



International Journal of  
Information and  
Cybersecurity  
DLpress is a publisher of  
scholarly books and  
peer-reviewed scientific  
research. With a dedication  
to academic excellence,  
DLpress publishes books and  
research papers on a diverse  
range of topics spanning  
various disciplines, including  
but not limited to, science,  
technology, engineering,  
mathematics, social sciences,  
humanities, and arts.  
Published 08, January, 2020

# Mitigating Security Threats in Service Function Chaining: A Study on Attack Vectors and Solutions for Enhancing NFV and SDN-Based Network Architectures

Arunkumar Velayutham <sup>1</sup>

<sup>1</sup>Cloud Software Development Engineer and Technical Lead at Intel, Arizona, USA

## RESEARCH ARTICLE

### Abstract

Service Function Chaining (SFC) is a networking architecture that enables the dynamic sequencing of network functions to manage various traffic types in contemporary software-defined infrastructures. It relies on two core technologies: Network Function Virtualization (NFV) and Software-Defined Networking (SDN), both of which provide programmability, flexibility, and scalability. However, these benefits come at the cost of increased security vulnerabilities across multiple layers of the SFC stack. This paper critically examines the possible vulnerabilities in SFC, focusing on key attack vectors such as man-in-the-middle (MitM) attacks, service chain manipulation, data leakage, and denial of service (DoS). The paper discusses the expanded attack surface introduced by the programmability of SDN, the multi-tenancy of NFV, and the dynamic nature of SFC orchestration. To counteract these threats, various mitigation strategies, including cryptographic safeguards, policy enforcement, secure service orchestration, and advanced traffic monitoring, are discussed.

Keywords: attack vectors, mitigation strategies, network function virtualization, security vulnerabilities, service function chaining, software-defined networking

## 1 Introduction

Service Function Chaining (SFC) represents an advancement in the design of network architectures, enabling the dynamic composition of Virtualized Network Functions (VNFs) to satisfy specific service needs [1]. By integrating NFV and SDN, SFC abstracts service functions from physical hardware and centralizes control over traffic flows. This flexibility allows for optimized resource utilization and improved operational efficiency. However, SFC also introduces security challenges due to its expanded attack surface at the control and data plane levels.

NFV decouples network functions from dedicated hardware, allowing multiple VNFs to run on shared infrastructure, increasing the risk of inter-tenant attacks and data breaches. SDN, by design, centralizes network control, making its controllers a high-value target for attacks. The dynamic nature of SFC orchestrations further complicates security, as service chains can be altered in real-time, creating potential for malicious manipulation [2]. As SFC becomes more widely deployed, its security becomes important.

This paper addresses the primary vulnerabilities in SFC, focusing on the key attack vectors that threaten the control plane, data plane, and VNFs. Each vulnerability is analyzed in detail, followed by technical mitigation strategies aimed at minimizing risks and safeguarding SFC infrastructures.

## OPEN ACCESS Reproducible Model

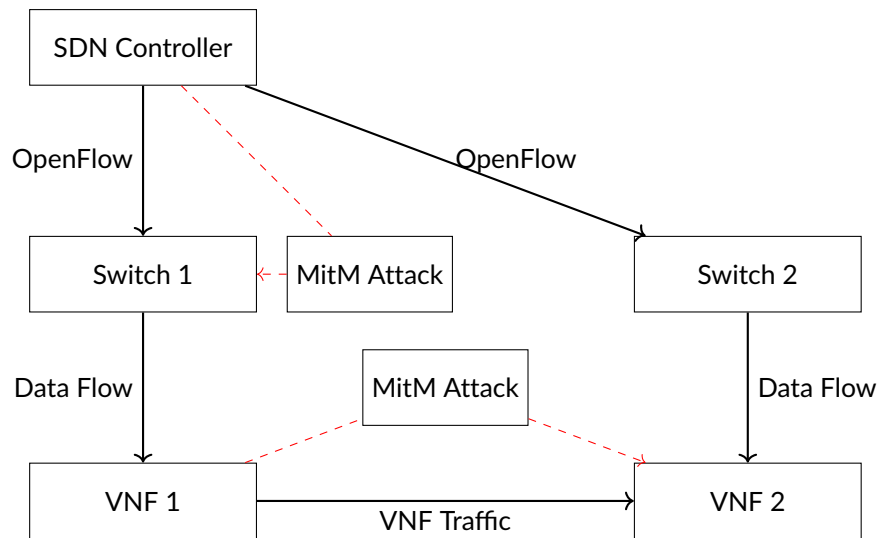
*Edited by*  
Associate Editor

*Curated by*  
The Editor-in-Chief

*Submitted* 18, October, 2019

*Accepted* 11, December, 2019

*Citation*  
Velayutham, A. (2020)  
Mitigating Security Threats in  
Service Function Chaining: A  
Detailed Study of Attack  
Vectors and Solutions for  
Enhancing NFV and  
SDN-Based Network  
Architectures



**Figure 1.** Man-in-the-Middle Attack in SFC Environment.

## 2 Security Vulnerabilities in Service Function Chaining

### 2.1 Man-in-the-Middle (MitM) Attacks

In modern Service Function Chaining (SFC) environments, the increasing reliance on Software-Defined Networking (SDN) introduces numerous security vulnerabilities in communication between the control and data planes. One of the most significant threats that can arise in this architecture is the Man-in-the-Middle (MitM) attack. These attacks exploit the separation between the control plane, where decisions are made, and the data plane, where packets are forwarded. The SDN controller, which governs traffic flow and communicates with virtual network functions (VNFs) and network devices, becomes a critical point of vulnerability. In an MitM attack, an adversary infiltrates the communication path between these entities, gaining the ability to intercept, manipulate, or even inject malicious traffic, potentially compromising the entire service chain [3].

The vulnerability to MitM attacks in an SFC context stems from the inadequately secured communication channels between the SDN controller and the VNFs, as well as between the VNFs themselves. The control-data plane communication is often transmitted via protocols like OpenFlow, which, if not properly secured, can become an attractive vector for interception. Attackers can exploit weak encryption or authentication methods in these protocols to reroute traffic, manipulate flow rules, or inject malicious traffic into the network. Specifically, control messages sent from the SDN controller to network switches, which determine the routing of packets, are vulnerable if they are not encrypted or authenticated. The modification of these flow rules can lead to significant disruptions in network behavior, allowing attackers to steer legitimate traffic through compromised nodes or introduce malicious traffic into the data stream.

One of the primary attack vectors in MitM scenarios is the lack of encryption in traffic exchanged between VNFs. In an SFC, VNFs often perform critical tasks such as firewalls, intrusion detection systems, or load balancers. If the communication between these VNFs is unencrypted, attackers can eavesdrop on sensitive information, such as security policies or user data, as it passes through the service chain. Furthermore, this lack of encryption not only exposes sensitive data to potential interception but also allows attackers to manipulate the traffic, introducing vulnerabilities into the network's operations and security posture.

