Detecting and Mitigating Distributed Denial of Service Attack with Transparent Security

A Technical Paper prepared for SCTE by

Randy Levensalor
Principal Architect
CableLabs
858 Coal Creek Circle Louisville, CO 80027
(303) 661-3455
r.levensalor@cablelabs.com

Chris Sibley
Senior Engineer
Cox Communications
6305 Peachtree Dunwoody Rd. Atlanta, GA 3034
404-269-6701
chris.sibley@cox.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

Transparent Security is an open-source solution for identifying and mitigating distributed denial of service (DDoS) attacks and the devices (e.g., Internet of Things [IoT] sensors) that are the source of those attacks. Transparent Security is enabled through a programmable data plane (e.g., "P4"-based) and uses in-band network telemetry (INT) technology for device identification and mitigation, blocking attack traffic where it originates on the operator's network.

Cox Communications and CableLabs conducted a proof-of-concept test of the Transparent Security solution in the Cox lab in late 2020. Testing was primarily focused on the following major objectives:

- Compare and contrast performance of the Transparent Security solution against that of a leading commercially available DDoS mitigation solution.
- Validate that INT-encapsulated packets can be transported across an IPv4/IPv6/Multiprotocol Label Switching (MPLS) network without any adverse impact to network performance.
- Validate that the Transparent Security solution can be readily implemented on commercially available programmable switches.

CableLabs and Cox completed the lab trial in conjunction with Intel® and Arista Networks. Transparent Security was able to identify and mitigate attacks in less than one second as compared to greater than one minute for the leading vendor. We also validated that inserting and removing the INT header had no observable impact on throughput or latency.

Distributed denial of service (DDoS) attacks and other cyberattacks cost operators billions of dollars, and the impact of these attacks continues to grow in size and scale, with some exceeding 1 Tbps. The number of Internet of things (IoT) devices continues to grow rapidly, many have poor security, and upstream bandwidth is ever increasing—this perfect storm has led to exponential increases in IoT attacks, by over 600% between 2016 and 2017 alone. With an estimated increase in the number of IoT devices from 5 billion in 2016 to over 20 billion in 2020, we can expect the number of attacks and the size of attacks to continue this upward trend.

Detecting attacks is difficult, but mitigating them is even harder, and several solutions have been proposed with varying degrees of success. Typically, these solutions focus on blocking attacks that originate offnet and which targets an on-net resource. Even solutions that do identify attacks that originate on-net are limited in that they cannot block the attack traffic until after it has already traversed the access network. As such an operator's access networks can still be seriously affected, resulting in connectivity loss and quality of service (QoS) issues for customers.

Enhanced device visibility and packet processing can be used to identify the sources of such attacks. Leveraging programmable ASICs and P4 applications provides support for these enhancements by making the behavior of the data plane expressible in software and customizable without affecting performance. Specifically, this project will pair a DOCSIS modem with a series of P4-enabled devices connecting back to the operator headend via the P4 Runtime or Barefoot Runtime Interface (BRI). Both the P4 Runtime and BRI leverage GRPC as the underlying protocol. This architecture allows for visibility throughout the access network.

Transparent Security can leverage a machine-learning controller that has been trained with patterns to identify conditions and to perform dynamic operations by deploying new packet processing behaviors in the network (e.g., DDoS mitigation, virtual firewall, QoS detection/enforcement, and DOCSIS data plane functions). All operations will be performed at line rate while leveraging P4 in-band network telemetry

(INT), which allows data to be collected for reporting and analysis without control plane intervention. By inserting telemetry into the packet header, telemetry can be added to all packets rather than simply a sampling, which significantly reduces the time required to identify and mitigate the attack.

2. Motivation

As the proliferation of IoT devices continues to increase, the number of devices that can be compromised and used to participate in DDoS attacks also increases. At the same time, the frequency of DDoS attacks continues to grow because of the widespread availability of DDoS for-hire sites that allow individuals to launch DDoS attacks for relatively little cost. These factors contribute to a trend of malicious traffic increasingly using upstream bandwidth on the access network.

Typical DDoS mitigation solutions use techniques such as BGP diversion and Flowspec to drop traffic at certain parts of the network. However, mitigating outbound attacks using these techniques isn't entirely effective because the malicious traffic will have already traversed the access network, where it has the greatest negative impact before the traffic can be diverted to a scrubber or dropped by a Flowspec rule. Additional information on Flowspec can be found in section 4.2.2.2.

Transparent Security offers the promise of near-instantaneous detection of outbound attacks, as well as the ability to mitigate that attack at the source, on the customer premises equipment (CPE), thereby preventing that traffic from using upstream access network resources.

In addition to Transparent Security's DDoS mitigation capabilities, there are additional benefits to network performance/visibility in general. Implementation of Transparent Security on the CPE means that network operators can derive the specific device type associated with a given flow. This allows the operator to determine the type of IoT devices being leveraged in the attack.

This also opens myriad other possibilities—for example, reducing truck rolls by enabling customer service personnel to determine that a customer's issue is with one specific device versus all the devices on the internal network. Another example would be the capability to track the path a given packet followed through the network by examining the INT metadata.

Consumers will see a direct benefit from Transparent Security. Once compromised devices are identified, the consumer can be notified to resolve the issue or, alternatively, rules can be pushed to the CPE to isolate that device from the internet while allowing the consumer's other devices continued access. Such isolation mitigates the additional harm coming from compromised devices. This additional harm can take the form of degraded performance, exfiltration of private data, breaks in presumed confidentiality in communications, as well as the access network bandwidth consumed through DDoS. Less malicious traffic on the network provides for a better overall customer experience.

