CISCO SECURE

White Paper
Cisco Public

Cisco Secure DDoS Edge Protection:

Security at the 5G Network Edge

January 2023

The bridge to possible

Contents

1. Introduction

1.1 Introduction

5G networks promise to reshape the future of communications by offering significantly faster data speeds, accommodating more devices than 4G networks, and enabling advanced applications like 3D video. To deliver these benefits, 5G applications must be located as close to customers as possible so that network operators can meet sub-10-ms 5G latency requirements. Unfortunately, the changes in network topology required by 5G also expose the network to devastating distributed-denial-of-service (DDoS) attacks intended to compromise security, disrupt network services, and render applications unavailable to users.

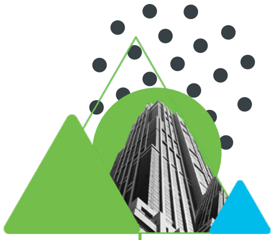

Overview of Cisco Secure DDoS Edge Protection

Cisco Secure DDoS Edge Protection® is an innovative software solution that stops cyberattacks at the service provider network edge. The Edge Protection solution consists of a controller and one (1) or more detectors. When deployed on Cisco® NCS 540 routers, Edge Protection detects and mitigates DDoS attacks at the cell site router. By moving DDoS protection to the network edge, service providers can meet the sub-10-ms latency requirements of 5G applications and ensure customer quality of experience (QoE).

Going deeper

With Cisco Secure DDoS Edge Protection, the security perimeter is pushed beyond the User Plane Function (UPF). This allows the cell site router to become the first line of defense against DDoS attacks from compromised User Equipment (UE) devices. Deployed on the Cisco NCS 540 router, Edge Protection defends against IoT (Internet of Things) and UE distributed attacks by providing not only full detection, but also granular mitigation capabilities in a small, lightweight containerized package, allowing service providers to ensure access to applications for legitimate users even while under attack.

This is a true on-box solution. The controller simply gives operational control to intervene in an attack and to manage configuration and deployment of the containers. All detection processing and blocking signature creation is done on-box using a containerized DDoS detection engine optimized for the access edge. Only a minimal dataset containing the attack profile and blocking rules are passed to the controller for operator validation/intervention. Once approved, the blocking rule is applied to the router directly for mitigation.

This is a radical departure from how DDoS has been managed in the past. NetFlow no longer needs to be passed to a massive, centralized processing infrastructure. Attack traffic can be blocked in-place and no longer needs to be backhauled to a dedicated scrubbing infrastructure. As we will demonstrate, this is a new approach to DDoS protection, designed to fully meet 5G's hyperscale and latency requirements while also being a lightweight, effective solution.

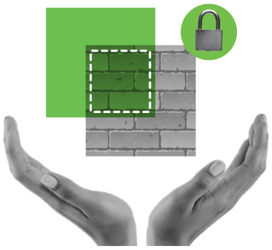

1.2 Trends in 5G service provider networks

Service provider networks continue to evolve to support the requirements of advanced, low-latency 5G applications. Protecting the network and applications must evolve as fast or faster than the evolution of the network itself. Ensuring the security and availability of the network and applications should be a natural outcome of the network design, automation, network operations, and security operations. This paper details Cisco Secure DDoS Edge Protection and its ability to deliver hyperscale infrastructure protection with the ability to match pace and scale of even the most aggressive 5G build-outs.

The development of 5G standalone networks and widely distributed architectures driven by cloud-native fabrics and edge computing are driving network transformation. As networks and applications continue to evolve, security must also evolve to ensure the security of modern networks and the availability of applications. We need to deliver modern outcomes while supporting migrations from existing networks and services used to protect them. A plan, then, is required to guarantee protection of the existing network while accommodating for the "new edges."

The challenges that service providers face driven by service and network evolution include, but are not limited to:

- Faster data speeds
- More devices on the network
- Enable low-latency apps, like interactive video

- Better performance
- Improved QoE
- Increased customer satisfaction

- Service growth
- Revenue growth
- Faster ROI on network investments

Figure 1. The benefits of 5G driving network transformation

DDoS protection is not a new or otherwise described "greenfield" service. However, as throughput demands increase based on the appetite of video-hungry users, IoT-based applications, V2X (vehicle to anything) or connected cars, virtual reality users, enterprise and commercial applications, military and defense applications, and game players, there must be a way to avoid security degradation due to limitations and cost. New applications and use cases are born every day, making securing those use cases a challenge. As Cisco provides industry-leading solutions to address customer challenges, we also need to account for continuity between legacy tools and modernizations. A quick explanation of existing architectures follows with addition of a critical detail: DDoS protection is "evolving" with the network and service evolution and therefore must seamlessly integrate into these existing architectures and APIs.

1.3 Current architectures

The following architectures are used by many customers today, and the new “edge” protection adds efficiency by offloading service providers core resources and, through this innovation, extends protection to access designs, previously cost prohibitive. Let’s examine the current architectures to see if they can effectively take care of this “new” protection. The answer is that current architecture, as detailed below, will have a role, but the Edge Protection solution is required to complete the protection required.

The Edge Protection solution allows us to add efficiency and automation to the current state of the art pictured below.

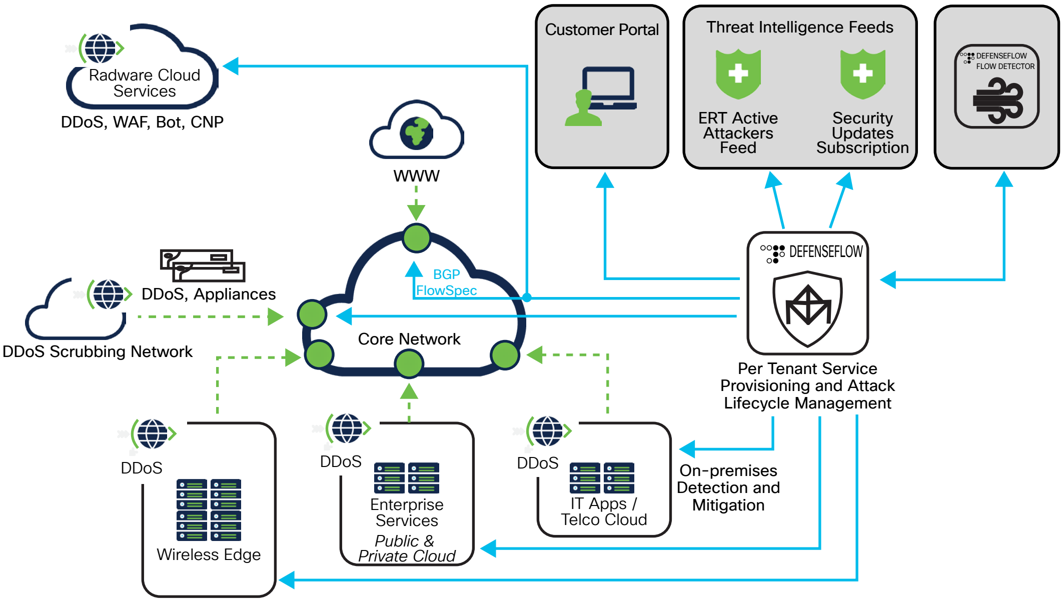

Figure 2. Network service design overview - multitenant network service infrastructure

The three DDoS architectures deployed today are described below:

- Scrubbing centers: Used for infrastructure protection of network services and used for MSSP services. This solution relies on NetFlow monitoring of distributed routers and a centralized detection engine identifying gross anomalies while steering attacks toward scrubbing centers for SOC-assisted mitigations. As core networks flatten by design, it becomes increasingly important to address events at the edge of the network to avoid latency insertion and removing the need to backhaul attacks. Cisco is adding more efficient/scalable capabilities by adding the “edge” protection to optimize quality of experience for customers.
- “Second line of defense”: Used for computing platform service assurance and high availability for network services, gateways, applications, and APIs. Defense in depth is one of the key architectural tenets applied in the “second line of defense” DDoS protection services today, providing real-time protection for key services such as DNS, VoIP/IMS, portals, VPN concentrators, and APIs. A centralized controller called DefenseFlow® enables the integration of the layers of defense into a network-wide system.

- Cloud DDoS Peak Protection: Service providers have varying levels of human resources to manage security infrastructure. Cloud DDoS Peak Protection can alleviate the burden of maintaining on-premises systems while partnering with Cisco to provide global coverage via cloud services. Service providers can combine any layer of protection offered to tailor protections addressing requirements in the most effective way possible.

1.4 New service provider network edges

What are the new "edges" of today's service provider networks? In 5G networks, components of the application and network are pushed out to the edge to account for requirements of "ultra-low-latency" applications. There are numerous examples of where Edge Protection compliments a defense-in-depth architecture. Protecting services at the converged access edge or the MEC edge (mobile edge compute), service aggregation edges, the new peering edge, the data center edge, and the application edge all require new agility and a more robust, more efficient method than what is in operation today.

We'll focus throughout the remainder of this white paper on addressing a major challenge in network security associated with the fast-paced progression of converged access designs. More specifically, what is the problem on the new "mobile edge" that needs to be solved? Let's start by reviewing the underlying challenges.

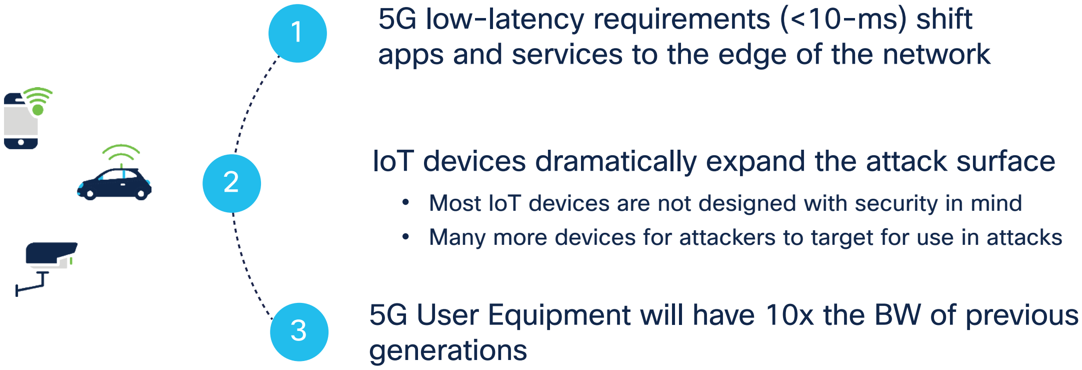

1 5G low-latency requirements (<10-ms) shift apps and services to the edge of the network

2 IoT devices dramatically expand the attack surface
- Most IoT devices are not designed with security in mind
- Many more devices for attackers to target for use in attacks

3 5G User Equipment will have 10x the BW of previous generations

Ultra-low-latency application SLAs are forcing service providers to reconsider how network access security is implemented. Let's review the benefits and challenges addressed by integrating DDoS services into the "mobile edge" on the cell site router.

