Which is more complex to design: Nanoseconds in hardware or microseconds in software?

Tommy Cook
CEO

Achieving nanosecond-level accuracy in hardware

Synchronization Network

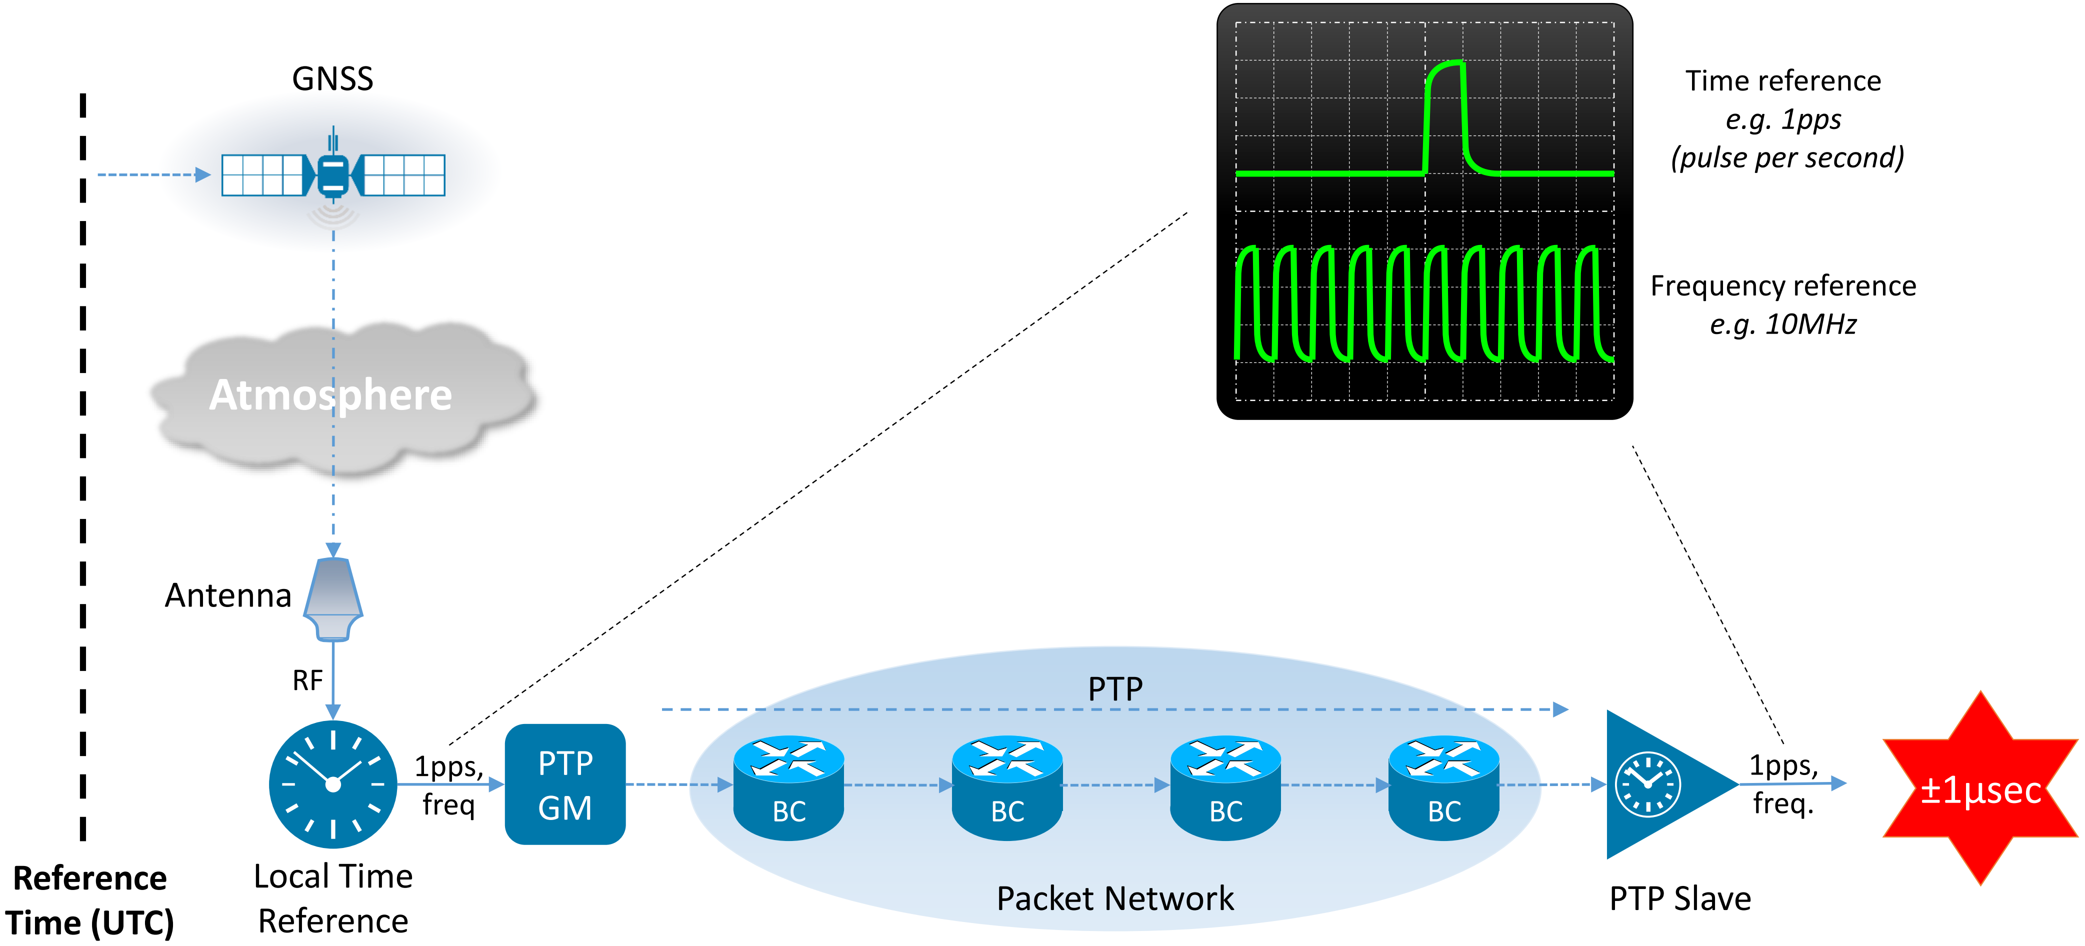

When you're looking for nanoseconds, everything counts

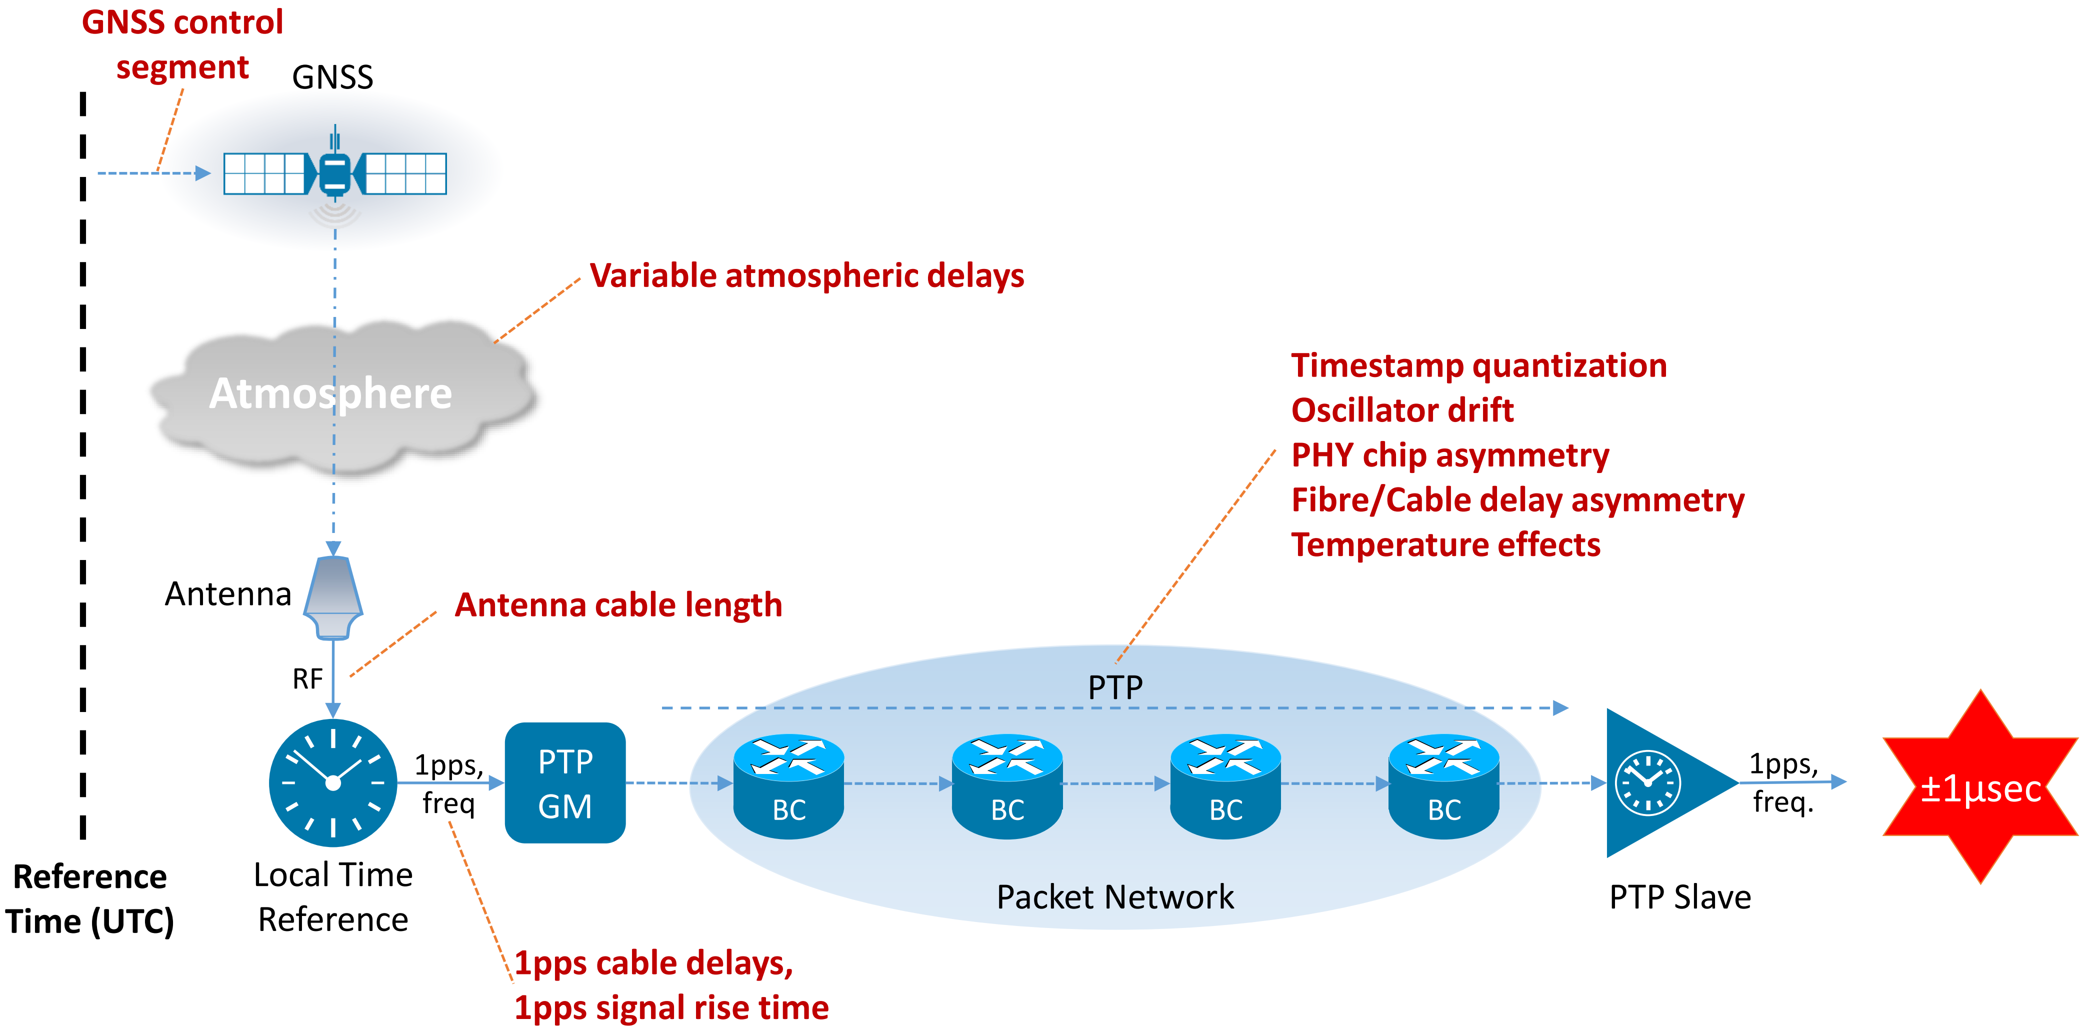

GNSS Receiver practices

Calnex

- Site antenna with a clear view of sky, away from other antennas that might interfere
- Atmospheric compensation:
