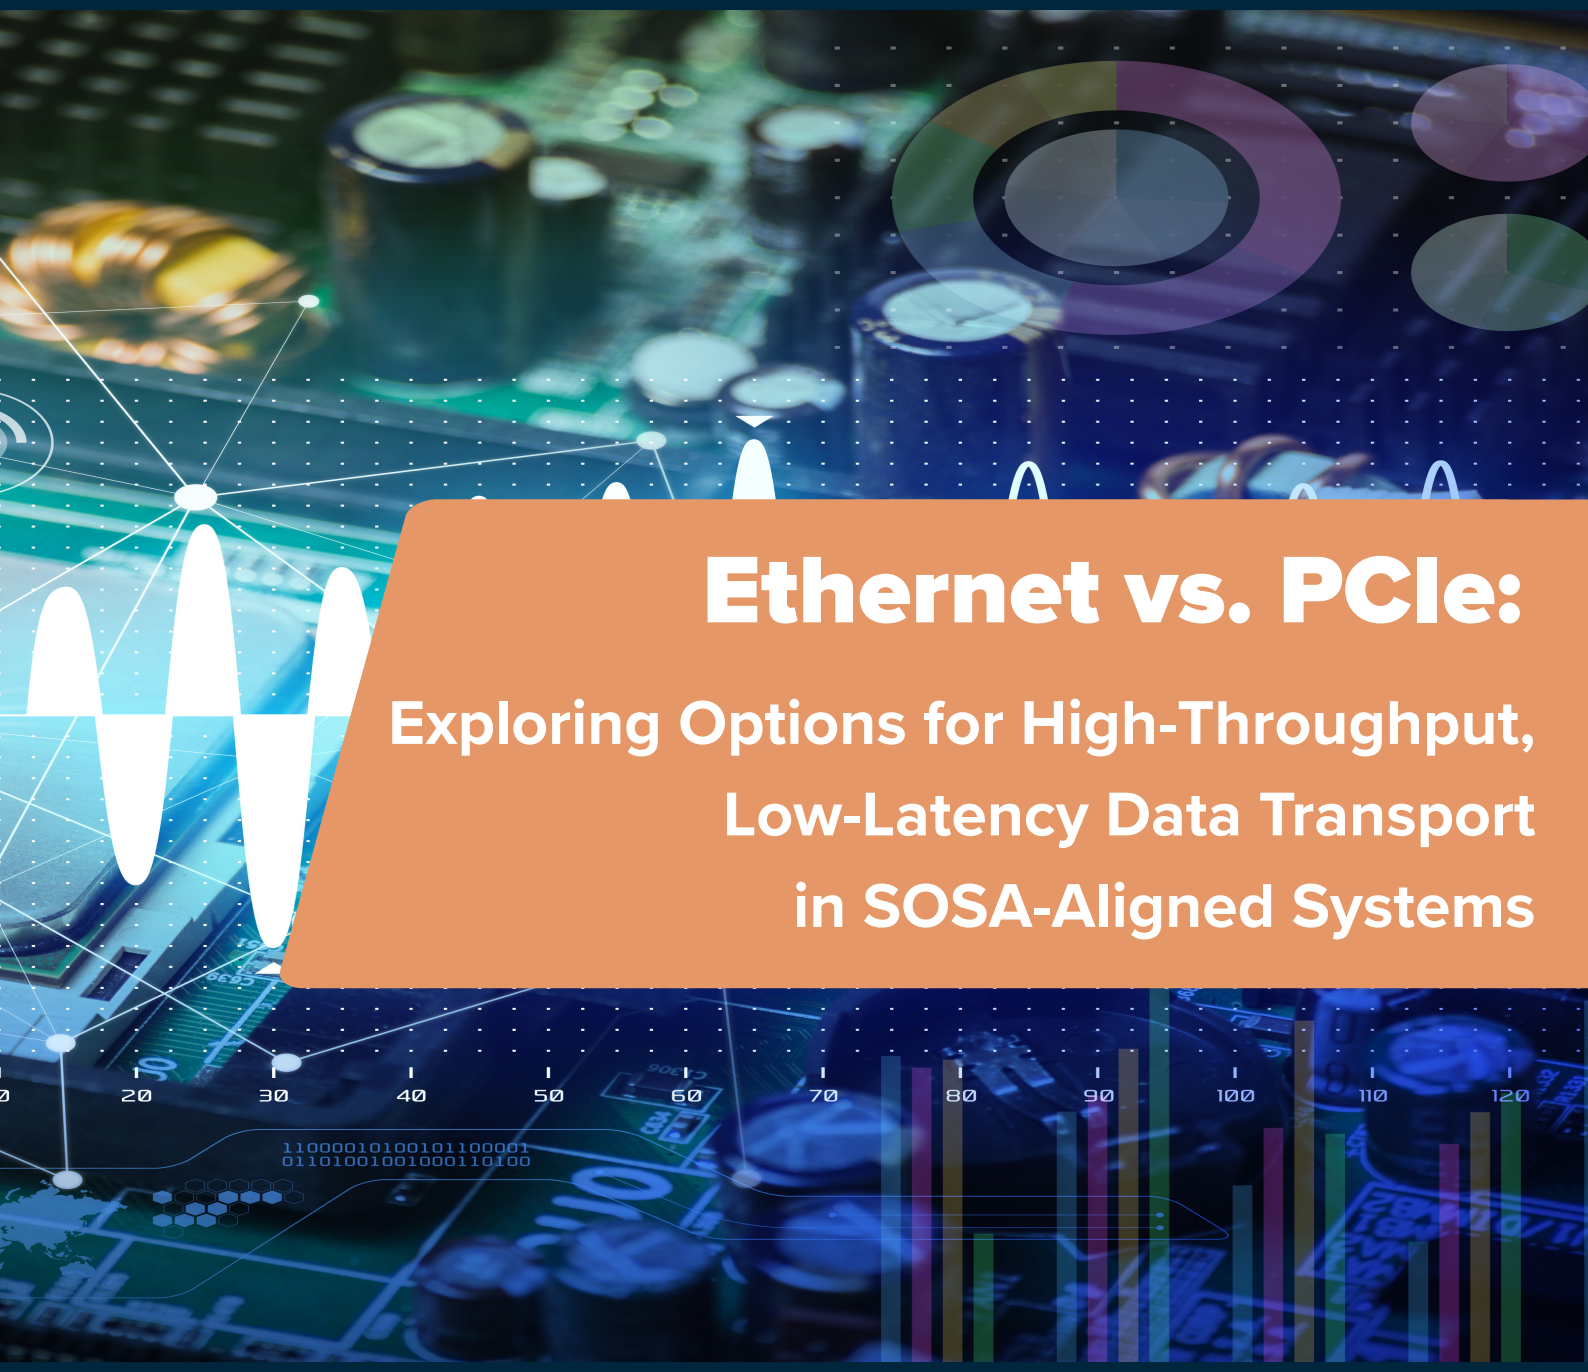

Ethernet vs. PCIe:

Exploring Options for High-Throughput, Low-Latency Data Transport in SOSA-Aligned Systems

EPIQ SOLUTIONS

Table Of Contents

As RF channel bandwidth increases, along with the number of channels being used, the systems capable of processing these signals also exponentially increase in complexity.

The Growing Challenge of Ever-Increasing Signal Bandwidth

The modern wireless landscape continues to evolve at a frenetic pace. Commercial wireless standards have come a long way since the days of 2G cellular where voice and text messages ruled the day and the needed spectrum fit within a few hundred KHz per channel. This is in stark contrast to 5G cellular which has a channel bandwidth capacity of 400 MHz per antenna, multiplied by four or more antennas, all of which make the requisite digitized data rate for these signals staggering. System architects developing solutions to process modern wireless signals are pushed to keep pace with this growth, and every aspect of the system must be up to the task.

In this whitepaper, we'll explore how both Ethernet and PCIe can be used as the transport to connect SOSA-aligned plug-in cards (PICs).