3. The History and Updates of Transparent Security

We initially released the Transparent Security architecture and open-source reference implementation in October 2019. Since then, we've achieved several milestones:

- Added source-only metadata to the P4 in-band telemetry specification, along with Transparent Security as an example implementation.
- Added support in the Telemetry Report 2.0 specification to collate multiple packet headers in a single telemetry report.
- Released a document titled "Transparent Security: Personal Data Privacy Considerations."

- Created a Transparent Security landing page.

Most proposed DDoS solutions fall into one of two categories, detection, or mitigation. This section discusses additional possible solutions in those categories and examines their advantages and disadvantages.

4. Existing Solutions

4.1. DDoS attack detection

DDoS attack detection is typically identified by an analytics engine identifying network traffic trends. These trends are based on routers sampling a summary of the packets as they enter the network. These samples can use IPFIX, NetFlow and SFlow to export this data.

The sampling technique cannot offer an operator a complete view of the network. Sampling is a statistics-based method for measuring network traffic by collecting, storing, and analyzing a sample of traffic data. This method has the advantage of allowing many interfaces to be monitored without significantly affecting network traffic.

IPFIX, SFlow and NetFlow are protocols used for sampling data. IPFIX and NetFlow are very similar, with IPFIX extending NetFlow v9 (https://www.oreilly.com/library/view/practical-network-scanning/9781788839235/1d6b69c7-62c3-40d0-be58-1ad82b22c115.xhtml). All these methods are limited to the data in the packet and do not contain any information about the network path used by the packet, validation for the source IP address or visibility to the source IP/device on a network behind a NAT.

Although the sampling method can effectively identify larger trends, it can miss anomalies that occur in a smaller portion of the traffic. That is, it can effectively detect a DDoS attack directed at a target, but it can miss or take longer to detect a DDoS attack originating at a given source. With source-based DDoS attacks, the system is trying to identify smaller anomalies, which are less likely to be captured in the sample.

4.2. DDoS Attack Mitigation

There are a wide variety of methods for mitigating DDoS attacks. A few of the common methods and the methods used by Transparent Security are highlighted below. This list is not exhaustive.

Also examined in this section are the points on the network where the mitigation is performed.

4.2.1. Network Locations for Mitigation

4.2.1.1. Out of Band

Out-of-band DDoS mitigations come in two flavors, appliances and scrubbing services. When an attack is detected against a host in the network, the traffic is routed through the out-of-band device, where it can remove or re-route the malicious traffic. An appliance routes traffic to the target to itself then removes the malicious traffic and provides a clean flow to the target. A service functions similarly except it routes target traffic to a mitigation center that cleans the traffic and forwards it to the destination.

However, both methods of out of band DDoS mitigation have downsides. The appliance is generally costly, and deployment can be complex. The service has difficulty defending low-bandwidth slow attacks, and every interface must be protected, or it can fail to stop some attacks. Both methods can add latency to network traffic during an attack, and they generally rely on other methods for attack detection.

4.2.1.2. In Band

Traditional in-band DDoS mitigations are appliance based and work in the network path, comparable to a firewall, allowing the method to see all network traffic and react accordingly. As a result, in-band mitigation can provide detection as well as mitigation. However, it is costly, adds to network complexity and introduces another point of failure in the network.

4.2.1.3. At the Source

Transparent Security enables network operators to provide DDoS detection and mitigation at the source. This method protects target organizations with little to no effort on their part. It can also provide insight into hacked devices on the customer premises. This information can go a long way into preventing future attacks. Mitigation at this location has the added benefit of limiting attack traffic on an operator's network by blocking malicious packets before they enter the core network. Transparent Security mitigates primarily at the source, but the method is compatible with additional in-band target-based mitigation methods.

4.2.1.4. At the Target

While DDoS mitigation can be deployed at the target, it is generally of little use unless the attack is small (less bandwidth than the access circuit and less packets per second than the CPE is capable of processing).

4.2.2. Methods for mitigating DDoS attacks

4.2.2.1. Scrubber

As the name suggests, a scrubber cleans traffic to a specific host that is under attack. When an attack is detected, the traffic to that host is diverted to the scrubber—malicious packets are removed, and clean packets are forwarded. Different types of scrubbers have issues with either high-volume attacks or low-volume attacks. See Section 2.2.2.1, "Out-of-Band," for explanations of the weaknesses of out-of-band packet scrubbing. Scrubbers also require additional hardware in the network, which adds capital and operational costs for the service provider.

4.2.2.2. Flowspec

BGP Flowspec is an IETF specification (https://datatracker.ietf.org/doc/html/rfc7674) for diversion and filtering of malicious traffic. With Flowspec, the DDoS mitigation solution sends a BGP message to the routers which are forwarding the attack traffic. That router then blocks, rate limits, or forwards the traffic matching the Flowspec pattern.

4.2.2.3. Proprietary Matching Engine

Switch vendors have started developing proprietary packet matching and manipulation engines. Many of these engines could be leveraged to support DDoS mitigation.

5. Transparent Security Architecture

Transparent Security uses programmable data plane capabilities to enable real-time packet processing, high-resolution packet inspection, and in-band network telemetry (INT). INT allows MSOs to identify the compromised devices in milliseconds rather than minutes. This technology is enabled by new, currently available programmable chips that can process packets at line speed and be deployed at any point in the network, from the core network to residential and business customer premises.

Transparent Security is focused on inspecting, finding, and blocking malicious packets as close to the source as possible by adding details about the packet's source device, exact route through the network, and travel duration. INT can then be used by upstream processes to identify traffic patterns and then act on that information.

