Who Are You?

Identity and Location in IP

Geoff Huston
Chief Scientist
APNIC

Addresses and the IP Architecture

Architecturally, IP addresses are:

- Drawn from a stable global space
- Intended to be used in a unique context

Within the IP architecture addresses are:

- Endpoint identifiers
- Routing objects
- Key lookup value for forwarding function

IP Addresses are:

A means of uniquely identifying a device interface that is attached to a network

- **Endpoint identifier**

A means of identifying where a device is located within a network

- **Location identifier**

A lookup key into a forwarding table to make local switching decisions

- **Forwarding identifier**

This deliberate overload of semantic intent has been a basic and consistent property of the IP architecture

Challenges to the IP Address Model

- Roaming endpoints - Nomadism
- Mobile endpoints – Home and Away
- Mobile Networks
- Session hijacking and disruption
- Multi-homed endpoints and “session” resiliency
- Scoped address realms
- NATs and ALGs
- VOIP
- Peer-to-Peer applications
- Routing Complexity and Scaling
- Middleware, DNS and Renumbering

Wouldn’t it be good if…..

Your identity was stable - irrespective of your location

You could maintain sessions while being mobile

You could maintain sessions across changes in local connectivity

That locator use was dynamic while identity was long-term stable

That the binding between locator and identity was verifiable

Anyone could reach you anytime, anywhere

You could reach anyone, anytime, anywhere

IPv6 and identity

This overt split of identification and location semantics does not happen today in IPv4 as we know it – a change of location implies a change of address and a change of address implies a change of identity

Wouldn’t it have been good if IPv6 had offered solutions in this space that allowed endpoint identity to be distinguished from location and forwarding functions

IPv6 digression

"Second-Comer" Syndrome:

This perspective can be phrased as: Unless IPv6 directly tackles some of the fundamental issues that have caused IPv4 to enter into highly complex solution spaces that stress various aspects of the deployed environment than I'm afraid that we've achieved very little in terms of actual progress in evolution of the IP architecture. Reproducing IPv4 with larger locator identifiers is not a major step forward – its just a small step sideways!

Another IPv6 digression

"We've Been Here Before" Warning:

Of course this burdens the IPv6 effort in attempting to find solutions to quite complex networking issues that have proved, over many years of collective effort, to be intractable in IPv4.

If the problem was hard in an IPv4 context it does not get any easier in IPv6! That should not stop further exploration of the space, but it should add a touch of caution to evaluation of solutions in this space.

What do we want from “Identity”?

Generally we prefer identity systems that have:

- Uniqueness
- Persistence
- Structure
- Clear scope of applicability
- Validity and Authenticity
- Clear lines of derivation authority

Identity is not a unilateral private assertion – it is better viewed as a public recognition of derived uniqueness within a commonly understood context

Mob Mentality

Given a sufficiently valuable problem, **everyone** wants to solve it!

But everyone wants to solve it **their** way!

So far we have:

- Mobile IPv4
- Mobile Ipv6
- AD Hoc Networking
- NEMO
- HIP
- SCTP
- SHIM6
- Teredo
- Dynamic DNS
- NAPTR and SNAPTR DNS RRs

Choices, Choices, Choices

Its possible to find an identity concept at every level of the network protocol stack model

- **Application Identities** shared across transport sessions
- **Transport Identities** to allow agility of stack location
- **Host identities** to allow agility of location of all hosted sessions

In this context an “identity” is a token to allow both parties to agree on some set of packet-level locators that are to be recognised as belonging to a single communication state at both (or multiple) ends of the communication

Choices, Choices, Choices

Identity at the Application level

- Use a stable name space that is mapped to a locator (using the DNS)
  - DNS dynamic incremental updates
- Allow indirection and referral via DNS NAPTR records
  - Generic identity ornamented with service-specific mappings
  - ENUM
- Use application agents to provide stable rendezvous points
  - For example: *sip:gih@voip.apnic.net*

Application level identity

- Issues:
  - Can the DNS support dynamic interaction at a suitable scale and speed?
  - Would applications converge to a single identity framework, or is divergence a natural tendency?
  - Can we stop application designers from creating NAT-agile locator-independent application-specific solutions that rely on an application-specific identity space?
  - Is the proliferation of families of diverse application-specific identities desireable? (what about cross-application referral and hand-over?)

