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

Robert DeWitt, Jr
robert.dewitt@exablaze.com

STAC Innovation Roundup, Fall 2016

ExaNIC family: network interface cards with ultra low latency kernel bypass, hardware timestamping & PTP

ExaNIC X10
Lowest latency (~780ns*)

ExaNIC X10-GM
PTP grandmaster

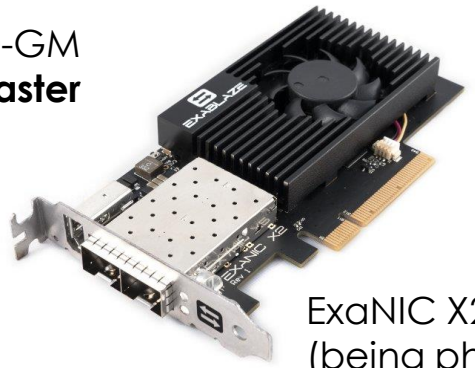

ExaNIC X2
(being phased out, replaced by X10)

ExaNIC X40
High port density (8x10G)

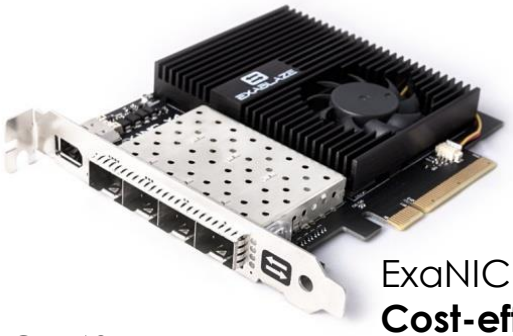

ExaNIC X4
Cost-effective 4 port capture

* NOT STAC BENCHMARK
Details on next slide

EXABLAZE

ExaNIC X10:

- Flagship low-latency card
- Shipping since October 2015
- 780ns raw latency* for small frames using API
  - 710ns achievable with pre-push
- 880ns UDP latency / 930ns TCP latency for small payload
  - measured using unmodified sockperf benchmark
- Further latency optimizations in development

* Quoted latency numbers are median app-network-app latencies (transmit + receive) on 3.5Ghz i7-3770K. Add ~40ns for newer Intel chips. Reduce slightly for faster clock rates. NOT STAC BENCHMARKS

ExaNIC X10-GM:

- 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 works as a normal X10 network card
- Now shipping

ExaNIC X40:

- 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)
- Now shipping

(available in various lengths)

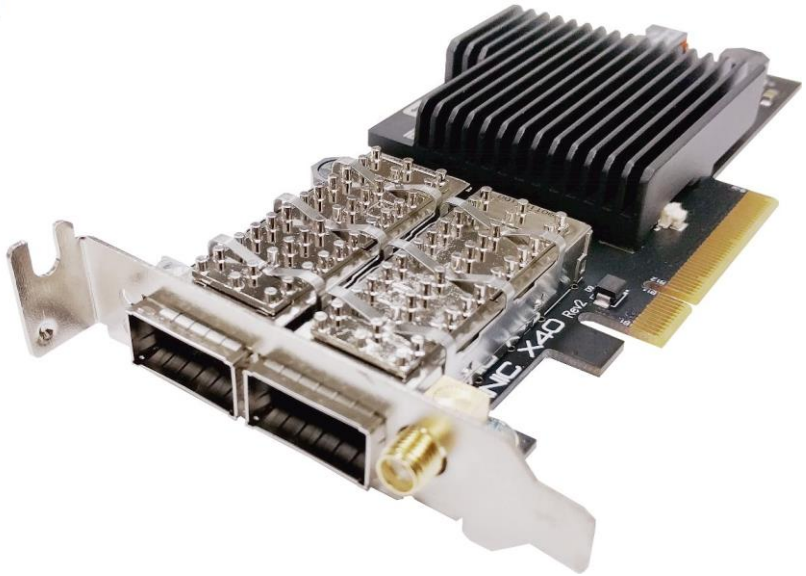

EXABLAZE

Other new features in current software release:

