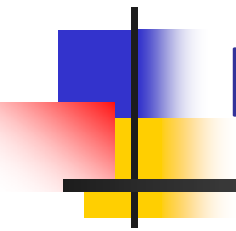

Architectural Approaches to Multi-Homing for IPv6

A Walk-Through of

draft-huston-multi6-architectures-00

Geoff Huston

June 2004

Recap – Multi-Homing in IPv4

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

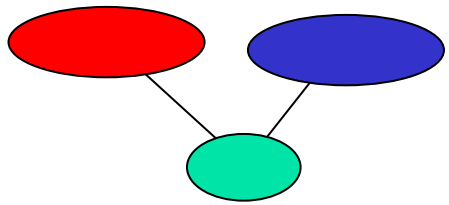

- Or:
  - Use PA space fragment from one provider
  - Advertise the fragment to all other upstream providers
  - Follow routing

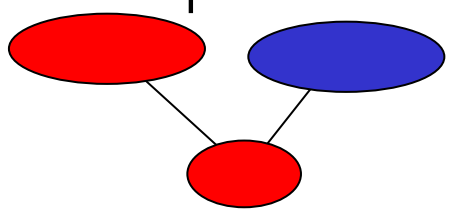

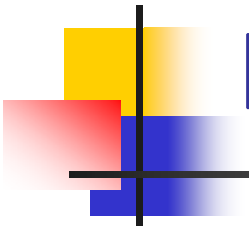

But…

- There are potentially millions of sites that would see a benefit in multi-homing
- It is assumed that routing table cannot meet this demand, in addition to other imposed loads on routing scaleability

- Is there an alternative approach that can support multi-homing without imposing a massive load on the routing system?

The objective…

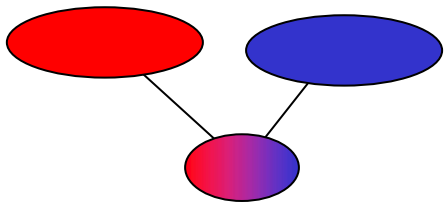

- The multi-homed site uses 2 address blocks
  - One from each provider
- No additional routing table entry required
- Data traffic uses either path depending on path availability and policy constraints

