Defining DDoS Mitigation Policies using BGP FlowSpec

Vinit Jain
CCIE# 22854
Twitter - @vinugenie

CP-1015

Agenda

- DDOS Overview
  - Requirements and Customer Feedback
- BGP Flowspec Overview
- Configuration
- Demo

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Public

BGP FLowSpec

DDoS Attacks

- ❑ Distributed denial-of-service (DDoS) attacks target network infrastructures or computer services by sending overwhelming number of service requests to the server from many sources.

- ❑ Server resources are used up in serving the fake requests resulting in denial or degradation of legitimate service requests to be served

- ❑ Addressing DDoS attacks
  - Detection – Detect incoming fake requests
  - **Mitigation**
    - Diversion – Send traffic to a specialized device that removes the fake packets from the traffic stream while retaining the legitimate packets
    - Return – Send back the clean traffic to the server

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Public

Remote Triggered Black Hole Filtering

Major Internet Outages

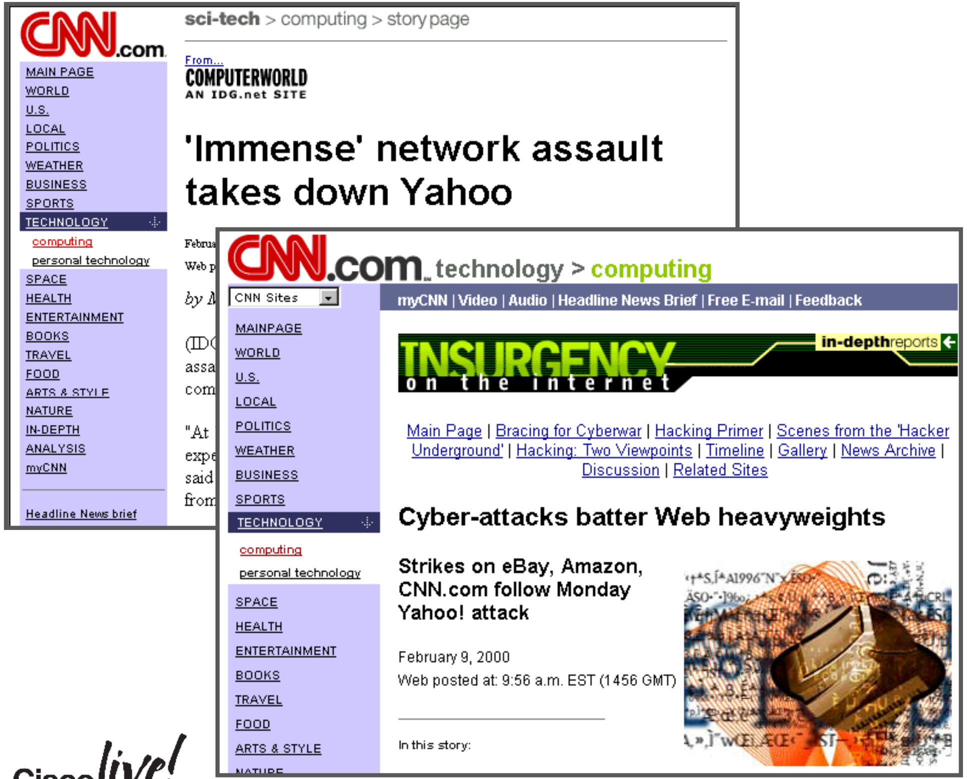

CNN.com

MAIN PAGE
WORLD
U.S.
LOCAL
POLITICS
WEATHER
BUSINESS
SPORTS
TECHNOLOGY
computing
personal technology
SPACE
HEALTH
ENTERTAINMENT
BOOKS
TRAVEL
FOOD
ARTS & STYLE
NATURE
IN-DEPTH
ANALYSIS
myCNN

Headline News brief

sci-tech > computing > story page

From...
COMPUTERWORLD
AN IDG.net SITE

'Immense' network assault takes down Yahoo

Februa
Web p
by M

(IDG
assa
com

"At
expe
said
from

CNN.com technology > computing

CNN Sites

myCNN | Video | Audio | Headline News Brief | Free E-mail | Feedback

MAINPAGE
WORLD
U.S.
LOCAL
POLITICS
WEATHER
BUSINESS
SPORTS
TECHNOLOGY
computing
personal technology
SPACE
HEALTH
ENTERTAINMENT
BOOKS
TRAVEL
FOOD
ARTS & STYLE
NATURE

Main Page | Bracing for Cyberwar | Hacking Primer | Scenes from the 'Hacker Underground' | Hacking: Two Viewpoints | Timeline | Gallery | News Archive | Discussion | Related Sites

Cyber-attacks batter Web heavyweights

Strikes on eBay, Amazon, CNN.com follow Monday Yahoo! attack

February 9, 2000
Web posted at: 9:56 a.m. EST (1456 GMT)

In this story:

YAHOO!

amazon.com

CNN.com

ebay™

E*TRADE℠

Remote Triggered Black Hole Filtering

The *Exodus* Requirement

> ***"We need a tool to drop packets based on source IP address that can be pushed out to over 60 routers with in 60 seconds, be longer than a thousand lines, be modified on the fly, and work in all your platforms filtering at line rate."***

