White Paper

Intel® Ethernet

Enabling RAN Evolution with Standards-Based Ethernet

Intel® Ethernet 800 Series Network Adapters provide high-precision timing capabilities that enable the network synchronization needed for vRAN deployments on open-standards-based general-purpose server hardware. These capabilities support greater choice in solutions and lower total costs, while making solutions more future-ready.

The evolution of radio access networks (RANs) from closed-ecosystem, single-vendor approaches to highly interoperable sets of open standards-based solutions continues to drive down implementation cost and complexity. Software-defined virtualized RANs (vRANs), consisting of disaggregated subsystems and elements, enable a greater level of choice as communication service providers (CoSPs) can mix and match elements from multiple vendors to suit customer requirements. Likewise, vRANs are deployed using general-purpose, commercial off-the-shelf (COTS) server hardware, enabling them to take advantage of the cost, performance and ecosystem advantages of Intel® architecture, including a diverse and growing selection of Intel Ethernet Network Adapters.

The evolution of Open RAN (O-RAN) architecture gives CoSPs the ability to choose the combination of hardware and software elements that best meet the needs of each individual deployment. They can select optimal components to balance factors such as capacity and power for improved total cost of ownership (TCO), which plays a critical role in enabling CoSPs to build out the increased number of base stations that fully realized 5G networks require. As technologies and solutions begin to gain acceptance and evolve in the O-RAN ecosystem, the power of choice in the marketplace drives even more ecosystem innovation, in a virtuous cycle of development. The open-standard O-RAN architecture allows operators to easily integrate new innovations via software updates or by replacing specific network components, rather than the forklift upgrades associated with traditional single-vendor solutions.

O-RAN disaggregates RAN components into centralized units (CUs), distributed units (DUs) and radio units (RUs). Intel is enabling O-RAN deployments across the stack, including the development of advanced network timing capabilities in Intel Ethernet 800 Series Network Adapters. Enhanced capabilities such as IEEE 1588 precision time protocol (PTP) and Synchronous Ethernet (SyncE) virtualize and replace the traditional cell site router while enabling the use of packet-based networking for fronthaul between the RU and DU on an add-in card or integrated COTS solution.

Table of Contents

Evolution of RAN fronthaul technologies

In parallel with other aspects of carrier networks, RANs have evolved to be instantiated in software using containerized network functions (CNFs), in cloud-native topologies built for openness and interoperability. For RANs, this shift requires a refactoring of fronthaul mechanisms from their traditional closed implementations to more open technologies.

CPRI and Radio over Ethernet

An industry consortium introduced the Common Public Radio Interface (CPRI) standard in 2003 as a specification for fronthaul of digitized radio-frequency data based on proprietary, vendor-specific implementations. In the context of a typical 4G deployment, CPRI might connect a cluster of remote radio units (RRUs) to the baseband unit (BBU) using Ethernet over optical fiber with a line rate of 10, 25 or soon 50Gbps. As a point-to-point interface based on vendor-specific antenna implementations, CPRI is oriented toward proprietary networking, rather than more forward-looking open models.

CPRI is a high-speed serial interface based on time-division multiplexing (TDM) that uses a separate timing module. To enable the use of a packet-based network, Radio over Ethernet (RoE) technology can encapsulate CPRI traffic into Ethernet frames at a gateway before transmission, with corresponding decapsulation at the other end, as shown in Figure 1.

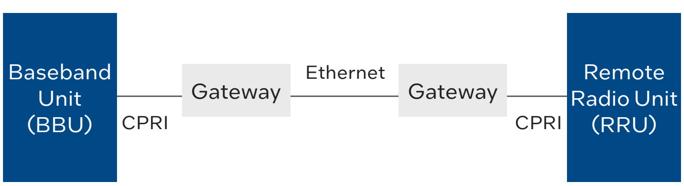

Figure 1. CPRI fronthaul uses proprietary connectivity, potentially with Radio over Ethernet.

The technology must ensure a constant transmission rate and delay by maintaining a predefined channel bandwidth, regardless of load. In practical terms, the need to accommodate usage peaks means that the network must always maintain transmission at peak bandwidth. The resulting inefficiency becomes more financially significant as the number of antennas in use and traffic volumes between RRUs and BBUs increase. As the industry moves to 5G, both those quantities rise dramatically, placing strain on the capacity of the fiber and making fronthaul costs potentially untenable.