- Enterprise applications such as interactive video, massive IoT, tactile internet, smart cities, and virtual reality are being migrated to edge computing infrastructure to improve QoE for users.
- Ultra-low-latency applications will not tolerate latency inserted by redirecting suspicious flows to scrubbing infrastructure.
- Cisco Visual Network Index predicts that by 2023, there will be approximately 30 billion devices/connections.
- 5G terminals' average broadband speeds will rise from 110 Mbps to 1 Gbps, increasing the disruption from attacks.
- Wireless operators must ensure network and application availability and target edge resources to provide tailored security service to end customers.
- The cell site router would be the first line of defense to guarantee the QoE of the network.

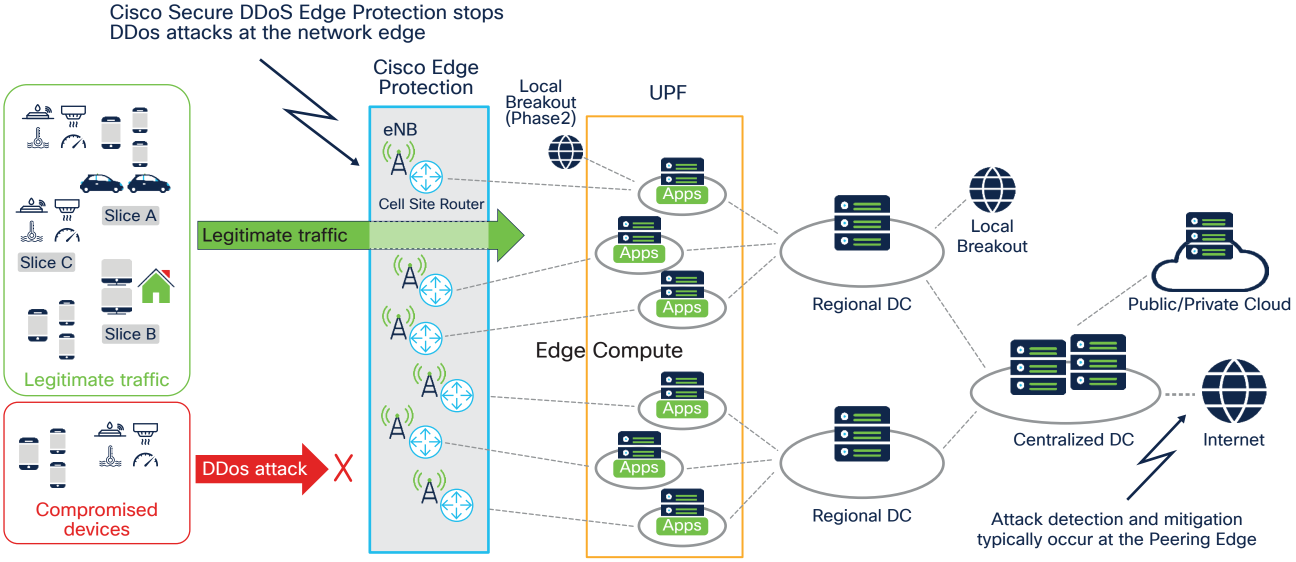

Figure 3. DDoS attack protection at the 5G network edge

The hyperscale evolution of the network and services requires solutions that are innovative and different. Cisco Secure DDoS Edge Protection is just that, a highly optimized network integrated security application. Here's what Cisco Secure DDoS Edge Protection delivers that is new and different:

- Combining the strength of network and security resources to see more than we could see in the past and take faster actions driving efficiency of network resources
- Leverage IOS® XR ability to deliver telemetry 5 times more efficiently to aggregate telemetry from a far greater number of edge devices
- Leveraging our behavioral DDoS methodology and adding auto-tunable access and peering-specific optimizations via machine learning, we are uniquely positioned to meet the efficacy requirements of today and tomorrow's network and service demands
- Leverage IOS XR ability to host an intelligent container application to add a new layer of localized detection and mitigation enforcement
- Agility to adaptively mitigate attacks leveraging IOS XR routers to maintain line rate performance as the network continues to scale
- Leverage a "centralized" controller to correlate aggregated telemetry from all of the "edges" delivering network-wide visibility, detector fleet control, and attack lifecycle management.
- Deliver a solution capable of leveraging new and existing capabilities to tailor efficient designs to customers

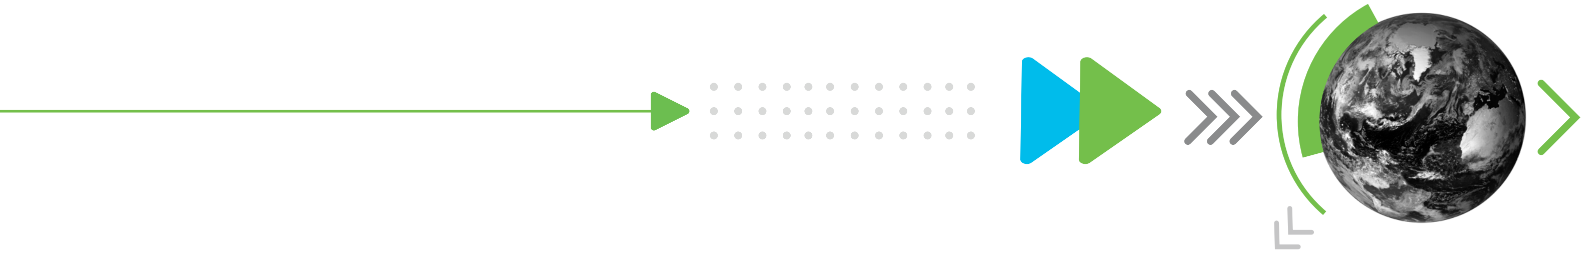

1.5 Architecture of SP network protection now adding "the edge" - Cisco Secure DDoS Edge Protection

Protection of hyperscale networks, the hyperscale infrastructure, and the services that will run on top of them require a disaggregated, distributed defense system that previously didn't exist. Let's take a quick look at how the solution works in the mobile edge deployment. First, a quick overview.

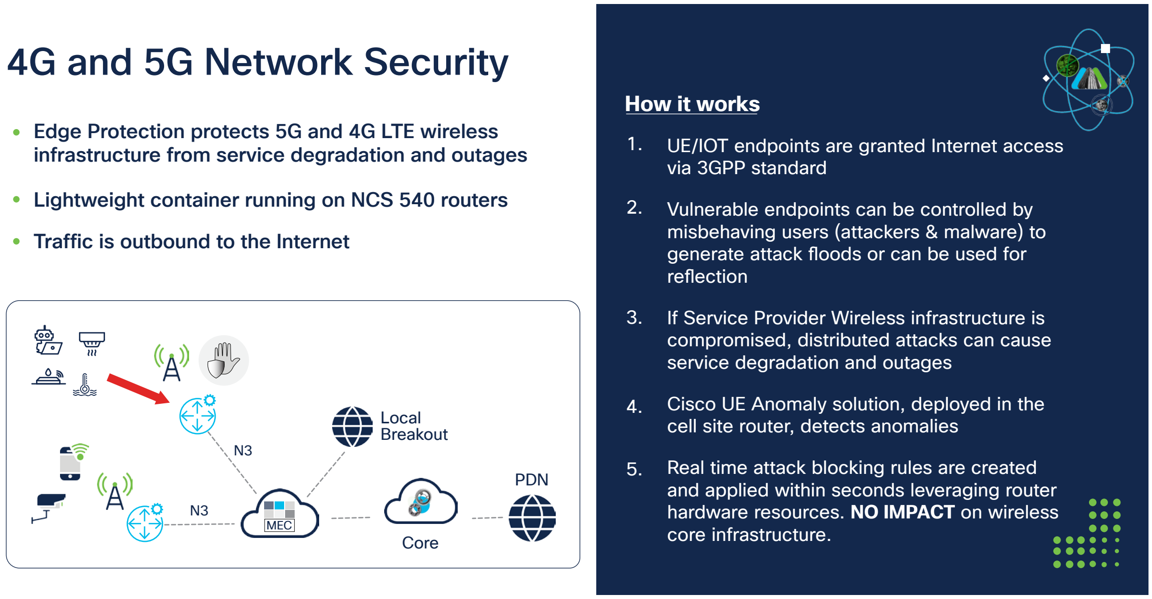

4G and 5G Network Security

- Edge Protection protects 5G and 4G LTE wireless infrastructure from service degradation and outages
- Lightweight container running on NCS 540 routers
- Traffic is outbound to the Internet

How it works

1. UE/IOT endpoints are granted Internet access via 3GPP standard
2. Vulnerable endpoints can be controlled by misbehaving users (attackers & malware) to generate attack floods or can be used for reflection
3. If Service Provider Wireless infrastructure is compromised, distributed attacks can cause service degradation and outages
4. Cisco UE Anomaly solution, deployed in the cell site router, detects anomalies
5. Real time attack blocking rules are created and applied within seconds leveraging router hardware resources. **NO IMPACT** on wireless core infrastructure.

Figure 4. 4G and 5G network security

2. Network integrated telemetry, attack detection, and automated remediation

The protection needed for the new "edge" requires increased agility and efficacy, resulting in finding the threat faster and more comprehensively and fixing the problem faster, dynamically, and network wide.

2.1 Network visibility and modern DDoS protection

When we examine the requirements driven by convergence of the mobile and broadband edge, using the network as a sensor and mitigator is the only way to scale protection with the requirements of the critical applications and workloads now running "at the edge." Using the network as a sensor is implemented with critical security capabilities including machine learning integrated into network devices at "the edge." In the case of the mobile edge, the Cisco NCS 540 with the Cisco Secure DDoS Edge Protection service container running inside of the IOS XR network operating system is one such implementation. What follows below is a summary and then a deeper dive into how this implementation works.

