Intel

Intel® Ethernet 800 Series Delivering High Timing Accuracy for 5G vRAN

CORPORATE PARTICIPANTS

Xiaojun (Shawn) Li
Intel – Sales Director, Next Wave OEM & eODM

Sean Lion
Intel – Job Title

Shachi Paithankar
Intel – Job Title

PRESENTATION

Xiaojun Li

Welcome everyone to the Intel Network Builders Insights Series. I'm Shawn Li, Sales Director, Next Wave OEM and eODM, Network Communications sales organization at Intel Corporation. I am your host for today's webinar.

Thank you for taking the time to join us today for our webinar titled Intel Ethernet 800 Series: Delivering High Timing Accuracy for 5G vRAN.

Before we get started, I want to point out some of the features of the BrightTALK tool that may improve your experience. There is a Questions tab below your viewer. I encourage our live audience to please ask questions at any time. Our presenters will hold answering them until the end of the presentation. Below your viewer screen, you will also find the Attachments tab with additional documentation and reference materials, which pertain to this presentation. Finally, at the end of the presentation, please take the time to provide feedback using the Rating tab. We value your thoughts and we will use the information to improve our future webinars.

Intel Network Builders Insights Series takes place live every month. So, please check the channel to see what's coming, and access our growing library of recorded content. In addition to the resources you see here, we also offer comprehensive NFV and 5G training programs through the Intel Network Builders University. You can find the link to this program in the Attachments tab, as well as a link to the Intel Network Builders Newsletter.

Today we are pleased to welcome Shachi Paithankar and Sean Lion. Shachi is a product marketing engineer in Ethernet Products Group. She is responsible for Ethernet network foundation NICs, with a focus on enabling customers to build OCP and OEM customer Ethernet adapters.

Sean is a business development manager for the Telecommunications segment within the Ethernet Products group. He has worked in both the wired and wireless sides of telecommunication and has been in the industry since the first generation of GSM with Omnipoint, Voicestream, T-Mobile, and Intel.

Welcome, Shachi and Sean, and thank you for taking the time to join us today. I will hand it over to Shachi to start off. Thank you.

Shachi Paithankar

Thank you, Shawn. vRAN is the biggest opportunity that we have seen in recent times. As cloud service providers and telco customers prepare for broad vRAN deployments, they are moving to open and disaggregated IP stacks. This includes Intel Xeon processors, enhanced timing, Ethernet adapters, and discrete FEC accelerators. These deployments have key advantages over the closed proprietary systems of the past, which involves scalability, that's being able to support these functions with standard server broadly-available NICs, and fixed-function FEC accelerators. That makes it easy to scale and customize deployments as demand increases or as the use cases evolve. Contrast this with the specialized timing cards and rack units being used in older deployments or FPGAs, and FEC acceleration in early vRAN deployments. Another crucial advantage is removing dedicated timing hardware that simplifies deployments. This helps in lowering the costs by removing this hardware from these deployments, and providing the same functionality on standard servers and standard NICs. Obviously, this is a simplified explanation, but it gives you a very good general idea of a shift that's taking place with 4G and 5G vRAN. With 5G vRAN, it presents an excellent opportunity for us to expand our region to Comms workloads. The advantages that we have to succeed in this space is our platform balance, OSV collaboration, and our products integrated with Intel's FlexRAN solutions.

Now, speaking about platform balance, with vRAN-specific optimizations in Intel Ethernet 800 series network adapters, and Intel ACC100 dedicated FEC accelerators, these platforms enable the CPU to focus on critical vRAN functions.

Regarding OSV collaboration, we have worked closely with the major operating environment vendors to integrate these layered vRAN solutions. There's no one-size-fits-all here, so it's important that customers know a particular vRAN solution will work in a specific operating environment. We are working to close gaps here because the more we do, the easier we make things for our customers. And finally, with our integrated products with Intel's FlexRAN solutions, which are completely validated top-to-bottom vRAN reference designs, customers using these designs can be confident that our adapters will integrate and operate correctly in these multi-ingredient stacks.

