Shielding Networks

Pioneering Defense Strategies Against Distributed Denial of Service (DDoS) Attacks

A technical paper prepared for presentation at SCTE TechExpo24

Carl Klatsky
Senior Principal Engineer
Comcast
carl_klatsky@comcast.com

Aju Kalarickal
Principal Engineer
Comcast
aju_francis@comcast.com

Ananda Mahalingham
Senior Engineer
Comcast
ananda_mahalingam@comcast.com

Jimit Salvi
Engineer III
Comcast
Jimit_salvi@comcast.com

Eswaramoorthy Subramaniam
Principal Engineer
Comcast
eswaramoorthy_subramaniam@comcast.com

Aditya Vallabhajosyula
Senior Engineer
Comcast
aditya_vallabhajosyula@comcast.com

Table of Contents

List of Figures

Title **Page Number**

1. Introduction

During the SCTE (Society of Cable Telecommunications Engineers) 2023 Cable TechExpo, the authors introduce a new network element called the Virtual Services Gateway (VSG) [1]. The VSG can be used to address several challenges faced by network operators, including collecting customer experience metrics and reducing packet attacks. The VSG integrates with the Cable Modem Termination System (CMTS) or Virtual Cable Modem Termination System (VCMTS) providing new services to both end customers and the network operator. The VSG is in the logical path of the customer data packets enabling byte usage and quality analysis use-cases. These initial use-cases are implemented using copies of the customer data packets. In this paper the authors propose a new use-case, providing security protection for the customer and the operator infrastructure by dropping or rate limiting malicious attack packets while transparently forwarding non-attack packets. The VSG can implement this use-case by virtue of its unique position in the data packet path.

2. System Overview

This section provides an overview of the VSG Ecosystem and how it fits into the path of providing Internet service to the customer.

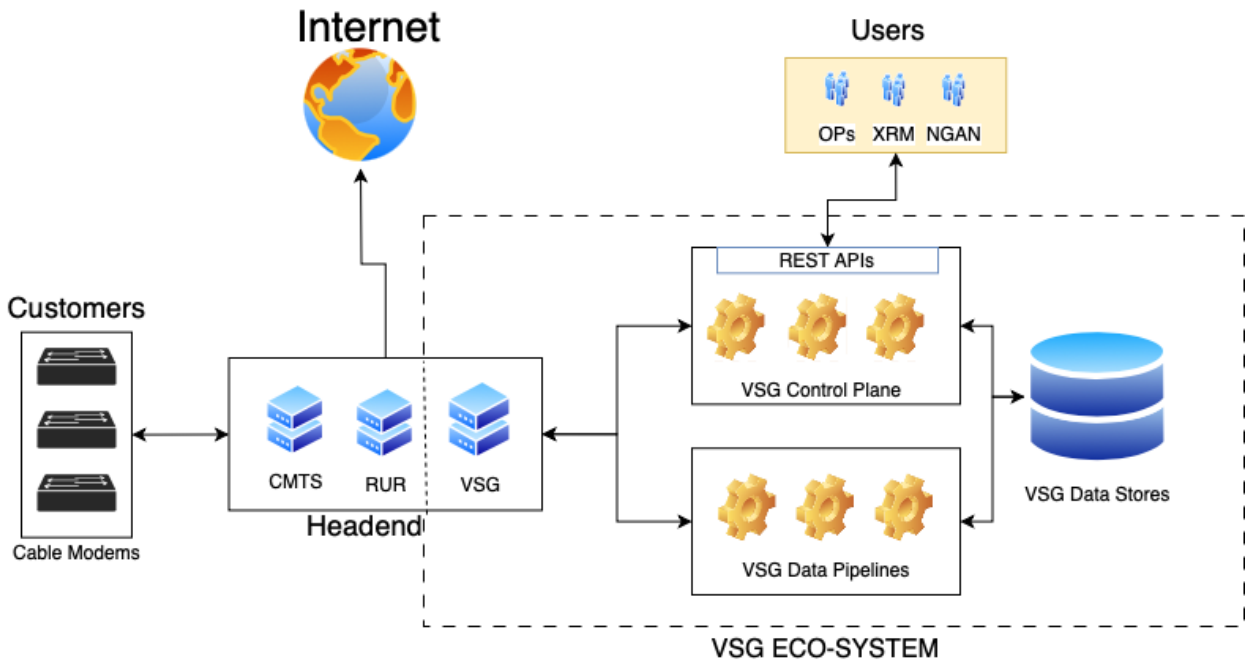

Figure 1 – System Overview

2.1. Internet Services

Internet services typically consist of several key components that collectively provide Internet connectivity, Voice over Internet Protocol (VoIP), Streaming Services, and related features to customers. The main elements of Internet Services are:

1. CMTS and / or VCMTS
2. Cable Modems / Routers
3. Set Top Boxes
4. VoIP Devices

The cable modem establishes a connection to the Internet through the CMTS, interconnecting the router that provides a local network connection for the end devices. The CMTS acts as a central hub for all cable modem connections, managing traffic and ensuring the smooth operation of the network and services.