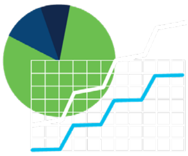

Telemetry, attack detection, and mitigation

Telemetry provides critical visibility into network operations and allows operators to optimize performance and minimize security risks, but analyzing metadata normally comes at the expense of computing resources and latency. IOS XR provides greater efficiency by exporting traffic attributes specifically required for anomaly detection. The efficiency of transforming telemetry to the containerized security application allows for long-term investment protection by offering highly scalable telemetry consumption.

Secure Edge Protection allows administrators flexibility using configuration templates designed for different use cases according to traffic profiles. This flexibility allows Cisco Secure DDoS Edge Protection to efficiently detect and sign attacks in real time while applying adaptive remediation policy to the router and leveraging line rate hardware-assisted mitigations via a TCAM-based access control list.

Leveraging extended telemetry features and an expanded data model, Cisco Secure DDoS Edge Protection employs algorithms from the proven market leader in DDoS solutions, Radware. With over two decades of production success in automating DDoS defense, Cisco customers benefit from disaggregated, distributed defense delivery where a lightweight, efficient containerized machine learning application is deployed on IOS XR platforms out of path and not impacting data plane processing, even under failure. With this approach, customers appreciate the best of both worlds, on-line mitigation and out-of-path sampling and the hyperscale performance of leveraging the network as a line rate enforcer. The solution is managed by a centralized controller designed to scale to the edge of the largest networks in the world, while automating container lifecycle management, centralized configuration, automated attack lifecycle, and reporting. This application is lowering time to mitigation, significantly reducing the need to carry attack traffic across core resources, optimizing existing security infrastructure, and helping security analysts to achieve additional focus for their most sophisticated events.

2.2 Cisco Secure DDoS Edge Protection – solution components and high level overview

The Edge Protection solution is implemented in two components:

- **A series of detectors** – Containers that run inside of the Cisco IOS XR NOS, one on each “mobile edge” router, performing local detection and mitigation “on-box”
- **A controller** – A service that manages the distributed network of edge detectors, analyzes trends across the network, and orchestrates cross-network visibility and mitigation. The controller also delivers a full system management lifecycle for the entire service. API integration is also delivered by the controller.

In the Architecture section earlier in this document, the “mobile edge” component of the solution is further described in the graphic below:

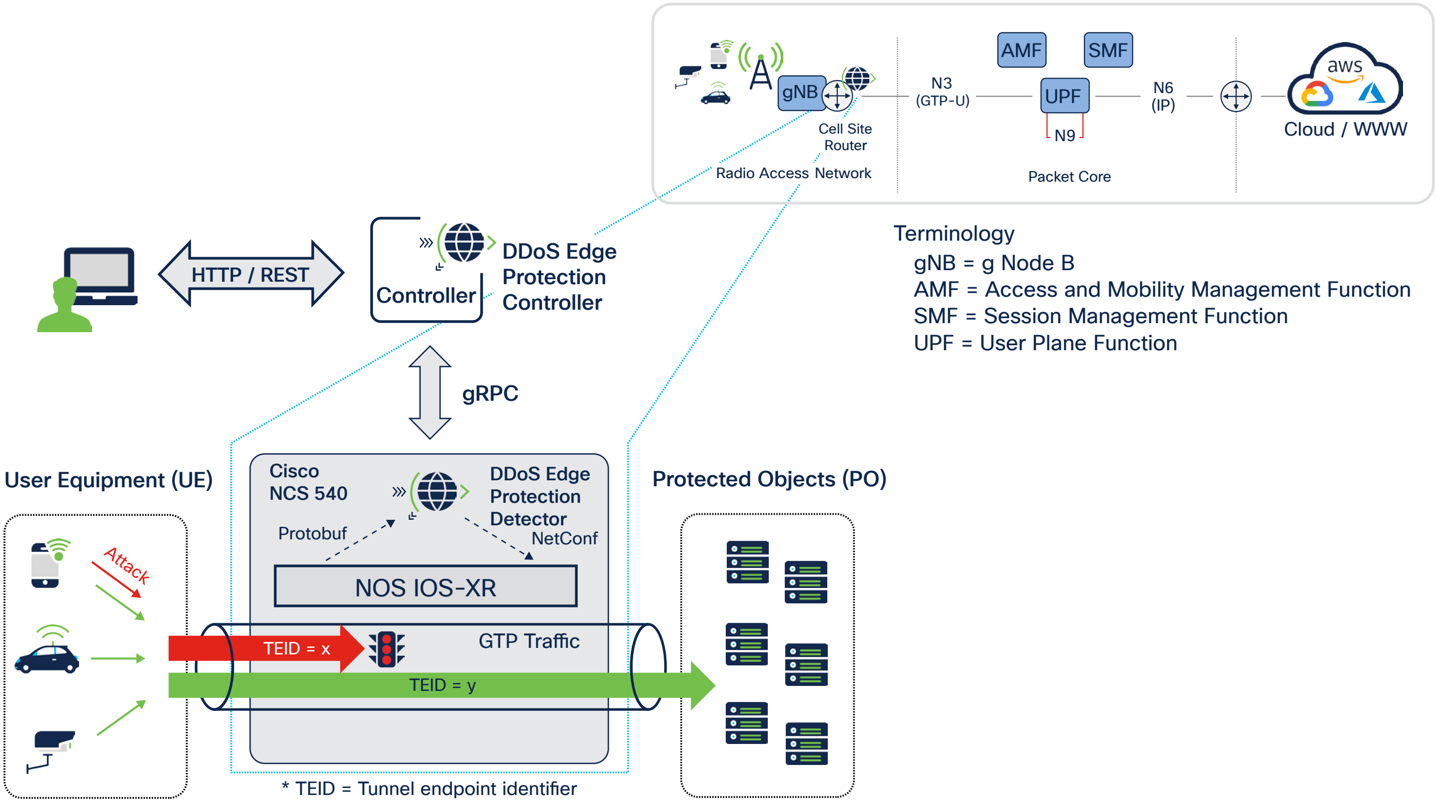

Terminology
gNB = g Node B
AMF = Access and Mobility Management Function
SMF = Session Management Function
UPF = User Plane Function

Figure 5. Cisco Secure DDoS Edge Protection – solution components and traffic flow

Detector – The localized DDoS detection function is implemented in a virtual microservice container running on spare management vCPU resources, avoiding any negative impact on the functionality of the router or traffic flow through the router. The detector performs DDoS attack analyses based on the traffic flow information received from the router's hardware. The information is passed to localized detection algorithms to generate alert indications when a DDoS attack is identified. The containerized application employs patented algorithms for characterizing cyberattacks within tunneled traffic (GTP), which is common in the mobile 5G edge, and provides enhanced visibility, enabling agility and comprehensive protection against attacks at the edge at scale. Localized telemetry processing allows for better sensitivity of distributed anomaly awareness and consumes far fewer computing resources given an optimized system interaction with the centralized controller. There is an additional win as it applies to minimizing energy consumption at the edge. There's less deployment hassle, less of a carbon footprint, and a significant step toward a "greener" internet.

This solution can be applied to existing and new routers with no additional hardware required. The detector can identify the attacking devices' IDs (TEID) and instruct the router to block them directly. This is unique, allowing very efficient mitigation of the offending traffic only while letting the legitimate traffic pass through the router without any hindrance or any degradation. The controller, using information from the container, dynamically characterizes the attack and finds its real-time signature, thus enabling a very efficient mitigation. The detector adapts to evolving threats continually during an attack to see whether attack vectors change and updates the mitigation accordingly.

Controller - A central interface management function for operator admin users to manage a fleet of detectors. The controller performs the actions listed below:

- Manages container lifecycle for a fleet of detectors (up to 50,000 per controller)
- Configures and edits detector profiles and security settings
- Checks the health of detectors
- Displays information about real-time attack forensics and threat intelligence analyses
- Mitigates DDOS attacks on the network
- Provides real-time and historical reporting of events
- Provides operational control/incident response – the operator can choose to start mitigation, clear the attack, and do nothing, or the operator can choose to let the system automatically mitigate a protected object

3. Addressing the quality of experience challenge for converged access

3.1 How To protect hyperscale networks – the mobile edge use case

5G is about use cases and the user experience. For example, delivery of services like IoT services, connected cars, virtual reality and many others are so exacting on latency minimalization and control that a natural evolution toward a hyperscale network to deliver the “edge” is a necessity. Quite literally, portions of the network such as the UPF or User Plane Function are moved out to the edge delivering use case-based outcomes while maximizing the user experience.

Increases in bandwidth and capabilities of the UE/IoT (mixed media endpoint) have led to compromises and vulnerabilities, bringing about the need to have an EDR-like (endpoint detection and response) service at the 5G edge for mobile UEs and delivering this on the cell site router to detect and mitigate attacks – thus removing the need to backhaul traffic for scrubbing unnecessarily and achieving the agility required to protect the 5G mobile “edge” at scale. Anomaly awareness is critically important to protect “the edge.” In fact, there are quite a few “edges” where this applies. Here the focus is the mobile “edge.”

Let’s now look at how to protect the “mobile edge.” First a little background and then a look into the methodology and tools required.

5G has not only brought hyperscale network evolution and a widely distributed network to serve a distributed set of network functions and services to deliver the ultra-low-latency-based services described above. It has also brought with it a set of new threats. These threats are referred to as the “5G threat surface.” In this paper, DDoS threats at the mobile “edge” are covered.

Today, in a service provider, state-of-the-art DDoS protection systems focus on ingress attacks. Traditional DDoS systems have evolved over the years as have the people, processes, and related tools to handle ingress DDoS attacks coming from the internet toward the service provider’s infrastructure. However, 5G introduces the egress DDoS attacks, highlighted in the graphics below. This new threat vector comes from highly capable UE (User Equipment) and IoT devices that are orders of magnitude higher in cores, memory, and most importantly bandwidth. 5G promises 10x+ the bandwidth of 4G per device, and as such the threat that a series of compromised UE devices can impact services is perhaps the biggest threat to 5G operators today. There is a need to see and mitigate these “egress” DDoS attacks as close to the "edge" as possible. With Edge Protection, we deploy in the cell site router (Cisco NCS 540) in a container on top of the Cisco network operating system, IOS XR, achieving agile visibility and mitigation at the "edge" and communicating back to a centralized controller for network-wide decision making and integration into existing DDoS systems already deployed.

