Davor Frank

STAC Summit Chicago

October 16th, 2018

ExtremeScale™ X2500 Highlights

- Next Generation Adapter
  - 10,25,40,50,100G Ethernet
  - 25% latency reduction
  - 23% frame rate increase
- Precision time on all adapter SKUs
  - Stable Stratum 3E oscillator
  - Hardware time stamps every frame
  - Time synchronization via PTP and PPS
  - Synchronises host clock and adapter clocks
  - Improved hardware timestamp precision
- Security
  - Tamper resistant, signed firmware
  - Integrated firewall
- Overlay network acceleration
  - VXLAN, NVGRE, Geneve, OVS

Ethernet latency improvement

Not a STAC Benchmark

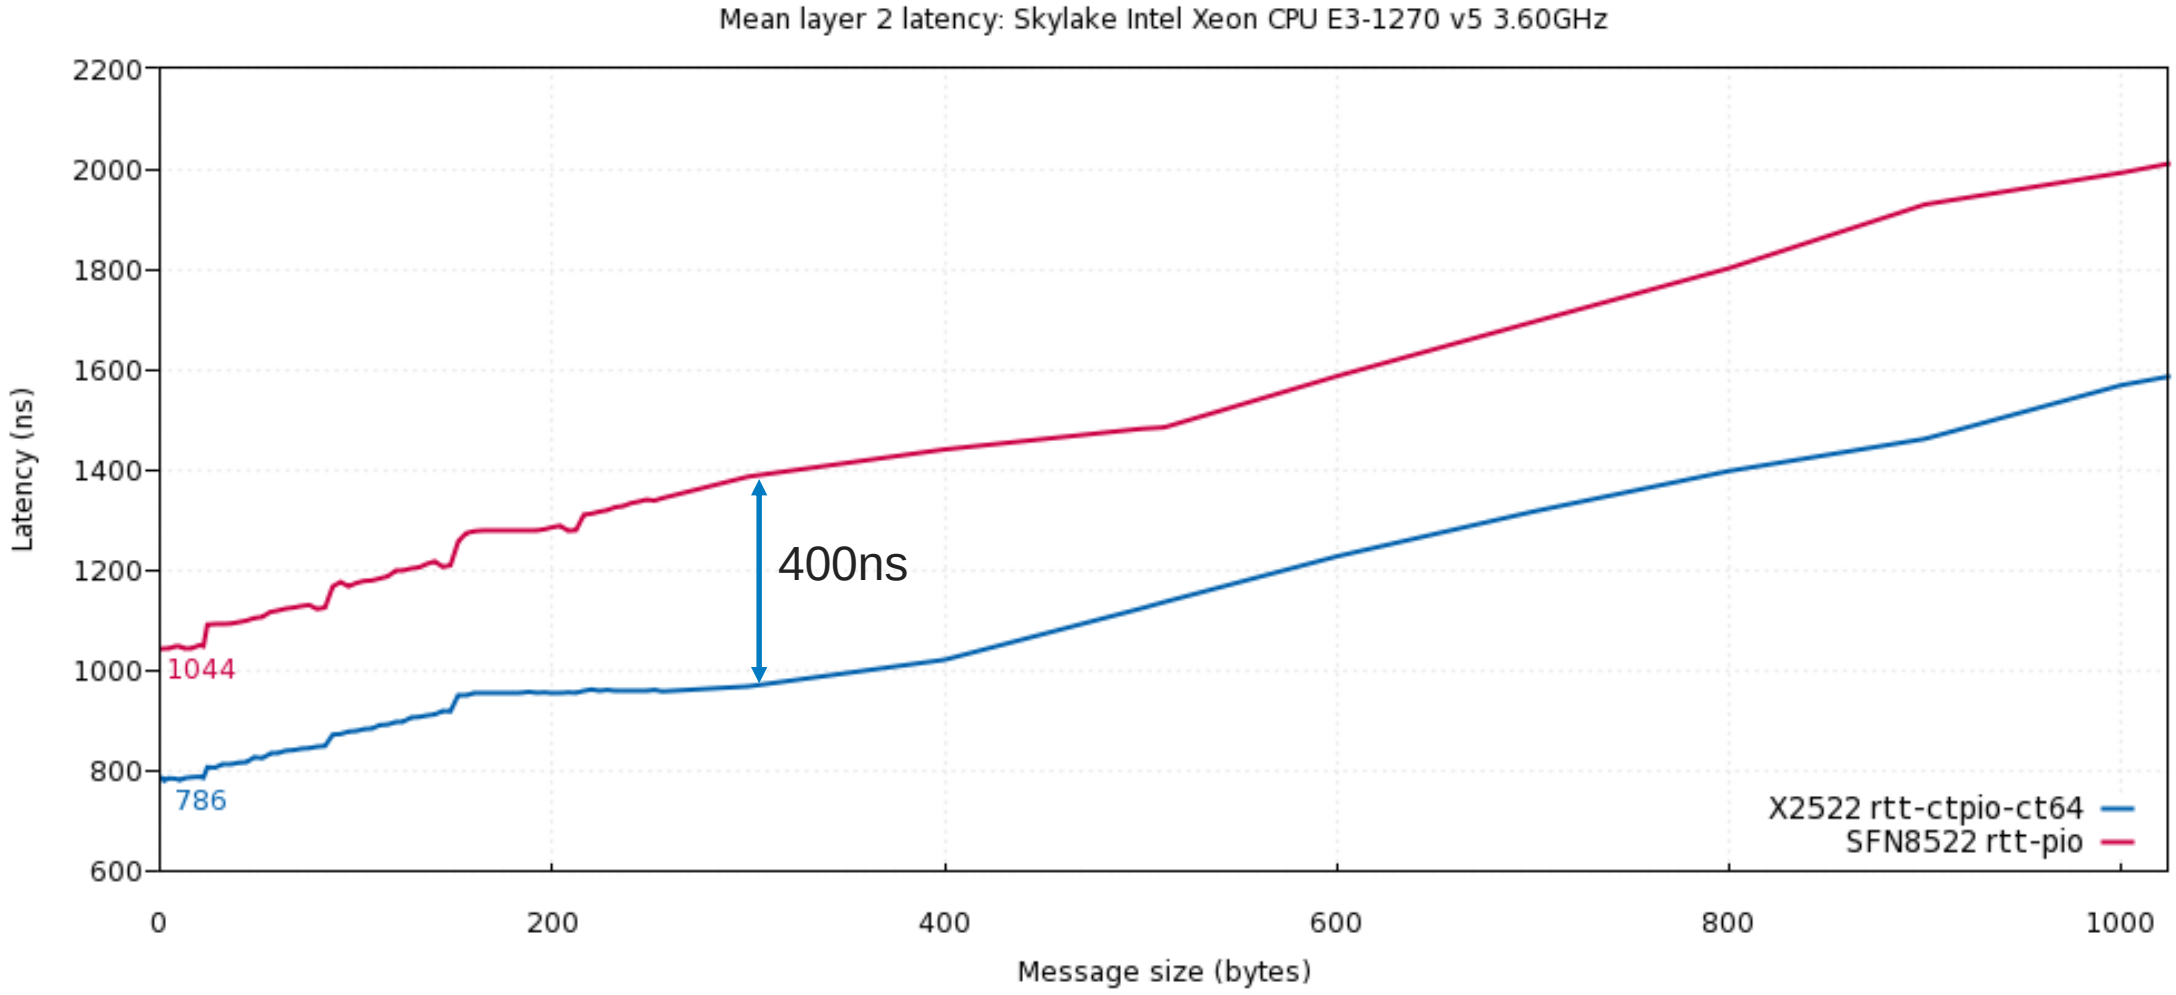

Onload

