

Low latency networking, time synchronization and packet capture using Exablaze NICs and switches

Peter von Konigsmark peter@exablaze.com

STAC Innovation Roundup, Spring 2016



Ultra-low latency network cards with timestamping & PTP





All ExaNICs support:

- Standard Linux network driver
 - Windows driver coming soon



- Transparent socket acceleration on Linux
 - In most cases no software modifications required
 - Sub-microsecond TCP/UDP latencies achievable
- For low-level users: Exablaze API library
 - Simple send/receive API
- For even more low-level users: customization of firmware with FPGA Development Kit
 - ~110ns tick-to-trade achievable for simple trigger (on X10)*
 - Some FPGA programming required, but examples provided

* Depends on format of incoming packet. NOT STAC BENCHMARKS



All ExaNICs support:

- Timestamping to 6.2ns resolution
- Time synchronization via any of:
 - hardware-assisted PTP
 - pulse-per-second input
 - built-in GPS (on X10-GM)
 - sync to host
- exanic-capture utility for simple line-rate capture
 - hardware filtering supported
- Hardware timestamps available through Linux APIs & LinuxPTP & tcpdump (as of software 1.8.0)





ExaNIC X10:

- Flagship low-latency card
- 780ns raw latency* for small frames (64 bytes) using API
 - 710ns achievable with pre-push
- 880ns UDP / 930ns TCP latency* for small payload (14 bytes)
 - measured using unmodified sockperf benchmark
- * Quoted latency numbers are median appnetwork-app latencies (transmit + receive) on 3.5Ghz i7-3770K. Add ~40ns for newer Intel chips. Reduce slightly for faster clock rates. NOT STAC BENCHMARKS





New ExaNIC X10-GM variant:

- Hardware PTP grandmaster, serves PTP even with host down
- Built-in GPS receiver & high-stability OCXO (<0.2ppb)
- Pulse-per-second (PPS) output
 - Allows synchronizing other cards more accurately than PTP
- Also operates as a normal X10 network card





New ExaNIC X40:

- Now shipping
- 2 QSFP ports in a half-height form factor
 - 8 ports of 10G via QSFP break-out cables
 - 40G support in development (later firmware update)





Ultra-low latency switches with 5ns tapping, timestamping and PTP



ExaLINK Fusion, available as:

- Layer 1 only configuration
- Layer 1 + Layer 1.5 (aggregation/mux)
- Layer 1 + Layer 1.5 + Layer 2

5ns patch/tap, 95ns mux, 110ns switching

Not STAC Benchmarks

FPGA Development Kit available for custom applications



Network visibility with ExaLINK Fusion:



Ports in 5ns 'patch' mode, or 95ns 'mux' or 110ns 'switch' mode



Data aggregated and timestamped to 3ns resolution

• No need for an external server if using internal blade module and internal storage (more on this shortly)

Not STAC Benchmarks





New in ExaLINK Fusion:

• LLDP transmit & receive to easily determine connected hosts

admin@EXALINK-FUSION> show lldp neighbors				
Local Port	Chassis ID	Port ID	System Name	
B1	BC:5F:F4:5B:1F:8B	64:3F:5F:01:21:E4	exablaze001	
B2	9C:5C:8E:BF:82:86	64:3F:5F:01:22:F0	exablaze002	
B3	9C:5C:8E:BF:82:86	64:3F:5F:01:14:A0	exablaze002	
B4	F4:6D:04:8D:F6:41	64:3F:5F:01:29:FC	exablaze003	

- VLAN tagging and trunking
 - Supported in both mux and switch modes
- Multiple users and roles
 - Local or TACACS+





New in ExaLINK Fusion (continued):

• Latency statistics

admin@EXALINK-FUSION> port B1 show latency Total packets : 5858442

Percentile Latency

maximum	107 ns
99.99	107 ns
99.90	107 ns
99.00	101 ns
95.00	101 ns
90.00	101 ns
75.00	101 ns
50.00	95 ns
25.00	95 ns
minimum	95 ns

Not STAC Benchmarks





New in ExaLINK Fusion (continued):

- Built-in GPS receiver & PPS in/out on front panel
 - Retrofittable to existing units



- "Append" mode now supported for timestamps
 - Fusion now supports both "replace" mode (overwrite CRC) and "append" mode (insert timestamp before CRC)





New in ExaLINK Fusion (continued):

- Optional processor blade module
 - 6th Generation Intel Core i7 (Skylake) quad-core
 - Entirely customer dedicated, not virtualized
 - ExaNIC X40 providing eight 10G ports "on motherboard"
 - 2 x mSATA SSD for root filesystem and storage
 - M.2 PCIe x4 NVMe SSD for high bandwidth data storage
 - Allows "in-switch" capture or other user applications with no extra rack space usage
 - Cross-connected to FPGA module bay for custom applications





Thank you

More information: www.exablaze.com info@exablaze.com peter@exablaze.com

EEXABLAZE

Copyright © 2016 Exablaze Pty Ltd