- Support for hardware timestamps with upstream LinuxPTP/ptpd (patched ptpd no longer required)
- Support for hardware timestamps with upstream tcpdump
- Support for hardware timestamps via exasock (our kernel bypass stack) using SO_TIMESTAMPING
- PPS out on X10/X10-GM/X40 cards

ExaLINK Fusion family: Ultra low latency Layer 1/1.5/2+ switches

Available as:

- Layer 1 only
- Layer 1 & Layer 1.5 (mux/timestamp/aggregate)
- Layer 1 & Layer 1.5 & Layer 2+ (full switching)

Modular device for upgradability and extensibility

- Pluggable line cards and internal datapath modules
- Optional 2nd FPGA module for customer applications
- Optional Skylake X86 module for customer applications

Skylake X86 module

- Allows in-switch capture or other control or data processing applications with no extra rack space usage
- 6th Generation Intel Core i7 (Skylake) quad-core
- Entirely customer dedicated, not virtualized
- Built-in ExaNIC X40 providing low latency connectivity to switch data plane
- 2 x mSATA SSD for root filesystem
- M.2 PCIe x4 NVMe SSD for high bandwidth data storage
- Cross-connected to FPGA module bay for custom applications

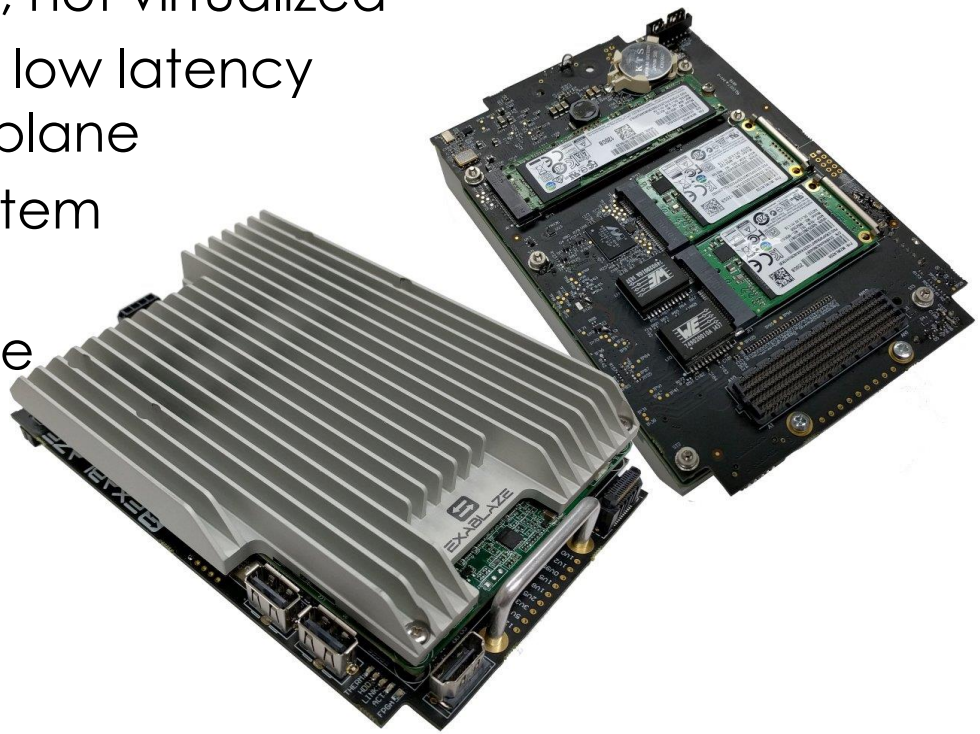

EXABLAZE

Continual latency reduction efforts

- Delivered in firmware updates on same hardware