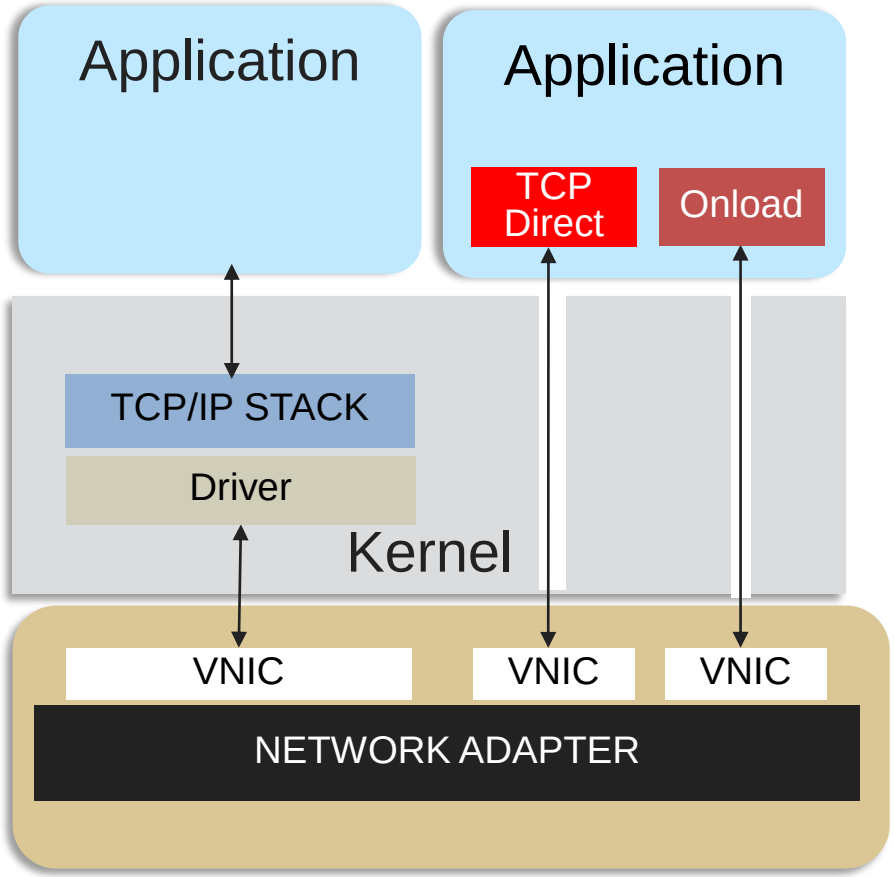

Onload accelerates:

- Latency sensitive applications AND
- Highly threaded and horizontally scaled applications
  - TCP & UDP (unicast & multicast)
  - Pipes
  - TCP and UDP loopback
  - Select(), poll() and epoll()
  - White/blacklisting/DDOS protection
  - Remote monitoring

Onload is compatible:

- Standard BSD sockets API
- No changes to applications needed
- Many languages (Java) / Middleware
- Standard protocols on the network

Onload supports:

- Precision timestamps
- Bonding (active-back and LACP) (teamd)
- VLANs and overlay networks
- Guest VM and Container acceleration
- FPGA offload

25G Latency Improvement (UDP)