Provided by Engineers at Exodus during the Feb 2000 DOS Post Mortem

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Public

BGP FlowSpec

Web Server

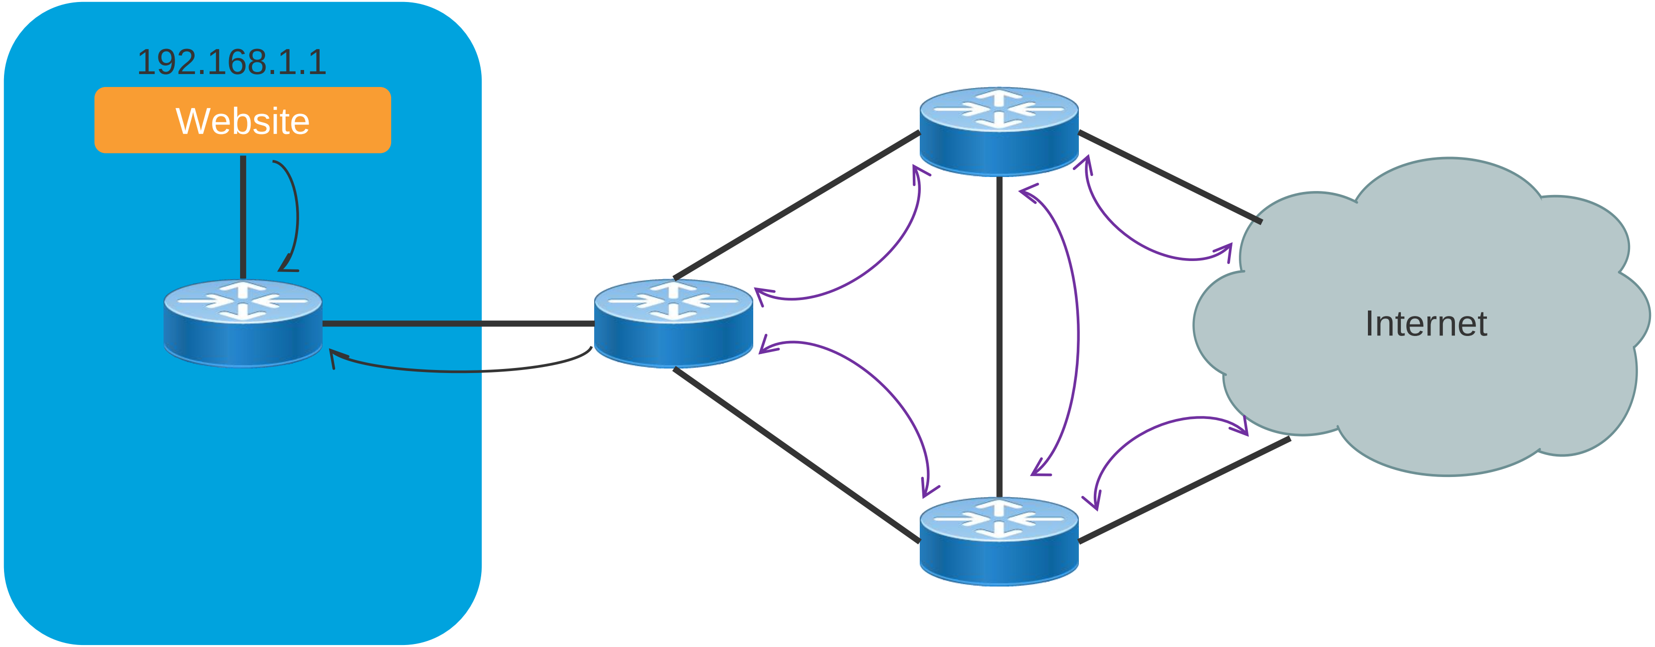

BGp FlowSpec

DDoS Attack