Improving flexibility and efficiency with eCPRI

With the advent of 5G networking, a number of shifts have combined to spur the development of enhanced CPRI (eCPRI). Higher data rates will quickly outstrip the capabilities of CPRI, while the expectation for interoperability among technology vendors requires a truly open approach. In keeping with the broader trend of disaggregation in 5G networks, eCPRI also enables functional partitioning and decomposition that allows task responsibility to be transferred from the BBU to the RRU, reducing requirements for higher bit rates and fidelity on the fiber connections between them.

eCPRI provides the full scope of functionality of CPRI (and more) using standards-based Ethernet as shown in Figure 2, so that commercial off-the-shelf (COTS) networking equipment can take the place of proprietary solutions. That reduces the cost of the solutions themselves while also implementing those solutions using common, ubiquitous Ethernet technologies and the ecosystem-driven innovations associated with them. It also reduces inflexible dependencies among different parts of the network, relieving CoSPs of the technical debt associated with predecessor approaches.

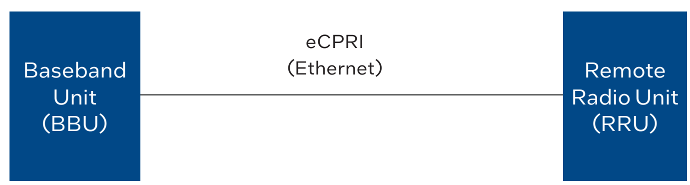

Figure 2. eCPRI fronthaul uses standards-based Ethernet.

The use of eCPRI with standard Ethernet reduces the complexity of the data path without gateway translation, and it does not require constant maintenance of peak bandwidth as CPRI does. These factors provide superior scalability and latency with eCPRI, making it more viable for 5G networks, including for vRAN and cloud RAN implementations. Accordingly, off-the-shelf Ethernet devices are evolving to meet emerging requirements of eCPRI usage. The proliferation of towers due to 5G's smaller cell size compounds the value of this change to CoSPs, enabling them to use open-standards networking equipment.

Tailoring Ethernet capabilities to emerging requirements

O-RAN is a multi-vendor, ecosystem-driven approach and initiative to building RAN technologies and solutions with an emphasis on interoperability. Disaggregation of network functions enables CoSPs to choose the ideal set of components for each deployment. By removing the need for single-vendor architectures, O-RAN spurs innovation and drives down costs. At the same time, this shift drives up the need for precise interoperation across vendors.

One aspect of that requirement as CoSPs build out the RAN deployments for their 5G networks is that accurate and reliable network synchronization is more critical than ever. The TDM mechanisms that underlie synchronization in the CPRI world are not present in standard Ethernet, and an alternate mechanism is required. Traffic such as high-bandwidth media content and secure protocols create tight timing requirements to ensure fidelity and user experience without service disruptions, dropped packets or loss of transmission quality. Synchronization between network adapters and network switches and radios is needed to maintain correct packet delivery order, in a secure and reliable way. The network must maintain that capability without data loss or corruption, across solution building blocks from multiple providers.

Timing and synchronization of fronthaul transport are maintained by the synchronization plane (S-plane), which draws on specialized capabilities for that purpose. High accuracy is required to support 5G processes, including multiple input/multiple output (MIMO), time-division duplexing (TDD) and aggregation of traffic from multiple RRUs. O-RAN implementations accomplish synchronization using standardized S-plane protocols that include the following:

- **IEEE 1588 PTP v1 and v2** enables precise timing synchronization among clocks in heterogeneous systems, which may have different characteristics in terms of precision, resolution and stability. The protocol enables a hierarchical architecture among clocks for a time distribution system. At the head of that hierarchy is a root timing reference known as the grandmaster clock, which is elected from among the member clocks.
- **SyncE** enables timing signals to be transmitted over the Ethernet physical layer, traceable to external clock data based on PTP. Thus, while PTP adjusts for the time differences between nodes, SyncE enables a device to consume that information by matching the speed of its upstream neighbor. Together the two protocols enable the S-plane to maintain lockstep in the delivery of packets through a distributed network.