The VSG provides a new platform where next-generation network services are deployed close to the customer edge. Automation and flexible deployment best practices are used to simplify the creation of new services and augment the capabilities of existing network services. Current state-of-the-art Software-Defined Networking (SDN) and Network-Function-Virtualization (NFV) technologies are used to enable agile implementation and orchestrated management and monitoring.

The VSG is a software application capable of running on general-purpose compute nodes, either virtual or bare metal. It can share platforms with the vCMTS or other virtual elements deployed at the network edge. It is a low-latency flow-aware platform that can forward traffic at a line rate, ensuring the introduction of the VSG in the network path does not cause a performance bottleneck. The VSG is connected between the CMTS/VCMTS and Residential Edge Router (RUR) in Headend or Hub sites.

2.2. Virtual Services Gateway – Network Element

The VSG is physically wired to the redundant routers with one 100 gigabits per second (Gbps) link to each router. Each link supports both upstream and downstream traffic. For the VSG to properly analyze the customer traffic, it is critical that the VSG application has some parameters to indicate upstream traffic from downstream traffic; this is achieved by Virtual Local Area Network (VLAN) ID or “VLAN tags” to distinguish the directionality of the traffic as it traverses the VSG. Traffic to and from the CMTS/VCMTS is marked with a unique VLAN tag. Using the router’s patch panel function, the VLAN tagged packets are re-directed through the router between the CMTS/VCMTS and the VSG. In the VSG system, a SmartNIC (Network Interface Card) is needed to perform the VLAN translation and packet copy functions, as seen in Figure 2 - VSG VLAN Interconnection and Figure 3- VSG NIC Internal View. At the VSG, the VSG NIC makes a copy of the packet to send to the VSG application and changes the VLAN tag of the original packet for transmission back to the RUR and onwards upstream to the Aggregation Routers (AR). From the perspective of the router, it is only aware of the VLAN tag value after it has been modified by the VSG. The router has a layer three sub-interface corresponding to this VLAN tag value and proceeds to route the traffic onto the network as per standard layer three routing functions. The same process occurs in reverse for downstream traffic.