Wireless Standard	Max Bandwidth per Antenna	Typical Quadrature Sample Rate	# of Antennas	Baseband I/O Bandwidth
GSM/2G	200 KHz	541.667 Ksps	1	21 MBps
UMTS/3G	5 MHz	7.68 Msps	1	30.72 MBps
LTE/4G	20 MHz	30.72 Msps	1-4	122.88 MBps – 491.5 MBps
802.11n	40 MHz	40 Msps	1-4	160 MBps – 640 MBps
802.11ac	160 MHz	160 Msps	1-4	640 MBps – 2560 MBps
5G-NR	400 MHz	491.52 Msps	1-4+	1964 MBps – 7856 MBps+

Data transport rates for modern digitized RF signals are increasing, creating challenges around how to digitize, transport, and process these signals.

Like wired networks themselves, wireless threats evolve with available technology, and the systems in place to mitigate those threats have to advance, as well. For engineering teams serving government organizations that need to understand and interact with their wireless environments, managing and processing these high data rate signals in SOSA-aligned systems is a growing challenge.

Electronic warfare (EW), spectrum monitoring, and signal intelligence (SIGINT) applications are requiring multi-channel, wideband software-defined radio (SDR) processing solutions to keep pace with technology. This, in turn, is generating a lot of data that needs to be processed as efficiently as possible. SOSA-aligned platforms bring together a collection of hardware and software elements that exemplify the modular open system architecture (MOSA) concept. This includes RF transceivers, FPGAs, GPUs, and CPUs, with each hardware element integrated onto a SOSA-aligned 3U VPX compatible plug-in card (PIC). With increases in antenna count as well as RF channel bandwidth, the interconnection between the PICs through the backplane becomes a critical system level consideration.

To connect SOSA-aligned plugin cards or PICs, there are two transport options: Ethernet and PCIe. Each option has its advantages, and we will help chart a course for system architects to determine which is the best fit for their system.

The Two Dominant Methods for Transporting Samples Today – PCIe and Ethernet

Two avenues for transporting data between PICs have emerged in the SOSA community – Ethernet and PCIe. While most people think of Ethernet as a cabled interface that uses RJ45 connectors or SFP/SFP+ connectors, in the context of SOSA, Ethernet signaling is provided between PICs through the backplane without additional cabling. PCIe, on the other hand, has its roots in the computer peripheral universe, where it has been used to interconnect a CPU with peripherals such as GPUs, hard drives, I/O cards, and more on a standard motherboard. These peripherals often need to exchange high rates of data with low latency and minimal overhead.

Each interface comes with strengths and weaknesses for transporting data around in a SOSA-aligned system. The question facing Department of Defense (DoD)-focused RF system designers is which method offers the most efficient avenue for transferring and processing all of these signals?

Historically, Ethernet has had an inside track over PCIe for a few reasons. As data centers became the hub of everything that happens online, Ethernet was tasked with moving all of that data around local area networks (LANs). Protocols were developed around Ethernet, new systems were expected to support it, and if part of your job was networking, chances are you've developed an understanding of how to use it. Through this exposure, Ethernet was pushed, pulled, and tested to the point where it became a very robust method for transporting data between physically separated devices.

Specific to RF systems, there are some cases where it is necessary to have a radio head physically separated from the backend processing. In these cases, Ethernet can be a good fit.

In these cases, Ethernet is a logical choice. Some RF transceivers provide an antenna interface on one end, and an Ethernet port on the other, allowing devices to connect to them over a standard network interface. The VITA 49 Radio Transport (VITA 49 aka VRT) specification helped further define how RF transceivers could format and transport digitized I/Q samples from an RF transceiver to a computing device. And while VITA 49 is technically transport-agnostic, most implementations utilize Ethernet, with the associated IP and UDP protocol layers on top of it, for transporting digitized spectrum between devices because of its proven robustness and reliability.

As a result, many software-defined radio (SDR) platforms have been architected with Ethernet as the interface between the RF transceiver and the compute/processing hardware. But for systems where both the RF transceiver and processing hardware live on a tightly coupled common backplane, Ethernet isn't necessarily the optimal solution. PCIe offers some compelling advantages in terms of reduced transport latency, minimized processing overhead, and optimized general system efficiency. This is especially true for SOSA-aligned platforms, which consist of a series of PICs interconnected using a common backplane in a chassis.