The integration of PTP and SyncE into industry-standard network adapters extends traditional telecom timing requirements and architectures to COTS hardware solutions and open software. That openness and choice are vital to the cost-effective, flexible development requirements of O-RAN solutions.

Timing and synchronization with Intel Ethernet

Some Intel Ethernet 800 Series Network Adapters, as shown in Figure 3, provide hardware-enhanced timing and synchronization capabilities. These capabilities deliver the phase accuracy that vRANs and other edge networks require, without the added cost and complexity of dedicated timing equipment.

The adapters incorporate precise timing data into traffic streams on a per-packet basis. They each provide two coaxial jacks for syncing or cascading one-pulse-per-second signals with other devices, as well as for auditing performance on each node. Flexible port configurations are available to support the density and small-form-factor requirements that are common in vRAN deployments.

- **E810-XXVDA4T network adapter** provides four SFP28 ports that are capable of 25Gbps, 10Gbps, or 1Gbps Ethernet connectivity.
- **E810-CQDA2T network adapter** features two QSFP28 100Gb ports that can be configured to deliver 100Gbps, 25Gbps or 10Gbps Ethernet.

Both adapters incorporate an onboard oscillator that maintains plus or minus 1.5 microsecond timing accuracy for up to four hours in the event of signal loss, improving the reliability of the system under challenging real-world conditions. Building further on solution resiliency, these adapters can support an optional, low-cost global navigation satellite system (GNSS) module for high-accuracy satellite-based timing references. The GNSS receiver supports frequency, phase and time-of-day synchronization with global navigation satellite systems, including GPS, Galileo, GLONASS, BeiDou and QZSS.

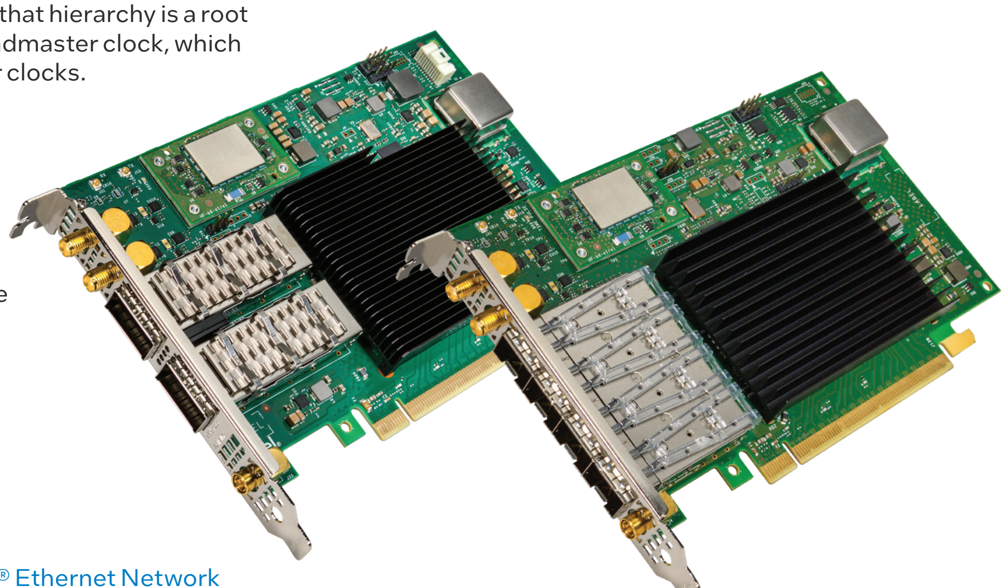

Figure 3. Intel® Ethernet Network Adapters E810-XXVDA4T and E810-CQDA2T.

Using these mechanisms, Ethernet can deliver high quality for calls, video streams and other sensitive traffic for subscribers. Likewise, the adapters are well suited to other critical network usages enabled by 5G, such as advanced location services, smart energy grids, financial services industry applications and industrial control systems. The precision timing characteristics of these adapters are complemented by additional technologies for deterministic performance, scalability and efficiency, including Dynamic Device Personalization (DDP). This feature defines workload-specific profiles with predetermined sets of protocols to tailor network operation and a programmable pipeline that optimizes throughput and latency for individual high-sensitivity network usages.

