The Architecture of the Internet

or

Waist Watching in IP

Geoff Huston

Executive Director,

Internet Architecture Board

Does the Internet Protocol even have an "Architecture"?

- One view is that there is no clear architecture
  - The Internet today is a product of a process of incremental short term feature creep rather than deliberate design
  - There is no process of imposition of architectural standards onto deployed networks
  - Each Internet provider is at liberty to deploy an architecture of choice (or use no coherent architecture at all!)

The “Hourglass” view of the IP protocol architecture

- Another view is that IP does have a consistent protocol architecture:

  a universal adaptation layer

  - IP sits above a large number of network media
    - SDN, SDH, Ethernet, DSL, Wireless, even carrier pigeon
  - IP provides a consistent addressing and transport service for a variety of application requirements
    - Reliable data transfer
    - Semi-Real time streams
    - High volume streams
    - Reliable Transactions
    - Multi-level Referrals

The Hourglass IP Model

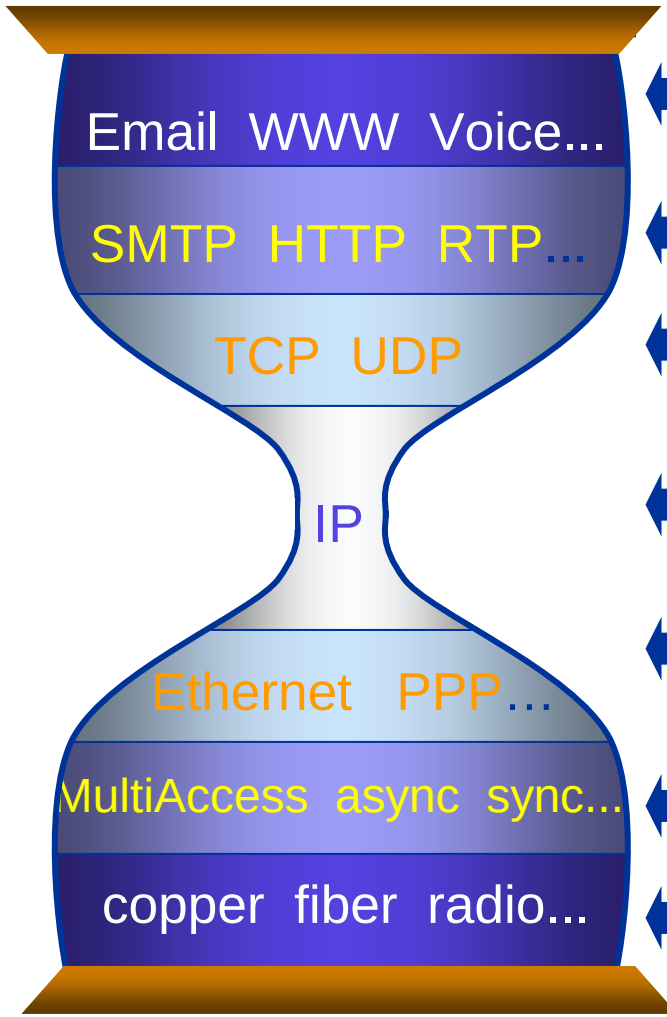

- **User Application**
- **End-to-end Application Protocol**
- **Transport Protocol**
- **Internet Layer**
- **Media Access Protocol**
- **Media Format**
- **Physical System**

Why use an IP adaptation layer?

- Why an IP layer?
  - isolate end-to-end protocols from transmission network details and changes
  - Add an overlay of consistent global addressing
  - make a bigger virtual network
- Why a *single* Internet protocol?
  - maximize interoperability
  - minimize number of service interfaces
- Why a *narrow* Internet protocol?
  - assumes least common network functionality to maximize number of usable networks
    - IP provides only unreliable, asynchronous datagram delivery

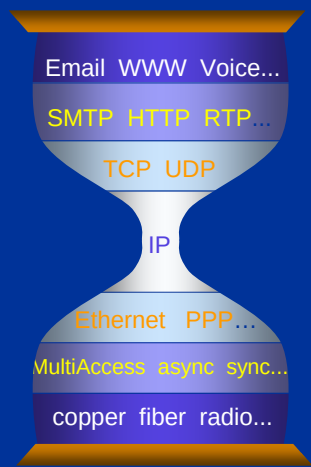

Why use an IP adaptation Layer?

