The NIST Internet Time Service

Michael Lombardi
NIST Time and Frequency Division
lombardi@nist.gov

Introduction

NIST maintains a Coordinated Universal Time scale, called UTC(NIST), that it distributes through a variety of free broadcast services that synchronize many millions of clocks every day. The largest service in terms of numbers of users is the Internet Time Service (ITS).

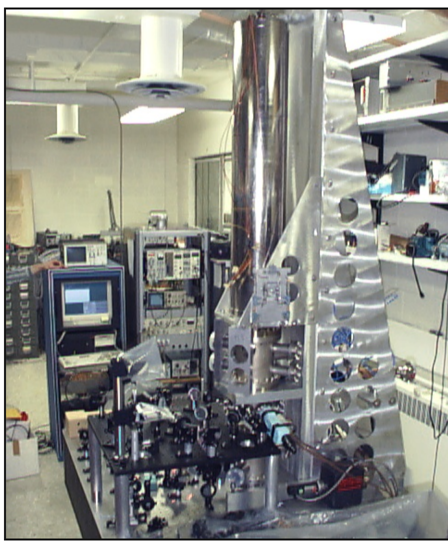

UTC(NIST) Time Scale

4 Cesium Beam standards

6 Hydrogen Masers

UTC(NIST)

Measurement System

Two-way satellite time & frequency transfer

GPS

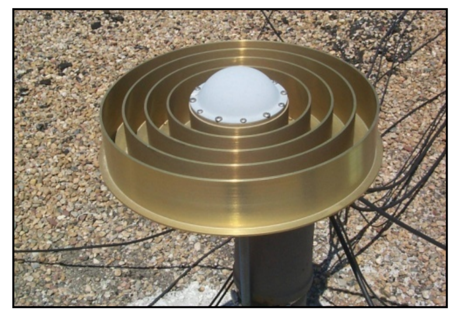

UTC(NIST) is continuously compared to other national time standards and the NIST clocks contribute to Coordinate Universal Time or UTC.

