EEC173B/ECS152C:

Tracking End-to-End Routes & Performance

Two Basic Monitoring/Measurement Techniques

- Active -- inject measurement probes
  - Ping: send ICMP packets
  - Traceroute: send UDP packets with increasing TTLs and collect path stats
  - Probing: use specific UDP or TCP packets from one host to the other (e.g. RTCP)
- Passive -- non-intrusive observations of the network
  - SNMP: collect aggregate statistics from routers (thru MIBs).
  - Netflow: collect aggregate flow stats from (Cisco) routers
  - "tap the link": OCxMON, packet sniffers, tcpdump

2

Time-To-Live (TTL) Field

- A value found in the header of an IP datagram
  - Max possible value: 255
  - TCP/IP spec stats that TTL should be 60, but many systems use smaller values (e.g., 30 or 15)
- Each router (intermediate nodes) decrements the TTL value
- If the TTL reaches zero, the router discards the datagram and sends an ICMP error message to the source
- Prevents a datagram from following a cycle of routes forever

3

Internet Control Message Protocol (ICMP)

- An error reporting mechanism
- Uses an 8-bit "type" field (0-255)
- Example types:
  - 0: Echo Reply
  - 3: Destination unreachable (Router cannot locate destination)
  - 8: Echo
  - 11: Time exceeded (packet discarded because TTL = 0)
  - 37-255: Reserved

4

Ping

- Using ICMP to testing reachability
  - Ping sends an IP datagram that contains an ICMP echo request message to a specified destination.
  - After sending the request, it waits a short time for the reply. If no reply arrives, it retransmits the request.
  - ICMP software on the remote machine replies to the echo request message.
  - If no reply arrives from retransmissions -> declares that the machine is not reachable
  - Round-trip time and packet loss statistics are computed

5

Why Round-Trip Time (RTT)?

- One-way delay measurements between A and B
  - Delay = *T_req_received_at_B – T_req_sent_by_A*
  - Need control over source/destination machines
  - Need clock synchronization between sender and receiver
- Round-trip delay between A and B
  - Delay = *T_replyfrB_received_at_A – T_req_sent_by_A*
  - No clock synchronization required

6

Time Synchronization

- NTP: The Network Time Protocol
  - http://www.ntp.org/
  - Protocol designed to synchronize the clocks of computers over a network
  - Provides accuracies typically within a millisecond on LANs and up to a few tens of milliseconds on WANs

7

Traceroute

- Tracing route to a destination
  - Exploit Time-to-live field
  - Send a series of datagrams to a given destination
    - Set TTL value for the first datagram to 1
    - First router that receives the datagram decrements the TTL, discard the packet (because TTL=0), and sends back an ICMP time exceeded message.
    - Traceroute extract the IP source address from the ICMP packet and announce the address of the first router along the path to the destination
    - After it discovers the address of the 1st router, it sends a datagram with TTL=2, then to 3, 4, and so on until it hits the destination

8

The Last Address Printed by Traceroute

- What happens when TTL is large enough for the datagram to reach its destination?
- Two implementations to make sure destination responds to the datagram
  - Microsoft tracert: send an ICMP echo request message; the destination host will generate an ICMP echo reply
    - Source address = IP address to which the request was sent
  - Unix version of traceroute: send UDP message to a nonexistent program on the destination machine. The ultimate destination will send a "ICMP destination unreachable message".
    - Source address = IP address of the actual interface of the router over which datagram arrived

9