Not a STAC Benchmark

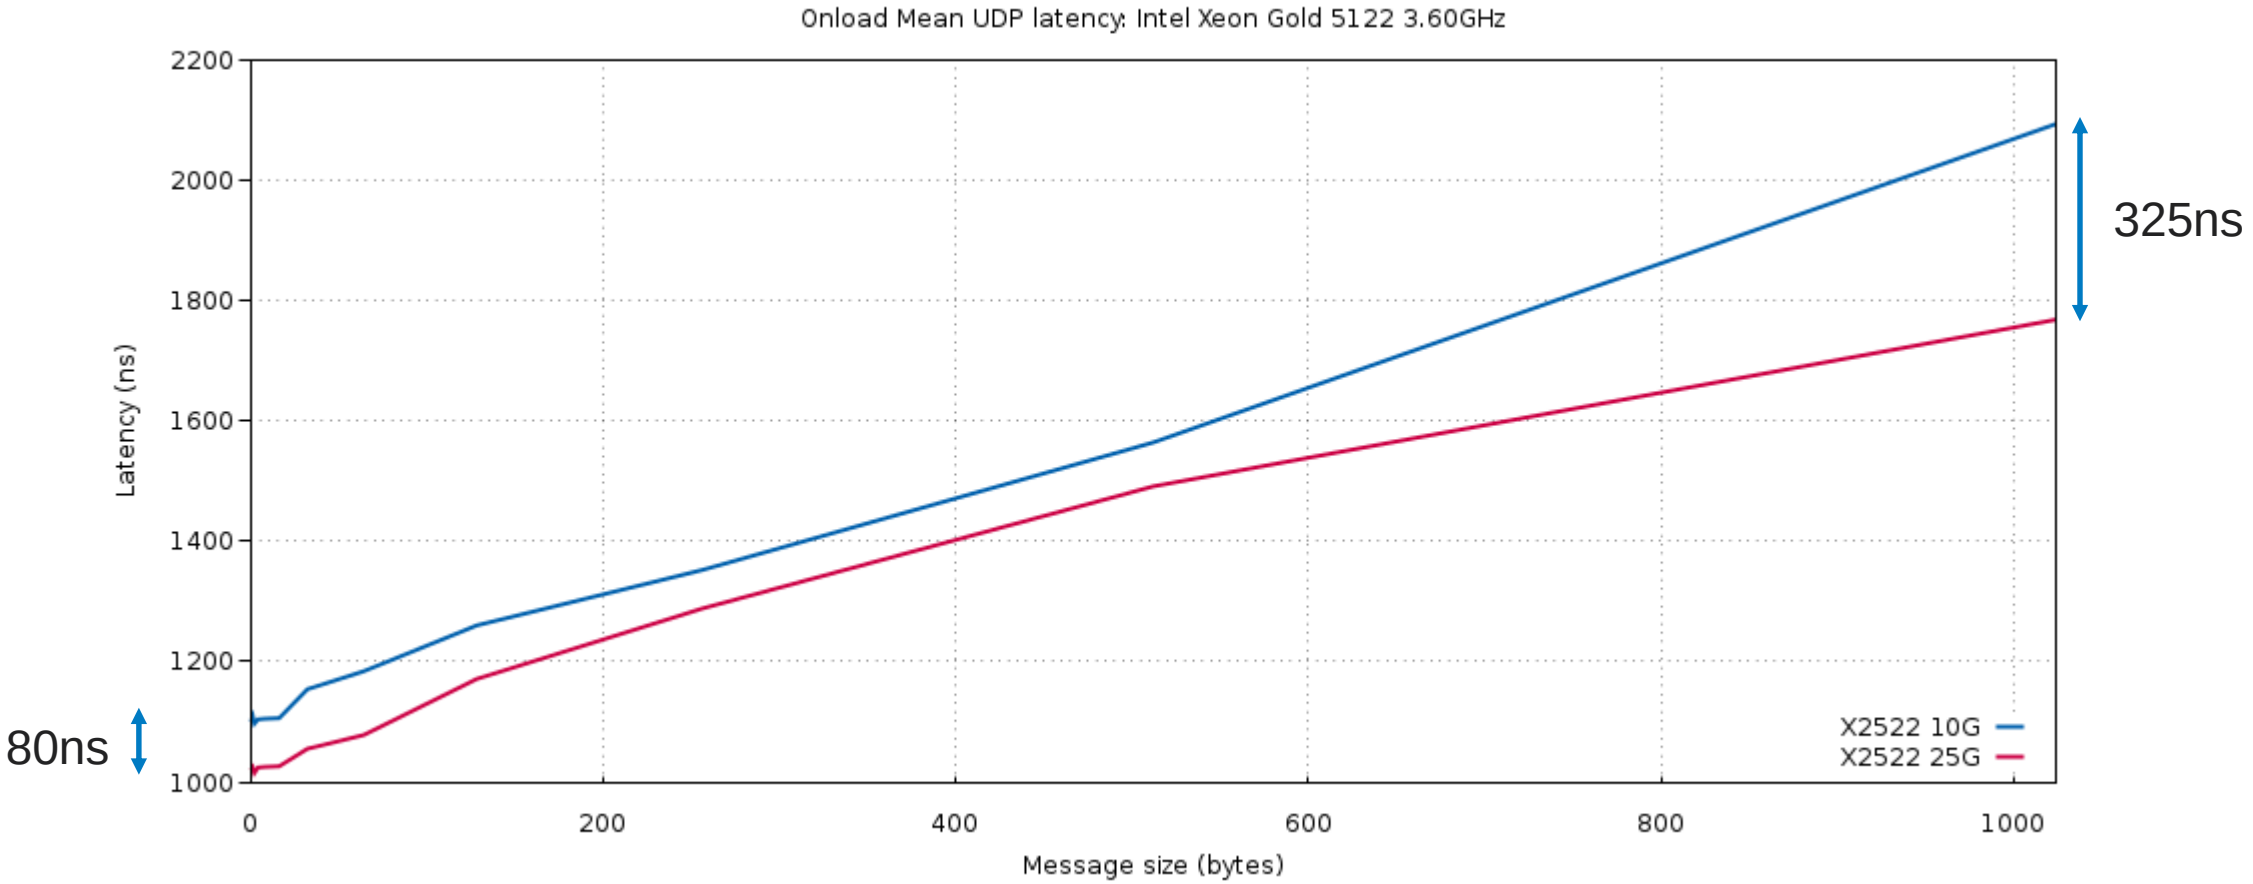

TCPDirect

- Sockets-like API supporting TCP and UDP (unicast and multicast)
  - Zero-Copy receive
  - Waitable FD' for integration with select/poll/epoll
  - Multiplexor for multiple sockets
  - Supports VLAN interfaces
  - Hardware timestamps
  - Simple bonding configurations
  - MSG_WARM and DMA fallback
- Much lower latency than Onload
- Much easier to use than ef_vi
- Stable API and ABI