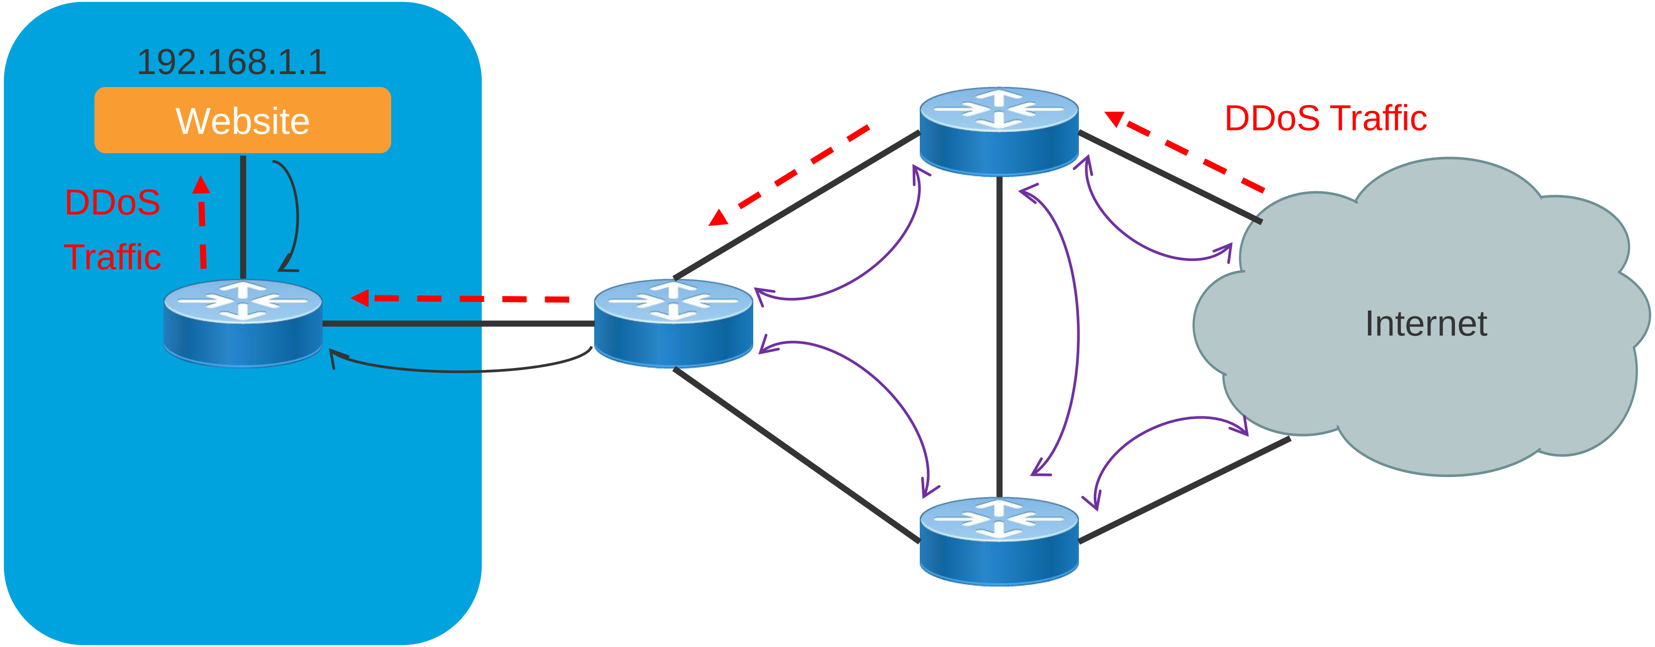

BGP FlowSpec

Black Hole Community Provided by Provider

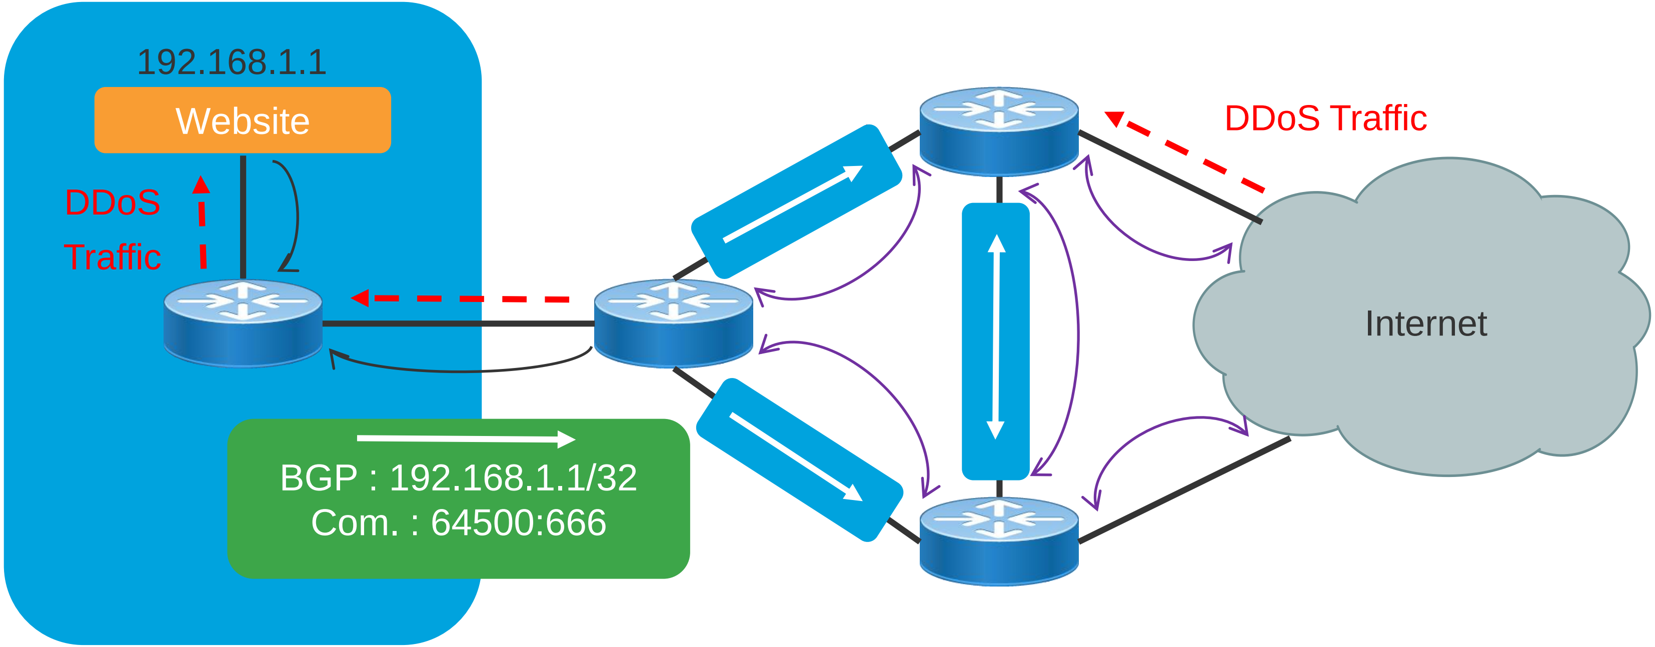

BGP : 192.168.1.1/32
Com. : 64500:666

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Public

BGP FlowSpec

