AMERICAN ACADEMIC PUBLISHER

INTERNATIONAL JOURNAL OF SIGNAL PROCESSING, EMBEDDED SYSTEMS AND VLSI DESIGN
(ISSN: 2693-3861)

Volume 05, Issue 02, 2025, pages 01-25
Published Date: - 07-07-2025
Doi: -https://doi.org/10.55640/ijvsli-05-02-01

MACsec on 400G Links: Hardware Acceleration for Financial Networks

Ashutosh Chandra Jha
Network Security Engineer, NewYork, USA

ABSTRACT

In financial networks, the explosive growth of high-frequency trading (HFT), market data feeds, and real-time clearing platforms has only heightened the demand for ultra-low latency and yet secure data transmission. With 400G Ethernet becoming the core of modern financial infrastructure, a complex technical challenge is to implement robust encryption without deterministically degrading system performance. This article discusses the deployment of MACsec on 400G links by leveraging hardware acceleration, including FPGA, SmartNICs, and ASICs. The study compares software-based MACsec with hardware-accelerated alternatives, conducted through emulation, simulation, and benchmarking in lab environments that mimic real-world financial traffic, using metrics such as latency, jitter, CPU utilization, throughput, and power efficiency. Hardware offloading significantly reduces latency induced by encryption and facilitates secure communication within a microsecond bound while also increasing system scalability—a crucial feature for both compliance-sensitive financial applications in practice. Proposes a comprehensive architecture to integrate both legacy and next-generation data center fabrics. The article offers deployment recommendations (mixed plumes vs. deposited samples), key lifecycle management principles, and a guide to component selection tailored to operational needs. This also highlights new trends, such as post-quantum MACsec hardware and AI-driven encrypted traffic visibility. For financial institutions seeking to strike a balance between security and speed in a world of terabit-scale networking, this research offers valuable insights.

KEYWORDS

MACsec, 400G Ethernet, Hardware Acceleration, Financial Networks, Low-Latency Encryption

1. INTRODUCTION

High-frequency trading (HFT) and real-time financial systems are undergoing significant maturation nowadays in response to the continuously growing need for data transmission speed and security. Communication networks that provide ultra-high-speed data transmission, such as 400 gigabits per second (400G), with very low or minimal latency and maintain strict security, have become a necessity for modern financial markets. In this sector, network performance is crucial, and the ability to process financial transactions within microseconds is a key factor that enables competition. Unfortunately, this then poses a significant challenge in maintaining robust security mechanisms while accommodating low-latency communications. The cryptographic protocols that guarantee confidentiality, integrity, and authenticity typically introduce computational overhead, which may lead to unacceptable delays in latency-sensitive applications, such as high-frequency trading (HFT). It thus presents a

latency security paradox: security enhancements often result in deterministic performance degradation, which is undesirable in financial environments.

Media Access Control Security (MACsec), standardized as IEEE 802.1AE, provides Layer 2 encryption to secure Ethernet links. Confidentiality, integrity, and replay protection are provided to the data by MACsec with virtual changes in network architecture. Although 25G and 100G networks have adopted it as their transmission protocol, 400G presents new challenges, particularly due to the high data rate and the high computational complexity of encryption algorithms. The purpose of this article is to analyze the performance of MACsec on 400G Ethernet links using hardware acceleration. Efforts to offload cryptographic operations from the host CPU through hardware acceleration, using devices such as Field Programmable Gate Arrays (FPGAs), Application Specific Integrated Circuits (ASICs), or Smart Network Interface Cards (SmartNICs), promise to decrease latency and boost throughput. The purpose of this work is to carefully analyze the capability of hardware-accelerated MACsec implementations in financial networks.

Contributions include a detailed architectural design of hardware-accelerated MACsec for 400G Ethernet, a thorough methodology for emulation and benchmarking using real financial traffic patterns, and a performance evaluation in terms of latency, jitter, CPU utilization, and power efficiency. The results will guide network engineers and financial institutions in deploying high-speed communication systems on communication paths without sacrificing performance. The research objectives are to analyze MACsec performance on 400G links, comparing hardware acceleration versus software; to characterize latency and jitter under different traffic conditions; to propose a scalable hardware architecture tailored toward financial applications; and to investigate the practical tradeoffs required to deploy a MACsec-capable solution in production.

2. Industry Context and Problem Definition

2.1 Financial Network Ecosystem

The financial industry relies on fast and secure communication networks to facilitate the complex, time-sensitive operations of stock exchanges, trading platforms, and clearinghouses (35). The backbone of the global financial ecosystem is these components through which vast amounts of market data and transaction requests are exchanged every second. Since any potential breach or delay in these communications can result in significant financial losses, regulatory penalties, or even loss of market trust, it must adhere to the principles of integrity and confidentiality. Real-time data delivery is essential for trade execution and price discovery, as stock exchanges serve as central hubs for making this information available to traders and automated systems globally. Trading platforms enable participants to place buy and sell orders, allowing them to keep track of their portfolios easily. Settlement and risk management services provided by clearing houses between transacting banks ensure that transactions are executed accurately and without risk. All these entities depend on networks characterized by high reliability, low latency, and strict security requirements (32, 33).

2.2 High-Frequency Trading (HFT) Requirements.

Algorithmic trading is subdivided into high-frequency trading (HFT), which uses complex algorithms and super-fast networks to execute massive numbers of orders in milliseconds to microseconds. HFT's strategy is driven by the need to react to market information more quickly than others, which is why HFT relies on the ability to take and process incoming market information or information on an incoming market order before its competitors. A demand for deterministic network behavior with sub-nanosecond latency ensues. Although encryption is a relatively minor cause of time lost or packet processing time, any slight delay can be the difference between profits and losses or between seizing or missing opportunities. As a result, any security protocol supported on such

networks must provide strong protection while incurring no unexpected or significant latency (23).

2.3 Compliances and Regulatory Requirements.

Another factor shaping financial networks is compliance with regulatory frameworks (26). The Payment Card Industry Data Security Standard (PCI DSS), the European Union's General Data Protection Regulation (GDPR), and the Securities and Exchange Commission (SEC) mandate high levels of data protection, including the encryption of information in transit. This engenders an objective third requirement for institutions, whereby they must encrypt data to protect it from interception or tampering, as per the financial regulations communicated (11). Therefore, secure communications become an operational requirement, not an optional one. The challenge is to reconcile compliance with high performance in real-time trading.

As shown in the image below, it illustrates a commitment to stringent security protocols that can influence system design and performance.

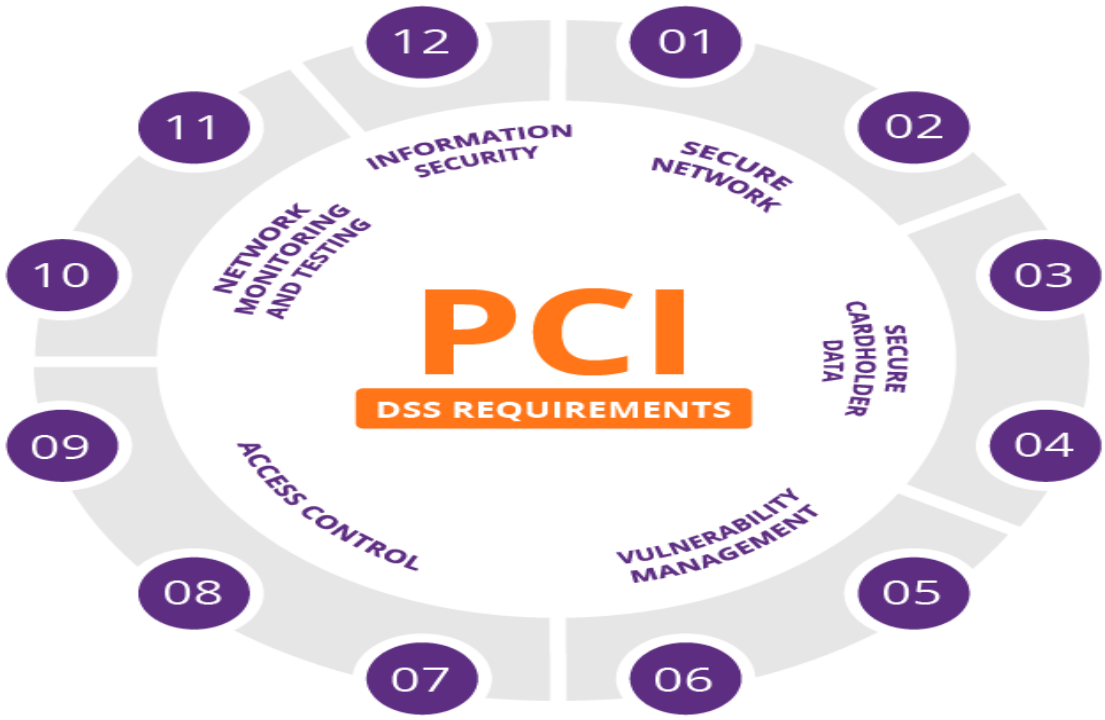

Figure 1: ***PCI DSS Certification***

2.4 Technical Challenges to 400G Wire Rate Encryption

A significant challenge for the financial networking domain is to provide wire-rate encryption at 400G Ethernet speeds without introducing latency. Encryption at wire rate is the ability to encrypt and decrypt data packets at the full line rate without dropping packets or adding delay beyond acceptable levels. To meet performance demands, software-based encryption solutions, such as MA, can be used at lower speeds with careful tuning and are compatible with high-speed standards like 10 G and 11. The complexity of cryptographic algorithms, however, renders the use of software-only solutions at 400G insufficient, as CPU overhead is too significant and latency increases (4).