  - Use ionospheric delay models to compensate for variable delay
  - Dual frequency receivers can measure ionospheric delay more accurately
- Compensate for fixed delays:
  - Measure signal delay through any cable amplifiers/filters
  - Measure RF cable length accurately
  - Measure 1pps cable length accurately
- Good design practices:
  - Use fast rise times on 1pps signal
  - Match impedances carefully for RF and 1pps cables
  - Avoid passive splitters ("trees") – they slow the signal edge
  - Low noise power supplies for accurate edge detection
- Use stable oscillator (OCXO, Rb):
  - Reduce GNSS receiver noise with low pass filter

PTP Clock Designs

Calnex

- Timestamp quantization:
  - Timestamp at PHY layer, not MAC layer
  - Use higher frequency clocks to drive timestamping
- PHY chip asymmetry:
  - Typical PHYs (especially high data rate) have longer receive delays than transmit
  - Need to compensate for the delay difference to the timestamp points
- PCB traces:
  - On FR4 board, signal propagation delay is about 70ps/cm
  - Keep traces short, transmit same length as receive
  - Use correct termination impedance to avoid reflections
- Oscillator drift:
  - For a SyncE-quality oscillator, phase drift can be up to 50 ns/s at constant temp.
  - At 16 PTP messages/s, the oscillator can drift 3.1ns between messages
  - Under variable temperature, could be much more…

