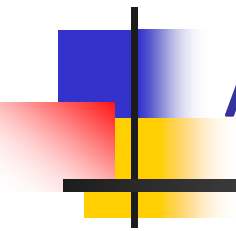

Approaches to Multi-Homing for IPv6

An Architectural View of IPv6 MultiHoming proposals

Geoff Huston

2004

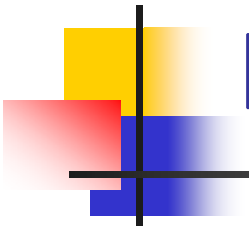

Resiliency in IP

- 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

How to resolve the Network Availability

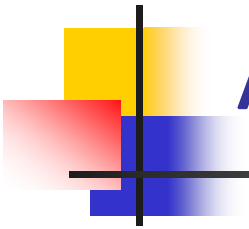

- Multiple connections to a single provider?
  - No – there's a single routing state that is vulnerable to failure
- Multiple Connections to multiple providers
  - More attractive, potentially allowing for failover from one provider to another in the event of various forms of network failure

How this is achieved in IPv4

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

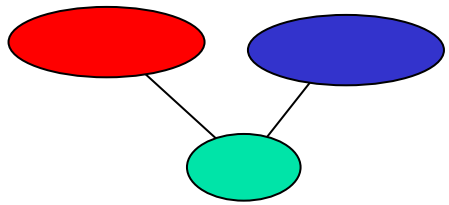

- Or:
  - Use PA space fragment from one provider
  - Advertise the fragment it to all other upstream providers
  - Follow routing

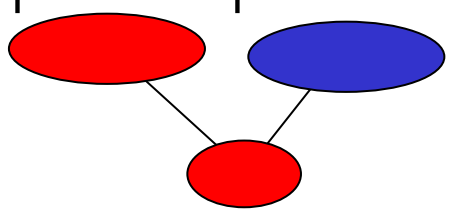

And the cost is:

