Networking Basics 1: TCP/IP Introduction

1. **TCP/IP Introduction**
   - **TCP/IP Protocol Suite**
   - **IPv4 vs. IPv6**
   - **IPv4 addresses**
   - **IPv6 addresses**
2. Routing and Ports
3. TCP/IP Protocols
4. Ethernet
5. DNS

The TCP/IP Protocol Suite

The **network protocols** that are used with the Internet are commonly known as **TCP/IP**.

TCP/IP is named for its most important parts:

- the Transmission Control Protocol (TCP)
- the Internet Protocol (IP)

Alternative, formal name: **Internet Protocol Suite**.

"Internet" here is not *the Internet*, but rather short for "inter-connected network" or internetwork (since predates the (global) Internet).

Goal: to be able to connect networks that were potentially using different physical technologies.

Layered Protocol

The functionality that makes up networking protocols is typically conceptualized in terms of being partitioned into *abstraction layers*.

TCP/IP's **network model** has *4 layers*:

- **application** – user-level applications (e.g., FTP, SSH)
- **transport** – end-to-end data transfer protocols (e.g., TCP and UDP)
- **network** – abstract (non-physical) network protocol (e.g., IP and ARP)
- **link/network interface** – physical network protocol (e.g., Ethernet, Token-ring, ATM)

Layered Protocol (contd.)

The more common **OSI network model** uses 7 layers:

- **application** – user-level applications
- **presentation** – data mapping (format, encryption)
- **session** – inter-host connection control
- **transport** – data transfer protocols for applications
- **network** – inter-host data transfer and addressing
- **data** –link physical data transfer and addressing
- **physical** – media and transmission

IPv4 vs. IPv6

There are two main IP standards:

- **IPv4** (IP version 4)
- **IPv6** (IP version 6)

IPv4 is currently the most widely used standard, so we will concentrate on it.

IPv6 is the upcoming/eventual standard, but its use on the Internet remains limited.

IPv6 has many advantages, including:

- much larger addresses
- support for larger packets
- improved security and **multicasting**

IP Addresses

IP addresses are one of the most visible and important elements of TCP/IP.

An **IP address** is intended to *uniquely identify* a **host** on the Internet.

IPv4 addresses are 32-bit numbers.

This means there are only 2^{32} or less than 4.3×10^9 addresses.

The available IPv4 addresses have basically all been allocated as of 2014!

IPv6 addresses use 128 bits, so there are 2^{128} or 3.4×10^{38} possible addresses!!

IPv4 Addresses

IPv4 addresses are often written as the values in *decimal* of each of the *four bytes*, separated by **dots**:
e.g., 192.168.0.101 or 131.230.133.154

This format is referred to as "**dotted decimal**" or "**dotted quad**."

Each byte/quad/octet can range from 0 to 255.

An IPv4 address consists of two conceptual parts:

- **network number** – assigned to the network owner by an Internet authority (or ISP)
- **host number** – assigned to a machine by the network owner

IPv4 Addresses (contd.)

How many of the bits in the address are used for the network number determines how many hosts can be supported within the network.

Originally, the **address class** determined this:

- **class A** – first bit 0, network number is 7 bits, so 16,777,214 hosts (last 3 quads)
- **class B** – first bits 10, network number is 14 bits, so 65,534 hosts (last 2 quads)
- **class C** – first bits 110, network number is 21 bits, so 254 hosts (last quad)
- **class D** – first bits 1110, multicast
- **class E** – first bits 1111, reserved

IPv4 Addresses (contd.)

The class method of determining network vs. host number proved to be too inflexible.

Initially the notion of a **subnet** was introduced, by dividing the host number into two parts.

So an IP address consisted of three parts:

`<network number><subnet number><host number>`

In 1993, classes were abandoned completely, replaced by **classless inter-domain routing** (CIDR).

With CIDR the network-host division need not be on quad/byte boundaries.

IPv4 Addresses (contd.)

A CIDR block of addresses is specified using:

1. an IP address, and
2. the network prefix length (in bits)

Thus the equivalent of a class B address block would be specified like: 130.201.0.0/16

An address block where the network-host division is not on a quad/byte boundaries would be something like: 192.168.1.128/26

The **Internet Assigned Numbers Authority** (IANA) gives out IP ranges using CIDR specs.

IPv4 Addresses (contd.)