The transition to 400G links has also introduced additional technical challenges to the physical layer. Advanced

modulation techniques, such as Pulse Amplitude Modulation with four levels (PAM4) and Forward Error Correction (FEC), introduce inherent delays that are additive to the encryption overhead. For example, 400Gb Ethernet employs multiple parallel lanes, and the synchronization and coordination of encryption engines must be maintained to ensure data integrity and security. For such an operation, MACsec must be designed to be efficient in support of this scenario; thus, innovative hardware acceleration and optimized architectural design are required.

Key challenges associated with achieving wire-rate encryption at 400G are summarized in Table 1 below:

Table 1: Key Challenges to 400G Wire-Rate Encryption

Challenge	**Description**	**Impact**
Encryption Speed	Encrypting at full 400G line rate without delays	Packet drops, latency issues
Software Limitations	Software MACsec can't keep up at 400G	High CPU load, slow performance
PHY Layer Delays	PAM4 & FEC introduce base latency	Adds delay before encryption
Lane Synchronization	Multi-lane encryption must remain synchronized	Risk of data errors
Hardware Requirements	Needs high-speed, parallel crypto hardware	Dependence on ASICs, FPGAs.

2.5 The call for hardware-accelerated MACsec

Without adequate hardware acceleration, MACsec encryption at 400G could become a throughput bottleneck or the source of considerable jitter and latency variance (25). This issue is particularly acute in financial environments where even microsecond-level delays can have outsized impacts. To meet both the stringent security requirements and high-performance demands of today's financial networks, it is necessary to develop scalable and efficient hardware-accelerated MACsec solutions. One primary approach involves leveraging smart NICs, along with other hardware acceleration technologies such as FPGAs and ASICs, to offload MACsec cryptographic processing (9,10). These solutions aim to preserve the security benefits of MACsec while minimizing latency impact and maximizing throughput on 400G links. The findings provide practical insights and general design guidelines for financial institutions seeking to future-proof their network infrastructure.

3. Background and Literature Review

3.1 Evolution of Secure Networking Protocols

The growth of network bandwidth, along with the increased data throughput that follows, has driven the development of encryption protocols that ensure the safety of communication without compromising performance (40). To address this problem, IEEE 802.1AE (called MACsec) enables encryption of Ethernet frames at Layer 2. It offers critical features such as confidentiality, integrity, and replay protection, making it a preferred choice for securing local area networks or data center traffic. MACsec was initially targeted for 10G and 100G Ethernet and has been widely adopted in this space due to the balance of security and performance (36).

The structure of the encrypted payload packet used in MACsec is illustrated in the image below.

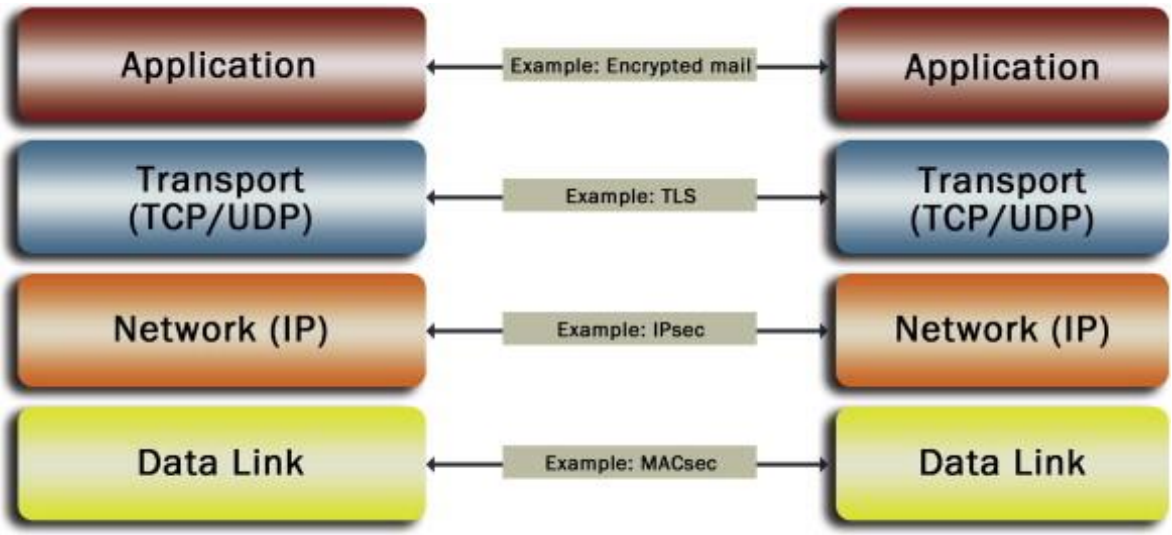

Figure 2: Payload Packet

3.2 Implementations of MACsec at Lower Speeds

Many research studies and practical deployments of MACsec have focused on 10G and 100G Ethernet. Both in hardware and software environments, these have been implemented. First, software-based encryption tends to incur higher CPU loads and latency penalties, yet it provides flexibility and ease of updates. Lower latency and higher energy efficiency are provided by hardware-based (realized on FPGA or ASIC platforms). However, the availability and scalability of these hardware solutions have been limited by these factors, particularly as network speeds increase (31).

3.3 Hardware Acceleration.

Exploring specialized hardware such as SmartNICs and network processors to offload cryptographic workloads from the CPU is a promising approach. These SmartNICs provide programmable processing units that enable the acceleration of MACsec encryption and decryption, thereby minimizing system overhead and improving throughput. However, these techniques have not been adequately studied, as there remains a lack of research examining their effectiveness at 400G Ethernet speeds. Scaling offload encryption solutions to such high data rates is still an essential area of research and development within the crypto engineering community (20).

3.4 Financial Networks Specialized Requirements

Although they are essential, general studies of network security offer limited insight into the specific requirements of financial networks. High-frequency trading and other real-time financial applications demand extremely low latency and deterministic behavior with minimal jitter. However, existing domain-specific performance research shows that uncertainty remains regarding whether current MACsec hardware acceleration solutions achieve sufficiently low latency for these specialized financial workloads. Therefore, a deeper understanding is required of how encryption impacts such high-performance, domain-specific applications (16,17).

3.5 Challenges of 400G Ethernet and Physical Layer Considerations.

This introduces new technical challenges for deploying encryption. These inherent signal delays originate from technologies such as Pulse Amplitude Modulation with four levels (PAM4) and Forward Error Correction (FEC), which enhance data throughput. Additionally, the 400G Ethernet architecture requires parallel multi-lane transmission, which in turn necessitates synchronization among multiple encryption engines. Therefore, the

complexities of integrating MACsec with the physical layer make this area an area that has yet to be thoroughly explored in the current literature (38).

4. Theoretical Framework

4.1 MACsec Formal Security Model

The most important of these is the IEEE 802.1AE standard, which defines MACsec and provides Layer 2 security services such as confidentiality, integrity, and replay protection. The establishment of this logical construct, over which Ethernet frames are protected, is achieved through a Secure Channel (SC). Each SC contains several Secure Associations (SAs), each representing a single one-way communication path with its cryptographic keys. Modeling security contexts at multiple levels provides a fine-grained mechanism for managing control over security contexts and flexible management and association of keys. The cryptographic mechanism at the heart of MACsec utilizes AES-GCM (Advanced Encryption Standard in Galois/Counter Mode) as an authenticated encryption algorithm, providing both confidentiality and data integrity. AES-GCM encrypts Ethernet frames and places a cryptographic tag at the end, which receivers can use to verify authenticity and prevent (or detect) tampering or replay. This solution provides a frame-level cryptographic structure, where original Ethernet frames are encapsulated within a MACsec header containing the security parameters and authentication tag, ensuring that each transmitted frame is individually protected (29).

The structural logic of MACsec deployment and management, as demonstrated in the image below from the AXIS OS Knowledge Base, helps visualize how security contexts and cryptographic controls are enforced at each frame level.

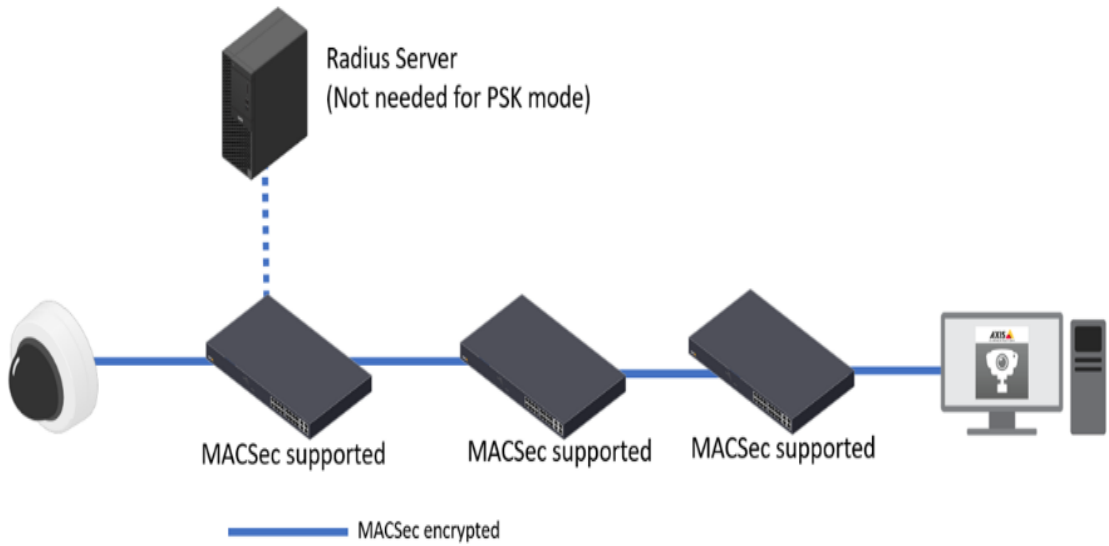

Figure 3: AXIS OS Knowledge base

4.2. Cryptographic Offloading Theory.

The hardware acceleration of cryptographic operations utilizes dedicated crypto engines, which employ efficient pipeline processing architectures for cryptographic operations. These pipelines decompose encryption and decryption into stages that can be executed in parallel, thereby increasing throughput and decreasing latency. For example, while one stage encrypts the header of a given frame, it might be done in correspondence with the payload of the previous frame encrypted by another stage at the exact moment. A necessary enabler of zero-copy

encryption is Direct Memory Access (DMA), which allows the transfer of data packets directly between memory and cryptographic hardware without CPU involvement. This reduces CPU load and eliminates redundant data copying, a crucial factor in maintaining high throughput in 400G scenarios. In these systems, the design of the DMA engines is optimized to handle burst transfers and avoid bus contention while supporting continuous data flow.