Mathematically, the security of the communication between the control and data planes depends on the robustness of the cryptographic protocols in place. Protocols like Transport Layer Security (TLS) 1.3 and Perfect Forward Secrecy (PFS) play a critical role in securing these communications by ensuring that even if a long-term key is compromised, past session keys remain secure. In

elliptic curve cryptography (ECC), which is often used in PFS implementations, the security relies on the difficulty of the Elliptic Curve Discrete Logarithm Problem (ECDLP). The hardness of solving the equation

$$P = k \cdot G,$$

where  $G$  is a known point on the elliptic curve,  $k$  is the secret key, and  $P$  is the public key, underpins the strength of ECC. Due to the exponential complexity of solving ECDLP, even with significant computational power, ECC provides a high degree of security in ensuring that session keys cannot be easily compromised.

TLS 1.3 is designed to enhance the security of communications by utilizing advanced cryptographic algorithms such as RSA, Diffie-Hellman, and ECC. In the case of RSA, the security is based on the factorization problem, which is the difficulty of factoring the product of two large prime numbers. Mathematically, if

$$n = p \cdot q,$$

where  $p$  and  $q$  are prime, the challenge is to determine  $p$  and  $q$  given only  $n$ . Similarly, in Diffie-Hellman, the security relies on the difficulty of solving the discrete logarithm problem. For example, if two parties agree on a prime  $p$  and a base  $g$ , and each chooses a secret exponent  $a$  and  $b$ , they can compute a shared secret  $g^{ab} \bmod p$  without revealing their private exponents.

In the context of securing SDN communications, these cryptographic mechanisms are indispensable. By utilizing TLS 1.3 with PFS, SFC systems can ensure that even if long-term encryption keys are compromised, session keys used for specific communications are not recoverable, preserving the confidentiality of past transmissions. Mathematically, this is achieved by generating a new key for each session, which is derived from a combination of short-lived, ephemeral keys. The key exchange process, often based on Diffie-Hellman or ECC, ensures that even if an attacker gains access to the long-term private keys of the communicating parties, they cannot retroactively decrypt previous communications due to the one-time nature of the session keys.

Another crucial aspect of mitigating MitM attacks is ensuring the integrity of the flow rules, which can be verified using cryptographic hash functions, such as those employed in hash-based message authentication codes (HMACs). These cryptographic hash functions are grounded in combinatorial mathematics and exhibit desirable properties such as collision resistance, preimage resistance, and second preimage resistance. A cryptographic hash function  $H(x)$  takes an input  $x$  of arbitrary length and produces a fixed-length output  $h$ . The security of the hash function ensures that it is computationally infeasible to find two distinct inputs  $x_1$  and  $x_2$  such that

$$H(x_1) = H(x_2),$$

a property known as collision resistance. This characteristic is essential for ensuring that flow rules between the SDN controller and network switches remain untampered, as any modification to the rules would result in a different hash value, alerting the system to potential tampering.

The verification process using HMAC can be described mathematically as follows. Given a message  $m$  and a secret key  $k$ , the HMAC is computed as:

$$\text{HMAC}(k, m) = H((k \oplus \text{opad}) \parallel H((k \oplus \text{ipad}) \parallel m)),$$

where  $H$  is a cryptographic hash function (such as SHA-256),  $\oplus$  denotes the bitwise XOR operation,  $\parallel$  represents concatenation, and opad and ipad are padding constants. The security of HMAC relies on the underlying strength of the hash function  $H$  and the secrecy of the key  $k$ . In the context of SDN, HMAC ensures that any changes to the flow rules are immediately detectable, as any tampering would alter the hash value, invalidating the HMAC.

To illustrate the impact of cryptographic protections in mitigating MitM attacks, consider the following table, which outlines a comparison between encrypted and unencrypted communications in an SDN-enabled SFC environment.

Scenario	Encryption Used	Vulnerability to MitM
Unencrypted VNF Traffic	None	High
TLS 1.3 with PFS	Yes (Elliptic Curve)	Low
Encrypted Control Messages (OpenFlow)	Yes (RSA)	Low
No Encryption on Control Plane	None	High

The table demonstrates that the use of encryption with modern cryptographic protocols like TLS 1.3 and PFS, significantly reduces the vulnerability to MitM attacks. Without encryption, attackers can easily intercept and manipulate traffic, but with proper cryptographic protections, these attacks become much more challenging.

To further illustrate the impact of cryptographic protections, consider the computational complexity of attacking an elliptic curve-based system versus an RSA-based system. The following table shows a comparison of the key lengths required to achieve equivalent security in RSA and ECC, emphasizing the efficiency of elliptic curve cryptography.

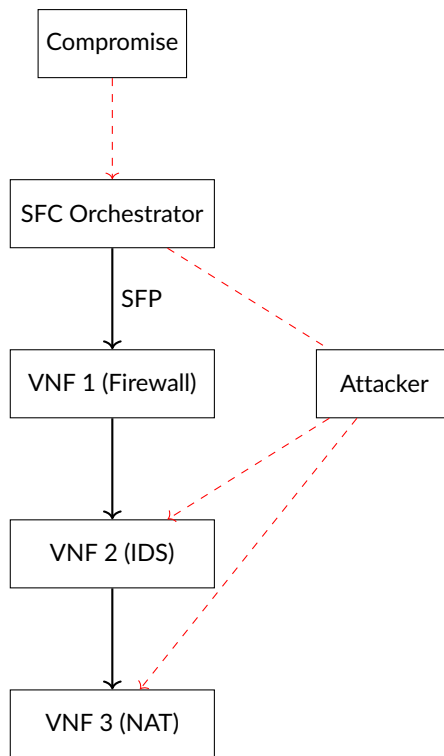
Algorithm	Key Length (bits)	Security Level
RSA	3072	128-bit security
ECC	256	128-bit security
RSA	15360	256-bit security
ECC	521	256-bit security

This table illustrates the efficiency of ECC compared to RSA. For the same level of security, ECC requires significantly shorter key lengths, which translates to faster computations and reduced resource consumption, making it a more attractive option for securing communications in resource-constrained environments such as SDN-enabled SFCs. In Figure 1, a Man-in-the-Middle (MitM) attack in a Service Function Chaining (SFC) environment is depicted. The SDN controller manages traffic through OpenFlow communication with network switches (Switch 1 and Switch 2) and subsequently with VNFs (VNF 1 and VNF 2). Normally, data flows seamlessly between the controller, switches, and VNFs. However, the figure illustrates two potential MitM attack scenarios. In one case, a malicious actor inserts themselves between VNFs (VNF 1 and VNF 2), intercepting and manipulating the traffic flowing between them. In another scenario, an attacker targets the control-data plane communication, intercepting messages between the SDN controller and the switch (Switch 1). These attacks exploit insecure communication paths, potentially leading to traffic manipulation, interception, or injection of malicious data [4].

