

DEVNET Demo Guide

Cisco Secure

DDoS Edge Protection with IOS XR



Contents

Goals of this document	3
Connecting to the Sandbox	4
Overview	4
Device credentials	5
Objectives	5
Gain familiarity with the Cisco Secure DDoS Edge Protection controller interface	5
Prerequisites	5
Terminology	5
<i>From this point on, you will need to be connected to IOS XR and Edge Protection Sandbox</i>	5
Step 1: Reviewing the controller	6
Step 2: Generating traffic	14
Step 3: Detecting an attack	21
Step 4: Stopping the attack	36
Resources	41



Goals of this document

The goal of this document is to demonstrate Cisco Secure DDoS Edge Protection® in service provider mobile and peering environments. After this lab, the user should be familiar with Cisco Secure DDoS Edge Protection, an industry-leading technology that blocks distributed-denial-of-service (DDoS) attacks at the network edge, enabling service providers to meet the sub-10-ms latency requirements of modern 5G networks and ensuring customer quality of experience (QoE).

This demo guide will provide an overview of the Edge Protection controller, which provides centralized management of one (1) or more Edge Protection detectors. The controller also provides the visualization of the status of each detector and any ongoing attacks in the network.

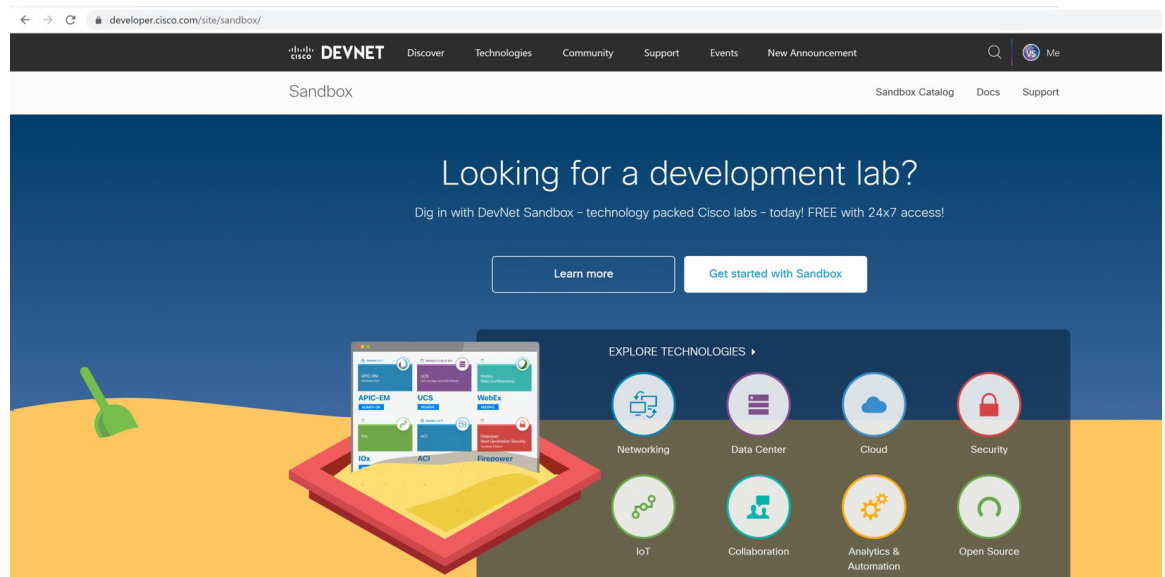
Upon completion of this lab, you should be able to demonstrate the value of deploying Cisco Secure DDoS Edge Protection in your customer’s network to protect them and their customers from service interruptions due to DDoS attacks.

Scheduling the Sandbox

This lab is meant to be run on the DEVNET Sandbox environment. To schedule the Sandbox, you will need to reserve the Cisco IOS® XR and Edge Protection Sandbox.

For information about how to reserve a DEVNET Sandbox, go to:
<https://developer.cisco.com/docs/sandbox/#!getting-started/try-it-out>.

To book a Sandbox, go to <https://developer.cisco.com/site/sandbox/>, click “Get Started with Sandbox,” and search for the “DDoS Edge Protection” sandbox.





Connecting to the Sandbox

Once the Sandbox has been scheduled, an automated email will be sent to the user with the VPN URL and the credentials to log in. The Cisco AnyConnect® client is required to be installed on the student’s computer to connect to the DEVNET Sandbox environment. If you don’t have AnyConnect installed, the email sent with the login information will also provide a link to download the AnyConnect client.

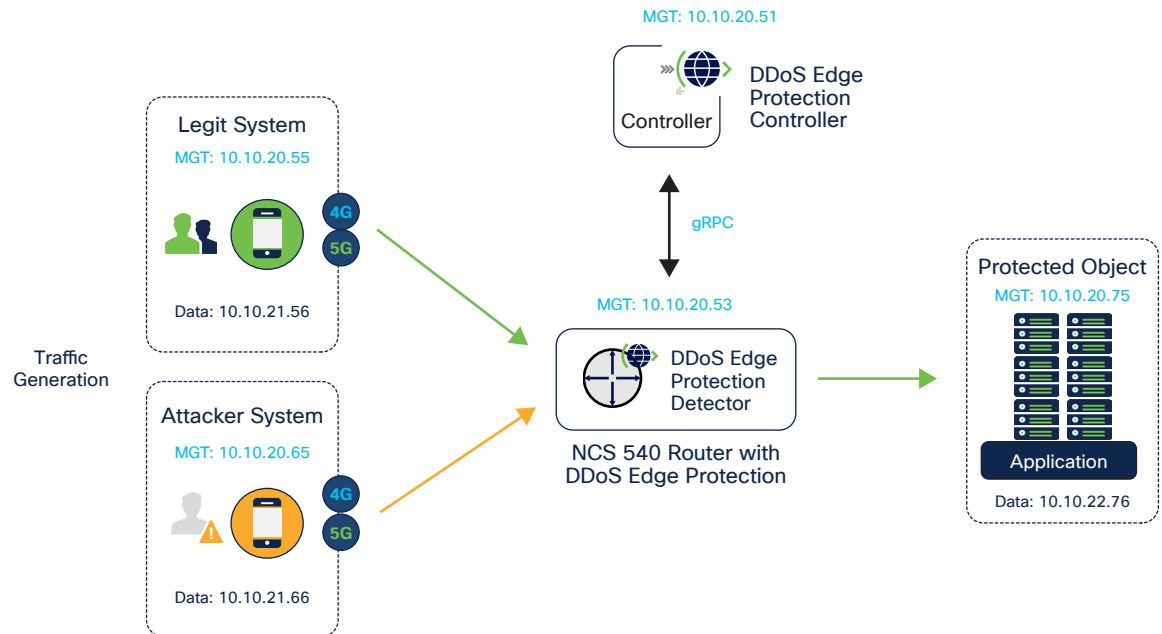
Overview

In this lab, we will be protecting a mobile edge network environment from malicious attacks that are hidden inside GTP packets.

We will review the Controller Dashboard and other menus. Once we have familiarized ourselves with the controller interface, we will launch traffic from a legitimate user and will review the relevant information and statistics. Then we will launch malicious traffic from another device and see how the Edge Protection solution almost immediately protects the network by identifying and mitigating this traffic.

The following describes the solution architecture:

In this lab, there are both legitimate and malicious (attacker) devices that connect to an application that’s behind a Cisco® NCS 540 router. On this NCS 540, there is a Cisco Secure DDoS Edge Protection detector deployed that analyzes all the flows of data being received by the router. The Edge Protection detector is controlled and managed by the Edge Protection controller. Visibility into the traffic passing through the detector can be observed on the controller graphical user interface.





Device credentials:

Device	Username	Password	MGT access
Cisco NCS 540	client	cisco123	10.10.20.53
Controller HTTP	admin@example.com	12345	http://10.10.20.51.nip.io
Legit user	demo	cisco123	10.10.20.55
Attack user	demo	cisco123	10.10.20.65
Protect application	demo	cisco123	10.10.20.75

Objectives

- Gain familiarity with the Edge Protection controller interface
- Launch legitimate traffic and see how it is viewed within the Edge Protection controller interface
- Launch attack traffic and see how it is easily and quickly detected by the Edge Protection detector
- Block attack traffic

Prerequisites

No prerequisites are required, but having knowledge of the challenges that the DDoS poses for enterprise and service provider customers and networks will help the student appreciate the value that Edge Protection offers customers.

Applications needed:

- Cisco AnyConnect client (link provided in email)
- Web browser
- SSH program, such as Putty

Terminology

This document uses the following terms with which you must be familiar:

- Controller - A central management function and a user Interface that manages a fleet of one or more detectors. The controller includes a GUI dashboard that presents real-time attack information for detector visibility, forensics, and threat intelligence analysis.
- Detector - DDoS detection and mitigation are functions implemented on a virtual Docker container as a microservice application. The function runs independently on the designated edge router.
- GTP - GPRS Tunneling Protocol defined by the 3GPP standards to carry packets between mobile function zones in 3G/4G and 5G.

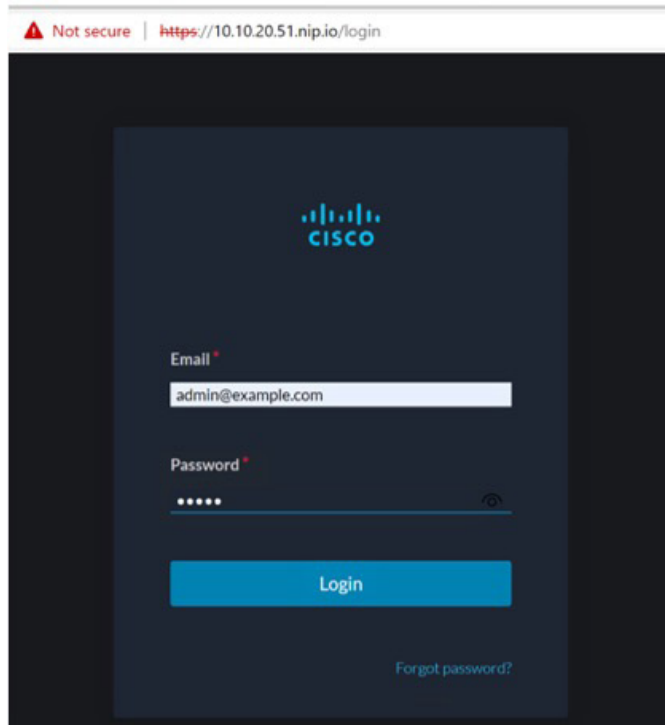
Note: From this point on, you will need to be connected to IOS XR and Edge Protection Sandbox.



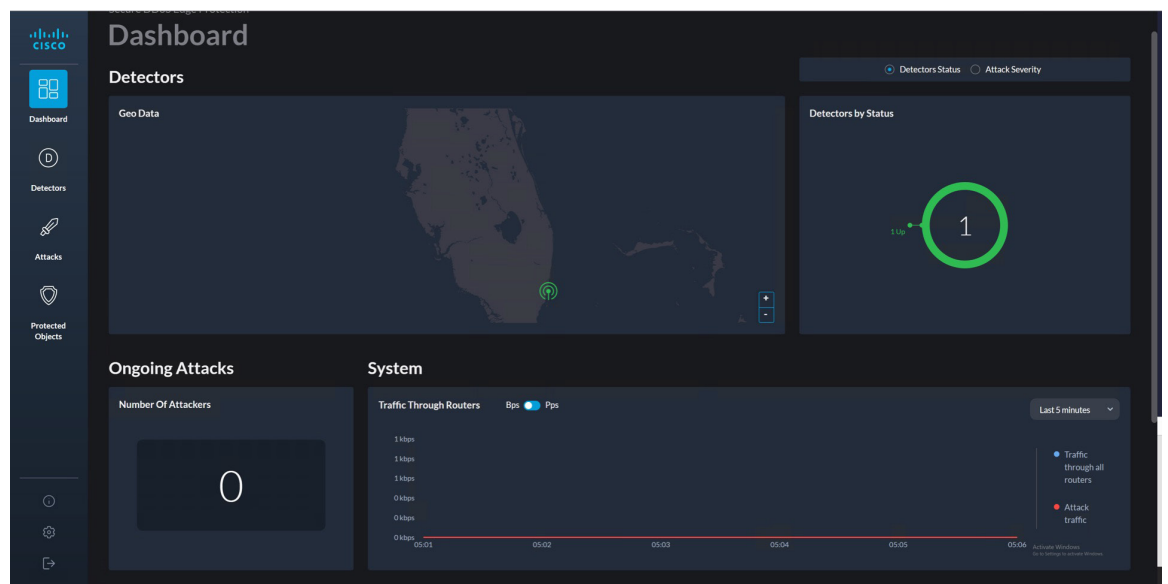



Step 1: Reviewing the controller


Open Chrome (or other browser) and connect to the Edge Protection controller, <http://10.10.20.51.nip.io>. For the email, enter admin@example.com, and for the password, use: [12345](#). Then click “Login.”



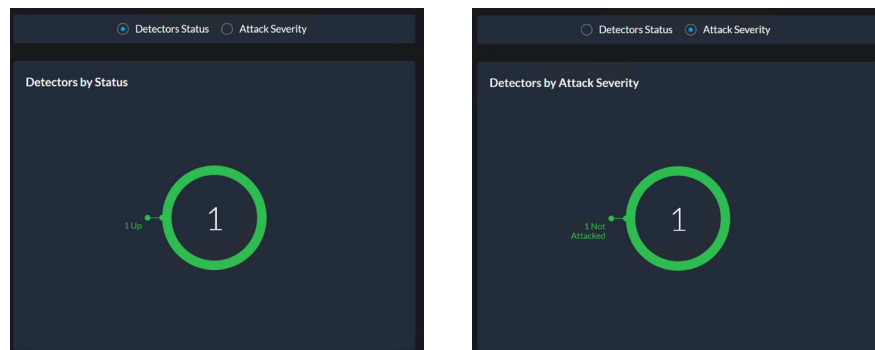
Once logged in, you will be in the Controller Dashboard.