5.1. Data Plane Architecture

Figure 1 shows a typical data plane architecture of customer premises connected to the access network into an operator's core network. Any combination of the customer gateways, switches, and routers can be P4 enabled or not. By enabling P4, the available use cases increase substantially, but this architecture can be implemented in stages. Since the INT metadata is encapsulated in a UDP shim, there is no impact to L3 routing. Packets with the INT header can traverse network devices which do not support INT header insertion. Firewalls and other devices which use L4 headers will need to be updated to support INT to inspect the original L4 header.

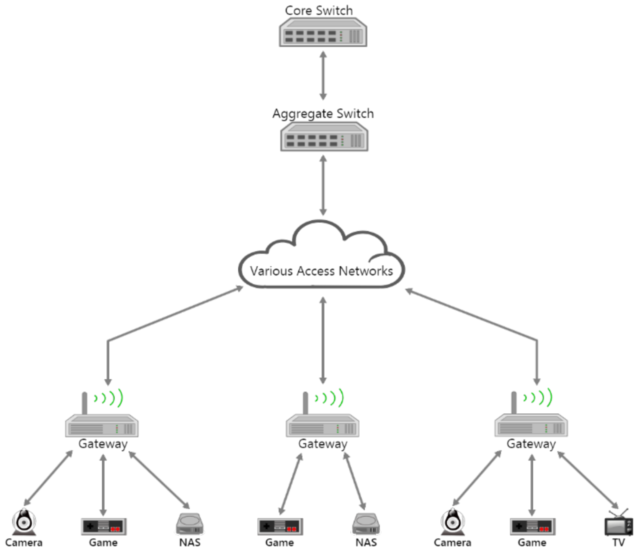

Figure 1 - General Architecture for the Data Plane in the Transparent Security Model

5.1.1. In-Band Network Telemetry

INT is a method for adding telemetry data to every packet at multiple points on the network. This approach can help determine things such as the path of a packet or the source device emitting packets. With INT, the packet only needs to be inspected at the edge of the network, which reduces the overhead and generates less traffic compared with sampling at multiple points in the network.

The P4.INT specification suggests a set of predefined fields:

- switch ID,
- control plane version,
- ingress/egress port,
- timestamp,
- RX packet count, and
- congestion status.

By leveraging P4 to implement INT, the data can be customized to meet the needs of a service provider. With this data, one can identify the source of a packet behind a firewall by including the originating MAC address in the INT header. With this data, one can obtain the specific source of a packet behind a firewall. Because customer premises and wireless networks are dynamic and multiple devices share the same ingress point on the gateway, knowing the ingress port is not sufficient when trying to identify the packet source. These extensions to the standard P4 INT data structure provide the packet source's MAC addresses, which are unique and are required to identify specific devices for customer premises and wireless networks.

One potential issue with INT is that it increases the size of the packet header. If this increase exceeds the frame size, it will cause packet fragmentation, which has a negative impact on network performance. This issue should not arise with Transparent Security, however, because the INT data is added at the gateway. With access networks, such as DOCSIS, the frame size is larger between the gateway and the core network than at the customer premises. Once the packet reaches the service provider core, INT metadata will not be added to the packet if doing so would cause the MTU size to be exceeded.

INT header uses UDP encapsulation and domain specific source only metadata. Transparent security is the reference example for source only metadata in the 2.1 version of the P4 INT specification (https://github.com/p4lang/p4-applications/blob/master/docs/INT_v2_1.pdf).

The INT header and metadata data will be removed after the telemetry report has been generated and before the packet leaves the network being monitored by transparent security.

5.1.2. Switch or Gateway with P4

The features available with networking hardware supporting the P4 language allow for the development of flexible, open, and consistent DDoS mitigation solutions. Using the information gleaned from the INT data, it is possible not only to detect an attack very quickly but also to identify the compromised device. Once identified, the switch or gateway can be notified to reroute or drop the problematic packets with a simple match rule performed at line speed when deployed with hardware acceleration in software on a gateway device containing a P4 chiplet. Additional information on P4 can be found at https://p4.org/.

5.2. Control Plane Architecture

Figure 2 depicts an example control plane architecture in which the analytics engine receives INT data from the core in-band as packets flow through the network. When malicious patterns are detected, the SDN controller is notified and updates the P4-enabled devices to handle the packets based on the pattern signature. The management interface between the controller and the P4-enabled devices can leverage a variety of protocols, including GRPC, Thrift, HTTP, or RPC. The protocol between the SDN controller and switches can vary, depending on the protocols supported by the switches and gateways. Telemetry data and alert notifications can optionally be sent to a dashboard or NOC server for integration with other analytics.

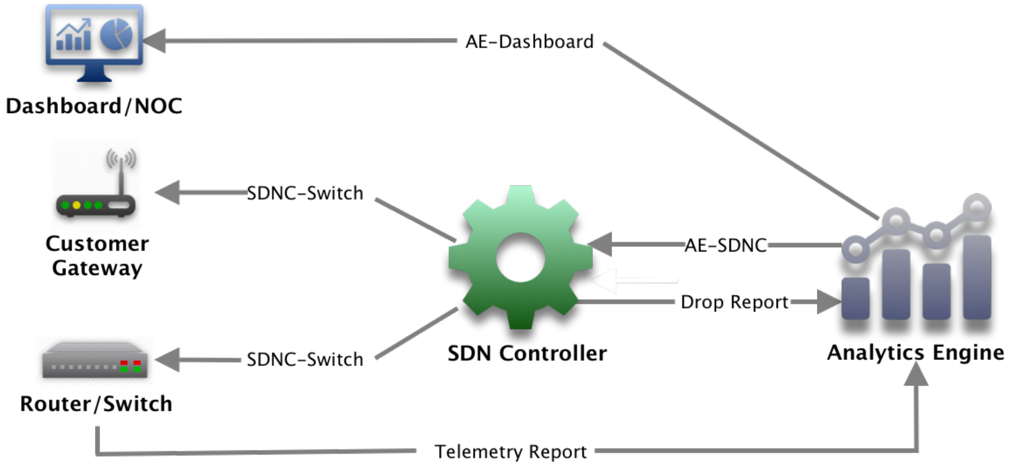

Figure 2 - General Architecture for the Control Plane in the Transparent Security Model

