IPV6 OUTLINE

I. IPv4 Depletion

-- The Internet is rapidly running out of globally-routable IPv4 addresses. the sort of IP addresses that have largely been used on the Internet to-date.

-- If an Internet Service Provider (ISP) is unable to obtain sufficient new globally-routable address space, it may be difficult or impossible for a new ISP to even get started, or for an existing ISP to readily grow its service, although obviously all currently assigned IPv4 addresses will continue to work.

-- IP addresses are assigned through a regional model, with ISPs getting IP addresses from the Regional Internet Registry (RIR) servicing their region. For example, North American ISPs ask for address space from ARIN, while European ISPs ask RIPE, Asian ISPs ask APNIC, Latin American ISPs ask LACNIC, and African ISPs ask AFRINIC.

-- The regional registries have differing inventories of assignable IPv4 address space, and serve communities with different rates of consumption, and hence will run out of assignable IPv4 address space at different times. RIPE and APNIC are now making IPv4 address assignments from their last /8, while ARIN is down to just 1.65 /8s, and LACNIC is down to 1.72 /8s. AFRINIC currently has the most remaining IPv4 address space, with 3.54 /8s remaining. (A /8 is 16,777,216 addresses). Geoff Huston's projected run out dates for each regional registry can be seen at http://www.potaroo.net/tools/ipv4/index.html

-- Once RIRs begin assigning address space from their last /8, RIR policies normally limit assignments from that last block to a single /22 (1,024 IPv4 addresses) per applicant, even if the applicant would formerly have been able to justify and obtain a far-larger assignment. A /22 will generally be too small to independently announce via BGP.

-- Organizations that need additional IPv4 address space, and which can justify the addresses they require, may be able to acquire additional legacy IPv4 address space on the secondary market, provided they can find someone with legacy space that's available. For example, in March 2011, Microsoft agreed to pay Nortel $7.5 million for 666,624 IPv4 addresses, or $11.25/IPv4 address. (see http://www.networkworld.com/community/blog/microsoft-pays-nortel-75-million-ipv4-address)

II. IPv4 Network Address Translation (NAT)

-- Given that the Internet has nearly exhausted its supply of globally-routable IPv4 addresses, some are using network address translation (NAT) to "stretch" the available supply of IPv4 addresses. When NAT is used, one (or a small number of publicly-routed IP addresses, when NAT is used for larger deployments) get shared across multiple devices located behind the NAT device, typically with the devices behind the NAT using non-routable RFC1918 private address space.

-- Consumers have most commonly encountered NAT in conjunction with home wireless "routers," customer premises equipment (CPE) which allow an entire household of computers to share an IP obtained from the household's ISP.

-- While NAT has historically been most notably a consumer technology, the shortage of publicly-routable IPv4 addresses has also caused some ISPs to employ NAT. When employed by ISPs, NAT is normally referred to as "large scale NAT" or "carrier grade NAT."

-- NAT might seem to be a clever way to make limited globally-routable IPv4 address space last longer, but NAT is not without its drawbacks. For example:

(a) public servers normally are not deployed behind NAT;
(b) NAT may make it difficult or impossible to use some existing applications such as SIP or H.323 video conferencing;
(c) NAT may inhibit the development and deployment of new/innovative applications;
(d) because NAT multiplexes multiple downstream customers across a single IP address, misbehavior by one customer may negatively affect the "IP reputation" of all the customers sharing the same public address;
(e) use of NAT may also complicate the process of identifying a specific problematic system, making it critical that abuse reports always include source port information as well as the usual source IP and accurate time stamp.

-- A nice discussion of the importance of end-to-end transparency (lost when relying on NAT) can be found in RFC 2775 and RFC4924.