Choices, Choices, Choices

Identity at the Transport level

- Can we provide a mechanism to allow identity / locator independence at the session level?
  - An application opens a session with a generated session identity token
  - The identity token is dynamically associated with locator pairs
  - Changes in locators do not change the session token
- Application of the layering approach
  - Allow applications to assume a framework of identity association
  - Perform identity / locator association at a lower level of the protocol stack
  - Use opportunistic identity values that have a limited context and role of supporting session integrity
  - Support legacy applications by providing a consistent API

Choices, Choices, Choices

Identity at the IP level

- Can we provide an identity / locator association that is shared across multiple sessions?
- Reduce the overhead of identity locator mappings to allow all sessions to a common endpoint to share a mapping state
- Want to provide a more comprehensive support of identity to support both session-oriented transport protocols and (potentially) datagram transactions
- Reduce the complexity of applications and transport sessions and place the per-endpoint mapping state in the IP level

Identity Issues

How could an identity mapping function?

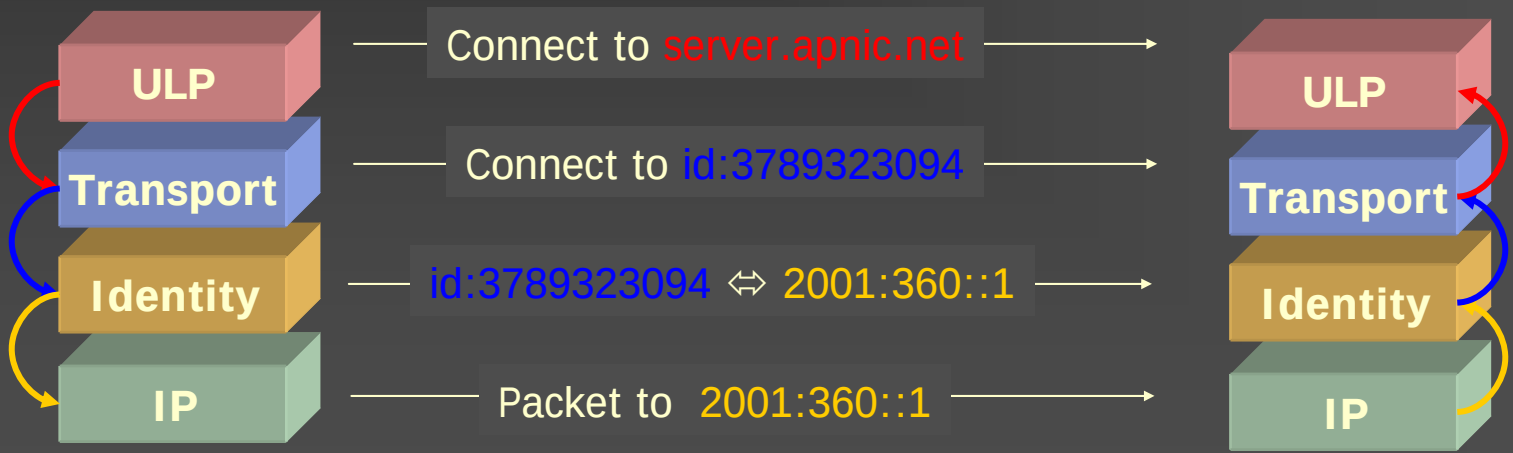

Identity Issues

How could an identity mapping function?

