

Date: December 5, 2024

Revision	Date	Changes
1.0	November 26, 2024	Initial release
1.1	December 5, 2024	Update the affected and fixed EOS versions

The CVE-ID tracking this issue: CVE-2024-6437

CVSSv3.1 Base Score: 5.8 (CVSS:3.1/ AV:N/AC:L/PR:N/UI:N/S:C/C:N/I:N/A:L)

Common Weakness Enumeration: CWE-1220: Insufficient Granularity of Access Control

This vulnerability is being tracked by BUG 962149

Description

On affected platforms running Arista EOS with one of the following features configured to redirect IP traffic to a next hop: policy-based routing (PBR), BGP Flowspec, or interface traffic policy -- certain IP traffic such as IPv4 packets with IP options may bypass the feature's **set nexthop** action and be slow-path forwarded (FIB routed) by the kernel as the packets are trapped to the CPU instead of following the redirect action's destination.

This issue was discovered externally and responsibly reported to Arista.

Vulnerability Assessment

Affected Software

Policy Based Routing (PBR)

EOS Versions

- 4.32.1F and below releases in the 4.32.x train
- 4.31.4M and below releases in the 4.31.x train
- 4.30.7M and below releases in the 4.30.x train
- 4.29.9M and below releases in the 4.29.x train
- 4.28.11M and below releases in the 4.28.x train
- 4.27.12M and below releases in the 4.27.x train
- 4.26.14M and below releases in the 4.26.x train
- 4.25.11M and below releases in the 4.25.x train
- 4.24.11M and below releases in the 4.24.x train
- 4.23.15M and below releases in the 4.23.x train
- 4.22.13M and below releases in the 4.22.x train
- 4.21.15M and below releases in the 4.21.x train

BGP Flowspec

EOS Versions

- 4.32.1F and below releases in the 4.32.x train
- 4.31.4M and below releases in the 4.31.x train
- 4.30.7M and below releases in the 4.30.x train
- 4.29.9M and below releases in the 4.29.x train
- 4.28.11M and below releases in the 4.28.x train
- 4.27.12M and below releases in the 4.27.x train
- 4.26.14M and below releases in the 4.26.x train
- 4.25.11M and below releases in the 4.25.x train
- 4.24.11M and below releases in the 4.24.x train
- 4.23.15M and below releases in the 4.23.x train
- 4.22.13M and below releases in the 4.22.x train
- Releases in the 4.21.x train between 4.21.3F and 4.21.15M

Interface Traffic Policy

EOS Versions

- 4.32.1F and below releases in the 4.32.x train
- 4.31.4M and below releases in the 4.31.x train
- 4.30.7M and below releases in the 4.30.x train
- 4.29.9M and below releases in the 4.29.x train
- 4.28.11M and below releases in the 4.28.x train
- Releases in the 4.27.x train between 4.27.2F and 4.27.12M

Affected Platforms

Policy Based Routing (PBR)

The following products **are** affected by this vulnerability when policy-based routing (PBR) is configured:

- Arista EOS-based products:
 - 7010 Series
 - 7010X Series
 - 7160 Series
 - 7050X/X2/X3/X4 Series
 - 7060X/X2/X4/X5/X6 Series
 - 7250X Series
 - 7260X/X3 Series
 - 7280E/R/R2/R3 Series
 - 7300X/X3 Series

- 7320X Series
- 7358X4 Series
- 7368X4 Series
- 7388X5 Series
- 7500E/R/R2/R3 Series
- 7800R3/R4 Series

BGP Flowspec

The following products **are** affected by this vulnerability when BGP Flowspec is configured:

- Arista EOS-based products:
 - 7280E/R/R2/R3 Series
 - 7500E/R/R2/R3 Series
 - 7800R3/R4 Series

Interface Traffic Policy

The following products **are** affected by this vulnerability when interface traffic policy is configured:

- Arista EOS-based products:
 - 7280E/R/R2/R3 Series
 - 7500E/R/R2/R3 Series
 - 7800R3/R4 Series

Not Affected

The following product versions and platforms **are not** affected by this vulnerability:

- Arista EOS-based products:
 - 710 Series
 - 720D Series
 - 720XP/722XPM Series
 - 750X Series
 - 7130 Series running EOS
 - 7150 Series
 - 7170 Series
 - AWE 5000 Series
 - AWE 7200R Series
 - CloudEOS
 - cEOS-lab