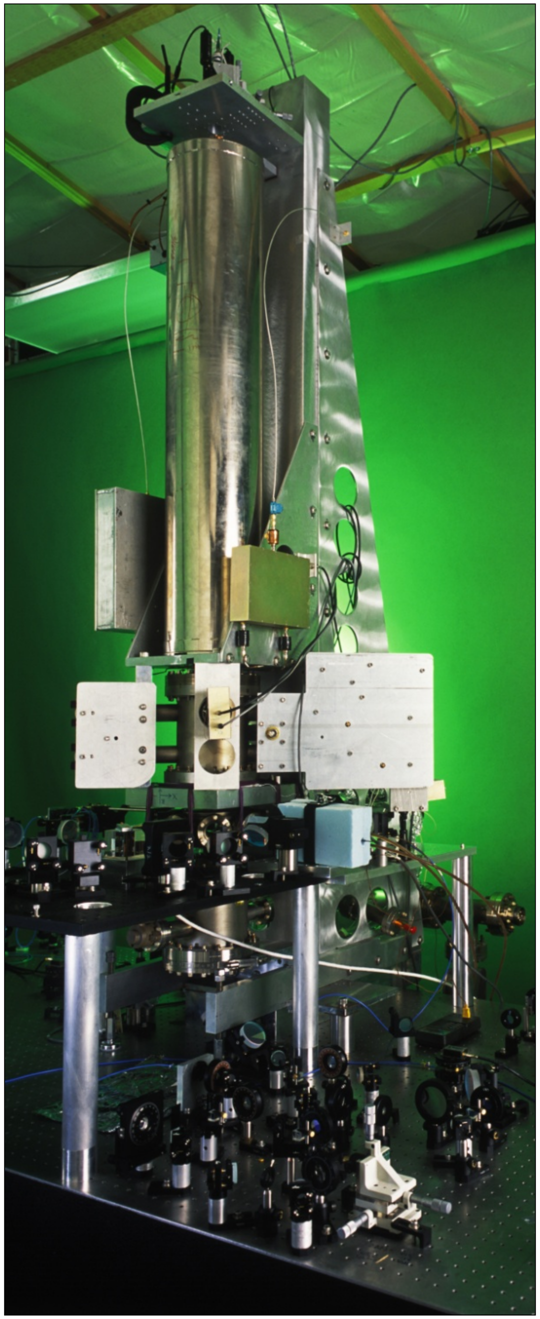

NIST-F1 Cesium Fountain Clock

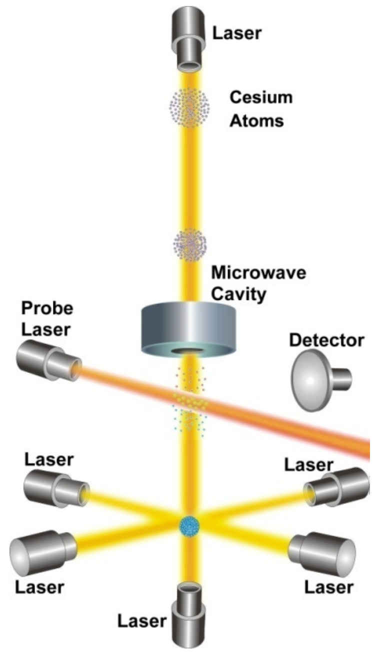

NIST-F1 laser-cooled fountain standard "atomic clock"

A cesium fountain frequency standard is used to calibrate the UTC(NIST) time scale. NIST-F1 produces the world's best realization of the definition of the second, which is the duration of 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the ^{133}Cs atom.

Current accuracy (uncertainty):

- 3×10^{-16}
- 26 trillionths of a second per day.
- 1 second in 105 million years.

Equivalent to measuring distance from earth to sun (1.5×10^{11} m or 93 million miles) to uncertainty of about 45 μm (less than thickness of human hair).

Computer clocks are more like wristwatches than atomic clocks

- Computer clock resolution is typically defined by the interrupt rate. This value limits the resolution of a time stamp and determines how accurately a clock can be set. Some common values:
  - 18.2 Hz (55 ms resolution, DOS computers)
  - 100 Hz (10 ms resolution, typical for Windows PCs)
  - 1024 Hz (1 ms, typical for Unix computers)
- The performance of the clock depends on the computer's internal timekeeping oscillator. This is typically a 32768 Hz (2^{15} Hz) quartz crystal oscillator, the same type of oscillator used in a wristwatch.
- The frequency of the quartz crystal is sensitive to temperature changes. They keep better time on a wrist, where the temperature is stable, than they do inside a computer.
- A good crystal inside a PC might keep time to within about 2 seconds per day (20 ppm). A bad crystal could be much worse (minutes per day).

Internet Time Protocols

- Time Protocol
  - 32-bit time code with 1 second resolution
- Daytime Protocol
  - no specified standard
  - NIST version: JJJJ YR-MO-DA HH:MM:SS TT L H msADV UTC(NIST) OTM
- Network Time Protocol (NTP)
  - 64-bit time code with 233 ps resolution
- About 95% of the requests to the NIST ITS are NTP requests, with the remaining requests divided about evenly between the Time and Daytime Protocols.

Network Time Protocol (NTP)

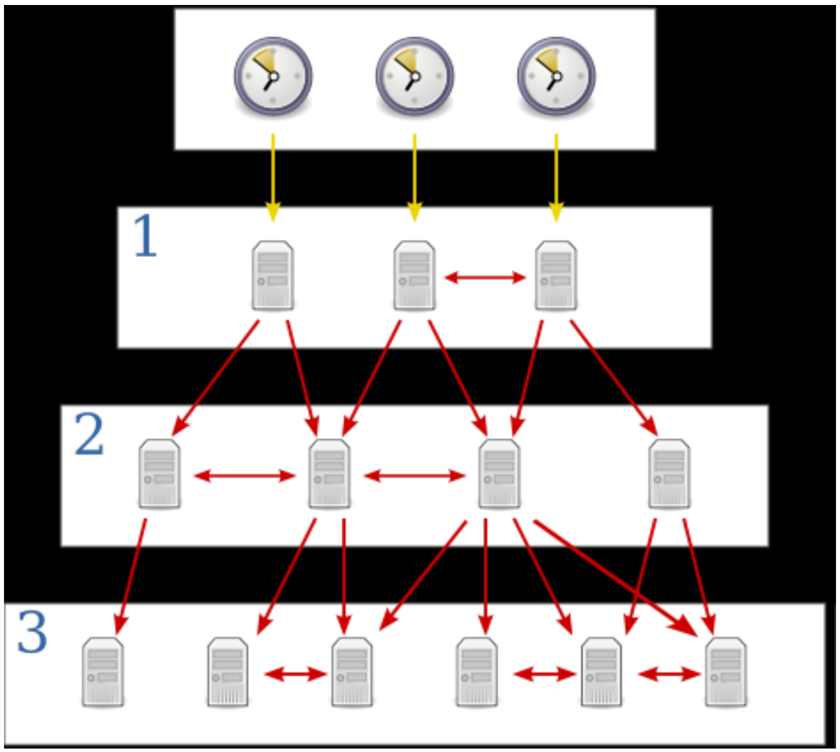

- The most widely used mechanism for time distribution via the Internet, defined by the RFC-1305 standard.
- Invented by David Mills of the University of Delaware with first implementation appearing before 1985.
- A hierarchical system consisting of levels of clock sources. Each level of this hierarchy is termed a *stratum.* Synchronization flows from primary servers at the lowest stratum to secondary servers at progressively higher stratums.
- Servers referenced to national time standards are called Stratum 1. There are about 40 to 50 Stratum-1 NTP servers in the U. S. that are open access, including the NIST servers (see ntp.org for a current list).

