



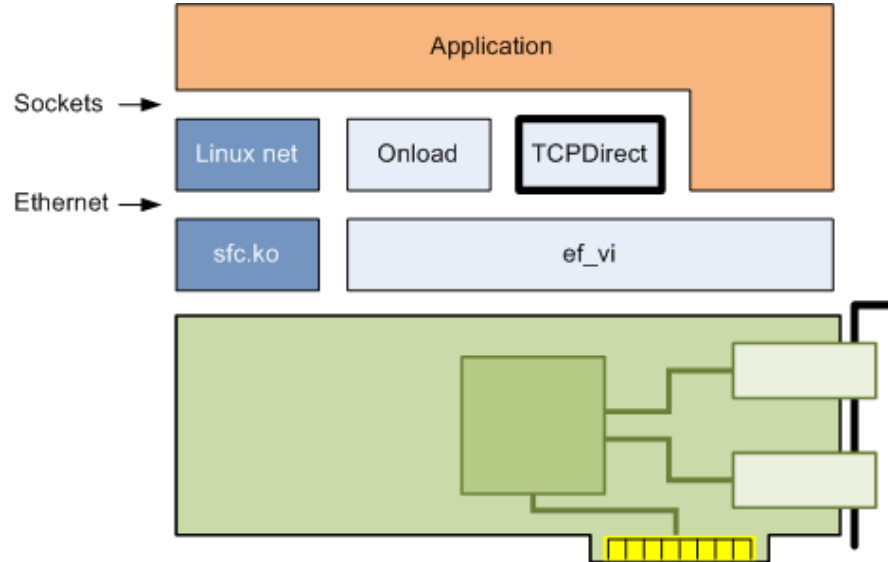
Faster and In-Sync

Davor Frank


TCPDirect: Ultra-low latency API



- Sockets-like API supporting TCP and UDP (unicast and multicast)
- Much lower latency than Onload
- Much easier to use than ef_vi
- Stable API and ABI



TCPDirect: Performance



Not STAC Benchmarks

TCPDirect: Coming soon



- By popular request...
 - ◉ Hardware timestamps
 - ◇ Tx and Rx
 - ◇ TCP and UDP
 - ◉ Bonding
 - ◇ Max two interfaces
 - ◉ MSG_WARM
 - ◇ Reduces send latency at lower rates

Timestamps and time synchronisation



- Solarflare adapters support:
 - ◉ High resolution hardware timestamps for all packets
 - ◉ Stable stratum 3 oscillator
 - ◉ Adapter clock synchronised to UTC

- sfptpd: Time synchronisation on Linux
 - ◉ PTP, PPS and NTP
 - ◉ Synchronises host clock (+/- 200ns) and adapter clocks (+/- 50ns)
 - ◉ Multiple time sources with best clock selection
 - ◉ Supports link aggregation and hot-plug


Not STAC Benchmarks

sfptpd v3.2 now available



New features:

- Machine-readable local logs
- Remote monitoring
- Configurable clock selection policy
 - With improved logging of *reason* when new clock is selected
- Support for hot-(un)plugging NICs when using bonding



sfptpd: Machine readable logs



- JSON format logs and long-term stats
- Easy to integrate with databases and monitoring tools


```
clock-name: phc0 (enp4s0f0/enp4s0f1)
clock-id: 000f:53ff:fe43:88f0
state: ptp-slave
alarms: none
control-flags: selected timestam...
interface: enp4s0f0 (enp4s0f0)
timestamping: hw
offset-from-master: -8.000
one-way-delay: 67.000
freq-adjustment-ppb: 82.576
in-sync: 1
ptp-domain: 0
steps-removed: 1
parent-clock-id: 000f:53ff:fe43:8860
parent-port-num: 1
delay-mechanism: end-to-end
...
```

```
{
  "instance": "slave100",
  "time": "2017-08-08 10:20:42.843704",
  "clock-master": {
    "name": "gm"
  },
  "clock-slave": {
    "name": "phc0(enp2s0f0/enp2s0f1)",
    "time": "2017-08-08 10:20:42.843991",
    "primary-interface": "enp2s0f0"
  },
  "is-disciplining": true,
  "in-sync": false,
  "alarms": ["no-follow-ups"],
  "stats": {
    "offset": 980.500000,
    ...
  }
}
```

sfptpd: Remote reporting




- Slave event monitoring
 - Uses PTP signalling messages
 - Defined in *draft* of next PTP spec
 - + Solarflare slave status extension
- Rx and Tx event timestamps
- Calculated offsets and delays
- Slave status
 - Sync state, alarms, error stats



sfptpd: Meinberg NetSync Monitoring



- Closed-loop
 - ◉ Detects problems even when slave believes everything is okay
- Monitoring system probes slave clocks using “reverse PTP”
 - ◉ Accurate measure of slave clock
- Status and stats piggy-backed on returned messages
 - ◉ Including MTIE



The Ultimate Trading Machine

World Record 98ns Tick-to-Trade Latency Based on STAC-T0™ Benchmark *



STAC-T0 measures the time between transmission of simulated UDP market data to the system and receipt of simulated TCP orders from it, without the system performing any trading logic or market-specific protocol handling.

Intel® Xeon® Scalable (Skylake) Processors

Intel Xeon Gold processors offer monumental leaps in I/O, memory, storage, and network technologies.



LDA Technologies LightSpeed TCP™ Cores

An ultra-light, ultra-high-speed, and ultra-low-latency FPGA-based distributed TCP offload with processing latencies under 20ns and thousands of TCP connections.



Solarflare XtremeScale™ Software Defined NIC

Leverages the Delegated Send™ capability of the Onload™ kernel bypass-enabled NIC—and Solarflare Application Nanosecond TCP Send (ANTS) technology—to maintain TCP connections that delivery blazingly fast network performance.



Xilinx Kintex® UltraScale™ FPGAs

Kintex UltraScale FPGAs are optimized for best-in-class performance per watt fabric in 10G to 100G networking applications.

Penguin Computing Relion® Server

Optimal performance through carefully selected and vetted processors, memory, bus, storage, and other options, architected into a 1U 19" EIA traditional form-factor.

Improving tick-to-trade latency is critically important in electronic trading because it means improving the queue position of trades, which increases the probability that trades will be executed. The Ultimate Trading Machine server configured by Penguin Computing, Solarflare, LDA and XILINX achieved 98-nanosecond tick-to-trade networking latency in STAC-T0 benchmark testing.

* Maximum value reported for STAC-T0.β1.MEDRATE.ACTIONABLE. See the full STAC Report™ at www.STACresearch.com/news/2017/10/13/SFC170831.



A large, abstract graphic in the top-left corner consisting of multiple overlapping, flowing blue lines that create a sense of motion and energy.

Thank you