Black Hole Community Provided by Provider

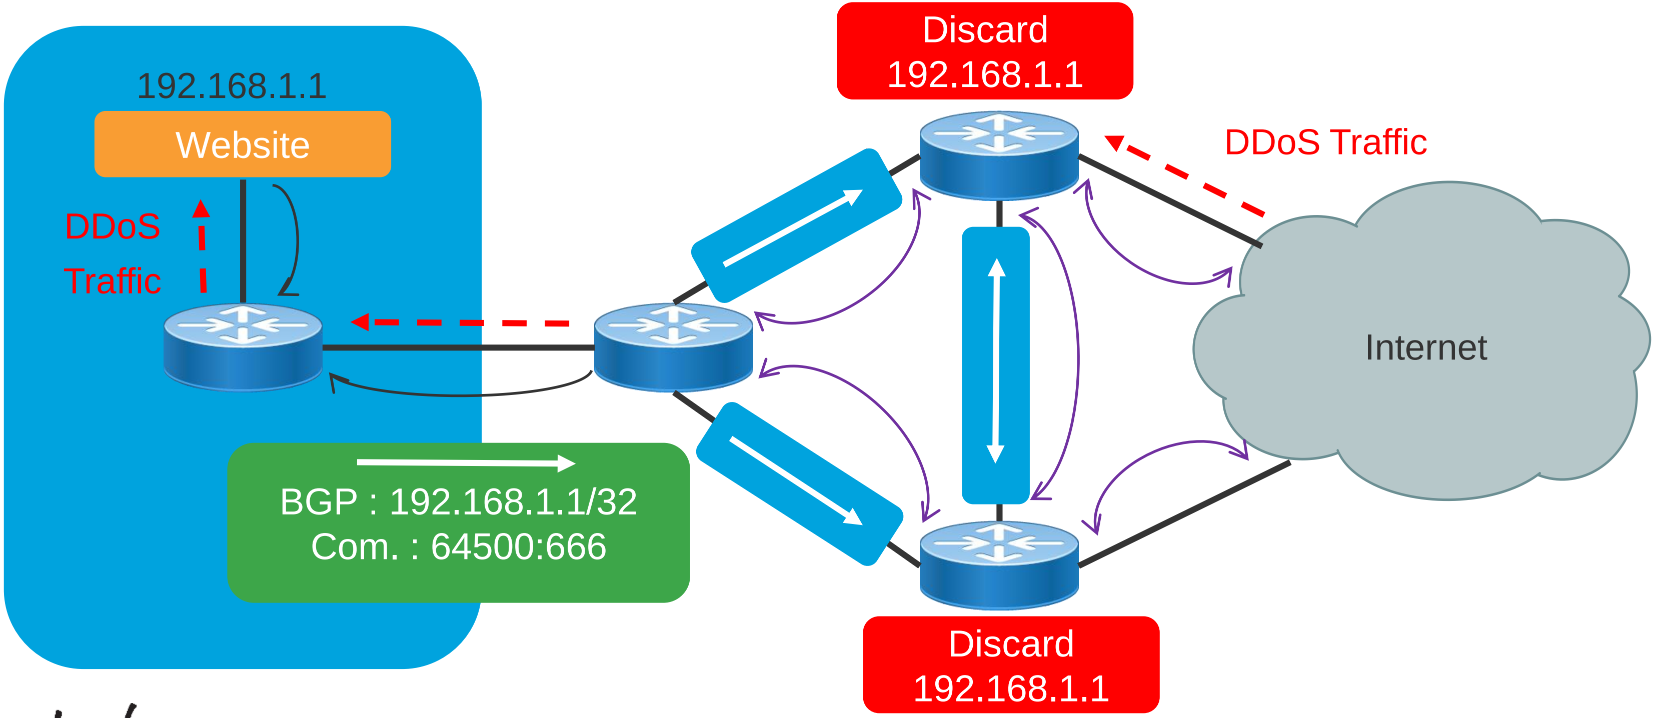

BGP FlowSpec

Drawback of RTBH

- Great, I have my website back online !
  - □ No more DDoS traffic on my network
  - □ But no more traffic at all on my website....