Architectures Typical of 3U VPX SOSA-Aligned Systems

SOSA-aligned RF systems generally use one of the following architecture options for signal processing:

Option #1: An RF signal goes from the RF transceiver into a local FPGA on the same PIC, where all of the required processing is performed

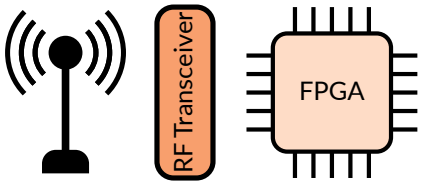

Option #2: An RF signal goes from the RF transceiver into a local FPGA on the same PIC, and then transported to a CPU on another PIC connected to the backplane for processing

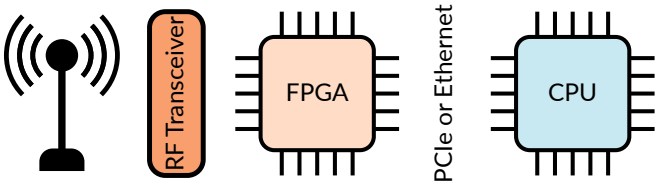

Option #3: An RF signal goes from the RF transceiver into a local FPGA on the same PIC, and then transported to a GPU on another PIC connected to the backplane for processing

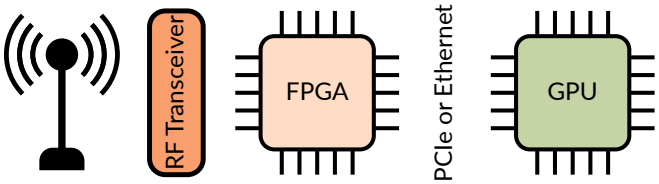

Option #4: An RF signal goes from the RF transceiver into a local FPGA on the same PIC, and then transported to an FPGA on another PIC connected to the backplane for processing

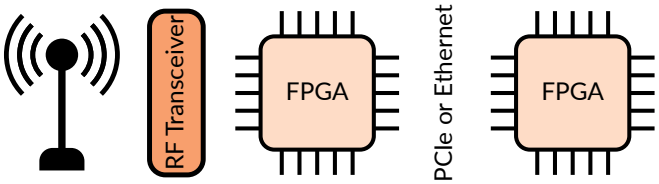

With Option #1, getting the radio signal into the FPGA is as far as it ever needs to go because the FPGA can process the data and reduce the sample rate needed to be transferred into the rest of the system. But for Options 2, 3, and 4, a significant rate of data transport is required to move the digitized radio spectrum from the local FPGA to the next processing PIC in the system,be it a CPU, GPU, or another FPGA. This is where a decision needs to be made between using Ethernet or PCIe.

Both Ethernet and PCIe are supported in current SOSA-aligned card profiles. Ethernet is typically supported through the "data plane" using four high speed serializer/deserializer (SerDes) transceiver pairs, accommodating up to 10 Gbps per SerDes pair. This yields a total transport of 40 Gbps. More recently, vendors of SOSA-aligned PICs and backplanes have been able to extend signal integrity performance of both the backplane and connectors to support 25 Gbps per SerDes, increasing the total theoretical aggregate throughput across four pairs to 100 Gbps.

Similarly, current SOSA-aligned card profiles support up to eight SerDes transceiver pairs on the "expansion plane," allowing for systems to utilize at least eight lanes (x8) of PCIe Gen4 (16 Gbps per lane) between the PIC and the backplane. This results in a theoretical throughput of 128 Gbps. The PCIe Gen5 specification has been approved, doubling the lane rate again to 32 Gbps. With the spec approved, the production of hardware is expected to be available in the coming years. And while lane rates up to 25 Gbps have already been validated for 100 GbE, additional work will be required to determine if extending this to 32 Gbps is plausible.