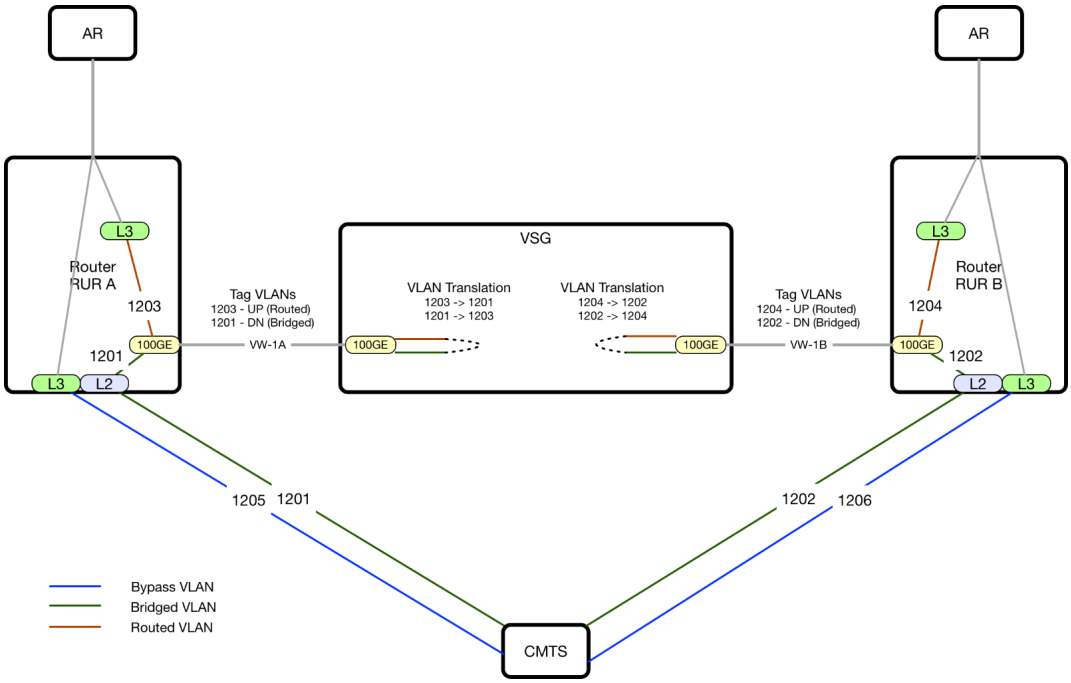

Figure 2 – VSG VLAN Interconnection

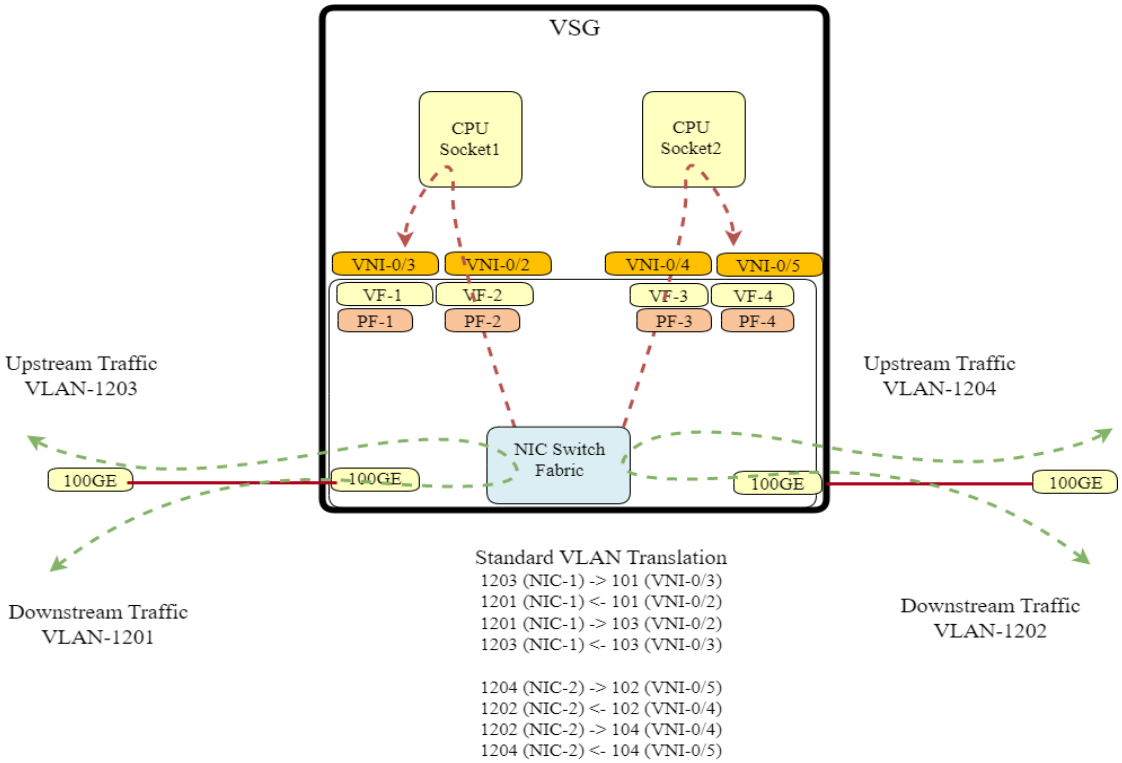

Figure 3 – VSG NIC Internal View

2.3. Virtual Services Gateway – Control Plane

The VSG Control Plane plays a pivotal role in managing and overseeing the operation of multiple VSGs. It serves as the backbone for administering subscriber data, applying network policies, and ensuring smooth communication between different components within the VSG eco-system. Its key responsibilities encompass the following:

1. **Subscriber Database Management**: The control plane maintains a comprehensive database of subscribers and their VSG mappings This mapping is for the control plane to track which VSG is serving a given subscriber. The control plane needs this information in order to apply a policy to the subscriber. This information is essential for policy provisioning and reporting.
2. **Policy Management**: It controls the policy information for both subscribers and VSGs, allowing administrators/users to define rules for traffic management, security, and quality of service. This ensures that network operations align with organizational policies and customer requirements.
3. **APIs and Infrastructure**: The control plane provides a robust set of Application Programming Interfaces (APIs) and the underlying infrastructure to manage VSGs and their associated databases of policies and subscribers. This facilitates interaction with the VSGs for configuration and monitoring purposes.
4. **Data Pipelines:** It offers numerous services to support the smooth operation of the VSG data pipelines. These services help streamline the flow of data and thereby enhance the reliability of upstream systems that heavily depend on this data.
5. **Representational State Transfer (REST) APIs for Various Functions**: The control plane provides RESTful APIs to support a wide range of functions, including:
    a. **Subscriber Policy Management**: APIs for setting subscriber policies, such as report frequency, priority, and access controls.
    b. **Virtual Network Functions (VNF) Policy Management**: APIs for managing VSG policies for traffic mirroring.
    c. **Data Collection and Interaction**: APIs that allow collection and interaction with VSG-generated data, such as subscriber usage and telemetry reports.
    d. **Exclusion List Management**: APIs to manage the exclusion list, which prevents certain subscribers from being subject to DDoS policy enforcement.

2.4. Virtual Services Gateway – Backend Processing

The backend processing system is responsible for processing a high throughput of streaming data at high volumes and low latency. The system uses data pipelines for data processing activities, which consist of series of interconnected activities where an activity's output is input to the next activity. The following data pipeline activities are performed by the system:

1. Report Generation from Network Element to Collectors
2. Report Streaming from Collector to Message Bus
3. Real Time Processing of Reports by a Streaming Application from Message Bus to Datastore

The data pipeline also addresses the availability, scalability, and durability system requirements by using the appropriate software frameworks or platforms for each activity in the data pipeline. The following are some of the data pipeline software components:

The data collection begins at the VSG, with the VSG collecting the defined metrics for its use-cases by inspecting the copies of the customer data packet. The VSG application software then assembles the data

into a predefined report structure and transmits the report using the IP Flow Information Export (IPFIX) format to an intermediate node called the Collector.

The Collector is an intermediate node in the processing chain serving an aggregate collection point for multiple VSG. The Collector receives the reports generated by the VSGs and converts the reports for transmission onto a Kafka message bus.

Apache Kafka is an open-source distributed message streaming platform to ingest the streamed telemetry reports in real time. This offers availability, scalability and durability requirements through its cluster, broker, and topic setups. The telemetry data is streamed to a Kafka topic used by external systems for data consumption.

Apache Flink is an open source distributed streaming analytical platform which ingests the data from the data stream sources such as Apache Kafka and performs data transformation and analytics. This offers the availability, scalability and durability requirements through its cluster setup and data stream processing APIs. The data is stored in analytical data stores for external system consumption.

Analytical datastores are to meet the availability, scalability and durability and the data freshness requirements for dashboards and alert management.

Grafana, an open-source observability application, is used for creating dashboards and alert monitoring systems. The alerts are sent to the alert manager on a notification channel used for subscription of the alerts from the external systems.

3. Challenges & Evolution

This section describes the challenge faced by the customer & network operator, and a summary of the proposed solution.

3.1. Distributed Denial of Service Attacks

A Distributed Denial of Service (DDoS) attack is a malicious attempt to disrupt the normal functioning of a targeted server, service, or network by overwhelming it with a flood of Internet traffic from multiple compromised systems [2]. The primary objective of a DDoS attack is to render the target inaccessible or unresponsive to legitimate users by exhausting its resources, such as network bandwidth, processing

power, or memory. An typical attack scenario is shown in Figure 4 – DDoS Attack Diagram.

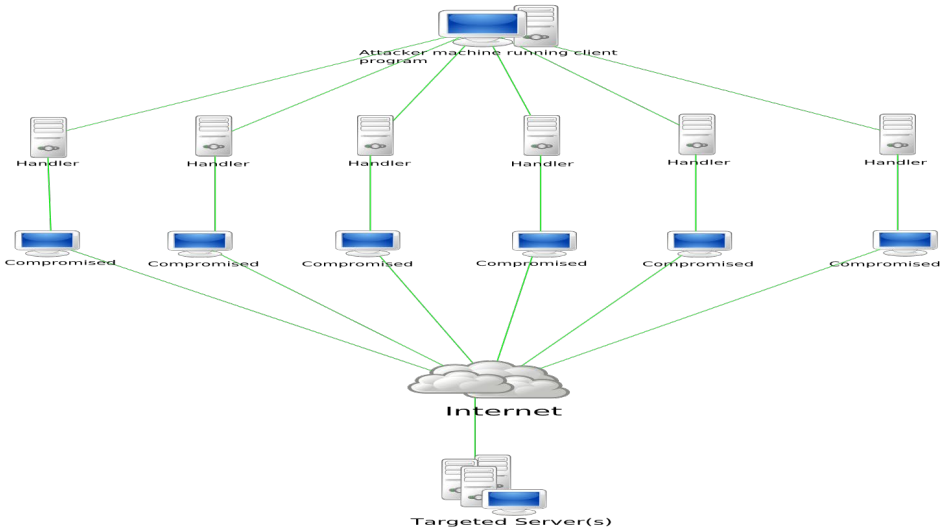

Figure 4 – DDoS Attack Diagram [3]