6. Components

6.1. Transparent Security Analytics Engine

The analytics engine (AE) serves as the intelligent core of the programmable data plane. Its purpose is to analyze telemetry reports and make inferences relative to generally defined patterns. For this use case, these patterns are limited to DDoS attack identification. The AE could be extended to manage Quality of Service (QoS), proactive network maintenance (PNM) or other use cases.

When an attack it identified, the AE informs the SDN controller, which is responsible for implementing network changes through the control plane.

The AE's include the following functions:

- identify DDoS attacks,
- mitigate the attacks, and
- remove inactive mitigations.

Multiple iterations of the AE have been tested as the project matures. The initial implementation was a Python service, which would parse the telemetry reports and identify the attacks. This solution was rigid and hard to scale. It was however able to detect and mitigate an attack in about 1 second and as such did meet our needs for detecting UDP flood-based attacks. This implementation can be found here: https://github.com/cablelabs/transparent-security/tree/master/tests/trans_sec/analytics

The second implementation leveraged a Painless pipeline to parse the telemetry reports and Elasticsearch was used to ingest the incoming data, visualize it and identity the attacks. This solution allows for multiple attack identification pipelines to use a shared parser. This architecture should scale without issues. The primary limitation of using Elasticsearch trend analysis is it took around 1 minute to identify an attack, which drove the minimum time to detect an attack outside of the target for transparent security. This implementation can be found here: https://github.com/cablelabs/transparent-security/tree/master/snaps-hcp

The current version of the AE is based on Siddhi. This solution provides modularity and scalability while still identifying attacks in about 1 second. Information on this implementation can be found here: https://github.com/cablelabs/transparent-security/blob/master/docs/SIDDHI_AE_SETUP.md

6.1.1. DDoS attack identification

Telemetry reports which include INT data are used to provide samples of the traffic running across the network. The telemetry reports use the telemetry report interface as noted in Figure 2 - General Architecture for the Control Plane in the Transparent Security Model. Telemetry reports are used instead of IPFIX since they contain the full packet header and a fragment of the data. With this richer data the AE will be able to employ attack identification techniques which are just not possible in DDoS identification applications that are limited to sampled flow data.

6.1.2. DDoS attack mitigation

The AE notifies the SDN controller of an active DDoS attack. This is sent over the AE-SDN interface as shown in Figure 2 - General Architecture for the Control Plane in the Transparent Security Model. This notification includes a signature of the attack, so that the SDN controller can mitigate it. The current implementation uses a rest call directly to the SDN controller. It is possible in production to have multiple SDN controllers. This interface may be updated to use a pub/sub model where SDN controllers can subscribe to message bus, such as Kafka (https://kafka.apache.org/protocol.html) or MQTT (https://docs.oasis-open.org/mqtt/mqtt/v5.0/mqtt-v5.0.html)

6.1.3. Remove inactive mitigations

The AE receives drop telemetry reports from the SDN controllers. These reports provide counters for all active mitigation rules on the network. These are sent over the drop report interface as shown in in Figure 2 - General Architecture for the Control Plane in the Transparent Security Model. The AE determines when the mitigation rule has not been used for a specific period which is an indication the attack has ended. Once an attack has ended, the AE will notify the SDN controller to remove the inactive mitigation rule.

6.2. SDN Controllers

The information used by the AE is captured by the P4 devices as part of the data plane and forwarded to the AE in line with the network traffic. When the SDN controller is informed by the AE that a DDoS attack has been detected, it updates action tables in the P4 devices through the control plane management interface (using protocols such as GRPC or Thrift). Thus, the P4 devices gain additional abilities in the data plane with a minimally intrusive controller activating them before consequences are experienced.

The controller’s functions include the following:

- manage the network configuration on switches and gateways,
- push DDoS mitigation to managed devices,
- track which devices are participating in an attack by querying counters of dropped packets based on DDoS mitigation from the device preforming the mitigation,
- send periodic drop reports to AE with the dropped packet counters for each mitigation, and
- remove DDoS mitigation from managed devices.

6.3. Programable Data Plane

The programable data plane consist of routers, switches and customer gateways used in Transparent Security to add/remove the INT header, generate telemetry reports, and mitigate DDoS attacks. To date, the testing and development has focused on performance on devices which support P4. These devices provide a great deal of flexibility to quickly develop the data plane for transparent security. It is possible to deploy transparent security using other types of programable devices. See 4.2.2.3 for additional information on these types of devices.

The programable data plane functions include the following:

- manage the traffic across the network,
- add P4.INT data (source port, time, queue),
- add Transparent Security data (source MAC),
- remove P4.INT headers before they leave the service provider network,
- send telemetry reports to the AE and,
- mitigate DDoS attacks from the core network and the aggregate network.

6.4. User Equipment

The end user devices are the typical end units used by end customers, including IoT devices, phones, laptops, and an ever-increasing variety of devices. Many are connected to the gateway over Wi-Fi or LTE. Implementing transparent security does not require any changes to user equipment. The user equipment should not be able to detect if transparent security is deployed or not.