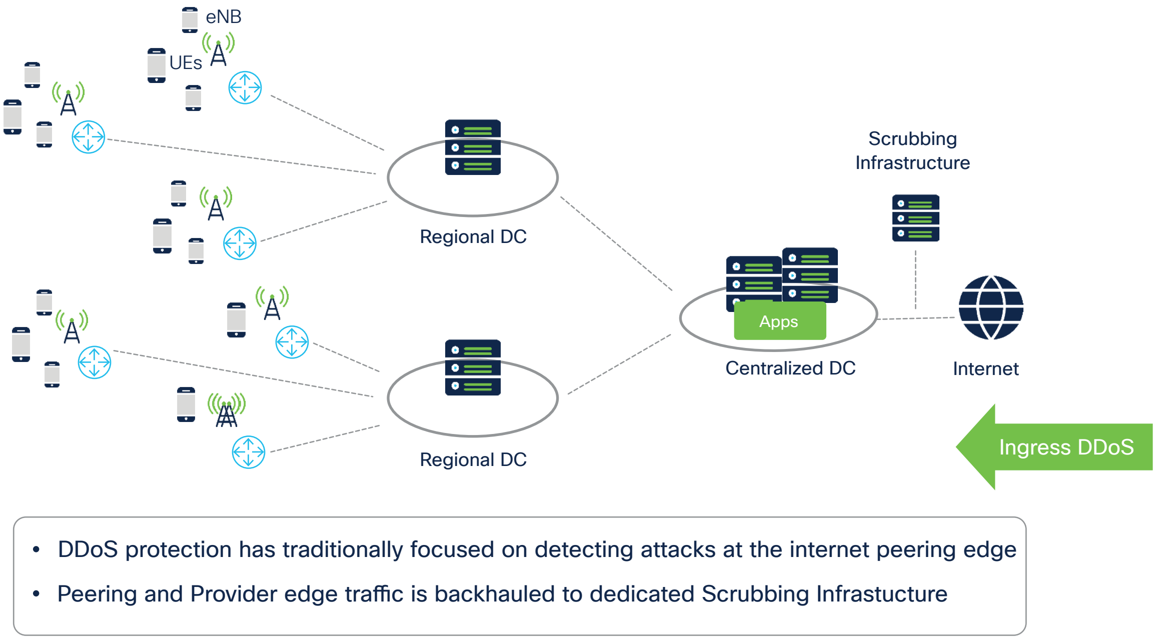

- DDoS protection has traditionally focused on detecting attacks at the internet peering edge
- Peering and Provider edge traffic is backhauled to dedicated Scrubbing Infrastucture

Figure 6. Traditional DDoS security

The newer threat vector for DDoS brought about by 5G is shown as "egress" in the graphic below.

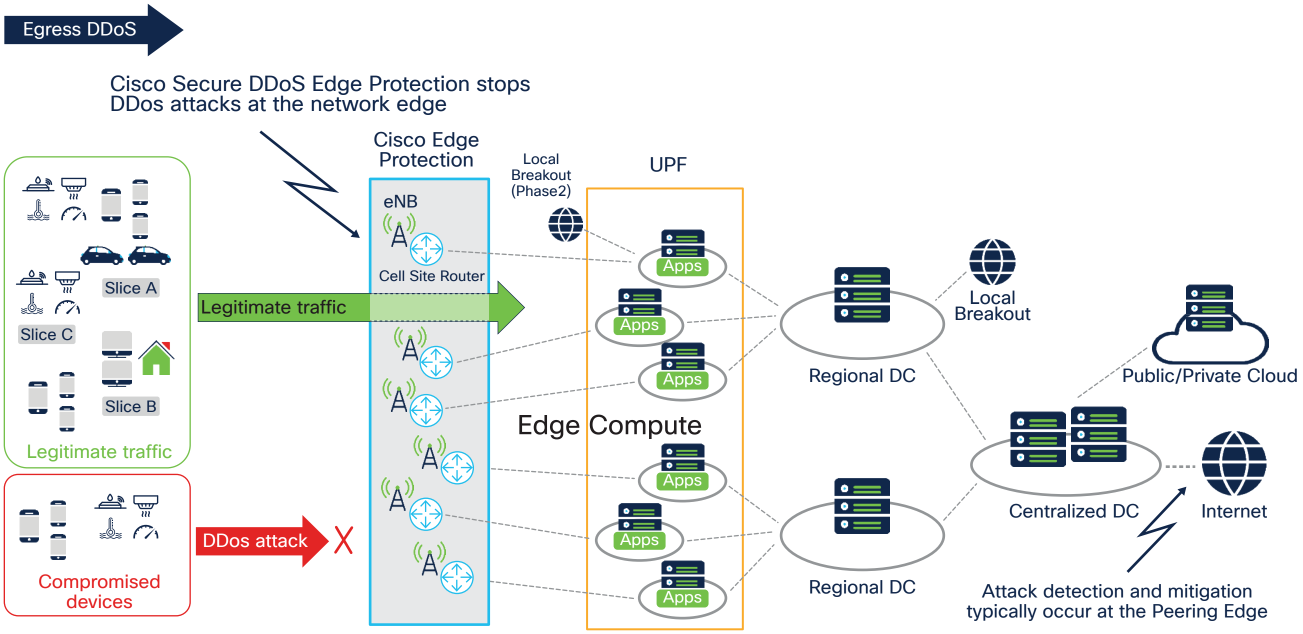

Figure 7. DDoS attack protection at the 5G network edge

The advantage of solving the DDoS attack protection at the 5G edge with the Cisco Secure DDoS Edge Protection solution includes:

- No additional hardware and equipment
- Line speed detection using less power, less energy, less carbon footprints; a step in the right direction for green internet
- A form factor that's cloud native
- Implemented "out of path" - not in the critical path for traffic flow through the router

What follows describes exactly how to achieve this outcome. Let's start with an examination of the relevant 5G and mobile architecture components and technologies.

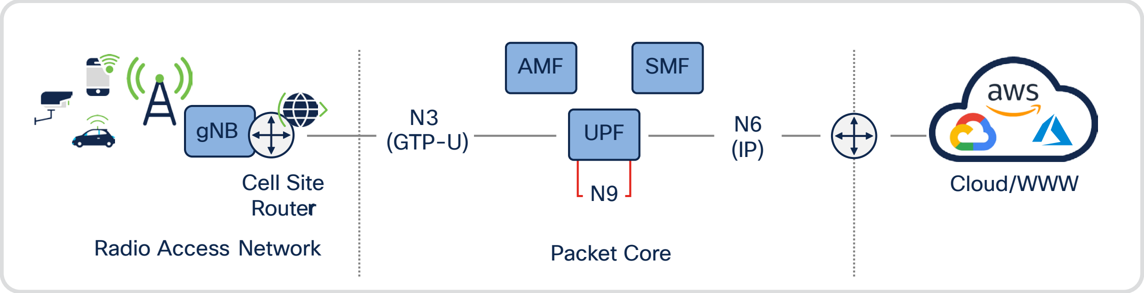

Terminology
gNB = g Node B
AMF = Access and Mobility Management Function
SMF = Session Management Function
UPF = User Plane Function

Figure 8. 5G mobile architecture components overview

When we examine the components of the mobile service provider, there are two aspects to focus on. First is the **access "edge"** of the 5G network. Each device (UE) will get an IP address from the service provider network and will communicate with the cell site tower. At the cell site tower, the cell site router will communicate with the second aspect, the **5G packet core**. The traffic from the mobile is encapsulated into a tunnel using the GPRS Tunneling Protocol or GTP for short. The traffic for GTP is broken up into two parts: GTP-U (for GTP on the User Plane), which is where the data is sent on its way out to the internet (or other destination, and GTP-C (for GTP on the Control Plane). One of the new innovations for 5G is that there are many breakout points for the traffic and those breakout points are often deployed very close to the edge to achieve the goal of ultra-low latency and optimal user experience. The traffic may be destined for an application deployed in a small edge data center known as a MEC or mobile edge compute node in the hyperscale infrastructure. Therefore, when there is an attack, it could go all the way out to the internet but could also go and impact an application in the MEC or even an application in the packet core itself. These attacks often go undetected as it takes a lot to monitor for them as they are encapsulated in GTP-U and require a deep analysis to monitor this traffic and detect threats. The current state-of-the-art DDoS systems can only detect such malicious activity once the traffic is decapsulated from the GTP-U tunnel, and by this time, it's too late. The issue is that this traffic is providing undue load and stress on all the components of the infrastructure that it passes through. By placing the control (security enforcement) at the mobile "edge," we solve this issue. As this is a new threat vector for 5G, this must change, and Cisco Secure DDoS Edge Protection incrementally adds this capability to the existing DDoS system.

To deliver ultra-low-latency outcomes for 5G, as described earlier, components of the packet core are “pushed out” to the edge, many times in the MEC. vUPF and vBNG are good examples of such functions. 5G also provides many different breakout points. Therefore, placement of a control to solve this problem as close to the “edge” as possible, at the cell site router in the mobile edge, makes sense. Now, we can see and mitigate the issue right away, gaining agility, efficiency, and protection of critical applications and services using Cisco Secure DDoS Edge Protection. The best practice for DDoS says you want to see and mitigate the attack as close to the source as possible to minimize the collateral damage. By seeing and mitigating the attacks on the cell site router, we do just that.

Now, we look at how the solution is implemented into the cell site router leveraging the Cisco network operating system, IOS XR.