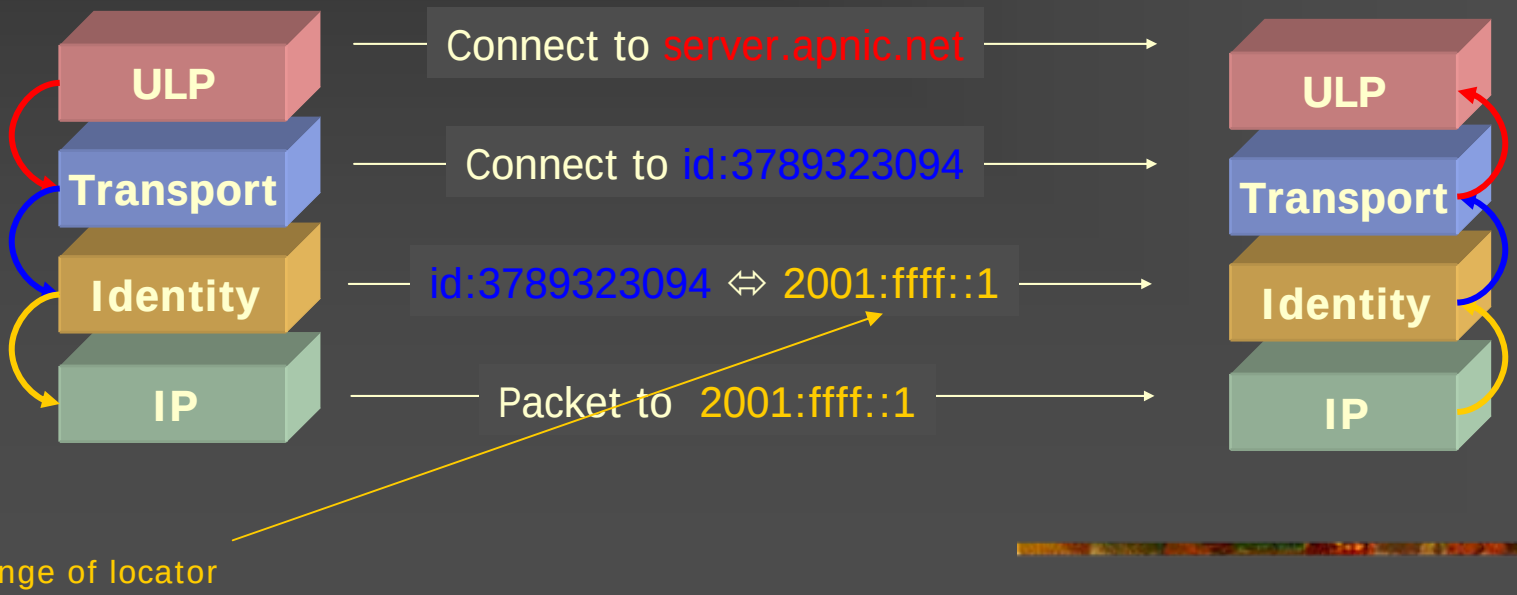

Identity Implementations

"Conventional"

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

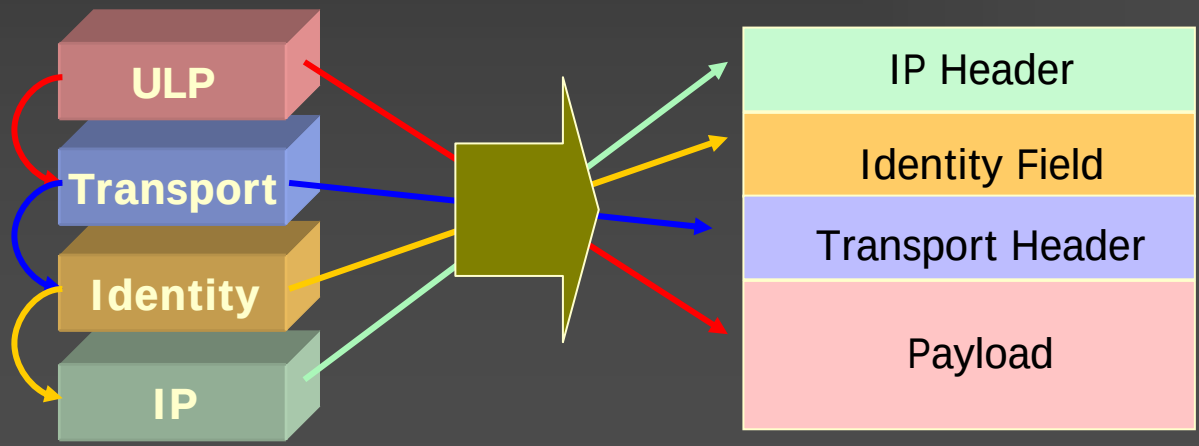

Protocol element

Packet Encapsulation

Identity Implementations

"Out of Band"