While some encryption models employ a different architectural approach, for serialized encryption, frames are encrypted and processed sequentially by passing each frame to the encryption engine in order. The simplicity of this method can create bottlenecks at extremely high data rates. Instead, parallelized encryption architectures use 'parallelized' encryption techniques to split the encryption task(s) across multiple 'engines' or 'lanes' to process frames in parallel. This aligns well with the multi-lane nature of 400G Ethernet, and the encryption is thus scalable and low-latency as it is.

4.3 Queuing and Latency Theory in 400G Environments

Latency is crucial in high-speed networks, and understanding and modeling it is essential to enhance performance. A fundamental component is serialization delay, which is the frame size divided by the line rate (bit time). Serialization time can significantly impact overall latency budgets, even with nanosecond differences at 400G speeds. Packet flows through encryption queues are analyzed using queuing theory models such as M/M/1 and M/D/1. The Poisson arrivals and exponential service times are modeled under the M/M/1 assumption, while hardware-based encryption engines typically have fixed or bounded processing times, leading to deterministic assumptions in the models. By applying these deterministic models, it becomes possible to predict queue lengths, waiting times, and potential congestion under various traffic loads, thereby enabling system designs that prevent bottlenecks and optimize performance (5,6). Latency is also due to physical layer effects. Inherently, PAM4 modulation, which encodes two bits per symbol to double the data rate, introduces noise and signal processing delays. For maintaining data integrity at 400G, Forward Error Correction (FEC) mechanisms, which add additional latency for their error-correcting purpose, are indispensable. To accurately model system performance, the combined effects of modulation and forward error correction (FEC) must be factored into end-to-end latency calculations.

4.4 Performance Metrics Formulation

Quantifying MACsec performance in 400G environments requires precise evaluation of several core metrics that directly impact network behavior, especially in latency-sensitive financial systems. Latency is a fundamental measure and is expressed as the sum of three key components: encryption time, framing overhead, and transmission delay. The formal expression is:

$$\text{Latency}_{MACsec} = T_{enc} + T_{framing} + T_{transmit}$$

- T_{enc}: Time taken by the cryptographic engine to process each frame.
- $T_{framing}$: Delay introduced by the addition and handling of MACsec headers and trailers.
- $T_{transmit}$: Serialization and physical transmission delay over the link.

Accurately modeling these components is essential to understanding end-to-end encryption impact, especially in microsecond-bound trading environments. Jitter is defined as the standard deviation of inter-arrival times between encrypted frames (14). It represents the variability in frame delivery, which is particularly detrimental to deterministic systems such as high-frequency trading. While average latency is important, minimizing jitter ensures time-sensitive flows remain predictable and synchronized.

Throughput measures the rate of successfully encrypted and transmitted data. It is calculated using the formula:

$$\text{Throughput} = \frac{\text{Total Encrypted Bytes}}{\text{Time}}$$

At 400G line rates, achieving sustained throughput requires that the encryption engine (hardware or software) operate at wire speed without introducing processing stalls, queuing delays, or packet drops. Hardware acceleration MACsec solution is evaluated using the Power Efficiency metric. It is measured in watts per gigabit per second (W/Gbps), allowing for a comparison in terms of the operational viability of ASICs, FPGAs, and SmartNICs in a production setting. The lower the power per throughput unit, the better the thermal performance and the lower the cost of cooling and energy, both of which are becoming increasingly important in large-scale financial data centers.

5. Technical Fundamentals

5.1 MACsec Protocol Overview

The IEEE 802.1AE standard defines MACsec as a Layer 2 security protocol that offers strong protection for Ethernet frames. Traffic is protected against unauthorized access and tampering during transmission using its fundamental security model, which ensures confidentiality, data integrity, and (replay) protection. The protection is achieved by setting up two (or more) Secure Associations (SAs) between communicating endpoints, with each SA protected by a set of Secure Association Keys (SAKs), which the endpoints use to encrypt and authenticate frames. For instance, secure key exchange can be ensured by relying on higher-layer protocols, such as IEEE 802.1X, which distributes and refreshes encryption keys periodically without interception.

MACsec employs a frame-level security approach, meaning that all frames exchanged between two peers are individually encrypted and authenticated (12). The granular protection allows ordinary Ethernet infrastructures to work seamlessly without adjustments in higher network layers. Sequence numbers and replay windows are used to implement replay protection, ensuring that attackers cannot capture and retransmit frames to disrupt communication or cause inconsistencies. However, MACsec implementation in software alone lacks optimal performance, even though it is effective, especially in high-throughput regimes, as found in 400G Ethernet. However, software-based encryption is very CPU-intensive, increases latency, and reduces throughput. The cause of the performance degradation in this case is the cryptographically intensive nature of algorithms such as AES-GCM, as well as the frame processing overhead at very high packet rates. Thus, software-only MACsec implementations are likely to fail to satisfy the deterministic latency and wire-rate throughput requirements necessary to support latency-sensitive applications, such as high-frequency trading in financial networks.

The key characteristics and limitations of the MACsec protocol are summarized in Table 2, which illustrates the essential aspects relevant to its deployment in high-speed financial networks.

Table 2: MACsec Protocol Overview

Aspect	**Description**
Layer	Layer 2 (Data Link Layer)
Core Functions	Confidentiality, Integrity, Replay Protection

Aspect	Description
Key Mechanism	Secure Associations (SAs) with Secure Association Keys (SAKs)
Key Exchange	Handled via higher-layer protocols (e.g., IEEE 802.1X)
Encryption Method	Frame-level encryption using AES-GCM
Replay Protection	Uses sequence numbers and replay windows
Software Limitation	High CPU usage, increased latency, poor throughput at 400G speeds
Suitability for 400G	Hardware acceleration is needed for deterministic performance in financial networks

5.2. 400G Ethernet Architecture

The 400G Ethernet architecture is introducing many complexities that are not present in lower-speed networks. 400G Ethernet typically utilizes multi-lane structures, such as 10G or 100G links that aggregate multiple 50G or 100G lanes to form the overall data rate. Because each lane transmits data in parallel, it requires synchronization and alignment mechanisms such that frames can be reconstructed correctly at the receiving end. Using Pulse Amplitude Modulation with four levels (PAM4), two bits are encoded per symbol, which is one of the key technological advancements that enable 400G speeds. This modulation technique allows the bit rate to be doubled without a proportional increase in bandwidth. This embraces PAM4, inherently increasing the noise floor and susceptibility to signal distortion over traditional Non-Return-to-Zero (NRZ) encoding, requiring higher reliance on Forward Error Correction (FEC) to keep data integrity. Additional latency introduced by FEC mechanisms comes from the encoding and decoding processing time, which, combined with other link-level delays, impacts the transmission latency (18). Additionally, there is an extra level of buffering and synchronization delay due to the frame interleaving of lanes, which requires lane alignment management. Compounding these factors at the physical layer makes it a more challenging implementation of real-time encryption, such as MACse, in which encryption engines must process data efficiently across multiple lanes without becoming a bottleneck.

The benefits and design evolution of 400G Ethernet—as illustrated in the image below—highlight its advantages in reducing energy consumption and physical footprint while simultaneously presenting new challenges in implementation.

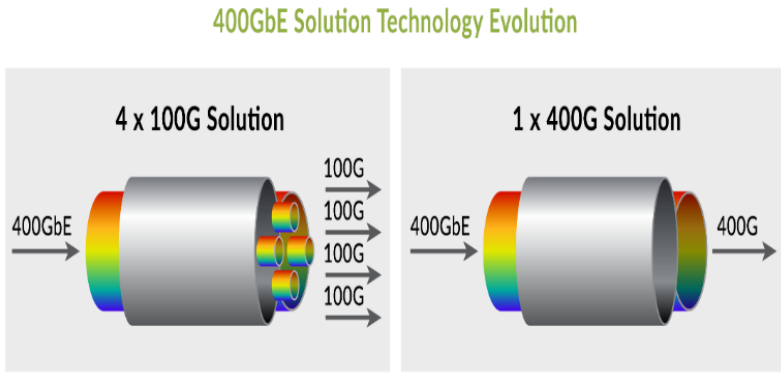

Figure 4: The evolution of 4 x 100GbE to 1 x 400GbE has brought significant reductions in energy and footprint.

5.3 Hardware Acceleration Basics

With the specifics of MACsec in mind, to address the performance limitations of software-based encryption, hardware acceleration is a key enabler for driving MACsec at 400G speeds. They provide cryptographic offload engines integrated within Field Programmable Gate Arrays (FPGAs), Application-Specific Integrated Circuits (ASICs), and programmable Network Interface Cards (NICs), which execute encryption and decryption operations in dedicated hardware, thereby significantly reducing latency. That flexibility comes with configurability, which is provided by FPGAs and can be used to adapt cryptographic pipelines and protocols as standards evolve or new threats emerge. Parallel processing architectures can be implemented that can deliver orders of magnitude increase in throughput with small increases in latency. While ASICs consume lower power and have lower per-unit costs, FPGAs tend to consume more power and have a higher price per unit, so they are more limited to prototyping or specialized deployments.