- Well, maybe it was not the solution I was looking for....

BGP FlowSpec

Policy Based Routing

- Identification of DDoS traffic: based around a conditions regarding MATCH statements
  - Source/Destination address
  - Protocol
  - Packet size
  - Etc...
- Actions upon DDoS traffic □ Discard
  - Logging
  - Rate-Limiting
  - Redirection
  - Etc...
- Doesn’t this sound as a great solution?

BGP FlowSpec

Pros n Cons..

- Good solution for
  - Done with hardware acceleration for carrier grade routers
  - Can provide chirurgical precision of match statements and actions to impose
- But...
  - Customer need to call my provider
  - Customer need the provider to accept and run this filter on each of their backbone/edge routers
  - Customer need to call the provider and remove the rule after!
- Reality: It won’t happen...

BGP FlowSpec

FlowSpec as Alternative

- Comparison with the other solutions
  - Makes static PBR a dynamic solution!
  - Allows to propagate PBR rules
  - Existing control plane communication channel is used

- How?
  - By using your existing MP-BGP infrastructure

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Public

BGP FlowSpec

Overview

- *RFC 5575 - A flow specification is an n-tuple consisting of several matching criteria that can be applied to IP traffic. A given IP packet is said to match the defined flow if it matches all the specified criteria*
- A flowspec is said to be n-tuple because there are multiple match cirteria's that can be defined and all the match criteria should be matched.
    - Traffic will not match the flowspec entry if all the tuples are not matched.
- BGP FlowSpec New NLRI – AFI=1 and SAFI=133

BGP FlowSpec

DDoS Mitigation Steps

- Mitigation of DDOS attacks is performed in two steps:
  - Diversion – Send traffic to a specialized device that removes the fake packets from the traffic stream while retaining the legitimate packets.
    - Define match criteria
    - Define action
  - Return – Send back the clean / legitimate traffic to the server.