How NTP transmits time information

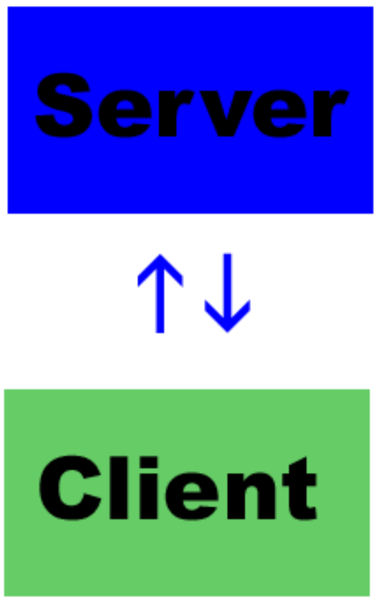

- A client can also be a server to computers at higher stratums.
- Client makes requests of a server at lower stratum via TCP/IP port 123
- A 48-bit packet is sent
- This packet contains four 64-bit timestamps:
  - Reference timestamp – local time at which system clock was last set or corrected
  - Originate timestamp – time when request left client for server
  - Receive timestamp – time when request arrived at server
  - Transmit timestamp – time when request left server for client

The NTP Packet

Packet header

Variables	Description
leap	leap indicator (LI)
version	version number (VN)
mode	protocol mode
stratum	stratum
τ	poll interval (16 s to 36 hours)
ρ	clock resolution
Δ	root delay, round trip delay to ref clock
E	root dispersion, max error
refid	reference ID – source of time
reftime	reference time stamp
T_1	originate time stamp
T_2	receive time stamp
T_3	transmit time stamp
MAC	MD5 message hash (optional)

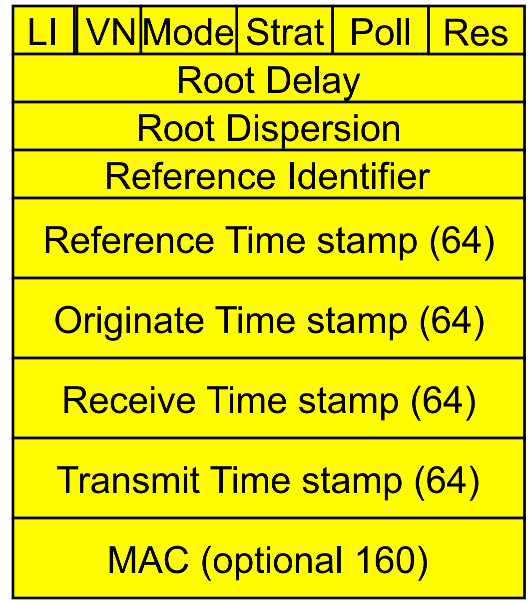

NIST time stamp format

The NTP time stamp is a 64 bit binary value with an implied fraction point between the two 32 bit halves. For example:

00000000000000000000000000000001 10000000000000000000000000000000

equals a decimal 1.5. The multipliers to the right of the point are 1/2, 1/4, 1/8, 1/16, and so on.

The 233 picosecond resolution is because:

$1 / 2^{32} = 0.00000000023283064365386962890625 = 233 \times 10^{-12}$ seconds

One picosecond is 1×10^{-12} seconds.

The epoch for NTP starts in the year 1900 while the epoch in UNIX starts in 1970. Therefore, the integer seconds will roll over in 2036. Thus, future versions of NTP might use 128 bit time stamps.