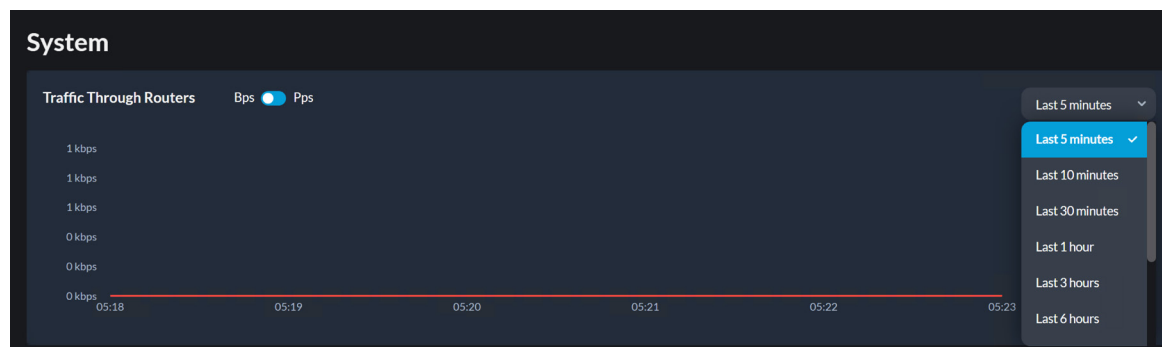
Along the left side of the page is the navigation bar, which allows access to different menus. The section being displayed has a light-blue background. As you can see, we are in the dashboard . The dashboard provides an overview of your environment. It shows the number of detectors, whether any detectors are under attack, and the total amount of both legitimate and attack traffic. Scroll down to see more information in the dashboard.




In the top right-hand corner of the dashboard, we see the number of detectors, and, by default, we also visually see their overall status. We also have the option to view by attack severity. In the example below, there is one (1) detector, and the status is green, indicating that it is active. When we click on the attack severity, it is also green, indicating that the network is not under attack.



Looking at the “Traffic Through Routers” graph below under “System,” we notice that no traffic is shown. The default time window is the last five minutes, but this can be changed by clicking on the dropdown menu on the top right of the graph.

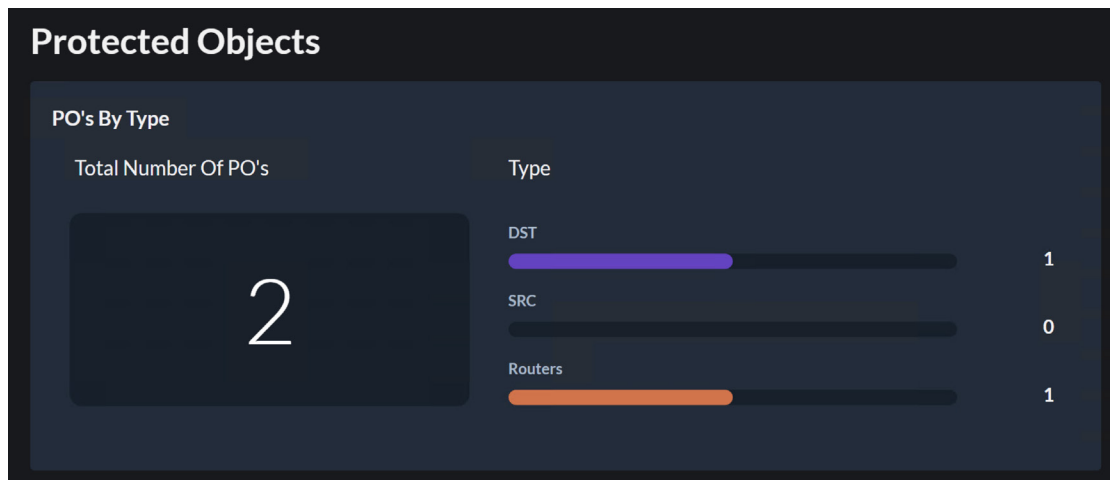


To the left of the “System > Traffic Through Routers” graph, we see the ongoing attacks graph. As shown below, there are also no attacks.




Scrolling down in the dashboard page, we get an overview of the protected objects (POs). Starting on the left, we can see the POs by type.

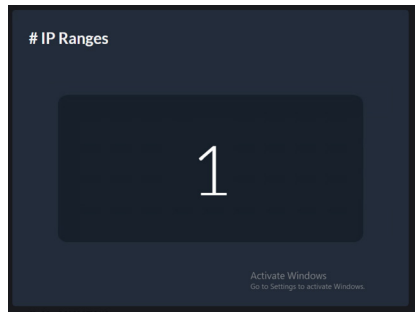
There are two (2) POs defined. One PO is defined by destination IP, none by the source IP, and one by router. We will get more details about this PO on the Protected Objects page.



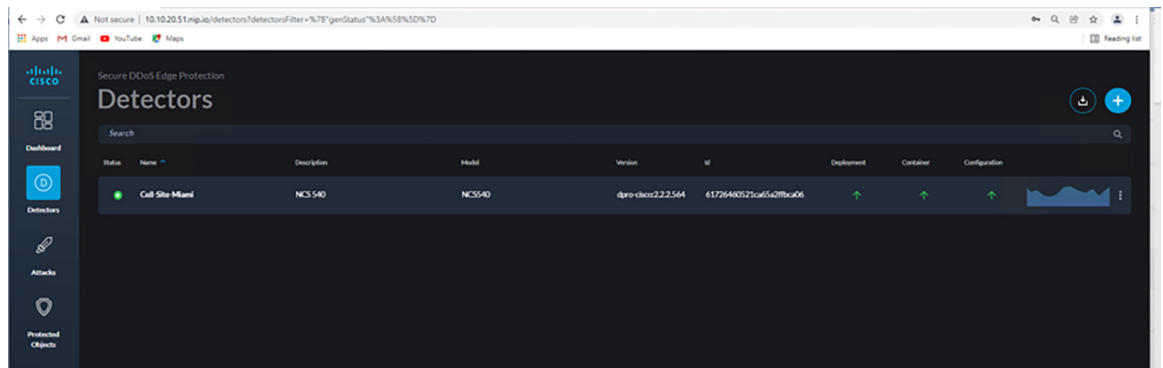
The "PO's Under Attack" widget provides information on how many of the POs are under attack. Presently there are none under attack.





Finally, on the dashboard under "Protected Objects," we can see the number of IP ranges in our POs. In this case, it is one.


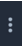



We will now review the Detectors menu page. Click on the Detector Icon  on the left-hand side of the screen.




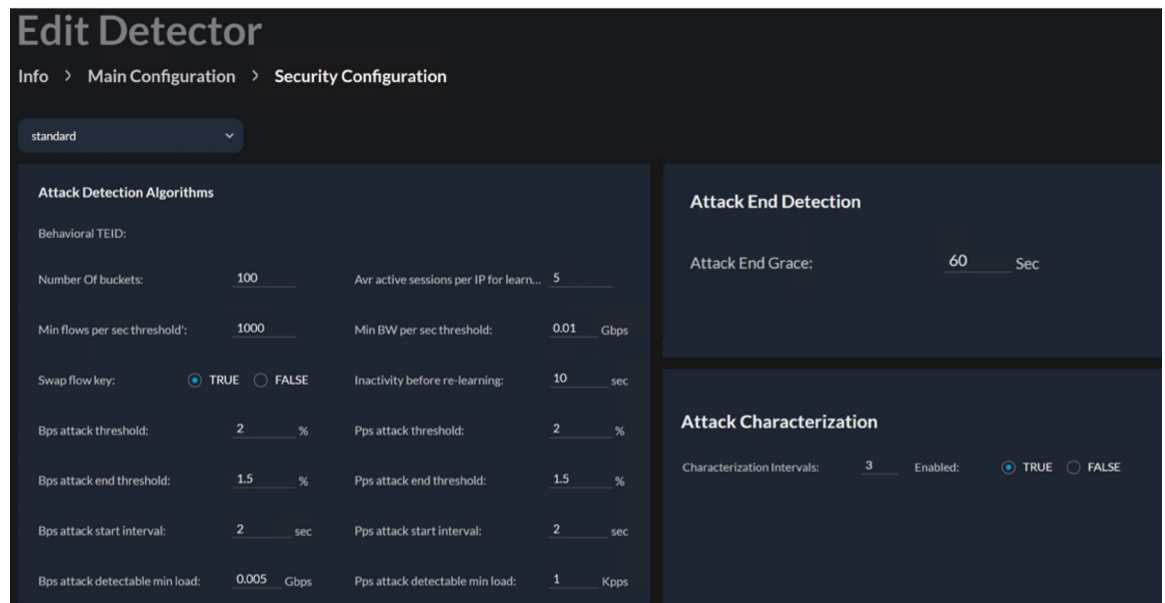
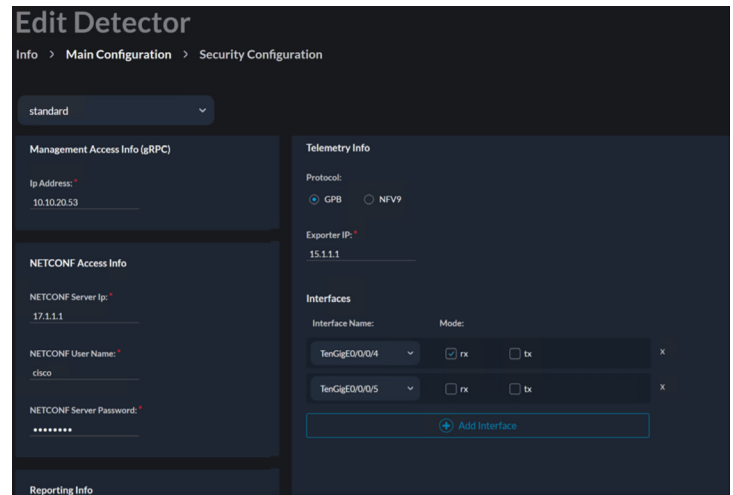
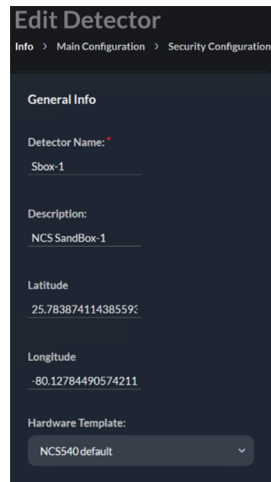
We see one detector with the status of green. Looking to the right, we can see the status for Deployment, Container, and Configuration. The green arrows indicate that all three are working correctly. See the enlargements below.



To add more detectors, one can click on the  icon and fill in the required fields for the router in which we would want to deploy the Edge Protection detector. We can also edit/view the existing detector details by clicking on the  icon to the right-hand side of the detector row and selecting “Edit.”



In the Edit Detector menu, we have three (3) tabs: Info, Main Configuration, and Security Configuration. Only the Info and Main Configuration need to be configured. The Security Configuration can be left at the default settings and only needs to be tweaked as necessary.

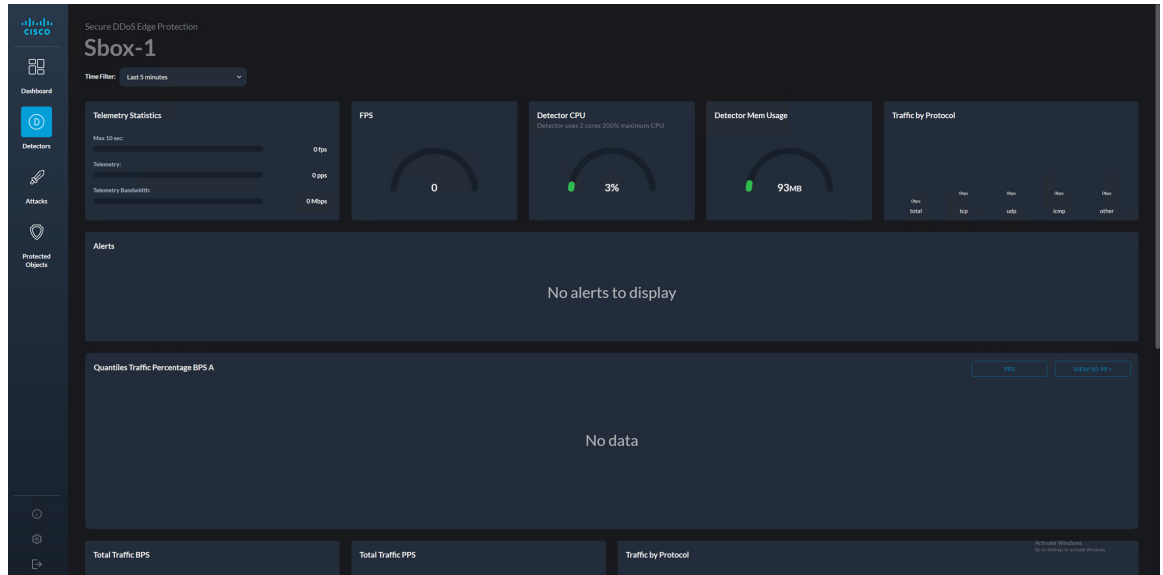


Click on the “Cancel” button on the bottom right.

