DPEN Compute Project®

OCP TAP Time-as-a-Service

Lasse Limkilde Johnsen



IEE1588/PTP is a Great Achievement!*

- The success of getting PTP support built-in to most commodity hardware network controllers MACs and PHYs is quite an achievement
 - Recently, I wrote a Linux kernel driver on behalf of Raspberry PI for the RPI CM4 first IEEE1588 driver in a PI.
 - Linux APIs have in-hardware and in-kernel timestamping in addition to things like one-step, extts and perout support for pulse-per-second input and output.
- The architects of IEEE1588-2008 have achieved some important milestones
 - Dishonest practices in broker-dealers has been significant curtailed by MiFID regulations in no small part because sequencing is now a recorded fact. A leads to B leads to C. It's recorded in real-time even in large distributed system. Front-running is much harder to get away with than it used to be because of the legal requirement for accurate timekeeping.
- Their work allows for great accuracy and the impact that this has had is tangible in many sectors, but no more so that in the financial services sector.
- The way PTP works is brilliant

Project

PTP solves real problems in the real world.



*However.... There are a few unresolved issues.... ③

...Let's talk about them...

Boundary clock distribution is literally nothing but single points of failure and unquantifiable error



Project

Everything green is a single-pointof-failure. Every :

- Clock
- Switch
- Host
- PTP Session / Connection between them!

The single-parent principle is everywhere

Boundary clock distribution and BMC

• Here is a typical PTP deployment

Compute Project®

- It is completely unknown if nodes have the same time
- In the real world, very (very) often, the answer when we measure is "no, they don't"
- The best master clock concept of self announcing how good the quality of the time distributed is "not a good strategy" ™
- It's hard not to think that BMC was an afterthought in PTP
 - It's also optional
- Problems changed when we went from ms->us->ns->ps
 - What used to be ignorable and insignificant is now the opposite

Does this server / have the same time as this server???





We thought about how to solve these problems...

Introducing PTP+Squared



Let's turn the tree upside down....

- Now any server has multiple paths to sources of time
- All of the paths can be compared in real time
- Error of variation can be communicated downstream in a cumulative fashion
- In that way if any part of the tree becomes congested it will be visible and time can be taken from a better path



How do we create the reverse tree / mesh structure?

• How do you do it?

Project

- What does it look like?
- We do it with something we call PTP+Squared
- It's a p2p distributed hash table type overlay network to PTP that sets up unicast sessions between PTP participants in the grid
 - we call them "nodes"
- Everything is completely dynamic and in real-time
 - D3.js is cool 😳



Error is quantifiable and we compare paths

• Squared nodes showing their cumulative error in the new alerts panel.

Project

- Above certain thresholds the bubbles grow and turn orange and red.
- This is 8x faster than realtime, so 5 1/2 minutes condensed into 40 seconds.
- Here the mesh goes from largely bad to largely good in that time.





Hybrid approaches welcome



Here are some things we are looking at next...

- Timebeat is currently closed source "freemium" software.
 - We feel that Open Source is a better model.

Compute Project®

- There are several paths we can chose going forward, but a reasonable guess it that in 2-3 months Timebeat will be "free-free" and open source.
- I will concern myself with the problem of synchronisation between the kernel and the ethernet controller.
 - Today doing sync across the PCIe bus between the CPU and the NIC
 4 centimeters apart introduces more error than between two NICs
 4,000 kilometers apart.
 - Daisy chaining NICs with PPS cable is a hopelessly uncool scheme.
- Scalable PTP deployments using commodity hardware and optimized software in kernel by-pass mode.
- We are launching devices pre-built with GNSS/OCXO/NICs to enable large scale PTP+Squared deployments.
 - We are aiming to significantly reduce the cost of injecting time from GNSS into a datacenter type environment.





Thank you for your patience!

If you want to get in touch :

Lasse Limkilde Johnsen lasse@timebeat.app +44 7832 335 392

lan Gough ian@timebeat.app +44 7989 140 622