Generic Problem Space

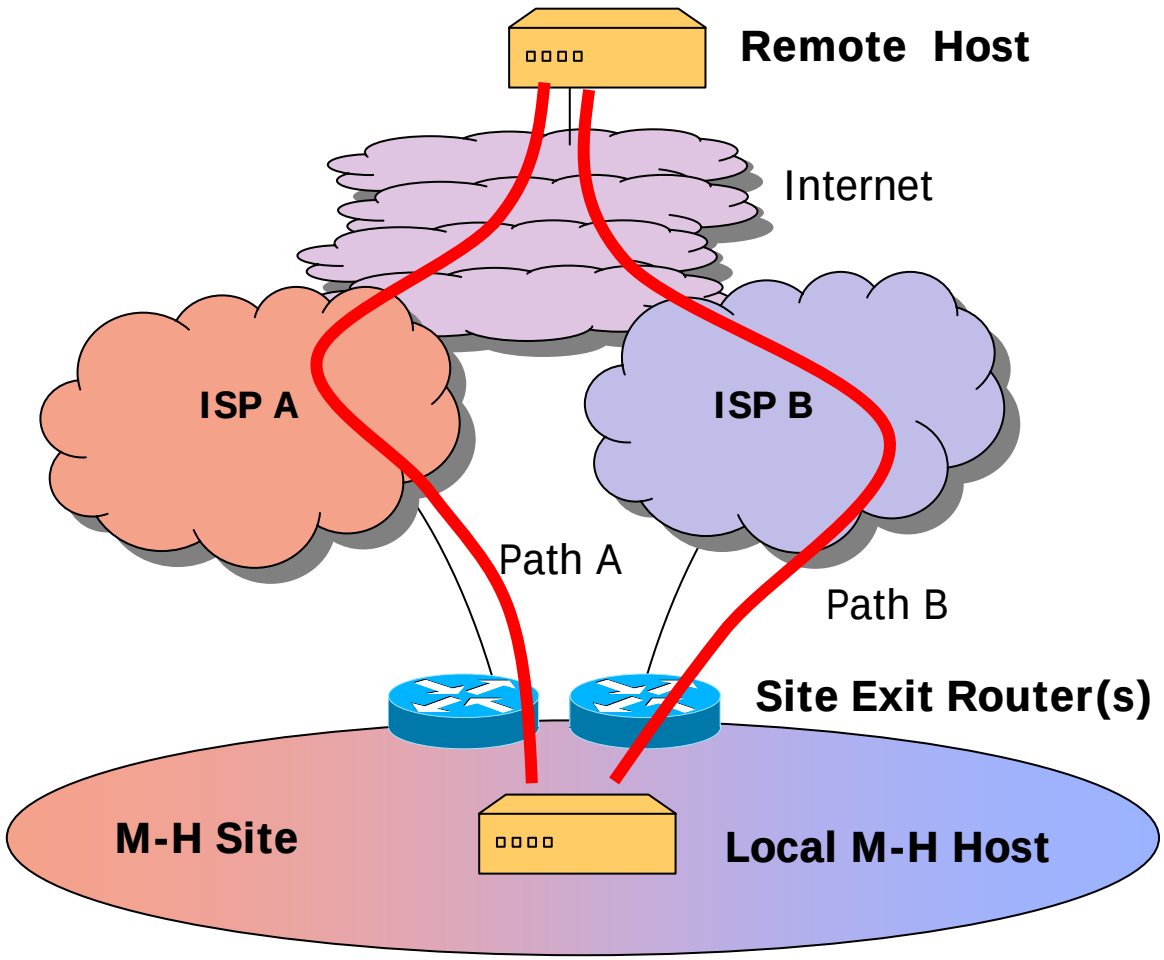

Functional Goals

- 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

Generic Approaches:

- Route each M-H site
  - IPv4 approach
- Introduce “Identity” into the protocol exchange
  - Insert a new element in the protocol stack
    - New synchronization element to exchange id/locator binding
  - Modify the Transport or IP layer of the protocol stack
    - Perform id/locator mapping within an existing protocol element
  - Modify the behaviour of the host/site exit router interaction
    - Modified forwarding architecture coupled with distributed state of identity / locator binding

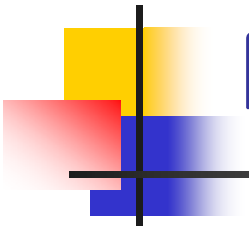

M-H via Routing

- Ultimately this recasts the definition ‘routing element’ to the level of a single site
- This has the potential to remove any structural hierarchy from the inter-domain system
- This would place significant scaling strains on the inter-domain routing system
  - There are significant doubts that a non-hierarchically structure routing space can scale in a viable and stable fashion

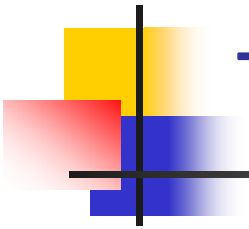

The M-H Identity Approach

- 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"

New Protocol Element

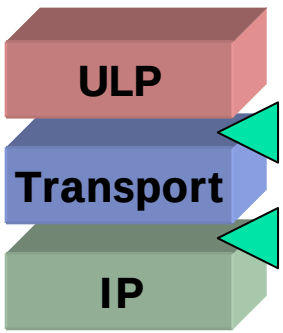

- Define a new Protocol element that:
  - presents an identity-based token to the upper layer protocol
  - Allows multiple IP address locators to be associated with the identity
  - Allows sessions to be defined by an identity peering, and allows the lower levels to be agile across a set of locators