Logical vs Physical Syntonization

Calnex

- Accurate time transfer also requires accurate frequency transfer

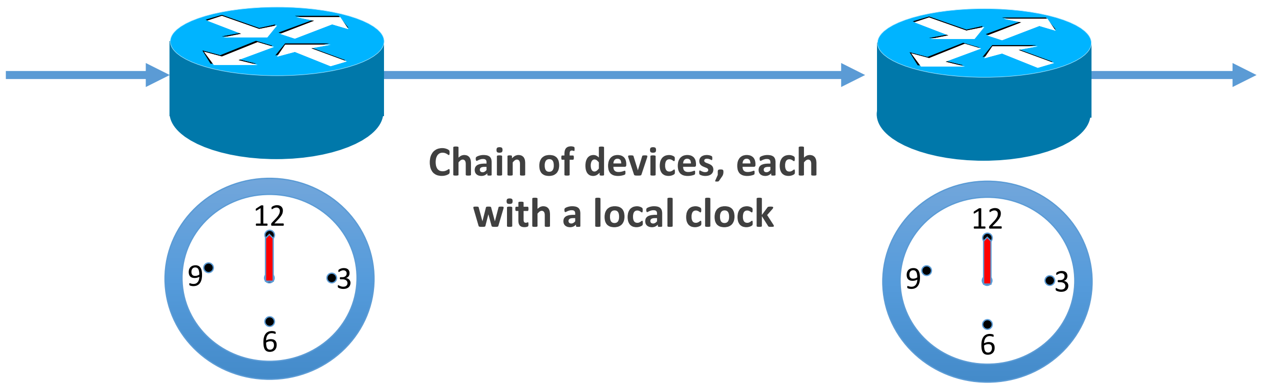

EITHER: Frequency of each clock locked to common reference (***physical syntonization***)

OR: Determine the frequency difference between successive devices and take account of that in the path delay calculations (***logical syntonization***)

- ***Logical syntonization*** (Cumulative Scaled Rate Offset, CSRO) is used in 802.1AS
- ***Physical syntonization*** to a SyncE reference is used in G.8275.1

Handling Link Asymmetry

Calnex

- PTP assumes symmetrical delays
  - Differences between upstream and downstream delay creates time error
- Upstream fibre may be different length to downstream fibre
  - Over 10s of kms, this difference can become significant
  - Operators have reported measuring up to 100ns delay difference over some longer fibre routes
- Upstream wavelength may be different to downstream wavelength
  - Propagation velocity at 1310nm is approx. 0.067% faster than at 1550nm
  - Equates to about 33ns difference over 10km of fibre
- In small networks or LANs, these factors might never be an issue, but in large telecom networks, they definitely will be
  - Either minimise the differences to keep them below a budget
  - Or measure each link individually and compensate for the difference

Achieving microsecond-level accuracy in software

Job Done?

- ***Hardware based networks*** can deliver 1 microsecond.

- That means delivering 100’s microseconds
  in ***software based systems*** must be easy, right?

Software Determinism – of lack of it!