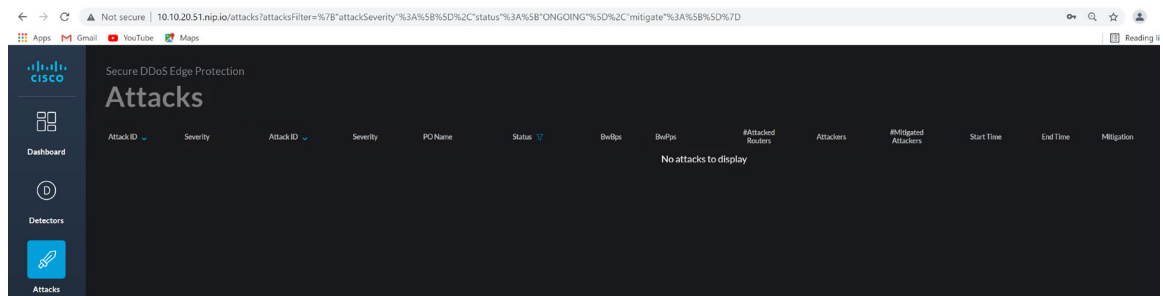


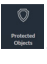

Now click on the graph icon to the left of the .


Here we get details specific to an individual detector. This page will be more interesting once we have both legitimate and attack traffic going through this particular router being monitored by the detector.

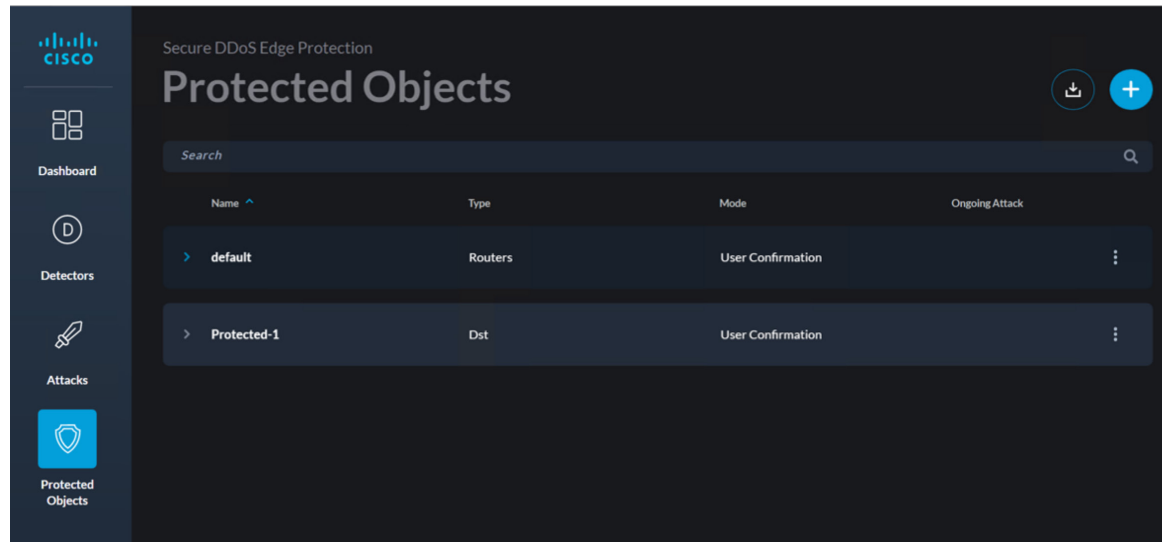



Click on the Attacks icon . Here we see the active attacks. This page is empty at the current time.




Next, click on the Protected Objects icon . Here we can view the two POs, the default PO and one manually configured object called Protected-1. To create more POs, we would click on the  located on the top right.

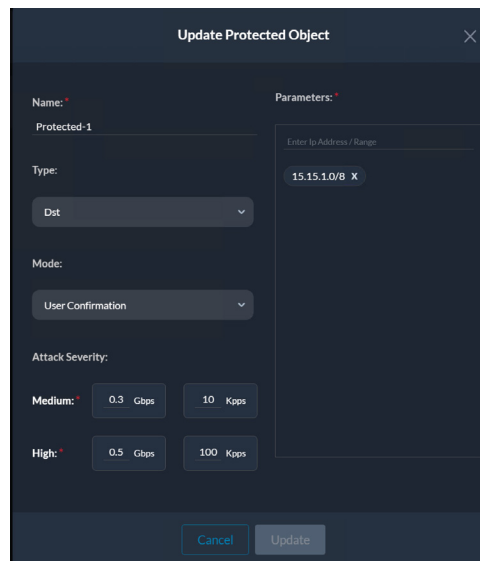





Click on the  at the end of the row for the Protected-1 object, and then click “Edit.”




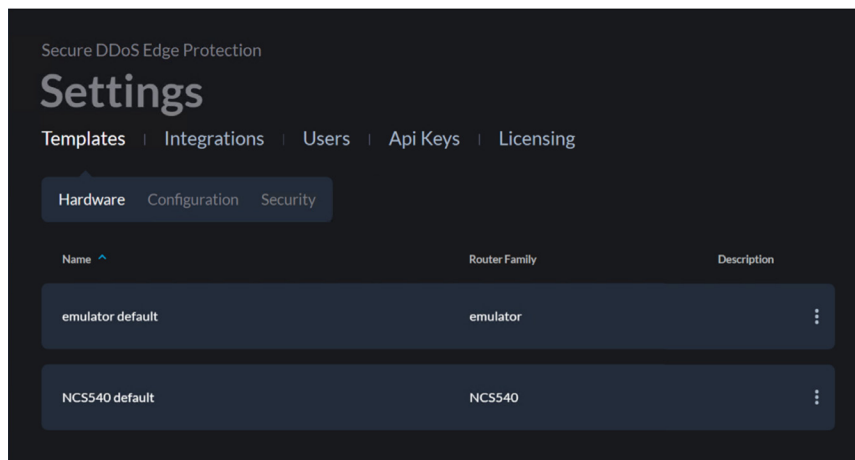
We can see the details for this PO. We see that the type is a destination IP (Dst). We can see that the value of the IP is 15.15.1.0/8 and, below that, the thresholds that classify an attack as medium or high severity. Click “Cancel.”



At the bottom left, we see information regarding the Edge Protection controller itself. Click on the  icon. Here we can see the details of the software that is installed on the various engines of the controller.




Click on the  icon. From the Settings page, you can access templates for the detectors, create users, and add the licensing.



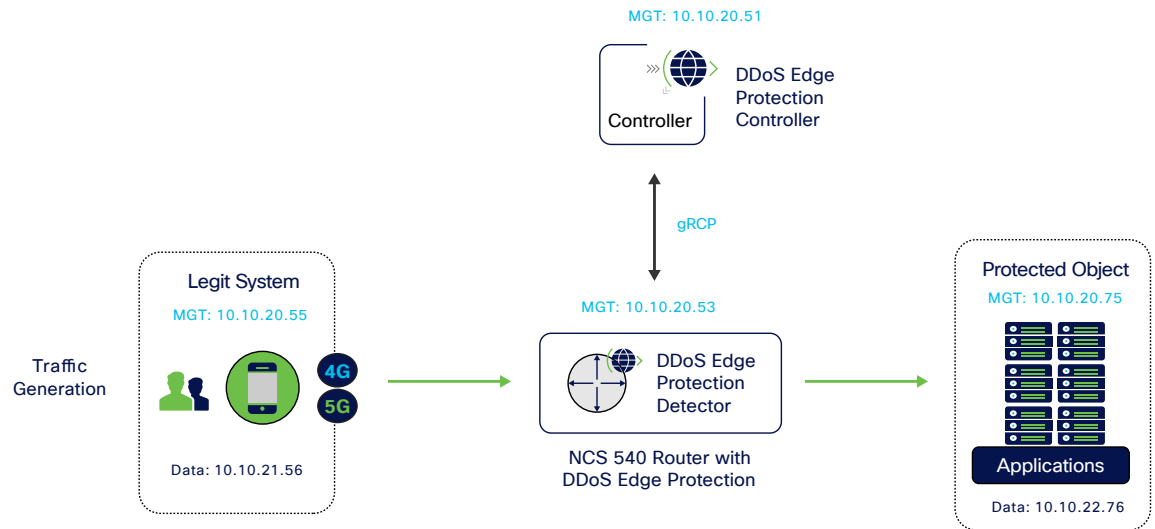
Finally, the last icon is the Exit Controller icon , which will log off the user from the controller.





Step 2: Generating traffic

Now that we are familiar with the Edge Protection controller interface, it is time to generate traffic from the legitimate user. We will be sending traffic to the protected application on the right of the diagram below.

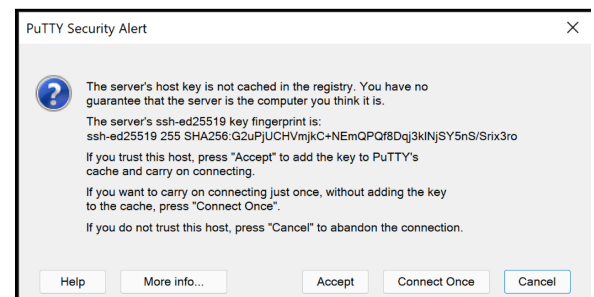
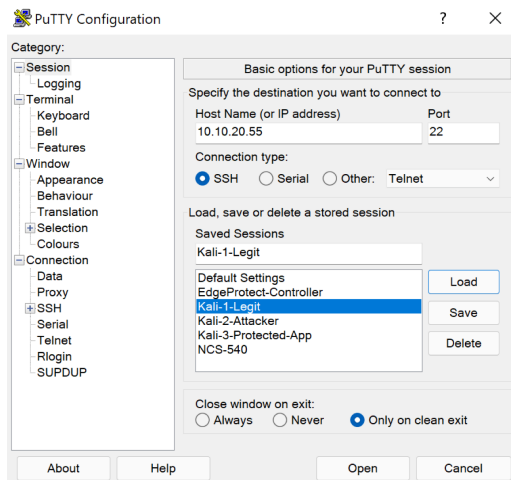


SSH into the legit system with the IP address of 10.10.20.55 and the following credentials:

Username: demo

Password: cisco123

If prompted with a security alert, click “Accept.” This happens because it is your first time connecting to this device securely.



Move into the ddos-work directory with the following command:

```
cd ddos-work
```

Now execute the script called run-legit-traffic.sh:

```
./run-legit-traffic.sh
```

The Kali Linux box will now generate traffic and send it to the protected object.

Note: Wait about 30 seconds to ensure that the traffic is being seen by the router and the detector.

```

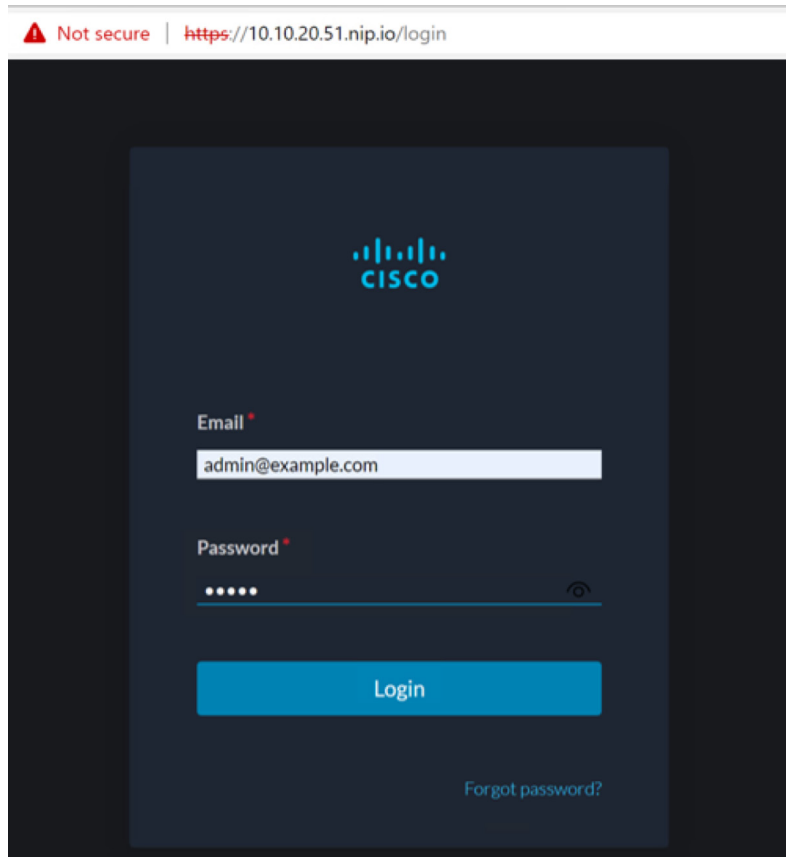
demo@kali-1-Legit: ~/ddos-work
demo@kali-1-Legit:~/ddos-work$
demo@kali-1-Legit:~/ddos-work$
demo@kali-1-Legit:~/ddos-work$
demo@kali-1-Legit:~/ddos-work$
demo@kali-1-Legit:~/ddos-work$ ./run-legit-traffic.sh
HPING 15.15.1.1 (eth0 15.15.1.1): NO TCP FLAGS are set (GTP Tunneling), 48 heads
rs + 1400 data bytes
ping in flood mode, no replies will be shown
  
```

Notice in the above image, the HPING command shows a destination IP address of 15.15.1.1. This is the final destination of the packet that has been encapsulated in the GPRS tunnel protocol (GTP). The GTP tunnel endpoint is 10.10.22.76, which can be clearly seen in the packet capture below.