- Use distinct protocol to allow the protocols element to exchange information with its peer

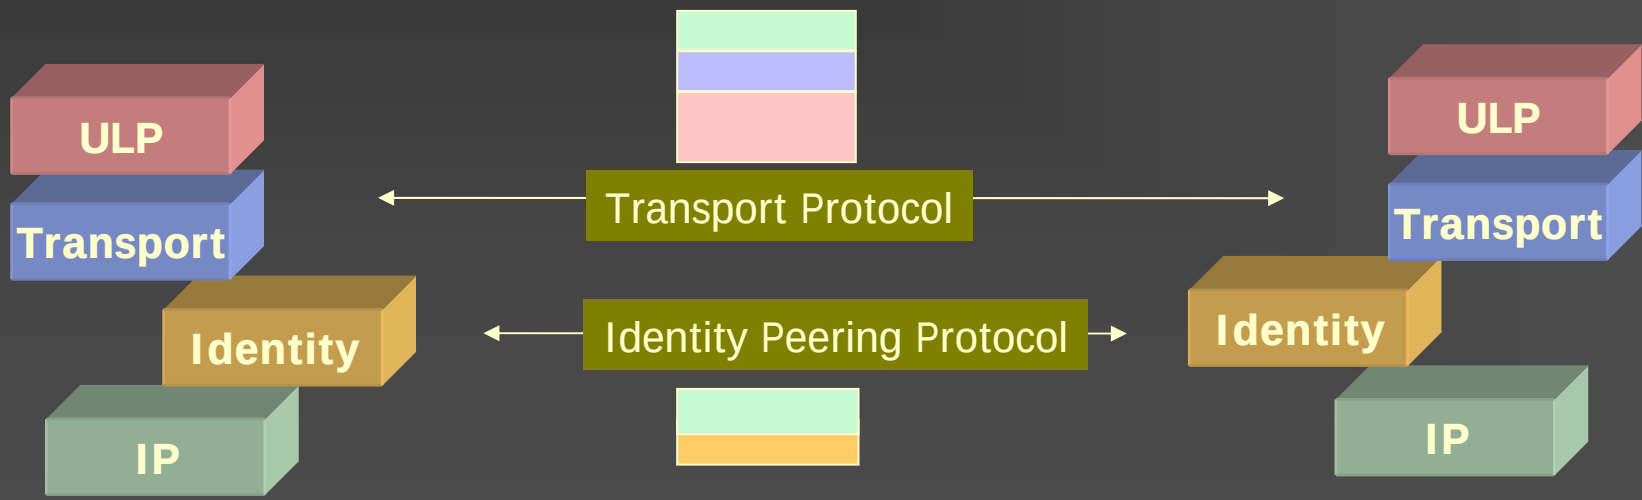