Ecosystem-driven innovation for the future

With traditional RANs, where a single vendor provides the aggregated solution stack of hardware, software and services, integration was part of the solution itself. Open RAN disaggregates the components and makes it possible to mix and match component offerings from multiple vendors. Along with the choice and flexibility this approach offers, it also moves integration responsibility from a single monolithic vendor to the broader ecosystem, including technology vendors, solution integrators and end customers, among others.

Collaborative co-engineering for shared benefit

Some aspects of vRAN integration are common to multiple deployments, across different end customers. That repeated work creates inefficiencies, increases risk and raises the bar for adoption. Intel works closely with major hardware, software and solution providers, including makers of operating system environments, network functions, equipment and more to help streamline integration. There can be no “one size fits all” for layered vRAN solutions, and this work helps give customers confidence that their vRAN implementations can draw on advanced ecosystem validation and integration work.

Intel’s engagement with ecosystem members begins early in the development of each product generation, to prepare makers of solution components to take advantage of new technologies and capabilities. At a practical level, this work calls for involving the ecosystem partners early in the planning process and providing early access to product samples for solution development, integration and validation. Intel also engages in joint work optimizing large deployments with CoSP and enterprise customers, which pays dividends to the broader industry as adoption matures.

Boundary clock testing with EANTC

The timing synchronization requirements typically fulfilled by separate timing appliances have been integrated into Intel Ethernet Network Adapters E810-CQDA2T and E810-XXVDA4T. In place of expensive dedicated devices, that functionality is available as a value-added feature of the adapter. To demonstrate the capability of this approach, Intel participated in testing by the European Advanced Networking Test Center (EANTC).

The results confirm Intel’s own testing outcomes and verify that Intel Ethernet Network Adapters E810-CQDA2T and E810-XXVDA4T meet or exceed key International Telecommunications Union (ITU) standards for network timing. This testing demonstrates that the adapters enabled COTS servers to fulfill all relevant roles in the timing topology as required for deployment in 5G RAN deployments.

Real-world deployment examples

Figure 4 depicts a cell site that uses a server equipped with two E810-XXVDA4T network adapters. The cell tower shown on the right side of the figure has six 10Gbps or 25Gbps fiber connections going up the mast, and two uplinks, all of which are split across the two adapters. In this scenario, only one of the adapters needs a GNSS module for them to synchronize with each other for fronthaul and backhaul.

