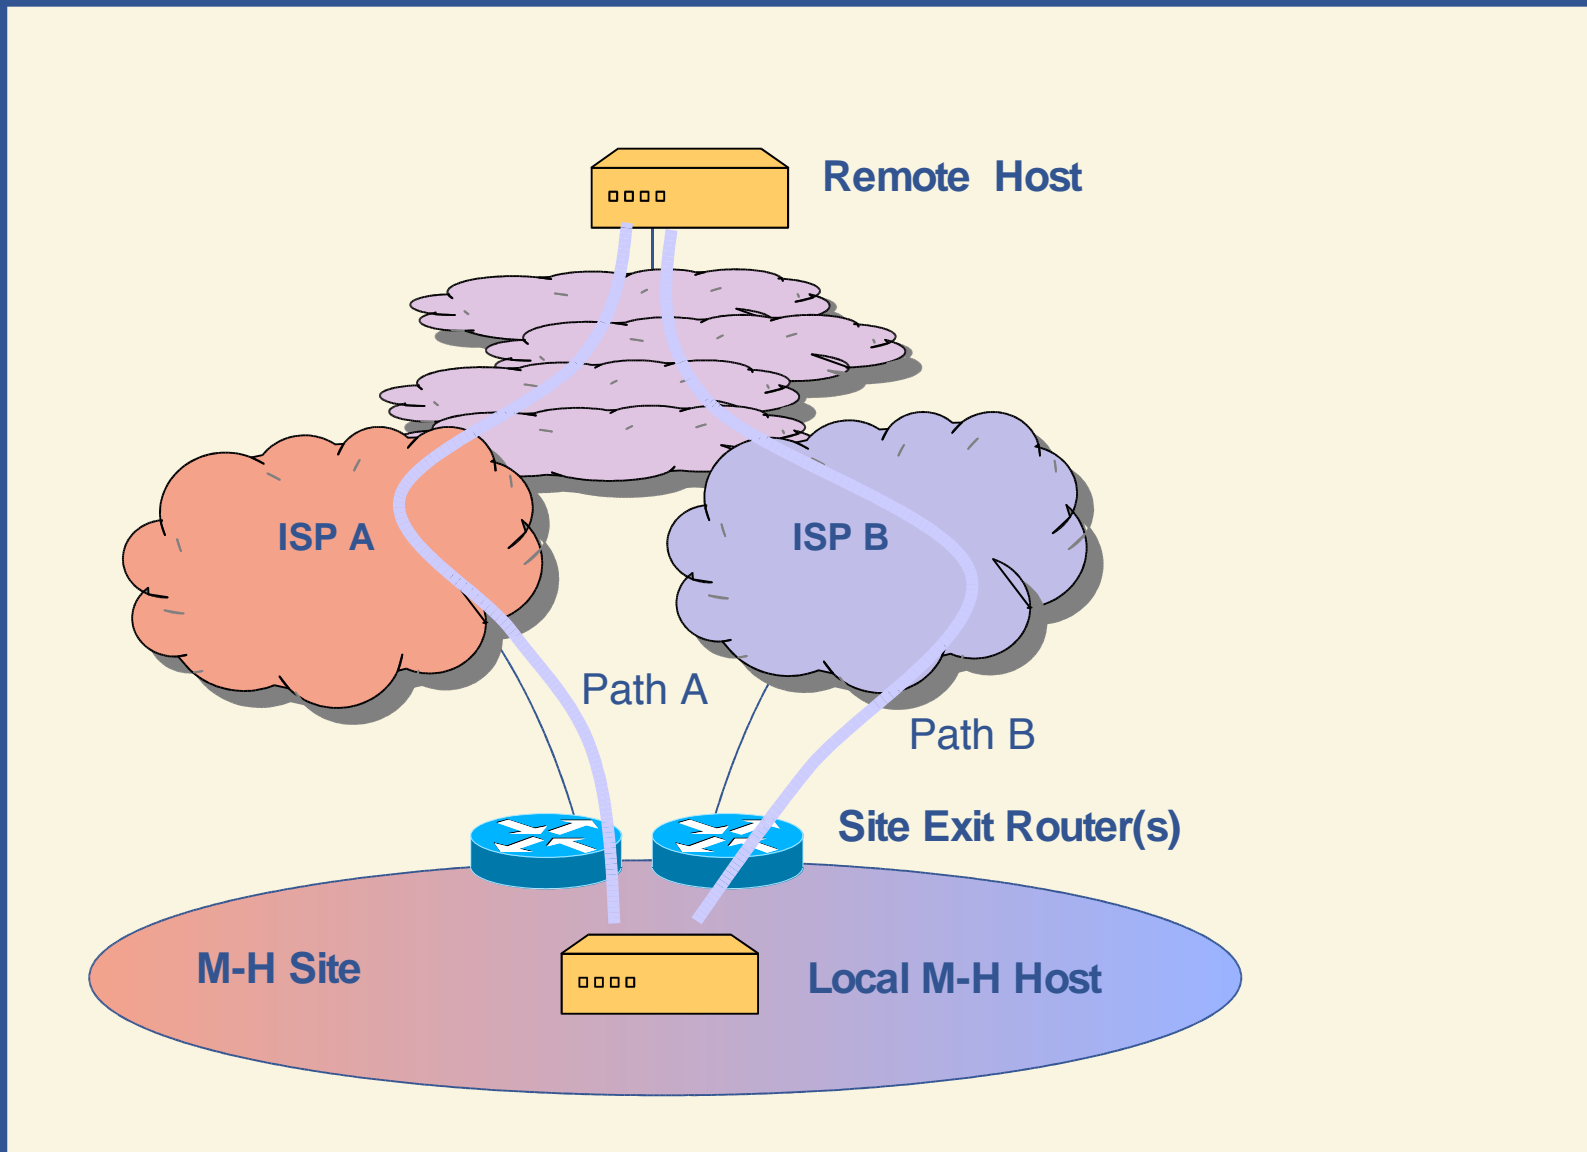
- This approach adds an additional entry into the routing system for each multi-homed end site
- The routing system is not an unbounded system
- Is there an alternative approach that can support multi-homing without imposing a massive load on the routing system?

