EXALINKFUSION

Connectivity redefined.

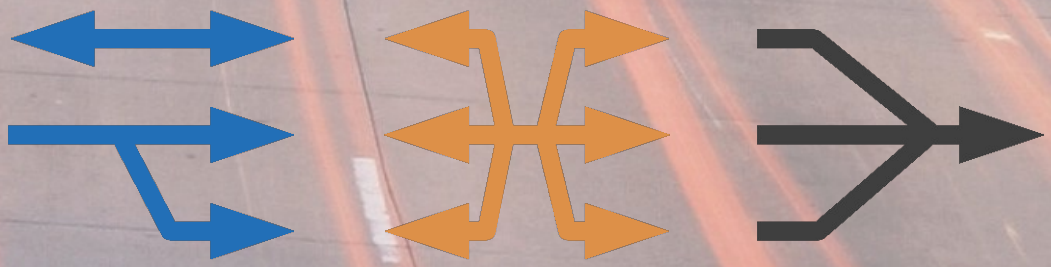

EXABLAZE

exablaze.com

EXALINKFUSION

Rethink your network architecture and application development using the ExaLINK Fusion.

Latency, reduced.

High frequency trading customers use the ExaLINK Fusion to seamlessly reduce the end to end latency in their networks.

Share a single network resource such as an order line by taking advantage of the Fusion's 95ns aggregation latency, many times lower than a traditional switch.

On the return path, optionally deliver order acknowledgements to clients by using layer 1 data distribution, replicating feeds at 5ns. In addition, use layer 1 fanout to distribute multicast to consumers with ultra low latency and negligable jitter.

This mode of operation reduces the full round trip time at the edge of your network to 100ns.

Patch and Tap

5 nanoseconds tapping and patching. Dynamically tap the feed on any input or output to any other port, replacing conventional optical taps. Dynamically patch any port to any other port to create bidirectional layer 1 links. Clock and data recovery circuitry ensures signal integrity is maintained throughout your network.

Aggregate

86 nanosecond aggregation. Aggregate multiple streams with extremely low and deterministic latency by bypassing MAC address lookups, sharing a single network resource fairly. Optionally, use layer 1 on the return path to further reduce latency and jitter and guarantee response times.

Switch

95 nanoseconds port to port latency with full layer 2 switching. The industry's lowest latency cut-through layer 2 switch. Supports ultra-low latency multicast and broadcast.

Redefine

It's your network, build it your way. The modular ExaLINK Fusion can be shipped with a Xilinx Ultrascale FPGA module that allows you to completely redefine the way your network operates. The layer 1 technology in the Fusion allows you to create dynamic circuits between any and all front panel ports and the FPGA. Build custom packet processing engines. Filter traffic. Switch based on packet content. Experiment with new protocols. The Fusion makes all of this and more possible.

Dual module bays

Fit two FPGA modules or x86 processors. Connect any high speed transceiver to the front panel by creating dynamic layer 1 circuits.

Redundant power supplies

Maximum reliability. Dual power supplies and fan modules protect against downtime.

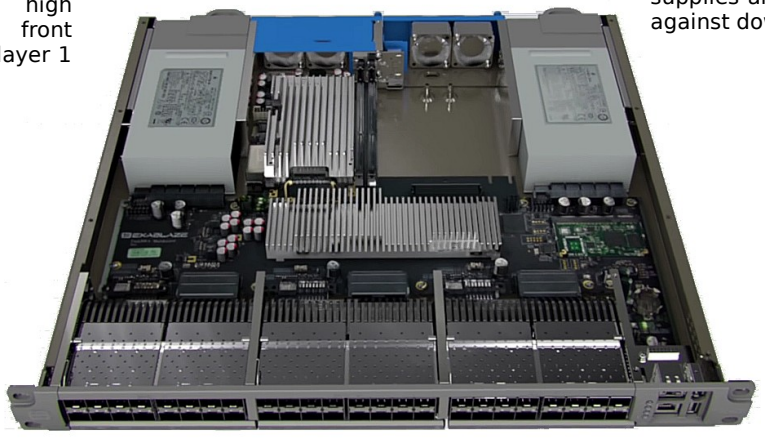

Reconfigurable front panel

Three line card bays provide front a flexible front panel, allowing for up to 48 SFP+ ports.