- vEOS-lab
- CloudVision eXchange, virtual or physical appliance
- Arista Wireless Access Points
- CloudVision CUE, virtual appliance or physical appliance
- CloudVision CUE cloud service delivery
- CloudVision Portal, virtual appliance or physical appliance
- CloudVision as-a-Service
- CloudVision AGNI
- Arista 7130 Systems running MOS
- Arista Converged Cloud Fabric and DANZ Monitoring Fabric (Formerly Big Switch Nodes for BCF and BMF)
- Arista Network Detection and Response (NDR) Security Platform (Formerly Awake NDR)
- Arista Edge Threat Management - Arista NG Firewall and Arista Micro Edge (Formerly Untangle)
- Arista NetVisor OS, Arista NetVisor UNUM, and Insight Analytics (Formerly Pluribus)

Required Configuration for Exploitation

In order to be vulnerable to CVE-2024-6437, **one** of the following **three** conditions must be met:

Policy Based Routing (PBR)

(1) A PBR policy must be configured with a rule which redirects to a next hop or set of next hops.

```
switch(config)#show policy-map type pbr
Service policy pmap1
  Configured on: Ethernet20/1
  Applied on:    Ethernet20/1
  10: Single match statement
    Match:
      0 permit ip any host 10.2.1.1
    Configured actions: set nexthop 10.20.1.1
    Active routing action:
      VRF default
        Route to nexthop 10.20.1.1 default
  20: Single match statement
    Match:
      0 permit ip any host 10.3.1.1
    Configured actions: set nexthop 10.20.2.1
    Active routing action:
      VRF default
        Route to nexthop 10.20.2.1 default
  30: Single match statement
```

```
Match:
  0 permit ip 10.50.1.0/24 any
Configured actions: set nexthop 10.20.3.1
Active routing action:
VRF default
  Route to nexthop 10.20.3.1 default
40: Single match statement
Match:
  0 permit ip any any
Configured actions: set nexthop 10.20.4.1
Active routing action:
VRF default
  Route to nexthop 10.20.4.1 default
```

With this configuration, any packet that does not match the PBR match rules will fall through and match the "default" match-all rule and should get redirected to the next hop **10.20.4.1**.

BGP Flowspec

OR

(2) A BGP Flowspec must be configured with a rule which redirects to a next hop or set of next hops.

```
switch#show flow-spec ipv4
Flow specification rules for VRF default
Configured on: Ethernet20/1
Applied on: Ethernet20/1
Flow-spec rule: 10.100.0.0/16;*;
Rule identifier: 1
Matches:
  Destination prefix: 10.100.0.0/16
Actions:
  Redirect: VRF default
           Route via next hop 10.20.4.1
Status:
  Installed: yes
  Counter: 0 packets, 0 bytes
```

With this configuration, all traffic ingressing **Ethernet20/1** with destination addresses in the **10.100.0.0/16** subnet should get redirected to the next hop **10.20.4.1**.

Interface Traffic Policy

OR

(3) An interface traffic policy must be configured with a rule that redirects to a next hop or set of next hops.

```
switch#show traffic-policy interface
Traffic policy foo
  Configured on input of interfaces: Ethernet20/1
  Applied on input of interfaces for IPv4 traffic: Ethernet20/1
  Applied on input of interfaces for IPv6 traffic:
  Applied on input of interfaces for MAC traffic:
  Configured on output of interfaces:
  Applied on output of interfaces for IPv4 traffic:
  Applied on output of interfaces for IPv6 traffic:
  Total number of rules configured: 3
    match rule1 ipv4
      Destination prefix: 10.100.0.0/16
      Actions: Redirect next hop 10.20.4.1
                Active routing action:
                  VRF default
                    Route via next hop 10.20.4.1 VRF default
    match ipv4-all-default ipv4
    match ipv6-all-default ipv6
```

With this configuration, all traffic ingressing **Ethernet20/1** with destination addresses in the **10.100.0.0/16** subnet should get redirected to the next hop **10.20.4.1**.

Indicators of Compromise

Exploitation of CVE-2024-6437 will lead to IP traffic following the FIB (regular routing table) route instead of the next hop defined by the **set nexthop** action.

