IAB Open Meeting:
IP Address Architectures

IETF 57

July 2003

IP Addresses are:…

- A means of uniquely identifying a device interface that is attached to a network
  - Endpoint identifier (who)

- A means of identifying where a device is located within the network
  - Locator identifier (where)

- A means of allowing intermediaries to pass a packet to a destination device
  - Forwarding identifier (how)

These roles are somewhat different and there is visible tension between the constraints of routing and the requirements of end-to-end service sessions

Address Types

IPv4 has two address types:

- Unique Structured Addresses == Global-Use Internet
- Private (Reused) Addresses == Local-Use Internets

IPv6 has three address types:

- Unique Structured Addresses == Global-Use Internet
- Site-Local Addresses == Scoped Local-Use Internet
- Link-Local Addresses == Very Local-Use

Address Realm Membership

- IPv4 address architecture assumes a unique 1:1 binding of a device interface to an IP address in a single realm membership model
  - although it is not strictly required
- IPv6 address architecture is exploring the notion of a device acquiring multiple addresses and being a member of multiple address realms simultaneously
  - Although its not entirely clear how this works in practice and what issues this architecture raises and what issues it addresses

IAB Open Meeting

- The IAB held an Open meeting this week to gather input on this topic
- Five Presentations, each exploring different aspects of this area:
  - The impacts of wireless models on the layered network architecture
  - The potential uses of explicit scoping in address realms
  - Transport considerations of multi-homed and mobile environments
  - Insights gained from ZEROCONF
  - Conflicting requirements placed on addresses

Observations

- This area of exploration of the semantics of an address within the IP architecture is not a new topic
- Recent (ish) activities include
  - Name Space Research Group
  - IAB output….

The IAB Time Machine said…

"As far as temporal uniqueness (identifier-like behaviour) is concerned, the IPv6 model is very similar to the current state of the IPv4 model, only more so…IPv6 will amplify the existing problem of finding stable identifiers to be used for end-to-end security and for session bindings such as TCP state.

The IAB feels that this is unfortunate, and that the transition to IPv6 would be an ideal occasion to provide upper layer end-to-end protocols with temporally unique identifiers. The exact nature of these identifiers requires further study."

- RFC 2101, February 1997

Areas of Relevance

- Addresses are used in many contexts within the IP environment. Particular contexts where address semantics have particular relevance include:…
  - Mobility in IPv4 and IPv6
  - Security associations
  - Routing architectures
  - Scoped address contexts
  - Transport protocols
  - Multi-Homing
  - Application Program Interface (API) to the network stack

  - And doubtless there are others….

Next Steps for the IAB

- The IAB Open Meeting presentations and meeting minutes will be included in the proceedings of the IETF
- Document the considerations raised at the meeting
- Create a moderated forum for further consideration of these issues

Objectives of this activity

- Definitely….
  - Gain a clearer understanding of
    - the roles of addresses and the attributes of addresses used within these various roles
    - the level of inter-dependency between these roles of addresses
- Possibly…..
  - Document some architectural considerations relating the distinguished use of addresses in various contexts

IP Address Policy evolution

1980's : Anyone can apply for an address block and they can obtain one

early 1990's : Anyone can apply for an address block, but they need to demonstrate that they can make 'good' use of it

late 1990's: You can apply for addresses for use in the Global Internet, and you will need to justify their use (and make a case why PA space is inappropriate)

IP Address Structure

- Why use any structure at all?
  - Blocked structure
    - Ease of administration
    - Aids the mapping of addresses into scaleable forwarding mechanisms
  - Hierarchical structure
    - Assist scaling of administration mechanisms through delegation levels
    - Assist scaling of computation within routing domains

IPv4 -- PI, PA or NAT?