Layer 1 crosspoint

The heart of the Fusion. A layer 1 crosspoint allows you to dynamically create circuits between any front panel port or any internal module bay

Smart management

With flexibility comes great complexity. The management interface on the Fusion provides an intuitive way to configure the device.

Multi-purpose trading toolkit

A typical trading environment consists of multiple servers, interfacing to one another and to an exchange via order lines and market data feeds. The ExaLINK Fusion can replace existing top of rack switches, optical taps and timestamping technology.

In the example shown, multiple trading servers receive redundant multicast market data via layer 1 taps from the exchange. These servers make order entry decisions based on this data and send them to an aggregation port on the ExaLINK Fusion, where it is on-routed to the exchange.

On a designated port the ExaLINK Fusion inserts nanosecond accurate timestamps into all frames and sends them to a logging server. This feed can be used to reconstruct trading activity and be used in post analysis.

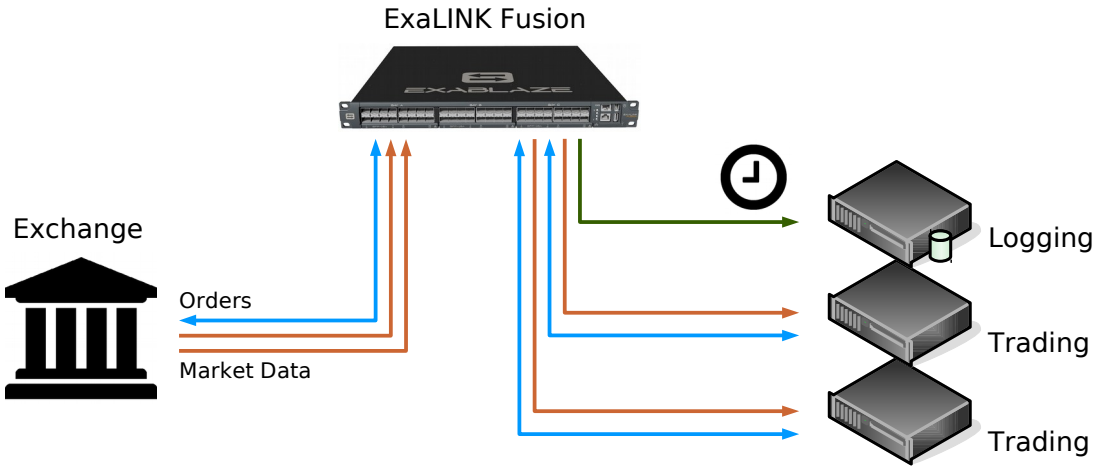

Why distribute data at layer 1?

When latency, jitter, or strict ordering matter, layer 1 fanout technology has several advantages.

Market data fanout allows you to 'tap' the market data feed prior to it reaching any local layer 2 hops. This instantly reduces latency by several hundred nanoseconds.

Layer one distribution also provides an upper bound on jitter. Since data is replicated at the signalling layer, there is no possibility of output contention or buffering. This means that the port to port latency is very deterministic.

Finally, as input and output queues are removed from the system, routing via Layer 1 ensures all end nodes receive data in the same order. This has advantages in many use cases.

Network application development platform

The ExaLINK Fusion architecture is ideal for all kinds of custom network application development. The central layer 1 crosspoint allows traffic from any of the line cards to be patched through to the internal module bays. These bays can be populated with FPGA and x86 processor modules and the connectivity configured dynamically.

