Scalable DDoS mitigation using BGP Flowspec

Wei Yin TAY
Consulting Systems Engineer
Cisco Systems

Agenda

- Goals of DDoS Mitigation
- Problem description
- Traditional DDoS Mitigation
- Scalable DDoS Mitigation

Goals of Scalable DDoS Mitigation

- Stop the attack
- Drop only the DDoS traffic
- Application aware filtering/redirect/mirroring
- Dynamic and adaptive technology
- Simple to configure
- Easy to disseminate

DDoD Scenario

Demonstration Scenario

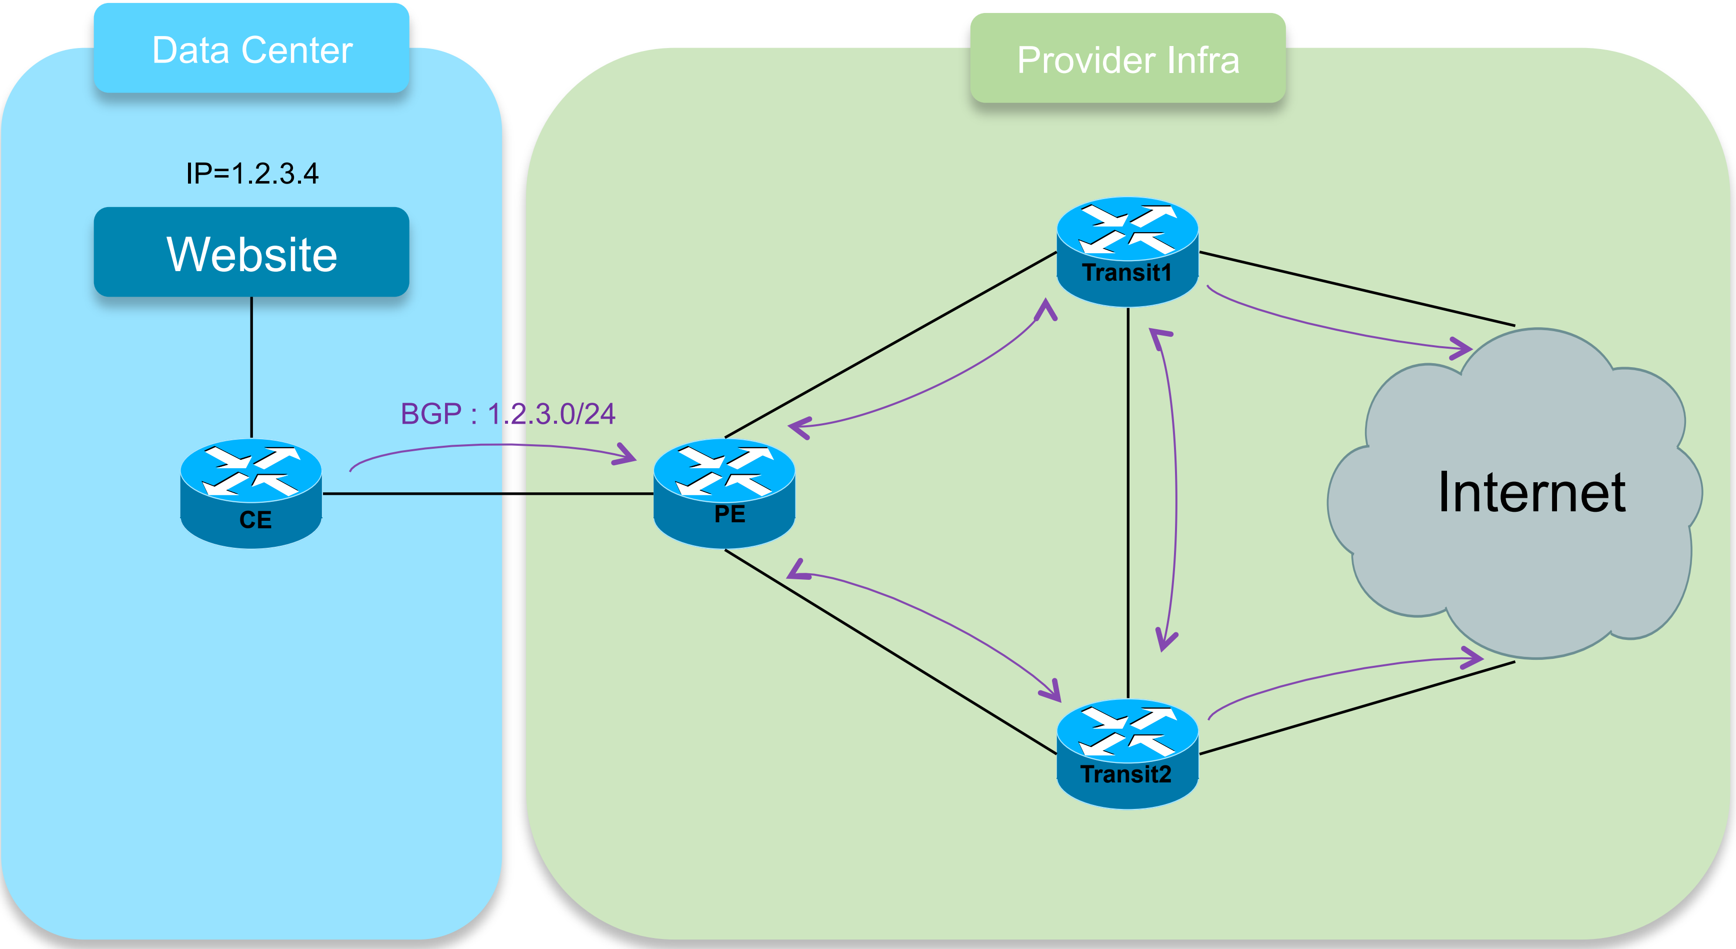

Demonstration Scenario

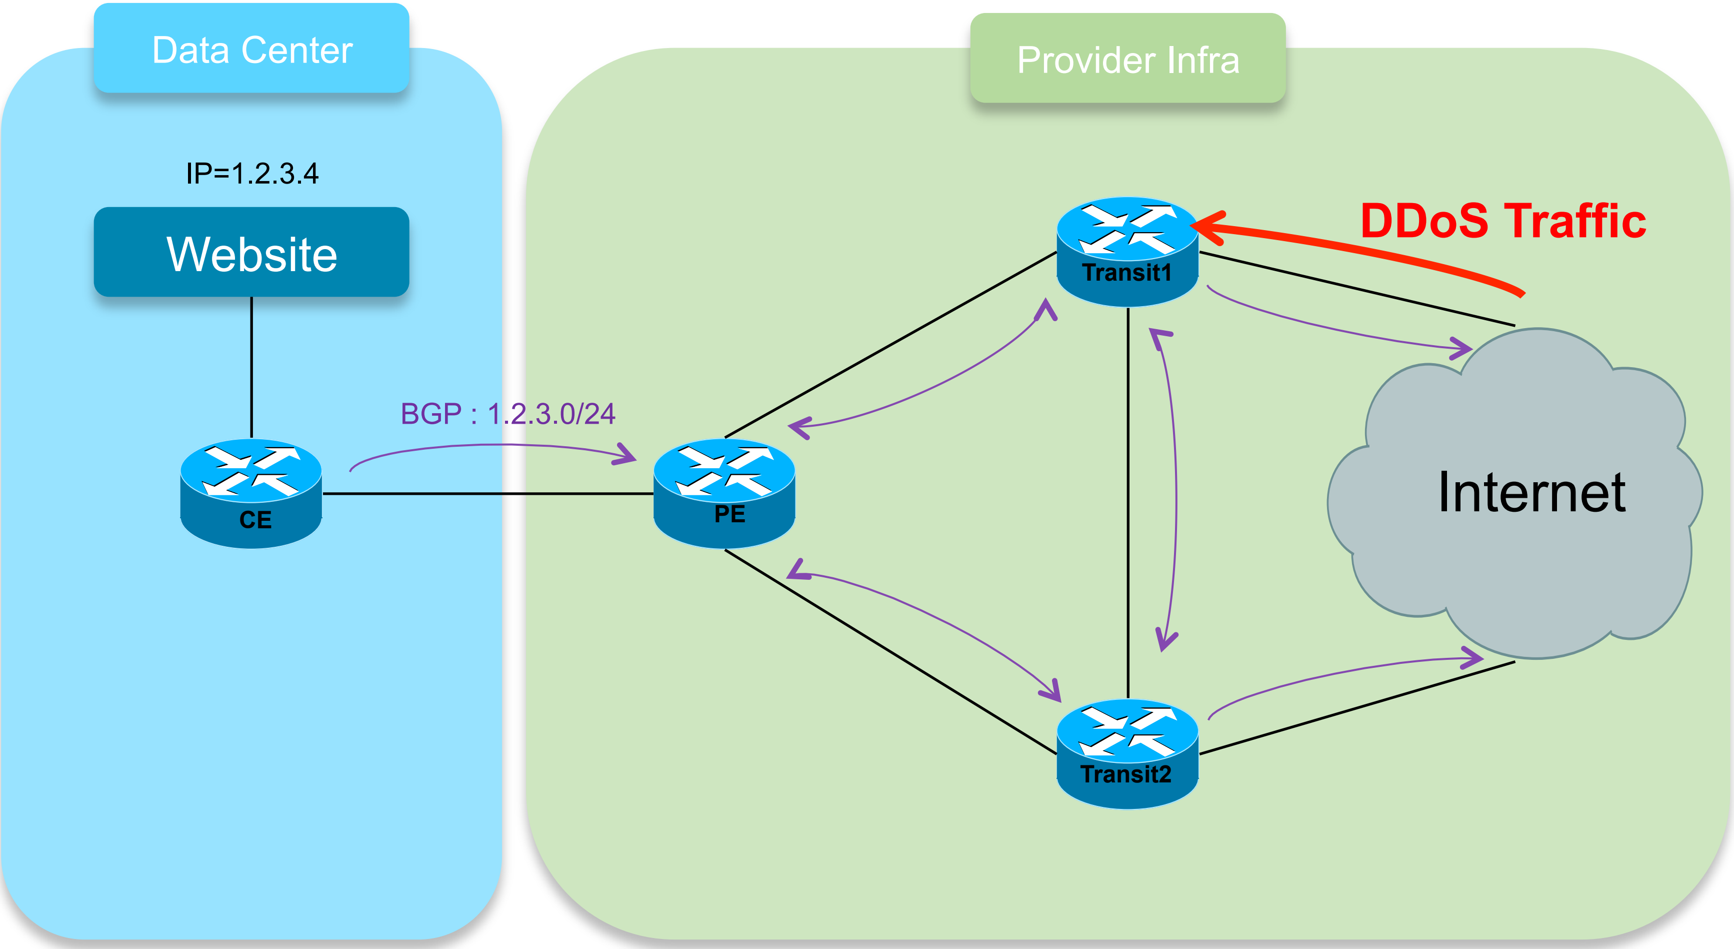