Calnex

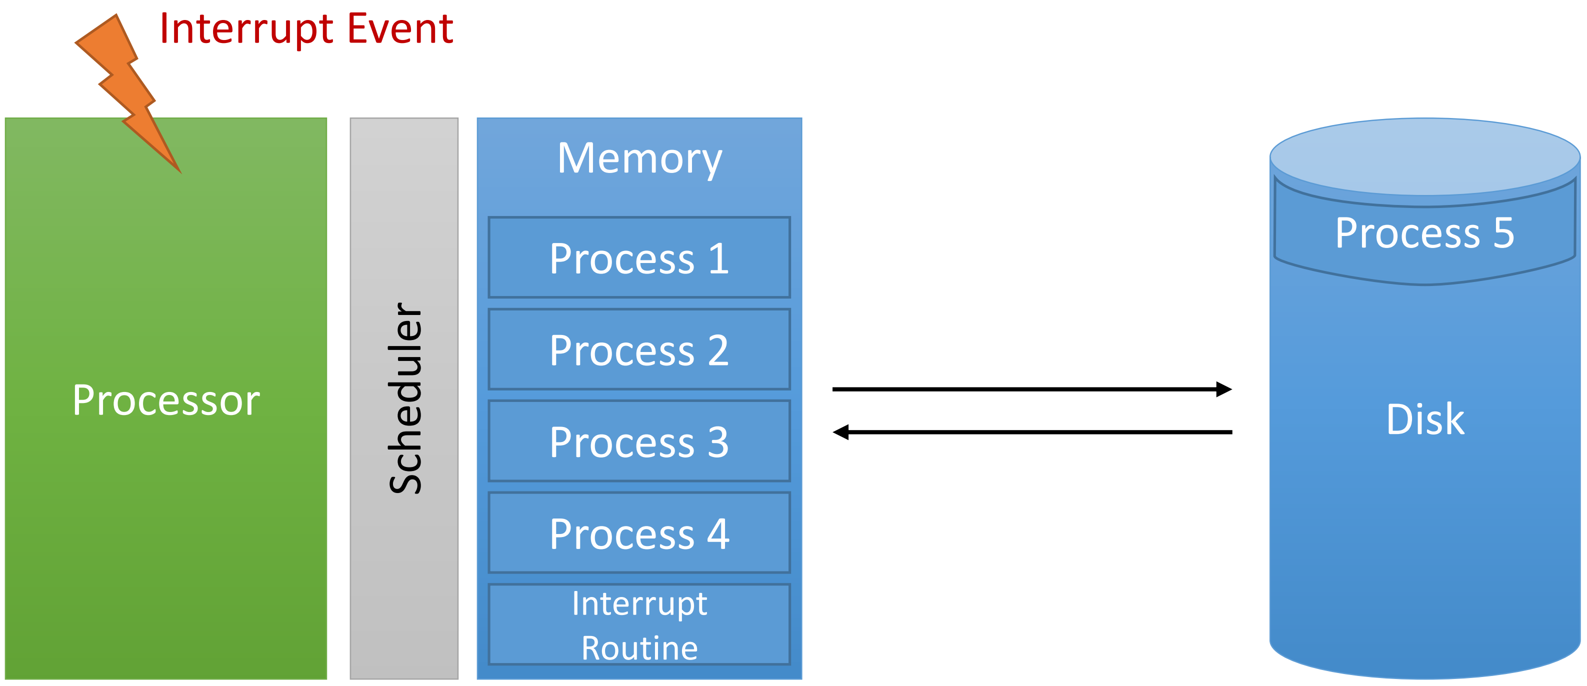

- The processor is shared between multiple processes.
- The number of active processes varies dynamically.
- Processes can be swapped out to disk to make space
- Interrupts can happen at any time!

Virtualization

Calnex

- How do you get accurate time into a virtual application?
  - With virtualization, the extra layers of software can make it even worse!