I'll hand it over to Sean.

Sean Lion

There are many different challenges in this RAN space, many different interconnects that could be part of a different solution. Thermal challenges. Many of the edge servers have limited PCIe slots for adding cards. There are new protocols at play, whether it be PPPOE or eCPRI. There are new layers of software with operators, containers, in-tree versus out-of-tree drivers, and real-time OSs that all seem to confound some of these different systems and these options for us.

Shachi?

Shachi Paithankar

Sorry, I was muted. So in order to meet the 5G network timing requirements, there is a significant challenge here. The challenge is 5G performance requirements driving precise timing needs from cloud to edge. The legacy proprietary solutions rely on high-cost, purpose-built appliances, and specialized network interface cards, with our Intel Ethernet solutions based on Westport Channel and Logan Beach adapters that we'll go over in upcoming slides, and delivers a Hardware-Enhanced Precision Time Protocol (PTP) and SyncE Ethernet NIC in standard servers. These Intel Ethernet NICs offer precision timing across an entire network at a cost-effective price point for 5G infrastructure scale-out. This helps in reducing solution cost versus the legacy hardware appliance and specialized NIC approach. This approach can also be used in other markets such as industrial, financial, energy, and many more.

Sean Lion

Here you see an illustration of how a more precise timing solution can benefit applications. Let's take a look here at emergency response services. By sensing a timing delay between a cell phone and the different timing towers using a technique known as trilateration or multilateration, the location of a caller's location can be determined. In the past, the tightest radius with which you could identify an emergency caller's location was approximately 100 meters, and that was improved over the generations to 50 meters. But by reducing the disparity in synchronization down to 75 nanoseconds between signals with 5G using the Intel Ethernet NIC, with Hardware-Enhanced PTP, it's possible to reduce this radius even further down to 30 meters. A huge difference when every second counts. It’s a clear illustration of how timing synchronization across the network benefits end-user applications. We'll have more information on trilateration in the back of this presentation.

Shachi Paithankar

This is one of our Intel Ethernet 800 Series NIC solution, formerly known as Westport Channel, with Hardware-Enhanced 1588 PTP and SyncE features. This is a PCIe Gen 4.0 quad-port adapter with maximum throughput of 100-gigabit and maximum speed of 25-

gigabit per port. This adapter has five hardware enhancements, which includes oven-controlled external oscillator that enables precise timing accuracy, and up to four hours of holdover time if the source timing signal is lost.

Synchronous Ethernet, that is the first Intel Ethernet NIC with SyncE. If the user application demands building on legacy hardware, or does not want to fully disaggregate, this gives them the ability to do so. This is how the partial disaggregation on legacy hardware can be achieved.

The third one is the GNSS mezzanine. This is an optional. It supports most GNSS satellite systems. Currently, NavIC is not supported, but we are working on getting it integrated soon.

Fourth is dual SMAs that enable connecting to external timing sources such as external GNSS receivers, or the user can connect to performance-testing equipment such as oscillators.

Fifth, we have the SMB connector that connects to antennas for the internal GNSS unit.

Along with the hardware enhancements, this NIC also has advanced features such as Application Device Queues (ADQ), Dynamic Device Personalization (DDP), and iWARP and RoCE version 2 RDMA support.

Sean Lion

This is a follow-on of the previous product that Shachi just mentioned. This is called the Logan Beach, also known as the Intel Ethernet 800 Series NIC solution, with integrated Hardware-Enhanced PTP and SyncE features. This NIC achieved ship release authorization status in mid-July and is now generally available in the market. It too is a PCIe Gen 4.0 adapter with a max throughput of 100 gigabits. It comes with dual QSFP ports with a max speed of 100 gigabits across those ports. It can support, using our Ethernet port configuration tool, eight ports of 10 gigabits for maximum density, using PSM4 modules and breakout cables. This adapter has the same hardware enhancements as previously mentioned for the Westport Channel NIC.