7. Message Flows

The following sections outline high-level message flows between the possible components.

7.1. Packet Flow and Alerts

7.1.1. Standard Packet Flow

Standard packet flow:

1. Customer device sends packets to the server provider network.
2. The packet will traverse a series of gateways, switches, and routers where some are aware of INT and others are not.
3. This first INT enabled switch will insert the INT UDP header and the metadata for that hop.
4. Incremental networking devices with INT support will add their node ID to the INT metadata.
5. Packet is forwarded to the final INT networking devices, such as a PE router.
6. Packet is cloned and sent to the analytics engine.
7. INT data are removed from the packet.
8. Packet is forwarded to the Internet.

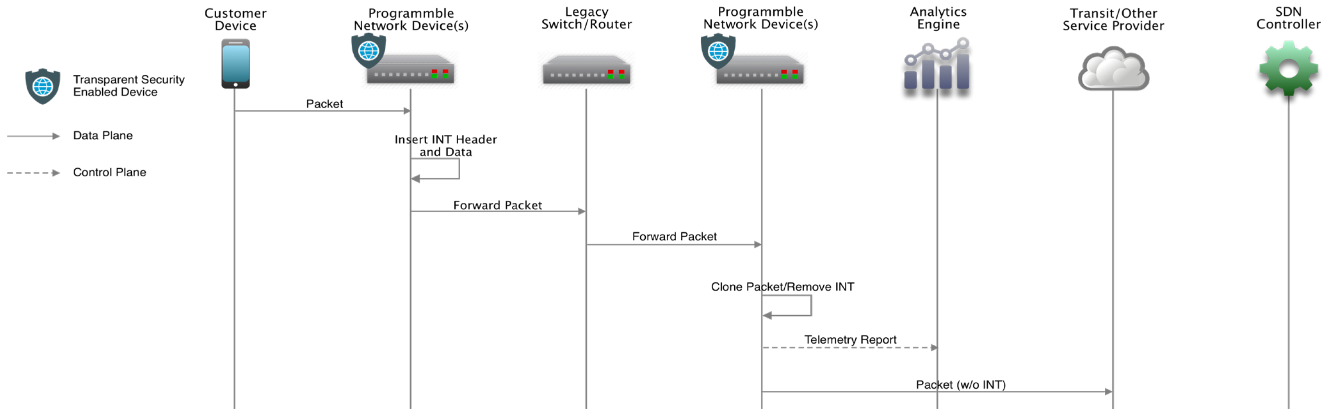

Figure 3 - High-Level Message Flow in the Control Plane and Data Plane when Packets Are Allowed Through.

7.1.2. DDoS Attack Identification and Mitigation

The following actions occur when an attack is detected:

1. Analytics engine recognizes an attack.
2. AE notifies the SDN controllers of the attack with its signature.
3. The SDN controllers notify the proper networking device to add an entry in the transparent security drop table for this specific attack signature.
4. The networking device add the entry into its transparent security drop table.

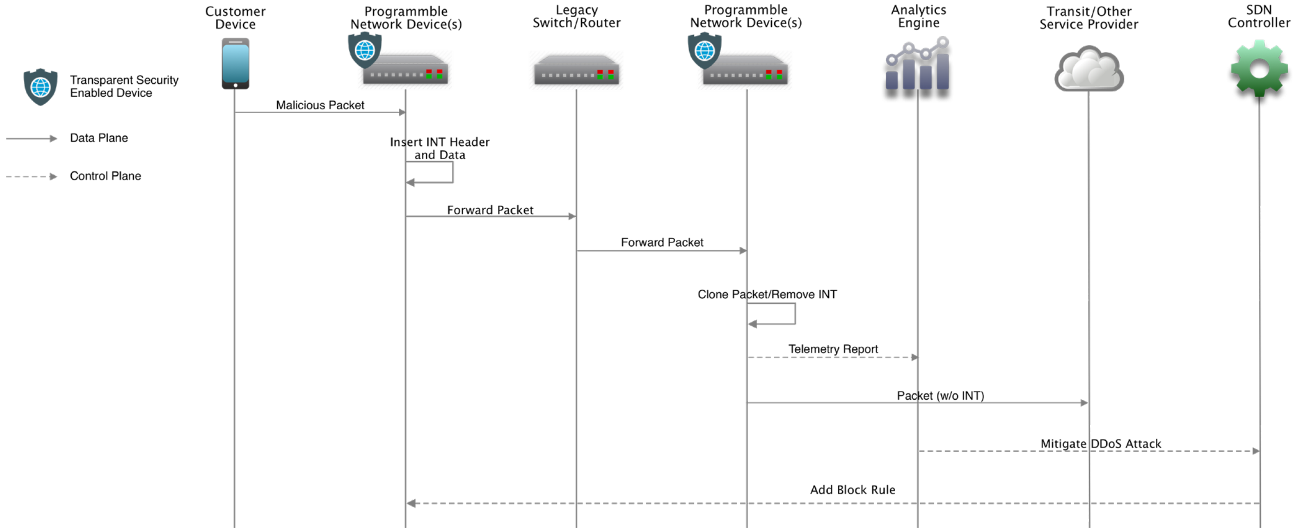

Figure 4 - High-Level Message Flow in the Control Plane and Data Plane during DDoS Attack Identification and Mitigation

7.1.3. Mitigated Attack Packet Flow

1. A packet as a part of the DDoS attack is sent to the service provide network.
2. The first transparent security enabled networking device with a drop rule drops the packet.
3. After the packet is dropped, the device increments the drop counter for that rule.
4. The SDN controller periodically collects the drop counters and sends them to the AE in drop telemetry report.