Example of NTP Client/Server Interaction

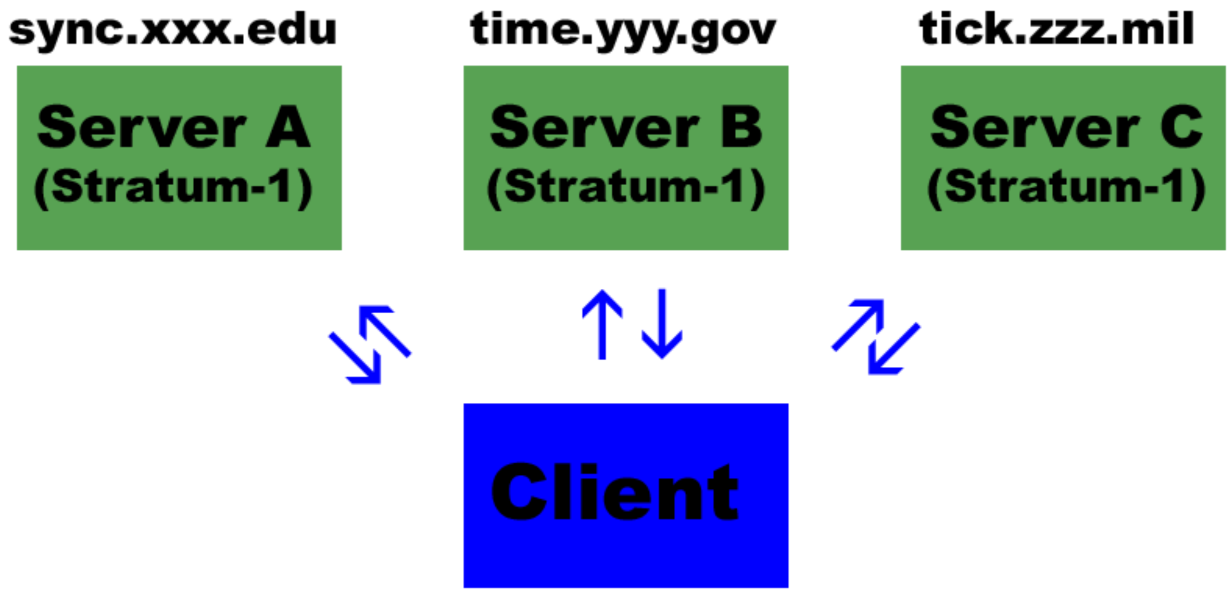

- Client can request time from several servers
- Client intercompares server times, network delays, etc.
- Looks for inconsistencies, throws out bad data
- Can repeat process to test network stability
- Uses data obtained to steer client clock

Clock Steering

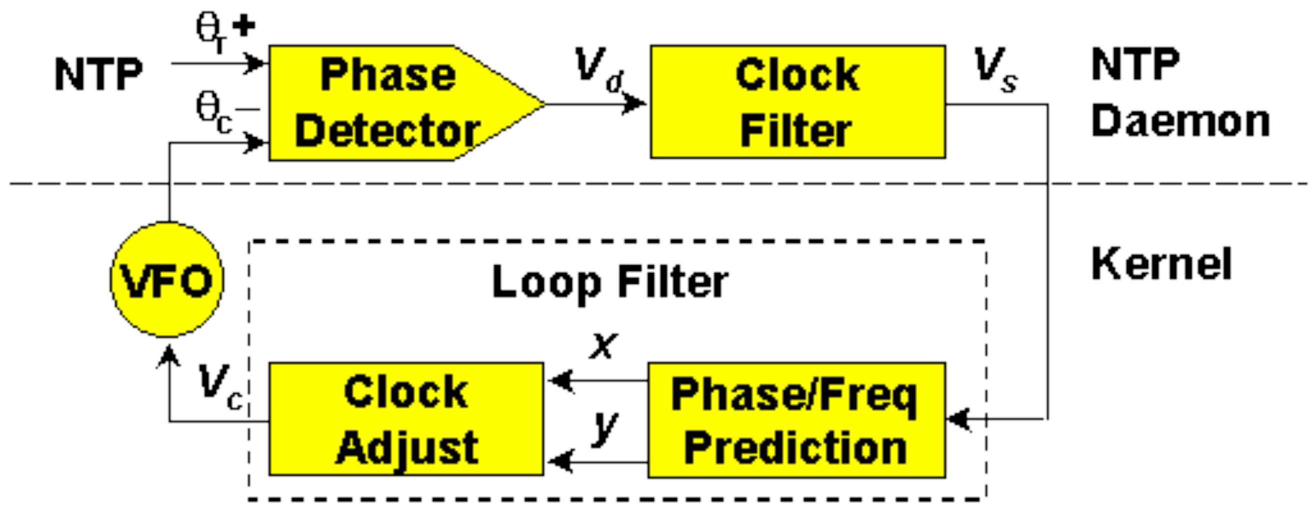

- Client software steers/corrects computer clock in between NTP requests
- Old versions of NTP used a phase locked loop (PLL) to minimize the time error
- Beginning with NTPv4, a combination PLL/FLL (frequency locked loop) algorithm is used to minimize time error. This algorithm improves performance over previous versions. Generally speaking
  - A PLL works best when system jitter dominates
  - A FLL works best when oscillator wander dominates

How accurate is NTP?

- Stratum-1 servers can keep time to better than 1 millisecond.