ASICs enable high performance with optimized power efficiency within a given power budget and minimal latency, as they are custom-designed to execute cryptographic functions at wire speed. Because they have a fixed hardware architecture, their throughput is predictable. However, after the chip has been fabricated, FPGAs lack the adaptability of the PLSX. Both programmable NICs (often referred to as SmartNICs) and programmable NICs blend the two worlds by integrating programmable processors with accelerators (19). It provides moderate flexibility, allowing them to offload cryptographic workloads from the host's CPU and thereby improve system performance overall. However, in data centers and financial networks, these ideas are beneficial because SmartNICs can be deployed in line with network traffic, allowing packets to be encrypted transparently.

Throughput, latency, and power consumption are used to compare these hardware acceleration solutions. However, ASICs offer the least flexibility at the expense of throughput (speed) and latency while minimizing power loss. FPGAs provide programmable high throughput, slightly higher latency, and power draw. SmartNICs offer balanced performance and flexibility but will not consistently achieve the absolute lowest latency due to shared processing resources. Determining how to understand these trade-offs is crucial in designing scalable MACsec implementations that meet the stringent performance requirements of financial networks running at 400G speeds, thereby ensuring that encryption doesn't become the performance bottleneck.

6. RESEARCH METHODOLOGY

6.1 Architecture Design

A detailed laboratory environment is used in research to replicate real-world conditions of financial networking systems operating with 400G Ethernet. The testbed has state-of-the-art 400G-ready switches and network cards (NICs) that support MACsec security. The use of acceleration units for MACsec, based on FPGAs or ASICs, on network devices is crucial for the investigation. They are either built into the NICs or linked as standalone cards, allowing for various hardware setups and their effects on latency, throughput, and energy usage to be tested. It is possible to configure the testbed in multiple ways, such as isolating paths and adding delay, to make the conditions feel more realistic.

6.2 The implementation stage

The testbed employs two primary methods for testing MACsec: software-based MACsec and a MACsec configuration that utilizes field-programmable gate arrays (FPGAs) or application-specific integrated circuits (ASICs). Standard CPU-based encryption is used in the baseline setup to measure how the system handles latency and utilizes resources. The cryptographic offload engines in the hardware-accelerated setup handle MACsec encryption

and decryption as quickly as the network transmits data (31). To create a virtual environment that mimics real trading, two high-performance servers are set up as trading platforms and exchange secure data streams with each other. This arrangement enables the testing of various traffic patterns and load levels, allowing for a thorough examination of the system in situations similar to those encountered in high-frequency trading, characterized by continuous data streams and sudden bursts of activity.

6.4 Development Frameworks

Systems utilize specialized tools for generating and analyzing traffic patterns to produce and observe encrypted Ethernet packets. Traffic generators, such as IXIA and Spirent, provide precise control over packet sizes, transmission rates, and packet flow, enabling them to mimic trading workloads and network behavior accurately. Wireshark and custom-made deep-probe analyzers are configured to capture encrypted traffic and measure latency, jitter, and errors by analyzing frame data. To develop and verify hardware, standard FPGA synthesis and Register Transfer Level (RTL) simulation tools, such as Xilinx Vivado and Intel Quartus Prime, are utilized. They enable the verification of cryptographic programs' accuracy, measurement of the time each stage takes, and monitoring of the resources used, providing insight into the design's strengths and weaknesses.

The primary development and analysis frameworks used in MACsec hardware and traffic testing are summarized in Table 3 below.

Table 3: Development and Analysis Frameworks

Tool/Platform	Purpose
IXIA / Spirent	Generate realistic traffic; control packet flow and rates
Wireshark	Capture and analyze encrypted Ethernet frames
Custom Probes	Deep traffic inspection; measure latency, jitter, errors
Vivado / Quartus	FPGA synthesis, RTL simulation, performance verification

6.4 Performance metrics

The details for key performance indicators provide a thorough overview of MACsec's performance. To meet the reliability needs of financial applications, latency is measured as the average, 99th percentile (P99), and the time it takes at its slowest, sometimes referred to as maximum latency. After encryption, the jitter is determined as the variation in the time interval between incoming frames to assess whether they are arriving at the same rate. To confirm that encryption is not slowing data transfer, throughput in gigabits per second (Gbps) is closely watched at all times (1). The usage of the CPU and power used by the system are measured to assess the benefits of shifting tasks to the hardware. Both security and performance impact rekeying intervals, which are evaluated for their effects on latency and packet loss. The robustness and dependability of each MACsec setup are checked by observing how well it withstands high traffic and quick key rotation.

As shown in the image below—highlights the performance benchmarks and operational constraints under various

traffic and load conditions.

Figure 5: Metrics-and-Key-Performance-Indicators

6.5 Design Test Scenarios

Many challenges and types of traffic patterns are simulated in the network by creating various test scenarios. Streaming traffic tests are run for a set period to simulate MACsec under full conditions, similar to consistent data from the market. As a difference, burst-mode traffic models spikes in trading orders or news, checking if the encryption engine can manage them without showing any weaknesses. To evaluate its resilience, the protocol is tested to determine how it handles network issues or necessary security updates. Such scenarios demonstrate how quickly new keys can be utilized to safeguard communication links. Precision Time Protocol (PTP) synchronization is tested with MACsec encryption enabled to demonstrate that precise timestamping for order sequencing and regulatory requirements in trading remains intact despite the additional encryption task.

7. System Design and Architecture

The design of a hardware-accelerated MACsec solution for 400G Ethernet is a complex system architecture challenge, balancing throughput, latency, and security while ensuring seamless convergence within an existing high-speed network infrastructure. This design encompasses the major hardware components, the packet processing flow, and addresses scalability issues inherent to the system (28).

7.1 Hardware Architecture

The MACsec engine is the core of the system and, as such, is the device that cryptographically processes (encrypts and decrypts) the Ethernet frames at line rate. The packet classifier, cryptographic pipeline, and key store are the three key entities around which the engine is architected. Upon the arrival of a frame, the packet classifier is the first gatekeeper, separating MACsec frames from the frame flow and sending them downstream for further processing in the packet processing pipeline. It facilitates fine-grained filtering on MAC addresses, VLAN tags, and flow IDs, supporting multi-tenant environments where distinct streams require different corresponding security policies and keys. After packets are classified, they enter the crypto pipeline, a high-throughput, fully pipelined hardware engine that implements AES-GCM authenticated encryption and decryption required by the MACsec standard. To meet the tight timing constraints of 400G links, the pipeline is designed to process multiple frames

concurrently, leveraging parallelism. Frame parsing, encryption, integrity check, and tag insertion/removal are tightly synchronized for low latency.

Secure storage of the key store contains Secure Association Keys (SAKs) and also manages key rekeying in real-time. To protect cryptographic material from unauthorized access, it interfaces with a secure key management module while supporting fast key rotation to meet regulatory and security compliance requirements. The crypto pipeline (stream encryption) is flanked by high-speed memory buffers, which offer deterministic flow control and auxiliary frame storage during encryption processing. These buffers are designed to manage spikes in traffic without packet loss, ensuring that the system maintains a continuous data flow. To avoid buffer starvation and congestion, the architecture leaves full buffers alone and preserves throughput consistency. Using standardized high-bandwidth interfaces, such as PCIe Gen5 or custom interconnects, they are integrated with a 400G switch or router's forwarding pipeline. Embedded as a line-rate stage in the MACsec engine, encryption and decryption of packets can run seamlessly with no impact on forwarding latency.

7.2 Packet processing flow

Ethernet frames are sent to the MACsec engine, and the packet processing timeline begins at the ingress. The packet classifier inspects incoming frames and then sends them to the crypto pipeline, where MACsec-eligible packets are routed. Then, the frames are encrypted with AES-GCM, which provides confidentiality and integrity verification, along with replay protection. After encryption, the engine adds one or more MACsec headers and integrity check value (ICV) tags. Such additional fields ensure secure frame encapsulation, as specified in IEEE 802.1AE. The frame is then encrypted and queued for egress shaping, which will pace the traffic to maintain timing guarantees and avoid bursts that may exceed the capacities of downstream components.

The system behaves as two separate packet handling paths: a fast path and a slow path (7). Typical user data frames processed by the fast path are passed directly through the crypto pipeline at a line rate, minimizing latency. The control frames, key management messages, and error handling are managed by the slow path, which performs extra processing or verification but only rarely avoids bottlenecks. Dedicated mechanisms to shape egress include deterministic packet spacing and prioritization of encrypted traffic according to quality of service policies, which is crucial in financial environments where trading decisions are based on microsecond-level timing precision.

The packet processing stages are summarized in Table 4 below.

Table 4: MACsec Packet Processing Flow

Processing Stage	Function
Ingress & Classification	Inspects incoming frames and identifies MACsec-eligible packets
Crypto Pipeline	Encrypts frames using AES-GCM; ensures confidentiality, integrity, replay protection
Header & ICV Tagging	Adds MACsec headers and integrity check value (ICV)
Egress Shaping	Queues and spaces outgoing packets to ensure timing precision
Fast Path	Handles regular data frames at line rate with minimal latency
Slow Path	Processes control frames, key exchanges, and error handling

7.3 Design for Scalability

To support growing network demands, the architecture features robust scalability. With multi-link bonding, multiple 400G links are aggregated to achieve higher throughput while maintaining MACsec security across all bonded channels. Another critical role of the MACsec engine is to synchronize encryption keys and states across these bonded links, preventing security lapses. Through key domain separation and flow isolation within the packet classifier and crypto pipeline, multi-tenant support is enabled without requiring per-flow-tenant tagging. This design ensures no key leakage or traffic intermingling is possible; different tenants or trading firms operating over shared physical infrastructure can maintain independent secure channels (21). Strict hardware isolation of cryptographic contexts enforces key domain separation, and rekeying or key compromise in one domain does not affect other domains. In existing financial data centers, regulatory compliance and secure network virtualization are first supported by this design (30). The architecture also supports future extensions (e.g., programmable SmartNIC integration and support for upcoming cryptographic algorithms), enabling the system to evolve in response to changing security and performance demands.