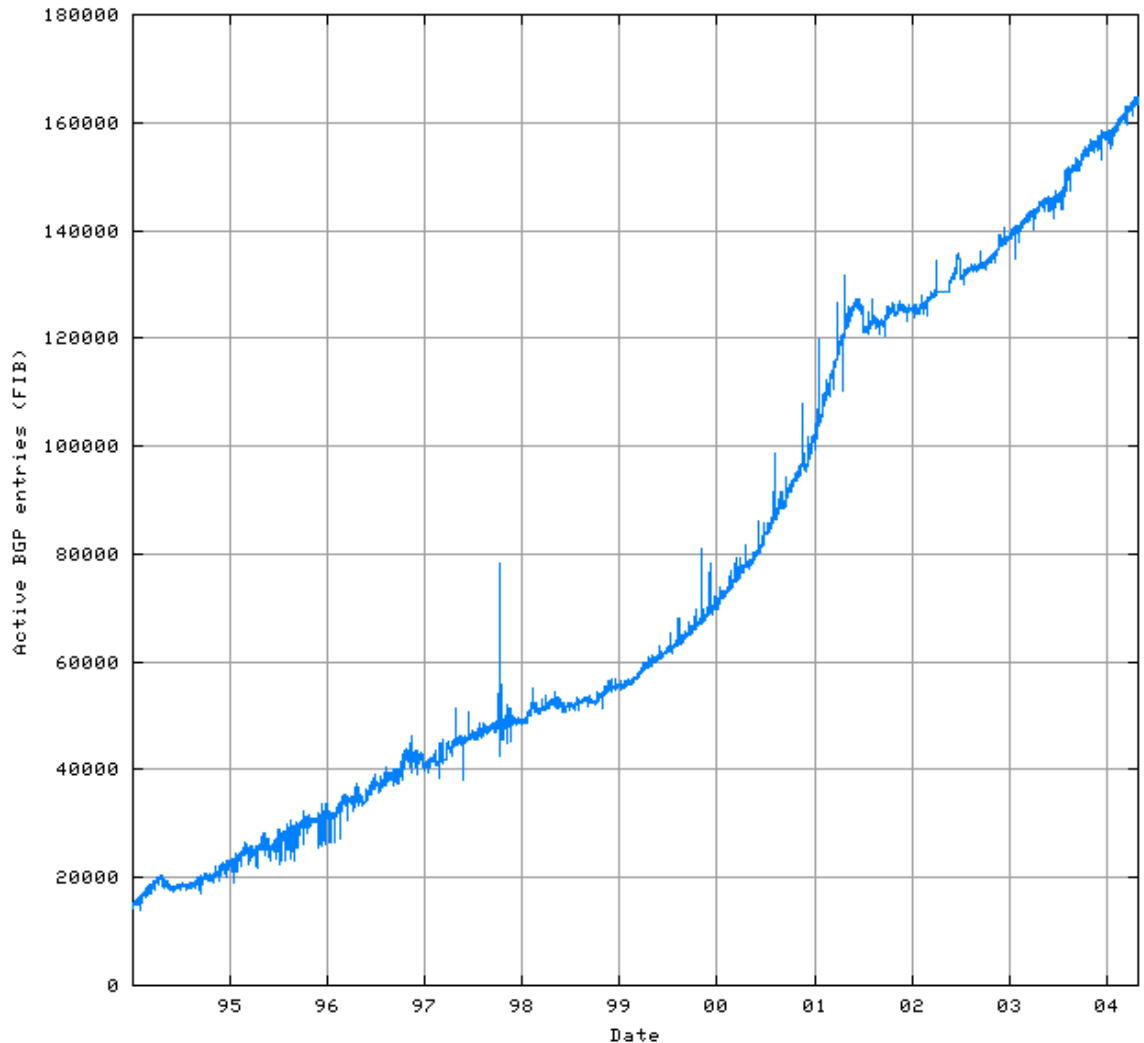

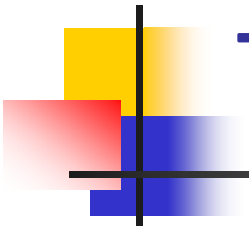

The Cost of IP Routing

- There are potentially millions of sites that would see a benefit in multi-homing
- The routing table cannot meet this demand
- Is there an alternative approach that can support multi-homing without imposing a massive load on the routing system?

What we would like…

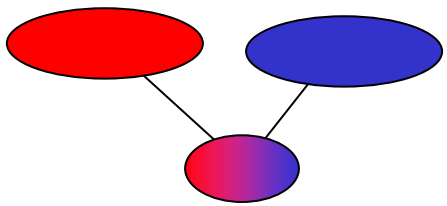

- The multi-homed site uses 2 address blocks
  - One from each provider
- No additional routing table entry required