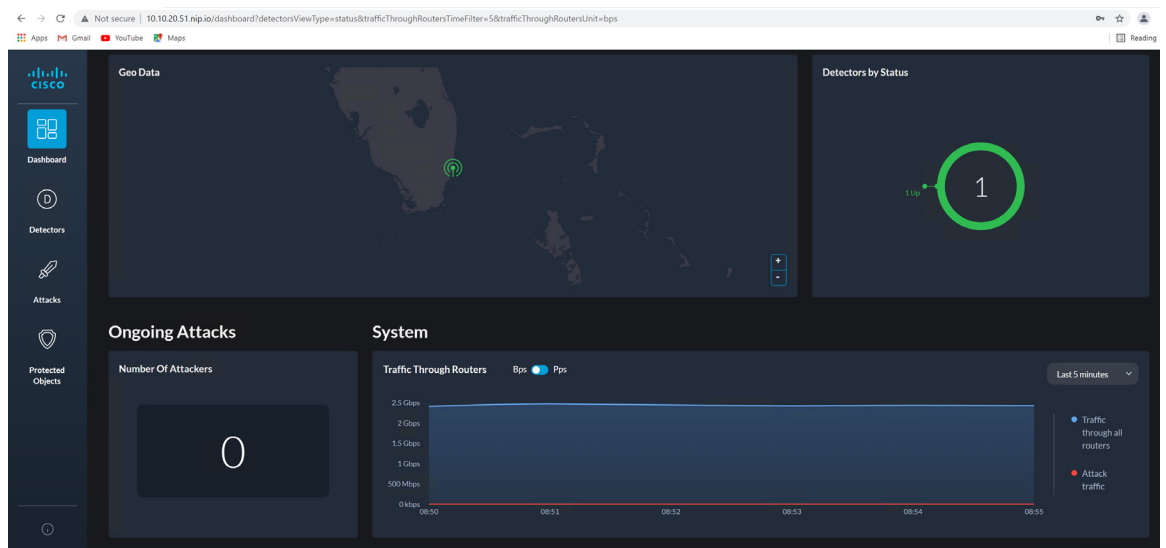
```

Wireshark · Packet 1 · mypcap2.pcap
> Frame 1: 1490 bytes on wire (11920 bits), 1490 bytes captured (11920 bits)
> Ethernet II, Src: VMware_bf:33:30 (00:50:56:bf:33:30), Dst: Cisco_06:e0:10 (d4:6a:35:06:e0:10)
> Internet Protocol Version 4, Src: 10.10.21.66, Dst: 10.10.22.76
> User Datagram Protocol, Src Port: 2152, Dst Port: 2152
> GPRS Tunneling Protocol
  > Internet Protocol Version 4, Src: 10.10.21.55, Dst: 15.15.1.1
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
    Total Length: 1440
    Identification: 0xbfce (49102)
  > Flags: 0x00
    ...0 0000 0000 0000 = Fragment Offset: 0
    Time to Live: 64
    Protocol: TCP (6)
    Header Checksum: 0x8639 [validation disabled]
    [Header checksum status: Unverified]
    Source Address: 10.10.21.55
    Destination Address: 15.15.1.1
  > Transmission Control Protocol, Src Port: 29033, Dst Port: 80, Seq: 1, Len: 1400
  
```

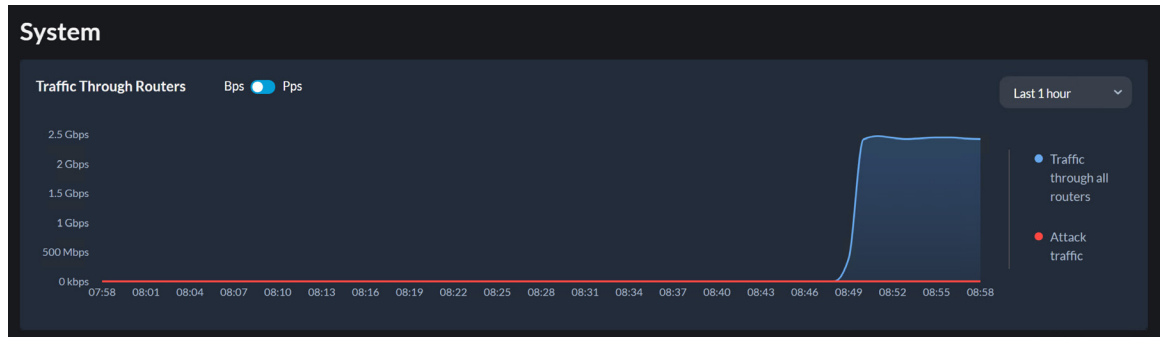
Now let's go back to the controller graphical user interface (GUI) portal and review some of the widgets. Open Chrome (or other browser) and connect to the Edge Protection controller: <http://10.10.20.51.nip.io>. For the email, enter `admin@example.com`, and for the password, use: 12345. Then click "Login."





After logging in, we can see traffic. See the blue line in the “System > Traffic Through Routers” pane two screenshots below.



We can change the time from “Last 5 minutes” to “1 hour” (top right of the widget) to see the traffic going from 0 to over 2 Gbps.



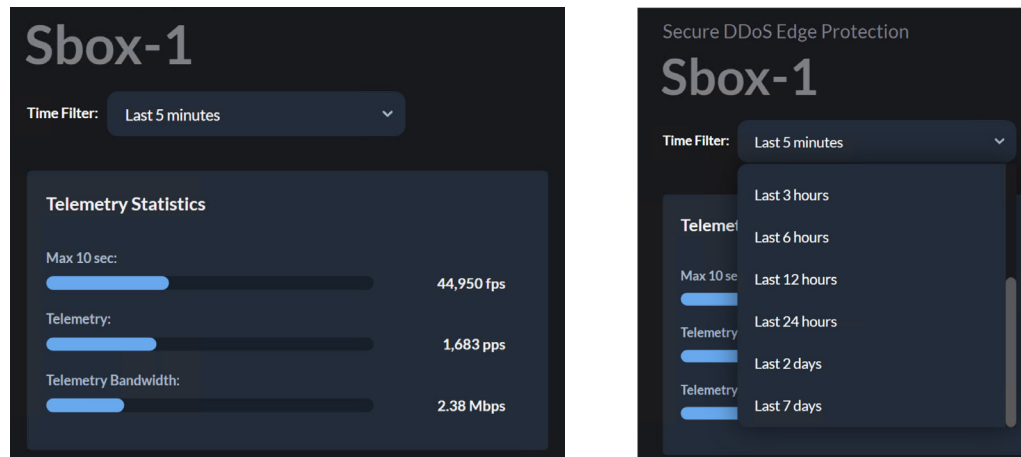
Click on the Detector icon  on the left-hand side of the screen. Then click on the graph icon  on the right-hand side of the detector row.

Status	Name	Description	Model	Version	Id	Deployment	Container	Configuration
OK	Cell-Site-Miami	NCS-540	NCS-540	dpro-cisco:2.2.2.564	61726460521ca65a2ffca06	↑	↑	↑

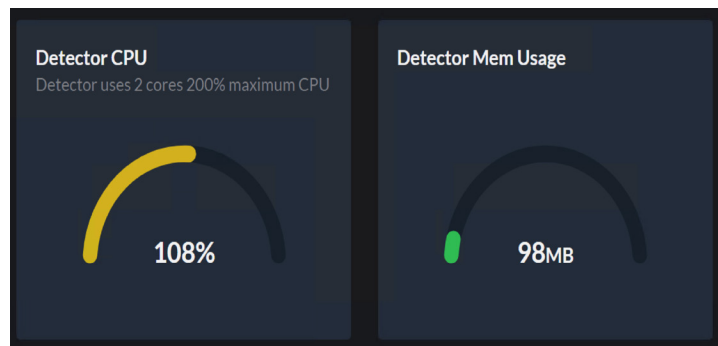
We can obtain a lot of information about what is being seen by the detector.

Starting from the top left, we can see telemetry statistics, which show the number of frames per second (fps), packets per second (pps), and the bandwidth that is used to collect statistics (telemetry) from the router.


Click on the time filter to notice that these statistics are for the last five minutes. We can change this range from one minute up to seven days.



We can also see the amount of CPU and memory being consumed by the detector. Note that the CPU can be over 100%. Since there are two (2) CPUs, the range is up to 200% (2 x 100%).




On the top right, the traffic is broken down by type. We can see that all the traffic is TCP.



Scrolling further down in the page, notice a bar graph with quantiles. This shows the traffic being broken down into blocks. These blocks are known as quantiles. The quantiles are then analyzed by the detector for anomalies and attacks.



Just below the quantile traffic graph we have the “Total Traffic BPS,” “Total Traffic PPS,” and “Traffic by Protocol” graphs.



Scrolling right to the end of the page we get a list of top talkers.

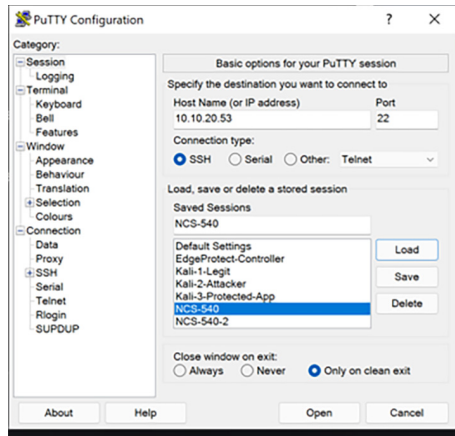
TIME	KEY	UDP BPS	OTHER BPS	TCP BPS
10/30/21 09:23:32	10.10.20.55	0	0	2,266,041,600
10/30/21 09:23:29	10.10.20.55	0	0	2,543,616,000
10/30/21 09:23:26	10.10.20.55	0	0	2,569,939,200
10/30/21 09:23:23	10.10.20.55	0	0	2,553,408,000

Next log in to the Cisco NCS 540 router and take a look at the flows being forwarded to the detector. SSH into Cisco NCS 540 with the IP address of 10.10.20.53 and use the following credentials:

Username: demo
Password: cisco123

If prompted with a security alert, click “Accept.” This happens because it is your first time connecting to this device securely.





Run the following command on the router to view the GTP access-list:

```
show flow monitor DetectPro_NFv9 cache brief location 0/0/CPU0
```

The output shows details of the flows being sent from the router to the detector.

```
RP/0/RP0/CPU0:NCS1#show flow monitor DetectPro_NFv9 cache brief location 0/0/C$
Wed Nov 24 14:58:48.029 UTC
Cache summary for Flow Monitor DetectPro_NFv9:
Cache size:                1000000
Current entries:           2349
Flows added:               6461486834
Flows not added:          0
Ager Polls:               2357805
- Active timeout          6461484485
- Inactive timeout        0
- Immediate               0
- TCP FIN flag            0
- Emergency aged         0
- Counter wrap aged      0
- Total                  6461484485
Periodic export:
- Counter wrap            0
- TCP FIN flag            0
Flows exported             6461484485

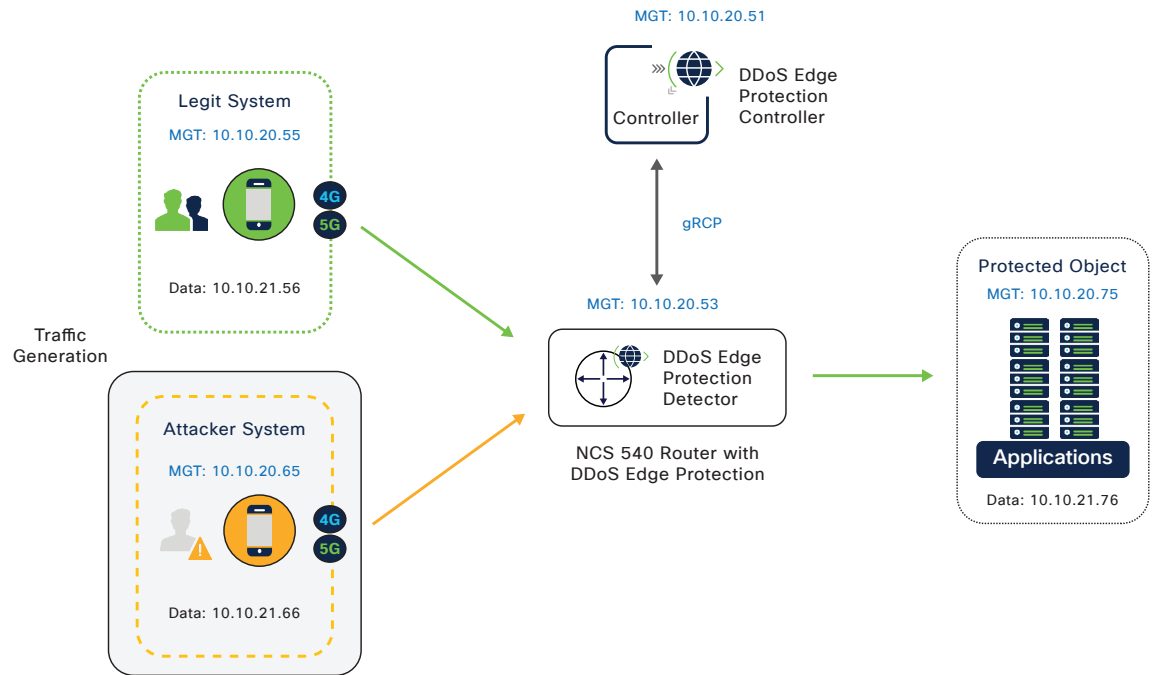
RecordType  IPV4SrcAddr      IPV4DstAddr      L4SrcPort  L4DestPort  IPV4Prot  IPV4
TOS  InputInterface  ByteCount      PacketCount  Dir
GTP  Tunneled Record  10.10.21.55    15.15.1.1    50401      80         tcp
0    Te0/0/0/4        1440           1            Ing
GTP  Tunneled Record  10.10.21.55    15.15.1.1    3783      80         tcp
0    Te0/0/0/4        1440           1            Ing
GTP  Tunneled Record  10.10.21.55    15.15.1.1    16166     80         tcp
0    Te0/0/0/4        1440           1            Ing
```

Let's launch an attack. Please go to the next step.