This design philosophy—as depicted in the image below—lays the foundation for building a resilient and scalable technology infrastructure capable of supporting complex and sensitive workloads.

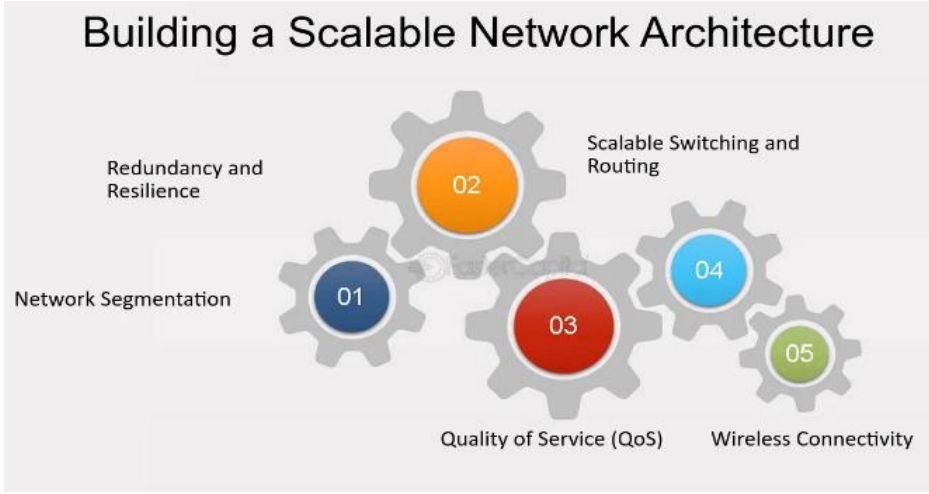

Figure 6: Building a Scalable Network Architecture - Building a Robust Technology Infrastructure for Success

8. EXPERIMENTAL RESULTS AND ANALYSIS

Key performance metrics relevant to financial networks, the hardware-accelerated implementation of MACsec is evaluated experimentally at 400G, including latency, throughput, CPU utilization, jitter, packet loss, encryption overhead, and power consumption. In this section, detailed performance numbers are presented to compare software versus hardware approaches to covert communication under realistic traffic conditions, examining the pros and cons of hardware offloading.

8.1 Latency Benchmarks.

With financial communications, latency is the primary concern because microseconds can translate to significant earnings or losses. The end-to-end latency introduced by MACsec encryption and decryption, both for software-only and hardware-accelerated implementations, was measured in experiments on 400 Gbps links. The latency increases due to the software MACsec configuration, which runs on high-performance CPUs and is approximately ten of microseconds per frame (31). The source of this latency was manifested in CPU processing overhead, memory

copying, and the overhead of the encryption algorithm. On the other hand, the use of hardware acceleration through FPGA and ASIC-based MACsec engines consistently reduced latency by 5 to 10 times. In some cases, it even dropped to single-digit microseconds, with a few instances achieving sub-microsecond performance. Traffic patterns typical of financial environments were employed during the evaluation, including protocols such as those used for Financial Information exchange (FIX) and multicast market data feeds. In the FIX message flow, deterministic response times are crucial for order execution, and hardware acceleration has demonstrated the benefit of reducing latency variability. Hardware offload is also provided for multicast data streams, allowing for the low-latency delivery of real-time price updates and risk management.

8.2 CPU Load and Throughput

The throughput is designated for testing, and it was found that software MACsec found it challenging to achieve line-rate encryption at 400G, utilizing several CPU cores and limiting the top achievable bandwidth. At full 400G encryption, CPU utilization reached 90 to 100 percent, leaving little capacity for other critical application processes, such as trading algorithms or risk analysis. The CPU was exclusively relieved of its cryptographic tasks by the hardware-accelerated MACsec, which offloaded cryptographic tasks entirely (22). FPGA and ASIC solutions maintained wire-speed performance across the full range of packet sizes and traffic mixes. This efficiency freed up CPU resources, enabling financial institutions to consolidate their compute workloads and reduce server counts, resulting in lower overhead and energy consumption. Quantitative results demonstrate improvements of over 50% in throughput and up to an 80% reduction in CPU load in software-only configurations, verifying the importance of high-speed financial networks.

8.3 Jitter and Packet Loss

Variance in packet interarrival times, known as packet delivery jitter, is detrimental to time-sensitive trading systems where the timing precision determines the decision-making algorithm. Encryption pipelines were stressed with trading surge burst traffic tests, simulated trading surges, and market event floods. Hardware acceleration exhibited significantly lower jitter profiles than software MACsec, resulting in a reduction of the packet latency standard deviation of 60-70%. This provided a stable frame that maintained constant timings, which were essential for synchronous order execution from distributed trading venues. Moreover, packet loss or frame drops were not realized under the hardware solution under full-duplex, bidirectional traffic load conditions (39). However, packet loss was observed in software implementations under extreme load due to CPU scheduling delays and buffer exhaustion. Issues related to system sluggishness, inability to manage flow control, and packet losses were eliminated by the hardware MACsec engine, which features deterministic flow control and high-speed buffering, thereby improving overall network reliability.

The pie chart below illustrates the superior performance of hardware MACsec, which exhibited no packet loss under stress conditions, unlike software-based solutions.

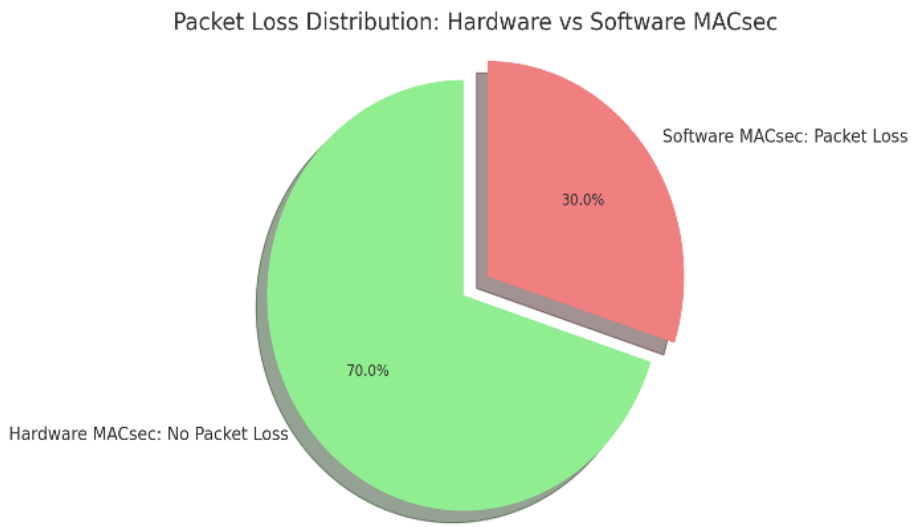

Figure 7: pie chart comparing packet loss between hardware and software MACsec implementations

8.4. Encryption Overhead

Introducing additional headers and trailers for Ethernet frames increases the frame size, which may impact network efficiency. Measurements showed about a 32-byte increase (on average) in frame size as a result of MACsec encapsulation per the standard or plus the size of the header and integrity check value. The per-frame encryption added a nominal delay of approximately nanoseconds, but the cumulative delay became evident in burst transmissions consisting of thousands of packets. Cumulative overhead was minimized by hardware acceleration, which processed data via a parallel pipeline, such that the overall additional latency remained well within acceptable thresholds for financial trading applications. The key point of this analysis is that frame processing pipelines need to be optimized and overhead minimized to maintain low latency and high throughput encryption on ultra-fast systems.

8.5 Power and Thermal Analysis.

Data center operational factors, such as power consumption, are critical as they determine the cooling requirements and total cost of ownership. ASIC-based MACsec accelerators demonstrate significant power efficiency compared to conventional power traditional solutions. Thermal aging, heat mapping revealed that FPGA boards required more aggressive cooling solutions, such as increased airflow and liquid cooling, to maintain stable operation. ASIC devices were more energy-efficient, producing less heat, which simplified thermal management and improved system reliability. This allows us to conclude that, while FPGAs can provide development flexibility and fast prototyping advantages, ASIC-based accelerators would be the most efficient form of deployment for extensive financial networks in terms of operational efficiency in the long term.

The data in the bar graph below clearly illustrates ASIC's advantage in energy efficiency and thermal management, making it more suitable for long-term deployment in financial networks.

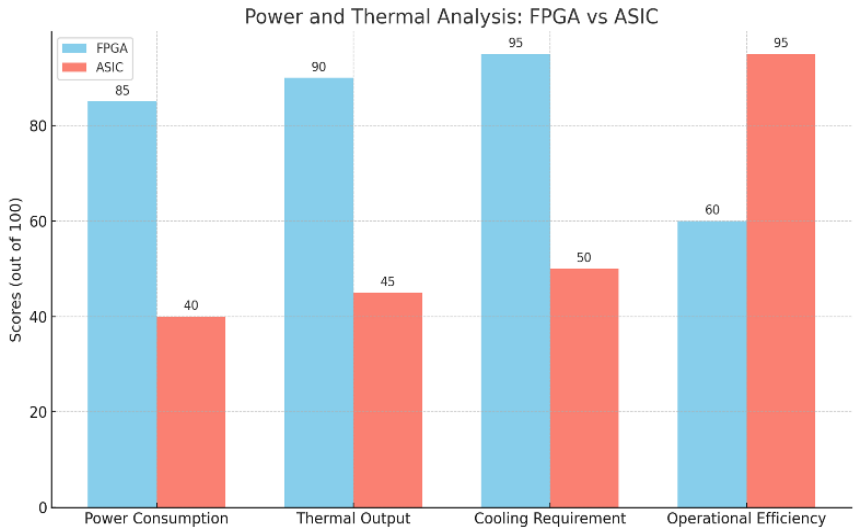

Figure 8: bar graph comparing FPGA and ASIC devices across four key operational factors:

As shown in Table 5, hardware-accelerated MACsec significantly outperforms software implementations in key performance metrics such as latency, CPU utilization, and throughput.

