Solution Brief

5G and Network Functions Virtualization (NFV)
Intel FPGAs

Intel's Accelerated Virtual Cell Site Router Solution on an Intel® FPGA-Based SmartNIC N6000-PL Platform Helps Communication Service Providers (CoSPs) Increase Monetization of Their Services

Intel adds an accelerated virtual cell site router (vCSR) with O-RAN-compliant precision timing on an Intel Agilex® 7 FPGA to its portfolio of virtualized workload solutions, enabling CoSPs to deploy and manage their 5G networks faster while achieving an optimal balance between network performance and cost effectiveness.

Authors

Darren van Wageningen
Business Development Manager
Intel Corporation

Supriya Velagapudi
Product Marketing Manager
Intel Corporation

Miroslaw Walukiewicz
Solutions Architect
Intel Corporation

Connectivity along with the number of connected devices is undergoing exponential growth. This is driving the demand for a flexible radio access network (RAN) architecture. An Open RAN (O-RAN) is a nonproprietary implementation of a RAN that allows interoperability between cellular network equipment provided by different vendors. To support the O-RAN concept, the O-RAN Alliance[1] is a world-wide community of mobile network operators, vendors, and research and academic institutions operating in the RAN industry.

The 3rd Generation Partnership Project (3GPP)[2] is an umbrella term for a group of Standards Organizations that develop protocols for mobile telecommunications. 5G NR (New Radio) is a new radio access technology (RAT) developed by 3GPP for 5G that is designed to be the global standard for the air interface of 5G networks.

Communications Service Providers (CoSPs) are virtualizing the RAN to achieve cloud-like agility and economics. A growing number of CoSPs look to the scalability and operational efficiencies that virtual RAN (vRAN) and the O-RAN collaboration bring to 5G deployments. This includes the concept of network slicing, whereby CoSPs are able to increase their revenues by delivering 5G services with guaranteed levels of performance and quality. The virtualization of the RAN network is expected to grow at a 19% CAGR (compound annual growth rate) from 2020 through 2030.[3]

A key part of an O-RAN solution is a cell site router (CSR), which aggregates mobile data traffic from the radio tower(s) before transporting it back to the CoSP's core network. One aspect of virtualizing the RAN is to virtualize the CSR. Intel's accelerated virtual CSR (vCSR) solution, which is implemented on an Intel Agilex 7 FPGA-based N6000 card with O-RAN-compliant precision timing, helps CoSPs increase monetization of their services.

Evolving from 4G LTE to 5G NR

There are many differences between the radio equipment used for 4G LTE versus 5G NR RANs. For example, there are differences in the protocols used on the fronthaul links between radio antennas and the base station.

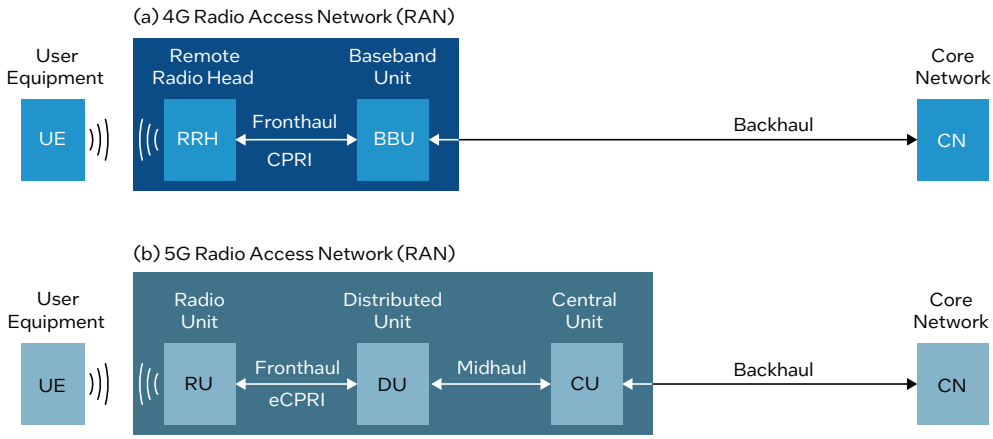

Figure 1(a) and (b). High-level representation of the difference between 4G and 5G RAN architectures. In the 5G RAN, the base station is split into three logical nodes: the RU, DU, and CU.