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Public

BGP FlowSpec

DDoS Mitigation

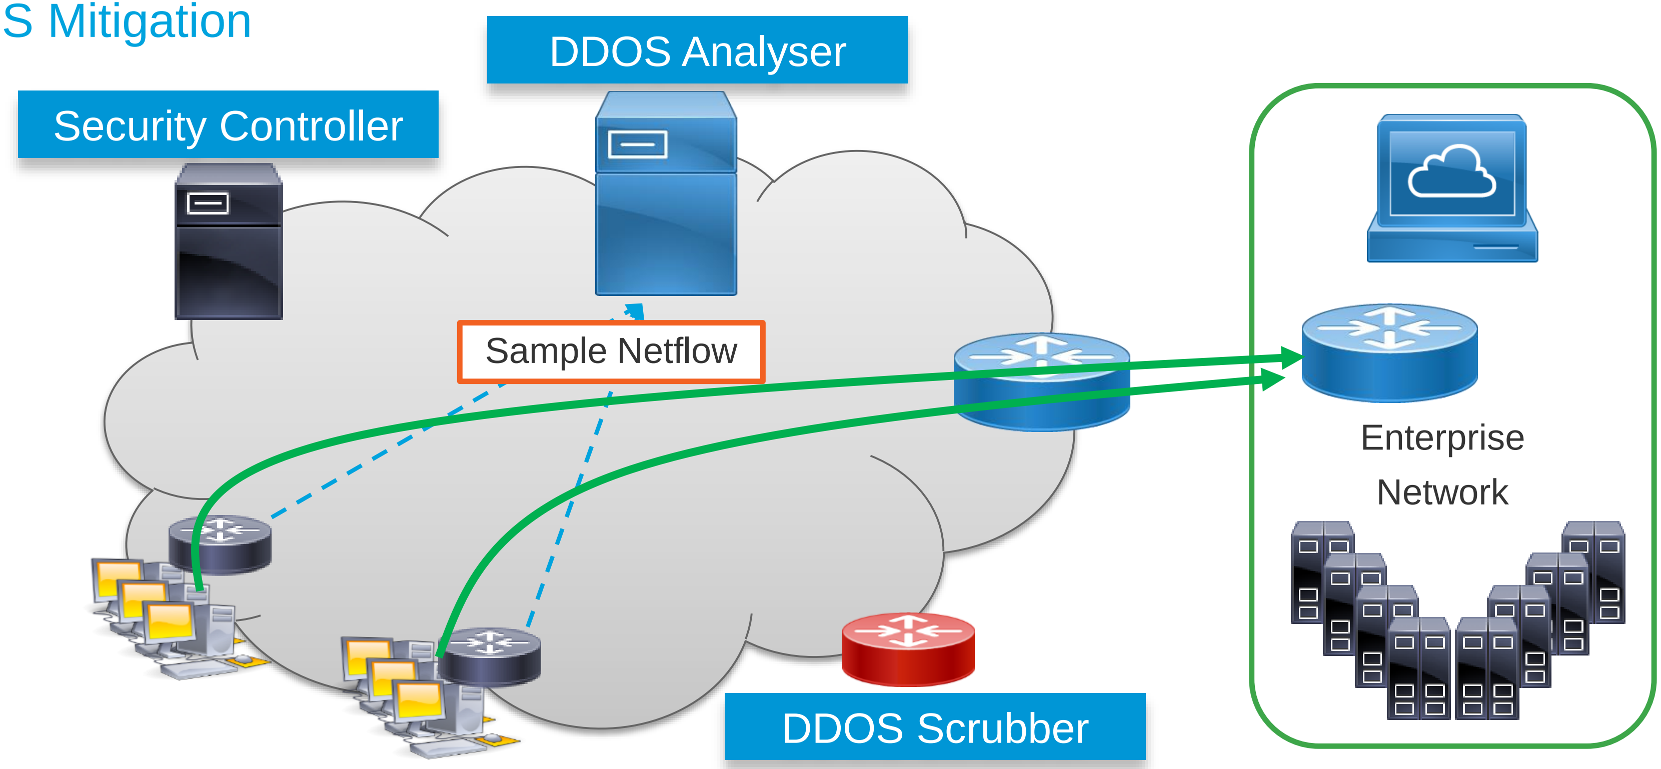

BGP FlowSpec

DDoS Mitigation

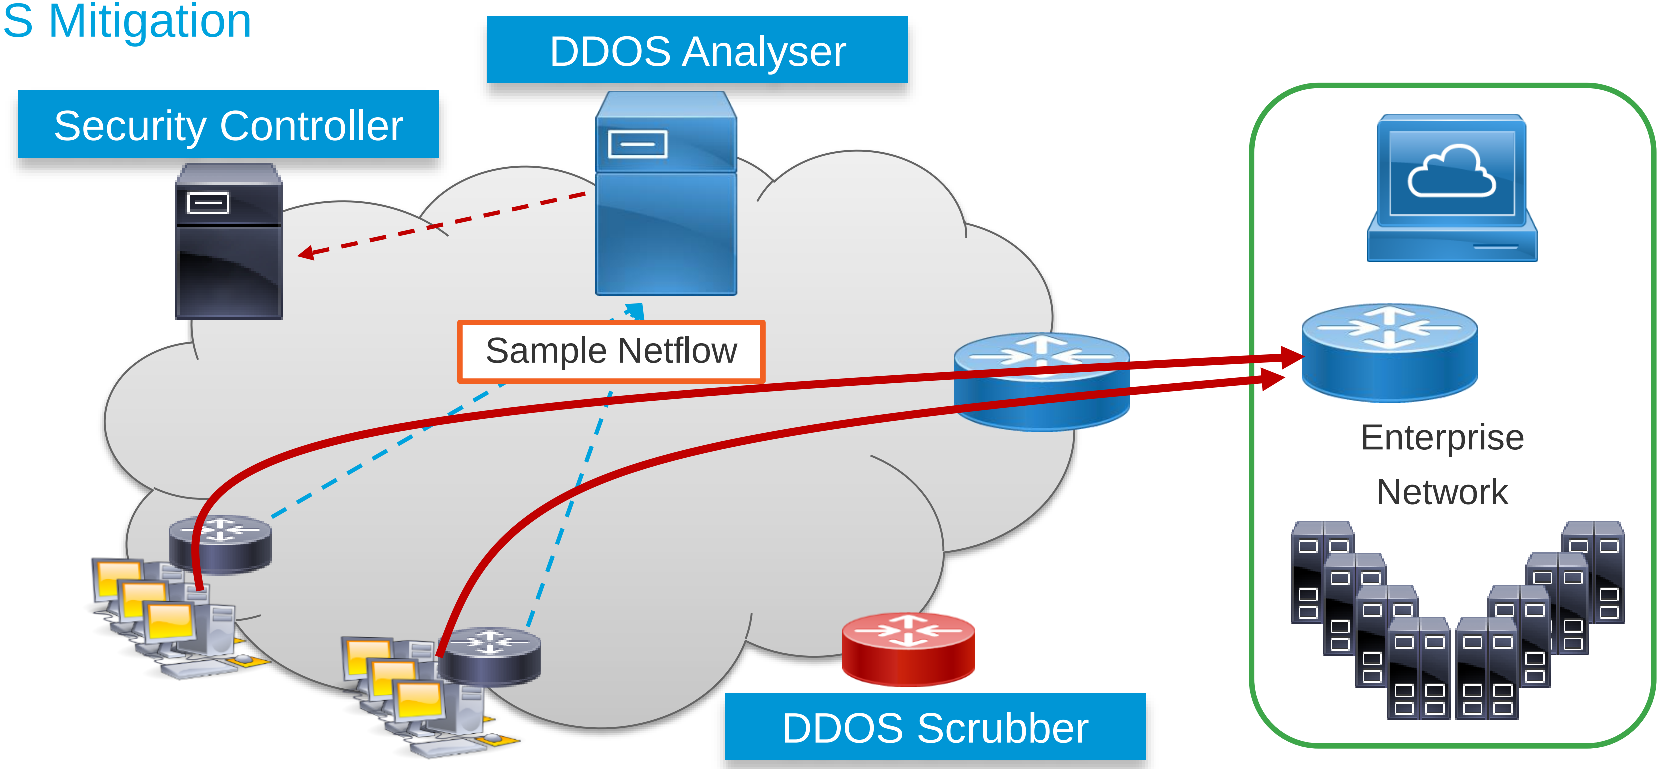

BGP FlowSpec

DDoS Mitigation

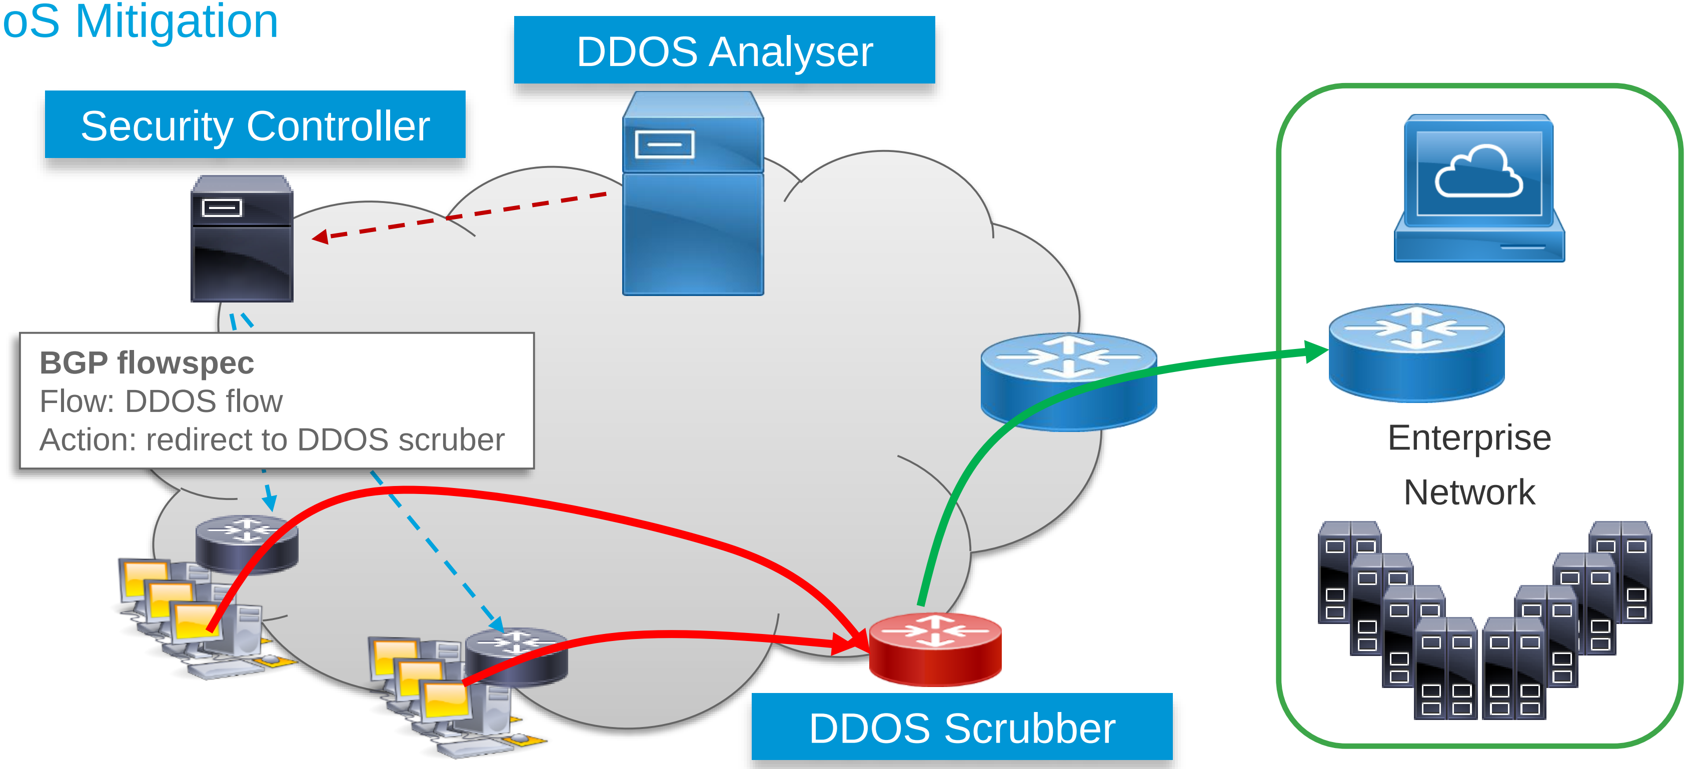

BGP FlowSpec – NLRI based on Match Criteria

BGP Flowspec NLRI Type	QoS Match Fields
Type 1	Destination IP / IPv6 address
Type 2	Source IP / IPv6 address
Type 4	IP / IPv6 Protocol
Type 4	Source or destination port
Type 5	Destination port
Type 6	Source port
Type 7	ICMP Type
Type 8	ICMP Code
Type 9	TCP flags
Type 10	Packet length
Type 11	DCSP
Type 12	Fragmentation bits