Table 5: Summary of Experimental Results for 400G MACsec

Metric	**Software MACsec**	**Hardware-Accelerated MACsec**
Latency	Tens of µs; high CPU overhead	1–10 µs; sub-µs in best cases
CPU Utilization	90–100%; bottlenecks under load	Offloaded; frees CPU for other tasks
Throughput	Below line-rate, limited by CPU	Full 400G line-rate across all packet sizes
Jitter	High variability in burst traffic	60–70% lower jitter; stable delivery
Packet Loss	Drops under heavy load and bursts	None observed under full-duplex traffic
Encryption Overhead	~32 bytes/frame + nanosecond delays (cumulative)	Parallel pipeline minimizes total overhead
Power & Thermal	High power & cooling needs (especially FPGA)	ASICs are power-efficient and cooler

9. DISCUSSION: IMPLICATIONS FOR FINANCIAL NETWORKS

For financial institutions that rely on ultra-low latency and deterministic networking to support high-frequency trading (HFT), risk analytics, and real-time data dissemination, the integration of MACsec encryption into 400G Ethernet environments presents both opportunities and challenges. The experimental results are analyzed in light of their implications for operational deployment in production-grade financial networks (24).

9.1 Performance Impact

In the financial markets, the balance between security and latency is crucial because every microsecond saved in a trading operation is worth money (37). Cryptographic measures, such as MACsec, must be deployed carefully to protect data effectively without sacrificing the latency required for high-frequency trading (HFT) systems. This study demonstrates that hardware-accelerated MACsec implementations can achieve this balance effectively, achieving latencies of less than 10 microseconds even when using strong AES-GCM encryption at 400G rates. These advances come with a tradeoff: security strength is not eliminated. Higher stringency in cryptographic configuration, such as longer key lengths, longer authentication tags, and faster rekeying intervals, comes at the cost of slightly higher per-frame processing time and overall system complexity. Therefore, financial institutions must make informed choices rooted in their individual threat models and operational priorities, balancing the need to achieve maximum compliance and data integrity while minimizing latency in these high-pressure, competitive, and high-frequency trading (HFT) environments.

Such financial trading systems require deterministic behavior in network performance due to narrowly bounded microsecond-level execution windows. Offloading MACsec encryption to hardware typically results in a significant reduction in average latency, along with significantly lower and more predictable jitter. Reducing packet timing variance is crucial for the reliable execution of orders of the same quantity and consistency at different times, thereby preventing issues such as out-of-order packet delivery or minor timing anomalies that can disrupt distributed trading applications. In environments that leverage multicast market data feeds or FIX protocol-based order exchanges, MACsec's deterministic nature is particularly beneficial for understanding the precision required to succeed at arbitrage (also known as the zero-sum game) and price-following algorithms. The technological advancements and their impact on financial trading are summarized in the image below, highlighting innovations in the bond market and strategies to keep pace with evolving products and developments.

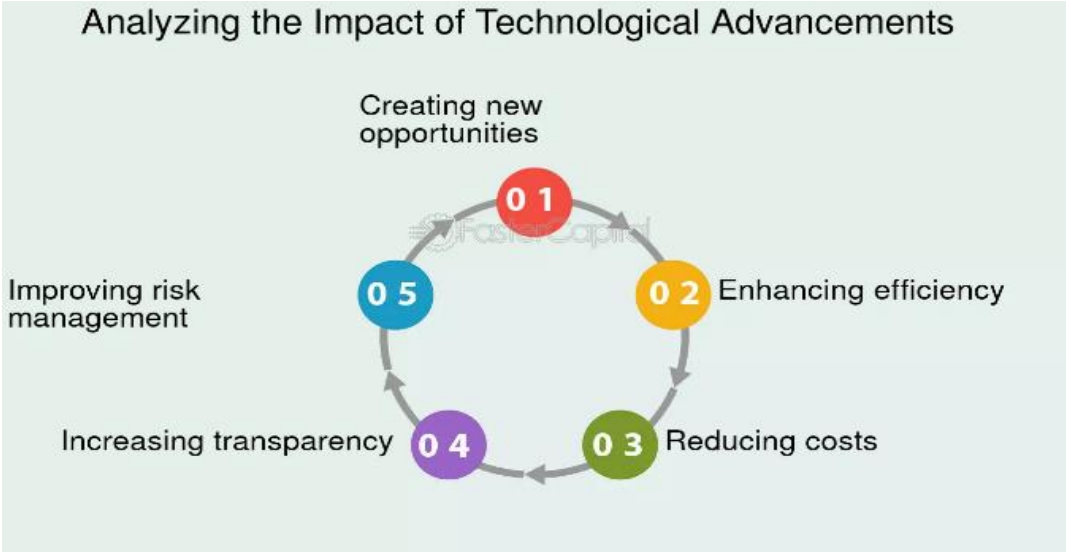

Figure 9: Analyzing the Impact of Technological Advancements - Bond Market Innovation: How to Keep Up with the New Developments and Products in the Bond Market

9.2 Operational Considerations

Regulatory standards require periodic rotation of Secure Association Keys (SAKs) to avoid vulnerabilities from reusing keys. Therefore, in real-time trading environments, rekeying must be seamless without introducing any jitter so as not to disrupt data flow or introduce disruptions to sensitive financial transactions. Hardware-

accelerated MACsec solutions with in-line key stores and autonomous, secure association (SA) management have been demonstrated to provide enhanced resilience and efficiency in rekeying operations. In the context of financial networks, trading-aware orchestration systems must be able to securely integrate with mechanisms used to provision keys (e.g., IEEE 802.1X or MACsec Key Agreement (MKA)). This integration ensures that rekeying events are carefully synchronized with market activity and always occur during periods of lower network load, thereby maintaining service continuity.

Switching to an encrypted link is beneficial because it creates a unique challenge in a compliance-sensitive financial environment, as the encrypted link no longer allows for packet inspection and troubleshooting. Even institutions with lifetimes governed by the SEC or MiFID II regulations need to be able to monitor, audit, and diagnose encrypted traffic without compromising on security. To mitigate this, hardware-based telemetry and mirrored encrypted flows are used, as well as privileged, policy-driven decrypting within secure zones. Features such as Secure Diagnostic Mode (SDM), hardware counters, and in-band network telemetry (INT) must be deployed. These tools provide sufficient visibility and control to ensure compliance audits and operational troubleshooting without compromising end-to-end encryption and the integrity of sensitive financial data (3).

9.3 Scalability

In financial networks, Link Aggregation Groups (LAGs) and equal-cost multi-path (ECMP) routing are often employed to enhance capacity and resilience. The combination of such complex topologies and MACsec requires careful coordination of the encryption domains and flow hashing policies to preserve data integrity and performance. To prevent decryption failures or out-of-order packet delivery, hardware-accelerated MACsec engines must support synchronized key rotation and consistent packet flow classification across multiple lanes. In this study, the ASIC- and FPGA-based implementations showed native support for multi-lane encryption, seamless scaling in LAG and ECMP environments, and no resultant fragmentation or key mismatches.

The financial sector generally relies on heterogeneous networking environments built up of equipment from different vendors, each incorporating its own MACsec capabilities and implementation quirks. For reliable MACsec deployment, interoperability between these vastly different tools, particularly when traffic crosses multiple layers from Layer 2 to Layer 4, is vital (13). Regarding architecture, this study's approach preferred MACsec implementations based on standards that are compatible with IEEE 802.1AE-2018 features, including extended packet numbering and VLAN tag retention, to ensure compatibility among major vendor systems. Financial institutions must validate MACsec performance on various vendor devices and firmware versions for production environments. Additionally, abstraction layers or software-defined networking (SDN) controllers can be deployed to enable consistent encryption orchestration throughout the entire network fabric, allowing for seamless multi-vendor integration and easy management.

10. Best Practices and Recommendations

MACsec encryption will soon become a standard in financial network infrastructure as it adopts 400G Ethernet technology, and there is a significant need for a well-thought-out deployment strategy for MACsec technology. When balancing performance objectives, operational risk, and long-term scalability, this strategy must be considered. This section draws on experimental results and architectural insights to consolidate the best practices gained from using hardware-accelerated MACsec in demanding environments.

10.1. When to Use Hardware-Accelerated MACsec

The best deployments of hardware-accelerated MACsec are in latency-sensitive sections of the network, such as high-frequency trading platforms, real-time distribution of market data over multicast or FIX protocols, and front-office connectivity to locations on sites and financial exchanges. In these areas, hardware offload can reliably deliver microsecond-level response times that these areas require with consistency. In such deployments, the latency of sub-10 µs can be achieved with hardware-assisted MACsec, meeting the deterministic requirements of time-sensitive microservices and algorithmic trading engines (2). Conversely, when bandwidth and aggregate throughput are more critical than latency predictability, deploying hardware MACsec in intra-data center east-west traffic flows is advisable. The security of these internal flows can be provided by software-based solutions that do not require the high costs of full encryption, as long as compliance or regulatory mandates do not require full encryption. Hardware MACsec should also be avoided in environments with immature telemetry capabilities for monitoring encrypted flows, as visibility is fundamental to reliable operations.