In the case of 4G RANs, radios on the cell tower are known as remote radio heads (RRHs). These are connected to a base station, also known as a baseband unit (BBU), which is itself connected to the core network. In this 4G context, the term fronthaul refers to the fiber-based connection between the RRH and the BBU, while the term backhaul refers to the connection between the BBU and the core network. As illustrated in Figure 1a, 4G fronthaul deployments typically employ the common public radio interface (CPRI).

In the case of a 5G RAN architecture, there are three logical nodes. Very often, these three logical nodes are implemented as three separate pieces of equipment. Some of the lower layer 1 (L1) functionality moves up into the radio unit (RU), while the remaining functionality is divided between the distributed unit (DU) and the central unit (CU). In this 5G scenario, fronthaul refers to the fiber-based connection between the RU and the DU, backhaul refers to the connection between the CU and the core network, and a new midhaul term is used to describe the connection between the DU and the CU.

Many 5G (and some new 4G) use cases demand a significant increase in fronthaul bandwidth. To address the increased cost of high bandwidth fronthaul, architects have suggested new functional partitions between the BBU and RRU. One way to reduce the fronthaul bandwidth is to move functions from the BBU into the RRU. Another is through data compression techniques. The 3GPP standards body defines the majority of mobile communications elements, testing, and interoperability. However, they have not defined the fronthaul interface. Previously, this gap was filled by common radio public interface (CPRI). For 5G and Ethernet based transport, several groups have stepped forward to fill the specification gap.

CPRI interfaces require a dedicated 'dark' fiber and carry 'basic frames.' These basic frames encapsulate user plane traffic (the time domain digitized radio traffic), control plane information, and timing information. This timing information is easily extracted and is highly deterministic. It is quite easy to synchronize the radio using this information to the required accuracy of around ±8 nanoseconds.

As illustrated in Figure 1b, 5G typically uses O-RAN with an enhanced CPRI (eCPRI) transport for fronthaul. O-RAN/eCPRI is transported over Ethernet. By using Ethernet, a very common type of transport, there are multiple technical and cost benefits. Ethernet, however, was not initially conceived to carry time-sensitive information and this presents a significant challenge. Ethernet packets can take different pathways through a network, they may pass through routers, and most of the elements in the network introduce some non-deterministic latency in the order of 10s of microseconds. To resolve this issue, the O-RAN interface supports several synchronization schemes. Primarily, these make use of IEEE standards in time sensitive networks (TSNs), such as the IEEE1588v2 "Standard for Precision Clock Synchronization Protocol for Network Measurement and Control Systems." By using a sequence of precision time protocol (PTP) messages over the same Ethernet interface/pathway, the two ends can accurately exchange what is known as the "time of day" (ToD). A primary clock, which is locked to a primary reference clock source (PRTC) (e.g., the global satellite navigation system (GNSS)), which distributes the ToD to other network entities that behave as subordinate clocks. In ORAN network topologies, these network entities are typically the radios (RU). These subordinate clocks compare their internal sense of time (ToD) and adjust their clocks to track the PRTC ToD. The PRTC may be co-located with the CU, DU, or core. While it is not mandatory to maintain end-to-end timing accuracy, switches and routers in the network should also be synchronized.

In the case of the fronthaul, typical 5G radios use eCPRI and the 3GPP-defined split point 7.2, whereas 4G radios use CPRI and split point 8. In cases where the RAN has to accommodate both 4G and 5G infrastructure, a device called a fronthaul gateway (FHGW) is often used to perform CPRI to eCPRI conversions, including the lower L1 signal processing (FFT and PRACH), which sits between split 7.2 and split 8.