- The ability to synchronize a clock is limited by the clock resolution discussed earlier (typically 10 milliseconds for Windows computers).

- The ability to keep time between NTP requests depends on the client hardware and the software implementation. The Windows Time Service (XP, Windows 7, etc.) only requests time (by default) from an NTP server once per week, and can keep time to about 1 to 2 seconds per week through clock steering.

- Better clock steering or more frequent timing requests allow some NTP clients to keep time to within a small fraction of a second at all times. In some cases, the maximum time error will be nearly the same as the clock resolution – for example, 10 milliseconds over the public Internet.

- For a detailed discussion of how NTP works and a discussion of how to model and analyze timekeeping errors, see:

  David L. Mills, "Computer Network Time Synchronization: The Network Time Protocol," CRC Press, 2006.

Internet Time Service (ITS)

- One of the world's most popular time distribution services. The ITS handles about 4 billion timing requests per day. The exact number of computer clocks synchronized daily is unknown, but is certainly more than 100 million.
- 24 servers located around the United States.
- Client software is built into common operating systems: Windows, Mac, Unix.

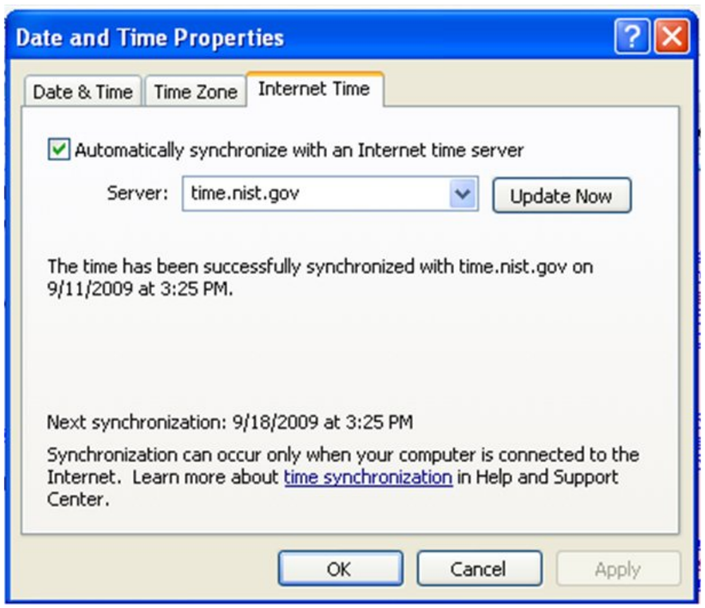

The Servers

- Dell Servers (several models in use, including 1850, 1950, 2950).
- The operating system is FreeBSD. The operating system kernel was modified at NIST to improve the timekeeping performance.
- All servers have external telephone modems to get time from NIST.
- Some sites have multiple servers on a load balancer.
- Some sites have small rubidium atomic oscillators to improve the short term stability of the server clocks.
- The time server software was written by Judah Levine at NIST, and can handle more traffic than most commercially available NTP servers. The busiest servers are easily handling about 10,000 NTP requests per second.

Server Sites and IP Addresses

- A list of the current ITS servers, including locations, IP addresses, and current status, can be found here:

  http://tf.nist.gov/tf-cgi/servers.cgi

- One third of the servers (eight) are located at NIST facilities (five in Boulder, one in Fort Collins, and two at the NIST headquarters in Gaithersburg, Maryland)
- The remaining 16 servers are located in 10 different states.
- One server is located at Microsoft in Redmond, Washington.