Step 3: Detecting an attack

We are now going to launch an attack. We want to keep the legitimate traffic going, so make sure you have not stopped the script from the legitimate server. If required, return to the previous section to see how to launch the legitimate traffic.

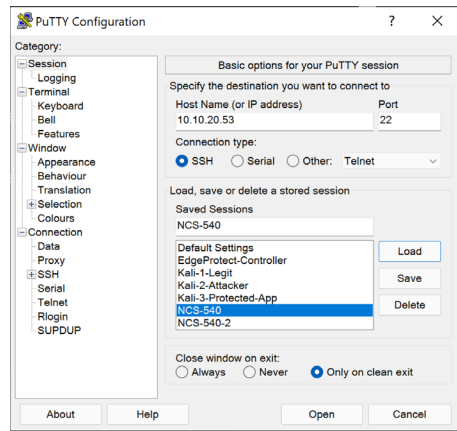


Log back into the Cisco NCS 540 Router and take a look at the access-list policy that has been preconfigured.

SSH into Cisco NCS 540 with the IP address of 10.10.20.53 and use the following credentials:

Username: client

Password: cisco123



Run the following command on the router to view the GTP access-list:

```
show access-list gtp
```

The output shows details of the flows being sent from the router to the detector.

Presently there are three entries in the GTP access-list.

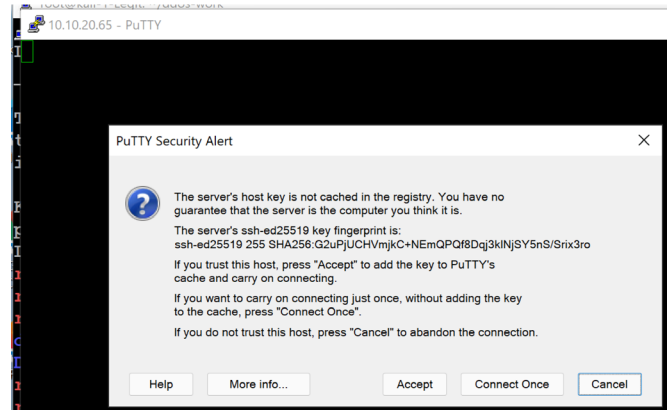
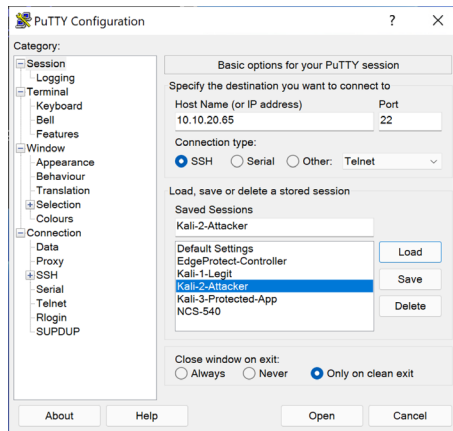
```
RP/0/RP0/CPU0:NCS1#
RP/0/RP0/CPU0:NCS1#show access-lists gtp
Mon Nov 1 18:51:13.106 UTC
ipv4 access-list gtp
 1000 permit udp any any eq 2152 capture (1175 matches)
 2000 permit ipv4 any any (4 matches)
 2001 permit icmp any any
RP/0/RP0/CPU0:NCS1#
```

Now SSH into the attacker system with the IP address of 10.10.20.65 and the following credentials:

Username: demo

Password: cisco123

If prompted with a security alert, click **“Accept.”** This happens because it is your first time connecting to this device securely.



Move into the ddos-work directory with the following command:

```
cd ddos-work
```

Now execute the script called ddos-udp-target.sh:

```
./ddos-udp-target.sh
```

The Kali Linux box will now generate traffic and send it to the protected object. Notice that the IP of traffic being sent is 15.15.1.1. This is the IP of the protected application server itself.

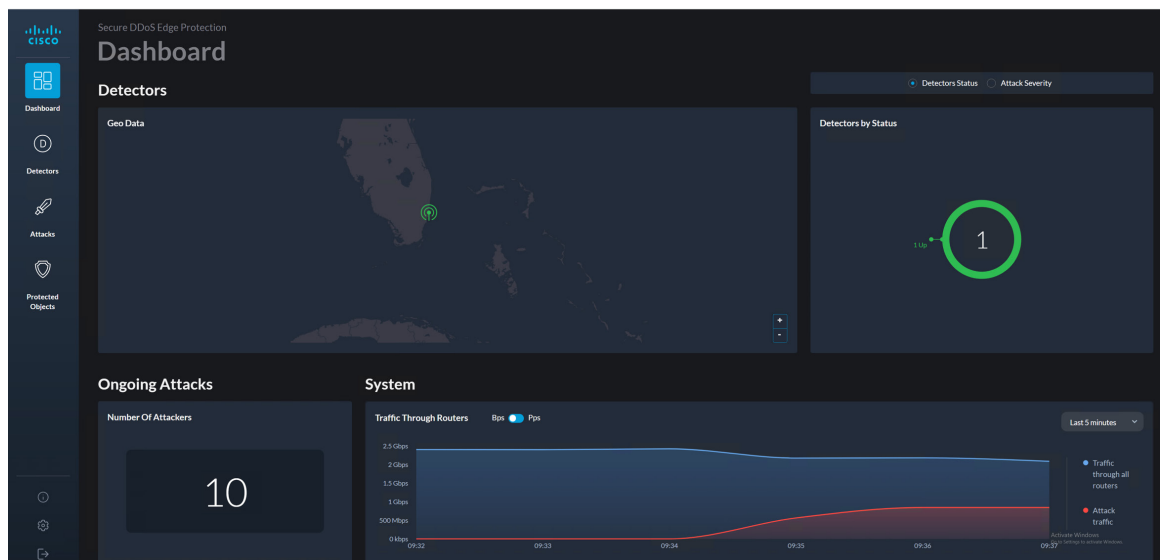


```
demo@kali-2-Attacker:~/ddos-work$
demo@kali-2-Attacker:~/ddos-work$ ls
ddos-tcp-Attack.sh ddos-udp-attack.sh
demo@kali-2-Attacker:~/ddos-work$ ./ddos-udp-attack.sh
HPING 15.15.1.1 (eth0 15.15.1.1): udp mode set (GTP Tunneling), 36 headers + 400
data bytes
hping in flood mode, no replies will be shown
```

Now let's go back to the controller GUI portal and review some of the widgets.

Open Chrome (or other browser) and connect to the Edge Protection Controller: <http://10.10.20.51.nip.io>.

For the email, enter admin@example.com, and for the password, use: 12345. Then click "Login."

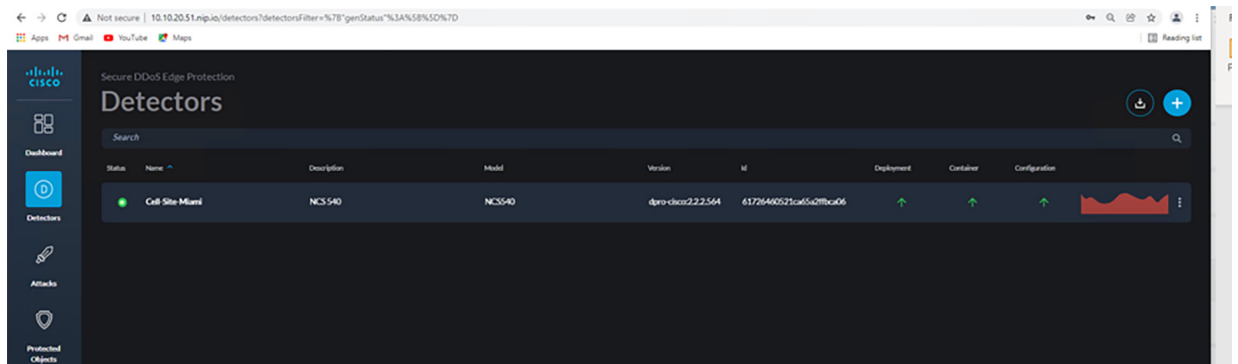


Click "Attack Severity."

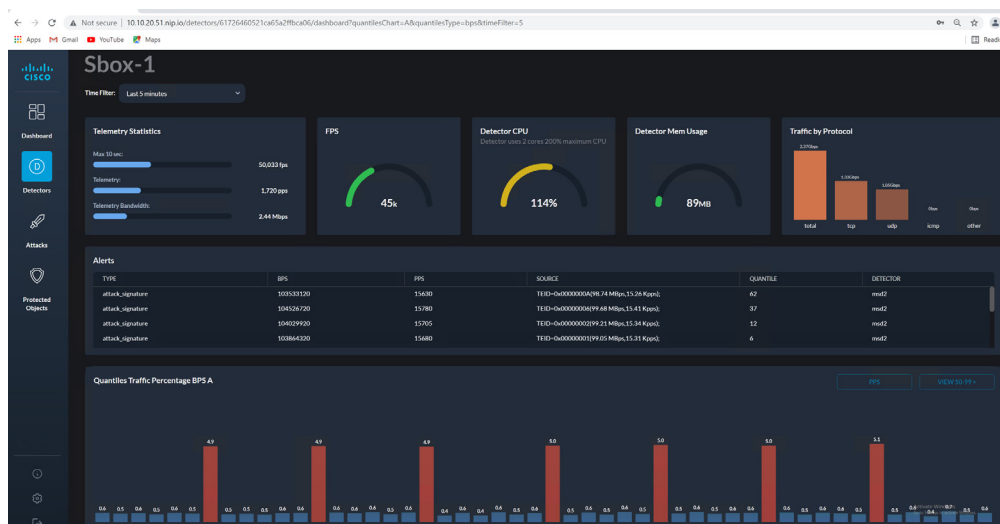
"Detectors by Attack Severity" is now red (if yellow, wait another minute), indicating that an attack has been detected.



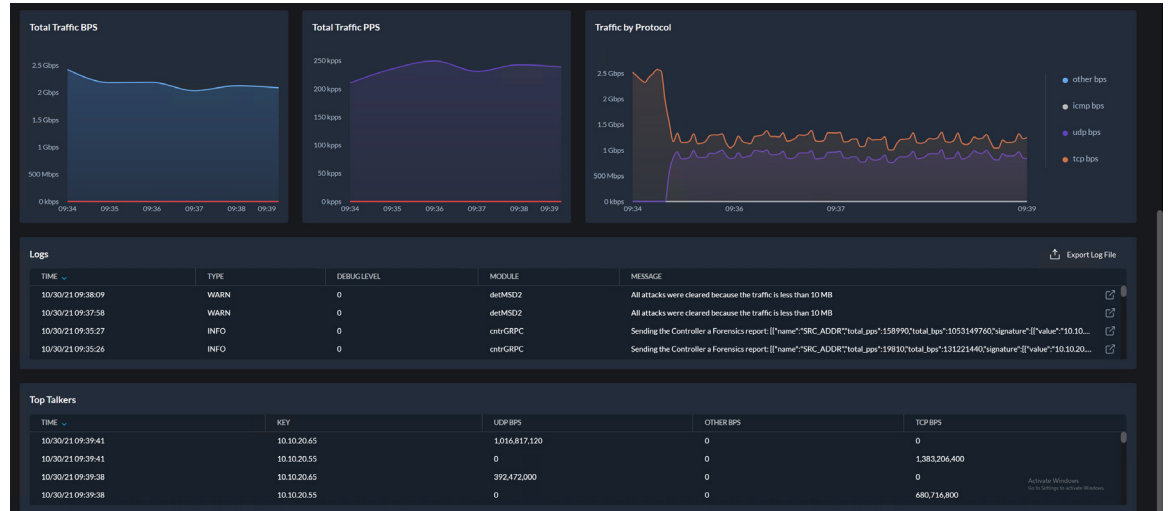
Click on the “1” and you will be redirected to the Detectors page. Notice that the graph is now red.




Click on the graph. Notice that, in the bar graph for “Quantiles Traffic Percentage,” we see some red quantiles. These are the quantiles where the attack has been detected.



Scroll down on your screen to see additional widgets that provide information about the network traffic.



In the "Traffic by Protocol" graph, we now see TCP and UDP traffic.



Below we now see logs indicating warning.