DDoS attacks pose significant challenges to the residential edge networks, which encompass Cable Modems and Customer Premises Equipment (CPE) devices, due to the substantial number of connected devices and the potential for extensive impact. These residential edge networks are particularly susceptible to DDoS attacks for several reasons:

1. **Insecure IoT Devices:** Many residential customers have a growing number of IoT devices, such as smart home appliances, security cameras, and smart TVs [4]. These devices often have weak security measures, default passwords, and unpatched vulnerabilities, making them easily compromised by attackers and recruited into botnets for launching DDoS attacks.
2. **Limited Bandwidth and Resources:** Residential networks typically have limited bandwidth and computational resources compared to enterprise networks. When subjected to a DDoS attack, the sudden surge in traffic can quickly overwhelm the network capacity, resulting in service degradation or complete unavailability for the affected customers.
3. **Amplification Attacks:** Attackers often exploit vulnerabilities in common protocols and services, such as DNS (Domain Name System), NTP (Network Time Protocol), and SNMP (Simple Network Management Protocol), to launch amplification attacks. These attacks leverage the asymmetry between small request packets and large response packets to amplify the attack traffic, making it easier to saturate the target's resources.
4. **Customer-Initiated Attacks:** In some cases, residential customers themselves may unknowingly participate in DDoS attacks. Malware-infected devices within the customer's network can be used as part of a botnet to launch attacks against other targets, consuming the subscriber's bandwidth and potentially leading to service disruptions.

3.2. Attack Targets

DDoS attacks can target a variety of devices connected to a home internet network. In a typical home internet setup, the most common targets for DDoS attacks are,

1. **Routers:** This is the central hub for all internet traffic in the home. By flooding it with requests, attackers can disrupt the entire home network.
2. **IoT devices (Internet of Things):** Devices like Smart TVs, Streaming devices, Smart thermostats, security cameras and voice assistants, often have weaker security measures and can be easily compromised.
3. **Computers & Mobile Devices:** While the attack to these is less common due to their stronger security features, they can be still targeted if they are compromised with malware or have open vulnerabilities.

3.3. Proposed Solution Summary

To address these challenges, a DDoS Protection feature is proposed as a solution implemented on the VSG, taking advantage of its unique position in the network as a flow aware network element coupled with the hardware processing power of its SmartNIC. For the VSG use cases, the VSG is receiving copies of the data plane (DP) packets to account byte usage and inform the network operator about the user experience, derived from telemetry data about the packet flow including the packet 5-tuple. The proposed enhancement to the VSG is to use the packet 5-tuple data, rate, and size to detect attack patterns compared against configured thresholds. When a customer's packet flow reaches the defined thresholds, the VSG dynamically reprograms the switching rules on the SmartNIC to redirect the original packet flow through the VSG application packet processing software where it can drop, or rate limit the attack packets while transparently forwarding on the non-attack packets. The next section of the paper captures further details on the solution.

4. Solution Details

This section delves into the details of the solution, covering the changes to each of the major system components: Network Element; Control Plane; Backend Processing

4.1. Virtual Services Gateway Changes – Network Element

The VSG requires all downstream and upstream traffic to be received on the SmartNIC. At the VSG, the VSG NIC makes a copy of the packet to send to the VSG application and changes the VLAN tag of the original packet for transmission back to the network. During VSG initialization, it will be configured with a pre-defined set of Security Policies which are defined based on the network industry's long history of learning attack signatures during large scale DDoS attacks. The VSG also provides a configuration utility to configure Security Policies during VSG run-time.

In steady state operation, the VSG utilizes this copy of the packet from SmartNIC to evaluate all the customers' traffic signature against these predefined policies. As soon as the VSG detects that a customer has violated a policy, VSG application will configure steering rules in the SmartNIC. These steering rules are responsible for steering the original customer traffic from the SmartNIC to VSG application. The VSG application will enforce traffic throttling or drops based upon the policy violation action for traffic flow that has violated the Security Policy, and allow customer traffic flows that do not violate the Security Policy to be sent back to the SmartNIC, which in turn allows this DDoS filtered traffic to flow back in the network. This is shown in Figure 5– DDoS Protection NIC View.