The MitM attack compromises the integrity and confidentiality of the service chain by allowing the attacker to modify or disrupt traffic without detection. The cloud symbols in the diagram represent the locations where attackers could insert themselves, exploiting vulnerabilities in either the VNF traffic or the OpenFlow-based control communication between the SDN controller and the switches.

## 2.2 Service Chain Manipulation

Service chain manipulation presents a significant threat in Service Function Chaining (SFC) environments, where an attacker can tamper with the sequence of Virtual Network Functions (VNFs) or the routing of network traffic through the service chain. This type of attack allows malicious actors to bypass critical security controls, such as firewalls or intrusion detection systems (IDS), thereby compromising the overall security posture of the network. The dynamic and programmable nature of SFC, which offers flexibility in orchestrating network services, also provides a potential vector for exploitation when orchestration mechanisms are insecure or insufficiently protected. Attackers who succeed in manipulating service chains can effectively neutralize the security functions embedded within these chains, leading to severe vulnerabilities in the network infrastructure.



**Figure 2.** Service Chain Manipulation in SFC Environment.

One of the primary vectors for service chain manipulation is the compromise of SFC orchestration platforms. These platforms, such as OpenStack or OpenDaylight, manage the creation and modification of service function paths (SFPs) that dictate how traffic flows through different VNFs in the chain. By exploiting vulnerabilities in these platforms, an attacker can reorder VNFs or alter SFPs, allowing traffic to bypass critical VNFs responsible for security, like firewalls or IDS systems. This enables the attacker to introduce malicious traffic into the network or circumvent key security controls [4].

Another significant vector is the modification of SFC descriptors, which define the logic and structure of the service chain. If attackers gain access to these descriptors, they can alter the intended functionality by injecting malicious VNFs or removing essential security components. As these descriptors are typically stored and exchanged in the orchestration layer, securing them is vital to maintaining the integrity of the service chain. A compromised SFC descriptor can result in a corrupted or improperly ordered chain that leaves the network vulnerable.

The use of cryptographic signatures provides a robust method for ensuring the immutability of service chain descriptors. Cryptographic signatures rely on public-key cryptography, specifically on mathematical problems such as the difficulty of factoring large integers (as used in RSA) or solving discrete logarithm problems (as used in Elliptic Curve Digital Signature Algorithm (ECDSA)). In RSA, the security of the system depends on the factorization of a large composite number. Given a number  $n = p \cdot q$ , where  $p$  and  $q$  are large primes, it is computationally difficult to factor  $n$  back into  $p$  and  $q$ , providing the foundation for RSA's security.

In contrast, ECDSA relies on the elliptic curve discrete logarithm problem (ECDLP), which is defined as follows: Given a point  $P$  on an elliptic curve and an integer  $k$ , computing  $Q = kP$  is straightforward, but determining  $k$  given  $P$  and  $Q$  is computationally infeasible. This one-way function is the basis for the security of elliptic curve cryptography, including ECDSA, which is widely used for digital signatures due to its efficiency and strong security properties. The digital

signature process can be modeled as:

$$\text{Signature} = (r, s) = \left( g^k \pmod p, (H(m) + r \cdot x)k^{-1} \pmod q \right),$$

where  $g$  is the generator of a cyclic group,  $p$  and  $q$  are large primes,  $H(m)$  is the hash of the message  $m$ , and  $x$  is the private key. Verification of the signature ensures that any unauthorized modifications to the service chain descriptors are detectable, preserving the integrity of the service chain.

In addition to cryptographic signatures, role-based access control (RBAC) plays a key role in preventing unauthorized changes to service chain configurations. RBAC models can be represented using mathematical structures such as graphs or matrices, which describe the relationships between roles, users, and permissions. Let the matrix  $A$  represent the access control policy, where each element  $a_{ij}$  indicates whether user  $u_i$  has permission for action  $a_j$ . This matrix can be analyzed for consistency and security by checking for permission conflicts or violations of least-privilege principles.

Moreover, real-time integrity checks for the service chain can further ensure that the VNF sequence and traffic paths remain consistent with the intended configuration. These checks involve continuously validating the integrity and sequencing of VNFs using cryptographic techniques, such as hash-based message authentication codes (HMACs). The HMAC of a VNF sequence can be computed as:

$$\text{HMAC}(k, m) = H((k \oplus \text{opad}) \parallel H((k \oplus \text{ipad}) \parallel m)),$$

where  $H$  is a cryptographic hash function,  $k$  is the secret key, and  $m$  is the message. Any discrepancies in the VNF sequence or unexpected path changes can trigger alarms, thereby preventing manipulation of the service chain in real-time.

The table below compares the use of cryptographic signatures with and without public-key cryptography for securing service chain descriptors:

Method	Cryptographic Basis	Security Level
Unsigned Descriptors	None	High vulnerability to manipulation
Signed Descriptors (RSA)	Factorization of large integers	Strong security
Signed Descriptors (ECDSA)	Elliptic Curve Discrete Logarithm Problem	Strong security with shorter keys

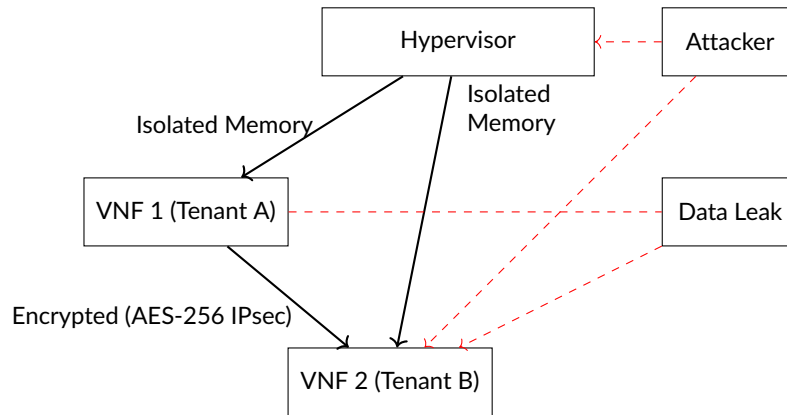
As shown, using cryptographic signatures significantly enhances the security of service chain descriptors by protecting against unauthorized modifications. The difference between RSA and ECDSA lies primarily in efficiency and key size, with ECDSA providing equivalent security with shorter keys and faster computations.