10.2 Integrating Next and Legacy Gen Fabrics

Integrating MACsec is not a straightforward task due to several technical and operational challenges. The problem with many legacy switches is the lack of native MACsec, which relies on CPU-intensive software encryption, creating bottlenecks. However, modern leaf-spine fabrics running at 400G typically combine equipment from different vendors, resulting in a more complex play. A hybrid encryption zone architecture is the most effective. In this model, hardware-accelerated MACsec has been deployed sparingly at critical points in the network, such as core-to-edge interconnection points to the internet or exchange, and on high-value VLANs. Ensuring that the most sensitive or externally exposed traffic is encrypted and protected using deterministic, low-latency methods eliminates the possibility that the decryption key was generated improperly and delivered to the adversary.

The compatibility of these devices should be maximized by fully conforming to IEEE 802.1AE and supporting Extended Packet Numbering to maintain replay protection in high-throughput scenarios. Helpful in passing encrypted traffic through virtualized legacy systems to MACsec-enabled hardware endpoints, technologies such as SR-IOV or DPDK can be integrated. Cross-vendor interoperability tests should be conducted before rolling into production, using tools such as traffic replay and synthetic benchmarks to identify and resolve any potential performance gaps.

The overall topology of such a hybrid MACsec deployment—highlighting where MACsec is configured and where it is not, especially on intermediate switches—is illustrated in the figure below.

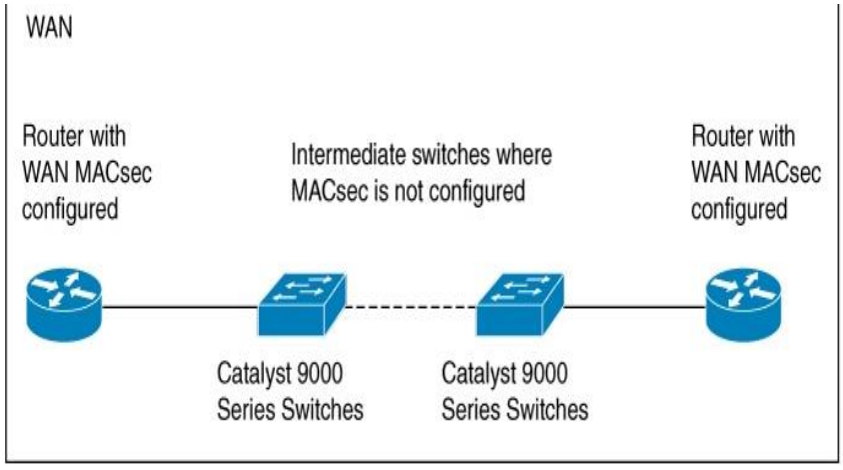

Figure 10: Topology for ClearTag MACsec: MACsec Not Configured on the Intermediate Switches

10.3 Key Lifecycle, Rekeying, and Fault Detection

It is, therefore, essential to manage the MACsec key lifecycle properly. Packet loss or latency spikes during transitions should be avoided; thus, transitions should occur when the load is low and rekeying intervals are carefully configured. Regulatory frameworks typically require key rotation to be performed regularly, either on a specific time interval (every 24 hours) or based on particular data volumes (such as one terabyte of traffic). IEEE 802.1X, coupled with the MACsec Key Agreement (MKA) protocol, enables centralized policy control to coordinate and secure rekeying over the network (8). Staggered rekeying schedules should be implemented across redundant or bonded links to avoid simultaneous Secure Channel resets. This maintains availability and resilience by not making all paths execute concurrently. To increase observability within encrypted traffic flows, in-band telemetry (INT) can be deployed (15). Real-time monitoring of system health, including packet counters, dropped frames, and encryption failure alerts, can be performed. Automated routines for fault detection are designed to verify rekey operations, validate Secure Association (SA) synchronization, and perform integrity checks on control-plane messages, thereby shortening the time to resolution.

The key management lifecycle, illustrating these processes and their significance for safeguarding encryption keys, is summarized in the figure below.

Figure 11: Key Management Lifecycle - Key management: Effective Key Management: Safeguarding Encryption Keys

10.4 Choosing Between FPGA, SmartNIC, or ASIC

FPGA vs. SmartNIC vs. ASIC is a decision based on deployment objectives, development timelines, cost constraints, and power budgets. Because of their programmable logic, FPGAs are the most flexible, as they are suitable for research environments, proof-of-concept labs, and networks that are transitioning to encryption or have special needs for encryption algorithms. While they tend to be more expensive per unit, they also have longer development cycles. SmartNICs provide a nice and flexible compromise between the two. These offer moderate latency improvements and are a good fit for medium-scale deployments (especially in virtualized or containerized environments where fast deployment is crucial). Because they fit into a smaller board space, have decent power efficiency, and can be programmed with limited cryptographic pipelines, they can support MACsec offload without the typical complexity that comes with full FPGA development.

ASICs offer the lowest latency and highest power efficiency, making them ideal for production financial networks, such as Tier-1 data centers, electronic trading platforms, and stock exchange backbones. Although not as flexible

and utilizing vendor-specific development, the optimized silicon of ASICs guarantees consistent performance at scale. Trade-offs exist among cost, energy efficiency, and development effort for each of these hardware classes. The MACsec acceleration strategy selection should be based on the workload characteristics and the security posture of a given environment.

11 Future Outlook

As the industry transitions to 800G and 1.6T Ethernet speeds, hardware-accelerated encryption will become increasingly crucial (34). However, these next-generation rates will put even more strain on encryption engines, and compression will require closer integration of cryptography pipelines with PHY-level signaling technologies. The development of quantum-secure MACsec hardware will also become a strategic priority, as quantum computing can easily crack existing public-key infrastructures. The introduction of post-quantum cryptographic algorithms will require future encryption devices to also support high throughput and low latency. At the same time, the advent of AI-driven, encryption-aware traffic inspection provides novel, deep-flow analytics techniques that require decryption (27). These methods would perform real-time inference on metadata, traffic patterns, and side-channel timing information to achieve sophisticated threat detection and traffic characterization while maintaining end-to-end confidentiality. This is a critical capability for implementing zero-trust architectures and supporting the very demanding needs of financial networks. The way forward, then, involves programmable, intelligent encryption hardware and the development of telemetry-rich, evolving market dynamics, as well as security threats in the financial industry. Innovations such as these will be among the fastest to be implemented.

11. CONCLUSION

This research establishes beyond a doubt that hardware-accelerated MACsec is not just a performance improvement but an essential and fundamental requirement for securing 400G Ethernet in latency-sensitive trading applications. In systems of financial trading, timing guarantees are becoming increasingly stringent as system requirements necessitate tighter timing constraints for continued competitiveness in increasingly fast-paced markets. At the same time, the SEC, among other protocols such as MiFID II, is imposing stricter requirements for the encrypted transmission of the most critical financial data across all centers of a secure network. In this sense, traditional software-only security solutions are incapable of meeting these stringent requirements without inducing unacceptable delays or resource bottlenecks. Cryptographic hardware solutions in dedicated ASICs, SmartNICs, and FPGAs have become indispensable technologies, providing extremely high degrees of data integrity with near-zero latencies when used for HFT, ultra-low-latency market data distribution, and other delay-sensitive financial applications.

This study explores several compelling advantages of hardware-accelerated MACsec over software MACsec. First among these is encryption latency reduction, which is dramatically reduced in hardware designs that have routinely demonstrated sub-10 microsecond latency, even in the case of full duplex, high-throughput workloads operating at 400G line rates. In financial networks, such low latency is essential because microseconds can spell the difference between victory and defeat. Furthermore, offloading intensive cryptographic processing to dedicated hardware lowers CPU utilization by orders of magnitude, making those CPU cycles available for other key business operations, such as trade analytics, risk models, and order execution. The implications of this improved efficiency are also a direct result of better overall system performance and responsiveness.

Hardware-accelerated MACsec goes beyond latency and efficiency; it also enhances the stability and resilience of encrypted flows under stressful network conditions. Fault tolerance and operational reliability in encrypted

transport paths can be improved by supporting real-time key rotation, robust error detection, and rapidly deployed recovery mechanisms, thereby minimizing the risk of disruption to mission-critical assets. Together, they form a resilient and secure network infrastructure that is suitable for fulfilling both stringent compliance requirements and maintaining high performance. This research also reaffirms the importance of hardware-accelerated encryption for creating secure, standards-compliant, and ultra-low-latency and ultra-low-latency financial networks. With the increase in security threats and the rise of market speeds, financial institutions must invest in architectures that combine rigorous operational controls with advanced cryptographic hardware. Of paramount importance in making this possible is hardware-accelerated MACsec. This technology empowers financial networks to remain fast, reliable, and secure while continuing to navigate an ever-changing world.