What we would like...



- The multi-homed site uses 2 address blocks
 - One from each provider
- No additional routing table entry required

The problem space



Functional goals

- RFC3582 enumerates the goals as
 - Redundancy
 - Load Sharing
 - Traffic Engineering
 - Policy
 - Simplicity
 - Transport-Layer Survivability
 - DNS compatibility
 - Filtering Capability
 - Scalability
 - Legacy compatibility
- Also we need to think about
 - Interaction with routing
 - Aspects of an ID/Locator split, if used
 - Changes to packets on the wire
 - Names, Hosts, endpoints and the DNS

But this is not IP as we knew it

- The IP protocol architecture has made a number of simplifying assumptions
- One major assumption was that IP hosts didn't move!
 - Your IP address is the same as your identity (who)
 - Your IP address is the same as your location (where)
 - Your IP address is used to forward packets to you (how)
- If you want multi-homing to work then your identity (who) must be dynamically mappable to multiple locations (where) and forwarding paths (how)
 - “its still me, but my location address has changed”


The multi-homing plan

- For multi-homing to work in a scalable fashion then we need to separate the “who” from the “where”
 - Or, we need to distinguish between the identity of the endpoint from the network-based location of that endpoint
 - Commonly termed “ID/Locator split”

Generic approaches

- **Insert a new level in the protocol stack (identity element)**
 - New protocol element
- **Modify the Transport or IP layer of the protocol stack in the host**
 - Modified protocol element
- **Modify the behaviour of the host/site exit router interaction**
 - Modified forwarding architecture

New protocol element



- Define a new Protocol element that
 - Presents an identity-based token to the upper layer protocol
 - Allows multiple IP address locators to be associated with the identity
 - Allows sessions to be defined by an identity peering, and allows the lower levels to be agile across a set of locators