- Simple to adapt to new media
  - IP Address to MAC address resolution protocol
  - IP packet framing definition
  - And its done!
- Simple to create composite networks
  - Ethernet - ATM – SDH – Ethernet – wireless
- Simple to scale
  - IP networks are composite networks
  - No single coordinated effort required
  - Minimal interdependencies between component networks
  - Very simple network-to-network interface
- Simple to create applications in IP
  - Applications do not need to understand or adapt to varying transport characteristics

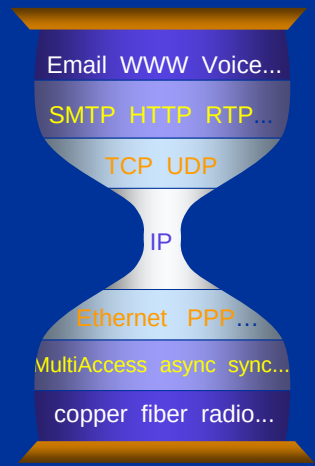

So Why Am I Talking About Watching the Waist?

- It happens on reaching middle age (me & IP)
- The IP layer is the only layer small enough for me to get my arms around
- I am worried about how the architecture is being damaged: the waste of the hourglass
- The hourglass theme offers some bad puns!

Putting on Weight!

Additional functionality within the IP layer requires greater levels of application complexity

Additional functionality within the IP layer requires more functionality and greater levels of coupling from underlying transmission networks

Mid-Life Identity Crisis

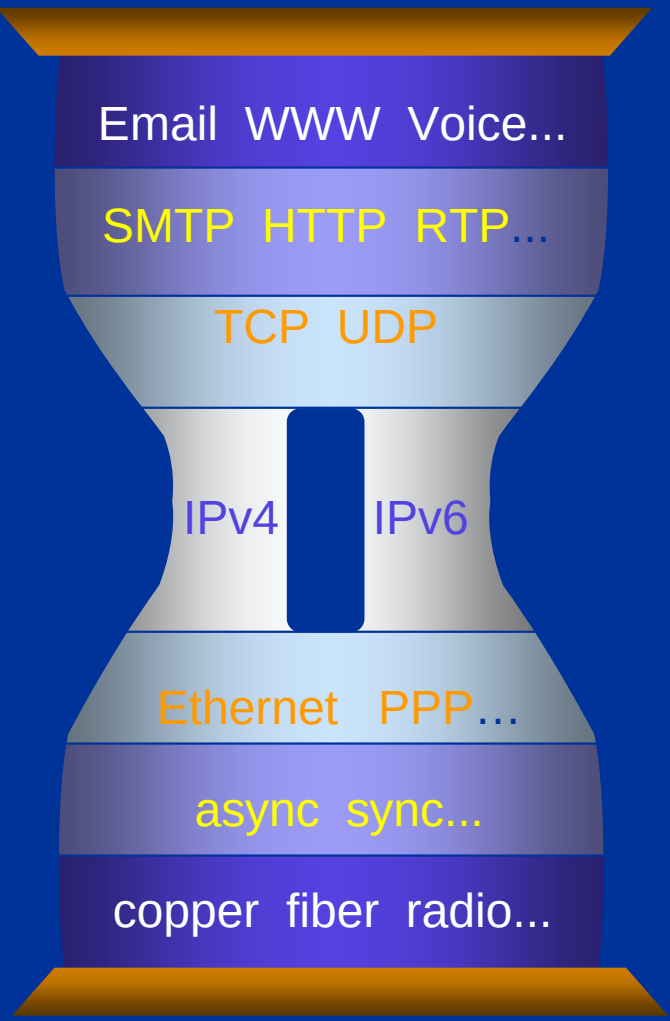

The introduction of a V6 transition into IP

- **Doubles the number of service interfaces**
- **Requires changes above and below the IP layer**
- **Creates subtle (and not so subtle) interoperability problems**
- **Does not appear to add new functionality or adequately address evolving requirements for IP**

Oops!

You can't take the falls any more without breaking something!