Modified Protocol Element Behaviour

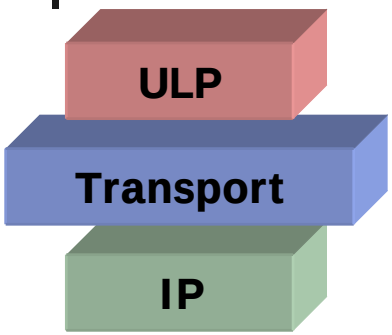

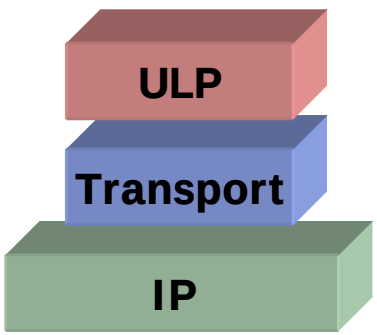

- Alter the Transport Protocol to allow a number of locators to be associated with a session
  - *e.g.* SCTP
- Alter the IP protocol to support IP-in-IP structures that distinguish between current-locator-address and persistent-locator-address
  - i.e. MIP6

Modified Host / Router Interaction

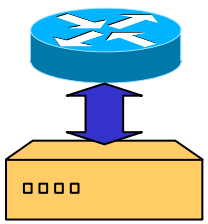

- Modify the interaction between the host and the Site Exit router to allow:
  - Source-based routing for support of host-based site-exit router selection
  - Site Exit router packet header modification
  - Host / Site Exit Router exchange of reachability information

Modified Host / Site Exit Router interaction

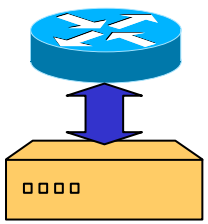

- Site Exit Anycast proposal
  - Allows local forwarding of outgoing packets to the ‘matching’ site exit router for the selected source address
- Local Site source locator-based forwarding
- Site Exit source address rewriting
  - May be used in combination with locator protocol element proposals
- Have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

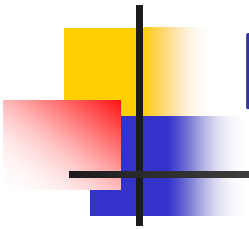

Identity / Locator Binding

- Allow a single transport session to be associated with multiple paths that transit the network
- One approach is to:
  - use the transport protocol to establish the session based on an “identity” token
  - Map this identity value to a valid locator
  - Use this locator in the packet on the wire as source / destination address