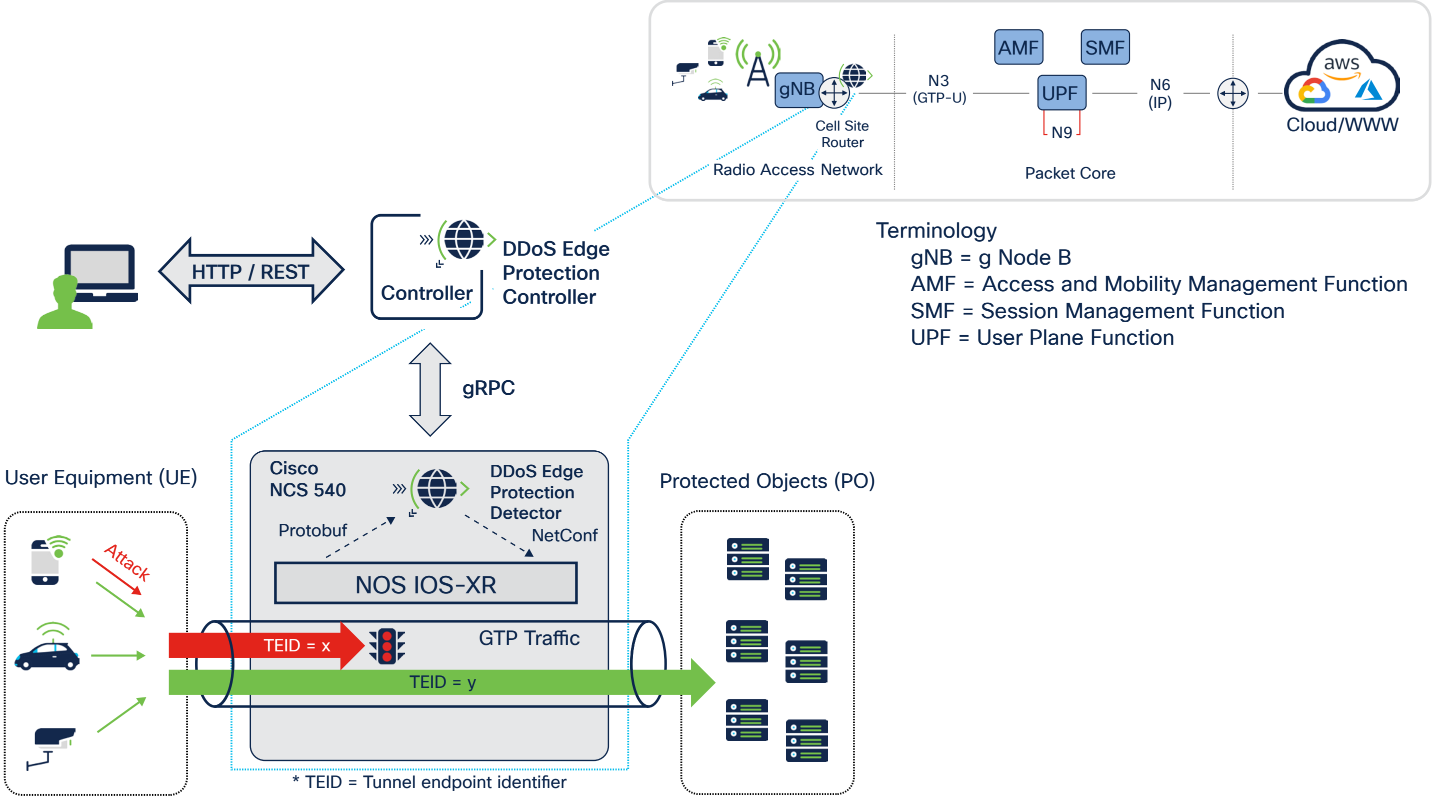

Figure 9. Cisco Secure DDoS Edge Protection – solution components and traffic flow

Starting from left to right, we have the UE devices and IoT sensors. The cell site router will route to the rest of the elements of the packet core. The assumption is that the UEs are usually sending legitimate, nonattack traffic as indicated by the green arrow in the graphic above, while certain UE devices will be sending attack or anomalous traffic as indicated by the red arrow in the generated graphic above. The traffic itself, albeit encapsulated in GTP, will be routed on the various interfaces of the packet core. From the traffic generated by the mobile UEs, telemetry informs the container that is deployed on to the cell site router and lives in the IOS XR application ecosystem.

Leveraging Protobuf technology, telemetry is collected as the first step in the process. The telemetry is collected from the appropriate interface(s) on the router itself then analyzed by the Cisco Secure DDoS Edge Protection solution to determine whether there is an attack or not. The container is not in line with the traffic. We do not need to check packet by packet, but instead samples are collected, and the behavior is examined to determine malicious traffic or the presence of an attack. Next, a Netconf interface is used to enable routers to block the offending UEs by Tunnel Endpoint Identifier (TEID) using a TCAM-based access control list (ACL). An example of an ACL blocking by GTP TEID is shown in the graphic below.

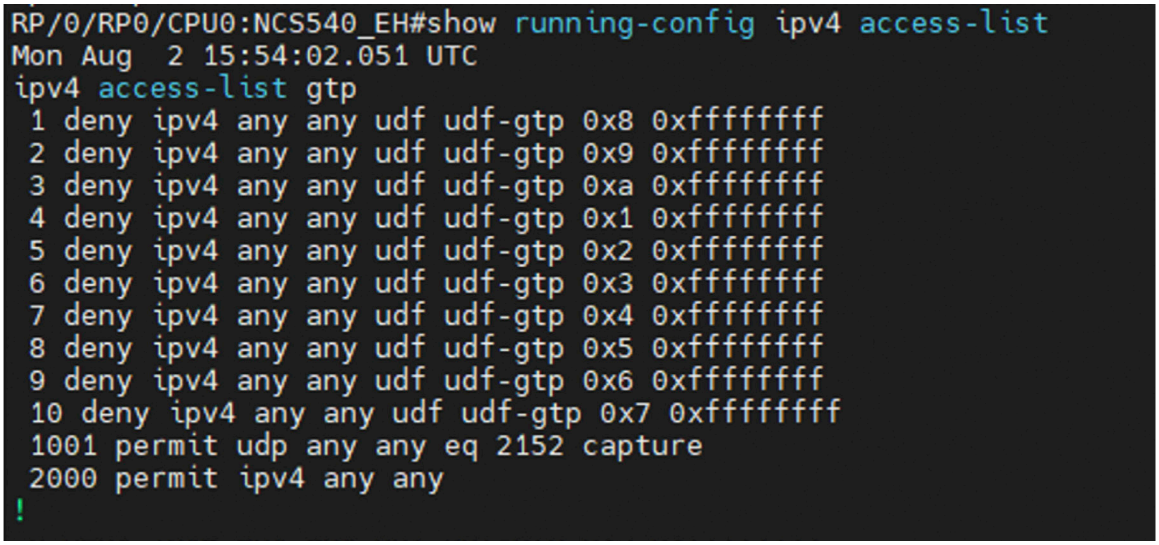

```
RP/0/RP0/CPU0:NCS540_EH#show running-config ipv4 access-list
Mon Aug  2 15:54:02.051 UTC
ipv4 access-list gtp
 1 deny ipv4 any any udf udf-gtp 0x8 0xffffffff
 2 deny ipv4 any any udf udf-gtp 0x9 0xffffffff
 3 deny ipv4 any any udf udf-gtp 0xa 0xffffffff
 4 deny ipv4 any any udf udf-gtp 0x1 0xffffffff
 5 deny ipv4 any any udf udf-gtp 0x2 0xffffffff
 6 deny ipv4 any any udf udf-gtp 0x3 0xffffffff
 7 deny ipv4 any any udf udf-gtp 0x4 0xffffffff
 8 deny ipv4 any any udf udf-gtp 0x5 0xffffffff
 9 deny ipv4 any any udf udf-gtp 0x6 0xffffffff
 10 deny ipv4 any any udf udf-gtp 0x7 0xffffffff
 1001 permit udp any any eq 2152 capture
 2000 permit ipv4 any any
!
```

Figure 10. DDoS mitigation with ACL on the router

One of the aspects of the Cisco IOS XR network operating system that we utilize here is the User-Defined Field or UDF. As you can see in the access control list above, the solution takes advantage of the UDF to add the ability to utilize TEID (tunnel endpoint identifier leveraged in the udf-gtp field) in the access list, giving the solution the ability to not only see the traffic, but also block in with the ability to see the GTP inner and outer parts.

Another important interface from the router is communication with the controller. The responsibility of the controller is to manage many routers, in some cases in the 10s to 100s of thousands of routers across the 5G infrastructure. Each of the routers (with its container) communicates with the controller. When there is an attack, the controller issues the command to block that attack on one or many routers, therefore fully protecting that "edge" and only blocking the specific (malicious) traffic in question. This makes sure that only attack traffic from that specific mobile(s) is blocked and not any legitimate traffic.

3.2 Defending against the Mirai botnet at the network edge

Now, let's examine how the Cisco Secure DDoS Edge Protection works against a Mirai botnet, a well-known DDoS attack. An illustration of how the Mirai botnet works follows below.

The Mirai botnet operates when a bot master activates many mobile UEs to attack. The Cisco Secure DDoS Edge Protection solution will detect the attack and mitigate the attack traffic only while allowing the legitimate, nonattack traffic to continue to its destination, performing all of this on multiple mobile edge nodes across the 5G network edge.

