The Precision Timing Protocol

What is it, and why does anyone need it?

Andrew Gallo

2016 Technology Exchange Meeting

09/28/2016

Miami, FL

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

GW | Division of Information Technology

CAAREN
Capital Area Advanced Research and Education Network

Precision Timing Protocol

A Brief History of Time (pun unintended, but unavoidable)

- Why is a common agreed upon time important?
  - Event correlation
  - Task sequencing
  - Causation (a future event can not cause an earlier event
- Methods of Timekeeping
  - Sundials (Egypt, 3500 BCE)
  - Gravity – water/sand (Greece and China)
  - Mechanical (Europe, c1500)
- Advances in human technology have required improvements in time accuracy

Improvements in Timekeeping

- Agricultural societies needed accuracy of seasons
  - Sundials were good enough
- Worshipping required additional accuracy and technology for alarm clocks
- Travel was a big reason for improvements in coordinated timekeeping
  - European Age of Discovery - Location finding (Longitude)
  - Trains - event concurrency, time zones
- The speed of information dictates the accuracy and precision of our clocks

A Note on Terminology: Accuracy and Precision

- Accuracy – how close a measurement is to a true value
- Precision – how close repeated measurements are to eachother

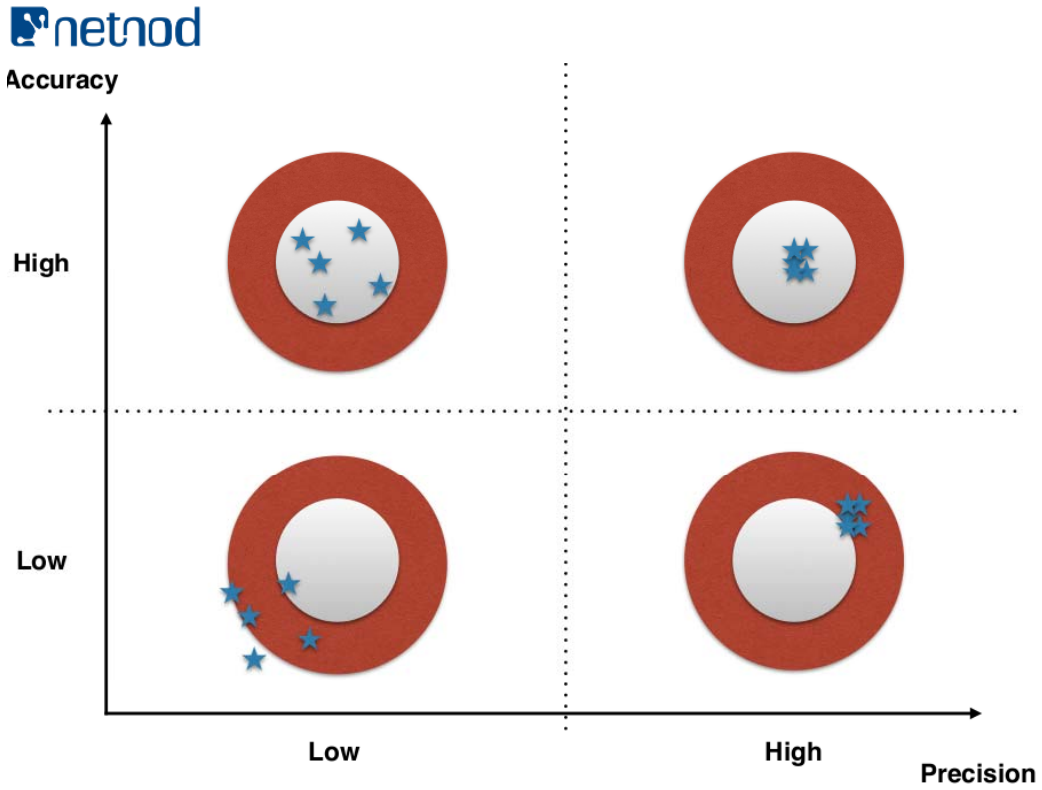

Source :https://www.netnod.se/sites/default/files/PAF%2020151014-Netnod-NTP.pdf

A Note on Terminology: Frequency, Phase, and Time

- Frequency
  - –Reference signal drives circuits to a common standard
    - "10 Mhz is the same everywhere"
- Phase
  - –making sure two systems understand when things start and stop- agree on milestones
    - "Everyone clapping together"
- Time of Day
  - –Wall Clock
    - NTP
- Time of Day and Phase are closely related

