

The triple-FPGA 7130 E Series: optimised for ultra-low latency, inter-FPGA communication

An Introduction to the Arista 7130E Series

The Arista 7130E Series supports three Xilinx UltraScale™+ FPGAs in a one rack unit (RU) form factor. The UltraScale™/UltraScale™+ series from Xilinx used in the 7130E devices are a significant step forward from previous solutions.

Key improvements over the Arista 7130K Series include:

- three times as many logic cells
- eight times as many DSP slices
- six times as much memory
- significantly lower-latency Ethernet transceivers
- the ability for logic to be clocked higher than ever
- significantly more IO bandwidth
- all the above while consuming less power than the previous generation

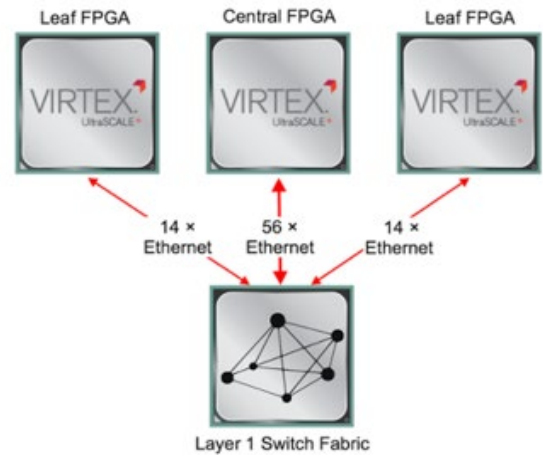
The 7130E Series features 48 10GbE SFP+ ports. Each device comes with a full Layer 1 switch fabric and contains a range of management and monitoring features including: per-port Ethernet statistics, monitoring telemetry and an x86 processor running MOS or EOS, a Linux-based operating system, offering enterprise-grade device and application management.

When ordered with three FPGA's, there are multiple dedicated communication paths between the FPGAs and the Ethernet network offering flexibility and the ability to optimise communication latency. As there are currently a number of permutations of the 7130E Series with slightly different physical characteristics, the characteristics described below refer to the 7130-48E variant (i.e.: 48-port, 1 RU devices with either 3 Ultrascale or Ultrascale + FPGAs).

Network Connectivity

The “Central” FPGA is connected to the Layer 1 switch fabric with 56 Ethernet transceivers. Each “Leaf” FPGA is connected to the Layer 1 switch fabric with 14 Ethernet transceivers. The Layer 1 switch fabric is also connected to the front-panel ports of the device for external connectivity. This Ethernet connectivity provides enormous flexibility in communication between the FPGAs and the external network.


Moreover, the Layer 1 switch can replicate data from incoming ports to multiple destination ports. For example, incoming UDP/IP multicast data coming into the Layer 1 switch from the network can be delivered to multiple transceivers (if desired) on one, two or all FPGAs in ~3 ns. One or multiple links can be configured between the FPGAs’ ports in any required combination. Any of the FPGAs can therefore communicate with each other, or any of the front-panel ports via an Ethernet interface with a latency of ~50 ns (depending upon which MAC + PHY is used).



The 7130 E Series Parallel Bus (MMP)

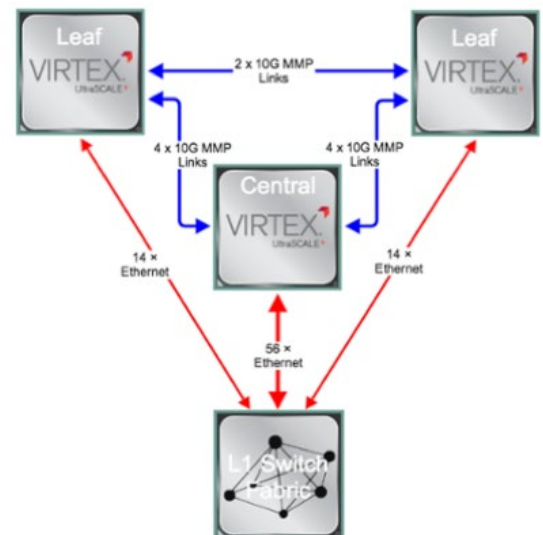
The 7130E Series parallel bus is a bus leveraging a large number of parallel IOs. Arista has released the Parallel Bus IP cores package allowing client applications to communicate between FPGAs with the absolute minimum of latency. The first IP core offers a unidirectional bandwidth of 10 Gbps per link and provides an AXI4-stream interface.