The Problem Space

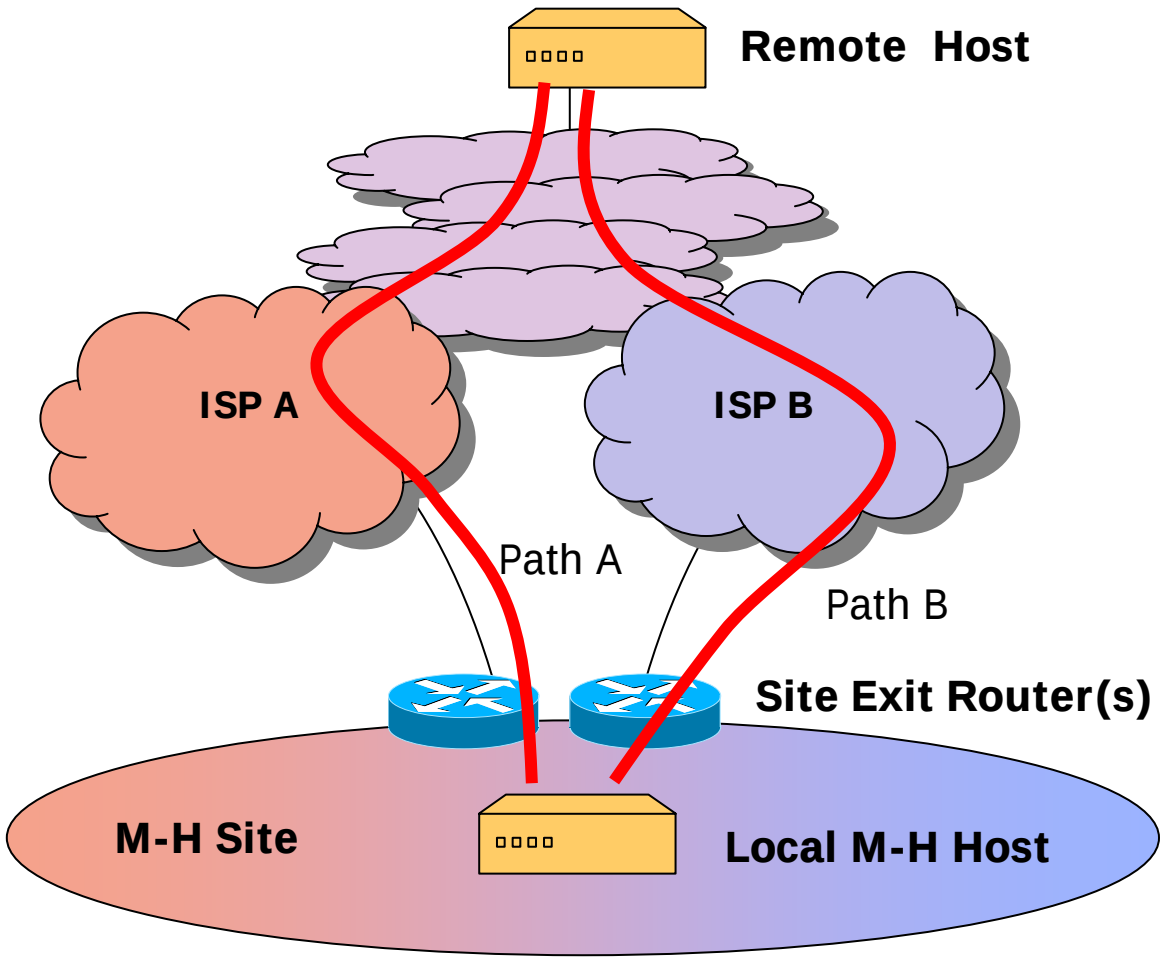

Functional Goals

- RFC3582 enumerates the goals as:
  - Redundancy
  - Load Sharing
  - Traffic Engineering
  - Policy
  - Simplicity
  - Transport-Layer Surviveability
  - DNS compatibility
  - Filtering Capability
  - Scaleability
  - 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

i.e. Do everything, simply, efficiently and cheaply with no other impact !

But this is not IP as we knew it

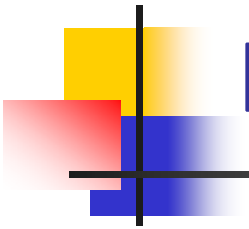

- The IP protocol architecture has made a number of simplifying assumptions
- One major assumption was that IP hosts didn't move!
  - Your IP address is the same as your identity (who)
  - Your IP address is the same as your location (where)
  - Your IP address is used to forward packets to you (how)

- 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"

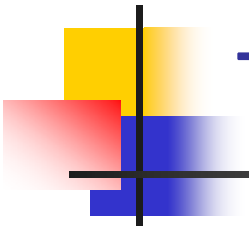

The Multi-Homing Plan

- 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”

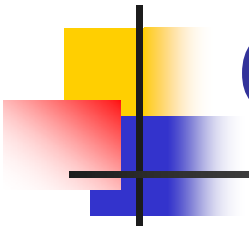

Generic Approaches:

- **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

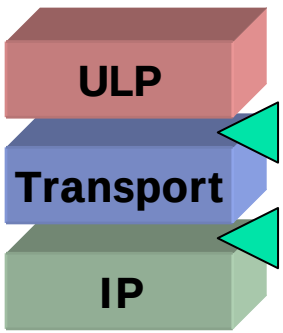

- Define a new Protocol element that:
  - presents an identity-based token to the upper layer protocol
  - Allows multiple IP address locators to be associated with the identity
  - Allows sessions to be defined by an identity peering, and allows the lower levels to be agile across a set of locators