Identity Implementations

Application Identity: Above the Session

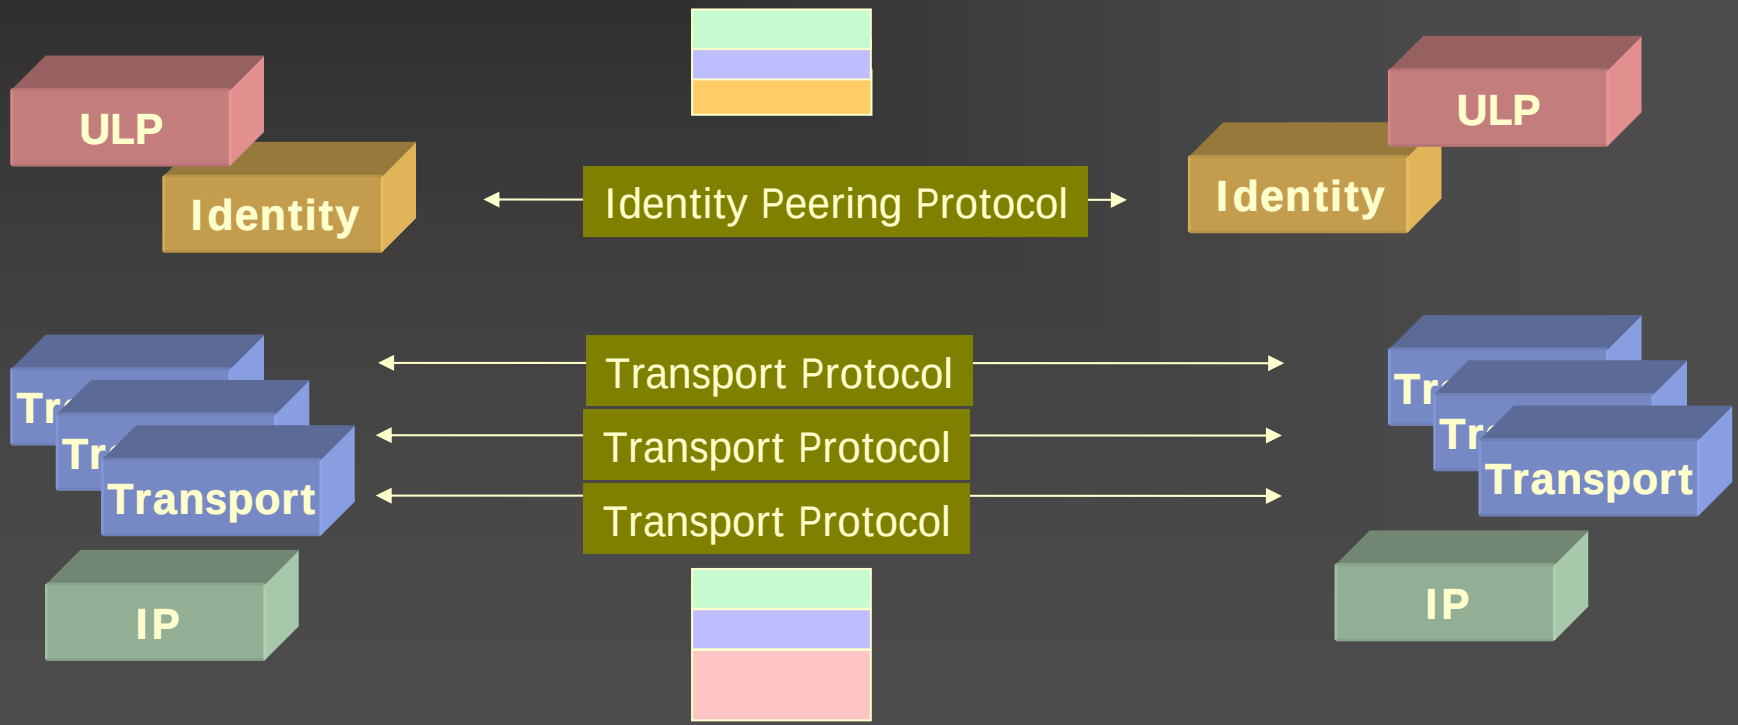

Identity Mapping

"Referential"