Shachi Paithankar

Here you can see our Intel Ethernet 800 Series options for vRAN. As you're now aware of Westport Channel and Logan Beach adapters, the other three options that Intel offers are Salem Channel, Tacoma Rapids, and Chapman Beach network interface cards. All 800 Series adapters support IEEE 1588 PTP. However, only Westport Channel and Logan Beach support features like SyncE. Also, oven-controlled external oscillator for tighter accuracy and extended holdover time, and comes with an option of GNSS module. These cards are best for deployments requiring high timing accuracy. Both these adapters, Westport Channel and Logan Beach, have support for a boundary clock implementation, which has the ability to adjust one pulse-per-second signal (1PPS signal) to a master timing source.

Sean Lion

End user devices perform a function called forward error correction. This is a type of encoding to mitigate the effects of noisy transmission channels. When data arrives at the Central Unit or Distributed Unit via the fronthaul network, the CPU handles most of the L1 pipeline functions. Most of these functions can be accelerated by CPU instructions. Dedicated accelerators are much more efficient at FEC processing than a host CPU. So the ACC100 adapter handles forward error correction functions like decoding and matching, for example. Forward error correction is critical for ensuring vRAN service reliability. It's a standard function with different algorithms for 4G and 5G, both defined by 3GPP standards. To reiterate, forward error correction introduces redundancy into a data-bearing signal so that the errors incurred during transmission can be corrected at the receiver. This solution is use-case agnostic, using the same algorithm for multiple workloads. Higher capacity 4G and 5G forward error correction requires some level of acceleration to monetize the CPU for other functionalities.

Shachi Paithankar

This is our Intel vRAN Accelerator ACC100 adapter, formerly known as Maclaren Summit. It is a dedicated 4G/5G FEC accelerator and supports concurrent 4G and 5G FEC acceleration. It's based on Intel vRAN dedicated accelerator ACC100, formerly known as Mount Bryce silicon. It has gone through the same thermal validation as our foundational NICs. It has manageability support such as SMBus, PLDM over MCTP over SMBus. It also supports PLDM 2.1 spec, and this manageability is consistent with traditional Intel Ethernet product lines. We have worked with OEMs to understand applications at the server level, and designed this product to meet their needs. FlexRAN software reference architectures include support for the ACC100 adapter. This support simplifies and speeds up platform validation.

Sean Lion

Many early FEC accelerators used FPGAs like the PAC N3000 card because of the flexibility these parts offered. Power and cost prevented broad scaling with these types of solutions. The dedicated fixed-function ACC100 delivers greater performance, lower power, cost-effectiveness, which makes scaling much simpler. The PAC N3000 and ACC100 share common interfaces: DPDK kernel drivers, FlexRAN code bases, which makes migration simple.

Shachi Paithankar

So in total, these are the Intel Ethernet Products Group vRAN platform ingredients where we have comprehensive Ethernet connectivity and FEC acceleration for vRAN implementations with integrated enhanced timing.

I would just like to hold on this slide just a little because we have gone over this, so just giving the audience some time to consume it.

Sean Lion

I mentioned many different challenges out there. As to the interconnects, we have a lot of new non-datacenter traditional link partners in use. Whether they're cell site routers, or new hubs and switches, or new link partners like radios, many of these have new features that we've not interoperated with heavily prior to this.

Thermal. There are thermal challenges out there. Many of these edge sites have less optimal cooling than typical data centers do. They experience higher daytime temperatures. They're further away from easy cooling sources. OCP versus PCIe form factors as well. Many edge servers now feature OCP slots in them as well. OCP cards can have different cooling profiles and different options in terms of where they are installed in the server, and that can make for some challenges in deployment. Also, 1U servers. 1U servers constrained thermal budgets as well. They could also prevent a large amount of flexibility in the cards that you can install. Therefore, we have come up with a lot of different options in terms of brackets, whether low profile, traditional standard profile, or low profile and high profile options for our various adapter cards.