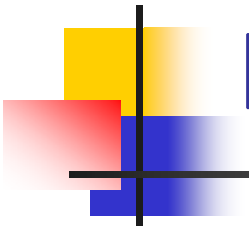

Benefits of Id/Loc Split

- 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 Location

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

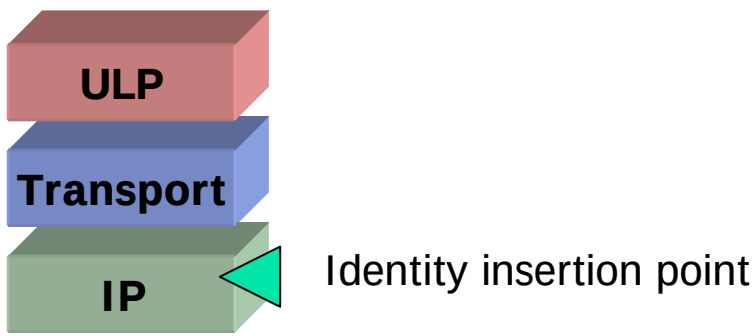

Identity Protocol Element

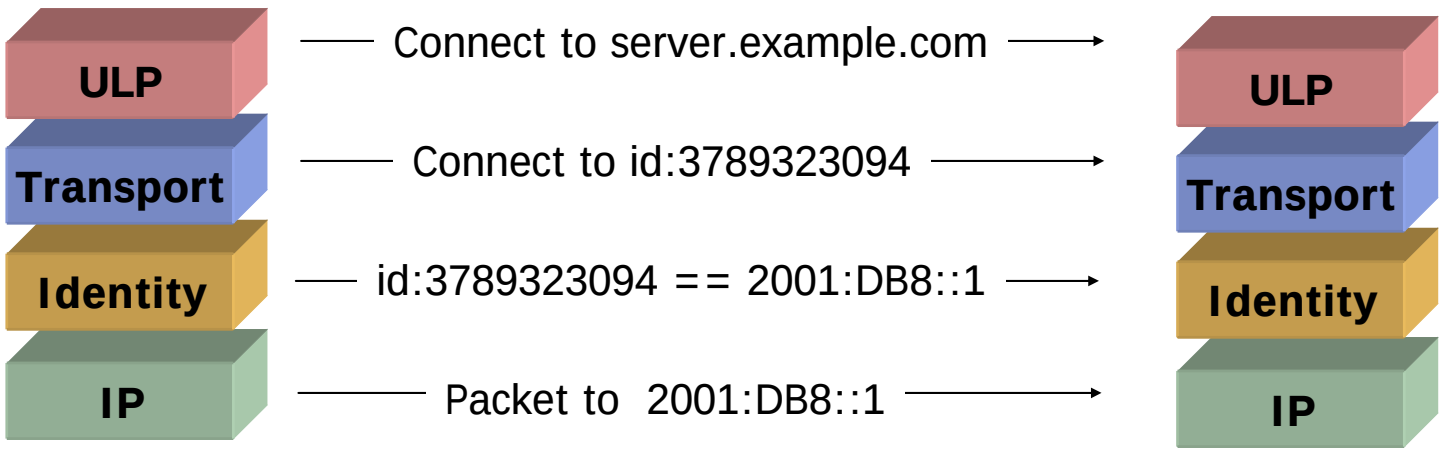

Protocol Element Implementation

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