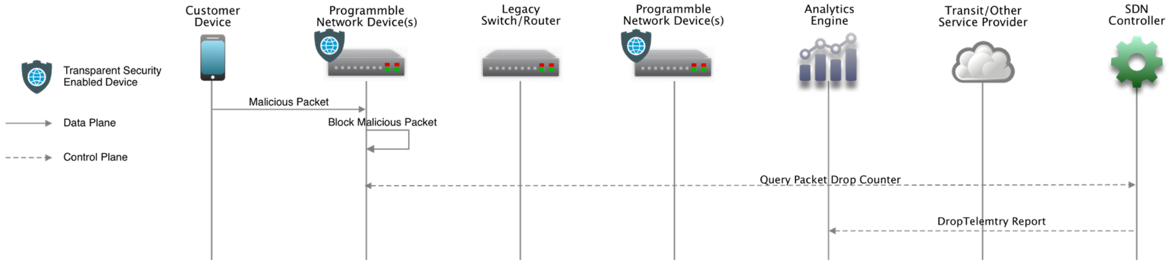
Customer Device
Programmble Network Device(s)
Legacy Switch/Router
Programmble Network Device(s)
Analytics Engine
Transit/Other Service Provider
SDN Controller
Transparent Security Enabled Device
Data Plane
Control Plane
Malicious Packet
Block Malicious Packet
Query Packet Drop Counter
DropTelemtry Report

Figure 5 - High-Level Message Flow in the Control Plane and Data Plane while Mitigating an Attack.

8. Deployment options

Transparent Security can be deployed across the service provider network in several phases. A phased deployment makes it easier to deploy and realize many of the benefits while postponing any changes to cable modems (CMs). The process of a phased deployment begins at the hub or head end and moves to the customer premises at later phases.

8.1. Core/Head End/Hub Deployment

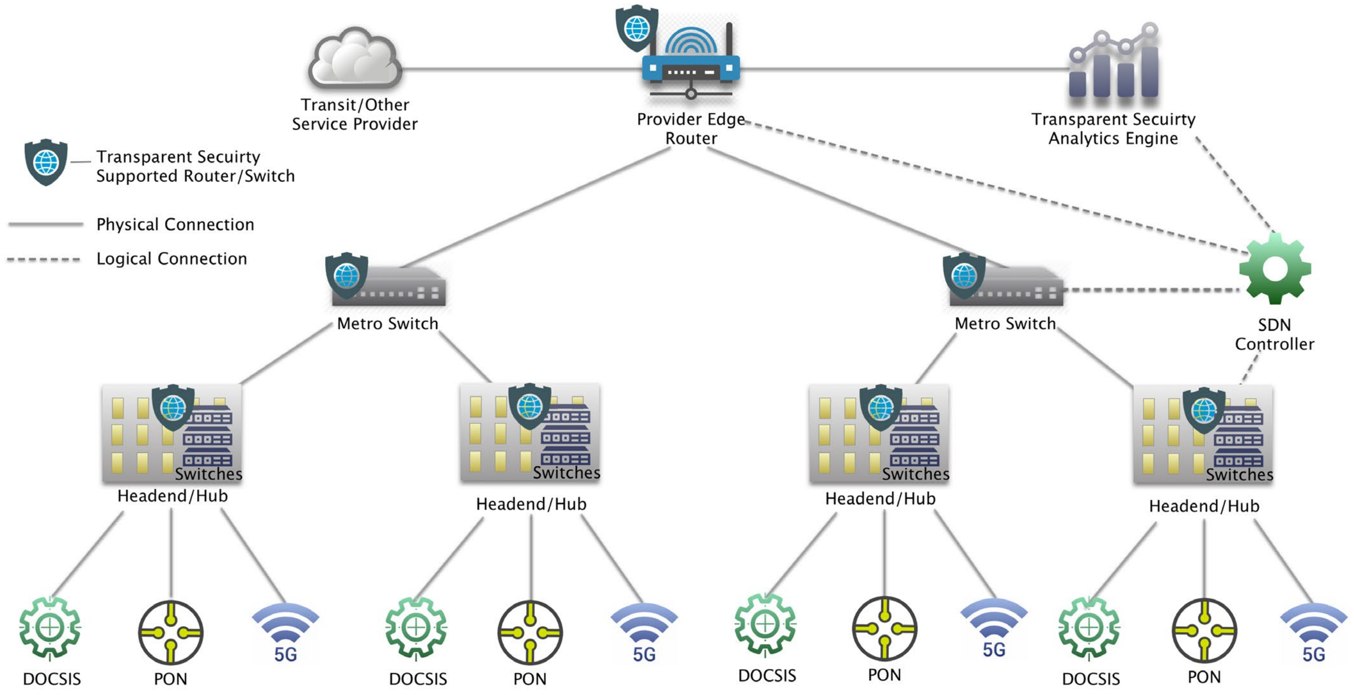

Figure 6 - Phased Deployment: Head End/Hub Updates

This deployment option focuses on the service provider's core network. Two classes of attacks can be mitigated. Attacks from this service provider's network against targets outside the network and attacks across the service providers network can be detected and mitigated with this model.

The first phase of deployment begins with upgrading the head end/hub by adding programmable switches with an analytics engine and an SDN controller. These switches can be included as part of distributed access architecture (DAA) upgrades. They provide the ability to mitigate a DDoS attack from a customer at the head end and identify which customer is the source of the attack.

This solution can also address ingress attacks from outside of the hub which is an advantage over a typical edge-based DDoS mitigation model.