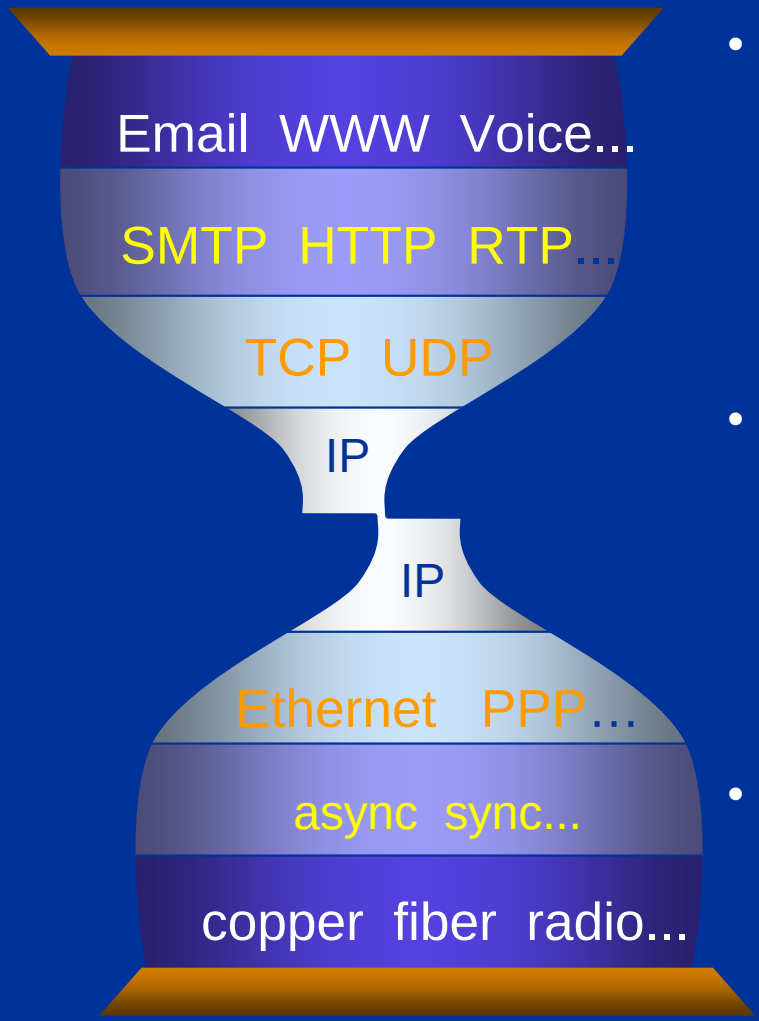

- **Network Address Translators (NATs) & Application Level Gateways (ALGs) used to glue together network domains**
- **lots of kinds of new glue being invented—ruins predictability and makes applications more complex**
- **some applications remain broken, since the NAT glue does not provide fully transparent connectivity**

Your body shape changes – with surprising results!

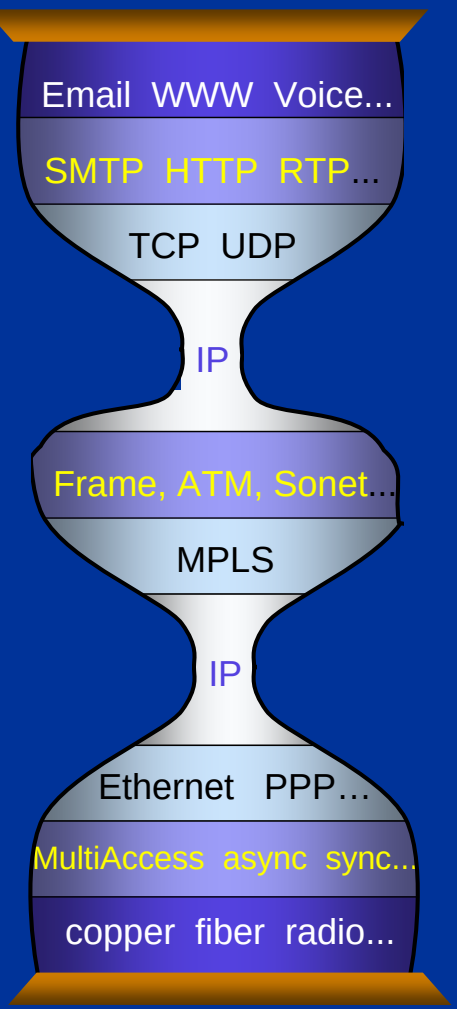

- **The addition of MPLS to the protocol model has caused some surprising outcomes in terms of using MPLS and IP as a substrate for emulated wire services**
- **It is not obvious this this form of complexity is a reliable foundation for a scaleable network architecture**

Your children now challenge your role!