Logs

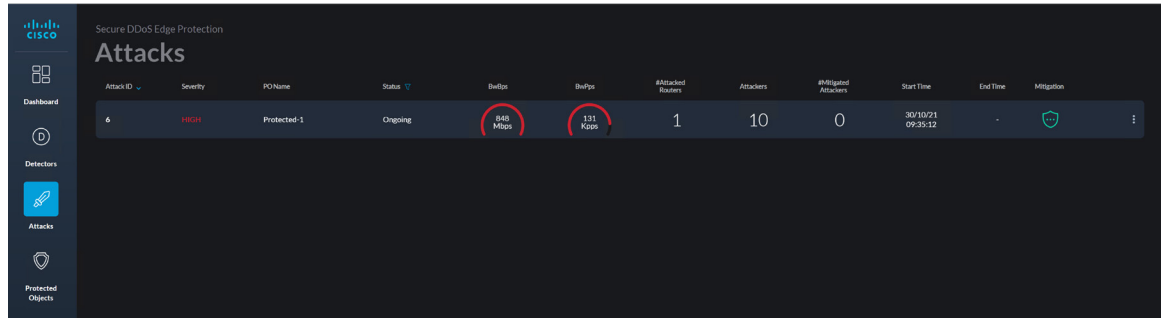
TIME	TYPE	DEBUG LEVEL	MODULE	MESSAGE
10/30/21 09:40:22	WARN	0	detMSD2	All attacks were cleared because the traffic is less than 10 MB
10/30/21 09:39:48	WARN	0	detMSD2	All attacks were cleared because the traffic is less than 10 MB
10/30/21 09:38:09	WARN	0	detMSD2	All attacks were cleared because the traffic is less than 10 MB
10/30/21 09:37:58	WARN	0	detMSD2	All attacks were cleared because the traffic is less than 10 MB

We now see the attacker machine as a top talker.

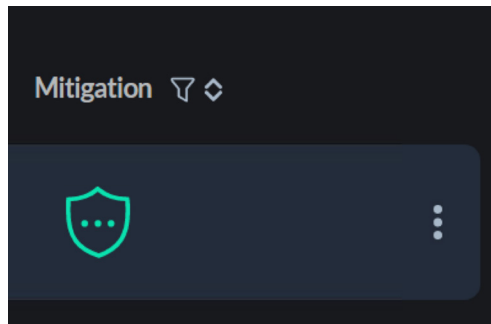
Top Talkers

TIME	KEY	UDP BPS	OTHER BPS	TCP BPS
10/30/21 09:41:02	10.10.20.65	888,775,200	0	0
10/30/21 09:41:02	10.10.20.55	0	0	1,125,964,800
10/30/21 09:40:59	10.10.20.65	987,009,120	0	0
10/30/21 09:40:59	10.10.20.55	0	0	1,421,798,400

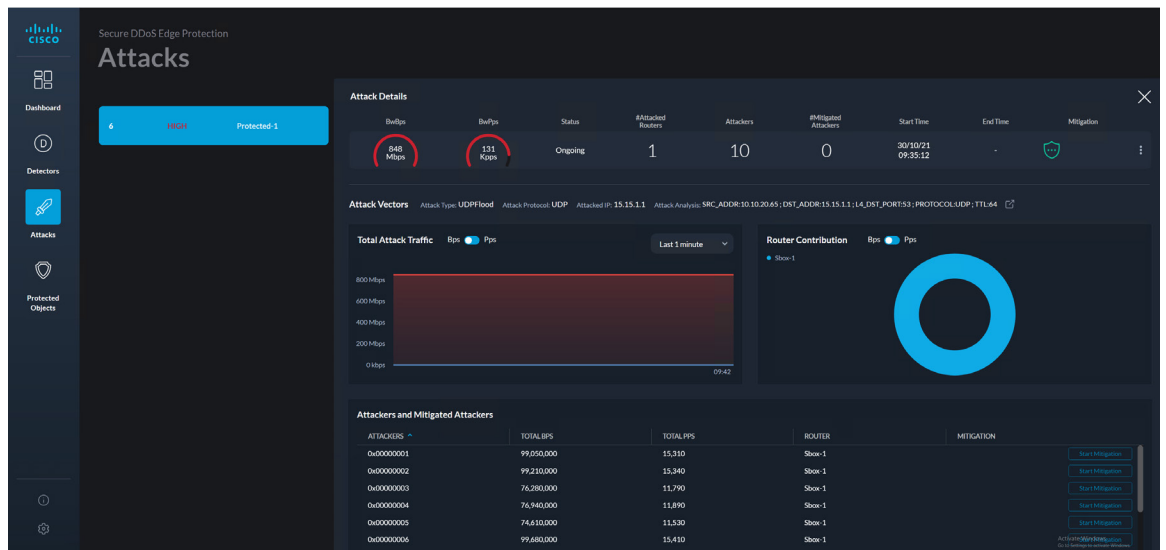
Click on the Attacks icon . We now see an active attack.



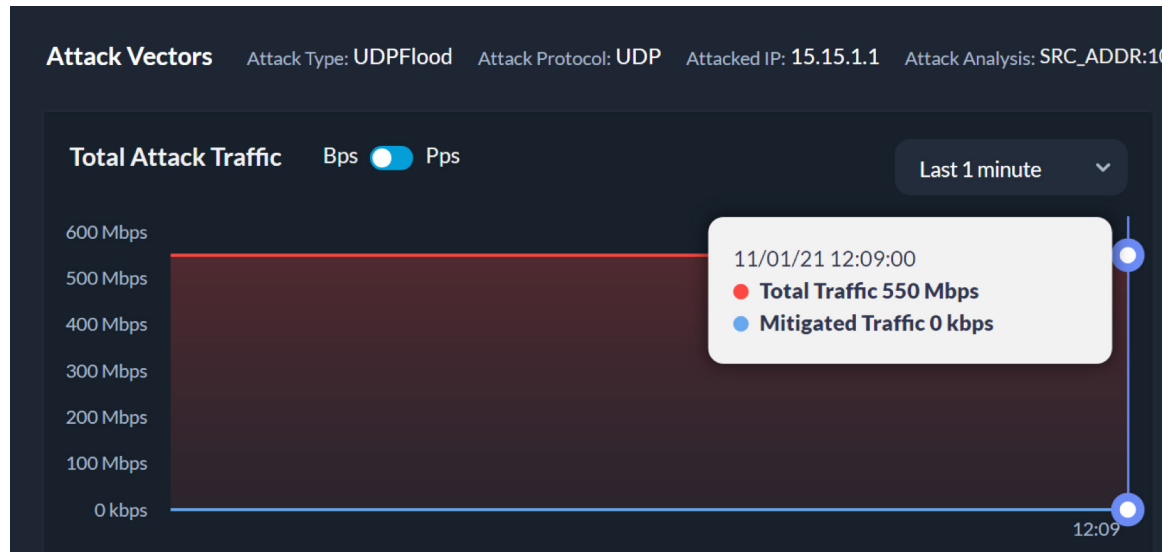
Note that the green shield in the upper right is flashing, indicating that an attack has been detected, but that the mitigation needs to be activated manually. This can also be set to automatic; however, mitigation would occur so quickly that the user would most likely miss viewing the attack. Therefore, for this lab, the mitigation has been set to user manual activation.




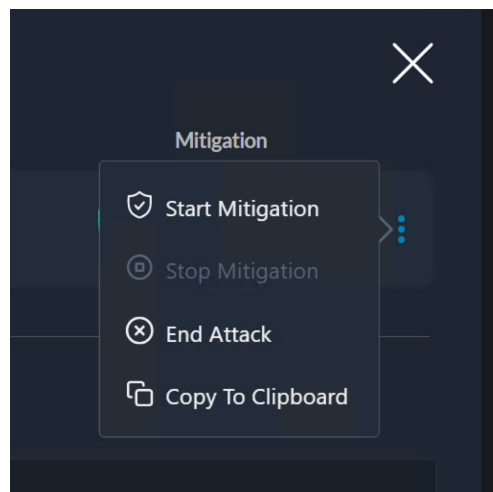
Click on the shield to view more details of the attack that has been detected. Analyzing this interface, we can see that there are ten (10) attackers and zero (0) have been mitigated. To the right of “Attackers and Mitigated Attackers,” we have the option of manually starting a mitigation on a particular attacker.



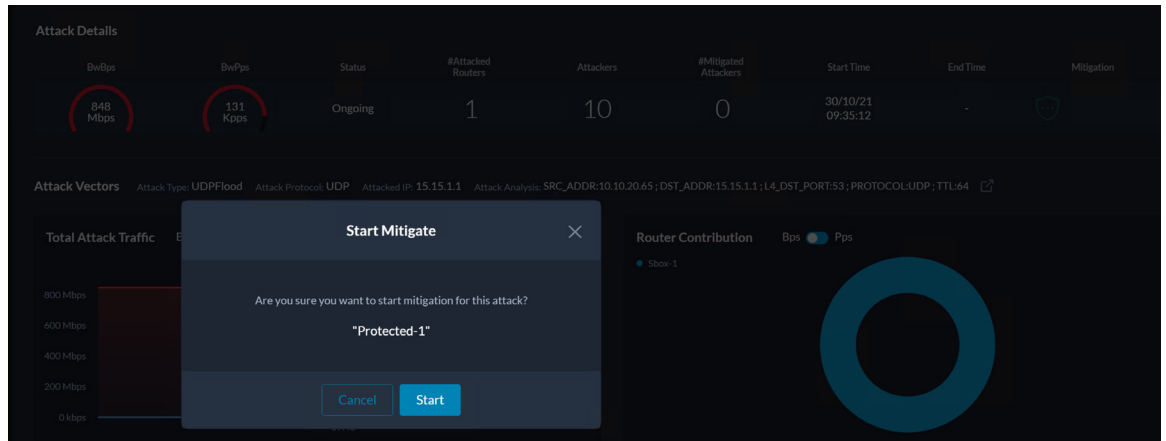
If you hover your mouse over the graph labeled “Total Attack Traffic,” you will see the amount of total attack traffic and how much of that traffic has been mitigated. As you can see, no traffic has been mitigated yet.




Click on the  icon in the upper right of the screen and click “Start Mitigation.”



Click the “Start” button. Now all attackers will be mitigated.



The green shield has stopped flashing, and we see that “# Mitigated Attackers” shows “10,” matching the number of attackers. This confirms that all attackers are now being mitigated. We can also stop a particular attacker from being mitigated by clicking the “Release Mitigation” button to the right of a particular attacker under the heading “Attackers and Mitigated Attackers” section.



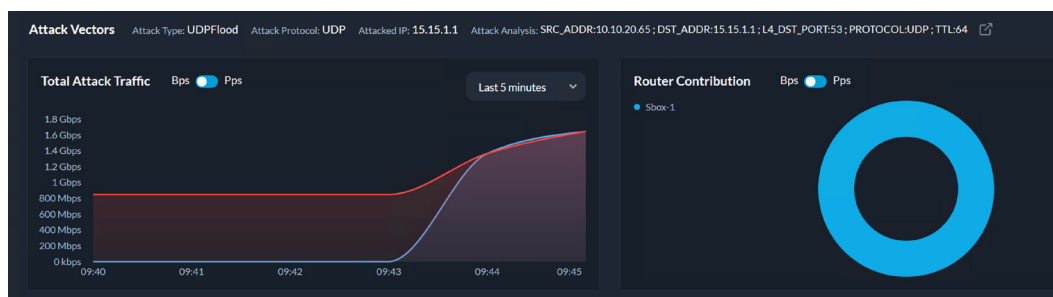
Attackers and Mitigated Attackers

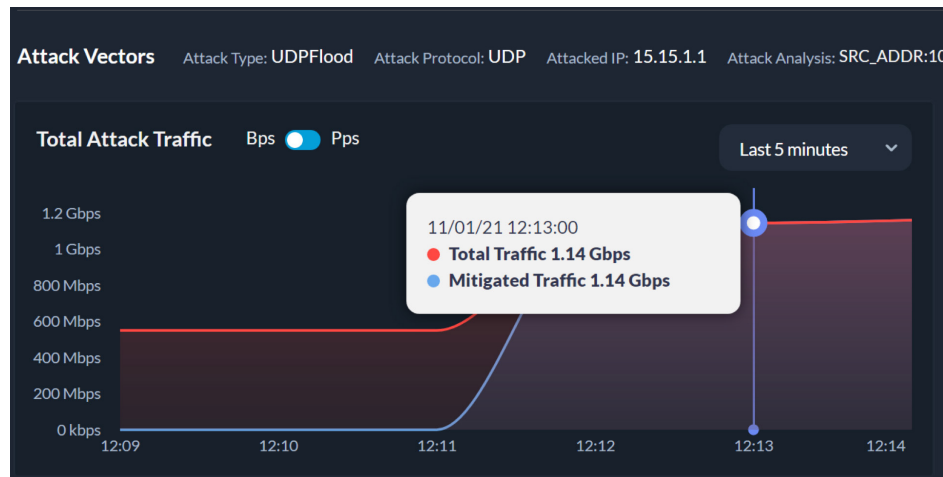
Attackers and Miti

ATTACKERS ^	TOTAL BPS	TOTAL PPS	ROUTER	MITIGATION
0x00000001	99,050,000	15,310	Sbox-1	Release Mitigation
0x00000002	99,210,000	15,340	Sbox-1	Release Mitigation
0x00000003	76,280,000	11,790	Sbox-1	Release Mitigation
0x00000004	76,940,000	11,890	Sbox-1	Release Mitigation
0x00000005	74,610,000	11,530	Sbox-1	Release Mitigation
0x00000006	99,680,000	15,410	Sbox-1	Release Mitigation

Activate Windows
Go to Settings to activate Windows.

Changing the view from “Last 1 Minute” to “Last 5 Minutes,” we can get a better view of the total attack traffic ramp-up.