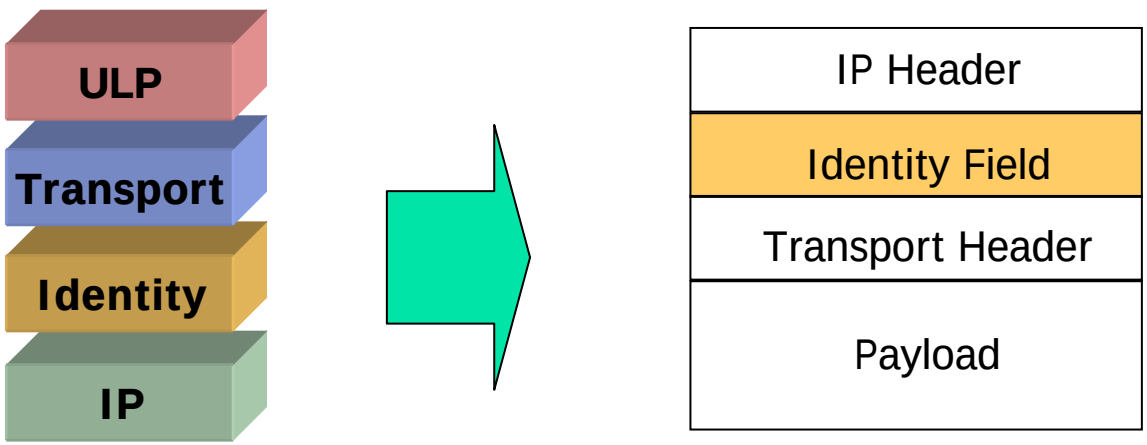

Protocol Element Implementation

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

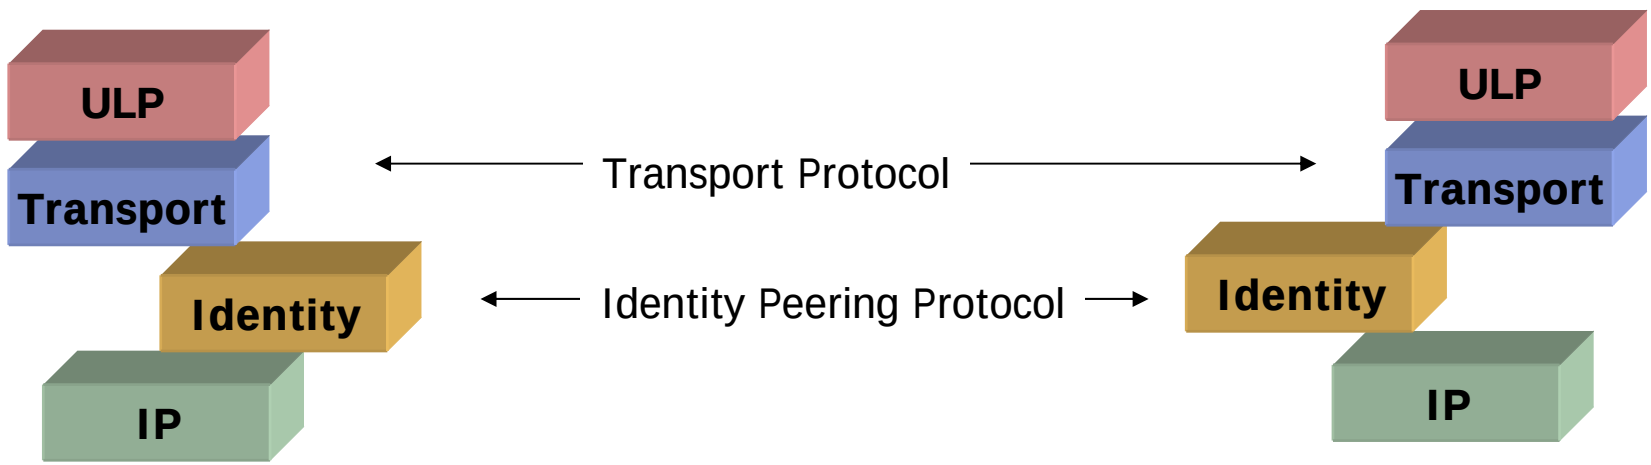

Protocol Element Implementation

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

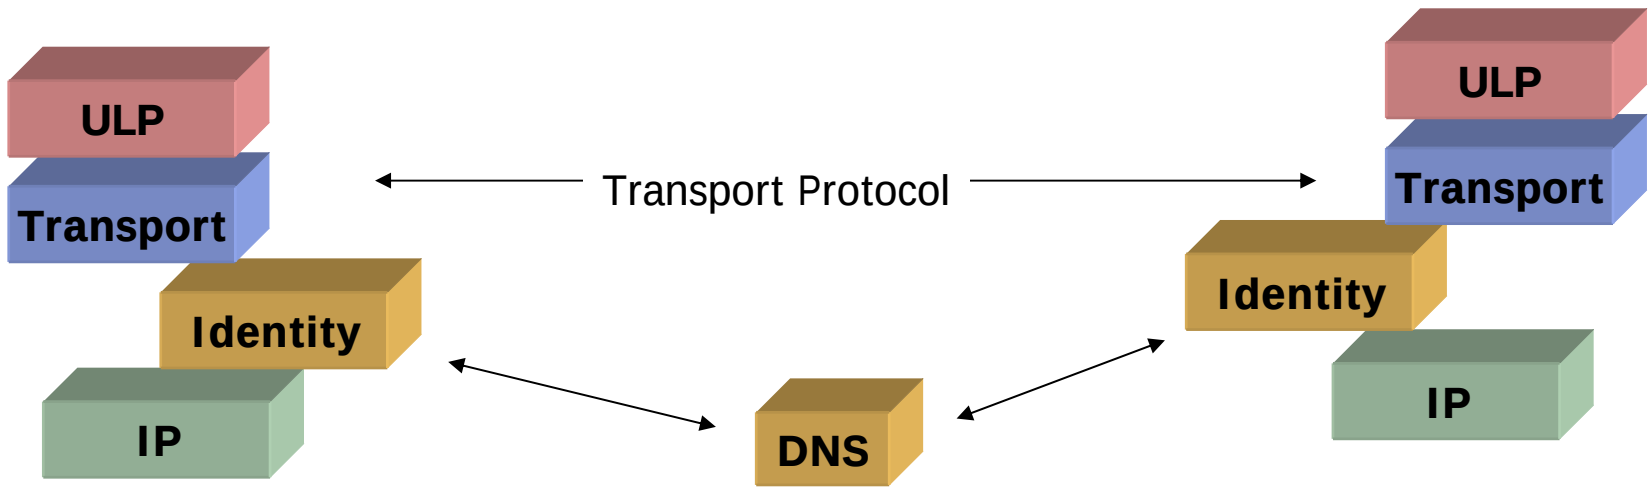

Proposals for an Identity Protocol Element

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?