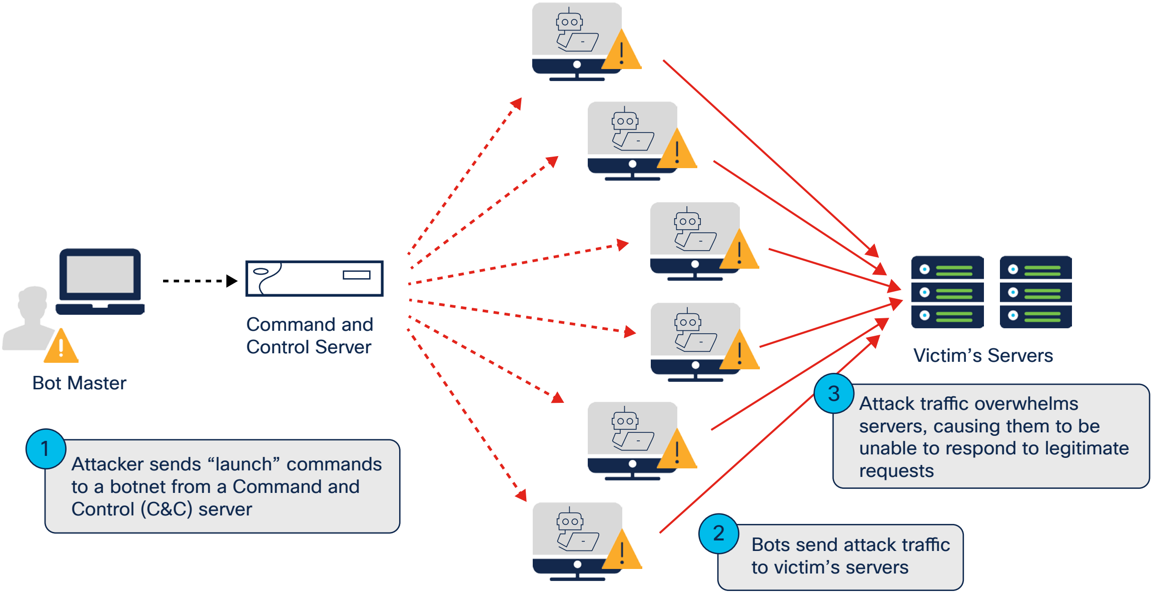
Bot Master
Command and Control Server
Victim's Servers
1
Attacker sends "launch" commands to a botnet from a Command and Control (C&C) server
2
Bots send attack traffic to victim's servers
3
Attack traffic overwhelms servers, causing them to be unable to respond to legitimate requests

Figure 11. The Mirai botnet

When we log in to the controller, we see a status map and a breakdown of detectors/containers by severity.

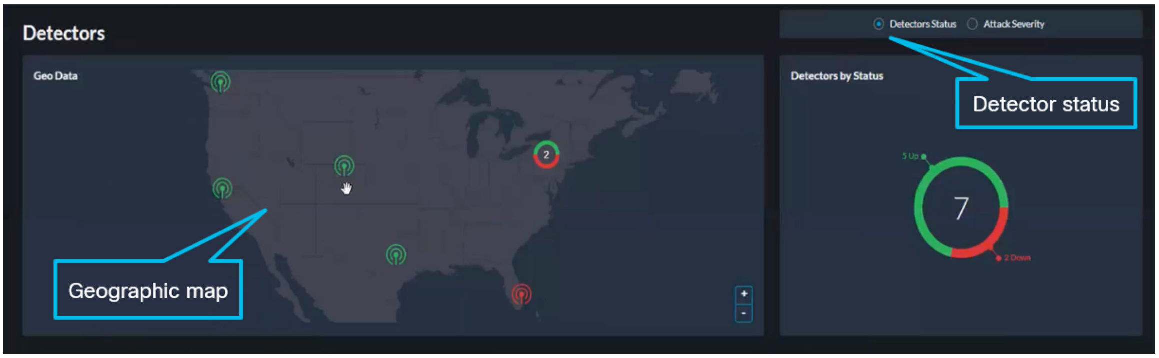

Figure 12. Geographic detector map and status

We can then drill down into each of the severity-specific icons to learn more about the attacks and act. Both manual and automated attack mitigation are offered as operational preference across service providers varies. Geolocation is also offered for each router allowing for groups of protected "edges" and is customizable to the liking of the service provider operator.

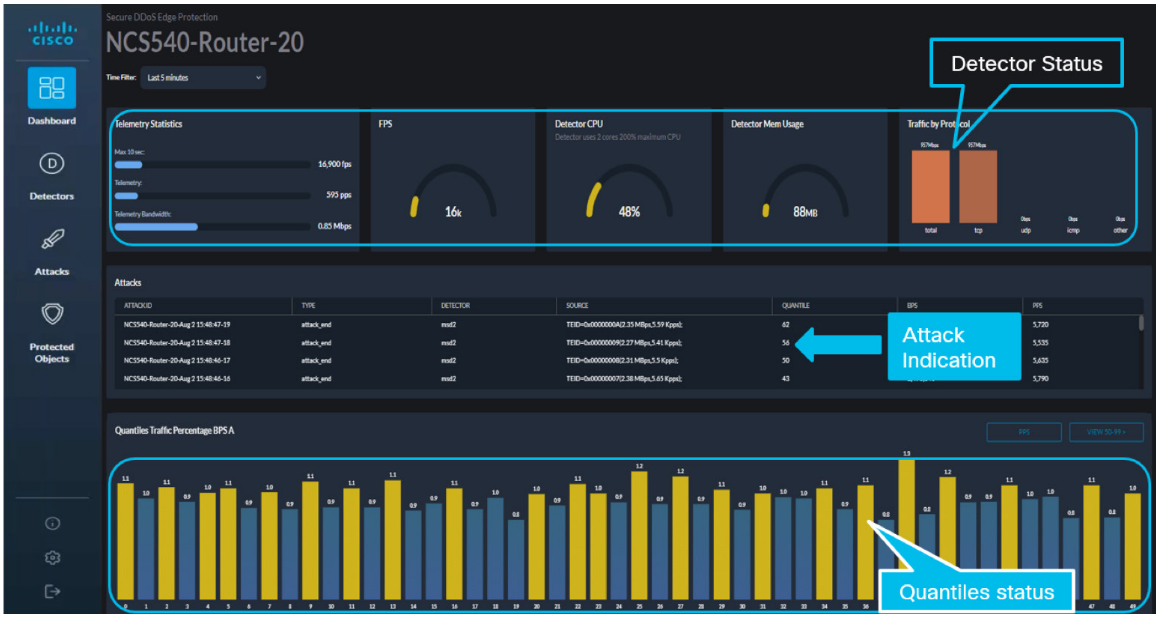

ATTACK ID	TYPE	DETECTOR	SOURCE	QUANTILE	BPS	PPS
NCS540-Router-20-Aug 2 15:48:47-19	attack_end	msd2	TEID=0x0000000A(2.35 MBps,5.59 Kpps);	62	[illegible]	5,720
NCS540-Router-20-Aug 2 15:48:47-18	attack_end	msd2	TEID=0x00000009(2.27 MBps,5.41 Kpps);	56	[illegible]	5,535
NCS540-Router-20-Aug 2 15:48:46-17	attack_end	msd2	TEID=0x00000008(2.31 MBps,5.5 Kpps);	50	[illegible]	5,635
NCS540-Router-20-Aug 2 15:48:46-16	attack_end	msd2	TEID=0x00000007(2.38 MBps,5.65 Kpps);	43	[illegible]	5,790

Figure 13. Controller graphical user interface – detector status/health, attacks, and quantiles

On the NCS 540, the detector/container is allocated two cores to operate and easily meets this constraint as tested across packet sizes and attack types.

As we scroll down the page, we find an analysis of the legitimate traffic through the container and then we will see indication of an attack differentiated from the legitimate traffic. First, the legitimate traffic is graphically represented by “Total Traffic BPS,” “Total Traffic PPS,” and “Traffic by Protocol.”

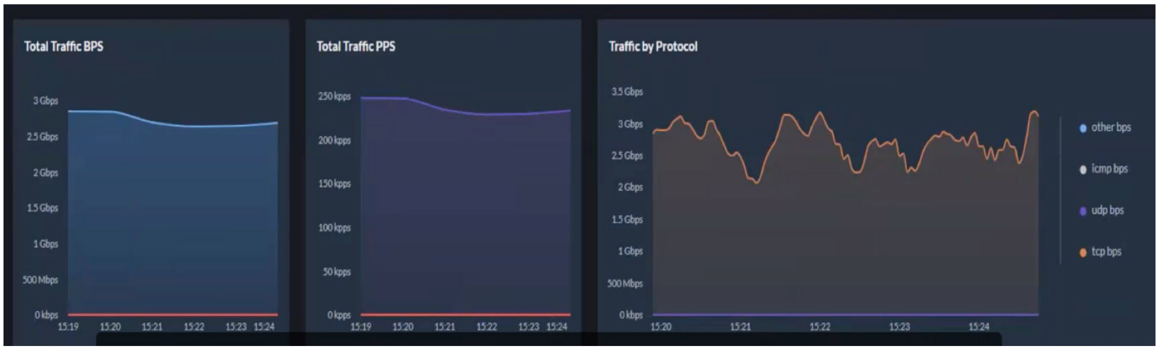

Figure 14. Analysis graphs – legitimate traffic

In the area of the screen in the graphic above titled “Controller GUI – Main Screen,” showing the bar-style histogram titled “Quantiles Traffic Percentage BPS A,” you see a visual representation of the behavioral algorithm applied. You might be wondering, “How does it work?” Let’s dig a little deeper to find out.

First, we distribute the overall traffic into quantiles (bar in the visualization above). Each quantile represents about 1% of the overall traffic. We then take all the mobile UEs that are sending traffic through the router and spread traffic from them out equally across the quantiles. Each mobile UE is represented by its tunnel endpoint identifier (TEID). For example, if we had 2000 total mobile UEs sending traffic through a specific router, then each quantile would have about 20 mobile UEs each. When a new mobile UE comes into a quantile, detection of an attack follows easily as the behavior of that mobile UE is compared to the rest of the UEs already in that quantile. In other words, when a mobile UE is attacking or has attack traffic, it will behave significantly differently than its peers within the quantile, forcing that attacker to stand out "in the crowd." From this point on, the solution will track the tunnel endpoint identifier (TEID) and instruct the router to block using the mitigation TCAM-based access control list shown above, by TEID. Otherwise stated, blocking is done by filtering on the router interface using the access control list.

Now we will move from how the system looks and operates at peace time to how it looks and operates when the attack(s) happens. The attack is a Mirai botnet with the attack mobile communicating with the command and control (C2), and after that, the botmaster will call for the start of the attack. The attack is a UDP-based "flood attack," generating a lot of attack traffic. You will see a few things that have changed in the graphics below. First, you see the red bars or quantiles indicating an attack is happening.