NIST ITS Server Locations

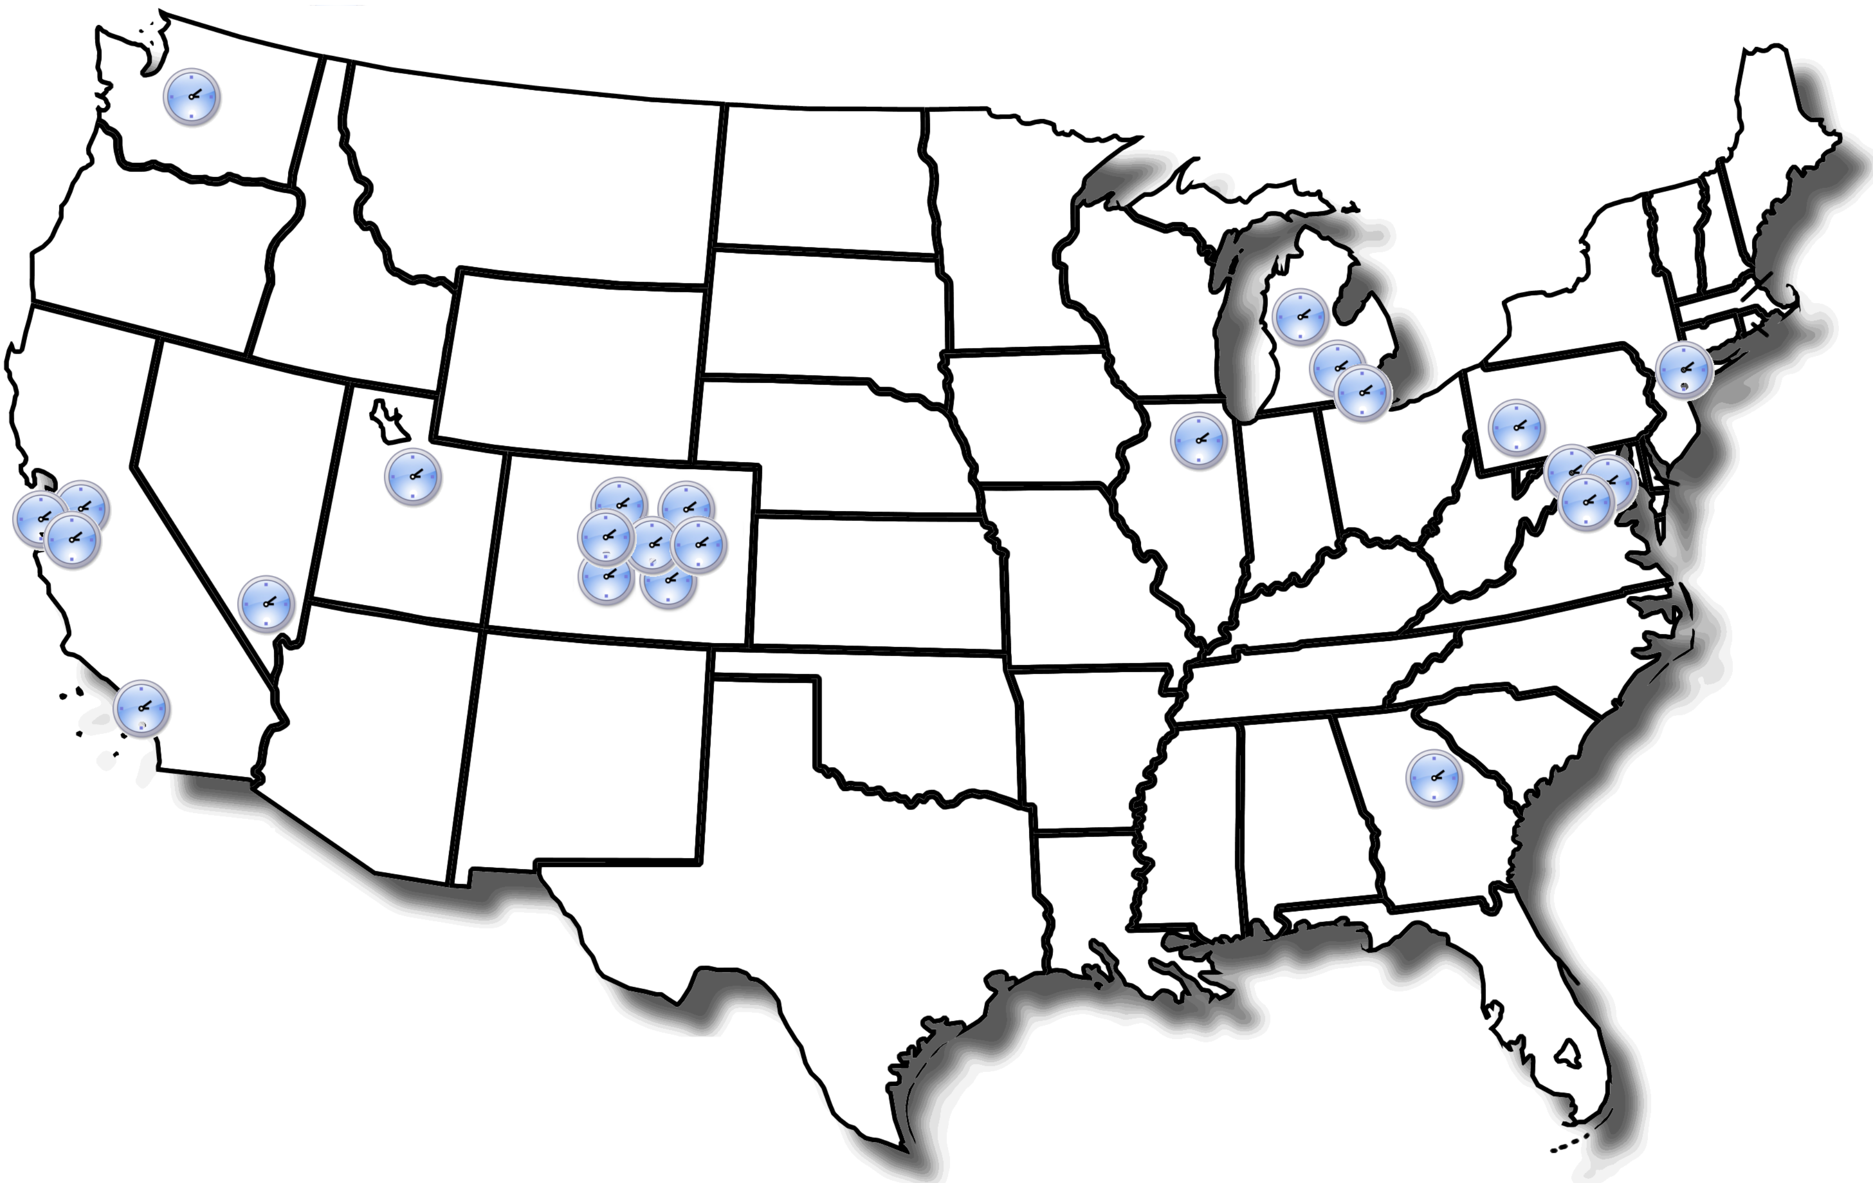

How the NIST Servers are Synchronized to UTC(NIST)

- Each ITS server makes a periodic phone call (about one call per hour) to servers in Boulder that are directly connected to the UTC(NIST) time scale. The call is made is over regular telephone lines using analog modems.

- The Boulder servers send an on-time marker (OTM) to the Internet time server. This marker is then sent back to Boulder and the round trip delay is measured. This is done several times to get a good estimate of the delay. Another OTM is then sent from Boulder that has been advanced by half the amount of the round trip delay.

- The OTMs sent from Boulder deliver UTC(NIST) to the servers. They OTMs provide the input to an algorithm called LOCKCLOCK that

  - Models the performance of the server clock

  - Steers the server clock to agree with UTC(NIST) using a frequency locked loop (FLL)

  - Makes periodic time corrections to the server clock based on a statistical evaluation of the clock's oscillator frequency

Telephone Time Code

```
52985 03-12-12 20:42:02 00 0 -.4 051.7 UTC(NIST) *