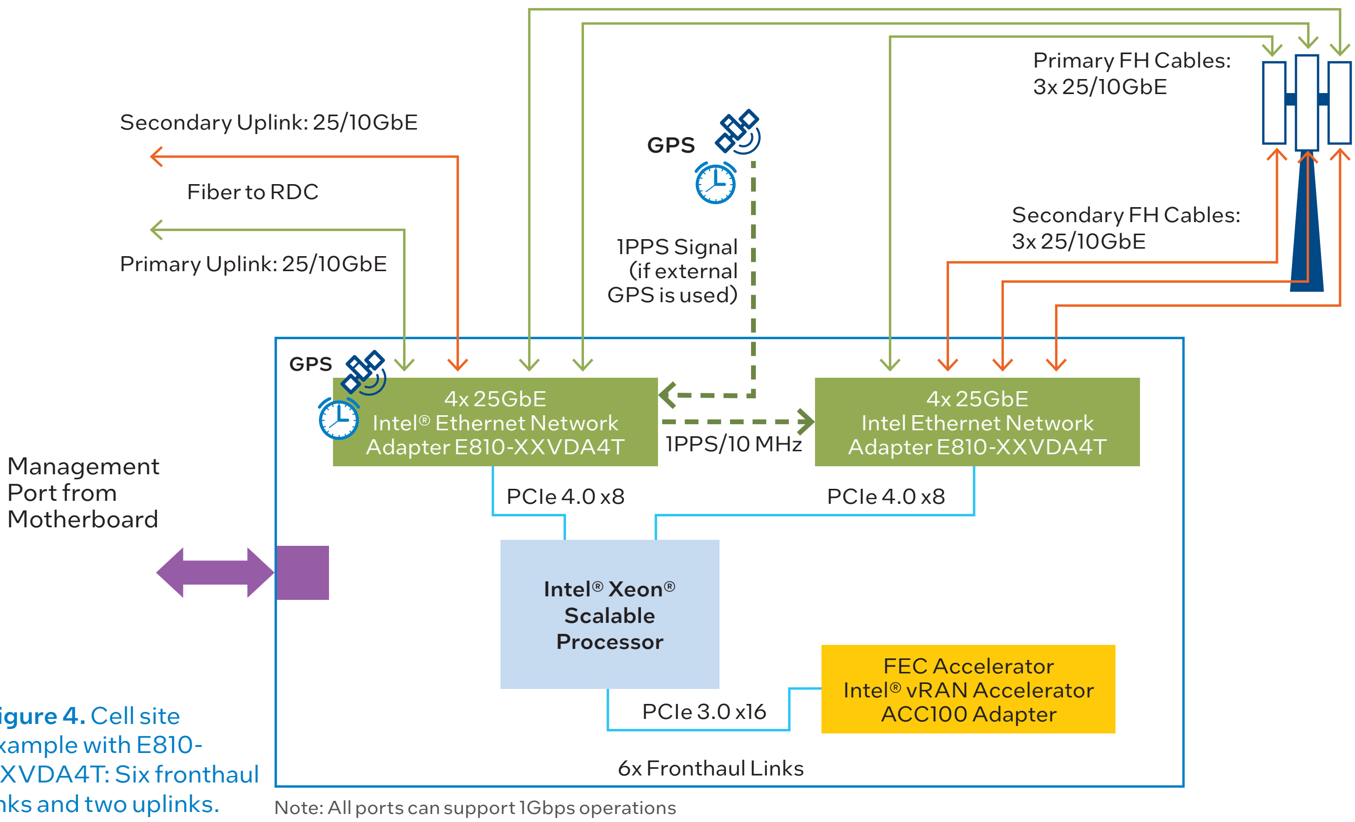

Figure 4. Cell site example with E810-XXVDA4T: Six fronthaul links and two uplinks.

Figure 5 illustrates a case with more 5G radio units on the tower. Here, there are nine single-mode fiber cables going up the mast and three links to backhaul, requiring a total of 12 Ethernet ports using three E810-XXVDA4T network adapters. Similar to the previous example, only one of the three adapters needs an integrated GNSS receiver to synchronize all three.