RAN deployments are divided into two main architectures. In C-RAN (Cloud RAN) deployments, the DUs are centralized and may be located 20 km or more from the RUs. This is a typical deployment for dense urban scenarios and allows baseband pooling, a technique to optimize processing needs. In D-RAN (Distributed RAN) deployments, the DUs are distributed at the cell sites. This is a typical deployment scenario for rural deployments.

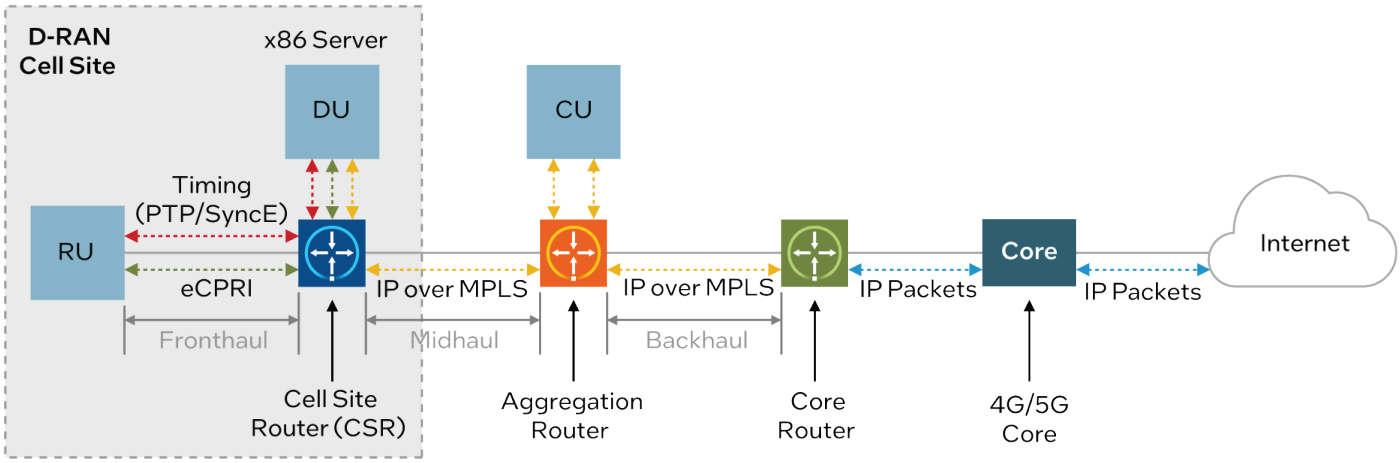

Figure 2. High-level representation of a cell site in the context of a 5G RAN.

A 5G D-RAN involves a combination of the RUs (radios and antennas), the DU, and any ancillary equipment (Figure 2). Here, a cell site router (CSR) connects the RUs with the DU and the DU with the rest of the mobile infrastructure where the CU resides. The fronthaul connection between the RU and the CSR/DU is usually fiber. The midhaul connection between the DU/CSR and the aggregation router/CU may be implemented using fiber, microwave, or satellite link using IP over multi-protocol label switching (MPLS).

Multiple RUs may be connected to a single CSR, and multiple CSRs in a mesh or ring network can be connected to a single aggregation router. At the time of this writing, most CSRs are implemented as stand-alone devices external to the DU. However, as the industry moves from 4G to 5G, there is an opportunity to merge the CSR function, an optional FHGW function for 4G radios, optional future proofing with 6G Numerology FEC acceleration, and the DU into a single commercial off the shelf (COTS) platform.