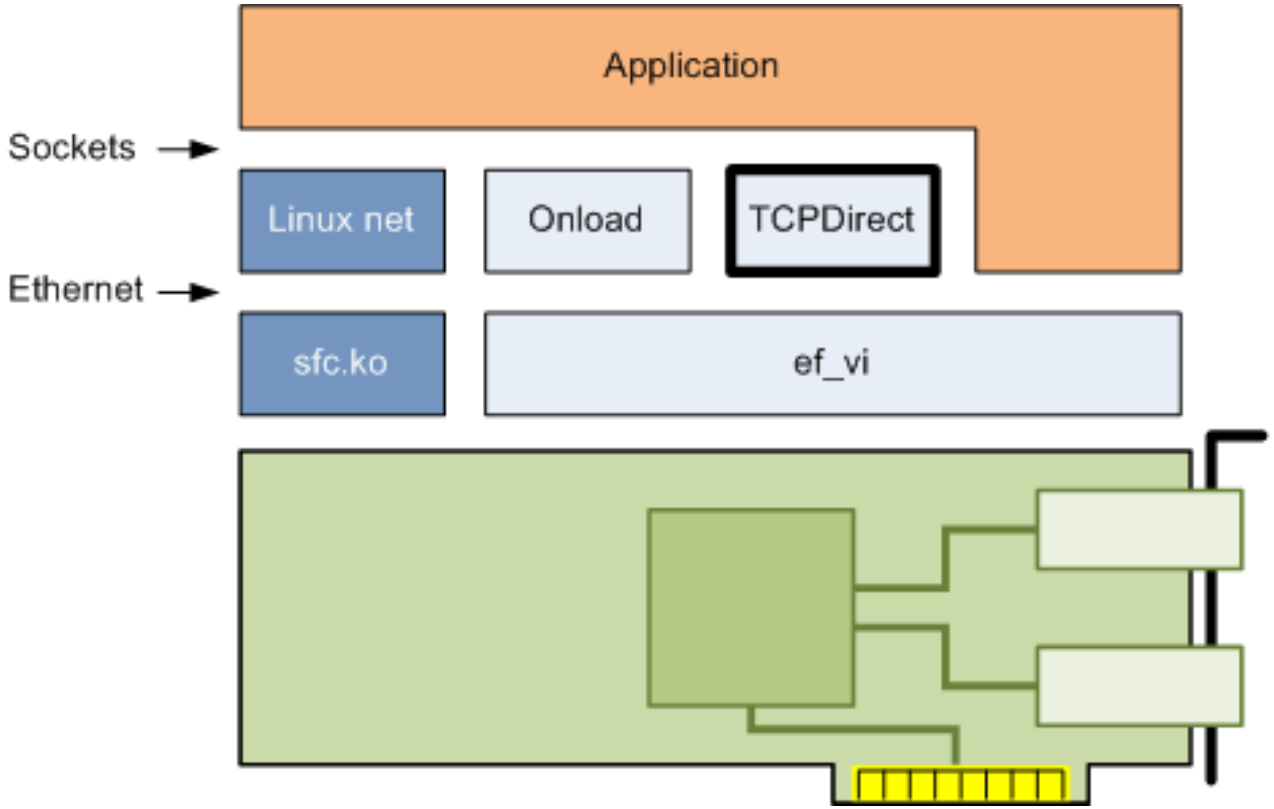

25G Latency Improvement (TCP)