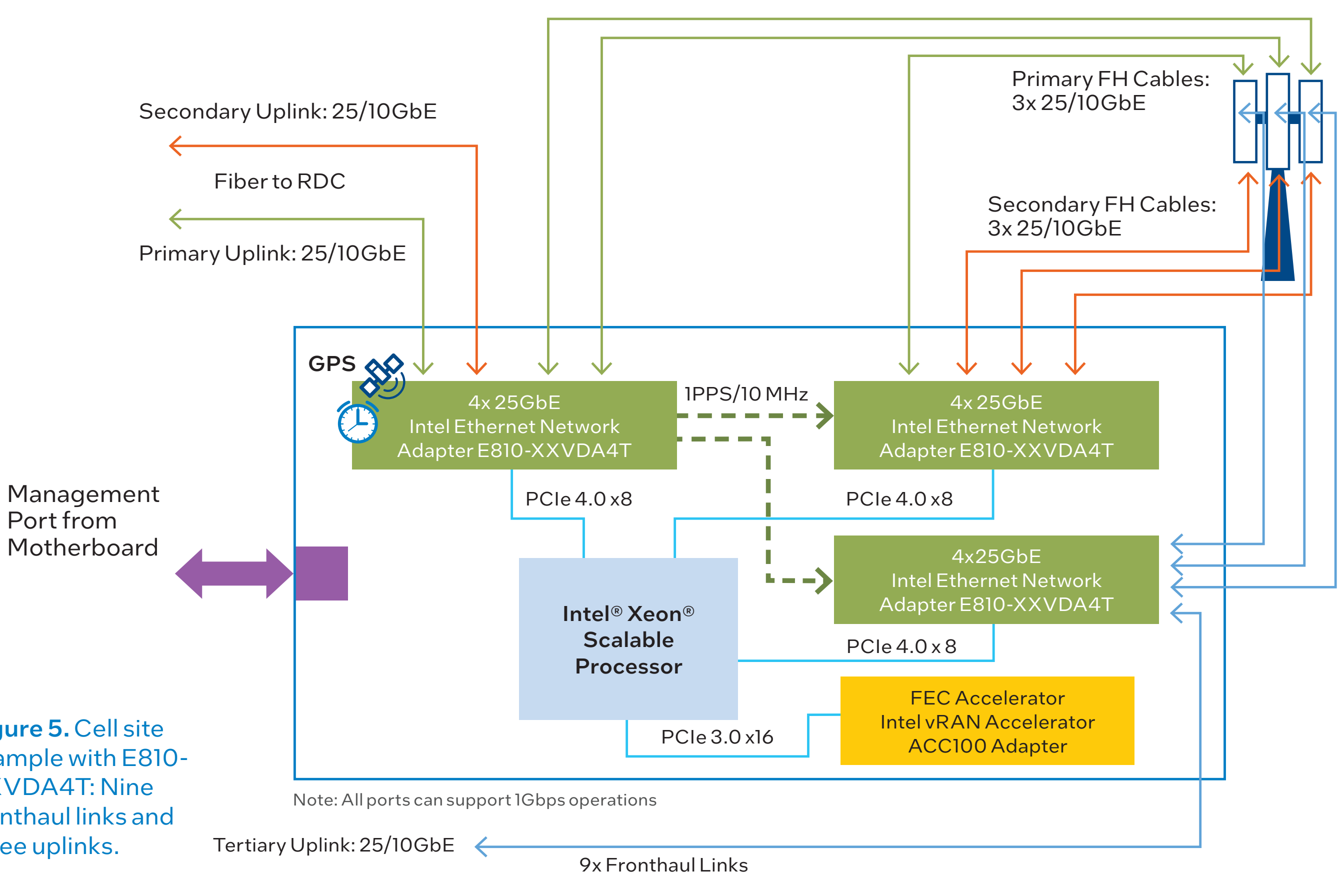

Figure 5. Cell site example with E810-XXVDA4T: Nine fronthaul links and three uplinks.

Figure 6 shows an implementation where the server is equipped with an E810-CQDA2T network adapter. Here, there are six 10Gbps single-mode fiber connections going up the mast and two single-mode 10Gbps connections providing uplink to the backhaul network. The single E810-CQDA2T network adapter can provide all eight of the ports required using PSM4 breakout cables or modules. In the illustration, one of the adapter's QSFP28 ports and its breakout provides four connections up the mast, while the other provides two connections up the mast and two backhaul connections.