The concept of a distributed base-station (CU and DU) coupled with an integrated vCSR has many advantages for operators. In the case of D-RANs, for example, the fronthaul network can be very short, thereby allowing the number of time-sensitive switches to be reduced or removed completely. As a byproduct of limiting the complexity of the fronthaul, the clock synchronization area is also limited. All this goes to help simplify the design of the hardware and reduce costs; for example, DUs can now be implemented using today's COTS servers from Supermicro, WNC, and Kontron, including the vCSR implemented as an Intel® FPGA SmartNIC N6000-PL Platform based acceleration card as discussed below. Since the fronthaul requires large data rates, terminating the link close to the antenna at the DU makes it possible to use significantly smaller rates on the midhaul where neither clock synchronization nor error correction is necessary. Also, disaggregating the base station makes it possible to employ multiple midhaul paths for the implementation of various network slices or to facilitate the sharing the mobile infrastructure between multiple operators.

Overview of Intel's Accelerated vCSR solution

The Intel® FPGA-based accelerated vCSR solution is an all-in-one architecture that features an integrated vRouter function (the vCSR) coupled with optional integration of a FHGW and optional integration of the baseband acceleration. As illustrated in Figure 3, all these functions are co-located on a single Intel® Agilex 7 FPGA-based SmartNIC—the Intel® FPGA SmartNIC N6000-PL Platform—which can be plugged into an X86 server acting as a virtual DU (vDU).

It's also possible to add additional user services, such as hierarchical quality of service (HQoS) because the Intel FPGA SmartNIC N6000-PL Platform based card appears as a foundational NIC with SR-IOV interfaces, which means there is no functional change with respect to applications running with or without the Intel FPGA SmartNIC N6000-PL Platform.

Figure 3. Using an FPGA SmartNIC accelerator to virtualize the CSR function

If the vDU is running on a 4th Gen Intel® Xeon® Scalable processor (SPR-EE), which contains a dedicated on-chip hardware accelerator called Intel® vRAN Boost. The vRAN Boost accelerator can be used to offload 4G (Turbo) and/or 5G (LDPC) FEC processing from the main CPU cores. For processors without Intel® vRAN Boost, a FEC can be implemented on the same Intel N6000-based SmartNIC card that is running the vCSR.

The integrated vRouter function will handle level 2 (L2) traffic management and level 3 (L3) routing, and also provide 1588/PTP support. This solution is designed to support the Juniper Networks Cloud-Native Router stack, which provides commercial-grade, high-performance, scalable routing. This provides a state-of-the-art solution for the routing plane integrated with the MPLS data-plane. Alternatively, open-source options such as the FRRouting (FRR) Project can also be used. These vCSR tasks are run on the Intel Agilex 7 SoC FPGA's hard processor subsystem (HPS) cores, thereby freeing up the host server's processor cores for revenue-generating functions.

The Intel Accelerated vCSR can support the midhaul (F1 interface) over MPLS or SR-MPLS. It also offers a unique capability to support future infrastructures (e.g., SRv6) with the same hardware. The solution also supports auto-provisioning, which is a key feature desired by CoSPs, to simplify the task of network slicing. Auto-provisioning avoids manual programming of each MPLS label, which slows deployment. For a CoSP to manually enter MPLS labels when used in network slicing, each must be multiplied by the number of slices. This can grow into tens of thousands of time-consuming and error-prone manual entries, all of which can be avoided by means of auto-provisioning.

The Intel® FPGA SmartNIC N6000-PL Platform supports all synchronization profiles defined by the O-RAN WG-4 specification, including LLS-C1 (the vCSR is GM/T-BC, the RU is synchronized over fiber), LLS-C2 (the vCSR is GM/T-BC, the RU is synchronized over the network), LLS-C3 (both the vCSR and RU are synchronized from the network), and LLS-C4 (both the vCSR and RU are synchronized separately). A high-level view is that Intel's Accelerated vCSR combines fronthaul to midhaul conversion with synchronization, midhaul transport capabilities with networking slicing support, and a management piece for server management and telemetry, which is a key requirement for network operators. Furthermore, Intel's vCSR supports the clock accuracy required for Class B and C requirements in 5G.