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Public

BGP FlowSpec

NLRI Type based on Action

Type	Description	PBR Action
0x8006	traffic-rate	Drop | Police
0x8007	traffic-action	Terminal Action + Sampling
0x8008	redirect-vrf	Redirect VRF
0x8009	traffic-marking	Set DSCP
0x0800	Redirect IP NH	Redirect IPv4 or IPv6 Next-Hop

BGP FlowSpec

Configuration – IOS XR

```
RP/0/0/CPU0:RR_R3(config)#class-map type traffic match-all FS_RULE
RP/0/0/CPU0:RR_R3(config-cmap)#match source-address ipv4 192.168.1.1/32
RP/0/0/CPU0:RR_R3(config-cmap)#match destination-address ipv4 192.168.5.5/32
RP/0/0/CPU0:RR_R3(config-cmap)#exit
RP/0/0/CPU0:RR_R3(config)#policy-map type pbr FS_POLICY_MAP
RP/0/0/CPU0:RR_R3(config-pmap)#class FS_RULE
RP/0/0/CPU0:RR_R3(config-pmap-c)#drop
RP/0/0/CPU0:RR_R3(config-pmap-c)#exit
RP/0/0/CPU0:RR_R3(config-pmap)#class class-default
RP/0/0/CPU0:RR_R3(config-pmap-c)#exit
RP/0/0/CPU0:RR_R3(config-pmap)#exit
RP/0/0/CPU0:RR_R3(config)#flowspec
RP/0/0/CPU0:RR_R3(config-flowspec)#local-install interface-all
RP/0/0/CPU0:RR_R3(config-flowspec)#address-family ipv4
RP/0/0/CPU0:RR_R3(config-flowspec-af)#service-policy type pbr FS_POLICY_MAP
RP/0/0/CPU0:RR_R3(config)#commit
```

Install the policies locally on the hardware

BGP FlowSpec

Configuration

- Policies are defined on RR or the controller
- Establish BGP peering with other routers in the network over **address-family flowspec**

```
R2(config)#flowspec
R2(config-flowspec)#local-install interface-all
R2(config-flowspec)#address-family ipv4
```

Demo

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Public

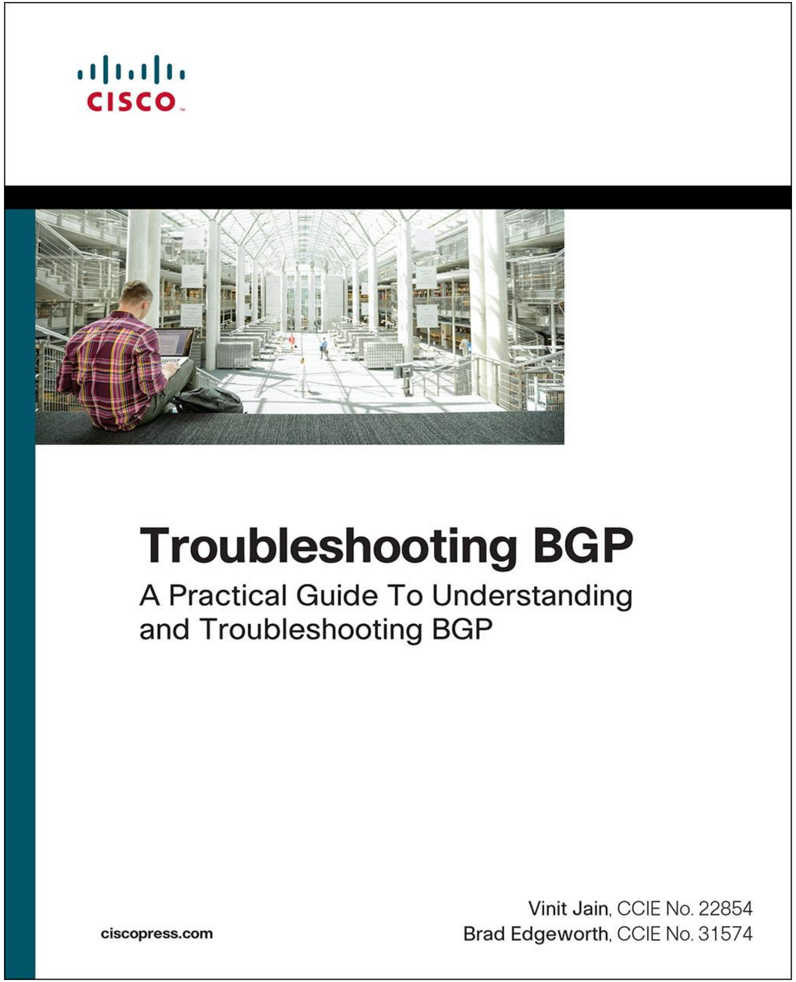

Troubleshooting BGP

A Practical Guide To Understanding and Troubleshooting BGP

Vinit Jain, CCIE No. 22854
Brad Edgeworth, CCIE No. 31574

ciscopress.com

Thank you

© 2016 Cisco and/or its affiliates. All rights reserved. Cisco Public