WHITE **PAPER**

cadence®

800G Ethernet MACsec Integration and Verification

By Rianta and Cadence

Ethernet is the interconnect technology of choice for a wide range of applications including (but not limited to) data centers, wireless backhaul, automotive, artificial intelligence (AI), and many other use cases. It is ubiquitous. As bandwidth requirements continue to scale to ever-dizzying heights due to the proliferation of data-intensive applications, developing, integrating, and validating Ethernet in SoC devices is becoming increasingly challenging. This white paper will explore the complexities surrounding the verification of an Ethernet subsystem that integrates MACsec and touch on the challenges of achieving first-time right silicon at current Ethernet nodes of 800G (and beyond). In addition, it will explore the collaboration of Rianta Solutions Inc. ("Rianta") and Cadence to build the verification system with Rianta™ design IP and Cadence® Verification IP (VIP).

Contents

Importance of MACSec in Ethernet

Integrating MACsec technology on Ethernet links provides confidentiality, integrity, and origin authentication using industry-standard AES-GCM-authenticated encryption and decryption algorithms employing 128-bit and 256-bit keys. Providing authenticated encryption at the link level is complementary to higher level end-to-end application-to-application security protocols, with MACsec providing line-rate encryption of network traffic over Layer 2 Ethernet transport services. MACsec is defined in the IEEE 802.1AE-2018 standard and is used in conjunction with higher level protocols providing authentication and cryptographic key distribution. Developing an Ethernet subsystem supporting MACsec is challenging to ASIC/SoC architecture, design, and verification. The proliferation of IEEE standards, IP solutions, domain knowledge, and verification state space that must be covered to ensure first-time right silicon requires careful engineering.

In order to tackle the challenge of Ethernet verification complexity, Rianta adopted Cadence 800G and MACsec VIP, hoping it would not only work as a stimulus generator but also perform the compliance checking for this complex system. It turned out to be an efficient engagement as the DIP+VIP environment was quickly brought up, traffic flowed in and out correctly, and security checking performed as expected. It saved quite a lot of time for engineers to code pin-level BFM and rewrite algorithms.

Challenges of 800G IP and MACsec Verification

A simplified picture of an 800G Ethernet / MACsec subsystem is shown in Figure 1. This subsystem consists of a set of SerDes macros capable of collectively providing 800G of Ethernet line bandwidth, 800G capable PCS/FEC IP, an 800G capable MAC, and 800G capable MACsec IP.

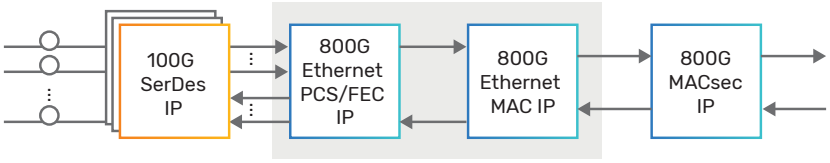

Figure 1: Simplified 800G Ethernet MACsec subsystem

Unencrypted data arrives from the host logic on the bottom right side of the diagram into the 800G MACsec IP where it is associated with a secure entity (SecY), a secure channel (SC), and a secure association (SA) defining the cryptographic policy data for use with this packet, including the cryptographic keys used for encryption and authentication. Aside from defining policy information necessary to facilitate encryption and authentication of the packet, the policy also specifies whether or not the packet should bypass cryptographic processing and be sent in the clear. The MACsec IP typically provides the packet parsing, classification, transformation, and statistics collection features required by the IEEE standard. The MACsec block drives a sequence of encrypted and unencrypted packets to the MAC IP as defined by the policy.