To further demonstrate the advantages of role-based access control (RBAC) in SFC orchestration platforms, consider the following access control matrix:

User	Modify Service Chain	View Logs	Delete Service Chain
Administrator	Yes	Yes	Yes
Operator	Yes	Yes	No
Auditor	No	Yes	No

This matrix ensures that only authorized roles, such as administrators, have the necessary permissions to modify service chains, while operators have restricted access, and auditors can only view logs without making changes. Implementing such RBAC policies ensures that unauthorized personnel cannot modify the service chain configurations, thus enhancing security and accountability within the SFC orchestration platforms.

### 2.3 Data Leakage



**Figure 3.** Data Leakage Vulnerability in SFC Environment.

Data leakage is a critical vulnerability in Service Function Chaining (SFC) environments in multi-tenant infrastructures where Virtual Network Functions (VNFs) often share the same physical resources. The absence of strong isolation mechanisms can result in sensitive data from one tenant leaking into another tenant's environment, leading to significant confidentiality breaches. Additionally, inadequate encryption of inter-VNF communication presents another major risk, as it allows attackers to intercept and access unencrypted data in transit. These vulnerabilities, if left unchecked, can lead to severe data exposure and unauthorized access in SFC environments [4].

One of the primary attack vectors in this context is weak isolation between VNFs, which attackers can exploit by targeting vulnerabilities in the hypervisor or containerization layers. Hypervisors, which manage the execution of virtual machines, and container runtimes, which handle lightweight virtual environments, are responsible for partitioning and isolating resources like memory and storage. However, if these layers are compromised, attackers can access the memory or storage of other VNFs, leading to data theft, privilege escalation, or even complete control over other tenants' VNFs. This is problematic in shared infrastructures where multiple tenants rely on the same physical hardware [5].

Another critical vector is the lack of encryption in inter-VNF communication. When traffic between VNFs is transmitted without proper encryption, it is susceptible to interception by attackers. This interception can lead to exposure of sensitive information, such as user data, security policies, or cryptographic keys. To mitigate this risk, encryption protocols such as IPsec are used to secure traffic between VNFs. IPsec operates at the network layer and ensures that data transmitted between VNFs is both encrypted and authenticated, preventing unauthorized access to sensitive information.

The encryption of inter-VNF communication typically involves the use of cryptographic algorithms like Advanced Encryption Standard (AES) with a key size of 256 bits (AES-256), combined with Galois/Counter Mode (GCM) for both encryption and integrity. AES-256 is a block cipher that operates on blocks of data, and its security relies on the difficulty of key recovery attacks due to its large key size. The encryption process can be modeled as follows: Given a plaintext  $P$ , AES-256 operates on blocks of 128 bits, transforming the plaintext into ciphertext  $C$  through a series of substitutions and permutations governed by the secret key  $K$ . Mathematically, the encryption of a block  $P_i$  is represented as:

$$C_i = \text{AES}_K(P_i),$$

where  $C_i$  is the ciphertext block, and  $K$  is the 256-bit key.

In Galois/Counter Mode (GCM), AES operates in conjunction with finite field mathematics to ensure both confidentiality and integrity. GCM uses a counter for encryption, where the counter is incremented for each block of plaintext, while a Galois field  $GF(2^{128})$  is used for integrity protection. The ciphertext  $C$  is accompanied by an authentication tag  $T$ , ensuring that any unauthorized

modification to the data can be detected. The integrity check relies on the mathematical structure of finite fields. In GCM, the tag  $T$  is computed using the following equation:

$$T = H \times (C_1 \parallel C_2 \parallel \dots \parallel C_n),$$

where  $H$  is derived from the encryption key, and  $C_1, C_2, \dots, C_n$  are the ciphertext blocks. The operation  $\times$  denotes multiplication in the finite field  $GF(2^{128})$ , and  $\parallel$  represents concatenation. This algebraic structure ensures that even if an attacker intercepts the traffic, any modification to the ciphertext or the tag will result in an invalid tag, allowing the receiving VNF to detect tampering.

To prevent data leakage due to weak VNF isolation, it is essential to deploy hypervisors with strong memory management features, such as Intel VT-x or AMD-V. These technologies provide hardware-assisted virtualization, enabling secure partitioning of memory and other resources between different VNFs. A secure hypervisor ensures that one VNF cannot access the memory space or storage of another, even if they share the same physical hardware. Formal methods can be used to analyze and verify the security of these isolation mechanisms. One such approach is to model resource isolation as a partitioning problem, where the set of physical resources  $R = \{r_1, r_2, \dots, r_n\}$  is partitioned into non-overlapping subsets  $R_1, R_2, \dots, R_k$ , where each subset  $R_j$  is assigned to a different VNF. Mathematically, the partitioning must satisfy the condition:

$$R_i \cap R_j = \emptyset \quad \text{for all } i \neq j,$$

ensuring that no two VNFs share the same physical resources. Secure hypervisors enforce this condition by managing the allocation of memory, storage, and CPU resources in a way that guarantees isolation between tenants.

In addition to encryption and isolation, regular security auditing and logging are crucial for detecting potential data leakage incidents. Implementing strict auditing mechanisms ensures that access to sensitive data is logged and analyzed in real-time. Logs should capture all access events, including memory reads and writes, storage access, and inter-VNF communication. By continuously monitoring these logs, administrators can detect abnormal behavior, such as unauthorized access or unusual traffic patterns, which may indicate an ongoing attack. Regular audits can help identify weaknesses in the system, enabling timely responses to mitigate potential data leaks.

The table below compares different encryption modes used for securing inter-VNF communication:

Encryption Mode	Key Size (bits)	Integrity Check	Security Level
AES-CBC	256	None	High encryption, no integrity
AES-GCM	256	Galois field $GF(2^{128})$	High encryption, high integrity
Unencrypted	N/A	None	No security

As shown in the table, AES-GCM provides both strong encryption and data integrity, making it the preferred choice for securing inter-VNF communication in SFC environments. The combination of encryption and integrity checks ensures that even if attackers manage to intercept the traffic, they cannot decrypt or modify the data without detection [6].