In The Beginning (beginning defined as mid-1990s)

- Most wide area communications was TDM based
  - Adaptation of voice technology to carry bursty packet traffic
    - T1, DS3
      - Primarily for private line & frame relay
    - SONET
      - Primarily for private line, ATM
- Most LAN technology was Asynchronous*
  - Ethernet
  - (others)

*It's more complicated than that….more in a bit

TDM goes away (temporarily)

- The boundary between enterprise networks (Ethernet) and wide area was a pain
- Growth of wide area data was explosive; voice traffic, declining
- Service providers slowly migrated to packet based wide area technologies
  - Done for the same reasons that Ethernet won in the enterprise
    - Cheaper, easier, faster….

TDM vs Ethernet

- TDM networks required more expensive hardware and careful planning to ensure proper synchronization design
- Ethernet had none of these constraints
- Ethernet is asynchronous network-wide
  - Buffer to managed congestion
  - No common clock
- But wait, don't all digital communications need to agree on clocking?
  - Each Ethernet receivers locks on to incoming bit stream (56 bit preamble) to train the receiver to the transmitters clock
  - This clock is *not* distributed to any other ports on the device

Ethernet Grows Up

- Ethernet has become the *de facto* enterprise and service provider transport
  - Industrial control
  - Video production
  - Automotive and train-borne
- Because so much has been converged onto a non-deterministic network originally designed for LANs, features need to be added
  - Far-end reporting and troubleshooting (Ethernet OAM)
  - Lossless fabric (for FCoE)
  - Multiplexing and subrating (FlexE)
  - Synchronization
- Service providers still need precise timing
  - Cellular backhaul, hand-off, LTE-A

Providing clocking on Data Networks

- Two standards bodies, multiple standards
- ITU-T
  - Synchronous Ethernet (SyncE)
    - Physical Layer, frequency distribution
- IEEE
  - Precision Timing Protocol (1588 and related standards and profiles)
    - Packet based, can provide frequency, phase, and ToD

Synchronous Ethernet

- Similar to previous service provider sync distribution
  - Network elements recover clock form the incoming bit stream
  - One (or more) ports designated to recover and distribute clock throughout the device
  - Each Ethernet link is now in sync with the common network clock
- Pros
  - Familiar design for service providers
  - “Easy” interworking with existing TDM networks
  - No additional traffic
- Cons
  - Requires every device in the path to be SyncE capable
  - Provides only frequency synchronization
- Common Uses
  - Mobile backhaul

IEEE 1588: The Precision Timing Protocol (PTP)

- A packet based, master-slave mechanism to distributed *both* frequency and phase (providing ToD)
- Pros
  - Can be incrementally deployed
  - Provides both frequency and phase sync
- Cons
  - More complex configuration
  - Additional network traffic

How is 1588 different than NTP

- Key differentiator- PTP is hardware assisted in the network elements that are PTP aware