There are no changes to the CM for this phase. The primary limitation to this solution is that it cannot identify the device originating the attack, only the customer. The compromised device remains active and continues to participate in ongoing attacks.

8.2. Customer Premisis

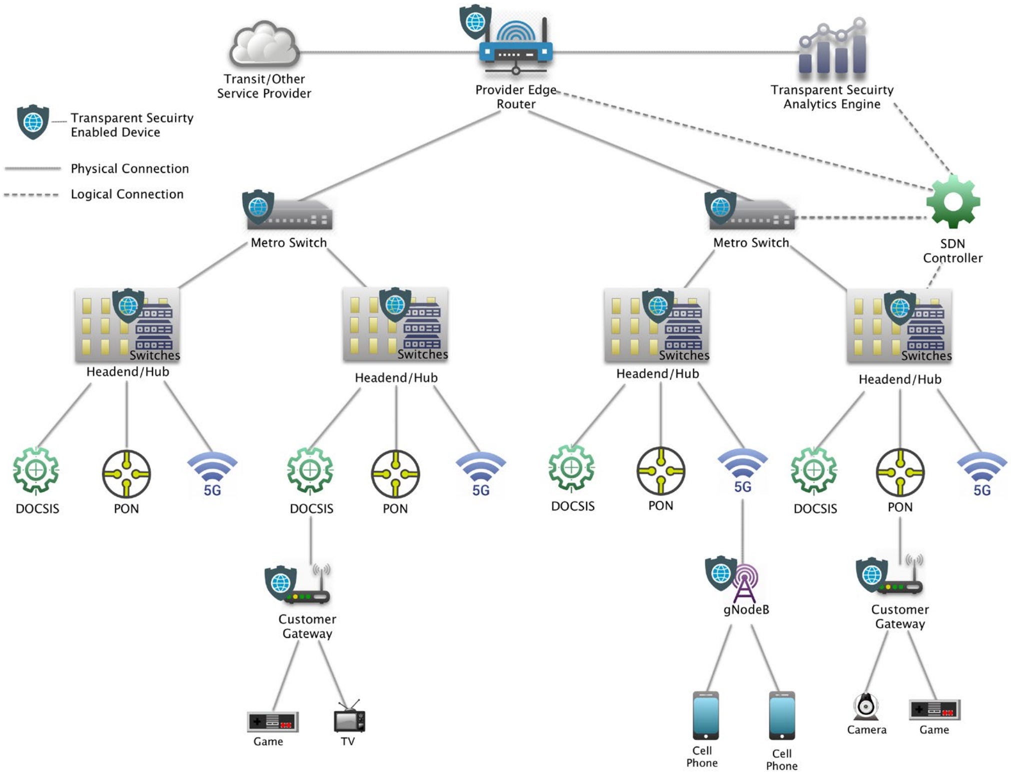

Figure 7 - Phased Deployment: Customer Premisis

In this deployment model, transparent security is deployed on a gateway device. For 5G, this can be deployed on a gNodeB. The initial INT insertion and DDoS mitigation is performed on the gateway device.

Once this model is deployed, the traffic on the access network will be scrubbed before it can impact any other customers, and customers will be able to address issues on compromised devices. The malicious traffic is dropped, and the count of blocked packets for each device is tracked. This model will not block benign traffic from the compromised device.

9. Lab Trial

To validate the viability and network impact for transparent security, Cox Communications conducted a lab trial with transparent security. This trial focused on a derivation of the Core/Head End/Hub deployment option. The goal of this trial was to validate that there is no adverse impact on network performance when adding transparent security. This includes latency, throughput and basic network routing.

9.1. Lab Trial Setup

The test environment was designed to simulate traffic originating from the access network, carried over the service provider's core backbone network, and targeting another endpoint on the service provider's access network in a different market (e.g., an "east-to-west" or "west-to-east" attack).