As an alternative to the Intel FPGA N6000-PL Platform, the Intel Accelerated vCSR solution can also be implemented on a custom board with any Intel Agilex 7 FPGA sized appropriately for a specific configuration. For example, a custom board could expand the number of fronthaul or midhaul ports beyond the six ports (three of each) that are currently supported on the Intel FPGA N6000-PL Platform. Additional customizations may include security features like internet protocol security (IPsec) or media access control security (MACsec) implemented in hardware. Also, the routing and control stack can be implemented using the hardened Arm processor cores embedded in Intel Agilex 7 SoC FPGAs, thereby allowing these functions to remain invisible to software applications running on the host.

Intel Agilex 7 FPGAs and SoC FPGAs include a hardened secure device manager (SDM) function, which allows the secure update of the FPGA in the field. This supports advanced features like secure Day-0/Day-1/Day-2 provisioning, which is a key element for system management allowing automatic remote updates.

It is difficult to overstate the value, power, and manageability that results from implementing all of these virtual functions—which have traditionally been realized as stand-alone external components or cards—into a single FPGA-based card. This results in greatly simplified component management, which traditionally involved multiple application programming interfaces (APIs) from multiple vendors. It also results in a dramatic reduction in the number of cables and optical fibers required to link everything in a large cell site.

Network Slicing Support

Implementing low-latency services at an edge compute location typically modifies the hardware and software requirements at that location. In the case of a 4G/LTE installation, since no user-specific functions were required, the hardware and software could be implemented as proprietary appliances. By comparison, to maximize 5G monetization, running low-latency user services with the DU, CU, and 5G user plane function (UPF) is a standard deployment model.

The most appropriate implementation is to use COTS servers accompanied by FPGA-based hardware accelerator cards, thereby allowing the maximum use of server CPU cores to run the user services. Very often, a user-service will be in the same location as the UPF, which can decide which traffic should be treated as low-latency to be handled by edge compute versus normal traffic that can be handled by core services. Figure 4 illustrates such possibilities.

The most appropriate implementation is to use COTS servers accompanied by FPGA-based hardware accelerator cards, thereby allowing the maximum use of server CPU cores to run the user services. Very often, a user-service will be in the same location as the UPF, which can decide which traffic should be treated as low-latency to be handled by edge compute versus normal traffic that can be handled by core services. Figure 4 illustrates such possibilities.