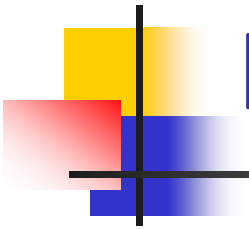

Benefits:

- 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

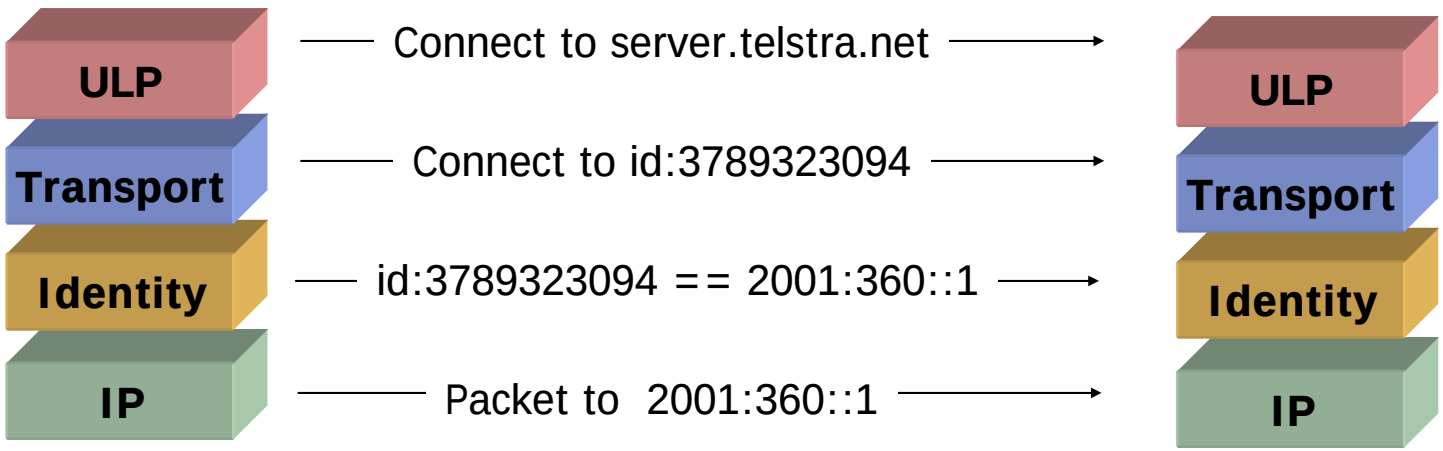

Protocol Element Implementation

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

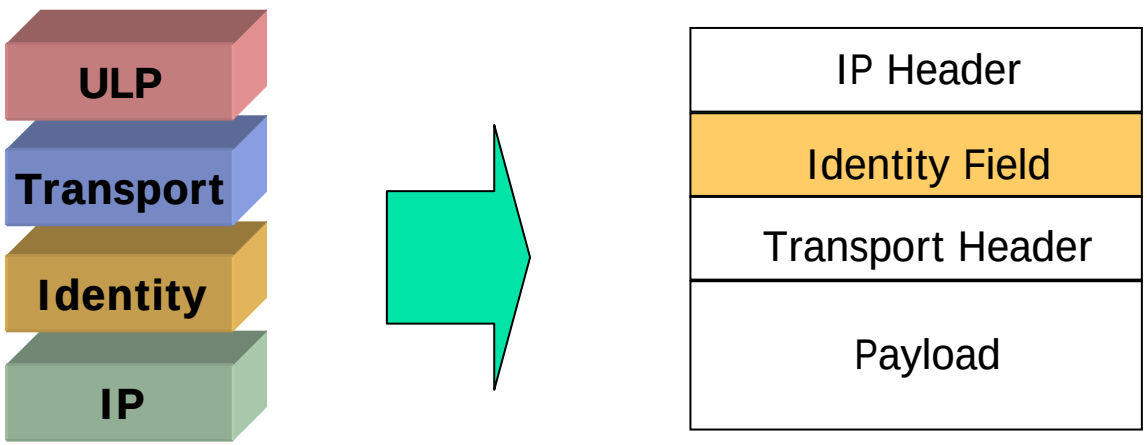

Protocol Element Implementation

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

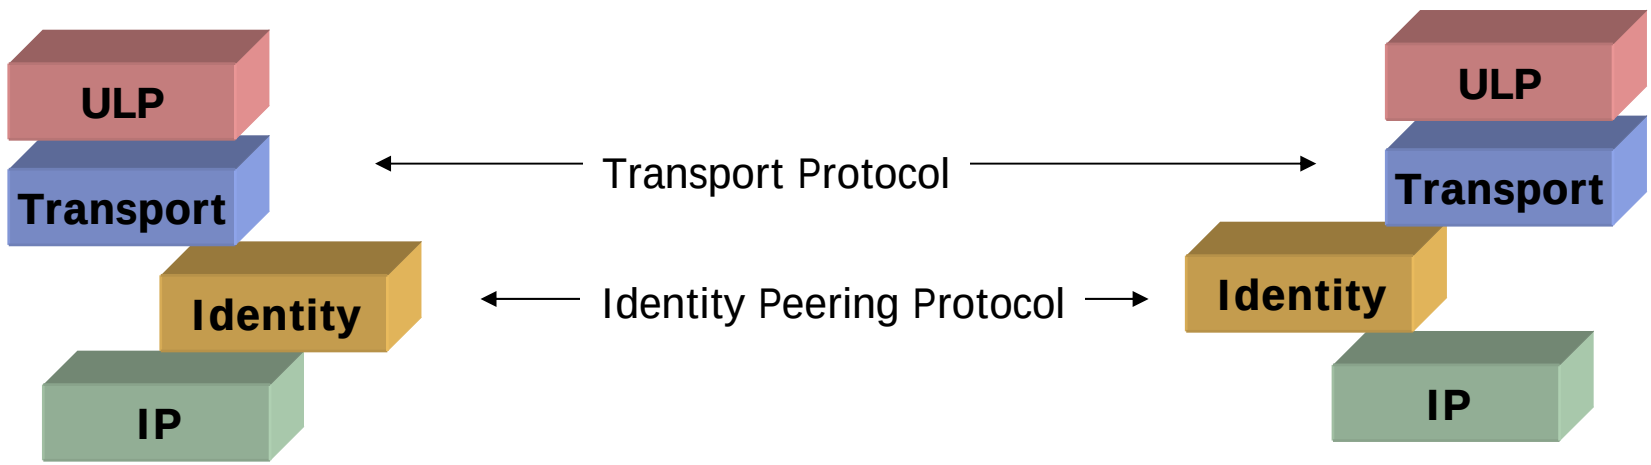

Protocol Element Implementation

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

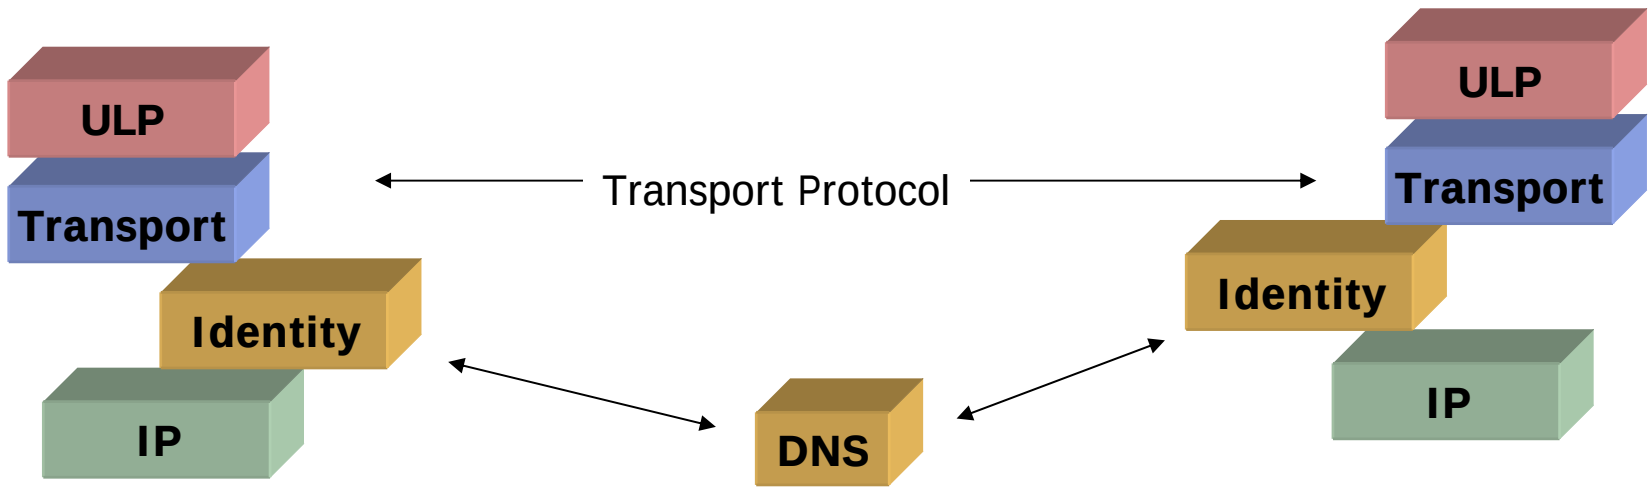