There are four links connecting each Leaf FPGA to the Central FPGA and two links connecting the two Leaf FPGAs together. The direction of each parallel bus link can be configured individually. The key advantage that the parallel bus offers over Ethernet is that, being a parallel rather than a serial bus, its latency is significantly lower. The 7130 Series parallel bus offers one-way transfer latencies between FPGAs of ~8 ns.

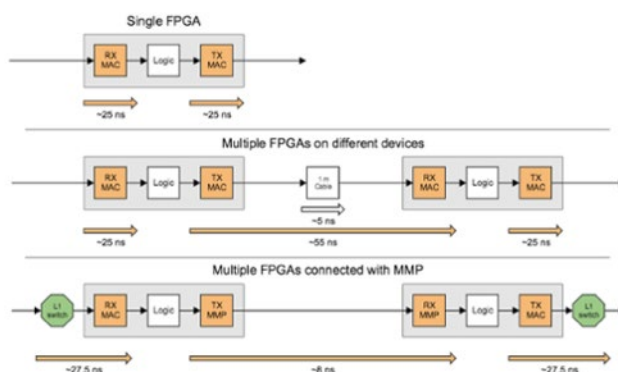


Inter-FPGA communication latencies compared

The visual latency breakdown below illustrates the latency differences (also refer to the graphic to the right) compared to a single FPGA reference. The latency benefits of inter-FPGA communication via parallel bus are significant in comparison to Ethernet.



In the following example, the latency savings of having FPGAs communicate locally via parallel bus links (~8 ns) versus having FPGAs on different devices communicating via 10GbE (~55 ns) are substantial. The parallel bus offers an 85% latency reduction compared to 10GbE.



What about PCI Express?

Plugging an FPGA PCI Express board into a server remains popular. The rationale is sound; software running on the server's CPUs has an efficient, low-latency path with which to communicate with the FPGA. This rationale however is predicated on the communication between the host application and the FPGA being latency-critical, otherwise why not simply use Ethernet to communicate between the host and CPU?

Architecturally, this may be due to a lack of FPGA capacity. It is not always possible to run all the latency-critical code on the FPGA, with part of it having to run on the host. The current generation of Xilinx UltraScale™/UltraScale™+ FPGAs now has sufficient capacity to implement, in many cases, the entirety of the latency-critical portion of a transaction on the FPGA (or FPGAs in the case of the 7130E Series).

The non-latency sensitive part of the trade can still reside on the server but now communicate with FPGAs over the Ethernet network with a significantly simpler interface. There is also the added benefit that all communication between the FPGA and the controlling host can be monitored and timestamped as it traverses the network. This can be extremely useful to meet requirements such as record-keeping and compliance.

Example use cases for the triple-FPGA 7130E Series in electronic trading

- One "stable" FPGA trading image can be retained with the remaining two FPGAs available for proving out strategies or experimentation
- Risk checks can be run on a different FPGA from the actual trading algorithms whilst communicating with significantly lower latency than previously possible
- Higher FPGA density means that multiple trading strategies can be implemented in as little as 1 RU, with the option to segregate them across three FPGA images
- Three FPGAs contain over 1 Gb of on-chip RAM which is a real alternative to externally connected QDR-II/QDR-II+ (using the Xilinx UltraScale™+ VU9P as an example)

Summary

- The Arista 7130E Series offers devices containing three Xilinx UltraScale™+ FPGAs.
- The 7130E devices offer the same enterprise-grade management and core features offered by all Arista devices
- All three FPGAs are connected to the device's Layer 1 Ethernet switch with as many as 84 transceivers
- The FPGAs are connected to each other via one (leaf-to-leaf) or two (leaf-to-centre) parallel bus links offering applications inter-FPGA communication in ~8 ns
- In an electronic trading use case benefiting from multiple-FPGAs, the 7130E Series offers significantly lower latency communication between FPGAs than Ethernet-connected FPGAs

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390

Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2021 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. Feb, 2021