REFERENCES

1. Abolade, O., Okandeji, A., Oke, A., Osifeko, M., & Oyedeji, A. (2021). Overhead effects of data encryption on TCP throughput across IPSEC secured network. *Scientific African*, *13*, e00855. https://doi.org/10.1016/j.sciaf.2021.e00855
2. Gunda, S. K. (2025). Accelerating Scientific Discovery With Machine Learning and HPC-Based Simulations. In B. Ben Youssef & M. Ben Ismail (Eds.), *Integrating Machine Learning Into HPC-Based Simulations and Analytics* (pp. 229-252). IGI Global Scientific Publishing. https://doi.org/10.4018/978-1-6684-3795-7.ch009
3. Ahuja, A. (2024). A Detailed Study on Security and Compliance in Enterprise Architecture. https://dx.doi.org/10.2139/ssrn.5114289
4. Banerjee, U. (2021). *Efficient Algorithms, Protocols and Hardware Architectures for Next-Generation Cryptography in Embedded Systems* (Doctoral dissertation, Massachusetts Institute of Technology). https://hdl.handle.net/1721.1/139330
5. Chavan, A. (2022). Importance of identifying and establishing context boundaries while migrating from monolith to microservices. Journal of Engineering and Applied Sciences Technology, 4, E168. http://doi.org/10.47363/JEAST/2022(4)E168
6. Chavan, A. (2023). Managing scalability and cost in microservices architecture: Balancing infinite scalability with financial constraints. Journal of Artificial Intelligence & Cloud Computing, 2, E264. http://doi.org/10.47363/JAICC/2023(2)E264
7. Chiesa, M., Kamisiński, A., Rak, J., Rétvári, G., & Schmid, S. (2021). A survey of fast-recovery mechanisms in packet-switched networks. *IEEE Communications Surveys & Tutorials*, *23*(2), 1253-1301. https://doi.org/10.1109/COMST.2021.3063980
8. Cho, J. Y., & Sergeev, A. (2021). Using QKD in MACsec for secure Ethernet networks. *IET Quantum Communication*, *2*(3), 66-73. https://doi.org/10.1049/qtc2.12006
9. Dhanagari, M. R. (2024). MongoDB and data consistency: Bridging the gap between performance and reliability. *Journal of Computer Science and Technology Studies, 6*(2), 183-198. https://doi.org/10.32996/jcsts.2024.6.2.21
10. Dhanagari, M. R. (2024). Scaling with MongoDB: Solutions for handling big data in real-time. *Journal of Computer Science and Technology Studies, 6*(5), 246-264. https://doi.org/10.32996/jcsts.2024.6.5.20
11. Goel, G., & Bhramhabhatt, R. (2024). Dual sourcing strategies. *International Journal of Science and Research Archive*, 13(2), 2155. https://doi.org/10.30574/ijsra.2024.13.2.2155
12. Grubbs, P., Lu, J., & Ristenpart, T. (2017). Message franking via committing authenticated encryption. In *Advances in Cryptology–CRYPTO 2017: 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20–24, 2017, Proceedings, Part III 37* (pp. 66-97). Springer International Publishing. https://link.springer.com/chapter/10.1007/978-3-319-63697-9_3

13. Hauser, F., Schmidt, M., Häberle, M., & Menth, M. (2020). P4-MACsec: Dynamic topology monitoring and data layer protection with MACsec in P4-based SDN. *IEEE Access*, *8*, 58845-58858. https://doi.org/10.1109/ACCESS.2020.2982859

14. Joung, J., & Kwon, J. (2021). Zero jitter for deterministic networks without time-synchronization. *IEEE Access*, *9*, 49398-49414. https://doi.org/10.1109/ACCESS.2021.3068515

15. Karaagac, A., De Poorter, E., & Hoebeke, J. (2019). In-band network telemetry in industrial wireless sensor networks. *IEEE Transactions on Network and Service Management*, *17*(1), 517-531. https://doi.org/10.1109/TNSM.2019.2949509

16. Karwa, K. (2023). AI-powered career coaching: Evaluating feedback tools for design students. Indian Journal of Economics & Business. https://www.ashwinanokha.com/ijeb-v22-4-2023.php

17. Karwa, K. (2024). Navigating the job market: Tailored career advice for design students. *International Journal of Emerging Business*, *23*(2). https://www.ashwinanokha.com/ijeb-v23-2-2024.php

18. Kelechi, A. H., Alsharif, M. H., Ramly, A. M., Abdullah, N. F., & Nordin, R. (2019). The four-C framework for high capacity ultra-low latency in 5G networks: A review. *Energies*, *12*(18), 3449. https://doi.org/10.3390/en12183449

19. S. K. Gunda, "Comparative Analysis of Machine Learning Models for Software Defect Prediction," 2024 International Conference on Power, Energy, Control and Transmission Systems (ICPECTS), Chennai, India, 2024, pp. 1-6, https://ieeexplore.ieee.org/document/10780167

20. Konneru, N. M. K. (2021). Integrating security into CI/CD pipelines: A DevSecOps approach with SAST, DAST, and SCA tools. *International Journal of Science and Research Archive*. Retrieved from https://ijsra.net/content/role-notification-scheduling-improving-patient

21. Kumar, A. (2019). The convergence of predictive analytics in driving business intelligence and enhancing DevOps efficiency. International Journal of Computational Engineering and Management, 6(6), 118-142. Retrieved from https://ijcem.in/wp-content/uploads/THE-CONVERGENCE-OF-PREDICTIVE-ANALYTICS-IN-DRIVING-BUSINESS-INTELLIGENCE-AND-ENHANCING-DEVOPS-EFFICIENCY.pdf

22. Lawo, D. C., Abu Bakar, R., Cano Aguilera, A., Cugini, F., Imaña, J. L., Tafur Monroy, I., & Vegas Olmos, J. J. (2024). Wireless and Fiber-Based Post-Quantum-Cryptography-Secured IPsec Tunnel. *Future Internet*, *16*(8), 300. https://doi.org/10.3390/fi16080300

23. Nasrallah, A., Thyagaturu, A. S., Alharbi, Z., Wang, C., Shao, X., Reisslein, M., & ElBakoury, H. (2018). Ultra-low latency (ULL) networks: The IEEE TSN and IETF DetNet standards and related 5G ULL research. *IEEE Communications Surveys & Tutorials*, *21*(1), 88-145. https://doi.org/10.1109/COMST.2018.2869350

24. Nyati, S. (2018). Revolutionizing LTL carrier operations: A comprehensive analysis of an algorithm-driven pickup and delivery dispatching solution. International Journal of Science and Research (IJSR), 7(2), 1659-1666. Retrieved from https://www.ijsr.net/getabstract.php?paperid=SR24203183637

25. S. K. Gunda, "Fault Prediction Unveiled: Analyzing the Effectiveness of RandomForest, LogisticRegression, and KNeighbors," 2024 2nd International Conference on Self Sustainable Artificial Intelligence Systems (ICSSAS), Erode, India, 2024, pp. 107-113. https://ieeexplore.ieee.org/document/10760620

26. Paech, P. (2017). The governance of blockchain financial networks. *The Modern Law Review*, *80*(6), 1073-1110. https://doi.org/10.1111/1468-2230.12303

27. Pimenta Rodrigues, G. A., de Oliveira Albuquerque, R., Gomes de Deus, F. E., de Sousa Jr, R. T., de Oliveira Júnior, G. A., Garcia Villalba, L. J., & Kim, T. H. (2017). Cybersecurity and network forensics: Analysis of malicious traffic towards a honeynet with deep packet inspection. *Applied Sciences*, *7*(10), 1082. https://doi.org/10.3390/app7101082

28. Raju, R. K. (2017). Dynamic memory inference network for natural language inference. International Journal of Science and Research (IJSR), 6(2). https://www.ijsr.net/archive/v6i2/SR24926091431.pdf

29. Sardana, J. (2022). The role of notification scheduling in improving patient outcomes. *International Journal of Science and Research Archive*. Retrieved from https://ijsra.net/content/role-notification-scheduling-improving-patient

30. Schulz, G. (2016). *The green and virtual data center*. CRC Press.

31. Shantharama, P., Thyagaturu, A. S., & Reisslein, M. (2020). Hardware-accelerated platforms and infrastructures for network functions: A survey of enabling technologies and research studies. *IEEE Access*, *8*, 132021-132085. https://doi.org/10.1109/ACCESS.2020.3008250

32. Singh, V. (2022). Advanced generative models for 3D multi-object scene generation: Exploring the use of cutting-edge generative models like diffusion models to synthesize complex 3D environments. https://doi.org/10.47363/JAICC/2022(1)E224

33. Singh, V. (2023). Enhancing object detection with self-supervised learning: Improving object detection algorithms using unlabeled data through self-supervised techniques. International Journal of Advanced Engineering and Technology. https://romanpub.com/resources/Vol%205%20%2C%20No%201%20-%2023.pdf

34. Singla, A., Mudgerikar, A., Papapanagiotou, I., & Yavuz, A. A. (2015, October). Haa: Hardware-accelerated authentication for internet of things in mission critical vehicular networks. In *MILCOM 2015-2015 IEEE Military Communications Conference* (pp. 1298-1304). IEEE. https://doi.org/10.1109/MILCOM.2015.7357624

35. Sriram, H. K., ADUSUPALLI, B., & Malempati, M. (2021). Revolutionizing Risk Assessment and Financial Ecosystems with Smart Automation, Secure Digital Solutions, and Advanced Analytical Frameworks. https://dx.doi.org/10.2139/ssrn.5213196

36. Sukhadiya, J., Pandya, H., & Singh, V. (2018). Comparison of Image Captioning Methods. *INTERNATIONAL JOURNAL OF ENGINEERING DEVELOPMENT AND RESEARCH*, *6*(4), 43-48. https://rjwave.org/ijedr/papers/IJEDR1804011.pdf

37. Thompson, G. F. (2017). Time, trading and algorithms in financial sector security. *New Political Economy*, *22*(1), 1-11. https://doi.org/10.1080/13563467.2016.1183116

38. Wang, D., Bai, B., Lei, K., Zhao, W., Yang, Y., & Han, Z. (2019). Enhancing information security via physical layer approaches in heterogeneous IoT with multiple access mobile edge computing in smart city. *IEEE Access*, *7*, 54508-54521. https://doi.org/10.1109/ACCESS.2019.2913438

39. Yadav, A., Dobre, O. A., & Ansari, N. (2017). Energy and traffic aware full-duplex communications for 5G systems. *IEEE Access*, *5*, 11278-11290. https://doi.org/10.1109/ACCESS.2017.2696822

40. Zou, Y., Zhu, J., Wang, X., & Hanzo, L. (2016). A survey on wireless security: Technical challenges, recent advances, and future trends. *Proceedings of the IEEE*, *104*(9), 1727-1765. https://doi.org/10.1109/JPROC.2016.2558521