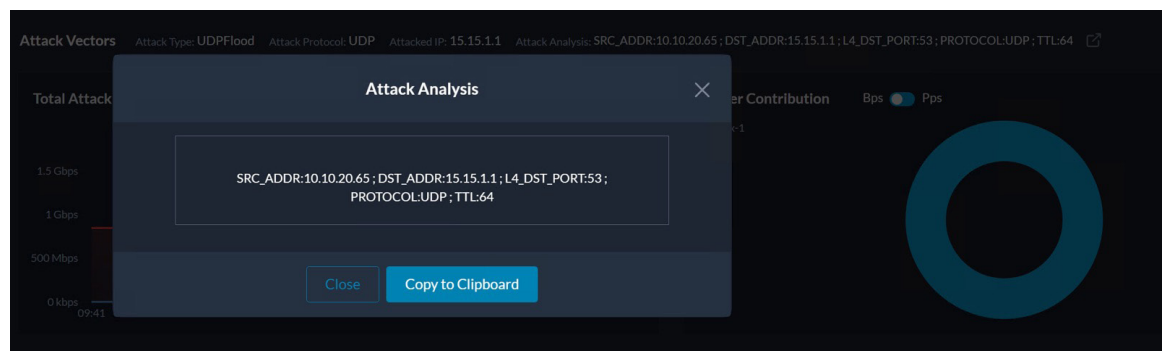




We can also view the attack signature that was created based on the attack traffic. The attack analysis shows a source address of 10.10.20.65, destination address of 15.15.1.1, protocol UDP port 53, and time to live (TTL) of 64.

```
Attack Analysis: SRC_ADDR:10.10.20.65 ; DST_ADDR:15.15.1.1 ; L4_DST_PORT:53 ; PROTOCOL:UDP ; TTL:64
```

Clicking on the box with the arrow pointing out (see image above) provides a cleaner view of the attack signature detected.




Now go back to your SSH session to the NCS 540 router (if closed, open it again) and look at the GTP access-list:

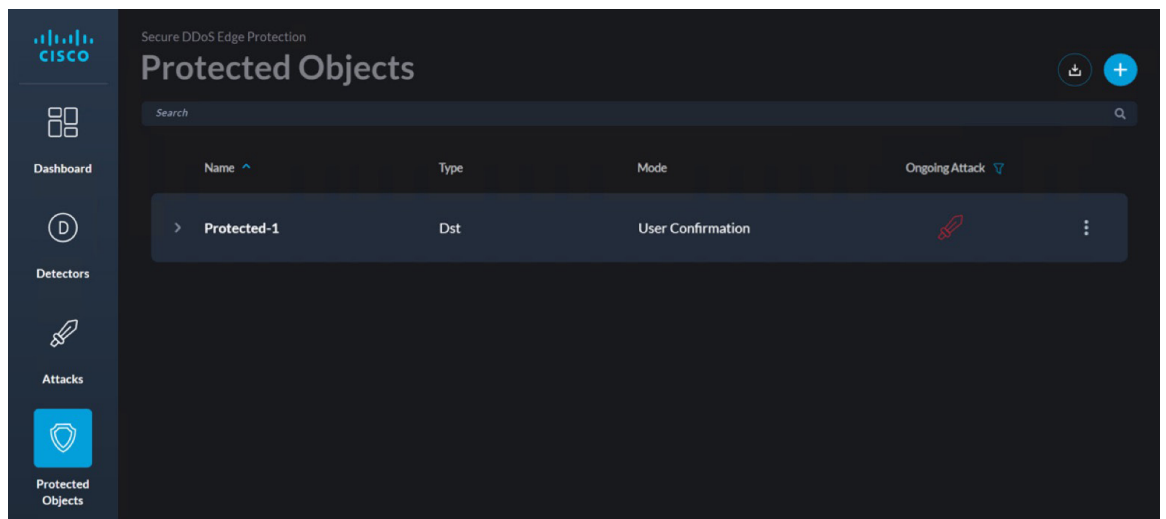
```
show access-list gtp
```

The GTP access-list has been updated with ten (10) new deny policies blocking the malicious traffic. Looking at each of the new lines in the ACL (below), each malicious GTP tunnel endpoint ID has been blocked (example: 0x8 GTP tunnel endpoint ID).


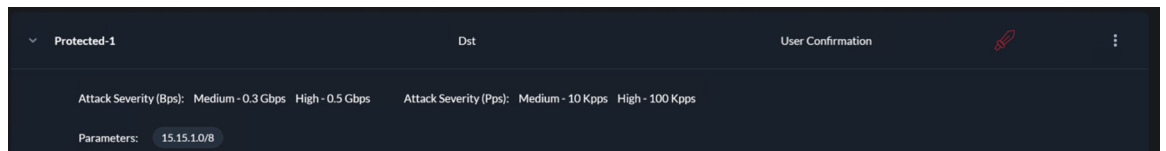
Click on the “1” in the “PO’s Under Attack” screen.




This will take you into a filter view of the Protected Objects page, showing only the PO that is under attack.

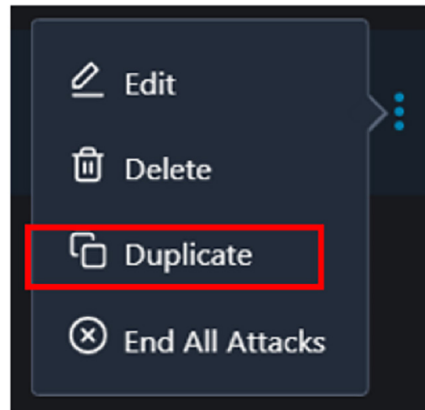
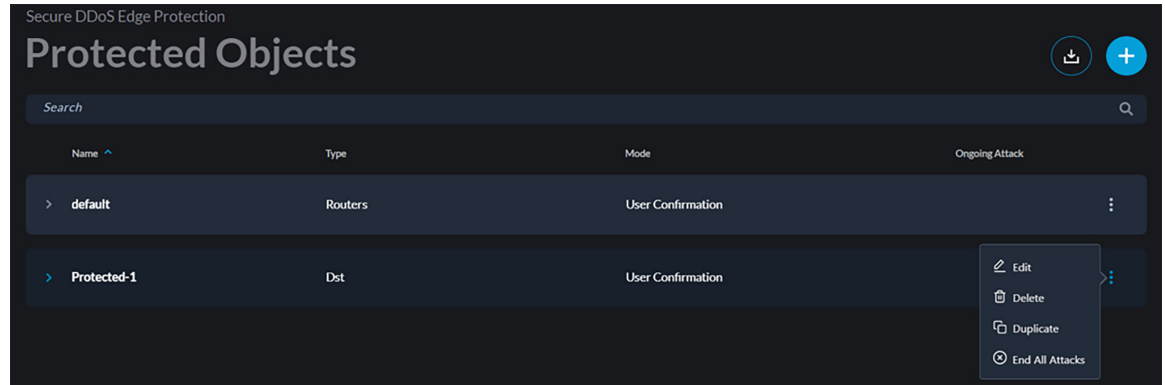


Click on the “>” symbol to view details of the PO.



Creating a New Protected Object (PO)

From the Protected Objects menu click on the icon  on the right of Protected-1



Click on Duplicate

Fill-in the form. For Name use DemoPO-2, click on the x next 15.15.1.0/8 under parameters heading to delete the entry and type 16.16.1.0/24 and press enter. Then click the Create button at the bottom.

New Protected Object

Name: DemoPO-2

Type: Dst

Mode: User Confirmation

Attack Severity:

Medium: 0.3 Gbps, 10 Kpps

High: 0.5 Gbps, 100 Kpps

Parameters: Enter Ip Address / Range
16.16.10/24 x

Cancel Create

You should now see 3 PO: default, Protected-1 and the newly created DemoPO-2

Secure DDoS Edge Protection

Protected Objects

Search

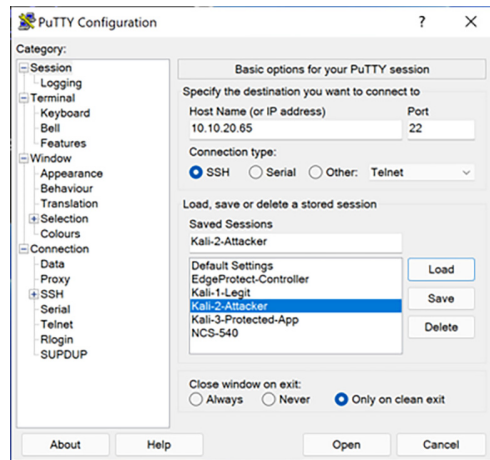
Name	Type	Mode	Ongoing Attack
> default	Routers	User Confirmation	
> DemoPO-2	Dst	User Confirmation	
> Protected-1	Dst	User Confirmation	

Now that we have the new PO created let's launch a second attack.

Open a second SSH session into the attacker system with the IP address of 10.10.20.65 and the following credentials:

Username: demo

Password: cisco123



Move into the ddos-work directory with the following command:

```
cd ddos-work
```

Now execute the script called ddos-tcp-target.sh:

```
./ddos-tcp-Attack.sh
```

The Kali Linux box will now generate a second attack.

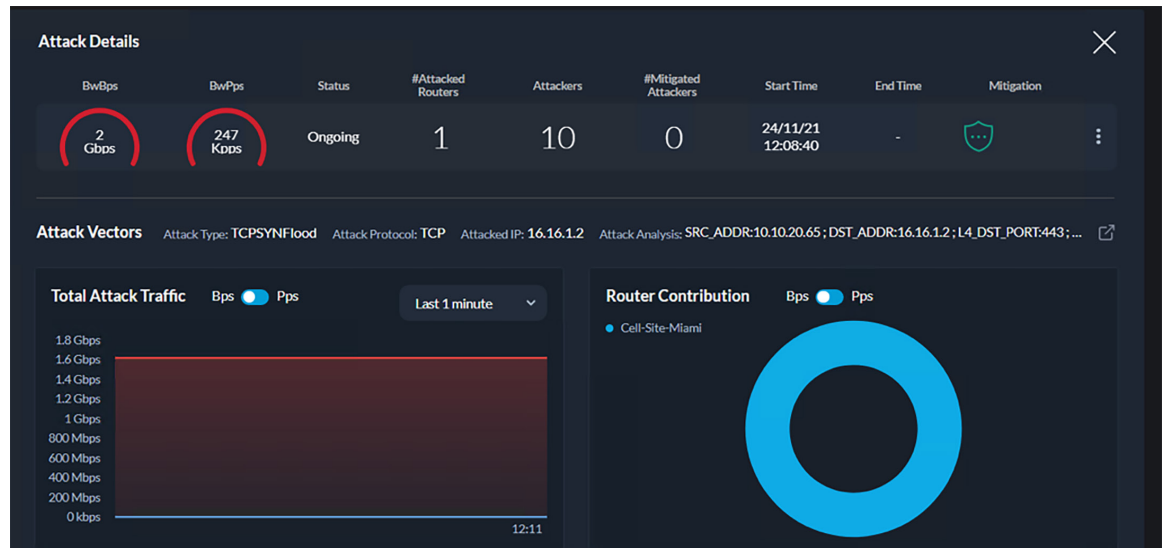
```
demo@kali-2-Attacker:~$
demo@kali-2-Attacker:~$
demo@kali-2-Attacker:~$
demo@kali-2-Attacker:~$ cd ddos-work/
demo@kali-2-Attacker:~/ddos-work$
demo@kali-2-Attacker:~/ddos-work$
demo@kali-2-Attacker:~/ddos-work$
demo@kali-2-Attacker:~/ddos-work$ ./ddos-tcp-Attack.sh
HPING 16.16.1.2 (eth0 16.16.1.2): S set (GTP Tunneling), 48 headers + 800 data bytes
hping in flood mode, no replies will be shown
```

Wait for about 15 seconds and then go back to the controller GUI portal and review some of the widgets.

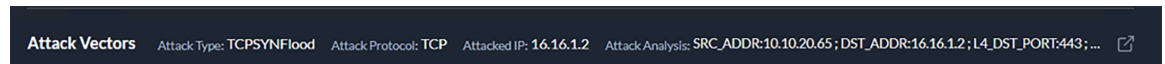
Review the Attack menu, we now see 2 attacks. The first attack shows 10 attackers and 10 mitigated attacks and a new attack, with a PO Name of DemoPO-2, shows 10 attackers and 0 mitigated.



Click on the DemoPO-2 attack to go into the Attack Dashboard



We can see that the attack is about 1.6Gbps(+ or -) in size and that the attack type is TCPSYNFlood



If you like, you can mitigate the attack as you did in the previous section and review how the ACL has been updated to protect against both attacks. Once mitigated the router ACL will be updated blocking the new attack

On the router command line type:

```
show access-list gtp
```

```
10.10.20.53 - PuTTY
RE/0/RP0/CPU0:NCS1#show access-lists
Tue Nov 30 19:33:30.897 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
11 deny ipv4 any any udf udf-gtp 0x20021b6 0xffffffff
12 deny ipv4 any any udf udf-gtp 0x20021b5 0xffffffff
13 deny ipv4 any any udf udf-gtp 0x20021b8 0xffffffff
14 deny ipv4 any any udf udf-gtp 0x20021af 0xffffffff
15 deny ipv4 any any udf udf-gtp 0x20021b7 0xffffffff
16 deny ipv4 any any udf udf-gtp 0x20021b2 0xffffffff
17 deny ipv4 any any udf udf-gtp 0x20021b1 0xffffffff
18 deny ipv4 any any udf udf-gtp 0x20021b4 0xffffffff
19 deny ipv4 any any udf udf-gtp 0x20021b3 0xffffffff
20 deny ipv4 any any udf udf-gtp 0x20021b0 0xffffffff
21 deny ipv4 any any udf udf-gtp 0x10009f 0xffffffff
1000 permit udp any any eq 2152 capture (1251 matches)
2000 permit ipv4 any any (7 matches)
2001 permit icmp any any
RE/0/RP0/CPU0:NCS1#
```

We can see that the ACL has been updated to protect the networks both attacks.