To further illustrate the benefits of secure hypervisors in preventing data leakage, consider the following comparison of VNF isolation methods:

Isolation Method	Isolation Mechanism	Vulnerability to Data Leakage
Containerization (Weak)	Namespace isolation	High
Hypervisor (Secure)	Hardware-assisted memory isolation	Low
Bare-Metal (No Virtualization)	Dedicated resources	None

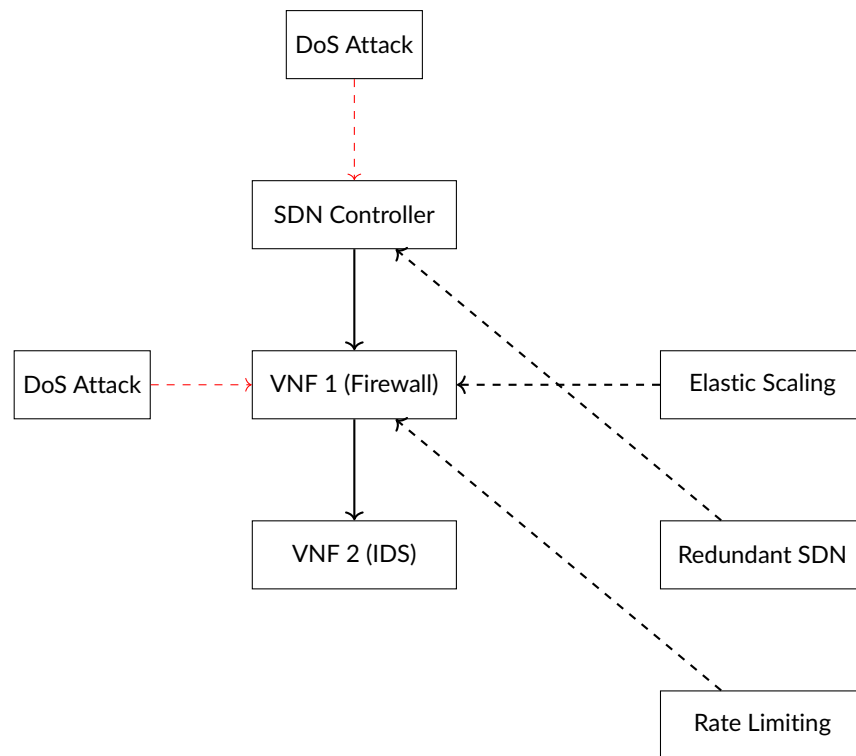


This table highlights that secure hypervisors, such as those using Intel VT-x or AMD-V, significantly reduce the risk of data leakage compared to weak containerization methods, where namespace isolation can be insufficient to prevent cross-tenant access.

In Figure 3, the data leakage vulnerability within a Service Function Chaining (SFC) environment is depicted, focusing on a multi-tenant infrastructure where VNFs from different tenants share the same physical resources managed by a hypervisor. VNFs (VNF 1 for Tenant A and VNF 2 for Tenant B) are isolated at the memory level through the hypervisor, but vulnerabilities in the isolation mechanism could be exploited by attackers. The cloud labeled "Data Leak" represents the potential for data leakage between VNFs if the isolation is weak, allowing sensitive information from Tenant A to be accessed by Tenant B or vice versa. This highlights the risk of cross-tenant data breaches when proper isolation techniques are not enforced [7].

The diagram also illustrates the importance of encrypting traffic between VNFs to prevent interception by attackers. The secure communication between VNF 1 and VNF 2 is indicated by the encrypted connection (using AES-256 with IPsec). However, if the encryption is not implemented, as represented by the dashed red lines from the attacker, malicious actors can compromise the hypervisor or intercept the unencrypted traffic, gaining unauthorized access to sensitive data. This emphasizes the need for both strong VNF isolation mechanisms and secure communication protocols to mitigate the risk of data leakage in SFC environments [8].

## 2.4 Denial of Service (DoS) Attacks



**Figure 4.** DoS Attack on VNF and SDN Controller in SFC Environment.

Denial of Service (DoS) attacks present a significant vulnerability in Service Function Chaining (SFC) environments, where attackers can target specific Virtual Network Functions (VNFs) or even the Software-Defined Networking (SDN) controller itself. The nature of these attacks involves overwhelming the target with excessive traffic, depleting its resources and leading to service degradation or complete failure. In extreme cases, the attacker may exhaust the resources of the SDN controller, leading to a network-wide service disruption. Understanding and mitigating these attacks require mathematical tools from queueing theory, traffic modeling, and control theory, which allow for the prediction and management of network load during high-traffic events.

A primary attack vector for DoS in SFC environments is the exhaustion of resources in resource-heavy VNFs, such as firewalls or intrusion detection systems (IDS). Attackers can flood these VNFs with a massive number of packets, consuming critical resources such as CPU, memory, or bandwidth. As the load increases beyond the capacity of the VNF, it fails to process legitimate traffic, leading to service degradation or complete failure. Mathematically, the load on a VNF can be modeled using queueing theory, where packets are queued for processing based on the VNF's available resources. Let  $\lambda$  be the arrival rate of packets and  $\mu$  the service rate of the VNF. The system can be modeled as an  $M/M/1$  queue, where the utilization  $\rho$  is defined as:

$$\rho = \frac{\lambda}{\mu}.$$

When  $\rho$  approaches 1, the system becomes saturated, leading to long delays and potential packet loss. In the case of a DoS attack,  $\lambda$  increases dramatically, overwhelming the VNF and causing it to fail.

Another common vector for DoS attacks is targeting the SDN controller. Attackers can generate massive amounts of flow requests, overwhelming the controller's processing capacity. This leads to delays in propagating flow rules to the network devices or, in the worst-case scenario, causes a complete failure of the control plane. The performance of an SDN controller under load can also be analyzed using queueing theory, where the controller must process incoming flow requests at a rate  $\mu$ . If the incoming rate of flow requests  $\lambda$  exceeds  $\mu$ , the controller's queue builds up, and delays in rule propagation increase. In such cases, elastic scaling and distributed control mechanisms are required to mitigate the impact of these attacks.

To defend against VNF resource exhaustion, elastic scaling is one of the most effective mitigation strategies. Autoscaling mechanisms dynamically adjust the resources allocated to VNFs based on the traffic load. These mechanisms rely on predictive algorithms rooted in control theory and statistical models, which analyze incoming traffic patterns and predict when additional resources are needed. Mathematically, autoscaling can be modeled as a control system, where the traffic load  $L(t)$  is the system's input, and the allocated resources  $R(t)$  are the output. The goal is to maintain a stable system where the available resources match or exceed the incoming load, preventing service degradation. A simple feedback control model can be expressed as:

$$R(t) = K_p \cdot (L(t) - L_{\text{target}}),$$

where  $K_p$  is the proportional gain, and  $L_{\text{target}}$  is the desired target load. The system dynamically adjusts  $R(t)$  based on deviations from the target load  $L_{\text{target}}$ , ensuring that the VNF does not become overwhelmed.

In addition to elastic scaling, rate-limiting and flow control mechanisms play a critical role in preventing DoS attacks. Rate-limiting is applied at the network ingress to restrict the amount of traffic that can be forwarded to specific VNFs. This technique prevents malicious traffic from overwhelming a VNF by capping the rate at which packets are processed. Rate-limiting can be mathematically modeled using token bucket algorithms, where tokens accumulate at a fixed rate  $r$ , and each packet requires a token to be processed. If the token bucket is empty, excess packets are either delayed or dropped. The system can be represented as:

$$\text{Rate Limit} = \min(B + r \cdot t, \text{max rate}),$$

where  $B$  is the burst size,  $r$  is the token generation rate, and  $t$  is time. This ensures that traffic flows at a manageable rate, preventing overwhelming spikes that could deplete VNF resources.