Modified Protocol Element Behaviour

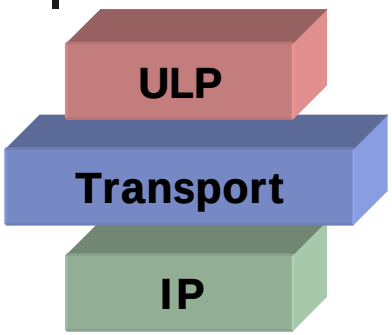

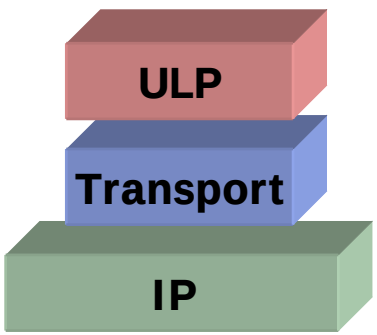

- Alter the Transport Protocol to allow a number of locators to be associated with a session
  - *e.g.* SCTP
- Alter the IP protocol to support IP-in-IP structures that distinguish between current-locator-address and persistent-locator-address
  - i.e. MIP6

Modified Host / Router Interaction

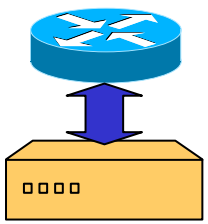

- Modify the interaction between the host and the Site Exit router to allow:
  - Source-based routing for support of host-based site-exit router selection
  - Site Exit router packet header modification
  - Host / Site Exit Router exchange of reachability information

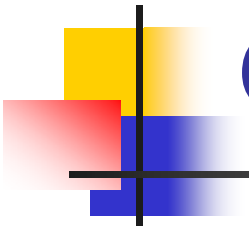

Common Issues

- Host based locator address selection
  - How to pick the “best” source locator for the reverse packet?
  - How to pick the “best” destination locator if there are more than one available?
- Detection of network element failure
  - How to detect reverse path failure?
- Session Persistence
  - How and when to switch locators for active sessions ?