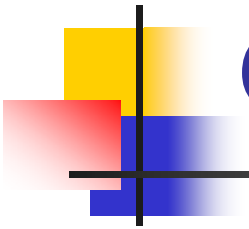

Common Issues

- 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

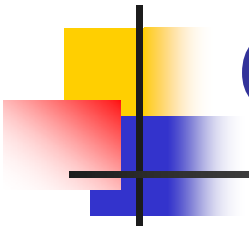

Common Issues

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

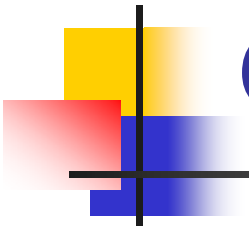

Common Issues

- 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

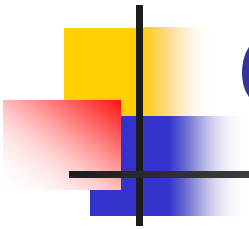

Common Issues

- 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?

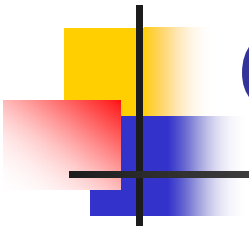

Common Issues

- 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

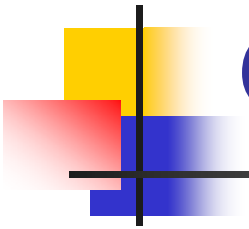

Common Issues

- 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?

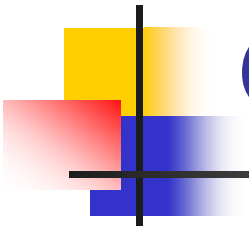

Common Issues

- 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?

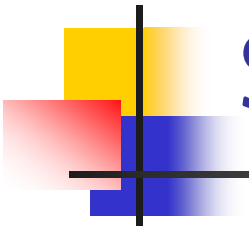

Security Considerations

- 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

Proposed next steps for the draft

1. Complete the proposal survey (attachment)
2. Analyse Identity properties in further detail
3. Examine some further open issues (next slides)
4. Make some tentative conclusions regarding the properties of a robust M-H approach
5. Submit to WG for adoption as a WG document

- Following slides have some details on steps 3 - 6

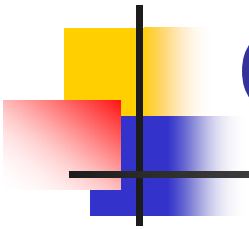

Open Questions

- Routing Questions
  - How serious a routing problem is multi-homing anyway?
  - Can routing scope be a better solution than complete protocol-reengineering?
  - Are there other approaches to managing the inflation rate of the Internet routing system?

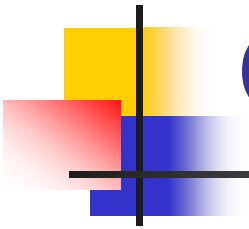

Open Questions

- Id/Loc questions
  - Is the specification of a structured identity space coupled with changes to the IPV6 protocol stack a case of solution overkill?
  - What additional infrastructure service overheads are required to distribute a structured identity space?
  - Is there an existing identity space that could be used for this purpose?
  - Is the identity point the device or the protocol stack?
  - Is per-session opportunistic identity a suitably lightweight solution?
  - Is this just multi-homing or a more generic id/locator discussion?

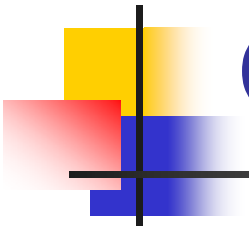

Open Questions

- Applications and Identities
  - Is a self reference within an application the identity value?
  - If so, then can opportunistic id values be used in this context?

Properties of an ID-based M-H Solution

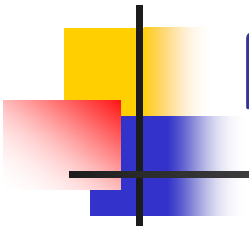

- ID/Locator split and associated stack modification appears to be a robust form of identity implementation
- Properties of a structured identity space
  - Creating yet another managed token space for a set of structured stack identities may be overkill
- Properties of opportunistic keys
  - The lack of persistence may make initial key association vulnerable to attack
  - Lack of support for referral function
  - Continuation of overloaded semantics of IPv6 addresses