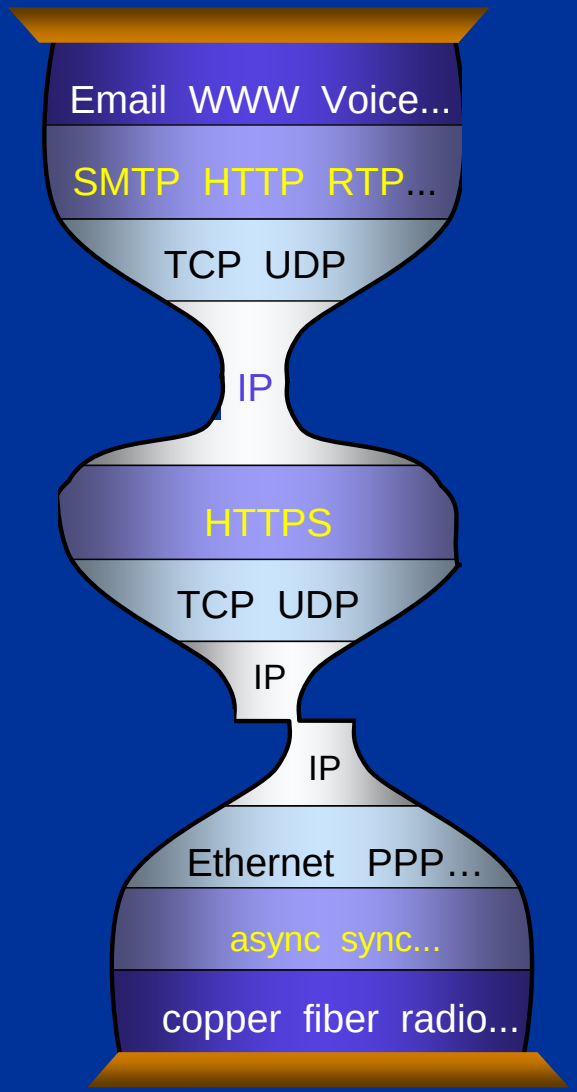

- IP over HTTPS is now a popular solution for firewall traversal
- Any level of a layered network model can be seen as functionally equivalent to any other layer – it all depends on the committee that standardized it
- The temptation to solve a problem by adding another layer of indirection is a fine example of computer science
  - it does not always create robust networking architectures!

Insecurities and Anxieties Appear

- IP networks today are plagued with hostile and annoying forms of traffic
- The End-to-End model of applications operating above the IP layer is causing a multitude of problems for end users, operators and IP itself
  - Firewalls, Application Level Gateways, Network mediation of traffic
  - Application servers are being embedded into the service provider's architectures
- Requirement for "robust" IP services

Your self-confidence is sagging …

- **IP alone is not enough any more**
  - **A crisis in confidence in "basic" IP as being a viable and sustainable platform for all forms of public and private communications services**
  - **there is a push to add "features" into the IP platform as a way of adding value to a basic IP service offering**
  - **This is leading to more complex and more expensive IP+ platforms**
    - **VPNs with QoS**
    - **Real Time support for multi-media delivery**
    - **Integration of content delivery services into the IP architecture**

And you recognize that you can't be the absolute best in everything…

- IP has some weaknesses in large scale environments that support high volume real time synchronous communications
- IP has some problems with wide area coverage radio environments
- IP has challenges in supporting provider-based VPNs with address and service quality partitioning

But IP is still supple!

- **IP-in-IP tunnelling offers a number of solutions that can support various forms of VPN architectures and provider-selection functions while still retaining much of the benefit of the thin adaptation function of IP**
- **IP-in-IP offers solutions to mobility of hosts and networks using discrete IP headers for identity and location of the mobile object**