To ensure the resilience of the SDN controller against DoS attacks, redundancy and distributed control mechanisms are critical. SDN controllers can be deployed in a distributed architecture, where multiple controllers work together to manage the network. In this configuration, consensus protocols such as Paxos or Raft are used to ensure that the controllers agree on network state and can tolerate failures. These protocols rely on distributed systems theory and are designed to be fault-tolerant and efficient even in the presence of network partitions or controller failures.

Paxos, for instance, ensures consensus among distributed controllers by proposing values (e.g., flow rules) and achieving agreement despite failures. Mathematically, Paxos can be described as a series of phases involving proposers, acceptors, and learners. In phase 1, a proposer sends a prepare request with a proposal number  $n$ , and the acceptors respond with their highest accepted proposal number  $n'$ . In phase 2, the proposer sends a proposal with value  $v$ , and the acceptors decide on this value if no higher proposal number  $n'$  has been accepted.

Another popular consensus algorithm, Raft, simplifies the consensus process by electing a leader that coordinates all changes to the system. Raft divides time into terms, and during each term, an election is held to select a leader. The leader handles all client requests and replicates the changes to other controllers (followers). Raft guarantees consistency and fault tolerance by ensuring that any committed log entry is durable even in the event of failures. These consensus protocols ensure that the SDN control plane remains operational and consistent, even under heavy load or in the event of failures.

The table below provides a comparison of the mitigation strategies for DoS attacks in SFC environments:

Mitigation Strategy	Technique	Effectiveness against DoS
Elastic Scaling	Dynamic resource allocation based on traffic load	High
Rate Limiting	Restricts traffic to prevent overload	Moderate to High
Distributed SDN Controllers	Consensus protocols like Paxos or Raft for fault tolerance	High

**Table 1.** DoS Mitigation Strategies

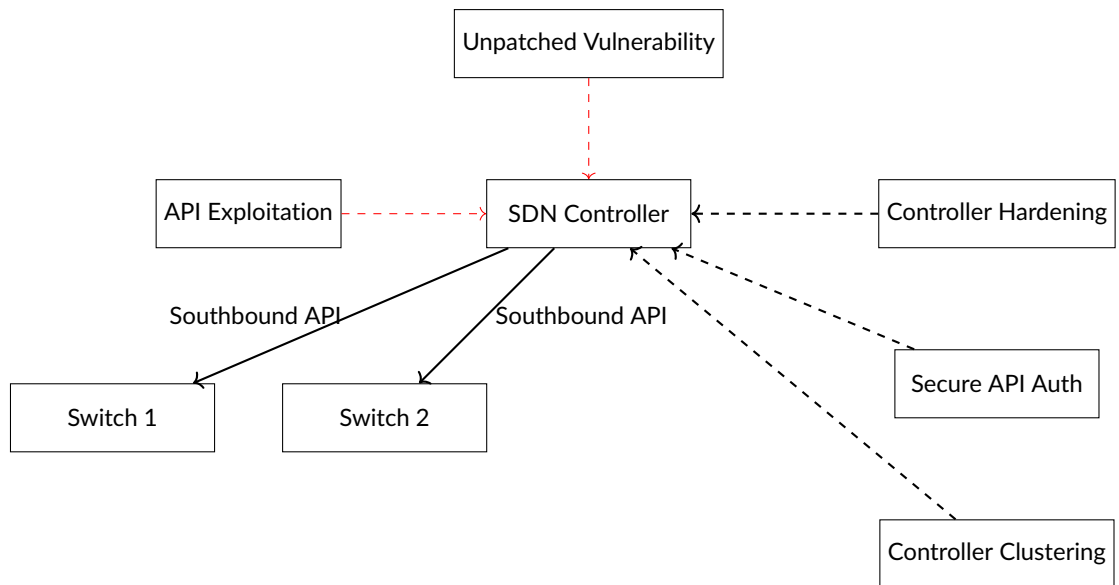
As seen in the table, elastic scaling and distributed control are highly effective in mitigating DoS attacks by dynamically adapting to traffic patterns and ensuring fault tolerance in the control plane.

Figure 4 illustrates a Denial of Service (DoS) attack in an SFC environment, targeting both a VNF and the SDN controller. The normal traffic flow is depicted between the SDN controller and two VNFs—VNF 1 (Firewall) and VNF 2 (IDS). In this scenario, a DoS attack can target VNF 1 by overwhelming it with excessive traffic, as shown by the dashed red line from the cloud labeled "DoS Attack" to VNF 1. This resource exhaustion causes the VNF to fail, disrupting its functionality and potentially causing service degradation across the network. Similarly, the SDN controller itself can be targeted by a DoS attack, as shown by another dashed red line, flooding it with flow requests and causing a breakdown in the control plane, affecting network-wide operations [9].

Mitigation strategies are also depicted to counter these attacks. Elastic scaling mechanisms are applied to VNF 1, allowing dynamic allocation of resources to handle traffic spikes. The implementation of redundant SDN controllers ensures control plane continuity even if the primary controller is overwhelmed, maintaining operational availability. Finally, rate limiting is applied at VNF 1 to control traffic flows and prevent the VNF from being overwhelmed by excessive traffic, thus ensuring that the service chain remains functional even during attack events. These combined strategies provide resilience against DoS attacks in SFC environments.

## 2.5 Compromise of SDN Controllers

The SDN controller is a critical component in a Service Function Chaining (SFC) environment, acting as the central point of control responsible for managing traffic flows, deploying service chains, and communicating with network devices. A compromised SDN controller allows attackers to gain full control over the network, which can have catastrophic consequences. Attackers with access to the SDN controller can reroute traffic, inject malicious flows, disable critical security functions, and severely disrupt network operations. Thus, securing and hardening the SDN controller is paramount to maintaining the integrity and availability of the network.



**Figure 5.** Compromise of SDN Controller in SFC Environment.

One common attack vector involves exploiting vulnerabilities in the APIs exposed by the SDN controller. The SDN controller interacts with various network entities through its northbound and southbound APIs, making these interfaces vulnerable. Attackers can exploit weak authentication or authorization mechanisms in these APIs to gain unauthorized access or execute privileged commands. In particular, the northbound API, which provides programmability to applications that interact with the SDN controller, can be targeted for unauthorized control of network resources. On the other hand, the southbound API, which communicates with network devices, is vulnerable to attacks that manipulate flow rules or interfere with the control-data plane communication.

