The How of Where

Some Observations on IPv6 Addresses

Geoff Huston
APNIC

The IPv6 Vision

Communications as a commodity service:

- anywhere, anyhow, anytime
- present-and-play auto-configuration
- every device with an IP protocol stack

  appliances, automobiles, buildings, cameras, control units, embedded systems, home networks, medical devices, mobile devices, monitors, offices, output devices, phones, robots, sensors, switches, tags, Vans ....

And every device will need an address...

What do we want from addresses?

- Assured Uniqueness
- Verifiable Authenticity
- Routeability
- Simplicity
- Stability
- Assured availability
- Low cost

What do we want from IPv6 addresses?

- Servicing Ubiquity
  - Global populations of people, places, activities, devices,…
- Simplicity
  - Easy to obtain, easy to deploy, easy to route
- Longevity
  - 70 - 100 year technology lifespan
- Commodity
  - Low cost per address
- Scaleability
  - Global end-site populations of the order of hundreds of billions of sites

3 Questions:

- Does the IPv6 address plan scale to meet these expectations?
- What forms of distribution are most appropriate here?
- Are addresses long-term stable?

Scaling: How many addresses?

IPv4: 32 bits

- IPv4 provides 2^{32} addresses

  = 4,294,967,296 addresses

  = 4 billion addresses

IPv6: 128 bits

- IPv6 provides 2^{128} addresses

  = 340,282,366,920,938,463,463,374,607,431,770,000,000

  = 340 billion billion billion billion addresses

Just how big is 2^{128}?

"If the earth were made entirely out of 1 cubic millimetre grains of sand, then you could give a unique [IPv6] address to each grain in 300 million planets the size of the earth" -- Wikipedia

IPv6 Address Structure

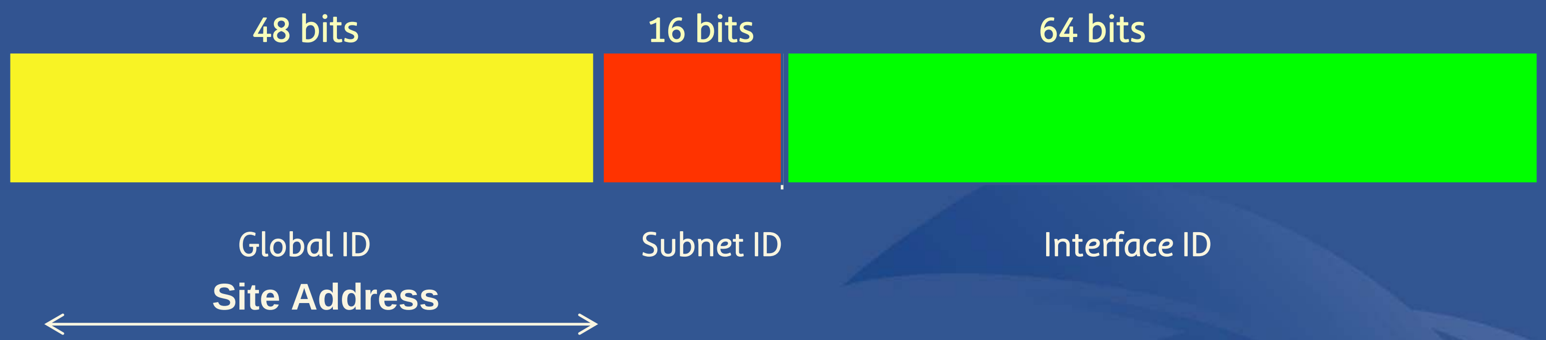

- IPv6 provides 2^{48} end site addresses
  - = 281,474,976,710,656
  - = 281 thousand billion end site identifiers

Address Utilization Efficiency

- Addresses utilized will be far fewer than addresses available
- Larger deployments are generally less efficient than smaller deployments
  - Because of hierarchical addressing architecture
- Host Density Ratio defines utilisation in hierarchical address space:

$$HD = \frac{\log(utilised)}{\log(total)}$$

- Value of 0.8 initially suggested for IPv6

- IPv6 will provide 0.0013 x 2^{48} site addresses

  = 362,703,572,709

  = 362 billion end site identifiers

Asia Pacific Network Information Centre

APNIC

Current Considerations

- Can this useable identifier pool be expanded without altering the address structure?
  - Consideration of higher values for the threshold value of the HD Ratio
    - 0.94 appears to offer a reasonable balance between address utility and higher efficiency
  - Consideration of a /56 end-site allocation for SOHO sector end sites
    - Allows for up to 256 distinct subnets per end site
    - More suitable for home, small office, small cluster networked sites than a /48

- IPv6 can provide 0.1×2^{52} site addresses

  = 450,359,972,737,050

  = 450 thousand billion end site identifiers

  = 4.5×10^{15} end site identifiers

The Demand Model

- The demand - global populations:
  - Households, Workplaces, Devices, Manufacturers, Public agencies….
  - Thousands of service enterprises serving millions of end sites in commodity communications services
  - Addressing technology to last for at least tens of decades, and perferably over a century
  - Total end-site populations of tens of billions of end sites

i.e. the total is order (10^{11} - 10^{12}) ?

So we need to have a useable end-site identifier pool of some 10^{13} identifiers.

3 Questions:

- Does the IPv6 address plan scale to meet these expectations?

  Yes
- What forms of distribution are most appropriate here?
- Are addresses long-term stable?