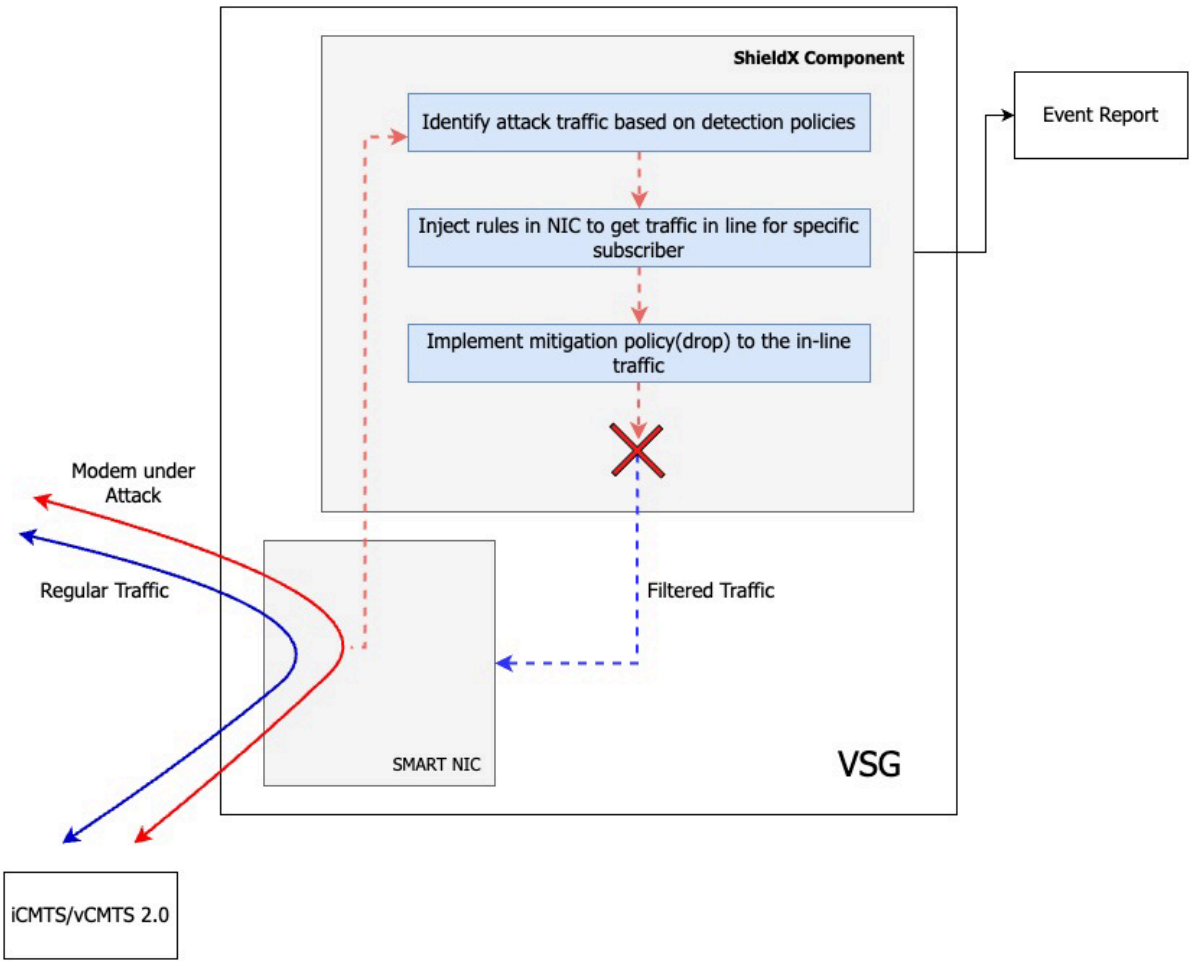

Figure 5 – DDoS Protection Feature - NIC View

The VSG will continue to enforce the DDoS Protection if the customer traffic pattern continues to violate one of the configured Security Policies, providing protection from the attack traffic while allowing non-attack traffic to flow seamlessly when the attack traffic subsides due to the attacker concluding attack attempts, the VSG application performs a Security Policy violation clear for that customer. The VSG application then updates the steering rules for the customer. This change reverts the customer data flow back to steady state operation with the VSG application again receives copies of the packets instead of the original packets. For consistent repeat offenders, the customer will be deny-listed. The customer will be removed from the deny-list once the pattern of attack traffic has been shown to stop over a period of time.