For the examples to follow, let's suppose the following FIB route is configured in the system, and the ARP entry is resolved.

```
switch#show ip route 10.100.1.1/32

VRF: default
Source Codes:
  C - connected, S - static, K - kernel,
  O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
```

```
E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,  
N2 - OSPF NSSA external type2, B - Other BGP Routes,  
B I - iBGP, B E - eBGP, R - RIP, I L1 - IS-IS level 1,  
I L2 - IS-IS level 2, O3 - OSPFv3, A B - BGP Aggregate,  
A O - OSPF Summary, NG - Nexthop Group Static Route,  
V - VXLAN Control Service, M - Martian,  
DH - DHCP client installed default route,  
DP - Dynamic Policy Route, L - VRF Leaked,  
G - gRIBI, RC - Route Cache Route,  
CL - CBF Leaked Route
```

```
S      10.100.1.1/32  
      directly connected, Ethernet10/1
```

Policy Based Routing (PBR)

The following example uses the PBR configuration from condition (1) in [Required Configuration for Exploitation](#).

Suppose the flow (source IP: **10.50.1.1** destination IP: **10.100.1.1**) ingresses **Ethernet20/1**, which has a PBR policy **pmap1** applied. The flow matches rule 40.

```
switch(config)#show policy-map type pbr  
Service policy pmap1  
  Configured on: Ethernet20/1  
  Applied on:    Ethernet20/1  
...  
40: Single match statement  
  Match:  
    0 permit ip any any  
  Configured actions: set nexthop 10.20.4.1  
  Active routing action:  
  VRF default  
    Route to nexthop 10.20.4.1 default
```

The expectation is that the flow gets redirected to **10.20.4.1**; instead, traffic gets redirected to the destination in the routing table, **Ethernet10/1**.

BGP Flowspec

The following example uses the BGP Flowspec configuration from condition (2) in [Required Configuration for Exploitation](#).

Suppose the flow (source IP: **10.50.1.1** destination IP: **10.100.1.1**) ingresses **Ethernet20/1**, which has BGP Flowspec policy applied to **Ethernet20/1**'s VRF. The flow matches **Flow-spec rule: 10.100.0.0/16;***.

```
switch#show flow-spec ipv4
Flow specification rules for VRF default
Configured on: Ethernet20/1
Applied on: Ethernet20/1
Flow-spec rule: 10.100.0.0/16;*;
Rule identifier: 1
Matches:
  Destination prefix: 10.100.0.0/16
Actions:
  Redirect: VRF default
           Route via next hop 10.20.4.1
Status:
  Installed: yes
  Counter: 0 packets, 0 bytes
```

The expectation is that the flow gets redirected to **10.20.4.1**; instead, traffic gets redirected to the destination in the routing table, **Ethernet10/1**.

Interface Traffic Policy

The following example uses the interface traffic policy configuration from condition (3) in [Required Configuration for Exploitation](#).

Suppose the flow (source IP: **10.50.1.1** destination IP: **10.100.1.1**) ingresses **Ethernet20/1**, which has an interface traffic policy, **foo**, applied.

The flow matches **rule1**.

```
switch#show traffic-policy interface
Traffic policy foo
  Configured on input of interfaces: Ethernet20/1
  Applied on input of interfaces for IPv4 traffic: Ethernet20/1
  Applied on input of interfaces for IPv6 traffic:
  Applied on input of interfaces for MAC traffic:
  Configured on output of interfaces:
  Applied on output of interfaces for IPv4 traffic:
  Applied on output of interfaces for IPv6 traffic:
  Total number of rules configured: 3
    match rule1 ipv4
```

```

Destination prefix: 10.100.0.0/16
Actions: Redirect next hop 10.20.4.1
         Active routing action:
           VRF default
           Route via next hop 10.20.4.1 VRF default
match ipv4-all-default ipv4
match ipv6-all-default ipv6

```

The expectation is that the flow gets redirected to **10.20.4.1**; instead, traffic gets redirected to the destination in the routing table, **Ethernet10/1**.