Distribution Mechanisms - Objectives

- Preserve valued attributes
  - Ensures that distributed addresses are assuredly unique, have clear lines of authenticity, and support routeability
- Maximize current utility
  - Readily available to meet network demand with low marginal cost of deployment
- Maximise future utility
  - Readily available to meet various future demand scenarios
- Minimize distribution overheads
  - Low cost of access

Distribution Mechanisms - Risks and Threats

- Any distribution system can fail – the forms of possible failure include:
  - Exhaustion
  - Induced scarcity
  - Hoarding
  - Fragmentation
  - Instability of supply
  - Pricing distortions
  - Forced renumbering
  - Speculative acquisition and disposal
  - Erosion of assured uniqueness and/or authenticity
  - Theft and Seizure

Potential Mechanisms – Characteristics

- Distribution
  - Allocations / Auctions / Markets
- Title
  - Freehold / Leasehold
- Circulation
  - Tradeable Asset / Restricted Use
- Structure
  - Uniform / Various
- Nature
  - Global / Regional / National / Industry
- Pricing
  - Asset-based pricing / Service-based pricing

Distribution Frameworks

- Allocation Scope
  - Global / Regional / National ?
  - Public / Private / Hybrid ?
  - Coordinated function / Multi-source competitive framework ?
- Supporting Authenticity
  - Trust points
  - Accuracy of information
  - Currency of information
- Supporting Routeability
  - Supporting an allocation framework that supports hierarchies of aggregation within the routing system
  - Service provider alignment

Some Lessons from IPv4

- Address distribution characteristics
  - simple, uniform and generic
  - consistent and stable
  - relevant
  - routeable
  - accurate and trustable
- Some useful considerations:
  - Be liberal in supply (but not prolifigate!)
  - Avoid “once and forever” allocations
  - Avoid creating future scarcity
  - Plan (well) ahead to avoid making changes on the fly

National Distribution Channels?

- To what extent would national regimes impose particular constraints or variations on address use conditions?
  - How would you put these constraints into your routers?
  - What additional overheads would ensure?
- What is the underlying network model?
  - National service operations interlinked by bilateral arrangements?
  - Heterogenous service industry based on private sector investments at the local, regional and global levels
- Are there end-user visible IP address semantics?
  - Toll or international address prefixes?
- Is there the risk of scarcity in IPv6 addresses?
  - At last count we appear to have provision for 225,179,981,368,525 useable end site address prefixes. This appears to be adequate for the most optimistic forecasts of IPv6 lifetime address consumption.

Competitive Distribution Channels?

- What would be the basis of competition?
  - Pricing, Policies, Use Restrictions, Local regulation?
    - It appears likely that competition would be based predominately on policy dilution in the distribution function.

- Would this enhance or erode address attributes?
  - Availability, Uniqueness, Stability, Routeability, Confidence?
    - A regime of progressive policy dilution would expose consequent risks of increased routing overheads address fragmentation and restricted address policies, dilution of authenticity and integrity, the potential for gains derived from hoarding and speculative pricing ,and erosion of confidence in the address distribution system

- Would this enhance or erode IPv6 viability?
  - Scaleability, Stability, technology lifecycle

What form of distribution is most appropriate for the future IPv6 commodity network?

- Accommodates multi-sector needs and interests
- Preserves strong address integrity
- Stays within technology bounds
- Highly stable
- Very simple
- Very cheap

Today's IP Address Distribution System

- Industry self-regulatory framework
  - Consensus-based, open and transparent policy development processes
  - Balancing of interests
- Reflective of global trend to deregulation and multi-sector involvement
  - Policy development process open and accessible to all interested parties
- Separation of Policy and Operation
  - Non-profit, neutral and independent operational unit
  - Consistent application of the adopted policy framework
- Structured as a stable service function
  - Self funded as an industry service function
  - Preserve address integrity

What are we really trying to achieve here?

The distribution of network addresses is an enabling function, and not an enduring value proposition in its own right. The enduring value proposition here lies in the exploitation of networked services to create value.

Asia Pacific Network Information Centre

APNIC

3 Questions:

- Does the IPv6 address plan scale to meet these expectations?

  Yes

- What forms of distribution are most appropriate here?

  Addresses multi-sector needs and interests, preserves address integrity, operates with low overhead and is highly stable

- Are addresses long-term stable?

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

Challenges to the IP Address Model

- Roaming endpoints - Nomadism
- Mobile endpoints – Home and Away
- Session hijacking and disruption
- Multi-homed endpoints
- Scoped address realms
- NATs and ALGs
- VOIP
- Peer-to-Peer applications
- Routing Complexity and Scaling

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
- Anyone could reach you anytime, anywhere
- You could reach anyone, anytime, anywhere

Wouldn’t if be good if…

- IPv6 offered solutions in this space that allowed endpoint identity to be distinguished from location and forwarding functions

1. “Second-Comer” Warning:

This perspective can be phrases 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 IPv6. Reproducing IPv4 with larger locator identifiers is not a major step forward – its just a small step sideways!

2. “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 very challenging in IPv4. If the problem was hard in an IPv4 context it will not get any easier in IPv6!

Where next?

- One view is that the overloaded semantics of IP addresses is not sustainable indefinitely
  - 128 bits of address space has not provided a new routing architecture
  - Hierarchical network-aligned addressing is the only way we know how to support large-scale inter-networks.
  - This constrains identity attributes in a “your address is your identity” realm
- If we want more natural identity attributes from IPv6 (persistence, reference, relevance and usefulness) then we need to consider further protocol refinements that treat endpoint identity and endpoint location as a dynamically discoverable association

3 Questions:

- Does the IPv6 address plan scale to meet these expectations?

    Yes

- What forms of distribution are most appropriate here?

    Addresses multi-sector needs and interests, preserves address integrity, operates with low overhead and is highly stable

- Are addresses long-term stable?

    We need to consider forms of identity / location splits within the protocol architecture. This is a current research topic

Thank You

Questions?