Another significant attack vector is unpatched SDN controller software. Popular SDN controller platforms, such as OpenDaylight and ONOS, may contain known vulnerabilities that attackers can exploit. Failing to apply regular security patches leaves these systems open to attacks, potentially leading to a full control plane compromise. Once an attacker gains control over the SDN controller, they can manipulate the entire network by injecting malicious flows or modifying service function chains to bypass security controls.

To mitigate these vulnerabilities, several strategies are employed. One of the most effective strategies is hardening the SDN controller itself. Controller hardening involves minimizing the attack surface by disabling unnecessary services, restricting API access, and applying regular security patches to eliminate known vulnerabilities. Additionally, firewall rules can be applied to control access to the northbound and southbound APIs, ensuring that only authorized devices and users can interact with the controller. These measures significantly reduce the potential for exploitation by limiting the ways in which an attacker can interact with the SDN controller.

A critical component of securing the APIs is the use of cryptographic authentication mechanisms. OAuth2, a widely used token-based authentication protocol, is an example of such a mechanism. OAuth2 relies on cryptographic primitives rooted in number theory and complexity theory to authenticate users and devices. In OAuth2, the client receives an access token after authenticating, which it uses to make authorized API requests. The security of OAuth2 depends on the difficulty of forging valid tokens, which is underpinned by the cryptographic strength of the algorithms used. For example, the token generation process can involve cryptographic hashing (e.g., SHA-256), which is computationally infeasible to reverse due to the preimage resistance property of hash functions:

$$H(x) = h, \quad \text{where it is infeasible to find } x' \text{ such that } H(x') = h.$$

In addition to OAuth2, role-based access control (RBAC) should be implemented to limit access

to sensitive API functions. In RBAC, access permissions are assigned based on the roles of users, with the access matrix  $A$  representing the relationship between users, roles, and resources. This access matrix can be modeled mathematically as:

$$A = \{a_{ij}\}, \quad \text{where } a_{ij} = 1 \text{ if user } i \text{ has permission for resource } j, \text{ and } 0 \text{ otherwise.}$$

This formal representation ensures that only users with the necessary roles can execute privileged commands on the SDN controller, thus reducing the risk of unauthorized actions.

To further enhance the resilience of the SDN controller, controller clustering and isolation are employed. Controller clustering distributes the control plane load across multiple controllers, thus preventing any single point of failure. In this architecture, multiple SDN controllers operate in unison to manage the network, and consensus algorithms, such as Paxos or Raft, are used to ensure that all controllers agree on the network state. These consensus protocols are based on distributed systems theory and provide fault tolerance by allowing the system to continue operating even if some controllers fail.

Mathematically, consensus algorithms like Paxos ensure that all non-faulty controllers agree on a single value, even in the presence of faults. Paxos operates through a series of phases involving proposers, acceptors, and learners. In Phase 1, a proposer sends a prepare request with a proposal number  $n$  to the acceptors, and in Phase 2, the proposer sends the proposal  $v$  to the acceptors, which decide whether to accept or reject based on the highest proposal number received. The mathematical guarantees of Paxos ensure that if a majority of acceptors agree on a proposal, the system reaches consensus:

If a value  $v$  has been accepted by a majority of acceptors, no other value can be chosen.

Raft, another consensus algorithm, simplifies this process by electing a leader responsible for managing changes to the network. Raft divides time into terms, and during each term, an election is held to select a leader. The leader handles all client requests and replicates the changes to the follower controllers. Raft guarantees consistency by ensuring that any committed log entry persists, even if the leader crashes, thereby maintaining the fault tolerance of the system.

Controller isolation adds an additional layer of security by separating the SDN controllers from the rest of the network. This isolation is achieved by deploying the controllers in secure, isolated environments and using secure communication channels for internal communication between the controllers. Transport Layer Security (TLS), a cryptographic protocol that ensures confidentiality, integrity, and authentication, is commonly used for this purpose. TLS uses cryptographic algorithms such as RSA or elliptic curve cryptography (ECC) for key exchange, where the security of the system depends on the mathematical hardness of problems such as the factorization of large integers (RSA) or solving the elliptic curve discrete logarithm problem (ECDLP) in ECC.

The effectiveness of TLS in securing communication between SDN controllers can be expressed mathematically through the use of public and private keys. For instance, in RSA, the public key  $(e, n)$  is used to encrypt messages, while the private key  $d$  is used for decryption. Given a message  $m$ , the encryption is:

$$c = m^e \pmod n,$$

and decryption is:

$$m = c^d \pmod n.$$

The security of RSA relies on the computational difficulty of factoring the modulus  $n$  into its prime factors, which underpins the confidentiality and security of the TLS protocol.

The table below compares key mitigation strategies for securing SDN controllers:

Mitigation Strategy	Technique	Effectiveness
Controller Hardening	Regular patches, API restriction	High
OAuth2 with RBAC	Token-based authentication, role-based access	High
Controller Clustering	Consensus protocols (e.g., Paxos, Raft)	High
TLS for Communication	Secure cryptographic communication	High

Figure 5 depicts the compromise of the SDN controller in an SFC environment. The SDN controller serves as the central point of control, managing the communication with network switches through the southbound API, represented by the connections between the SDN controller and Switch 1 and Switch 2. This figure highlights two key attack vectors. The first is API exploitation, where an attacker leverages vulnerabilities in the SDN controller's API to gain unauthorized access, depicted by the dashed red line labeled "API Exploitation." The second attack vector involves unpatched vulnerabilities in the SDN controller software, allowing attackers to take advantage of known security flaws, as shown by the cloud labeled "Unpatched Vulnerability." Both attack vectors can result in full control of the SDN controller, leading to unauthorized traffic manipulation, injection of malicious flows, or disabling critical network functions.

To mitigate these threats, three key strategies are shown in the diagram. Controller hardening involves regularly applying security patches and reducing the attack surface, depicted by the dashed line connecting "Controller Hardening" to the SDN controller. Secure API authentication, such as OAuth2, is depicted as a second mitigation, where robust authentication mechanisms ensure that only authorized entities can interact with the controller. Finally, controller clustering is represented as a method of distributing control plane responsibilities, ensuring operational resilience even in the event of a single controller being compromised. These strategies collectively strengthen the security posture of the SDN controller, reducing its susceptibility to exploitation and ensuring network stability.

### 3 Conclusion

The objective of this research is to systematically examine and evaluate the security vulnerabilities and associated threats within Service Function Chaining (SFC) architectures. SFC, which integrates Network Function Virtualization (NFV) and Software-Defined Networking (SDN), introduces new attack surfaces due to its dynamic and programmable nature. This study aims to explore key attack vectors, such as man-in-the-middle (MitM) attacks, service chain manipulation, data leakage, denial of service (DoS), and SDN controller compromise. By identifying these vulnerabilities and assessing their potential impact, the research seeks to inform the development of technical solutions and mitigation strategies that can improve the security posture of SFC-based systems. The findings of this study will contribute to the ongoing efforts to enhance the robustness and reliability of modern, virtualized network infrastructures.