Limited slots on edge servers. Many of these cards that support multiple applications, or cards with maximum I/O density, can be required to fulfill a given need on a given CU or DU in this RAN space.

The new vRAN-specific protocols that I mentioned earlier like eCPRI and PPPOE have existed for a while in the aggregated RAN system, but now we are dealing with them on the Ethernet basis, and features like DDP can significantly help us here.

There are many layers of software that I mentioned earlier as well, whether they're operators, containers, in-tree versus out-of-tree software solutions, real-time OSs, containerization, immutable OSs, new applications, non-hardened applications, applications that haven't been designed necessarily for open-source solutions, and real-time OSs can all cause problems. Integration is key, and integration partners are key here, and that's what we're trying to get to, is a space where we have good, integrated solutions with each of our ecosystem partners.

Some links here to some of the resources that we mentioned. There are many more available and many more coming down the pipe, but this is a good place to start. The vRAN ecosystem is complex. We've got various ISVs, OSVs, ODMs, OEMs, integrators, suppliers, manufacturers all having to work together. We have to work together to re-aggregate some of these solutions for consumption by our key service providers. The way we arrive there will be by utilizing the best-known methods that we are developing right now, and best configurations for the most stable experience by our customers. Intel Network Builders is now, and will continue to be, a critical part of this solution as we go to market, and as we broaden our base.

Are there any questions for Shachi, or myself?

Xiaojun Li

Thank you, Shachi and Sean, and let's go to the question session. We have some questions here. First, what is the primary driving business case for disaggregated RAN?

Sean Lion

The primary business case for disaggregated RAN is really total cost of ownership. With the footprint that 5G needs to be truly successful, it requires more base stations than we did in the first generation of GSM, more than we did in the second generation. There is a… There's a phenomenon out there known as base station shrinkage, and as we've gone from 1G to 2G, to 3G, 4G, and now 5G, the requirements put on us by each successive generation basically caused the ability of a base station to serve a smaller and smaller footprint, because of things like latency, the latency requirements. Obviously, you can't serve very low latency over very long distances. That's a good example. But that base station requires shrinkage, each generation to perform better and better. So more base stations are required to fill in those gaps in the footprint, which means that the operators have to consider electricity, the capability to run those base stations out to larger and larger audiences of customers, and so, therefore, they need to put more-- they're very cautious about the amount of money they want to spend on electricity-- on numbers of buildings they have to maintain. All those edge data centers. So that's why density is becoming a driving factor, and of course, the more dense you're talking, the more total cost of ownership is key to the disaggregated RAN space.

Xiaojun Li

Great, thank you. And next question. When did Intel first include 1588 PTP?

Sean Lion

Intel's been doing PTP for quite a while. The Intel Ethernet division had products back as far as 2006 that included 1588 PTP. It's been largely used and largely internalized within Intel, and the Internet of Things, and a lot of our Timing-Sensitive Networking (TSN) groups as well, and they've been productized into various products across the industry, but it's not really taken off until this opportunity that we have now in front of us with disaggregated RAN.

Xiaojun Li

Okay, great. Thank you. The next question. How many years has Intel been in the Ethernet business?

Sean Lion

Intel's been in the Ethernet business for over 30 years. It may not seem like that, but many of us that are currently in the Ethernet business, even myself who's been here for 26 years, Ethernet came along within Intel before I did. It's been around for I think 31 years now, and we have been a founding member of the IEEE, and we are a current member of the IEEE standard organization, and many of the different groups within Intel actually have representation within that IEEE as well.

Xiaojun Li

Great. That's all the questions we have today. And thank you, Sean and Shachi, for sharing great information. Thank you all for joining us today, and please give our team a rating for the live recording so we may continually improve the quality of our webinars. Please be sure to join us next time, Wednesday, September 7th, at 8:00 a.m. Pacific time. You can find the link for our Insights Series in the Attachments tab.

Thank you, again. This concludes our webcast. Thank you.