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

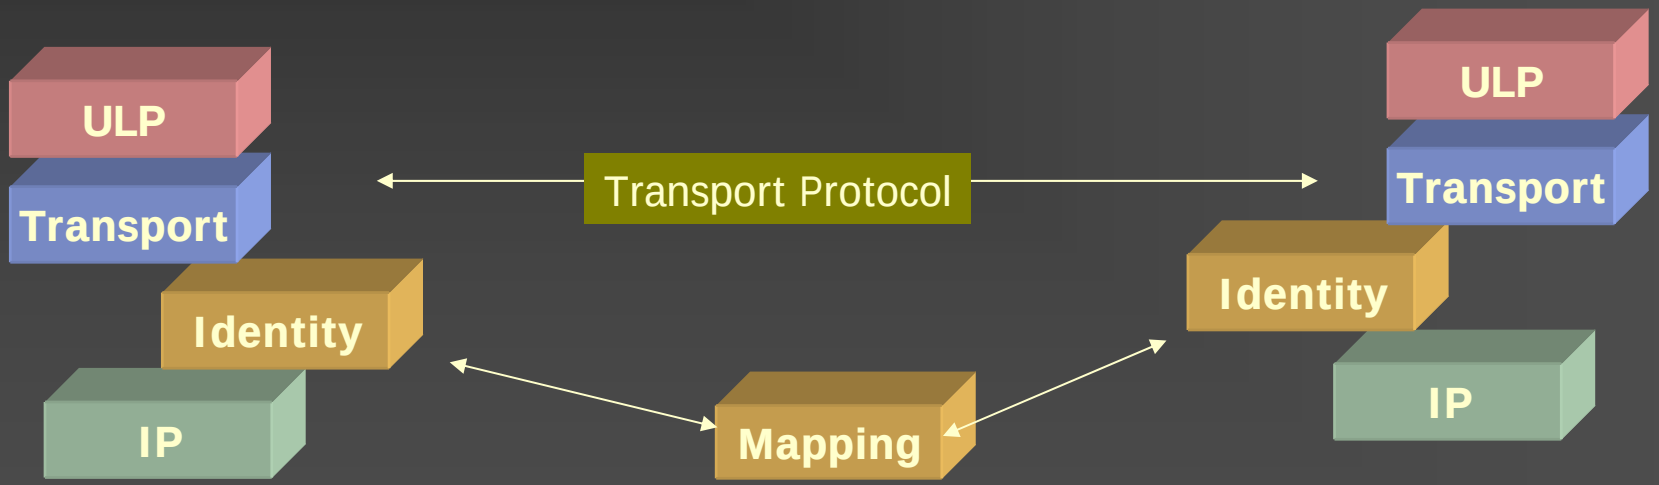

Identity Implementations

Self-Referential

- Use an opportunistic identity as an equivalence token for a collection of locators