Proposals for a new Protocol Element

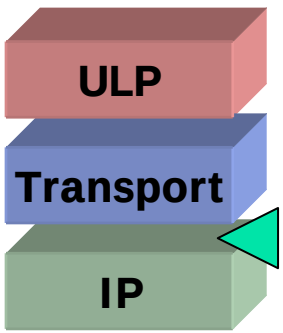

- HIP:
  - Shim between Transport and IP layer
  - Presents a stable identity to the transport layer (cryptographic hash of local identity key)
  - Allows multiple locators to be bound to the identity, and communicates this binding to the remote end (HIP protocol)
  - Allows the local host to switch source locators in the event of network failure to ensure session surviveability. The crytographic function is used to determine if the new locator is part of an already established session. (“same key, same session”)

Proposals for a new Protocol Element

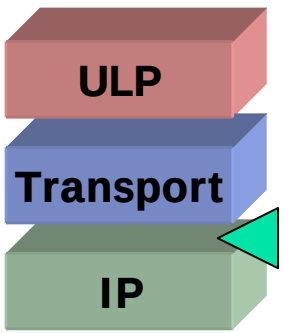

- NOID +
- SIM (CBID 128) +
- CB64:
  - Addition of an identifier shim layer to the protocol stack.
  - The identifier / locator mapping may be contained in the DNS (NOID) or may be contained within a protocol exchange (SIM), or a hybrid approach (CB64)
  - Permits Site Exit routers to rewrite source locators on egress
    - (i.e. includes elements of host / Site Exit Router interaction)

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)

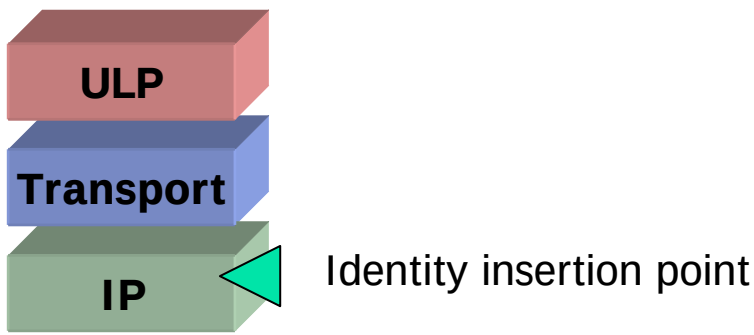

Proposals for an Identity Protocol Element

Hierarchically Structured Space

- 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

- Unstructured token
  - Allows for self-allocation of identity tokens (opportunistic tokens)
  - How to map from identity tokens to locators using a lookup service?

Proposal for a Modified Transport Protocol

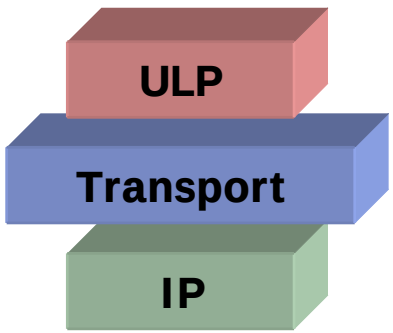

- SCTP:
  - Host-based solution that sets up multiple locators for a session
  - Changes locators on end-to-end heartbeat failure
  - Depends on IPSEC for operational integrity of locator exchange

Proposal for a Modified IP Layer

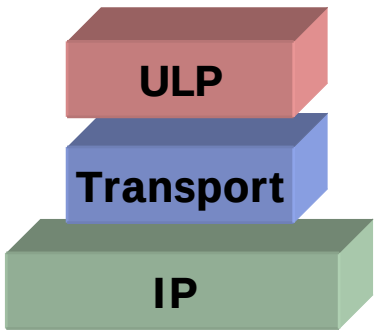

- MIP6:
  - Use one locator as the home address
  - Allow a dynamic switch to an alternate locator as a session surviveability response
  - An instance of a generic approach of packet encapsulation, where the outer encap is the current locator binding and the inner packet is the identifier peering.

Modified Host / Site Exit Router interaction

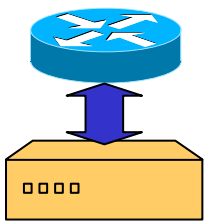

- Site Exit Anycast proposal
  - Allows local forwarding of outgoing packets to the ‘matching’ site exit router for the selected source address
- Local Site source locator-based forwarding
- Site Exit source address rewriting
  - May be used in combination with locator protocol element proposals
- Have upstream accept all of the site’s sources and use host-based source locator selection

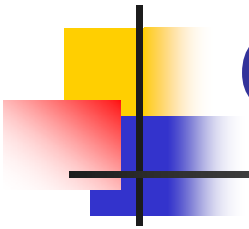

Common Issues

- Picking the ‘best’ source locator

  (how do know what destination works at the remote end?)