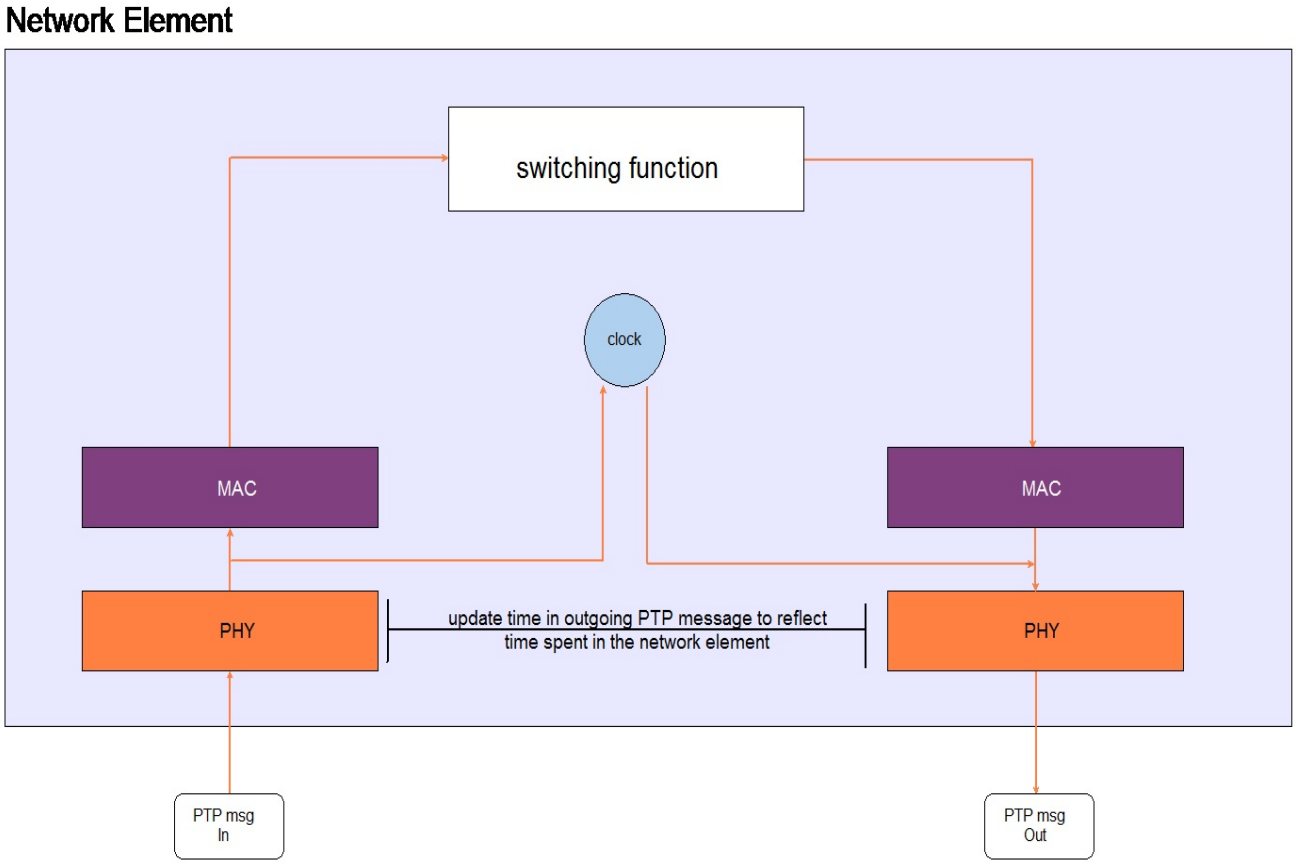

Precision Timing Protocol

1588 vs NTP performance

Protocol	Accuracy
NTP	50-100 milliseconds[1] .1 to 10s milliseconds[2,3]
PTP	20-100 nanoseconds[1]
IRIG-B	1-10 microseconds[1]

Sources:

[1] http://literature.rockwellautomation.com/idc/groups/literature/documents/wp/enet-wp030_-en-e.pdf
[2] https://www.eecis.udel.edu/~mills/ntp.html
[3] https://www.youtube.com/watch?v=250reOmrN70 (Cisco SPAG: Clocking & Sync Part 2/3: IEEE 1588 and PTPv2)

The Case of nanosecond accuracy

- Who needs more than NTP?
  - Service Providers
    - Mobile backhaul (roaming and handoff)
    - “HetNets” (small cells, reduction of inter-cell RF interference)
    - FCC studying over reliance on GPS for SP timing (CSRIC V, WG4, subgroup B)
  - Video production and editing
  - High Frequency Trading
  - Electrical generation (voltage can fluctuate, frequency can not)
  - Laboratory and scientific (replace IRIG?)
- Enterprise use cases
  - Cisco ACI
    - Fine grained data center performance monitoring (packet health scores)
  - ERSPAN Type III
    - Mirrored packet has timestamp in ERSPAN header
  - Spidercloud
    - Cellular rebroadcast, proprietary implementation of PTP
  - Service provider monitoring
    - Y.1731 monitoring of metro Ethernet circuits

Thank You!

Andrew Gallo
agallo@gwu.edu

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

GW | Division of Information Technology

CAAREN
Capital Area Advanced Research and Education Network