Coping with Overload on the Network Time Protocol Public Servers

David Mills, U Delaware
Judah Levine, NIST
Richard Schmidt, USNO
David Plonka, U Wisconsin

From NBS Special Publication 432 (1979 edition, now out of print)

Introduction

- The Network Time Protocol (NTP) synchronizes clocks of hosts and routers in the Internet.
- NTP provides nominal accuracies in the low tens of milliseconds on WANs, submilliseconds on LANs, and submicroseconds with appropriate hardware.
- NTP software has been ported to almost every workstation and server platform available today - from PCs to Crays - Unix, Windows, VMS and embedded systems, even home routers and battery backup systems.
- The NTP architecture, protocol and algorithms have been evolved over the last two decades to the latest NTP Version 4 software distribution.

The Sun never sets on NTP

- NTP is argueably the longest running, continuously operating, ubiquitously available protocol in the Internet
  - USNO and NIST operate a total of about three dozen Internet primary servers directly synchronized to national standard cesium clock ensembles and GPS
  - Over 200 Internet primary servers are in other countries, including Australia, Canada, Chile, France, Germany, Isreal, Italy, Holland, Japan, Norway, Sweden, Switzerland, UK, and US.
- NIST estimates 10-20 million NTP servers and clients deployed in the Internet and its tributaries all over the world.
  - Agencies and organizations: US Weather Service, US Treasury Service, IRS, PBS, Merrill Lynch, Citicorp, GTE, Sun, HP, etc.
  - Private networks are reported to have over 10,000 NTP servers and clients behind firewalls; one (GTE) reports in the order of 30,000 NTP workstations and PCs.
  - NTP has been on the NASA Shuttle and in Antarctica and planned for the Mars Internet.

On the hazards of serving time

- With potential client populations in the millions, there is a very real vulnerability to grossly overload the public primary server population.
- The public NTP client software exchanges packets with the server on a countinuous basis in order to discipline the computer clock time and frequency.
    - This software has been carefully designed to be a good network citizen and ordinarily does not exceed a rate of one packet every fifteen minutes.
- Defective NTP client implementations have appeared that exhibit gross violations of the Internet social contract.
    - An example is the U Wisconsin incident reported in the next slide.
- The sheer weight of numbers threatens to overwhelm at least some of the current NIST and USNO servers.
    - Other incidents reveal really bad network engineering and counterproductive parameter selection, especially poll interval.

The U Wisconsin incident

- U Wisconsin operates a number of time servers for campus access.
- A home router came on the market that
  - had the address of one of these servers hard-coded in firmware and could not be changed,
  - could send packets continuously at one-second intervals under certain conditions when service was interrupted.
- This would not be a problem if only a small numbers of these routers were sold.
  - However, eventually 750,000 routers were sold and most could not be recalled, updated or even reliably found.
  - The resulting traffic overwhelmed the server, university network and service provider.
- There has been no wholly satisfactory solution to this problem other than to insure continuous service and to educate the manufacturer about socially responsible product design.

Conditions at USNO

- USNO operates about 20 NTP servers in the US, Alaska and Hawaii
  - Three of the busiest servers are in Washington, DC.
  - They share an aggregate load of 3,000-7,000 packets per second (PPS).
  - While these three servers are at only 10 percent capacity, the 10-Mbps network is badly overloaded, leading to significant packet loss and badly degraded time quality.
- Much of the traffic is from clients sending at unrealistic rates.
  - In one case the client is spraying at 14 PPS, a rate equivalent to 731 properly configured NTP clients.
  - In another case a university firewall has channeled 2,000 campus clients separately to the USNO servers when it should synchronize to USNO and have all campus clients synchronize to it, as NTP is designed to do.

Conditions at NIST

- NIST operates about a dozen NTP public time servers in the US.
  - Three of the busiest servers are in Boulder, CO
  - They share an aggregate load similar to USNO, but the NIST network infrastructure is far more resilient than USNO.
- An experiment collected statistics in a nine-second window on each machine using a sampling technique which captured about 13 percent of the arrivals.
  - The results revealed over 500 clients with polling intervals of 5 seconds or less and 15 with poll intervals less than one second. Well behaved NTP clients send at rates usually at intervals of fifteen minutes or more.
  - Most incidences involve packet bursts lasting from a few seconds to multiple days and separated by minutes, hours or days.
  - The most bizarre observation is the lengths of the bursts; 20 percent last over one minute and a few over two days, presumably continuously.
  - One particularly offensive elephant is sending continuously at two packets per second.

Things to do about it

- Some things are obvious
  - Rig the host name/address translation (DNS) to lie, cheat and steal; that is, randomize the addresses over a geographically dispersed server population (e.g., NTP pool scheme).
  - Find ways to deflect traffic from congested servers (e.g., time.nist.gov) to less busy servers closer to users (BGP Anycast).
  - Never ever carve an unconfigurable server address in the firmware.
- Educate potential implementors about best and worst practices in the selection of protocol parameters, especially poll interval.
  - Produce updated protocol specification including clearly defined best practices which minimize network and server impact (RFC submitted!).
  - Boil violators of best practices in oil. Smelly, stinky, public-ridicule oil.
- More aggressive pro-active nastiness may be necessary.
  - The Kiss-o'-Death packet is designed to disable clients as necessary.
  - The Call-Gap scheme finds the elephants and shoots them.