Not a STAC Benchmark

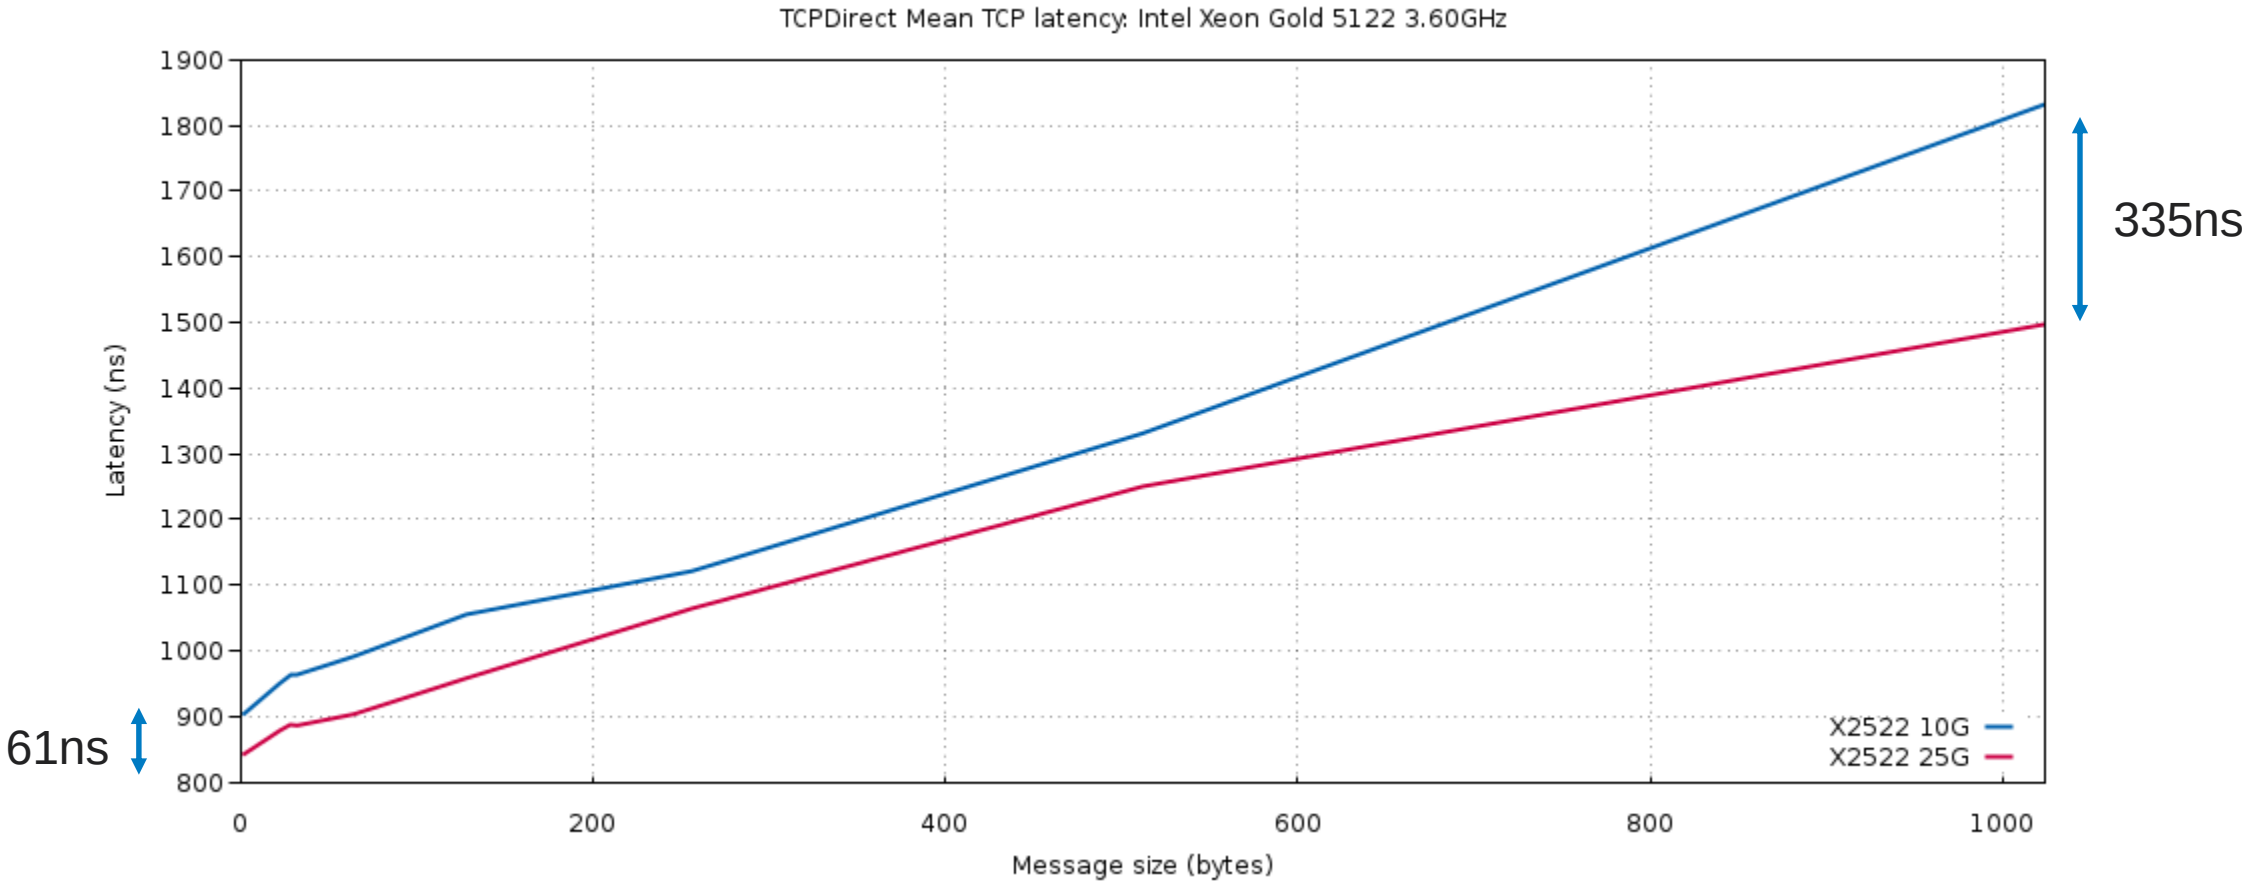

25G Cables

Forward Error Correction (FEC)

- To achieve low BER on longer cables 25G links use FEC
- Two types of FEC
  - BASE-R FEC (Firecode) and RS-FEC (Reed Solomon)
- FEC increases latency
  - RS-FEC ~250ns and BASE-R FEC ~80ns