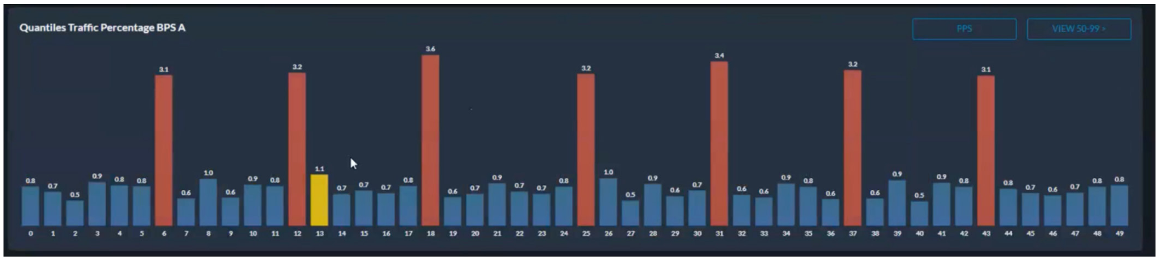

Figure 15. Quantiles – visualization of an attack

We have a total of 10 attackers (you see only 7 of them in the graphic above) indicated by the red "outliers" you see in the graphic above. We also see graphic indication of the attack traffic in the graphic below, and we see the total attack traffic compared to the total legitimate traffic, just below the quantile visualization for that attack.

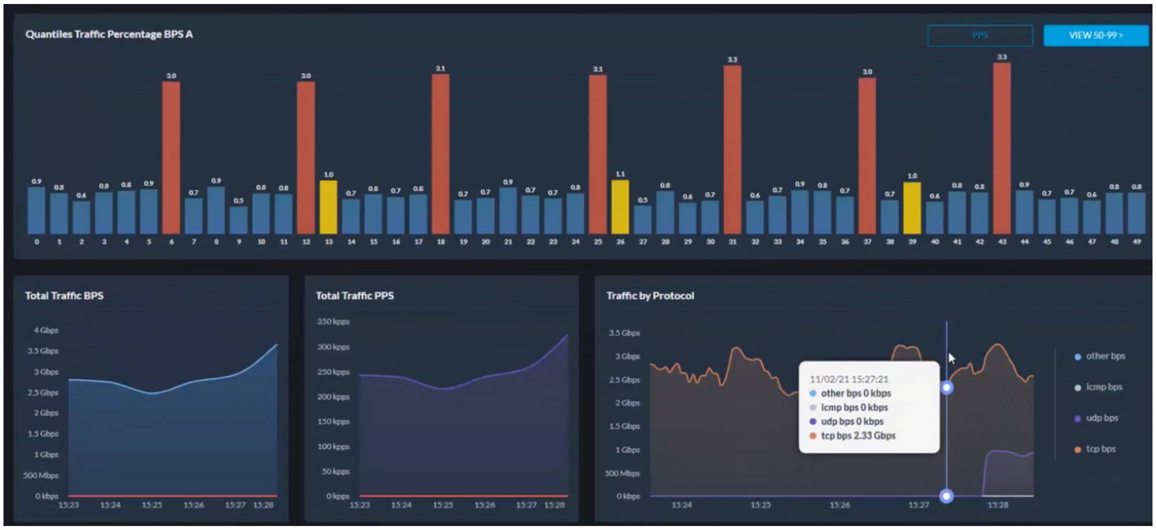

Figure 16. Quantiles – visualization of an attack with attack traffic against legitimate traffic

If we examine the attack dashboard below, we can get more specific information about the attack in question. In this case, the attack belongs to a specific protected object, which is a managed entity representing a range of IP addresses belonging to a specific customer or group. When the system detects an attack, it will alert these specific organizations.

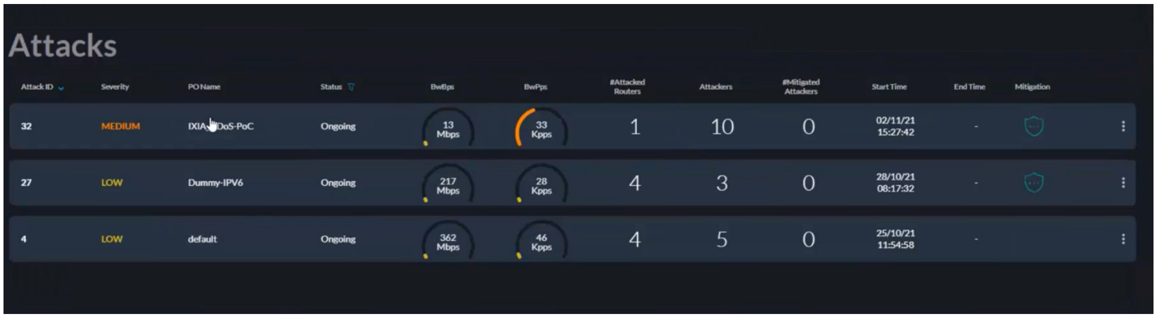

Attacks

Attack ID	Severity	PO Name	Status	BwBps	BwPps	#Attacked Routers	Attackers	#Mitigated Attackers	Start Time	End Time	Mitigation
32	MEDIUM	IXIA_DDoS-PoC	Ongoing	13 Mbps	33 Kpps	1	10	0	02/11/21 15:27:42	-	
27	LOW	Dummy-IPV6	Ongoing	217 Mbps	28 Kpps	4	3	0	28/10/21 08:17:32	-	
4	LOW	default	Ongoing	362 Mbps	46 Kpps	4	5	0	25/10/21 11:54:58	-	

Figure 17. Dynamic attack dashboard with mitigation status and actions

As we investigate further by clicking on the attack in question (three are indicated above), we arrive at a much more detailed description of what is happening in the attack.

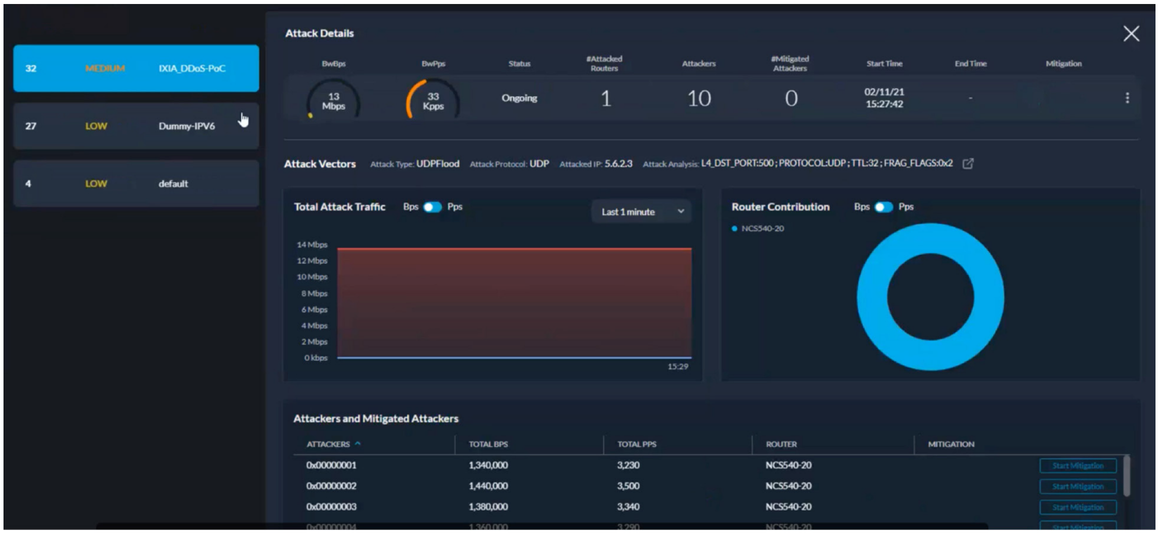

32	MEDIUM	IXIA_DDoS-PoC
27	LOW	Dummy-IPV6
4	LOW	default

Attack Details

BwBps	BwPps	Status	#Attacked Routers	Attackers	#Mitigated Attackers	Start Time	End Time	Mitigation
13 Mbps	33 Kpps	Ongoing	1	10	0	02/11/21 15:27:42	-	

Attack Vectors Attack Type: UDPFlood Attack Protocol: UDP Attacked IP: 5.6.2.3 Attack Analysis: L4_DST_PORT:500 ; PROTOCOL:UDP ; TTL:32 ; FRAG_FLAGS:0x2

Attackers and Mitigated Attackers

ATTACKERS	TOTAL BPS	TOTAL PPS	ROUTER	MITIGATION
0x00000001	1,340,000	3,230	NCS540-20	Start Mitigation
0x00000002	1,440,000	3,500	NCS540-20	Start Mitigation
0x00000003	1,380,000	3,340	NCS540-20	Start Mitigation
0x00000004	1,360,000	3,290	NCS540-20	Start Mitigation

Figure 18. Dynamic detailed attack visualization

We can see the level of the attack is about 30 Mbps, that there are 10 attackers, and how the attack progressed over time. Since the protected object (PO) is set to “user confirmation,” in the graphic below you will see the “options” presented to the operator.

The choices presented to the operator include, but are not limited to:

- Start mitigation – Implement the TCAM-based access control list into the router.
- Clear the attack – The operator has determined that this is not, in fact, an attack.
- The operator could have chosen automatic mitigation as well on the protected object.

