Advanced Unix System Administration

Lecture 12
April 2, 2008

Steven Luo
<sluo+decal@OCF.Berkeley.EDU>

IP: the Internet Protocol

- IPv6
  - IPv4 address exhaustion a looming issue (possibly as early as 2010)
    - Is already an issue in Japan and developing countries, which got less of the address allocation than existing developed countries
  - Complex CIDR network topology means that the IPv4 routing table is growing exponentially
  - IP needs to catch up with advances in routing technology and with the modern security environment

IP: the Internet Protocol

- IPv6 con't
  - Addressing: same idea as IPv4, but 128-bit addresses instead of 32-bit
    - Usual notation: 8 parts in hex separated by colons
  - By convention, high 64 bits of the IP used for the network ID, low 64 bits for the host ID
    - The (excessively) large host ID allows for stateless autoconfiguration via MAC or EUI-64 addresses
    - Router just announces prefix, clients automagically configure themselves
    - DHCPv6 available for more complex setups

IP: the Internet Protocol

- IPv6 con't
  - The original concept: hierarchical allocations and hierarchical routing
    - Regional Internet Registries get very large blocks (/16) from IANA
    - RIRs allocate /32 or larger blocks to ISPs
    - ISPs allocate /48 to organizations, /64 to known small subscribers
    - Prefixes to be allocated to facilitate prefix aggregation; would make the IPv6 routing table much smaller and easier to manage

IP: the Internet Protocol

- IPv6 con't
  - Failure of the hierarchical allocation model
    - Large organizations insisted on provider-independent (PI) address space
      - Renumbering is difficult, so PI addresses make switching ISPs easier
      - Also makes multihoming (connecting to more than one ISP) easier
    - Large endpoints can therefore now get /32 allocations on their own
    - With this allocation model, the IPv6 routing table may get to be just as large – or larger! – than the IPv4 routing table

IP: the Internet Protocol

- IPv6 con't
  - IPv4 and IPv6 can coexist, but v6 is not backwards-compatible
  - Transition mechanisms:
    - IPv6-in-IPv4 tunnels
    - Mapping the IPv4 space into the IPv6 space
    - Protocol translation mechanisms
  - Transition is going slowly, but Catch-22 involved: no one will deploy until commonly used, no one will use until commonly deployed

The Transport Layer

- User Datagram Protocol (UDP)
  - Extremely simple, but provides basically no features
  - Unreliable, message-oriented, stateless
  - Each message is packaged into a single datagram and sent over the network
  - No guarantees on delivery, order of arrival
  - UDP packet: source port (16 bits), destination port (16 bits), length (16 bits), checksum (16 bits), data

The Transport Layer

- Ports
  - In TCP and UDP, each socket is assigned a port number to identify the traffic to it
  - Kernel examines destination port to decide which process to give data to; source port used to determine destination of replies
  - By convention:
    - 1-1023 are “Well Known Ports” for services, 1024-49151 are “Registered Ports” for services,
    - 49152-65535 are dynamic ports used as source ports