3 types of DAC cables

- CA-25G-L: up to 5m, requires RS-FEC
- CA-25G-S: up to 3m, lower loss, requires either RS-FEC or BASE-R FEC
- CA-25G-N: up to 3m, even lower loss, can work with RS-FEC, BASE-R FEC, or no FEC

Better

USE THIS

Length	AWG	Type
0.5	30	CA-N
1	30	CA-N
2	30	CA-S
2	26	CA-N
3	30	CA-L
3	26	CA-N
5	26	CA-L

Amphenol SFP28

Onload for Core applications

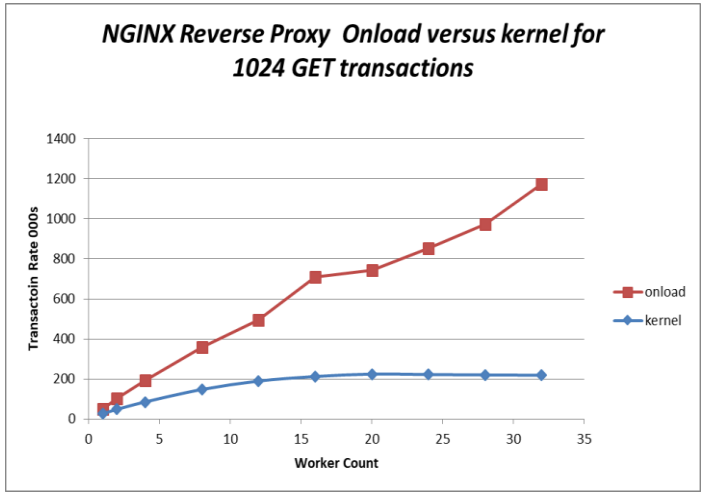

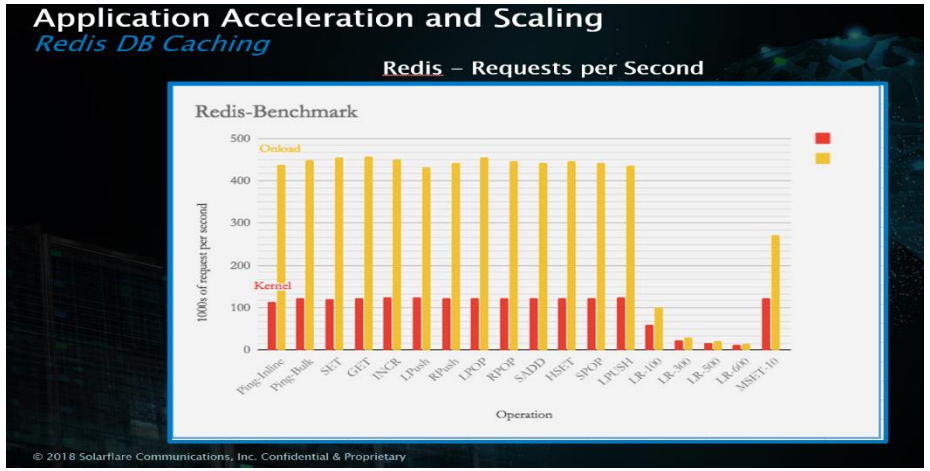

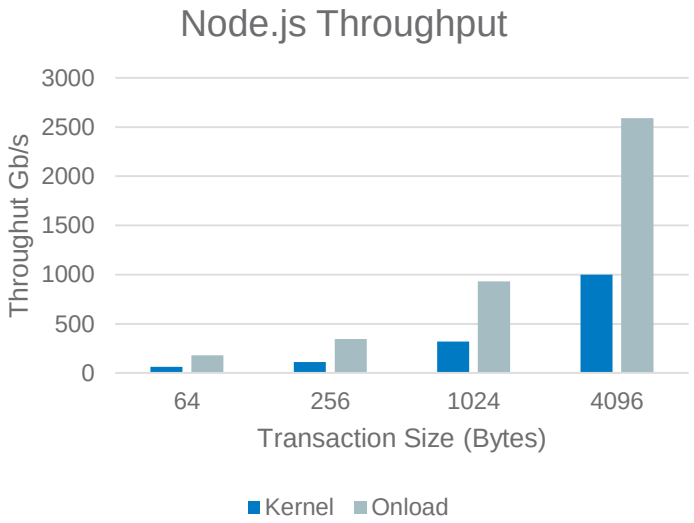

- Not just trading applications
- Scales to support large numbers of processes, sockets and flows
- Transparent acceleration for critical infrastructure components
  - e.g NGINX, Redis, Node.js

Integrated Security

- NIC resident secure enclave
- Control
- Telemetry
- Anomaly Detection
- Policy Enforcement
- Low-Latency

