Advanced Unix System Administration

Lecture 11
October 20, 2008

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

Networking Intro

- The OSI model
  - Seven layers that conceptually separate the different functions of a network stack very cleanly
  - No practical modern network stacks actually implement the full separation model
- Physical layer (layer 1)
  - Specifies the actual communications hardware
  - Fiber, copper twisted pairs, wireless, SCSI . . .

Networking Intro

- Data link layer (layer 2)
  - Specifies details of over-the-wire communication and error correction between physical hosts
  - We'll focus on Ethernet in this class
- Network layer (layer 3)
  - Routes traffic from Point A to Point B, possibly with QoS considerations
  - IP is the only important example nowadays

Networking Intro

- Transport layer (layer 4)
  - Provides facilities for a link between hosts, such as flow control, error correction
  - TCP and UDP are the most important
- Session layer (layer 5)
  - Provides the link between hosts (connections, ports, state)
  - TCP performs the functions of the session layer; applications running on UDP must provide their own session facilities

Networking Intro

- Presentation layer (layer 6)
  - Data representation for the application
  - Almost always performed by the application nowadays; protocols are usually considered part of layer 6
  - We'll cover some important examples
- Application layer (layer 7)
  - Provides a useful service to users
  - The application implementing a protocol is in layer 7

Ethernet

- What is Ethernet?
  - Ken Thompson: “Robert Metcalfe says that if something comes along to replace Ethernet, it will be called 'Ethernet'”
  - The physical layer has varied enormously over the last 30 years: 10 Mbit coaxial cable, 100 Mbit twisted pair, 10 Gbit fiber-optic
  - Actual electrical (optical) standard has varied over the years

Ethernet

- Ethernet topology
  - Classic Ethernet is a bus topology – devices share one medium
    - Only one device can transmit at one time
    - Devices use CSMA/CD to handle access to the medium
    - All devices examine all transmissions to determine whether or not it's intended for them
    - Devices identified by MAC address
    - Doesn't scale well with large networks
  - Switched Ethernet only forwards transmissions to the wanted device

Ethernet

- The Ethernet frame
  - Each Ethernet transmission is packaged into a “frame”
  - Several different Ethernet frame types, which can coexist on the same network
  - “Ethernet II” is most common
    - 14-byte header containing source, destination MAC addresses, and an EtherType
    - Payload; 46-1500 bytes (hence the 1500 byte MTU)
    - 4 byte CRC checksum

Ethernet

- The EtherType
  - Supposedly was originally a length field
  - In Ethernet II, is a 16-bit value greater than 0x0600 used to indicate the network layer protocol contained in the payload
  - Common values:
    - 0x0800 – IPv4
    - 0x0806 – ARP
    - 0x86DD – IPv6

IP: the Internet Protocol

- Packet-switched data protocol for use on unreliable networks
- Responsible for identifying computers and getting packets from A to B
- Addressing
  - 32-bit addresses (4.2 billion possible unique addresses)
  - Usually written in four decimal quads of 8 bits each, for convenience
    - This form is frequently misleading!

IP: the Internet Protocol

- Address resolution
  - There has to be a way of mapping IP addresses to hardware devices
  - The Address Resolution Protocol (ARP) deals with this
    - Devices broadcast an ARP packet: either a request for the physical address of a particular IP address
    - The matching device broadcasts a reply
    - Devices may also broadcast “gratuitous” ARP announcements
    - Address mappings are (or should be) cached by devices

IP: the Internet Protocol

- Networking
  - The IP address is divided into a network portion (n high bits) and a host portion (remaining low bits)
  - Hosts only speak directly to other hosts with the same network ID
  - Network mask (bitmask ANDed with the IP to obtain network ID) is usually written in dotted-quad or in CIDR notation (/n, where n is the number of bits)