To allow recovery of the original packet at the receiver, the outgoing MACsec packet includes a standard 8B or 16B SecTag and a 16B authentication tag. The data contained within these tags is critical to the overall MACsec operation and successful decryption of the packet and results in an expansion of each packet by 24B or 32B. The encrypted packet from the MACsec block proceeds to the MAC/PCS/FEC IP, where the packet is transformed further to meet Ethernet frame requirements on the line. This includes the insertion of pause frames, preamble and IFG characters, handling of line encoding and alignment markers, and FEC transcoding as required to meet 800G Ethernet standards. The Ethernet frames are finally transmitted on the physical medium by the SerDes macros (See Figure 2).

Figure 2: Ethernet frame format with embedded MACsec

On the receive path, the processing proceeds in reverse order to recover the original unencrypted packet. Packets received on the Ethernet line may or may not be encrypted. Encrypted packets will have an 8B or 16B SecTag inserted into the L2 frame and a 16B authentication tag (ICV) appended to the end of the packet before the Ethernet FCS and padding. This information, along with the packet itself and the incoming port number, is used by the MACsec IP to map the packet to an appropriate SecY, SC, and SA policy. The policy is used to determine if the packet needs to be processed as a MACsec packet or bypassed. MACsec processing on the receive side includes a number of operations as defined by the IEEE 802.1AE-2018 standard

including the actual decryption and authentication process, as well as the control flow required by the standard, anti-replay checking, and statistics. The embedded SecTag and ICV tags may optionally be removed based on the policy before the final unencrypted packet is forwarded towards the host logic.

Developing such a subsystem to operate at 800Gbps in each direction is, indeed, a challenging endeavor. Not only must the design engineer understand all the aspects and complexities inherent in 800G Ethernet, but the designer must also be knowledgeable in cryptographic technologies and their application to Ethernet links. The IEEE standards defining these protocols, while exhaustive, are extremely complex and require years of experience to master. Achieving non-blocking bidirectional performance at 800Gbps throughput is especially difficult given the high data and packet rates that must be dealt with. Assuming an average Ethernet overhead on the line, the maximum packet rate that needs to be processed approaches 1.2 billion packets per second, leading to critical architectural tradeoffs related to packet-per-cycle performance and clock rate. Additionally, the sheer volume of data that must be decrypted on the receive path and encrypted on the transmit path leads to complex tradeoffs between area and power that must be explored. Designing PCS, FEC, MAC, and MACsec IP that operates correctly at 800Gbps is challenging, time consuming, and error prone if not done properly.

Verifying this type of subsystem also requires careful environment planning and execution to ensure first-time success and to minimize schedule duration. It is important to consider the following items when planning this type of verification activity:

- What language will be used to create the verification environment and write the required tests?
- How much third-party VIP can be used to build this verification environment?
- How do I ensure that this subsystem complies with the required standards?
- How extensively do I need to test each piece of IP if it is third-party IP?
- How many resources are required and how long will the subsystem take to complete?
- Do I have the right domain knowledge to ensure first-time success?

Creating a complete environment from scratch to validate the complete subsystem, even if that testing is limited to integration-style testing, is extremely challenging and time consuming. The verification engineer must not only have deep knowledge of a large number of technologies but must also be able understand the interactions and verification state space that exists both within and between the major functional blocks of the subsystem. To succeed, especially at bleeding-edge rates of 800G and beyond, requires a deep knowledge of the protocols and an understanding of the verification methodologies necessary to complete this task in a timely manner. The latest MACsec standards also require more complex knowledge of Ethernet frame formats, as some information in the form of VLAN tags are left in the clear and require knowledge of cryptographic keys and their associated protocols. Developing a verification environment that is capable of stressing the critical aspects of such a system is just as challenging as its design.

The Cadence and Rianta Collaboration

To help expedite the development and validation of an 800G Ethernet subsystem with integrated MACsec, Cadence in collaboration with Rianta have developed a fully integrated IP solution consisting of both the IP design solution itself and an extensive verification suite that can be used for vertical integration into subsystem and full-chip SoC validation environments. The complete solution consists of required components necessary to develop a fully integrated secure Ethernet link at speeds up to 800G (and beyond) for embedded SoCs and fully integrated solutions like PHYs, switches, and routers, but also applicable to lower rate solutions in the automotive and wireless spaces. The design IP consists of fully synthesizable RTL complete with all required lint, coverage waivers, synthesis sdc/tcl files, and documentation collateral as well as SerDes macros. When coupled with the Ethernet and MACsec verification IP provided by Cadence and Rianta, time-to-market concerns are mitigated and development efforts are reduced.