Transport Type	Transport Data Rate per SerDes Transceiver	# of SerDes Transceivers	Total Transport Throughput
PCIe Gen3	8 Gigabits/sec	8	64 Gigabits/sec
PCIe Gen4	16 Gigabits/sec	8	128 Gigabits/sec
PCIe Gen5 (spec approved but hardware not available yet)	TBD (at minimum 25 Gigabits/sec, though spec allows 32 Gigabits/sec)	8	TBD; likely to be 256 Gigabits/sec
10 Gb Ethernet	10 Gigabits/sec	1	10 Gigabits/sec
40 Gb Ethernet	10 Gigabits/sec	4	40 Gigabits/sec
100 Gb Ethernet	25 Gigabits/sec	4	100 Gigabits/sec

Table 1: Comparison of throughput for standard PCIe and Ethernet configurations

In the aforementioned SDR architecture options, the benefits of PCIe can be significant. Digitized radio spectrum can be transported from an RF transceiver card to the next processing card element (FPGA, GPU, or CPU) using direct memory access (DMA), ensuring that latency and overhead are minimized. The host CPU where the PCIe root complex resides in the system is responsible for setting up the DMA transactions over PCIe, allowing one "source" PCIe endpoint to communicate with another "destination" PCIe endpoint directly. This destination PCIe endpoint could be either a CPU, GPU, or FPGA. Once the DMA transaction has started the data transport process, the host CPU is not involved, and thus can be used for other processing purposes.

For systems built around VITA 49, it is quite plausible to leverage this standard using PCIe to get the best of both worlds. The DMA packets transported over PCIe can be formatted

For tightly coupled systems with compute elements connected over a common backplane such as signal processing in signal intelligence applications, PCIe provides more flexibility, lower latency, higher data transport throughput, and improved system efficiency.

such that the standard VRT header is inserted at the beginning of each data packet, and the I/Q portion of the data packet would follow the header as specified in the VITA 49 standard. Context packets and control packets can also be utilized in a similar manner, with each being sent over PCIe DMA just like data packets. The main difference between how VITA 49 would be transported in Ethernet versus PCIe comes down to the packet size. Ethernet has substantial flexibility in the packet size, supporting small Ethernet frames of just 64 bytes all the way up to so-called "jumbo" frames that can be as large as 9,000 bytes. On the other hand, PCIe generally operates on fixed size packets that can not easily be changed on the fly. Further, DMA transactions operate most efficiently when transporting data into the native memory page size of the compute platform (typically 4 kilobytes per page for modern operating systems). But for RF processing systems, run-time flexibility in the packet size doesn't typically yield value, since digitized samples are continuously streaming. Thus, PCIe can be combined with VITA 49 to provide an efficient transport while maintaining a familiar data/control standard.

For system architectures leveraging NVIDIA GPUs to perform signal processing tasks, NVIDIA provides GPUDirect RDMA, a software solution to enable direct DMA transport from a PCIe source endpoint (such as an FPGA) to a PCIe destination endpoint in a GPU. At the 2019 GNU Radio Conference, Epiq Solutions presented the results of data rate transport improvements between FPGA and GPU using GPUDirect RDMA compared to a typical architecture with a CPU in the middle. This analysis provided insight regarding the substantial throughput enhancement that can be achieved with GPUDirect RDMA.