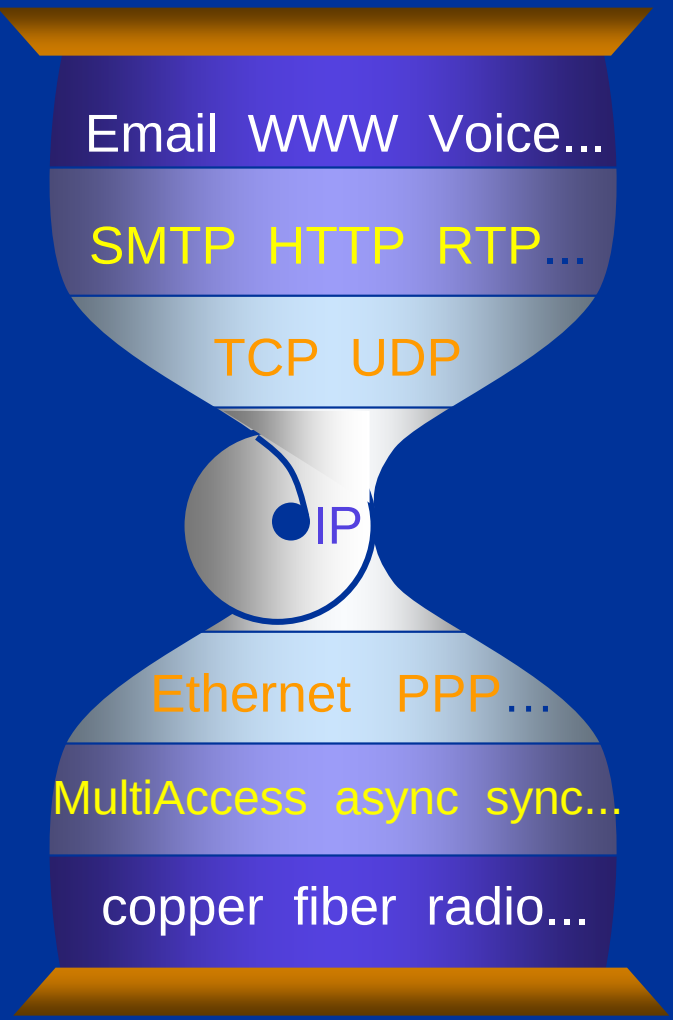

Entropy or Evolution?

- It looks like the normal ***entropy*** (decay) that besets all large, engineered systems over time
- I don't know where/how to reapply energy to fight the entropy
- Its less worrisome to view this process as ***evolution*** instead
  - the Internet as an evolving lifeform or ecosystem?
  - just let nature (the market) take its course
  - though result is undesigned and unpredictable, should not be viewed as decay. Its adaptation.

Multi-Homing – A Case in Point

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 target

- 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**

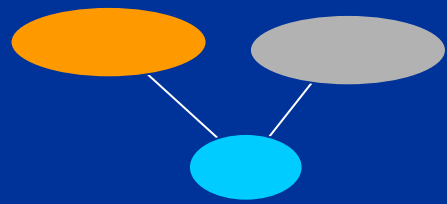

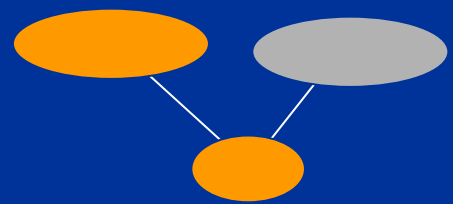

And the cost is:

