ICS 351: Today's plan

- SNMP
- Performance
- Peer-to-Peer Networking

Simple Network Management Protocol

- SNMP uses the network to report status information and alerts about remote systems
- SNMP messages are carried over UDP
- values can be loaded on demand (pull model), but when needed and configured appropriately, alerts are sent independently by the systems being managed (push)

SNMP
Management Information Base

- SNMP needs a machine-independent way to indicate which item of information is being requested or sent
- logically, the entire universe of information that can be accessed is built into a large tree: the Management Information Base or MIB
- the tree is extensible so individuals and organization can add their own subtrees -- private MIBs
- the tree is universal and known to all

navigating the MIB

- the path through the tree is sufficient to indicate one specific item (corresponding to a variable in a programming language)
- the path through the tree can be indicated by a sequence of numbers, the number of left siblings of the path being taken
- for example, 0.2.7.5.14.1.7.0 is such an Object Identifier (OID)
- OIDs are useful for enumerating arrays of objects, e.g., network interfaces, routing table entries

SNMP programs

- a network management station is used by the system administrator to monitor multiple systems
- a management agent must run on every managed device, get the required information, and provide it on request

SNMP basic operation

- the network management station may send GET requests to get one or more objects from specifc agents
- the network management station may also send SET requests to modify one or more objects on specifc agents
- agents will send TRAP or INFORM alerts to network management stations that they have been configured to alert
- because it uses UDP, SNMP (like DNS) cannot assume that its operations will be successful

SNMP examples

- sample MIB:

  http://www.net-snmp.org/docs/mibs/NET-SNMP-EXAMPLES-MIB.txt
- MIB table retrieval example: http://etutorials.org/Networking/network+management/Part+II+Implementations+on+the+Cisco+Devices/Chapter+4.+SNMP+and+MIBs/MIB+Table+Retrieval+Example/

Performance

- ultimately, performance is given by how long it takes a user to complete an action
- performance is measured as both latency and throughput
  - latency is the time to send 1 bit (or one packet)
  - throughput is how many bits can be sent in a given time
- *n* round-trips sending *m* bits might take time

  *n**latency + (*m* / throughput)
- in practice, it gets more complicated:
  - bits might be sent in both directions
  - operations don't always take the same amount of time
  - both throughput and latency vary

Latency

- latency is measured in seconds:
  - one-way (useful, but hard to measure)
  - round-trip (also useful, and easy to measure)
- usually measured by sending the smallest possible transmission unit
- latency matters when there are lots of round-trips
  - web transfers (less so with HTTP/2)
  - remote login
- latency matters when things must happen in real time
  - telephony, gaming, trading, remote control, telemedicine

Throughput

- throughuput is measured in bits/second or Bytes/second
  - network equipment generally rated in bits/second
  - network downloads generally list Bytes/second
  - in this class and in much (but not all) of the world:
    - b/s stands for bits/second
    - B/s stands for Bytes/second
- throughput matters when lots of data must be sent
- `ttcp`, `netperf`, and `iperf` measure and report throughput

Considerations

- at each level in the network, it is usually cheaper to provide less performance than more
  - i.e. more profitable to serve more people given the same infrastructure
- HTTP headers consume multiple hundred bytes, as opposed to about 40-60 bytes for TCP/IP headers
- but until HTTP/2, nobody was concerned enough to try to compress them
  - so the inefficiency must be negligible compared to the available throughput

A more advance performance test

- send two successive pings as fast as possible
- at the slowest link in the path, the second will be delayed in proportion to the size of the first
- this allows us to estimate the speed of the *bottleneck link*
- `bwping` performs this test

Peer-to-peer Networking

- in this class, it has been clear that a PC can perform as a router, and vice-versa
- why not use our PCs as routers?
- for many purposes this is adequate

Content-Addressable Networking

- if we are not restricted to forwarding to IP addresses, there is less need for hierarchical control
- in fact, we can have content-addressable networks
- perhaps, to increase the incentive, every node that has content and makes it available might get priority for obtaining other content
- the decentralization and lack of control are one of the attractions for many of the users

Managing Peer-to-peer Networks

- management of a network requires some authorities that cooperate
- in a peer-to-peer network, both the management and the authority are made as small as possible
- because there is no hierarchical assignment of addresses, each peer can decide what content to provide, i.e. what "address" to use
- this minimizes management

Connecting Peer-to-peer Networks

- the network is connected by having one peer exchange addresses about other peers
- then, if the original peer stops collaborating, the other peers can be used to connect to the network
- the expectation is that some useful data exchange will take place, not necessarily that there will be continuous end-to-end connectivity with all parties involved

Examples of Peer-to-Peer

- BitTorrent
- Distributed Hash Tables
- Freenet
- AllNet
  - maybe: Skype

Peer-to-peer vs. Client-Server

- a server is a program that provides a service
- typically, a server is found at a given address and port number, and is maintained by an individual or an organization
- clients may be anonymous or unidentified
- peer-to-peer (P2P) often brings to the infrastructure the anonymity of clients
- can imagine authenticated P2P networks, set up for reliability rather than anonymity

Peer-to-peer vs. Client-Server II

- for a very different example, routing is a peer-to-peer process
  - routers identified by IP address or ID
- the ultimate appeal of peer-to-peer for some networking people is a self-organizing, self-managing scalable network
- peer-to-peer networks seem to be widely accepted