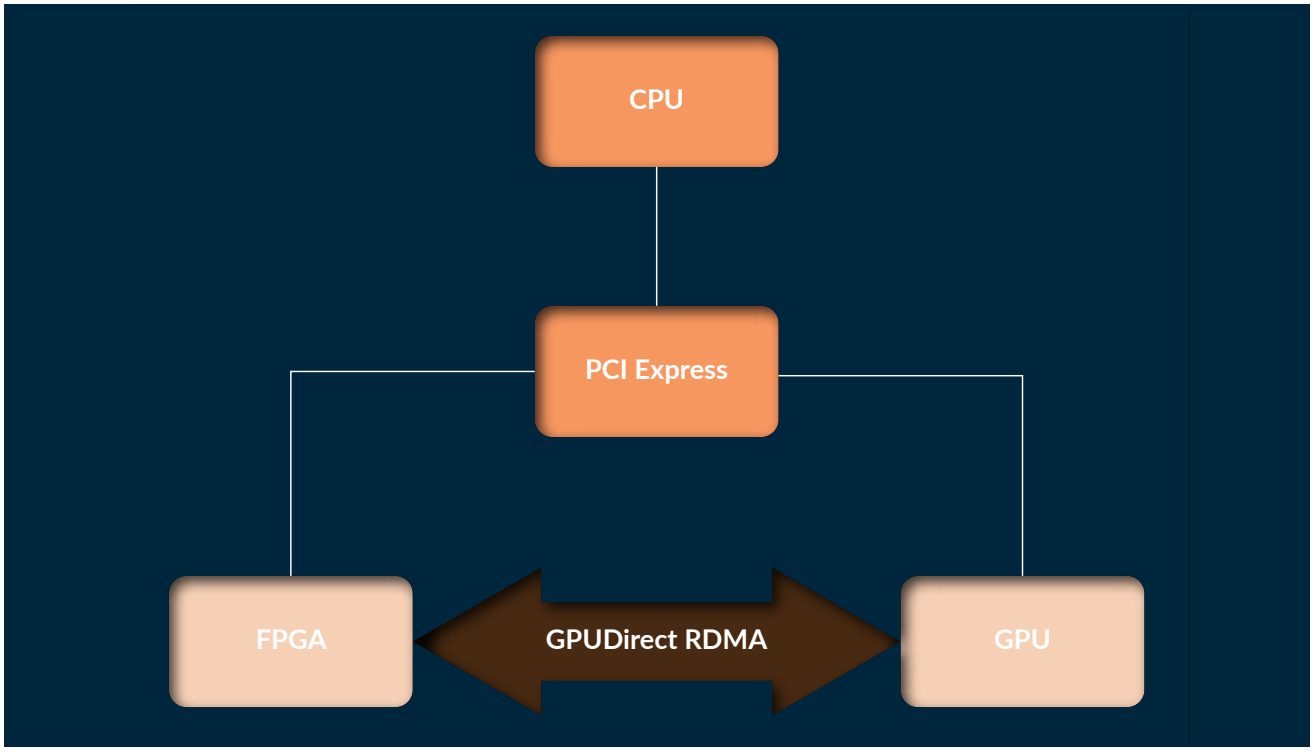

GPUDirect RDMA standard architecture

Ethernet, on the other hand, needs to be processed through a network protocol stack originally designed to be implemented in software, with each layer of the stack (Ethernet, IP, UDP, VITA 49) adding to the processing overhead and increasing latency. Despite its drawbacks, this protocol processing is well suited for software running on a CPU, and is typically implemented by the underlying operating system kernel itself (Windows, Linux, etc). However, in scenarios where the destination computing device is an FPGA or GPU, standard network protocol processing can present challenges. There have been efforts to utilize software packages such as the Data Plane Development Kit (DPDK) to allow for more customized network protocol processing to be implemented in userspace outside of the operating system's kernel, and in some cases DPDK has shown significant improvement compared to the typical kernel-based network stack.

PCIe vs Ethernet: At-a-Glance

PCIe	Ethernet
Designed to be as efficient as possible with minimal overhead	Designed for data centers and network-focused devices with distributed components
Architected for a backplane, not physical separation	Architected for connecting over distance
Rapid evolution to support increased capability in commercial computing systems	Rapid evolution to support increased capability in commercial networking systems
Compatible with VITA 49	Strong association with VITA 49
Run time packet size adjustment not trivial	Run time packet size adjustment trivial
Extremely low latency, especially compared to Ethernet	Reputation as “the networking standard”
Can support both copper and optical interfaces (though copper is far more common)	Can support both copper and optical interfaces (both are commonplace)
Seamless interface with CPUs, GPUs, etc.	

Conclusions and Future State

Both Ethernet and PCIe are viable options for use in SOSA-aligned systems. Each transport mechanism supports comparably high data rates per transceiver as shown in Table 1. However, with a tightly coupled system consisting of a series of PICs interfaced to a common backplane in a single SOSA-aligned chassis, PCIe can be a more efficient transport requiring much less overhead and with much lower latency transport. This mirrors the same architecture heavily leveraged in modern compute platforms that reside within a single chassis, where PCIe provides the crucial interconnect between CPU, GPU, hard drives, and more. Additionally, for applications where minimizing latency and processing overhead is essential, PCIe has significant advantages over Ethernet, since there is no protocol stack to manage, and DMA transactions can be set in motion with minimal overhead.

Epiq Solutions has a portfolio of SDR solutions that leverage PCIe as well as Ethernet. If you're interested in learning more about how PCIe or Ethernet can be leveraged in your SOSA-aligned SDR platform, please reach out to us.

LEARN MORE

EPIQ SOLUTIONS

epiqsolutions.com