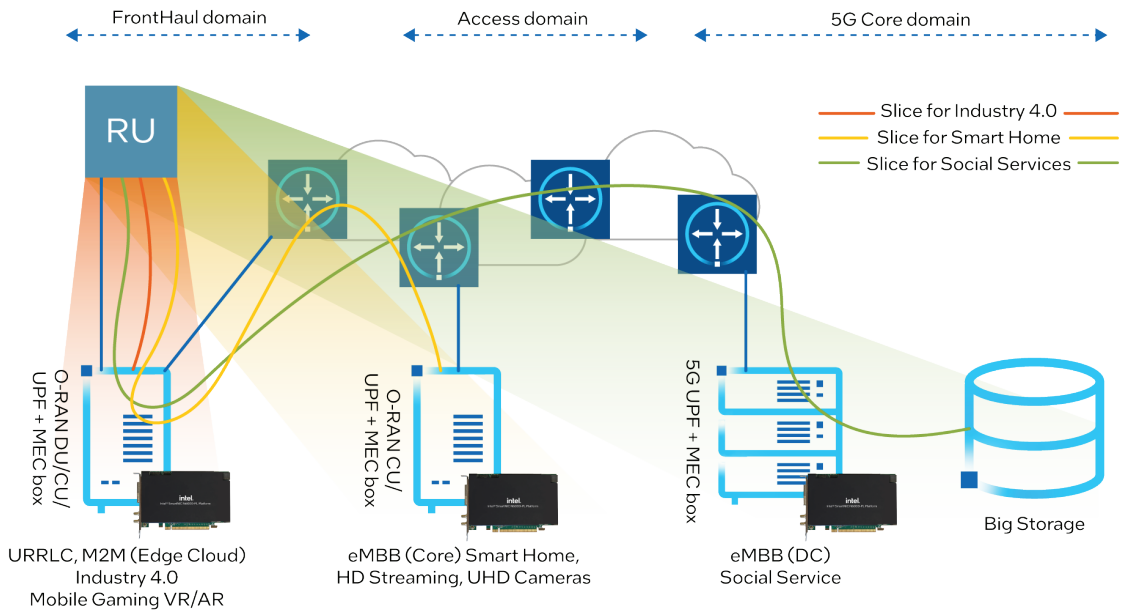

Figure 4. Network slicing examples

Depending on network slicing requirements, the UPF and multi-access edge computing (MEC) could be located as close as possible to the fronthaul, in the access network, or in the 5G core domain. COTS servers can be used in every location. This means that the mobile infrastructure can be formed with a single COTS server in which the DU, CU, and UPF functions are all running as applications consuming CPU cores. It's possible that a larger DU could be run using a single server or a set of servers, or it could be co-located with other functions. In all cases, the functionality is the same, but the traffic handling capability will scale depending on the deployment model.

Today's telecom operators seek optimal ways of connecting edge services with the rest of the mobile infrastructure. Many technologies are considered, from traditional label distribution protocol (LDP) MPLS to various options related to segment routing protocols enhanced with IPsec and in-situ operations, administration, and management (iOAM) capabilities. Every operator wishes to deploy an optimum solution to provide the best connectivity options for its customers. FPGAs provide a cost-effective way to implement high-performance hardware-accelerated solutions.

Overview of the Intel® FPGA SmartNIC N6000-PL Platform

The Intel FPGA SmartNIC N6000-PL Platform (Figure 5) is a 3rd generation Intel Agilex 7 FPGA family-based SmartNIC for network acceleration. It supports 2x100 Gbps Ethernet connectivity and offers higher performance, total cost of ownership (TCO) optimization, and scalability compared to previous generations.

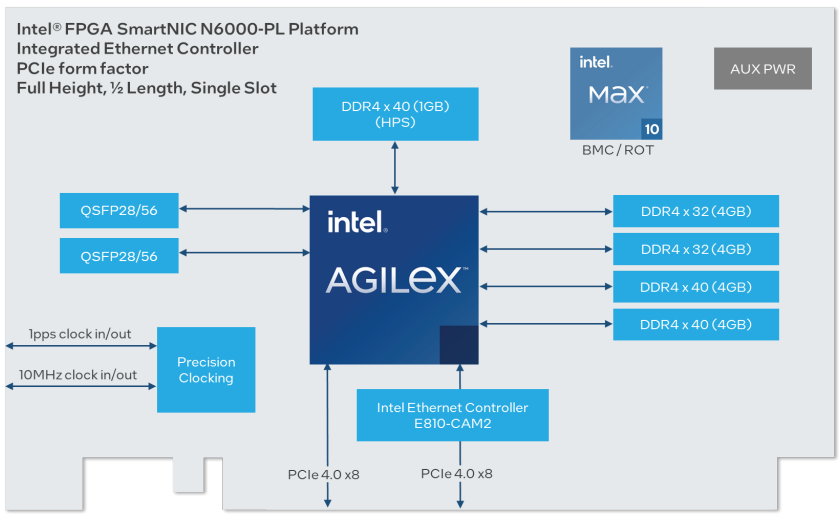

Figure 5. Block Diagram of Intel FPGA SmartNIC N6000-PL Platform

Intel Agilex 7 FPGAs are built with Intel's advanced 10 nm SuperFin technology and a second-generation Intel® Hyperflex™ FPGA Architecture. Intel Agilex 7 devices deliver ~2X better fabric performance per watt compared to competing 7 nm FPGAs. This high-performance FPGA-based SmartNIC is a versatile card supporting hardware programmable acceleration of communication workloads, such as 4G/5G vRANs, vCSRs, 5G UPFs, Contrail2 (CN2), SMPTE ST2110 professional media over managed IP networks, and more.