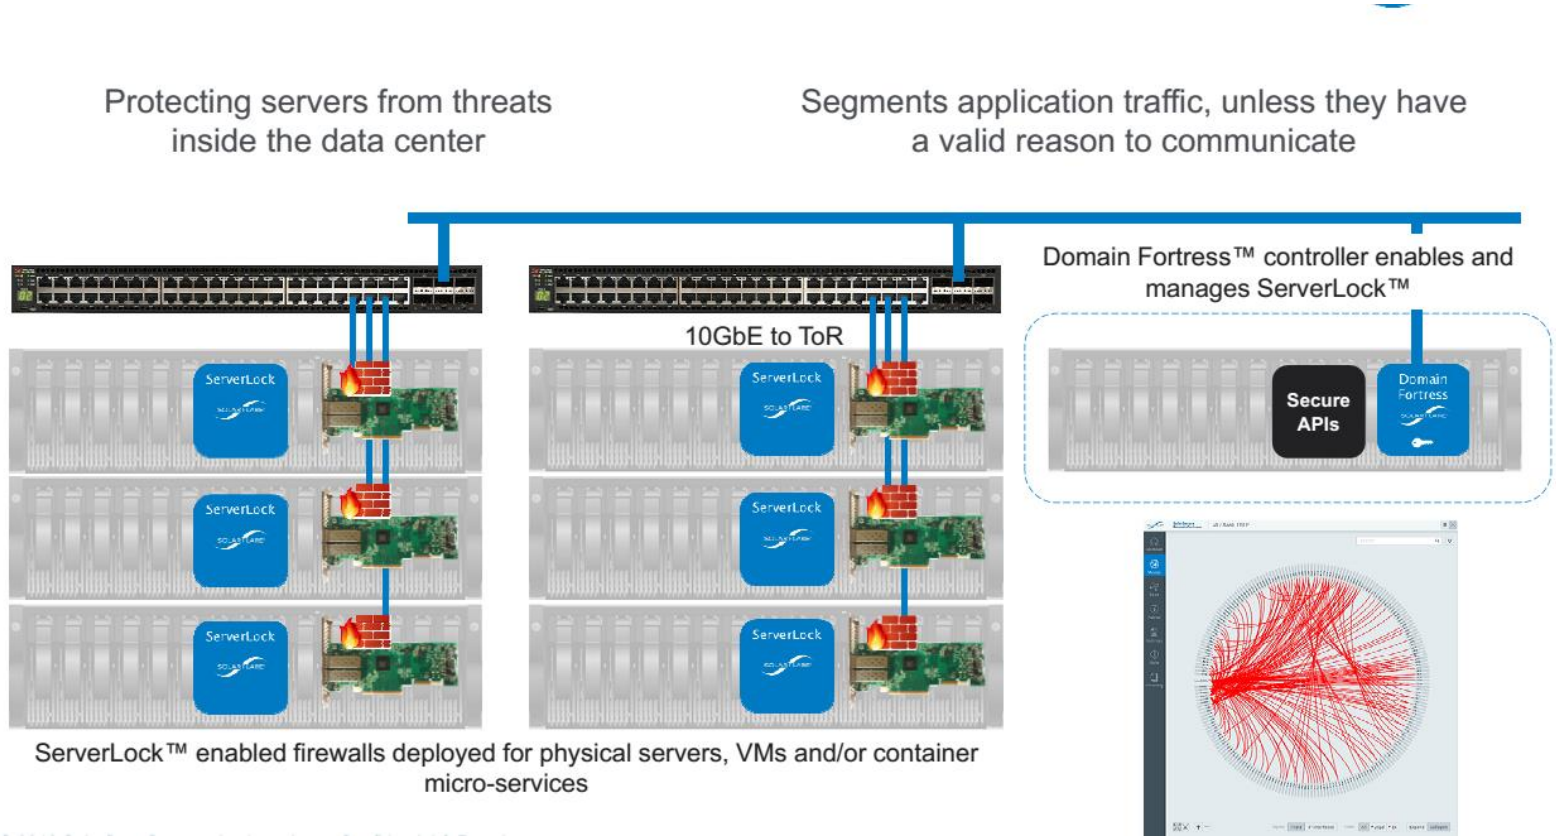

Solarflare ANTS Technology

Algo-Logic's Trading Solutions:

- Detect a trigger event and send out an order in nanoseconds
- Significantly reduce time to market with pre-built solutions allowing trading firms to instantly start trading

ULTRA-LOW LATENCY SOLUTIONS FOR TRADING BASED ON LDA LIGHTSPEED TCP

NOVASPARKS and SOLARFLARE: A Full FPGA Tick-to-Trade Development Platform