Consider the address block 192.168.1.128/26:

- 26 bits means the high order 3 quads/bytes plus the *high order 2 bits* in the last/lowest order quad/byte
- $128_{10} = 10000000_2$, so the two network bits are 10
- so all addresses in this block must have their high order 3 quads/bytes be 192.168.1 and then have the high order 2 bits in the last quad be 10: 192.168.1.(10xxxxxx)
- thus the address block range of IP addresses is:
  - 192.168.1.128 [which is 192.168.1.(10000000)] through
  - 192.168.1.191 [which is 192.168.1.(10111111)]

IPv4 Addresses (contd.)

The network vs. host parts of IP addresses are also often specified via a **subnet mask**:

- 32 bit number like an IP address, written similarly
- 1 bit means network part
- 0 bit means host part
- e.g., 255.255.255.192:
- $255_{10} = 11111111_2$ and $192_{10} = 11000000_2$
- so high order 26 bits are network part
- low order 6 bits are the host part
- thus subnet mask 255.255.255.192 is equivalent to a CIDR prefix of 26

IPv4 Addresses (contd.)

IP addresses are one of:

- **unicast** – intended for a single machine
- **broadcast** – intended for all (local) machines
- **multicast** – intended for multiple machines

Broadcast addresses have *all 1's in host portion*.

E.g., for 192.168.1.120/26 the local network broadcast address is 192.168.1.191 [192.168.1.(10111111)]

The *special broadcast address* 255.255.255.255 is interpreted as broadcast for the *local network*.

IPv4 Addresses (contd.)

An IP address with *all 0's in the host part* means *this machine*.

E.g., if machine is 192.168.1.129/26 then 192.168.1.128 also means this machine since $129 = 10000001_2$ and $128 = 10000000_2$.

Having *all 0's in the network part* means *this/local* network.

E.g., 0.0.0.1 is equivalent to 192.168.1.129 given the network defined above.

Special address 0.0.0.0 means this machine on this/local/zero network.

Multicast addresses start with 1110:
224.0.0.0 to 239.255.255.255

IPv4 Addresses (contd.)

Several ranges of IPv4 addresses are "special:"

- **loopback** – routed directly to the local machine
  - 127.0.0.0 to 127.255.255.255
  - 127.0.0.1 is the typical loopback/lo address
- **private** – not routable on Internet
  - 10.0.0.0 to 10.255.255.255
  - 172.16.0.0 to 172.31.255.255
  - 192.168.0.0 to 192.168.255.255
- **link-local** – not routable, automatically configured
  - 169.254.1.0 to 169.254.254.255
- **reserved** – currently unused
  - 240.0.0.0 to 255.255.255.254

IPv6 Addresses

As noted, IPv6 addresses use 128 bits, so the set of possible IPv6 addresses is vast.

Advantages over IPv4:

- will avoid **address exhaustion** forever
- large enough to allow every device to have an IP address ("**Internet of Things**")
- will allow pure hierarchical allocation of addresses, simplifying routing and changing ISPs
- **stateless address autoconfiguration** (SLAAC) possible (automatic address configuration for hosts)
- provides hosts both unchanging link-local address and global (routable) address at the same time

IPv6 Addresses (contd.)

128 bit address is written as 8 groups of 4 hex digits separated by colons:
e.g., 2001:0db8:0000:0000:5e72:287a:ccd4:00d1

Many groups/digits will be 0's, so a sequence of 0000 groups can be replaced by double colons (::) but only one :: per address.

Leading 0's in groups can be dropped:
e.g., 2001:0db8:0000:0000:5e72:287a:ccd4:00d1
or 2001:db8:0:0:5e72:287a:ccd4:d1
or 2001:db8::5e72:287a:ccd4:d1

IPv6 Addresses (contd.)

Typically first 64 bits used for network part and last 64 bits for host part.

Final 4 bytes can also be written in dotted quad:
e.g., xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:83e6:8552
or xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:131.230.133.82

Loopback address (localhost):
0000:0000:0000:0000:0000:0000:0000:0001
or 0:0:0:0:0:0:0:0:0:1
or just ::1

IPv4 Mapped Addresses (transition address):
::ffff:0:0/96 (96-bit prefix on IPv4 32-bit address)

IPv6 Addresses (contd.)

Unique local addresses (ULA's), a type of private (not globally routable) address: fc00::/7.

Multicast (set of hosts) addresses: ff00::/8

Multicast to all local hosts (i.e., broadcast): ff02::1

Can automatically generate 64-bit host part of address from a 48-bit NIC interface address (e.g., Ethernet MAC address):

- insert fffe between 3rd and 4th bytes
- complement 2nd lowest order bit of highest order byte
- e.g., 00:60:a8:52:f9:d8 becomes 0260:a8ff:fe52:f9d8