4.2. Virtual Services Gateway Changes – Control Plane

In the context of DDoS Protection, a false positive occurs when legitimate traffic is erroneously identified as attack traffic. This misclassification can happen when legacy applications, internet services, or faulty client software generate traffic that seems suspicious, exhibits unusual patterns, deviates from established best practices, or violates protocols.

The DDoS Protection feature might then flag this traffic as malicious and implement mitigation actions. If this traffic is legitimate and not part of an attack, these mitigation actions could cause service disruptions for the affected customers. To address this issue, the VSG Control Plane provides a mechanism to create an Exclusion List for subscribers whose traffic has been erroneously identified as attack traffic. The DDoS Protection feature will not apply the mitigation actions to customers on the Exclusion List.

The VSG Control Plane offers REST APIs to manage the DDoS Protection feature, allowing the following functions:

- **Create Exclusion List**: Establish an Exclusion List with one or more customers, ensuring that their traffic is exempt from mitigation actions.
- **Delete Exclusion List:** Remove customer(s) from the Exclusion List, enabling mitigation actions to be applied if their traffic is flagged as attack traffic in the future.
- **Get Exclusion List**: Retrieve a list of customers currently on the Exclusion List.
- **Get DDoS Policy**: Retrieve a list of customers who are subject to mitigation actions.

This allows administrators to effectively manage false positives in DDoS protection, minimizing disruptions to legitimate users while maintaining robust security.

4.3. Virtual Services Gateway Changes – Backend Processing

The DDoS Protection feature backend processing system is responsible for providing near real time visibility of problem identification and remediation. This involves processing customer violation reports that meet the violation criteria or policy configured on the Network Element (refer section 2.2). The key features of the backend processing system are violation report analytics and storage, alert management. The following is the overview of the key features in the backend processing system.

4.3.1. Violation Report Analytics and Storage

The streaming application is responsible for processing violation reports. These reports are enriched with the customer specific contextual information which includes geolocation and headend information. The enriched violation reports are serialized into specific data formats and stored in analytical data stores for efficient storage and retrieval. The business dashboards are configured on top of the analytical datastore to report on the violation report trends on hourly, daily, and weekly timelines at various aggregation levels such as geolocation, headend, and customer. The violation reports provide insights to track the security violations happening at several levels to gain near real time visibility for problem identification purposes.

4.3.2. Violation Report Alert Management

The Alert Management is responsible for providing alert notifications and subscription capabilities based on the violation report thresholds configured at several aggregation levels. The threshold configured is based on the historical data analysis of violation reports. The alerts are generated whenever the threshold criteria are met and are reported as notification events with corresponding severity. The external systems

subscribe to the alert event notifications to perform appropriate remediation actions necessary to address the alert notification.

5. Future Plans

The Edge Cloud Platform can have a tremendous amount of computing and network resources. The next level of sophistication is to build the system such that it is elastic and can auto-scale. These aspects enable the systems to automatically adjust the provisioned network and compute resources to match the current demand.

Future ideas for integrating the VSG with the VCMTS may involve "containerizing" the VSG application to run on the Edge Cloud Platform. The VSG application function may be disaggregated into its constituent functions, potentially with each function running in its own container. There are several possibilities on how to distribute the VSG container workload across the control and data compute nodes of the edge cloud. The key requirements for this solution are:

- The ability to create and destroy the VSG application containers automatically and/or on request, scaling the system as demand grows and subsides.
- The ability to copy and forward DP packets to the appropriate VSG application container
- The ability to copy and forward VSG control plane packets amongst the VSG application containers

One potential solution is shown in Figure 6 - Containerized VSG & VCMTS Deployment View. In this solution, the VSG pod supports all the needed VSG application containers running using available compute from a VCMTS "worker" node. The VSG pod supports both the VCMTS pod co-resident on the same compute and a VCMTS pod running on separate compute nodes. A variation on this solution is to only copy the packet headers instead of the full packet. This potential solution will reduce the overall link input/output bandwidth needed to support the packet copy function.
The key takeaway for these solutions is the capability of the SmartNIC to modify, copy, redirect, or forward packets both across and within the supported compute nodes. The SmartNIC in use also provides packet encryption using hardware offload, so it is possible to encrypt the VXLAN (Virtual eXtensible LAN) tunneled traffic between the VCMTS and VSG if desired.