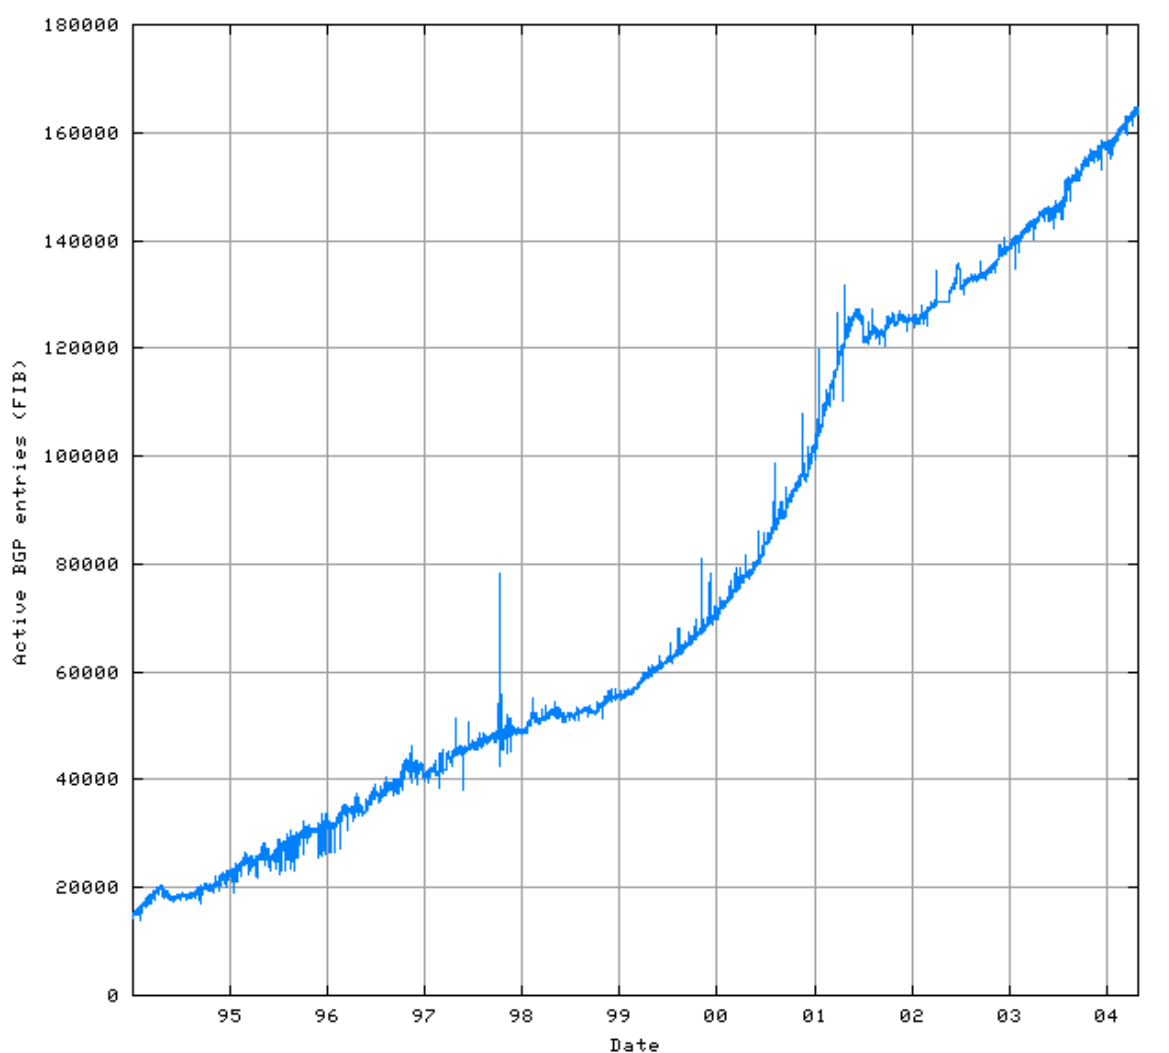

The Cost of IP Routing

- 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...

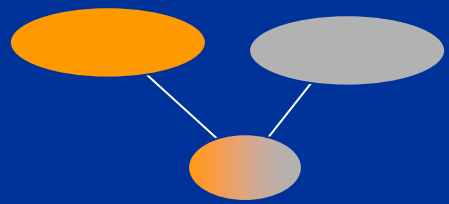

- The multi-homed site uses 2 address blocks
  - One from each provider
- No additional routing table entry required

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 to include identity / locator mapping

New Protocol Element

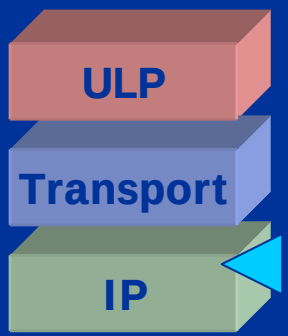

- Define a new Protocol element that:
  - presents an identity-based token to the upper layer protocol
  - Allows multiple IP address locators to be associated with the identity
  - Allows sessions to be defined by an identity peering, and allows the lower levels to be agile across a set of locators