Service Function Chaining is increasingly being adopted in modern networking environments due to the flexibility and efficiency it offers. By enabling the chaining of VNFs such as firewalls, load balancers, and intrusion detection systems (IDS), SFC optimizes resource allocation and simplifies network management. However, the introduction of programmable and virtualized network functions presents new security challenges that must be carefully addressed. While SFC offers clear operational benefits, its architecture can introduce vulnerabilities at both the control plane (managed by SDN) and the data plane (handled by VNFs).

The significance of this research lies in its contribution to the understanding of these specific security challenges. As networks shift towards increased reliance on software-defined and virtualized components, a nuanced understanding of the security risks posed by SFC becomes necessary. This study provides an analysis of vulnerabilities that can lead to data breaches, service disruptions, and unauthorized manipulation of network flows. Such incidents may have material consequences for service providers, including regulatory penalties, financial loss, and damage to reputation. By providing insights into these vulnerabilities and exploring possible mitigation

strategies, this research informs the design of more secure SFC implementations and contributes to ongoing security standardization efforts in the networking field.

One of the primary limitations of this research stems from the complexity of real-world SFC deployments. While the theoretical vulnerabilities associated with Service Function Chaining can be well-documented, the specific configurations and use cases in different network environments add significant variability. In practice, the architecture, placement of VNFs, types of services, and traffic patterns vary considerably between deployments. This diversity in implementation can affect both the identification of vulnerabilities and the effectiveness of mitigation strategies [10].

For instance, service providers may implement custom configurations of NFV and SDN, employ proprietary orchestration platforms, or use a hybrid mix of physical and virtual network functions, which complicates the generalizability of the findings. The research focuses on common SFC implementations, but it may not fully capture the nuances of every unique deployment scenario. Therefore, the applicability of some of the proposed mitigation strategies may be limited in scope, requiring further adaptation or customization in certain contexts.

Additionally, the dynamic nature of SFC, with frequent changes to service chains and network function placements, introduces additional challenges. Security vulnerabilities may emerge or evolve as the network adapts to different traffic conditions and operational requirements. This variability makes it difficult to provide a one-size-fits-all security solution, and ongoing monitoring and adaptation will be necessary in real-world scenarios.

Another limitation concerns the assumptions made regarding the threat models and capabilities of potential attackers. While this research identifies several prominent attack vectors in SFC, such as MitM attacks, service chain manipulation, and denial of service, the scope of the research is constrained by the need to focus on specific, well-known attack methods. The study assumes a general adversarial model where attackers possess certain capabilities, such as network access, sufficient knowledge of SFC mechanisms, or control over compromised VNFs or SDN components [10].

In reality, attackers may possess varying levels of sophistication, and emerging threats may exploit vulnerabilities that have not been considered in this research. For example, attackers with advanced capabilities may use zero-day exploits or novel attack techniques that circumvent the proposed mitigation strategies. Additionally, insider threats, which may involve actors with privileged access to orchestration platforms or network configurations, introduce a distinct class of risks that require different security approaches.

Furthermore, the study does not fully account for multi-stage or combined attack scenarios, where multiple vulnerabilities are exploited in sequence or in combination. In complex networks, attackers may use a blend of techniques, such as initiating a denial of service attack to distract from more covert attempts to manipulate service chains or exfiltrate data. The research highlights specific vulnerabilities, but future work would benefit from a more holistic approach to threat modeling, which accounts for a wider range of potential attack strategies and adversarial capabilities.

The research primarily focuses on the security aspects of SFC and its associated vulnerabilities, but it does not comprehensively address the impact of proposed mitigation strategies on the overall performance and interoperability of the system. SFC environments are complex, often involving multiple vendors, platforms, and protocols. The introduction of security mechanisms, such as encryption, authentication, and traffic monitoring, may have an impact on the performance, scalability, and interoperability of the network [6].

For example, while the implementation of Transport Layer Security (TLS) between VNFs can prevent data interception, it may also introduce latency and overhead that can degrade network performance in high-throughput environments. Similarly, rate limiting and flow control mechanisms to prevent denial of service attacks may reduce the risk of VNF exhaustion but could negatively impact legitimate traffic, leading to performance bottlenecks.

The interoperability of SFC components is also a key concern when considering the use of cryptographic techniques or secure communication protocols. In heterogeneous environments where VNFs and SDN controllers from different vendors are deployed, ensuring that security mechanisms are compatible and interoperable is a non-trivial challenge. The research acknowledges the importance of security but does not go deeply into the trade-offs between security and performance, nor does it explore the potential difficulties in ensuring seamless integration of security features across different platforms and devices.

## References

- [1] Gittler F, Toumani F. Service-oriented computing. ICSOC/ServiceWave. 2009.
- [2] Bari MF, Chowdhury SR, Ahmed R, Boutaba R. On orchestrating virtual network functions. In: 2015 11th international conference on network and service management (CNSM). IEEE; 2015. p. 50-6.
- [3] Han B, Gopalakrishnan V, Ji L, Lee S. Network function virtualization: Challenges and opportunities for innovations. IEEE communications magazine. 2015;53(2):90-7.
- [4] Hawilo H, Jammal M, Shami A. Network function virtualization-aware orchestrator for service function chaining placement in the cloud. IEEE Journal on Selected Areas in Communications. 2019;37(3):643-55.
- [5] Zhang Q, Xiao Y, Liu F, Lui JC, Guo J, Wang T. Joint optimization of chain placement and request scheduling for network function virtualization. In: 2017 IEEE 37th international conference on distributed computing systems (ICDCS). IEEE; 2017. p. 731-41.
- [6] Mijumbi R, Serrat J, Gorricho JL, Bouten N, De Turck F, Boutaba R. Network function virtualization: State-of-the-art and research challenges. IEEE Communications surveys & tutorials. 2015;18(1):236-62.
- [7] Sherman AT, McGrew DA. Key establishment in large dynamic groups using one-way function trees. IEEE transactions on Software Engineering. 2003;29(5):444-58.
- [8] Kiriansky V, Bruening D, Amarasinghe S. Secure execution via program shepherding. In: 11th USENIX Security Symposium (USENIX Security 02); 2002. .
- [9] Yi B, Wang X, Li K, Huang M, et al. A comprehensive survey of network function virtualization. Computer Networks. 2018;133:212-62.
- [10] Luizelli MC, Bays LR, Buriol LS, Barcellos MP, Gasparly LP. Piecing together the NFV provisioning puzzle: Efficient placement and chaining of virtual network functions. In: 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM). IEEE; 2015. p. 98-106.