III. IPv6 Deployment

-- The alternative to limping along with NAT is to begin to use IPv6. Abundant supplies of globally routable IPv6 address space are readily available for assignment to ISPs, but unfortunately, IPv6 deployment, and traffic levels flowing over IPv6, remain quite low. For example, as of November 1st, only about two percent (2%) of all traffic seen by Google was over IPv6 (see http://www.google.com/intl/en/ipv6/statistics.html), and there's still a tremendous number of "red cells" in Mark Prior's excellent IPv6 survey (see http://www.mrp.net/ipv6_survey/)

-- We postulate that in general, even most of those who are reading this document today will not be routinely using IPv6 connectivity (you can check if you are using a free IPv6 testing site such as http://ipv6-test.com/).

-- Low empirical levels of IPv6 network traffic may be due to many factors, including:

-- Adequate *local* IPv4 availability (the world may be running out of IPv4, but you/your ISP may still have plenty)

-- Assuming IPv6 transition mechanisms aren't in use, transmission of native IPv6 traffic requires an end-to-end IPv6-clean path. While a growing number of ISPs offer IPv6-native commodity transit, a non-trivial number of other ISPs still don't (http://en.wikipedia.org/wiki/Comparison_of_IPv6_support_by_major_transit_providers)

-- In order for IPv6 to be used, other critical support services also need to be IPv6-aware. In particular, registrars and name servers need to support use of IPv6 quad A records, and ideally, should offer access over IPv6 transport.

-- IPv6 support will also be needed on the operating system and the applications running on both endpoints. While modern operating systems and contemporary web browsers (and many other applications) now support native IPv6, legacy IPv4-only CPE devices (such as many home wireless "routers") may represent a particularly nettlesome end-site hardware obstacle to getting an IPv6-clean end-to-end path.

-- A fifth factor is that there is no meaningful population of Internet sites currently available *only* via IPv6, while all popular Internet sites remain reachable over IPv4 by default (to forgo offering IPv4 access would be to forgo an unacceptable fraction of a site's potential visitors, a choice that no sane Internet business will voluntarily make)

-- When a dual-stack-enabled (IPv4 *and* IPv6-connected) user connects to the (relatively rare) dual-stack-enabled Internet site, the endpoints face a choice: should they connect via IPv4, or via IPv6? While IPv6 adherents might champion a philosophical choice of IPv6 when that option exists, the pragmatic option may simply be to use IPv4 by default (for example, Firefox continues to prefer IPv4 over IPv6 by default on a dual stack-enabled system).

.

-- IPv6-only ISP eyeball networks are rare. To understand why, note that ISPs with IPv6-only connected customers need to provide some translation mechanism (or gateway device) that allows those IPv6-only customers to reach IPv4-only content (e.g., perhaps NAT64, c.f., RFC6145 and RFC6146). If deploying an IPv6-only eyeball network means that you need to make an investment in a large NAT64 gateway to reach the majority of Internet content -- content that's *only* available via IPv4 -- at least some providers may prefer to simply deploy traditional NAT (e.g., NAT44) with private IPv4 address space, thereby foregoing the need to work with IPv6 entirely.

-- IPv6 as deployed has also deviated from the original IPv6 vision in some important ways. For example:

-- Historically, IPSec support was an integral part of the IPv6 architecture, but Section 11 of RFC6434 (December 2011), clarifies that "Previously, IPv6 mandated implementation of IPsec and recommended the key management approach of IKE. This document updates that recommendation by making support of the IPsec Architecture [RFC4301] a SHOULD for all IPv6 nodes." Loss of mandatory IPSec support removes one potential incentive (e.g., improved security) for potentially transitioning from IPv4 to IPv6.

-- In addition to IPv4 address exhaustion, another important reason for deploying IPv6 was a desire to reign-in out-of-control growth in the number of prefixes carried in the global routing table. Originally, IPv6 was supposed to address this through use of strictly hierarchical (and thus easily aggregate-able) address assignments from each IPv6 transit provider, even if that meant that each endpoint received multiple IPv6 addresses. After years of work devoted to avoiding replication of the traditional IPv4 multihoming model (which relies on obtaining a provider independent prefix, and then announcing that prefix from multiple providers), the classic multihoming model appears to still reigns supreme in today's IPv6 networks, just as it did/does in IPv4 networks.