Figure 3 provides an overview of an 800G secure Ethernet IP and verification solution. The overall solution consists of a set of Cadence and Rianta hard and soft IP blocks, which interact with each other seamlessly, and a layered UVM verification environment consisting of Cadence Ethernet VIP and Rianta MACsec VIP. This environment facilitates the complete integration testing required to fully validate a secure Ethernet link. Both the Cadence and Rianta VIP are capable of interoperating with other third-party IP, as well. The complete solution involves the DUT itself, which includes the various IP blocks that comprise the design, and the verification components, which provide packet generation, configuration, prediction/scoreboarding, and self-checking. The figure illustrates a mix of Cadence and Rianta IP and VIP that can be used to implement the complete subsystem.

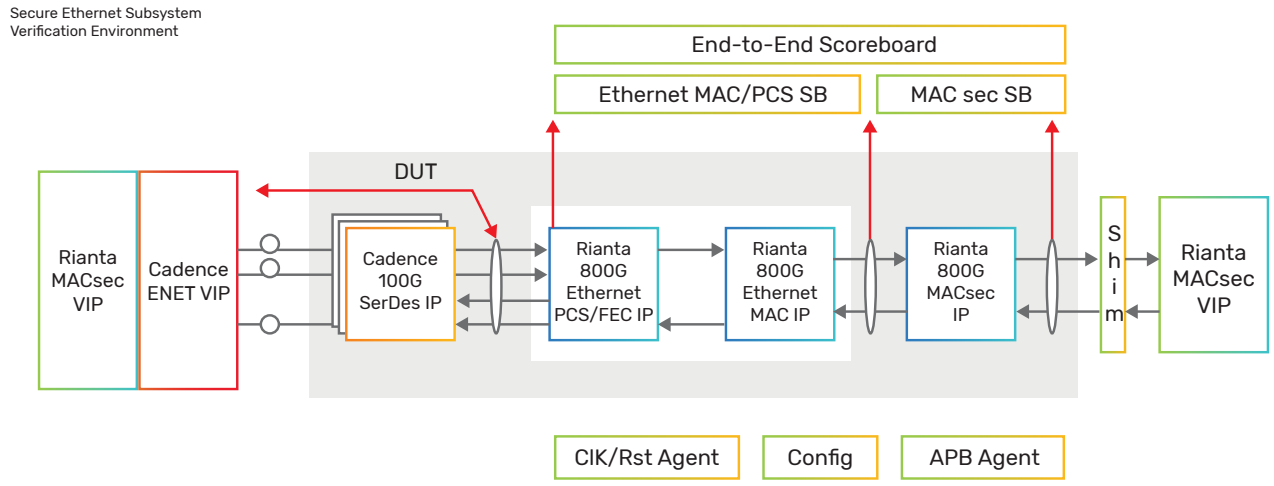

Figure 3: Secure Ethernet subsystem verification environment

The Cadence Ethernet VIP bolts seamlessly to the Rianta MACsec VIP to provide a layered approach for verification purposes. The Rianta MACsec VIP augments Cadence's Ethernet VIP to add necessary encryption, authentication, rekeying, replay protection checks, etc. that are required for integrated MACsec validation. The Cadence Ethernet VIP is capable of driving the serial SerDes directly or the parallel interface into the PCS. On the host-side to the right in the figure, the Rianta MACsec VIP can be configured to provide a loopback function or can be used to actively drive and monitor packets at the MACsec IP host-side interface through an interface shim layer. Additionally, end-to-end scoreboarding and predictors as well as clock/reset agents, configuration, and APB agents are required (and available) to complete the overall verification subsystem.