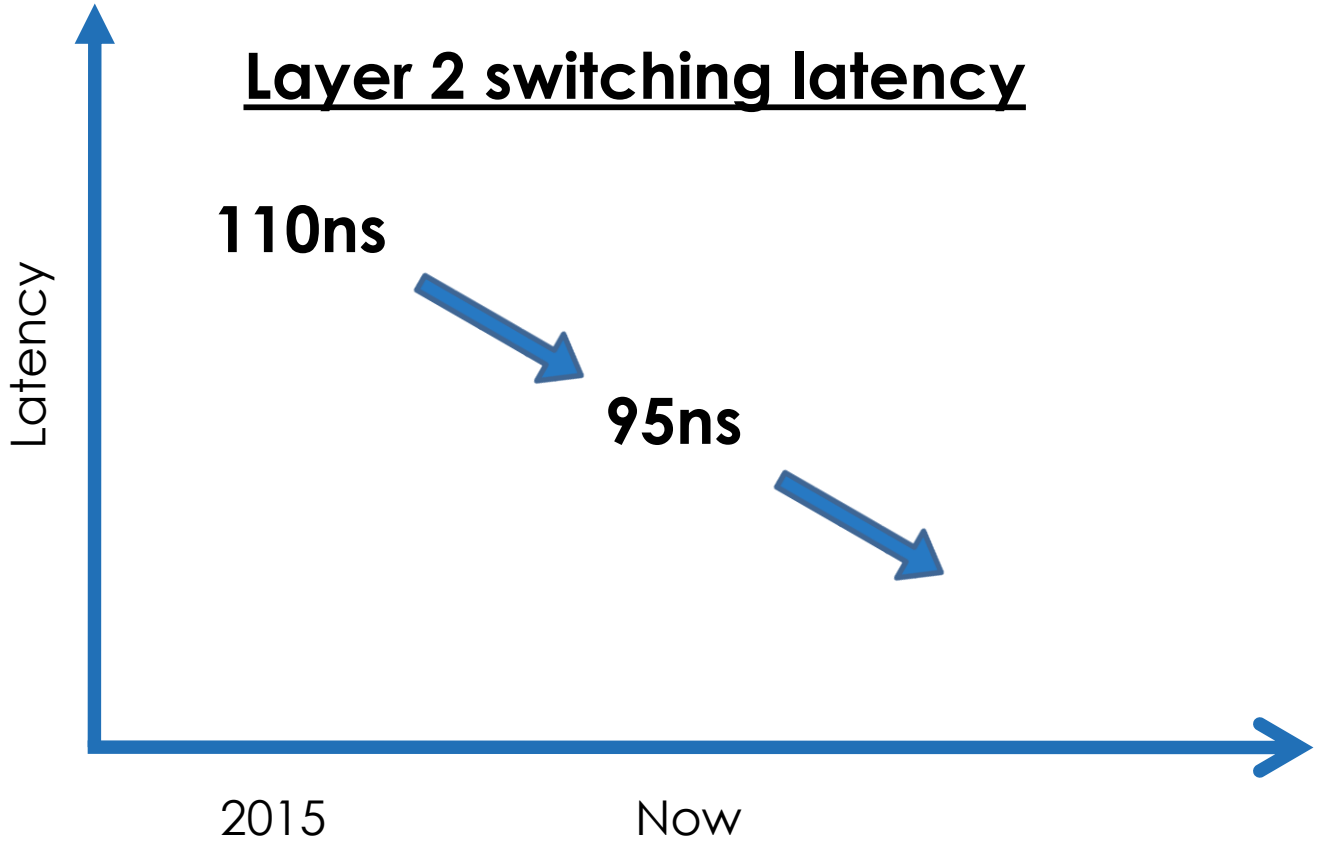

Not STAC Benchmarks

EXABLAZE

Continual latency reduction efforts

- Delivered in firmware updates on same hardware

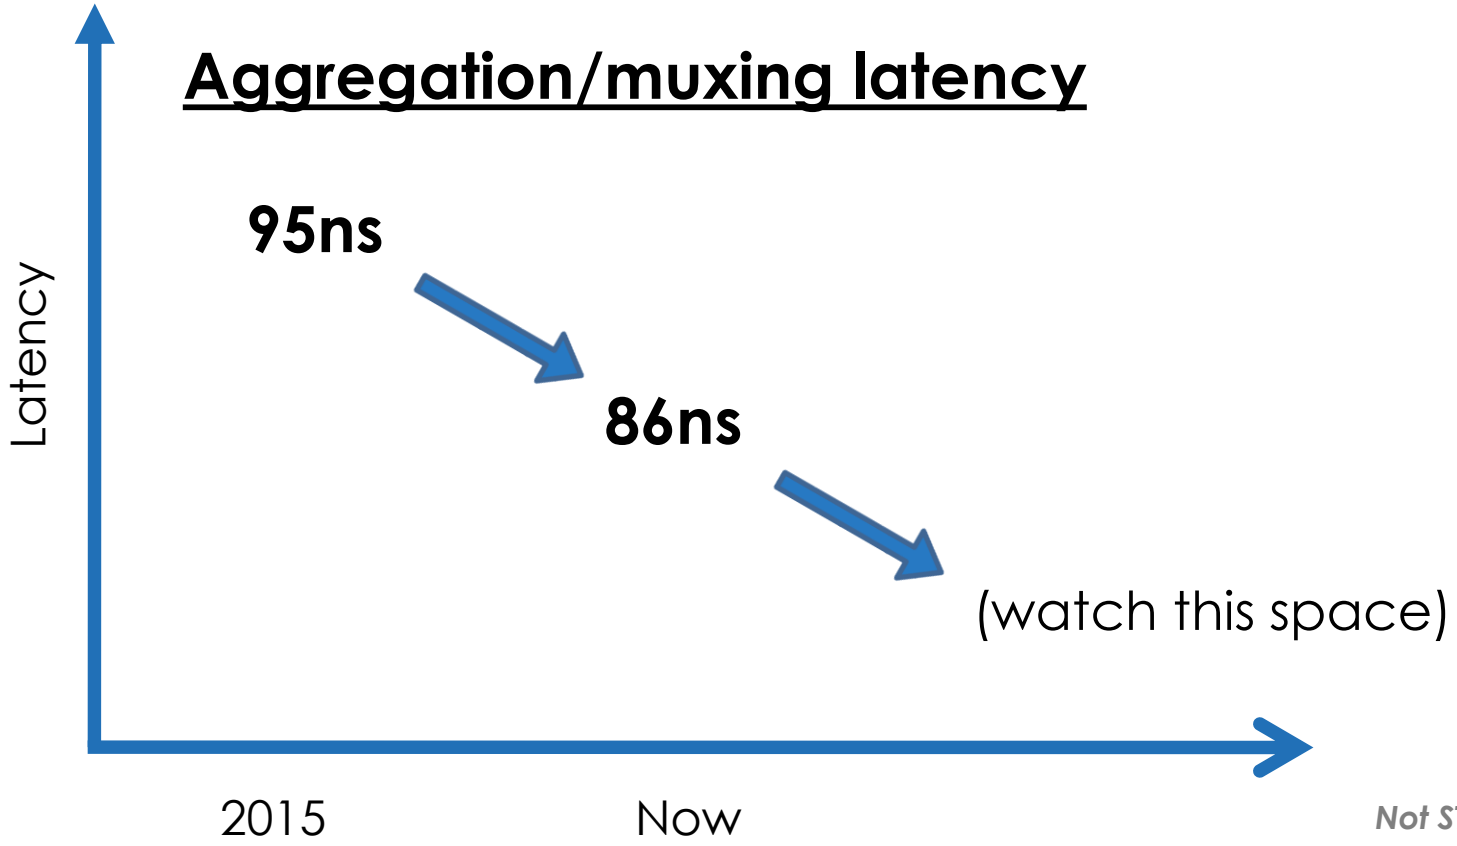

EXABLAZE

Other features shipping in current firmware:

- Layer 2 switching (MAC learning, etc.)
- VLAN tagging and trunking
- IGMP snooping, multicast filtering
- Multiple users and roles
- Latency statistics
- Timestamping to 2.9ns with choice of formats
- SNMP or JSON-RPC API for remote management

(… and more …)

Thank you

More information:

www.exablaze.com
info@exablaze.com
robert.dewitt@exablaze.com

EXABLAZE