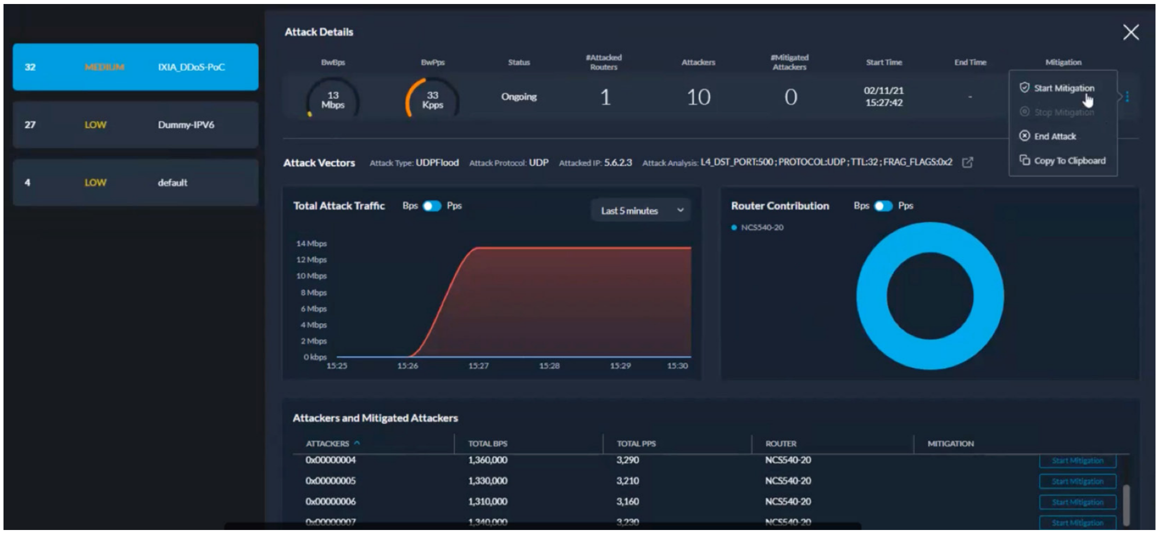

32	MEDIUM	IXIA_DDoS-PoC
27	LOW	Dummy-IPV6
4	LOW	default

Attack Details

BwBps	BwPps	Status	#Attacked Routers	Attackers	#Mitigated Attackers	Start Time	End Time	Mitigation
13 Mbps	33 Kpps	Ongoing	1	10	0	02/11/21 15:27:42	-	

Attack Vectors Attack Type: UDPFlood Attack Protocol: UDP Attacked IP: 5.6.2.3 Attack Analysis: L4_DST_PORT:500 ; PROTOCOL:UDP ; TTL:32 ; FRAG_FLAGS:0x2

Attackers and Mitigated Attackers

ATTACKERS	TOTAL BPS	TOTAL PPS	ROUTER	MITIGATION
0x00000004	1,360,000	3,290	NCS540-20	Start Mitigation
0x00000005	1,330,000	3,210	NCS540-20	Start Mitigation
0x00000006	1,310,000	3,160	NCS540-20	Start Mitigation
0x00000007	1,340,000	3,230	NCS540-20	Start Mitigation

Figure 19. Dynamic detailed attack visualization - 2

The attack vector shows detailed information about the attack. This information can be exported or simply copied from the controller graphical user interface and pasted to another system, manually or via API. This information includes attacked IP, source, attack type, and attack analysis, providing deeper detail on the attack.

When we click start mitigation and we look at the configuration of the router, the following access is applied to mitigate the attack for only the 10 attacking mobile UEs. The format of the TCAM-based access control list implemented to mitigate the attack is shown in the graphic below.

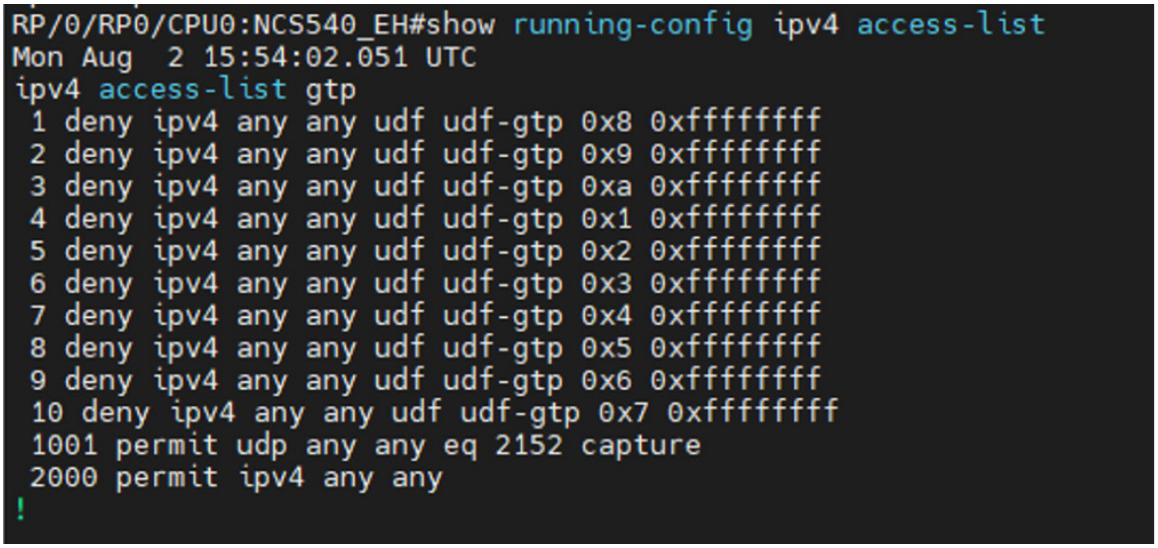

```
RP/0/RP0/CPU0:NCS540_EH#show running-config ipv4 access-list
Mon Aug  2 15:54:02.051 UTC
ipv4 access-list gtp
 1 deny ipv4 any any udf udf-gtp 0x8 0xffffffff
 2 deny ipv4 any any udf udf-gtp 0x9 0xffffffff
 3 deny ipv4 any any udf udf-gtp 0xa 0xffffffff
 4 deny ipv4 any any udf udf-gtp 0x1 0xffffffff
 5 deny ipv4 any any udf udf-gtp 0x2 0xffffffff
 6 deny ipv4 any any udf udf-gtp 0x3 0xffffffff
 7 deny ipv4 any any udf udf-gtp 0x4 0xffffffff
 8 deny ipv4 any any udf udf-gtp 0x5 0xffffffff
 9 deny ipv4 any any udf udf-gtp 0x6 0xffffffff
 10 deny ipv4 any any udf udf-gtp 0x7 0xffffffff
 1001 permit udp any any eq 2152 capture
 2000 permit ipv4 any any
!
```

Figure 20. DDoS mitigation with ACL on the router

Now that the attack is mitigated, as you see in the graphic below, the red quantiles are gone, and the system is back to "peace time."

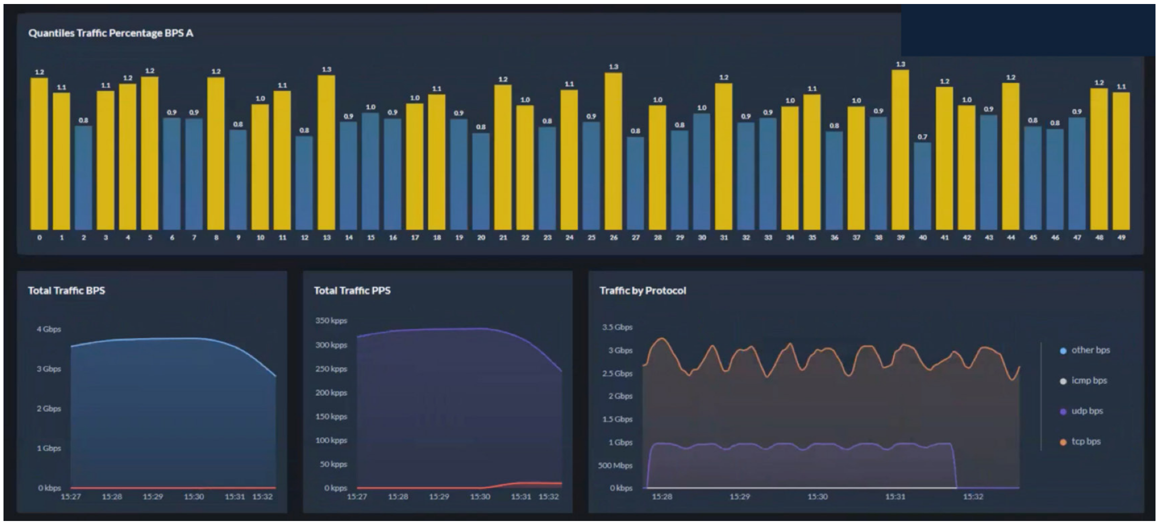

Figure 21. Back to "peace time"

The system continues to monitor the TCAM-based access control list to make sure that the attack has finished. If the "hits" on the access control list continue to increase, there is evidence that the attack is still happening, and one of two outcomes will happen. First, the access control list will remain in place and continue to block only the attacking traffic, or second, the operator will select to "end the attack" and that's what will happen with the access control list being removed from the router(s) in question.

4. Conclusions

The Cisco Secure DDoS Edge Protection solution is built for 5G. Deploying the solution at the cell site router preserves 5G sub-10-ms, low-latency requirements. The auto-tuning quantile algorithm, designed specifically for DDoS protection at the "edge," reduces security operations (SECOP) lifecycle overhead and is the foundation of the solution. Deployment "out of path" avoids impact on data plane resources. Communications between solution components, the controller, and the detectors rely on mutual authentication for 3GPP standard-based communication.

The Cisco Secure DDoS Edge Protection is "smarter" in its deployment at the network "edge." Because the solution provides line rate DDoS protection at the cell site router, this means no traffic gating at the User Plane Function (UPF), resulting in better quality of experience for user access to applications. At the mobile "edge," granular visibility accurately distinguishes malicious traffic from legitimate traffic by looking at User Equipment (UE) sessions within GTP tunnels, providing comprehensive anomaly awareness.

The Cisco Secure DDoS Edge Protection is highly optimized. The solution is highly optimized by leveraging a data model that provides visibility far beyond what is available today due to the use of Google Protobuf (GPB) and the application of User-Defined Fields in the Cisco network operating system, IOS XR. Although the mobile "edge" was described in depth in this white paper, the technology is applicable to other platforms that run IOS XR, including the family of Cisco routers and white-box systems that run IOS XR as the network operating system, allowing for application of this solution to other important places in the architecture such as a newly optimized methodology to solve for today's hyperscale infrastructure security protection. Legitimate traffic from other UEs is uninterrupted during an attack, protecting the quality of experience (QoE) for the entire network. This allows service providers to maximize efficiency and availability while delivering new 5G outcomes.

5. Learn more

Use the links below to learn more about DDoS and Cisco Secure DDoS Edge Protection:

- Cisco Secure DDoS Edge Protection on DEVNET: https://developer.cisco.com/docs/secure-ddos-edge-protection
- Cisco Secure DDoS Edge Protection AAG: www.cisco.com/c/en/us/products/collateral/security/secure-ddos-edge-protection-aag.pdf
- Cisco Secure DDoS webpage: www.cisco.com/go/secure-ddos
- Edge Protection support email alias: secure-ddos-edge-protection@external.cisco.com
- "What Is a DDoS attack?" on Cisco.com - www.cisco.com/c/en/us/products/security/what-is-a-ddos-attack.html

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte.Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at https://www.cisco.com/go/offices.

For information about Cisco security solutions that enable Any Device, go to: www.cisco.com/go/security