Rianta IP Features

Rianta 800G MAC/PCS/FEC IP (RSm800™)

- IEEE 802.3-2018 and Ethernet Technology Consortium specs
- 1x400G R4/R8 and 1x800 R8
- RS-528, RS-544, and RS-272 (LL-FEC)
- Fully featured MAC/PCS design
- Supports IEEE 802.3br IET enabling Time Sensitive Networking (TSN)
- IEEE 802.3 Pause + IEEE 802.1Qbb Priority Flow Control (PFC)
- 1588 1- and 2-step timestamping

Rianta 800G MACsec IP (RS_IMCS800™)

- IEEE 802.1AE-2018 standard-compliant core
- 1G-800G bidirectional core supports 100% line-rate authenticated encryption/decryption
- Up to 64 ports and 4K SA scale
- Flexible mapping between ports, SecY, SC, and SA
- AES-GCM 128/256-bit keys + XPN ciphers
- Fully parameterized, pipelined, and channelized cut-through architecture
- Extensive parsing, TCAM-based classification, and statistics support
- Integrated IEEE 802.3br IET support
- MACsec resource pooling across ports
- Constant latency
- Certified

Rianta MACsec VIP (RSVP1000™)

- Complies with IEEE 802.1AE-2018 and IEEE 802.1X (for rekeying)
- Provided as a standalone layer that can bolt onto customer drivers/monitors
- Supports flexible parsing (up to 2 VLANs + SecTag) and TCAM-based classification
- Supports flexible mapping of virtual ports to SecY / SC
- Provides 2 sets of TLM ports for controlled and uncontrolled packets
- Expected MACsec counter increments provided
- Provides default sequences for regular traffic, SA rekey, and error injection
- Integrates seamlessly with Cadence Ethernet VIP

Cadence VIP Features

Cadence Ethernet VIP

- Speeds from 10Mbps to 800Gbps
- Maximum per lane speed 100Gbps/lane
- All MII variants, BaseR, BaseX, BaseT1 Phys
- Multi-port interfaces USGMII, USXGMII, FlexE
- Traffic generation and extraction for range of Transport, Network, and Application protocols
- TSN protocols
- Adopted by leading companies for hundreds of successful verification project signoffs
- Flexible architecture supports Standalone, Full-stack, Subsystem-level verification
- Supports System Verilog, UVM, and e user APIs
- Provides consistent solution (easy to use, integrate, maintain, and upgrade)

Cadence 800G VIP

- Supports 800GBase-R Dual-PCS 32 lanes (25Gbps)
- Programmable PMA bus width allows bypassing SerDes
- Supports PMA (4:1 Mux) -8x106.25G and 16x53.125G
- Supports PMD interfaces -2x400GBASE-DR4 modules
- Multi-lane distribution (MLD) to distribute data from a single media access control (MAC) channel across 2x16 PCS lanes
- Supports NRZ, PAM4, and PAM16 support (supports PMA (4:1 Mux) :: 8x106.25G, 16x53.125G)
- Supports lane reordering and lane skew
- Supports Loopback mode and PRBS pattern generation
- Error injections for most of data items at all sublayer
- Provides consistent solution (easy to use, integrate, maintain, and upgrade)
- Provides compliance checking

Conclusions

This white paper discussed the complexities around developing a secure 800G Ethernet subsystem and the associated project and schedule risks associated with getting it right the first time. The challenge in achieving first-time success is a real concern and can be mitigated by using design and verification IP that is available today from both Cadence and Rianta.

For More Information

Contact - Tim Webster: twebster@rianta.ca and Krunal Patel: krunalku@cadence.com

cadence®

Cadence is a pivotal leader in electronic design and computational expertise, using its Intelligent System Design strategy to turn design concepts into reality. Cadence customers are the world's most creative and innovative companies, delivering extraordinary electronic products from chips to boards to systems for the most dynamic market applications. **www.cadence.com**