NTP (Network Time Protocol)
Why it important for network measurement

Jeff Boote (boote@internet2.edu)

Jason Zurawski (zurawski@internet2.edu)

Network Performance Workshop

05-Apr-2006

NTP

- Used to synchronize a group of servers to UTC
  - Attempts to keep time monotonically increasing while minimizing offset and skew
  - These goals contradict
    - Stability vs Accuracy
- RFC 1305

NTP

- Caveats:
  - Does not work very well for synchronizing to some (common) time
  - Must be the correct time
  - NTP does not deal with time zones (that’s a time display problem)

Utility for Measurement

- Scheduling requires coarse grain agreement on time (lets start/end together)
  - Agreement must be “global” in scope - UTC
    - Individual servers communicate with multiple other hosts
- One-Way latency requirements
  - Jitter (requires stability of offset within sample)
  - Latency (requires accuracy)
- Sensible compromise
  - Well defined error representation

NTP Basics

- Each NTP node has a stratum
  - Stratum is an integer between 0 and 16, inclusively
  - Stratum 0 means a physical clock, never a computer
    - Cesium oscillator: definition of time (subject to relativistic effects)
    - Rubidium oscillator: found in cell towers, very stable
    - GPS receiver: accuracy circa 10 ns
    - CDMA receiver: accuracy circa 10 μs

NTP Basics

- Each NTP node has a stratum (cont.)
  - Stratum 16 is reserved for devices that are not synchronized
  - The stratum of any NTP-synchronized device is the stratum of the device it is synchronized to plus 1, e.g.:
    - GPS receiver: stratum 0
    - Computer connected to it by a serial line: stratum 1
    - Client that gets the time from that computer: stratum 2

NTP Basics

- NTP servers form a loosely coupled network
  - Each node decides which server to use for synchronization based on complex selection algorithm (voting-like)
    - –Selection algorithm run repeatedly for protection against falsetickers
  - Sanity checks
  - Error estimates
  - Resiliency to clock failure
  - Resiliency to network failure
  - Enforces the global nature of NTP

NTP Limitations

- Needs external servers to work well, even with a local clock
- Can produce systematic errors with asymmetric paths
- Can have problems with asymmetric congestion
- (And no time zones, remember?)

NTP and Temperature

- Oscillator frequency depends on temperature
  - Typical correspondence: 1PPM (part per million) of clock rate for 1 c
  - NTP can resolve rate differences of .001PPM
- For comparison: the temperature inside a modern computer will vary by 10 c depending on CPU load
- NTP could notice human movement around the host (!), and certainly open windows or A/C failures
- NTP will compensate, but best servers sit in constant-temperature machine rooms

Sensible NTPD policy

- Open firewalls and open querying
  - Let others know your notion of time
- Good error determination requires 4 clocks (4 peers)
- Resilient setups will attempt to have the paths to all peers be as divergent as possible
  - Asymmetric paths to peers will cause offsets

Sensible NTPD policy

- Only servers of stratum n − 1 (where n is my stratum) fully matter
- Important to select servers of the same strata in all cases
  - Use all stratum 1 peers for stratum 1 configurations (unless you can have multiple physical time sources)
  - Use all stratum 1 servers for stratum 2 configurations
  - Use all stratum 2 servers for stratum 3 configurations
  - But don't do stratum 3 configurations for measurement
- Do not mix strata

Sensible NTPD deployment

- Stratum 1 if you need maximum possible accuracy (single microseconds)
- Stratum 2 is suitable for machines that serve time to others or are used in measurements (accuracy can be better than 1ms)
- Stratum 3 is a reasonable end-user stratum
- Most measurement nodes will probably use NTP stratum 2
  - Best accuracy bang for the buck
  - OK accuracy for most needs
  - Can be set up on most machines with no new hardware

Sensible NTPD deployment

- Hardware/network requirements are minimal
- Minimize temperature variations to minimize clock wander
- Use “real” NTP - not OpenNTP or some MS thing
    - Linux/FreeBSD release distributions should be fine
    - BUT YOU MUST CONFIGURE IT (Redhat default config is BAD!)

Additional Resources

- http://www.internet2.edu/%7Eshalunov/talks/20050322-Atlanta-PerformanceWorkshop-NTP.pdf
- Man ntpd
- Man ntp.conf
- /usr/share/doc/ntp
- http://e2epi.internet2.edu/owamp/details.html#NTP
- Twiki.ntp.org/bin/view/Support/SelectingOffsiteNTPServers

Internet2 Sample NTPD Config

- http://e2epi.internet2.edu/owamp/ntp/ntp.conf

Verify NTPD

- Ntpq
- Ntpd statistics (log files)

Demo of clock querying

```
$ ntpq -p
     remote           refid      st t when poll reach   delay   offset  jitter
==============================================================================
+GPS_PALISADE(0)              0 l    5   16  377    0.000    2.377   0.800
oPPS(0)          .PPS.        0 l    6   16  377    0.000    2.466   0.854
-nms4-chin.abile .PPS.        1 u   47   64  377    4.151    3.485   6.883
-nms3-atla.abile .PPS.        1 u   54   64  377   11.272    3.878   6.571
-nms4-kscy.abile .PPS.        1 u   56   64  377    9.510    1.305   4.192
+caspak.cerias.p .GPS.        1 u   49   64  337    6.825    2.463   5.926
```

www.internet2.edu