- PI - Provider Independent Address blocks
  - ☺ Not tied to a particular provider
  - ☺ Readily supports various forms of multi-homed configurations
  - ☹ Adds non-aggregateable entries into the routing system
- PA - Provider Assigned Address blocks
  - ☺ Aggregateable in the routing system
  - ☹ Requires renumbering on provider switch
  - ☹ Multi-homing is challenging
- PU - Private Use Address Blocks + NAT
  - ☺ Avoids consuming global common resources and bypasses global address allocation policies
  - ☺ Can support connectivity services through forms of network address translation at the connectivity interfaces, mapping into external PA space
  - ☹ But such supported services are limited in scope and the approach introduces additional points of vulnerability

Options for IPv6

- What options are available for IPv6 for private use
  - no need to apply 'strong' address conservation mechanisms to this space
  - but would prefer to avoid the implicit creation of a routing swamp of unaggregatable PI routing elements

IPv6 site-local addresses

- Why?
  - Auto configuration of non-connected networks
  - Private use realms that do not reference the global Internet, bounded by a 'site' scope
  - Access to private use address space without reference to public use policy and public use distribution mechanisms
  - Define implicit 'scoping' of service visibility through use of site-local addresses

  - But are not unique
  - And site 'edges' are often unclear

IPv6 Global non-PA Local Use

- Why?
  - Address availability for non-connected networks
  - Private use realms that do not reference the global Internet, bounded by route propagation controls
  - Assured access to private use address space without reference to public use policy and public use distribution mechanisms
  - But require a distribution mechanism to provide assurance of unique access to each local use block

IPv6 Link Local Addresses

- Why?
  - Ad hoc networks
  - Bootstrap point for Neighbor Discovery
  - But cannot scale to cover a site

Non-Unique Site Addresses

- Ambiguity
  - cannot be used as referrals,
  - can't detect if an address have crossed the 'zone' boundary
- DNS
  - DNS is global naming infrastucture, so direct conflict
  - requires split DNS ("views" in BIND9)
- unidirectional communication (NAT)
  - hope there's no IPv6 NAT...
- If scoping a 'site' is related to network topology
  - how do you do address selection?
- Must nodes be aware of the available address types, network topology and connectivity policies?
  - in order to make 'correct' choices about what address to use for each network transaction
  - i.e. should network topology and connectivity be explicitly visible to higher levels in the protocol stack?
  - How should the DNS interact with applications?

Globally-unique but not globally routable addresses

- no ambiguity
- As it shouldn't leak 'out' in any case, will the cost for "uniqueness" pay off?
  - who is to maintain "unique" property?
  - how?
    - Probability or record keeping?
- not globally routable == not usable for global communication
  - is this a realistic expectation?
- nodes must be aware of the address types
  - and make 'correct' choices about what address to use for each network transaction
  - Same considerations as previous

IPv6 link-local addresses

- a node itself knows the boundary
  - interface/set of interfaces will identify boundary
- ambiguity
- DNS
  - Link Local Multicast Name Resolution?
- node must be aware of the address types, if Link Local and global addresses are used concurrently
  - in IPv6 concurrency is the case

Scoping the 'zone' boundary

- how do we define a scoped zone boundary?
- how does one scoped zone interface with another?
  - how many scoped zones can a node claim simultaneous membership?
  - if more than one, how does it tell them apart?

Network Address Translation

- ☺ Address use on the "inside" is independent on the "outside"
- ☺ Disjoint "insides" can reuse the same address block
  - 😐 And, with (a lot of) care, disjoint insides can communicate with each other directly
- ☺ NATs can be incrementally deployed
- 😐 NATs are very widely used and are part of the current IP landscape
- ☹ Some protocols use IP address and port numbers
  - ☹ which means the data inside the protocol must be changed, as the conversation ends up being different depending on "what view" the peers have of each other, which means that the NAT must look inside the protocol conversation, not just in the outer IP header
- ☹ NATs weaken strong security models
  - ☹ Changing the packet header breaks any form of authentication header guard (IPSEC AH)
  - ☹ Changing the payload in flight is hard to secure (signing the payload, and then verifying the data is not possible)
- ☹ NATs re-introduce network state. Failure of a NAT unit breaks all active sessions using that NAT unit.
- ☹ NATs induce protocol complexity.
  - ☹ The implicit 'one-way' model of transaction initiation is very limiting. Solutions to circumvent this to allow external initiation of transactions add considerable complexity and inter-dependencies
  - ☹ New applications require specific NAT-traversal consideration (ALGs, etc)
- ☹ NATs have limited topologies of application – they work predictably as stubs to the global 'core'. They are less predictable in other contexts.