Common Symptoms

Increments in the egress interface counters for the FIB route destination may indicate that traffic is not getting redirected according to the **set nexthop** destination.

```

switch#show interface counters | nz

```

Port	InOctets	InUcastPkts	InMcastPkts	InBcastPkts
Et20/1	108160	1040	0	0
Ma1	512502	4220	331	8


```

switch#show interface counters | nz

```

Port	OutOctets	OutUcastPkts	OutMcastPkts	OutBcastPkts
Et10/1	108160	1040	0	0
Ma1	550930	3017	0	0

For all three vulnerabilities on 7280E/R/R2/R3, 7500E/R/R2/R3, and 7800 devices, the **show cpu counters queue | nz** command will also display increments in packet counts for traffic destined to the CPU. Increments in the **CoppSystemL3Ttl1IpOptUcast** counter may indicate that a lot of TTL0/1 and/or IP options unicast traffic is being trapped to the CPU. All such traffic will bypass the **set nexthop** action in all of the affected features.

```

switch#show cpu counters queue | nz
<Chip Name>:

```

CoPP Class	Queue	Pkts	Octets
DropPkts	DropOctets		

```

Aggregate

```

```

-----
CoppSystemL3Ttl1IpOptUcast
TC0                1040                108160                0
0
    
```

For all three vulnerabilities, on

- 7010X Series
- 7160 Series
- 7050X/X2/X3/X4 Series
- 7060X/X2/X4/X5/X6 Series
- 7250X Series
- 7260X/X3 Series
- 7300X/X3 Series
- 7320X Series
- 7358X4 Series
- 7368X4 Series
- 7388X5 Series

The **L3 Slow Path** counter in **show cpu counters queue | nz** indicates that traffic is being trapped to the CPU for slow path forwarding. All such traffic will bypass the **set nexthop** action in all of the affected features.

```

switch#show cpu counters queue | nz
-----
                                <Linecard Name>
-----
Queue                            Counters/pkts                Drops/pkts
-----
L3 Slow Path                        1040                          0
    
```

Mitigation

For all affected systems, the suggested mitigation for all three affected features is to drop all IPv4 options traffic via the **ip software forwarding options action drop**, available in 4.32.2F and later releases in the 4.32 train, 4.31.5M and later releases in the 4.31 train, and 4.30.8M and later releases in the 4.30 train. The command installs an **iptables** rule that drops all IPv4 options traffic in the filter table of the **FORWARD** chain.

```
switch(config)#ip software forwarding options action drop

# Below is shown to illustrate what the rule does. This is not a command that needs to
be run.

switch(config)#bash sudo iptables -vnL EOS_FORWARD
Chain EOS_FORWARD (1 references)
pkts bytes target      prot opt in      out     source        destination
  0     0 DROP          all  --  *       *       0.0.0.0/0     0.0.0.0/0
u32 ! "0x0>>0x18=0x45"
  0     0 REJECT        all  --  *       fwd+    0.0.0.0/0     0.0.0.0/0
u32 ! "0x0>>0x18=0x45" reject-with icmp-admin-prohibited
  0     0 DROP          all  --  *       ma+     0.0.0.0/0     0.0.0.0/0
  0     0 ACCEPT        all  --  *       *       !127.0.0.0/8  !127.0.0.0/8
```

Additionally, in 7280R3, 7500R3, and 7800R3 systems, the **system-rule overriding-action redirect** command (present in EOS-4.28.0F and newer releases) can be used to allow for all of the affected features' **set nexthop** action to take precedence over the system ACL's trap action to CPU. See [TCAM redirect action overriding system rules - TOI](#) for more information.

Resolution

The recommended resolution is to upgrade to a remediated software version that contains the **ip software forwarding options action drop** CLI command, and configure the command at your earliest convenience. Arista recommends customers move to the latest version of each release that contains all the fixes listed below. For more information about upgrading see [EOS User Manual: Upgrades and Downgrades](#)

CVE-2024-6437 has been fixed in the following releases:

- 4.32.2F and later releases in the 4.32.x train
- 4.31.5M and later releases in the 4.31.x train
- 4.30.8M and later releases in the 4.30.x train
- 4.29.10M and later releases in the 4.29.x train

Hotfix

No hotfix exists for this issue.

For More Information

If you require further assistance, or if you have any further questions regarding this security notice, please contact the Arista Networks Technical Assistance Center (TAC) by one of the

following methods:

Open a Service Request

Contact information needed to open a new service request may be found at:

<https://www.arista.com/en/support/customer-support>