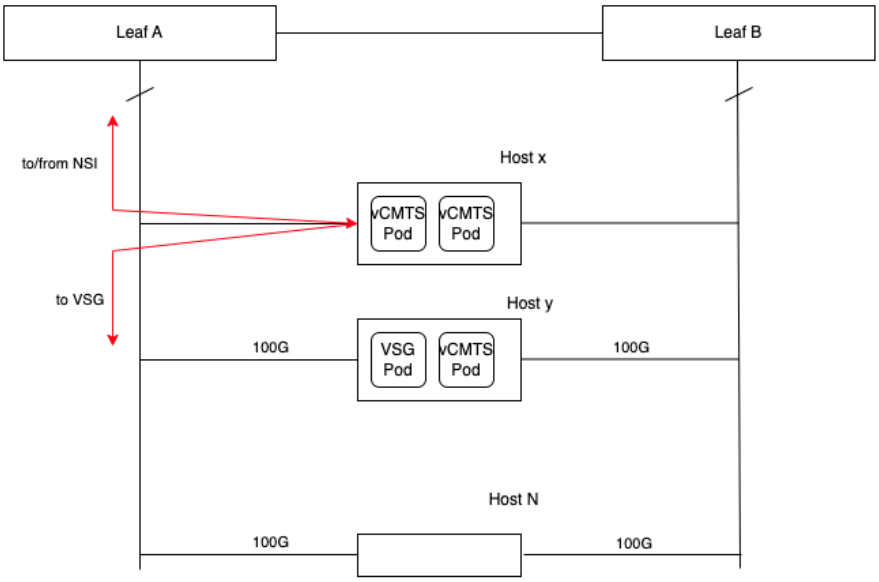

Figure 6 – Containerized VSG & VCMTS Deployment View

6. Conclusion

In the paper, we have proposed a new and unique solution for blocking or limiting malicious attack packets while preserving non-attack packets, providing a non-impactful DDoS Protection service to the customer. This DDoS Protection is engaged and disengaged dynamically as the attack packets begin and subsequently subside. Policy creation & deletion support is provided through the VSG management plane. Backoffice processing systems serve to refine and store the packet 5-tuple data for retrieval and analysis. Future enhancements include using AI / ML models to proactively detect new attack signatures and automatically install appropriate Security Policies. Further optimization of the overall system may be achieved by migrating the solution to a containerized system, allowing a more flexible deployment model.

Abbreviations

ACM	association for compute machinery
AI / ML	artificial intelligence / machine learning
API	application programming interfaces
AR	aggregation router
COTS	commercial-off-the-shelf
CCAP	converged cable access platform
CPU	central processing unit
DAA	distributed access architecture
DOCSIS®	data over cable system interface specification
DP	data plane
FDX	full duplex docsis
Gbps	gigabit per second
HSD	high speed-data
IEEE	institute of electrical and electronics engineers
LAN	local area network
M-CMTS	modular-cable modem termination system
MHA	modular head-end architecture
NFV	network function virtualization
NIC	network interface card
OFDM	orthogonal frequency division multiplexing
OVS	open vswitch
PCIe	peripheral component interconnect express
PF	physical function
QAT	quick assist technology
QOE	quality of experience
REST	representational state transfer
RF	radio frequency
RUR	residential u-ring router
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined network
SG	service group
SR-IOV	single root - input / output virtualization
TC	traffic control
TCP	transmission control protocol
UBB	usage based billing
UDP	user datagram protocol
VCMTS	virtual cable modem termination system
VF	virtual function
VLAN	virtual local area network
VNF	virtual network functions
VoIP	voice over internet protocol
VSG	virtual services gateway
VXLAN	virtual extensible lan
XRM	xfinity Resource Manager

Bibliography & References

[1] *Hyperscale Virtual Services Gateway*, Klatsky, C., Li, D., Combs, J., Grichina, A., Levy, A.; Society of Cable Television Engineers CableTEC Conference; October, 2023.

[2] https://www.cloudflare.com/learning/ddos/what-is-a-ddos-attack/

[3] https://commons.wikimedia.org/wiki/File:Stachledraht_DDos_Attack.svg ,Everaldo Coelho and YellowIcon, LGPL <http://www.gnu.org/licenses/lgpl.html>, via Wikimedia Commons

[4] https://www.trendmicro.com/vinfo/fr/security/news/internet-of-things/smart-yet-flawed-iot-device-vulnerabilities-explained