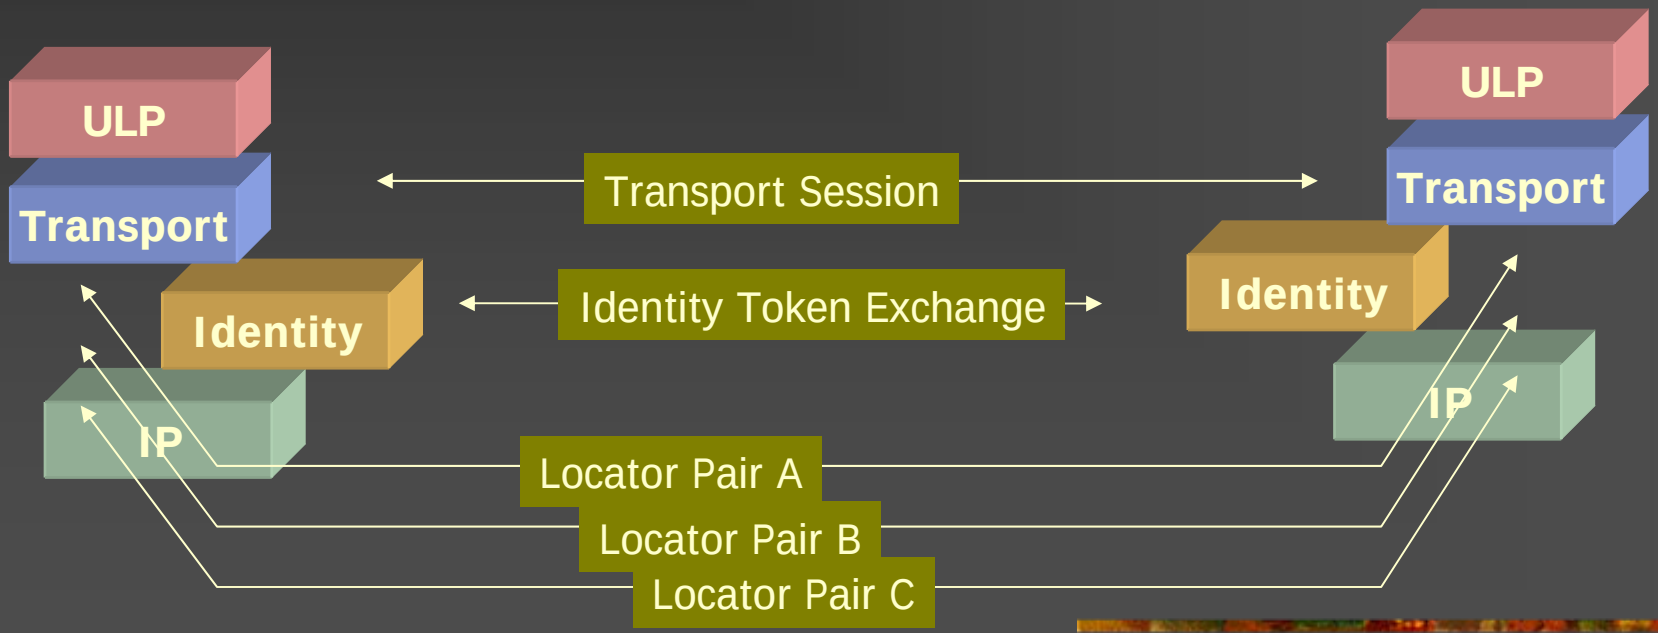

Identity Types

Use identity tokens lifted from a protocol's "address space"

- DNS, Appns, Transport manipulate a "distinguished address"
- IP functions on "locators"
- New Protocol Stack 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 new token space that distinguishes it from the above?

Unstructured token

- Allows for self-allocation of identity tokens that may not globally assuredly unique (opportunistic tokens)
- How to map from identity tokens to locators using a lookup service? Or how to avoid undertaking such a mapping function

Some Identity Suggestions

- IPv4 Address
- Centrally Assigned IPv6 Unique Local Addresses
- A crypto hash of your public key
- A crypto hash of a set of locator values
- The IPv6 address used to initiate the communication
- IPv6 Address
- DNS names
- URIs
- Telephone numbers

Identity 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
- Context of use determining semantic interpretation

Upper Level Issues of Identity Realms

- The significant effort and cost of supporting a new global unique token distribution system as an endpoint identity system
  - Uniqueness is not cheap!
- The side-effects of reusing some other existing token set as an identity set
  - Recycling identity is generally dangerous! (EUID-64?, E.164?...)
- The issue of support of dynamic identity to locator binding
  - Speed vs accuracy
- The protocol overhead of identity handshake for datagram transactions
- The security issues in maintaining integrity of identity

IPv6 and Identity

Is the 64bit Interface Identifier a rich location for carrying opportunistic identity?

Can the Flow-Id field be exploited?

Are header extensions and options useful?

Is packet inflation necessary?

Is IPv6 the only protocol for consideration of IP level identity approaches?

- Is there any leverage for transport session approaches?
- Can such approaches be IP version agnostic?

*

Our current direction appears to be developing solutions in **all** of these spaces simultaneously:

- Multi-Party Applications
- Application Agents
- Rendezvous protocols
- DNS Incremental Updates and DNSSEC
- DNS Indirection and Referral
- SCTP, HIP at the transport-layer
- Shim6
- Mobile IPv6
- Mobile IPv4
- And probably many more!

* *Let a hundred flowers bloom: let a hundred schools of thought contend*
Mao Zedong, 1956

*

- Should there be just one SINGLE identity model?
  - Can we coerce all the requirements of identification into a single identity realm?
  - Or are the various desireable properties of identity so mutually contradictory that there is no single solution?
  - Many of the approaches we've see assume a single identity mapping function that wants to ignore any side effects from any other simultaneous identity scheme in operation
  - One the other hand, retrofitting any form of identity function into today's IP architecture is bad enough in terms of legacy requirements without having to factor in other identity functions as well!

** Let one flower bloom: let one school of thought prevail*

Where Now?

We appear to have a lot more to learn here

- Its more than scaling routing or avoiding renumbering or multihoming
- Its more than IPv4 and IPv6
- If we want a more agile and flexible model of packet networking to support diverse communications environments and services then we need to understand how split apart the semantics of who wants to talk from the mechanics of how they may talk

Thank You