  - Use each locator in turn until a response is received
  - Use a identity peering protocol to allow the remote end to make its own selection from a locator set

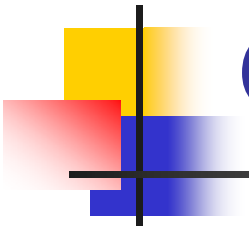

Common Issues

- Picking the ‘best’ destination locator
  - Longest match
  - Use each in turn
- Picking the ‘best” source / destination locator pair
  - As these may be related choices

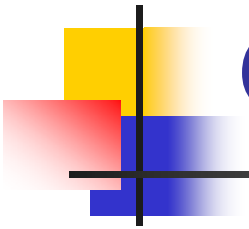

Common Issues

- Detecting network failure

  (How does a host know that its time to use a different source and/or destination locator?)

  - Heartbeat within the session
  - Modified transport protocol to trigger locator change
  - Host / Router interaction to trigger locator change
  - Application timeframe vs network timeframe
  - Failure during session startup and failure following session establishment

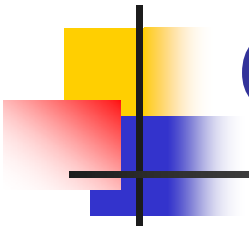

Common Issues

- Network layer protocol element
  - How do you know a session is completed?
    - The concept of session establishment and teardown is a transport concept, not an IP level concept
  - What do you need to do to bootstrap?
    - Are there ‘distinguished’ locators that you always need to use to get a session up?

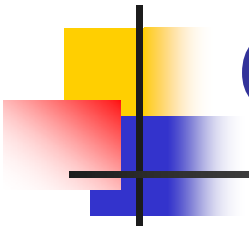

Common Issues

- Session Persistence
  - Use one locator as the “home” locator and encapsulate the packet with alternative locators
  - Set up the session with a set of locators and have transport protocol maintain the session across the locator set
    - Optionally delay the locator binding, or allow the peer dynamic change of the locator pool
  - Use a new peering based on an identity protocol element and allow locators to be associated with the session identity

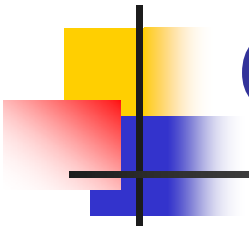

Common Issues

- Identity / Locator Binding domain
  - Is the binding maintained per session?
    - In which case multiple sessions with the same endpoints need to maintain parallel bindings
  - Is the binding shared across sessions?
    - In which case how do you know when to discard a binding set?

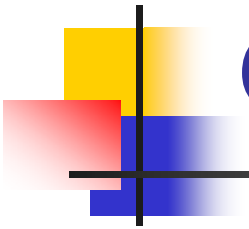

Common Issues

- Bilateral peer applications vs multi-party applications
  - What changes for 3 or more parties to a protocol exchange?
- Application hand-over and referral
  - How does the remote party identify the multi-homed party for third party referrals?

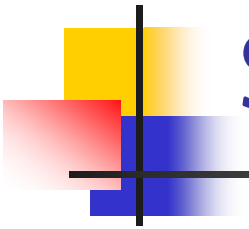

Security Considerations

- Major agenda of study required!
- Not considered in the scope of this work
- Worthy of a separate effort to identify security threats and how to mitigate these threat

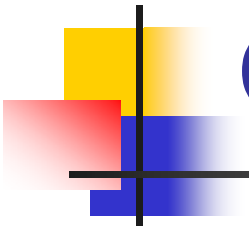

Questions

- 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
- Is per-session oppostunistic identity a suitably lightweight solution?
- Whats a practical compromise vs an engineered solution to an ill-defined problem space?

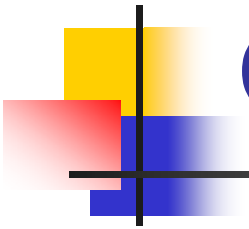

Questions?

- Your turn!

Thank You