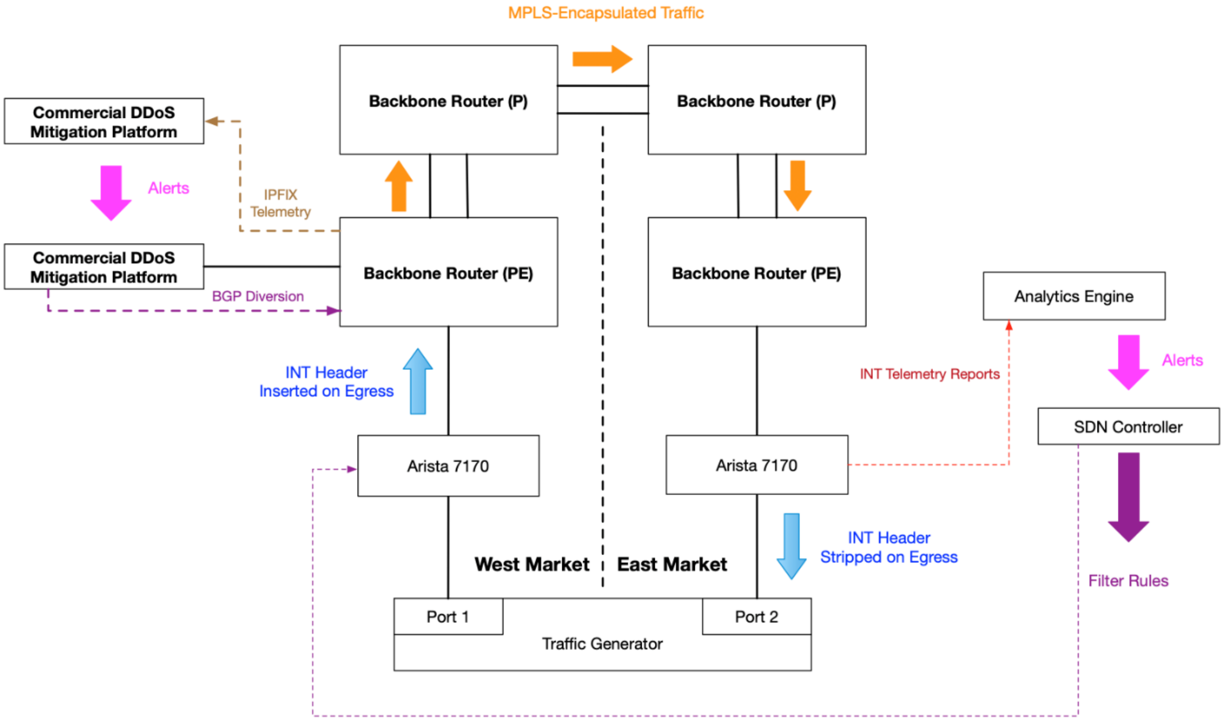

Figure 8 - High-level overview of the lab test environment

In the lab trial, various types of DDoS traffic (UDP/TCP over IPv4/IPV6) were generated by the traffic generator and sent to the West Market Arista switch, which used a custom P4 profile to insert an INT header and metadata before sending the traffic to the West Market PE router. The traffic then traversed an MPLS label-switched path (LSP) to the East Market PE router, before being sent to the East Market Arista, which used a custom P4 profile to generate INT telemetry reports and to strip the INT headers before sending the original IPv4/IPv6 packet back to the traffic generator.

9.2. Results

When comparing the performance of the Transparent Security solution against that of a leading commercially available DDoS mitigation solution, the lab test results were very promising. Detection and mitigation of outbound attacks was rapid, taking around one second. The commercial solution took 80 seconds to detect and mitigate the attack. Randomized UDP floods, UDP reflection and TCP state exhaustion attacks were identified and mitigated by both solutions. In this trial, only packets related to the attack were dropped. Packets not related to the attack were not dropped.

The Transparent Security solution was implemented on commercially available programmable switches provided by Arista. These switches are being deployed in networks today. No changes to the Networking Operations System (NOS) were required to implement Transparent Security.

The tests validated that INT-encapsulated packets can be transported across an IPv4/IPv6/MPLS network without any adverse impact. There was no observable impact to throughput when adding INT headers, generating telemetry reports, or mitigating the DDoS attacks. We validated that the traffic ran at line speed, with the INT headers increasing the packet size by an average 2.4 percent.

Application response time showed no variance with or without enabling Transparent Security. This suggests that there will be no measurable impact to customer traffic when the solution is deployed in a production network.

10. Conclusion and Next Steps

Transparent Security uses in-band telemetry to help identify the source of the DDoS attack.

This trial focused on using Transparent Security on switches inside the service provider's network. For the full impact of Transparent Security to be realized, its reach needs to be extended to gateways on the customer premises. Such a configuration can mitigate an attack before it uses any network bandwidth outside of the home and will help identify the exact device that is participating in the attack.

This testing took place on a custom P4 profile based on our open-source reference implementation. We would encourage vendors to add INT support to their devices and operators to deploy programmable switches and INT-enabled CPEs.

Take the opportunity today to explore the opportunities for using INT and Transparent Security to solve problems and improve traffic visibility across your network.

Distributed denial of service (DDoS) attacks and other cyberattacks can cost operators billions of dollars. With more and more devices coming into customers' homes and businesses, many of which with less-than-optimal security, this problem will only get worse. By quickly identifying attacks and blocking them before they reach the access network, Transparent Security can:

- reduce the operations impact of large-scale attacks by mitigating closer to the source within seconds of an attack starting,
- eliminate the risk of revenue impact resulting from a failure to meet service-level agreements,
- protect and enhance customer sentiments by avoiding large-scale DDoS attacks within a network, and
- provide a data stream from which new analytics and innovations can be built.

10.1. Alternate Applications for Programmable Data Plane

The Transparent Security source-based DDoS use case is just one of many that can benefit from the programmable data plane. Once deployed, this architecture can improve many of the operations performed today. It will also open networks to new waves of innovation and allow operators to provide such things as network optimization and additional customer services.

With the programmable data plane, it is possible to change the behavior of the network after hardware has been deployed. Use of an analytics engine and controller, as is done with Transparent Security, can provide a closed-loop automation. Some of the network management capabilities available with the programmable data plane are listed below:

- packet flow tracking and optimization,
- prioritization on low-latency flows,
- active network monitoring/management, and
- traffic shaping.

Providing new services frequently requires installing new purpose-built hardware or deploying a virtual machine. With the programmable data plane, some of these services can be deployed on existing switches with very good performance. Some of the services that can be deployed with the programmable data plane include the following:

- firewalls,
- managed router as a service,
- Layer 4 load balancing, and
- SD-WAN (software-defined networking in a wide-area network).

Micronets is another CableLabs project that can leverage the programable data plane. Micronets creates additional layers of security within the customer premises and helps protect devices by using OpenFlow for the L3 traffic management. Transparent Security is focused on identifying devices participating in a DDoS attack and mitigating the attack at the source. The programable data plane, analytics engine, and controller used in Transparent Security can be leveraged by Micronets to improve packet processing performance and provide access to additional fields in the packet header. For more information on the Micronets project, visit www.cablelabs.com/micronets.

Abbreviations

INT	In-band telemetry
CPE	Customer premises equipment
DDoS	distributed denial of service
IoT	Internet of Things