Overlapping Address Realms

- In order to avoid boundary address translation mechanisms, and be able to map some local elements into a larger connectivity domain, the local domain may need to support multiple address realms
  - This implies that applications may need to make explicit choices about which address realm to use to address the remote device, and which address to use to represent themselves
  - When is this necessary?
  - What network level information is required to pass to the application layer to allow it to make the realm choice for each network transaction?

Related Notes – Site Local (PAF)

- Selection
  - An application do not know which one of a number of possible addresses to use, especially it can not guess "the best one" given network topology. The application layer simply doesn't know about network topology.
- Reference
  - If A communicate with B, and then B is to communicate with C, it might be that A and B, and B and C (respectively) is in the same "site local scope" but A and C is not. So, B can not tell A where C is.
  - Many protocols do pass around IP addresses, like ftp and sip and dns, and this is bad enough given NAT and RFC 1918 addresses. We can not get something worse when you have different scoped addresses, because one loose the coherence in addressing we have in the baseline architecture of he Internet. "If A is connected to the internet as 'a' and B as 'b', then we know 'a' can address B as 'b'. And if A can address B as 'b', then C should as well.

- Basically, I think site local is so stupid I don't understand why it needs to be written down..... ;-)

- Yes, there are arguments for when "scoped addresses" and "RFC 1918 like addresses" are good, BUT, my view is that they most have to do with (a) lack of a session layer and (b) people think policy issues regarding routing and reachability should be built into the address architecture. I.e. we can not do much about (a) and regarding (b) it should be solved by other tools. Like having "normal" packet filters in the "firewall" one should have at the edge of the site one own anyway.

Related Notes – Site Local – (SRA)

- Scopes And Borders
  - Scopes (which imply borders)
    - node
    - link
    - site
    - global
  - Things that change at borders
    - routing
    - security
    - naming
    - addressing
  - Is single "site" border a good place to put a border for all of these things?
- Applications and Scope
  - Some applications are intrinsically scoped (eg: RA, ND)
  - Most applications have no concept of scope
    - Globally scoped by design
  - Most applications have no way of expressing scope
    - Scope constrained by mechanisms external to the protocol
  - Stuff leaks across the borders
    - Names leak (mail, web, files)
    - Addresses leak (early name->address binding)

Related Notes –Site Local (SRA)

- One Size Does Not Fit All
  - Site border sounds at first like a nice simple approach
- ...But it's wrong
  - Are these the same border?
    - Autonomous system
    - Address realm
    - Two-faced DNS border
    - Firewall
    - Demarcation point
- Private addresses do not enhance security
  - Attacks via a border machine
  - Attacks via leakage
- Weakened node security due to false sense of security
  - Firewalls have to filter bad global stuff anyway
  - Private addresses are just one more thing to filter
  - Private addresses do not make filtering easier

Related Notes – Site Local (SRA)

- Reachability versus Ambiguity
  - Firewalls limit reachability
  - But if you do get through, it's not ambiguous
  - Private address realms also limit reachability
  - But if you do get through, it is ambiguous
  - This is not an improvement
    - draft-ietf-dnsop-dontpublish-unreachable
- Multiple sites
  - Devices that have to live in multiple sites are hard
    - Multiple routing tables
    - Mulitple naming realms
    - Multiple (potentially colliding) addressing realms
    - Complex forwarding and leakage rules
- Recommendations
  - If we have to keep site-local at all, only use in disconnected case
  - Globally unique addresses would be better even in disconnected case