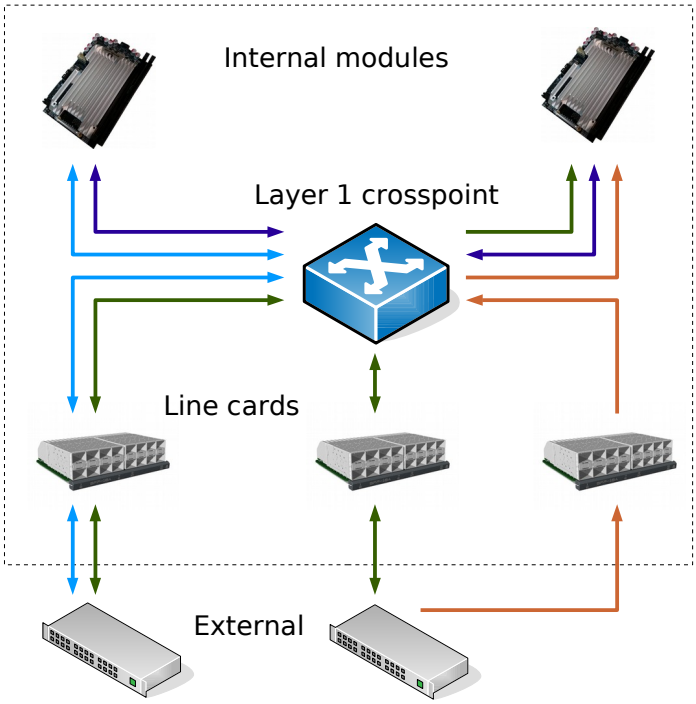

As an example, it is possible to transparently monitor traffic flows by replicating both sides of a bidirectional front panel 'patch' to an internal module. The user's FPGA application can then analyse, filter and act on the tapped data.

Using the layer 1 crosspoint to patch the internal modules to the front panel ports allows bidirectional communication between the FPGA application and the outside network. This allows the user to implement custom switch architectures, limited only by their imagination.

Furthermore, when both internal modules are installed they can also communicate via the layer 1 crosspoint. This allows any number of high speed links between the two module bays.

All 48 front panel ports can be dynamically patched to either internal module bay.

SPECIFICATIONS

Latency

Measured front panel port to port:

- 5 ns (Layer 1 tap/patch)
- 86 ns (Layer 2 aggregation)
- 95ns (Layer 2 switching)

  Further latency reductions will be delivered as firmware upgrades

Timestamping

- Any flow through the device can be mirrored & timestamped to nanosecond resolution
- Multiple keyframe formats supported
- Onboard GPS receiver for time and PPS
- Pulse per second input or output for time synchronization
- 0.2 parts per billion holdover drift with loss of external timing (with optional oscillator upgrade)
- Local time synchronized via GPS, NTP or PTP

Switch/Aggregation Features

- IGMP snooping, multicast filtering
- Static routes (MAC)
- VLAN trunking/tagging
- LLDP, including layer 1 paths
- LCAP
- Port stats including queuing latency

Physical, power, cooling

- 19” 1RU, rack mount
- Weight 11kg (24lbs)
- Dual, hot-swappable supplies
- Standard: AC 90-264V, 47-64 Hz, included IEC C13-C14 cables
- Optional: DC 40-72V
- Maximum consumption: 150W
- Dual hot-swappable fan modules
- Optional airflow direction

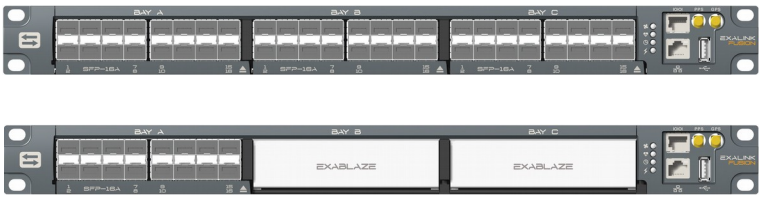

FPGA development

- FPGA board installable in both module bays
- Xilinx Ultrascale KU115 FPGA with 52 high speed transceivers
- 48 transceivers can be directly connected to front panel ports via internal layer 1 crosspoint
- 4 transceivers available on internal header, for communication with adjacent module
- 288 Mbit QDR4 SRAM onboard, 2 x DDR4 DIMM slots

Optional x86 CPU Module

- Intel Core-i7 Skylake CPU, 4 cores @ 2.8GHz
- 32GB DDR4
- Dual mSATA, up to 1TB ea
- M.2 PCIe SSD, 512GB for line rate capture to disc
- Onboard low latency ExaNIC X40

Connectivity

- 3 x 16 SFP+ line cards, up to 48 ports
- SFP+ Fiber (10GBASE-SR, 10GBASE-LR, 10GBASE-LRM, 1000BASE-SX, 1000BASE-LX)
- SFP+ Copper Direct Attach

Management

- Ethernet 10/100M management port
- Industry standard serial port
- Intuitive CLI via SSH / Telnet / Serial
- External JSON RPC scripting interface
- TACACS+ and multi user support
- ACL's on management interface
- SNMP, local and remote syslog
- Firmware upgradeable via SFTP, TFTP or USB
- Custom FPGA firmware loaded and managed by onboard processor

EXABLAZE