As illustrated in Figure 6, the Intel FPGA SmartNIC N6000-PL Platform is presented in two ways:

- **As Production-ready Solutions**: Customer who wishes to deploy commercial off-the-shelf N6000 based-SmartNIC "as-is" can buy N6000-based production cards and workloads from Intel's partners. Open FPGA Stack (OFS) and BMC design files are available from these partners to accelerate custom workload development.
- **As a Platform Design:** Customers who wish to accelerate their own custom board design by leveraging the N6000 board design and customizing it to add their own differentiation can use the Intel FPGA N6000-PL Platform, consisting of board design files, the Open FPGA Stack (OFS), BMC design, workloads, documentation, and a pre-production board (contact your local Intel sales representative to learn more).

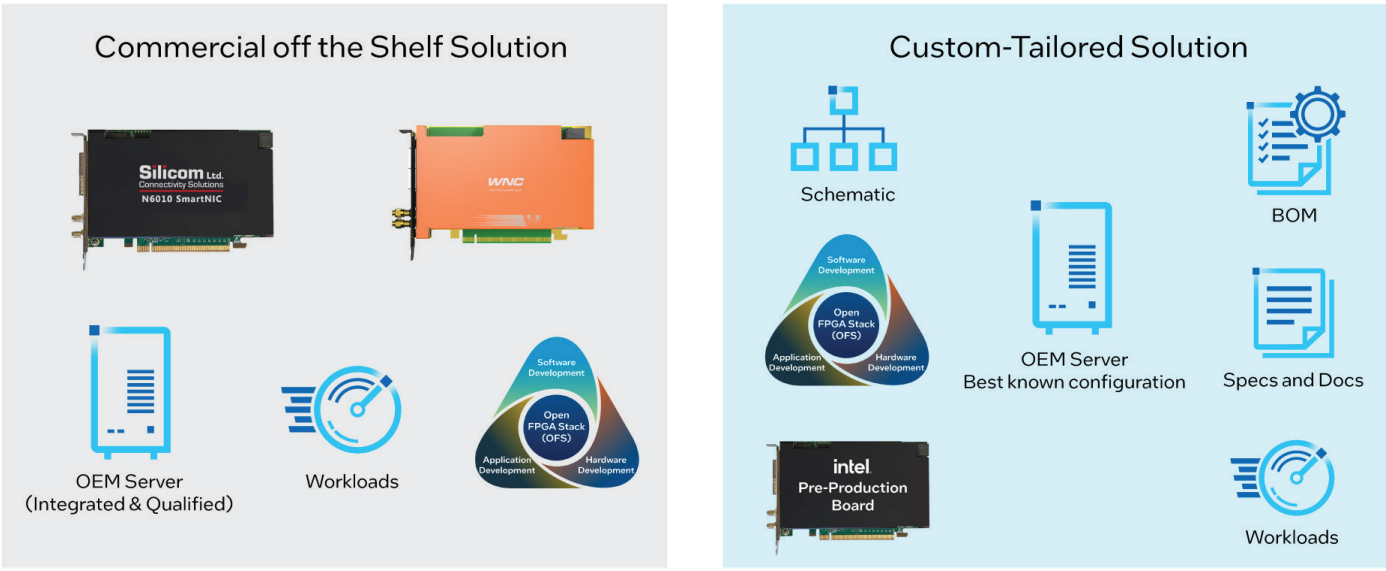

Figure 6. Intel FPGA N6000-PL Platform is available as COTS solution from Intel's partners or as platform design from Intel for custom-tailored solution