Benefits

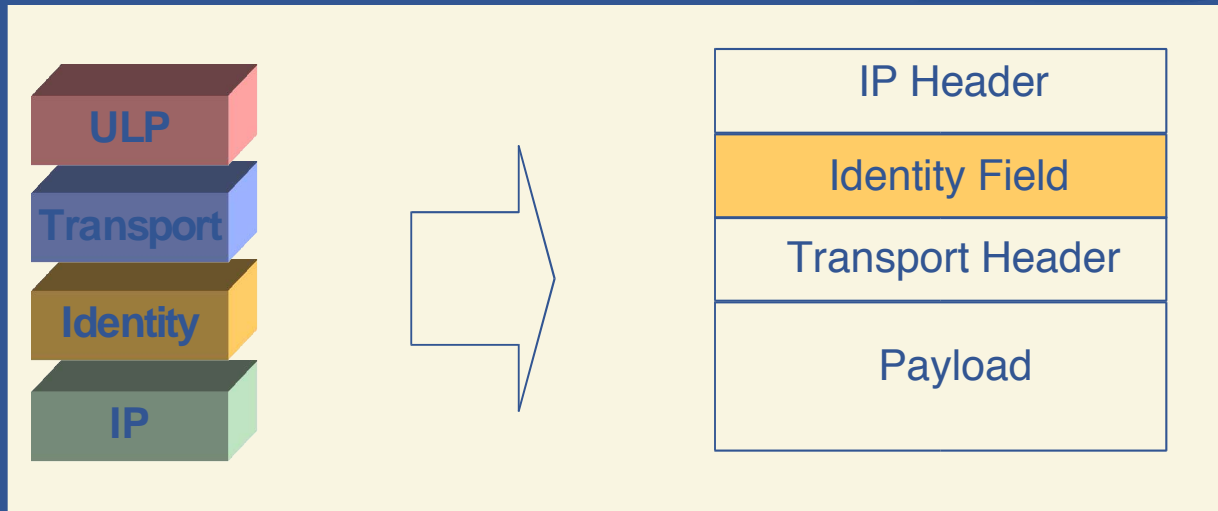
- Allow indirection between identity and location
- Provide appropriate authentication mechanisms for the right function
- Allow location addresses to reflect strict topology
- Allow identities to be persistent across location change (mobility, re-homing)