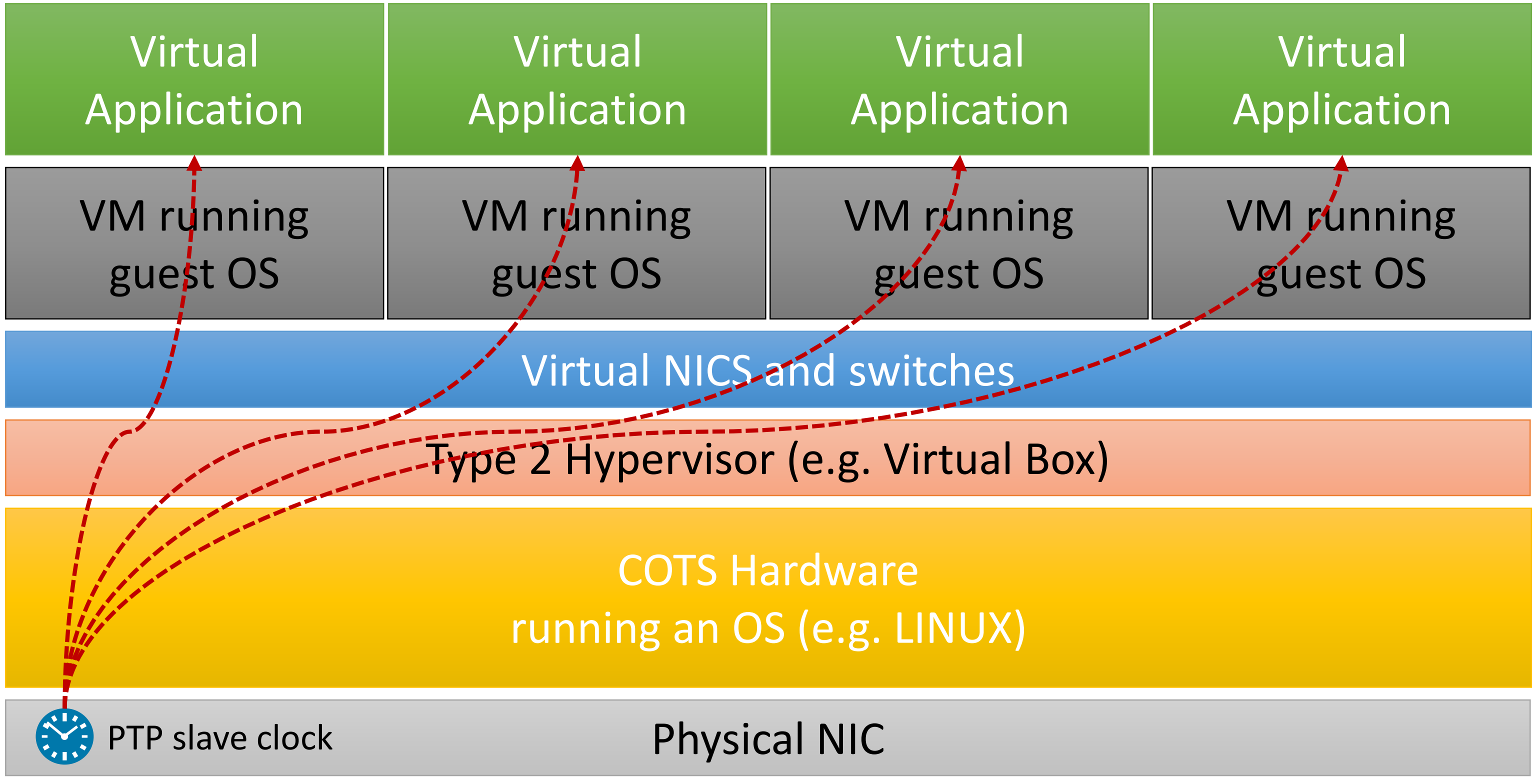

Logical Syntonization – again!

Calnex

- Borrowing the logical syntonization idea
  - Use the CPU system clock as the final clock in the chain
  - Measure the frequency of the system clock against the accurate frequency distributed using PTP
- Requires either
  - Hardware support in the CPU to compare the two frequencies
  - Or a low-level software task to constantly adjust the comparison data
- Virtual application uses system time as a proxy for PTP time, and adjusts it with the logical syntonization information
  - Still needs a fast system call to get system time
- OS still needs to support
  - Time as a construct, synchronized to an available time reference
  - Logical syntonization to an available frequency reference
  - Fast access to time information for virtualized processes

So where does this all get us?

Benefits of accurate time

Calnex

- The requirement for time is inherent in many industries
  - **Automotive** – sensor timestamping, plus infotainment and audio applications
  - **Financial trading** – now requires 100µs accuracy for record keeping
  - **Industrial machinery** – the more accurate the timing, the more accurate the process, and often the faster the operation
  - **Power generation** – accurate monitoring of power flow requires accurate time
  - **Telecoms** – accurate scheduling for transmission/reception improves efficiency and throughput

- Time is not just required in hardware, these processes are all software controlled
  - Increasing levels of virtualization happening in all industries

Summary

Calnex

- The world is going virtual, like it or not
  - Many processes now implemented in hardware will become software in future
  - The need for accurate time is still there, but won’t stop the virtual wave

- Nanoseconds in hardware, or microseconds in software?
  It’s not ***“either/or”***, but ***“both/and”***
  - The ***distribution*** of time will still require hardware
  - The ***use*** of time information will be in software

- **Both situations require careful system design.**

Insight and Innovation

calnexsol.com

Tommy Cook
CEO
tommy.cook@calnexsol.com
+44-770-4894-128