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

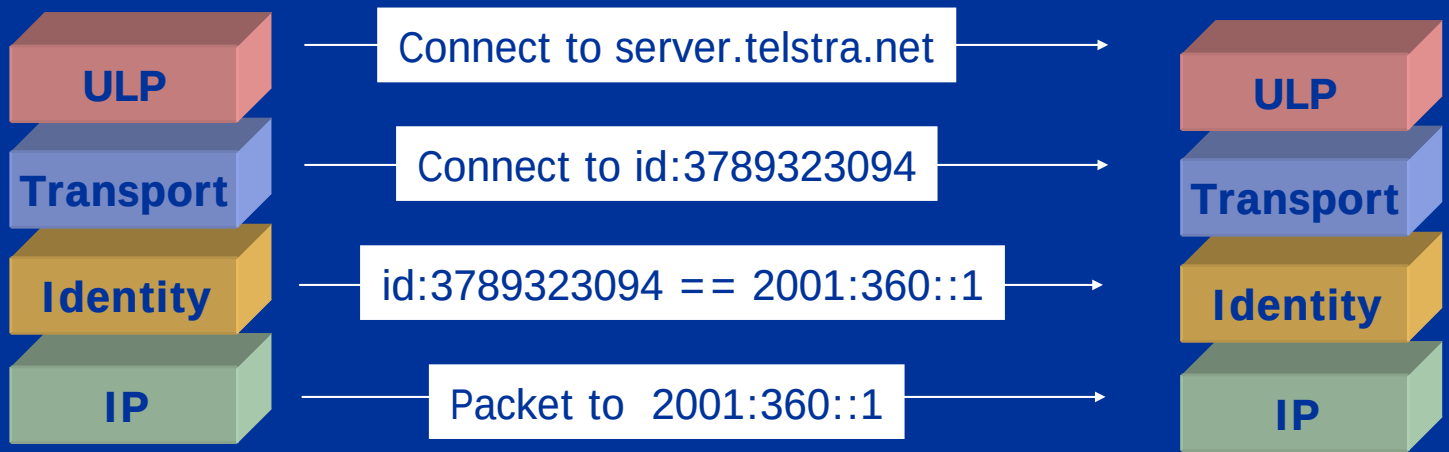

Protocol Element Implementation

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

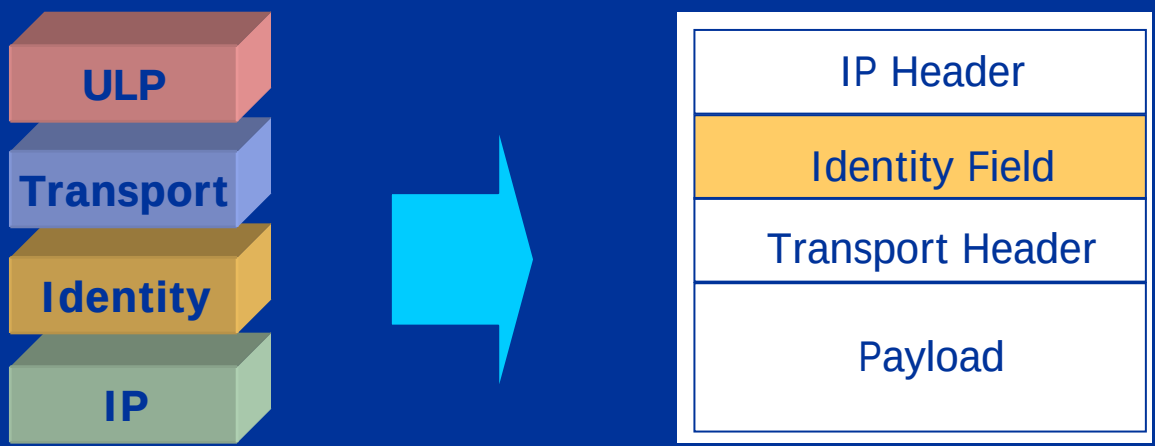

Protocol Element Implementation

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

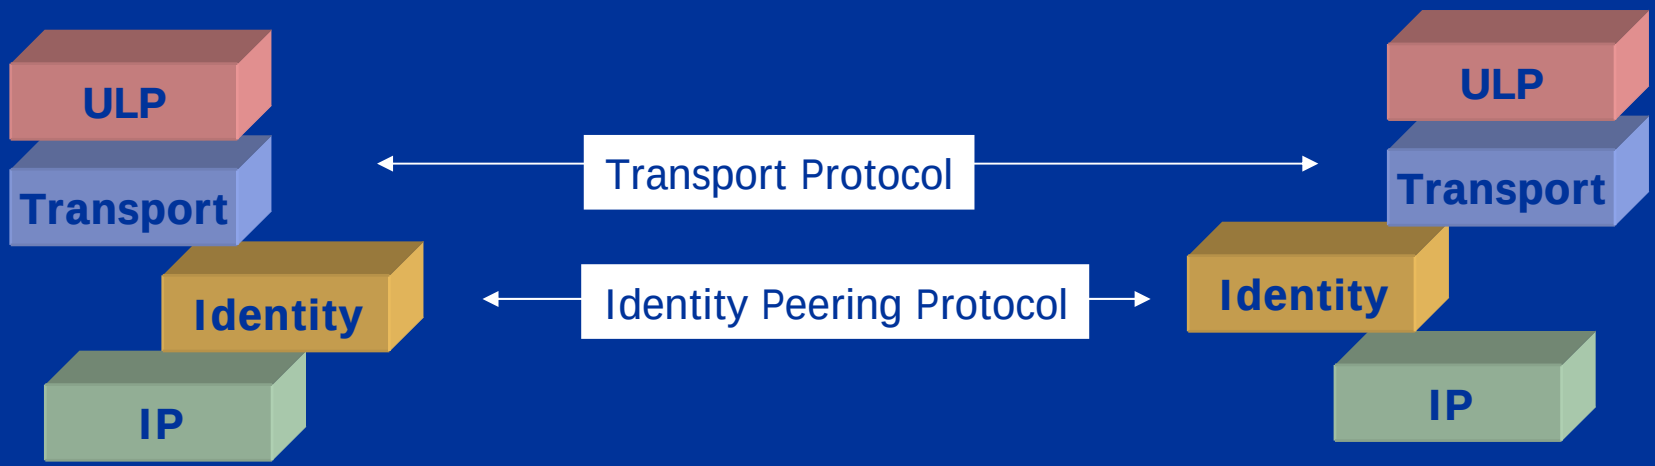

Protocol Element Implementation

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

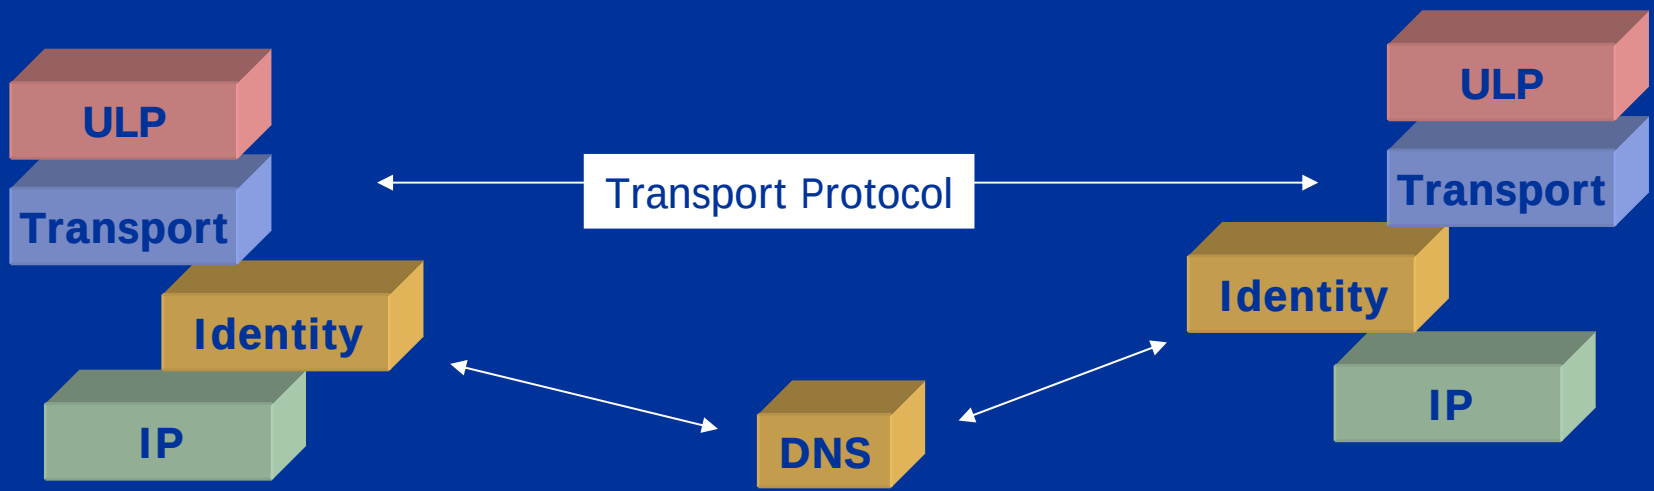

Modified Protocol Element Behaviour

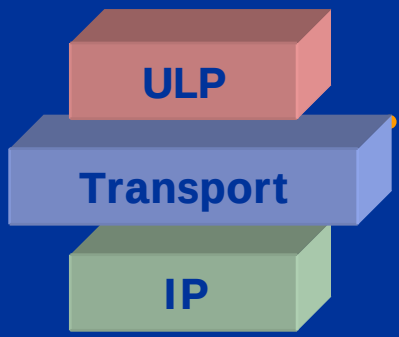

- Alter the Transport Protocol to allow a number of locators to be associated with a session
  - e.g. SCTP

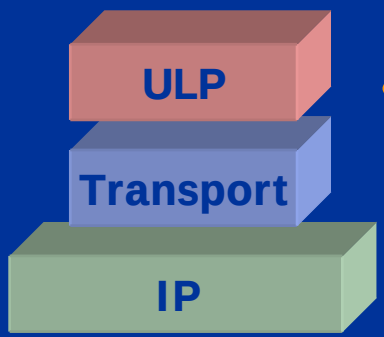

- Alter the IP protocol to support IP-in-IP structures that distinguish between current-locator-address and persistent-locator-address
  - i.e. MIP6

Whats Next?

- Lots of ways we COULD do this
  - whats the BEST approach?
  - is this a solution for just Multi-Homing?
    - Or are we talking about mobility, NAT traversal and the general model of identity-based transport, locator-based packets?
  - whats the minimal possible change that creates the best benefit?

Survival of the Fittest

- Often it's the most adaptable creation that survives
  - Adaptability implies making minimal demands on others in order to reduce complex interdependencies
  - Adaptability implies being able to create outcomes that are valued in any environment
- The essential combination for IP to survive and thrive is that of simplicity and functionality

Thanks

- **Questions?**