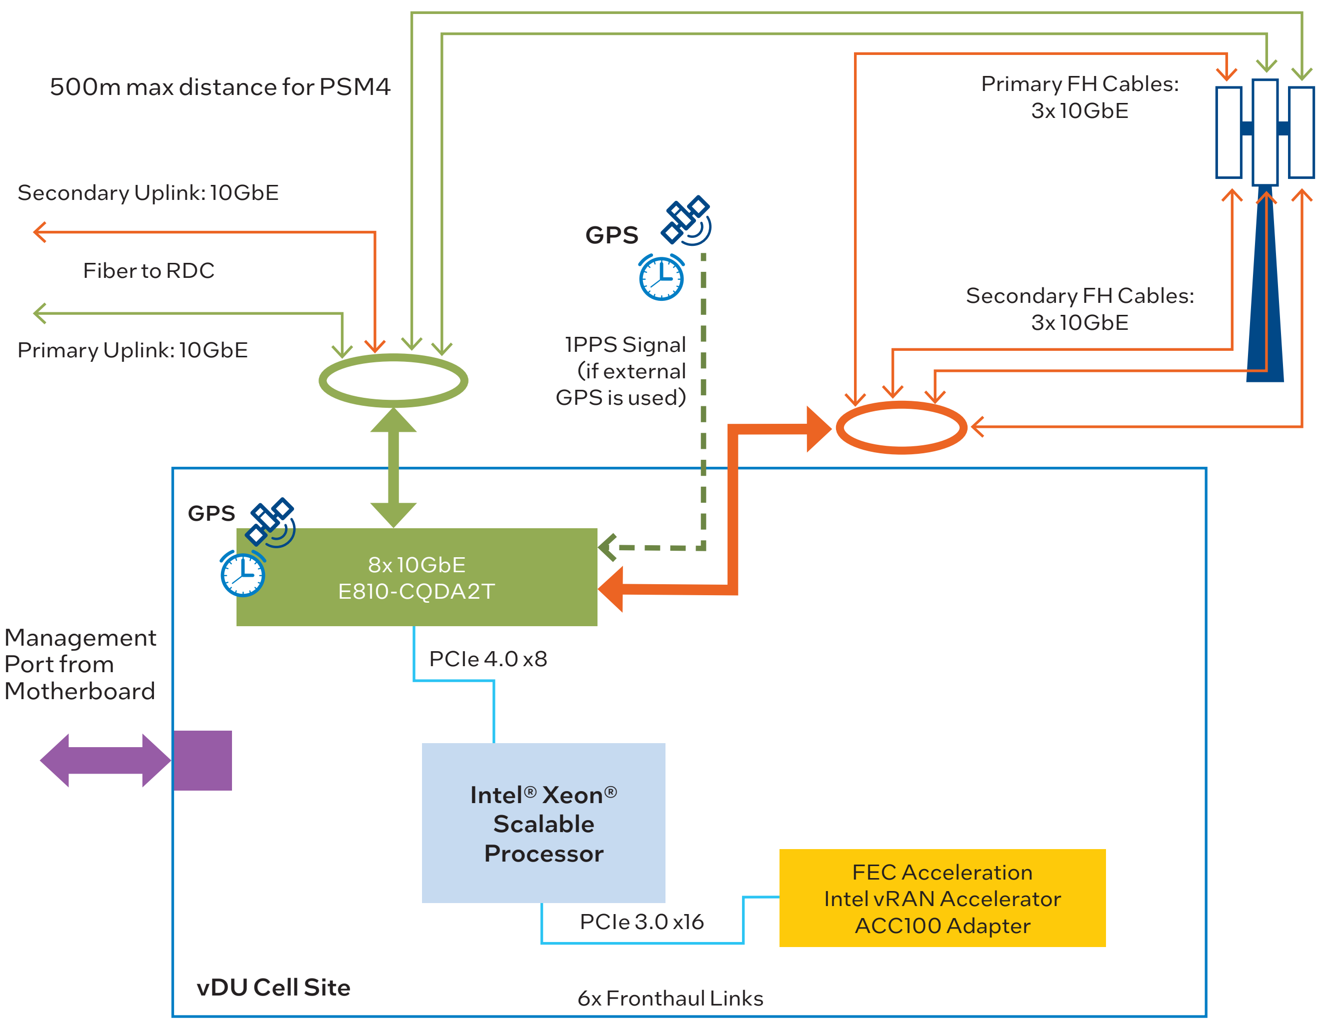

Figure 6. Cell site example with E810-CQDA2T: Six fronthaul links and two uplinks.

When timing precision saves lives: Emergency response location services

A common method of finding the locations of emergency callers to services such as 911, 999 or 112 is to calculate it based on the relative time delays from the end-user device to multiple cell towers. The more precise that timing information is, the more precise the location information based on it can be. In the past, timing delays of up to 250 nanoseconds limited location accuracy to within 100 meters.

Today, timing accuracies are commonly about 125 nanoseconds, providing location accuracy of about 50 meters. CoSPs are currently targeting reduction of the timing disparity to 75 nanoseconds, which could close location accuracy to 30 meters. This improvement will help first responders find emergency callers more quickly when lives are at stake.

Conclusion

Timing synchronization is a critical capability for 5G RAN solutions to enable open, multi-vendor deployments that safeguard the economic viability of CoSP buildouts. Intel Ethernet Network Adapters E810-XXVDA4T and E810-CQDA2T support IEEE 1588 PTP v2 and SyncE to ensure high-precision timing accuracy. These features enable CoSPs to embrace open, flexible networking while delivering outstanding user experiences with high call quality, less delay and better streaming.

Learn More:

www.intel.com/vran-ethernet

www.intel.com/vran