52985 03-12-12 20:42:03 00 0 -.4 051.7 UTC(NIST) *
52985 03-12-12 20:42:04 00 0 -.4 051.7 UTC(NIST) #
52985 03-12-12 20:42:05 00 0 -.4 051.7 UTC(NIST) #
52985 03-12-12 20:42:06 00 0 -.4 051.6 UTC(NIST) #
52985 03-12-12 20:42:07 00 0 -.4 051.8 UTC(NIST) #
52985 03-12-12 20:42:08 00 0 -.4 051.9 UTC(NIST) #
52985 03-12-12 20:42:09 00 0 -.4 051.8 UTC(NIST) #
52985 03-12-12 20:42:10 00 0 -.4 051.8 UTC(NIST) #
52985 03-12-12 20:42:11 00 0 -.4 051.8 UTC(NIST) #
```

JJJJ YY-MM-DD HH:MM:SS TT L DUT1 msADV UTC(NIST) OTM

Bandwidth Usage

- Bandwidth requirements are relatively small when considered in terms of the scale of Internet. Traffic is distributed across 24 servers, although much more than 20% of the traffic is handled by the servers in Boulder, Colorado. The amount of bandwidth used is unknown, but consider a very rough estimate:
  - 4×10^9 requests per day
  - Assume 100 bytes per request, with overhead, then
    $\sim 3.2 \times 10^{12}$ bits per day, or 3.2 billion kilobits
  - A rough estimate of bandwidth is 37,000 kbits/second, or about 1,500 kbits/second per average server