Intel FPGA SmartNIC N6000-PL Platform Specifications and Ordering Information

	Hardware
FPGA	Intel Agilex FPGA F-Series (AGF014) • High-performance F-Series FPGAs, multi-gigabit SERDES transceivers up to 58 Gbps • 1,437K logic elements • 190 Mb on-chip memory • 4,510 DSP blocks
Onboard Memory	16 GB DDR4 to FPGA 1 GB DDR4 to HPS
PCIe Interface	PCIe 4.0 bifurcated x8/x8
Other Interfaces	• On-board Intel E810-CAM2 Ethernet Controller – Extensive OS support and easier system integration – 100 Gbps pipeline On board Intel E810-CAM2 Ethernet Controller • 2X QSFP56 with up to 2x100 Gbps Ethernet support (2x1x100G, 2x2x50G, 2x4x25G, 2x4x10G) • Support for SyncE, CPRI, eCPRI • Front Panel SMA for IEEE1588v2 1pps/10MHz, • O-RAN S-Plane PTP support (G.8275.1) for LLS-C1, -C2, -C3, -C4
Form Factor	Full Height Half Length (FHHL)
Thermal/Power	• Passively cooled • Thermal design power(TDP*) – 125W • Scenario design power (SDP) <100 W (Power consumption dependent on workload) • NEBS Class 1 compliance support
Board Management	• Intel® MAX 10 FPGA BMC • Full security implementation using Intel MAX 10 FPGA as RoT • Remote update capabilities for FPGA flash memory and BMC • Full card BMC solution host communication via SMBus and PCIe VDM
Power Management	Intelligent system power management with real-time telemetry and system health monitoring
	Software
Software	• Data Plane Development Kit (DPDK) • FlexRAN™ software (BBDev (pf-bb-config)) for vRAN only • Open Programmable Acceleration Engine (OPAE) • Open FPGA Stack (OFS)
	Design Entry Tools
Design entry tool	Intel® Quartus® Prime Pro Edition Software
	Ordering Information
Ready-to-deploy COTS Board	Buy now from: • Winston NeWeb Corp (WNC) (OPN: **WNC FPGA SmartNIC WSN6050**) • Silicom Inc (OPN: **Silicom FPGA SmartNIC N6011**)
Custom tailored design	**For Intel FPGA SmartNIC N6000-PL Platform design** (OPN: Intel SmartNIC Platform N6000-PL) Contact Intel Sales representative.

*TDP not relevant to deployment power. TDP only used for server thermal design.

Conclusion

A growing number of CoSPs are looking to the scalability and operational efficiencies that virtual radio access networks bring to 5G deployments. One aspect of virtualizing the RAN is to virtualize the CSR, which helps CoSPs increase monetization of their services. Intel's Accelerated vCSR solution is implemented on an Intel Agilex® 7 FPGA-based N6000-PL Platform with O-RAN-compliant precision timing. This vCSR solution is available in COTS servers from Intel partners such as SuperMicro, Kontron, and WNC.

Compliant with the O-RAN LLS-C1, LLS-C2, LLS-C3, LLS-C4 synchronization profiles, this hardware-accelerated vCSR can provide Class B and C precision timing accuracy while increasing interoperability. The CSR routing and control stack run on the Intel Agilex® 7 SoC FPGA's hard processor system without the need for CPU cores to be provisioned on the host.

By employing common software and by being based on standard, open APIs, this vCSR solution removes vendor lock-in, gives customers more choices, and simplifies deployment and troubleshooting. Furthermore, Intel's vCSR solution also supports network slicing, thereby allowing operators to monetize individual services.

Learn More

- Intel Agilex 7 FPGAs
- Intel FPGA Programmable Acceleration Cards
- Intel FPGA SmartNIC N6000-PL Platform
- Silicom FPGA SmartNIC N6011
- WNC SmartNIC N6050
- Juniper Cloud-Native Router stack
- WNC Intel Xeon-D based servers
- For Kontron Multi-Edge Server/Platforms, click below:
  - ME1310 High Performance Multi-Edge Platform
  - RS1310 High Performance Outdoor IP65 Multi-Edge Server

References

1 www.o-ran.org/

2 www.3gpp.org/

3 www.transparencymarketresearch.com/virtualized-radio-access-network-market.html

SS-1161-1.0