Identity protocol element



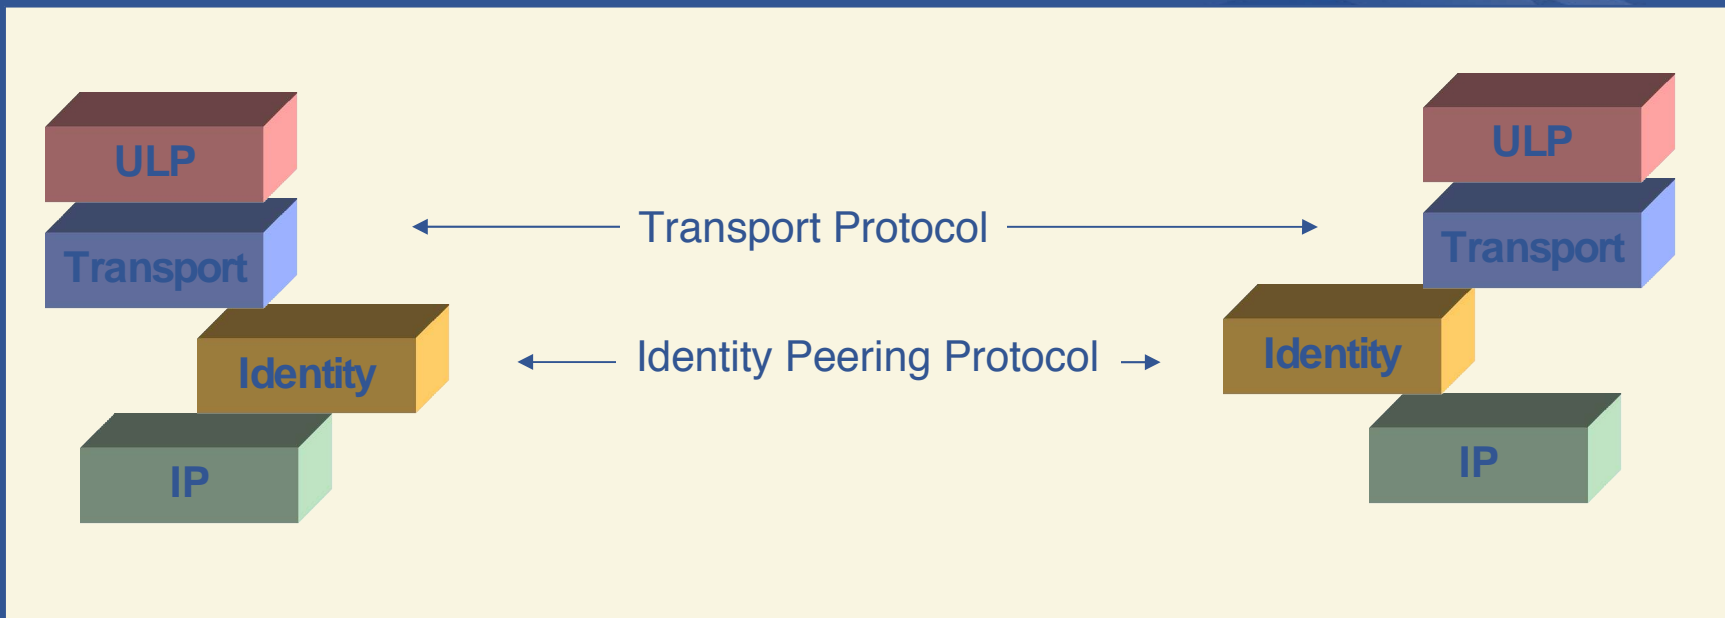
Protocol element implementation

- “Conventional”
 - Add a wrapper around the upper level protocol data unit and communicate with the peer element using this “in band” space




Protocol element implementation

- “Out of Band”
 - Use distinct protocol to allow the protocols element to exchange information with its peer





Protocol element implementation

- “Referential”
 - Use a reference to a third party point as a means of peering (e.g. DNS Identifier RRs)



Modified protocol element behaviour



- Alter the Transport Protocol to allow a number of locators to be associated with a session
 - e.g. SCTP
- Alter the IP protocol to support IP-in-IP structures that distinguish between current-locator-address and persistent-locator-address
 - i.e. MIP6

Modified host / router interaction

- Modify the interaction between the host and the Site Exit router to allow
 - Source-based routing for support of host-based site-exit router selection
 - Site Exit router packet header modification
 - Host / Site Exit Router exchange of reachability information

Identity protocol element location

- It appears that the proposals share a common approach
 - Above the IP forwarding layer (Routing)
 - Below IP fragmentation and IPSEC (IP Endpoint)



Proposals for an identity protocol element

- Use identity tokens lifted from a protocol's "address space"
 - DNS, Appns, Transport manipulate an "address"
 - IP functions on "locators"
 - Stack Protocol element performs mapping
- FQDN as the identity token
 - Is this creating a circular dependency?
 - Does this impose unreasonable demands on the properties of the DNS?
- Structured token
 - What would be the unique attribute of a novel token space that distinguishes it from the above?
- Unstructured token
 - Allows for self-allocation of identity tokens (opportunistic tokens)
 - How to map from identity tokens to locators using a lookup service?

Issues

- Identity / Locator Binding domain
 - Session or host?
 - Dynamic or static?
 - Configured or negotiated?
- Scope of identity role
 - Locator independent identity
 - Equivalence binding for multiple locators
- Locator Selection
- Application visibility of identity capability
- Scoped identities
- Identity Referrals and hand-overs
- Third party locator rewriting
- Security of the binding

Open questions

- Are structured identity spaces a heavy weight solution to a light weight problem?
- How serious a routing problem is multi-homing anyway?
- Can routing scope be a better solution than complete protocol-reengineering
- What's a practical compromise vs an engineered solution to an ill-defined problem space?
- Is per-session opportunistic identity a suitably lightweight solution?

Thank you!

- Questions