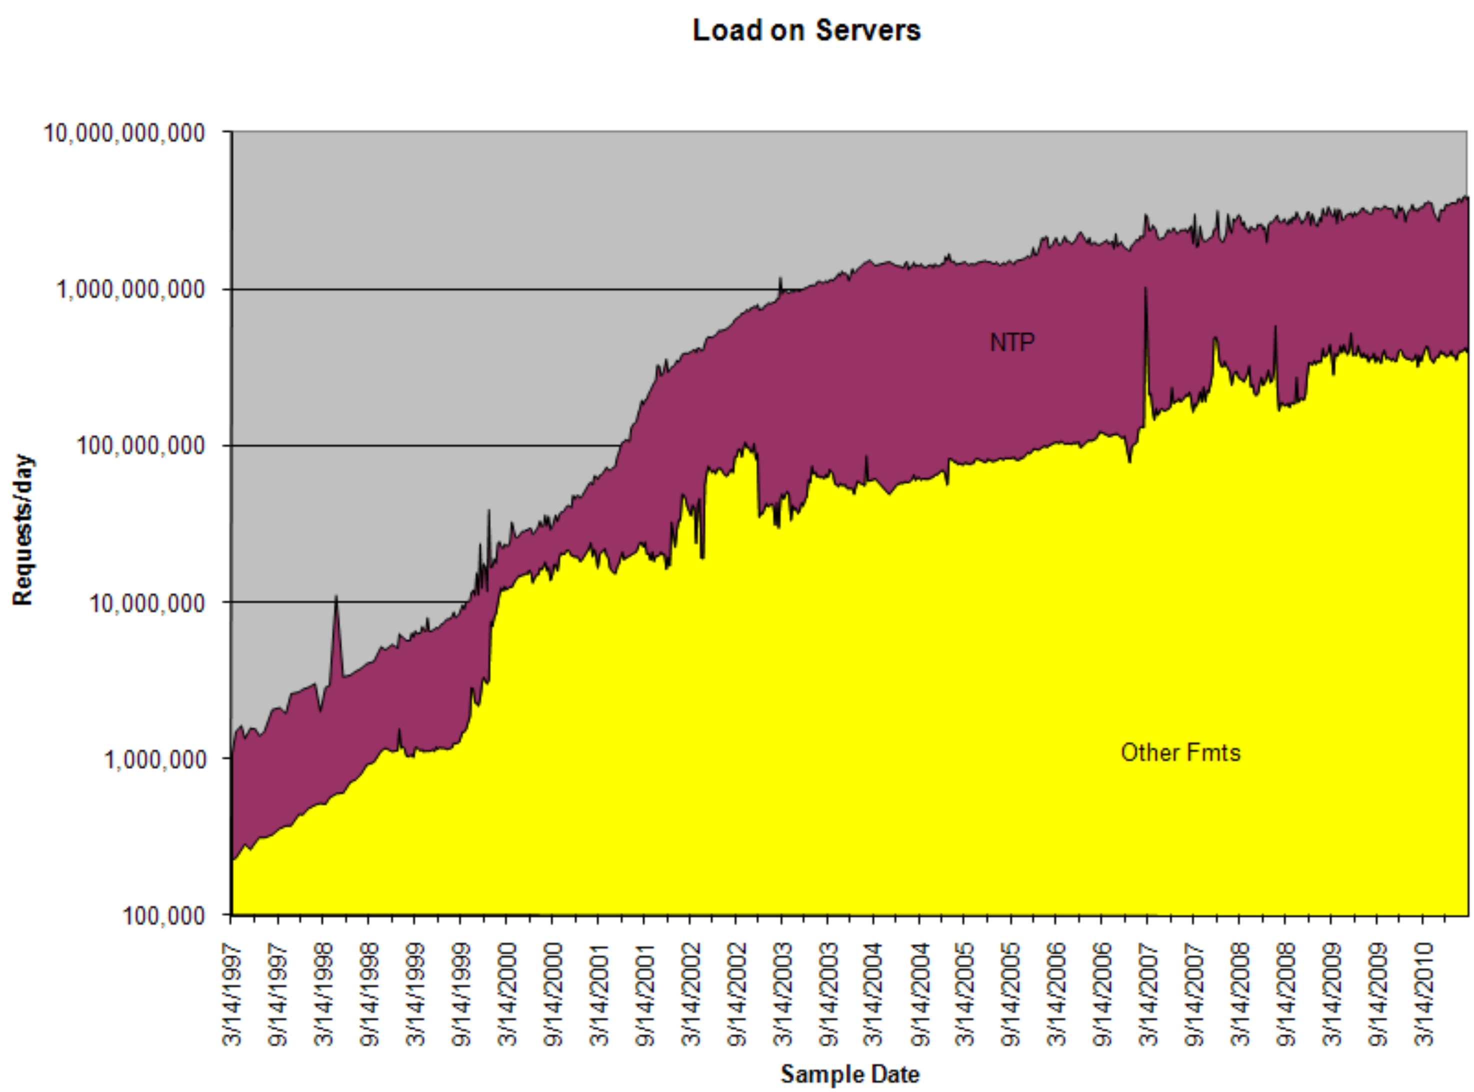

- One week averages since 1997 showing total number of requests to all NIST time servers
- Scale is logarithmic because variation is so large
- Increase is roughly 5% per month


Questions?