Step 4: Stopping the attack

Now we would like to stop the attack. Use CTRL+ A to stop the attack for both attacks launched.

```
demo@kali-2-Attacker:~/ddos-work$ ./ddos-udp-attack.sh
HPING 15.15.1.1 (eth0 15.15.1.1): udp mode set (GTP Tunneling), 36 headers + 400
data bytes
hping in flood mode, no replies will be shown
^C
--- 15.15.1.1 hping statistic ---
276760956 packets tramitted, 0 packets received, 100% packet loss
round-trip min/avg/max = 0.0/0.0/0.0 ms
demo@kali-2-Attacker:~/ddos-work$
```

```
demo@kali-2-Attacker:~/ddos-work$ ./ddos-top-Attack.sh
HPING 16.16.1.2 (eth0 16.16.1.2): S set (GTP Tunneling), 48 headers + 800 data b
ytes
hping in flood mode, no replies will be shown
^C
--- 16.16.1.2 hping statistic ---
99283962 packets tramitted, 0 packets received, 100% packet loss
round-trip min/avg/max = 0.0/0.0/0.0 ms
demo@kali-2-Attacker:~/ddos-work$
```

Please make sure that the attack script has stopped. Sometimes the script continues to run even though the terminal shows stopped. If you don't see the attack traffic stop in the Controller Dashboard, it is possible that the script is continuing to run and the process needs to be manually stopped (killed). The first step is to find the process ID.

To find the process ID of the hping4 attack script, use the following command:

```
command: ps -a
```

Look for the PID that says hping4. In the example below, it is 1127. Then kill that process by running the kill command:

```
kill -9 <PID>
```

In the example below:

```
kill -9 1127
```

```

root@kali-2-Attacker: ~
x86_64

The programs included with the Kali GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Kali GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
Last login: Mon Nov  1 15:49:48 2021 from 192.168.254.15
root@kali-2-Attacker:~#
root@kali-2-Attacker:~#
root@kali-2-Attacker:~#
root@kali-2-Attacker:~# ps -a
  PID TTY          TIME CMD
 1126 pts/1        00:00:00 sudo
 1127 pts/1        00:40:26 hping4
 1235 pts/1        00:00:00 tcpdump
 1258 pts/3        00:00:00 ps
root@kali-2-Attacker:~# kill -9 1127
root@kali-2-Attacker:~# ps -a
  PID TTY          TIME CMD
 1235 pts/1        00:00:00 tcpdump
 1289 pts/3        00:00:00 ps
root@kali-2-Attacker:~#

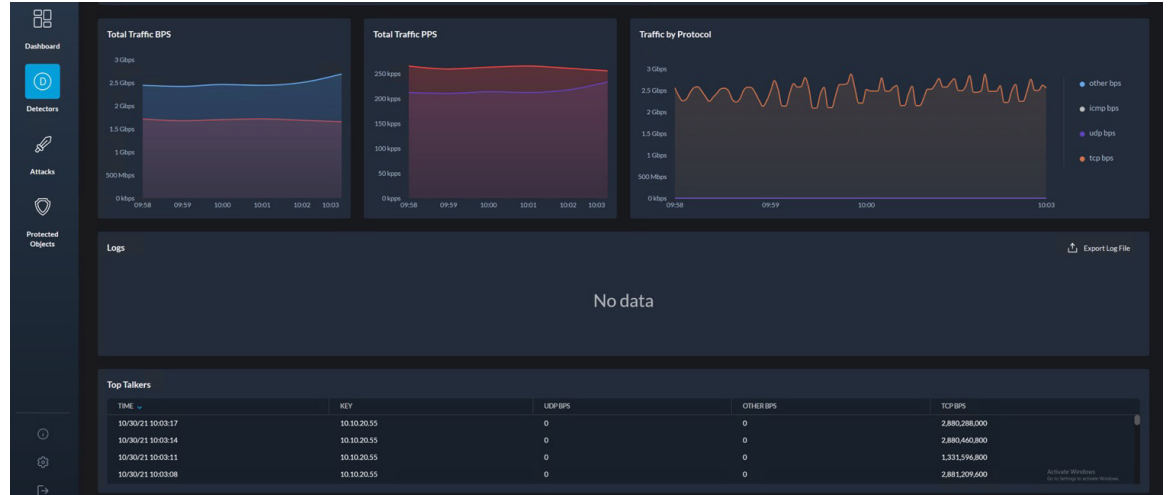
```

Once the attack has stopped, the controller will still report the router as under attack for a few minutes until it is certain that the attack has been terminated.

Going back to the Detectors page and clicking on the detector, we can now see that there no longer is any UDP traffic. The quantile bar graph is no longer showing any red bars, indicating that none of the quantiles are malicious.



Scrolling further down on the Detectors page, we also see that we no longer have any logs and that the top talkers are now all coming from our legitimate user.



Click on the Protected Objects icon



The screenshot shows the 'Protected Objects' configuration page. It includes a search bar and a table with the following columns: Name, Type, Mode, and Ongoing Attack. The table contains two entries:

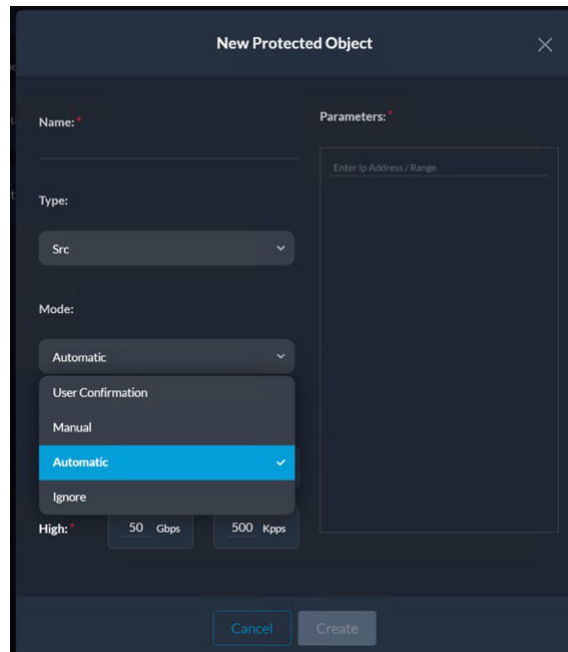
Name	Type	Mode	Ongoing Attack
default	Routers	User Confirmation	
Protected-1	Dst	User Confirmation	🔴

Notice that “Mode” is set to “User Confirmation.”

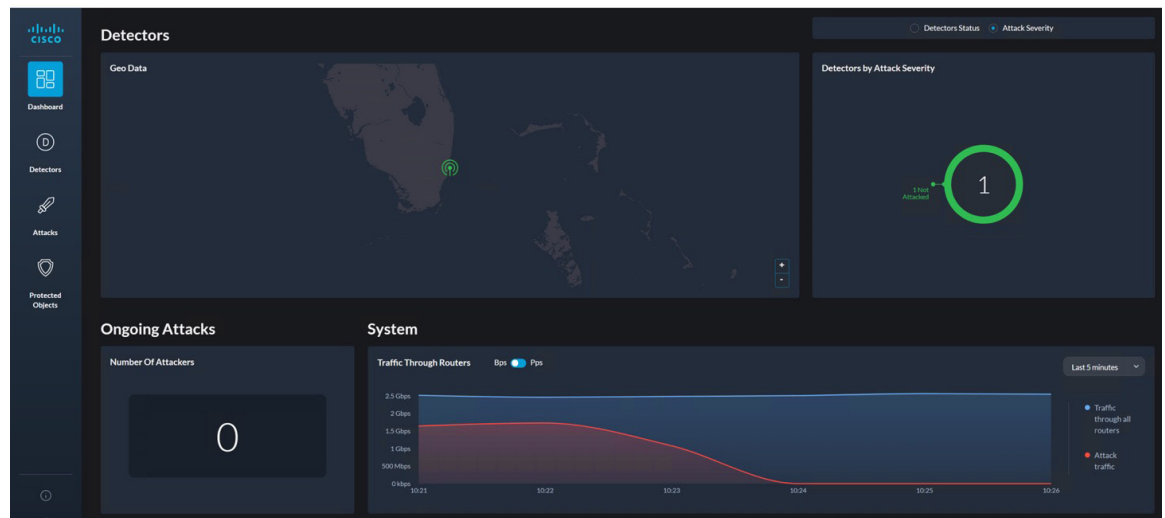
A close-up of the 'Mode' dropdown menu showing the following options:

- Mode
- User Confirmation
- User Confirmation

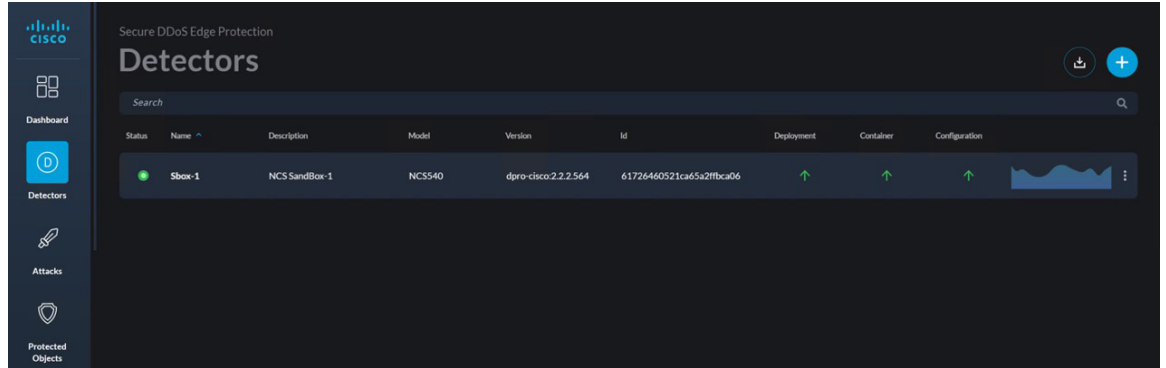
In a live environment, this would be changed to automatic, as we want to minimize the time between when the attack is detected and when it is stopped, or mitigated. When set to automatic, the attack is stopped within seconds.



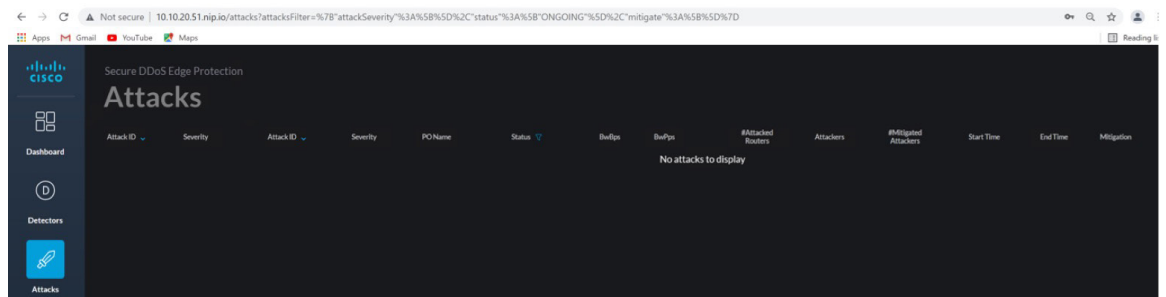
Let's go back to the dashboard. Be sure to wait at least five minutes after the attack is stopped. We now see that the router is longer under attack, and looking at "System > Traffic Through Routers," we see the red line (attack traffic) going down to zero (0).



If we go back to the Detectors screen, we see the graph on the right-hand side has changed from red to blue, indicating that there is no attack on this detector.

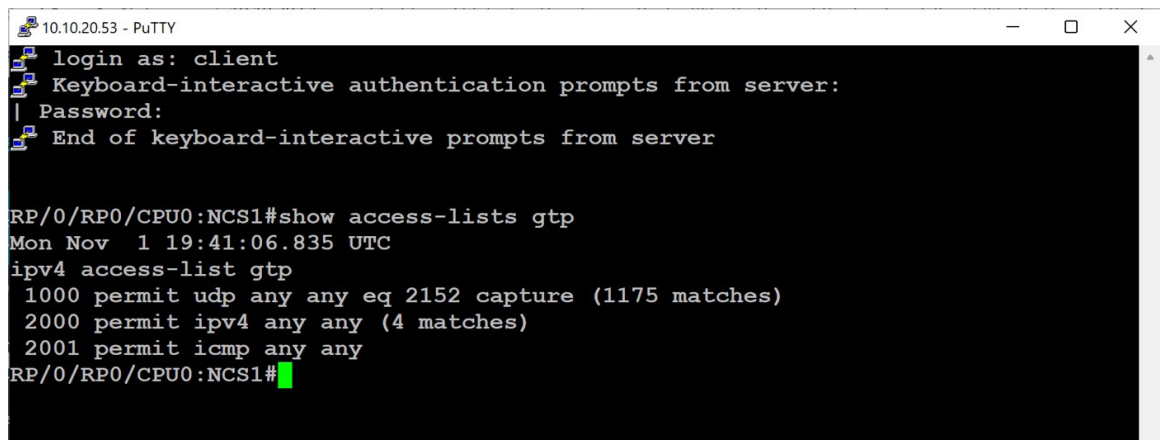


Going to the Attacks screen, we also see no attacks listed on this screen.




Finally, checking the Cisco NCS 540 GTP access-list, we see that the deny statements have been removed since there no longer is an attack on the network:

`show access-list gtp`



You have now completed the Cisco Secure DDoS Edge Protection Lab. Feel free to explore more of the user interface.



Resources

For more information about Cisco Secure DDoS Edge Protection, see:

- 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

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