Demonstration Scenario

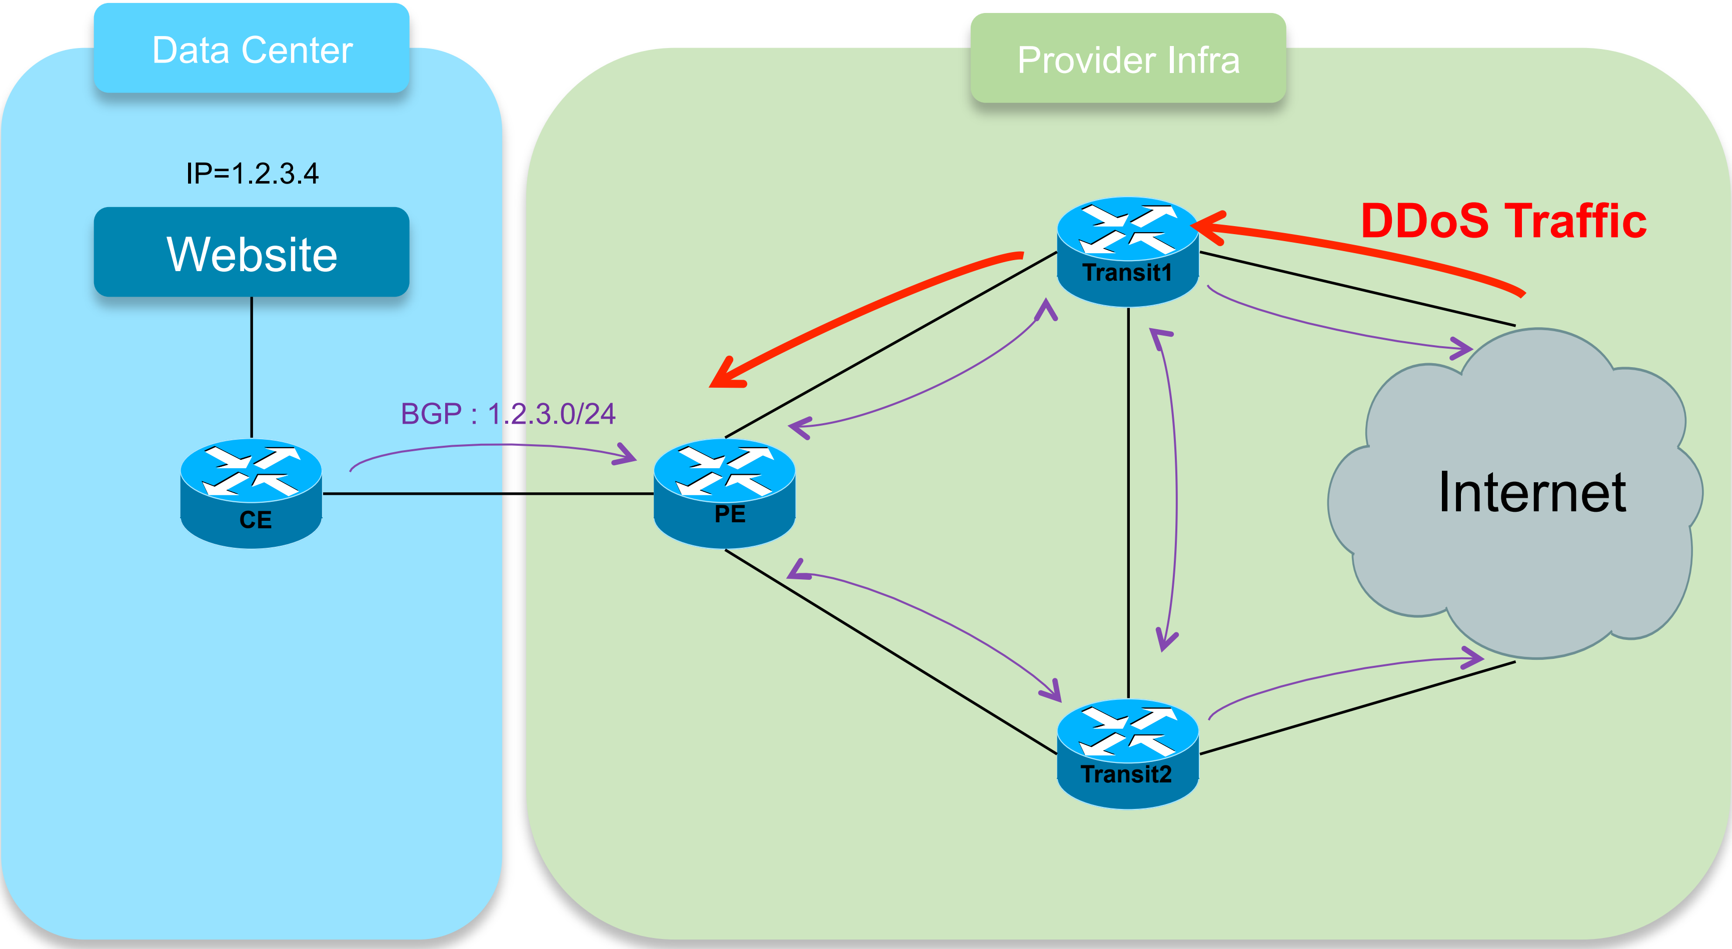

Demonstration Scenario

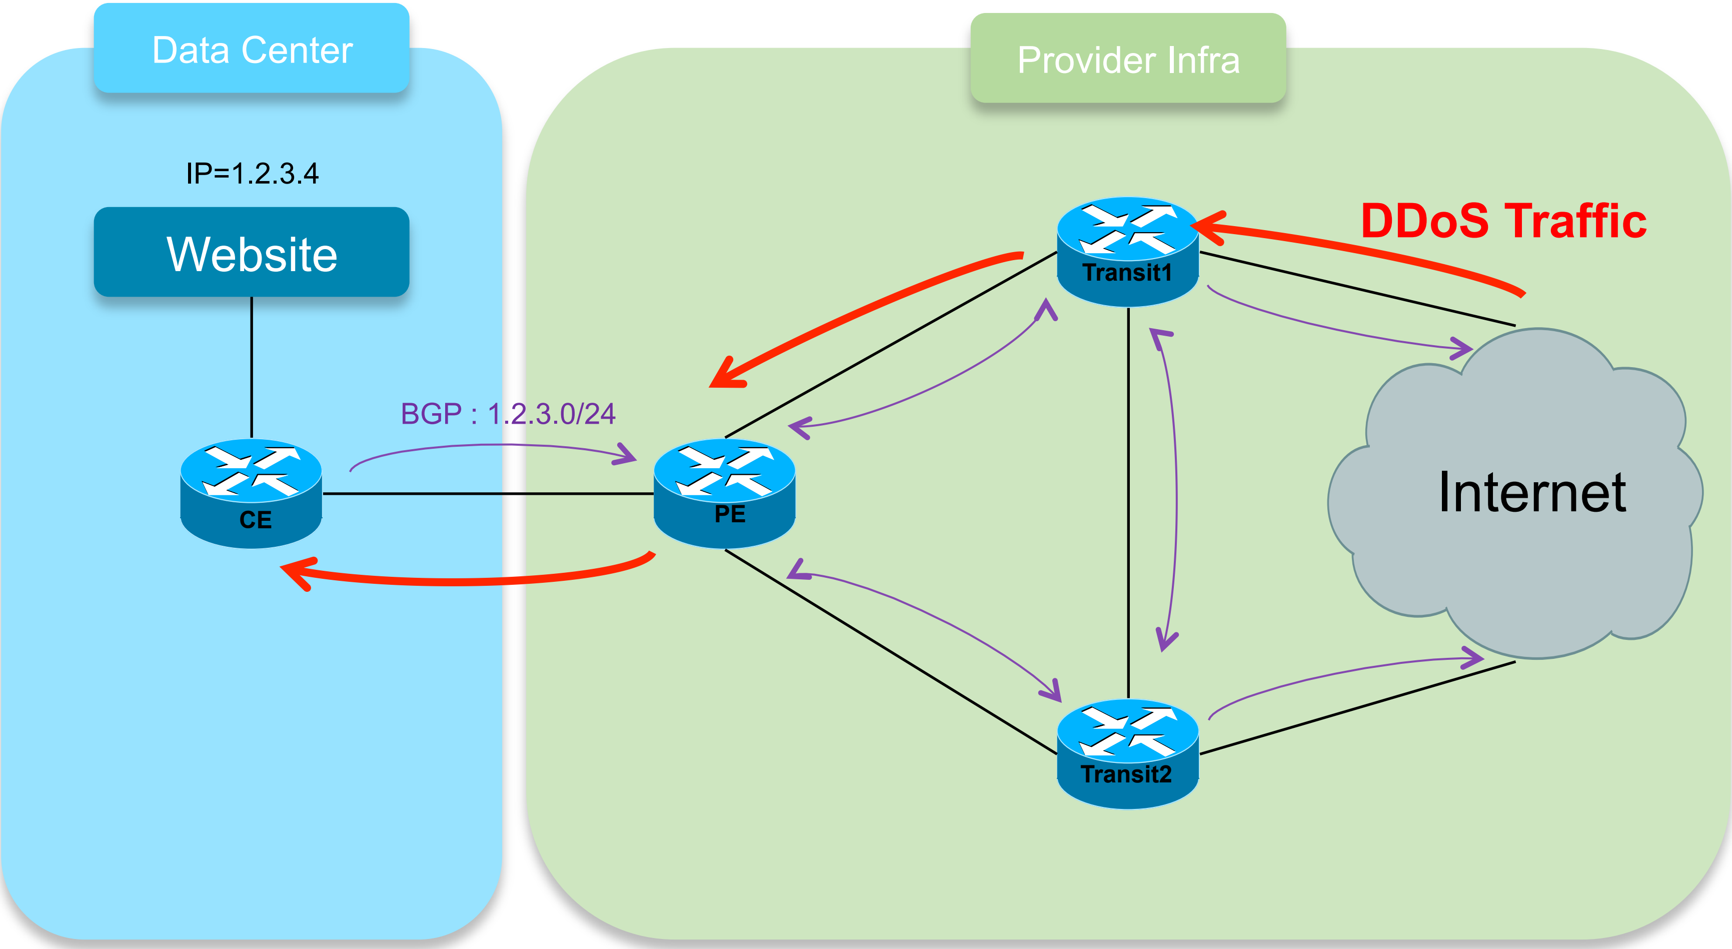

Demonstration Scenario

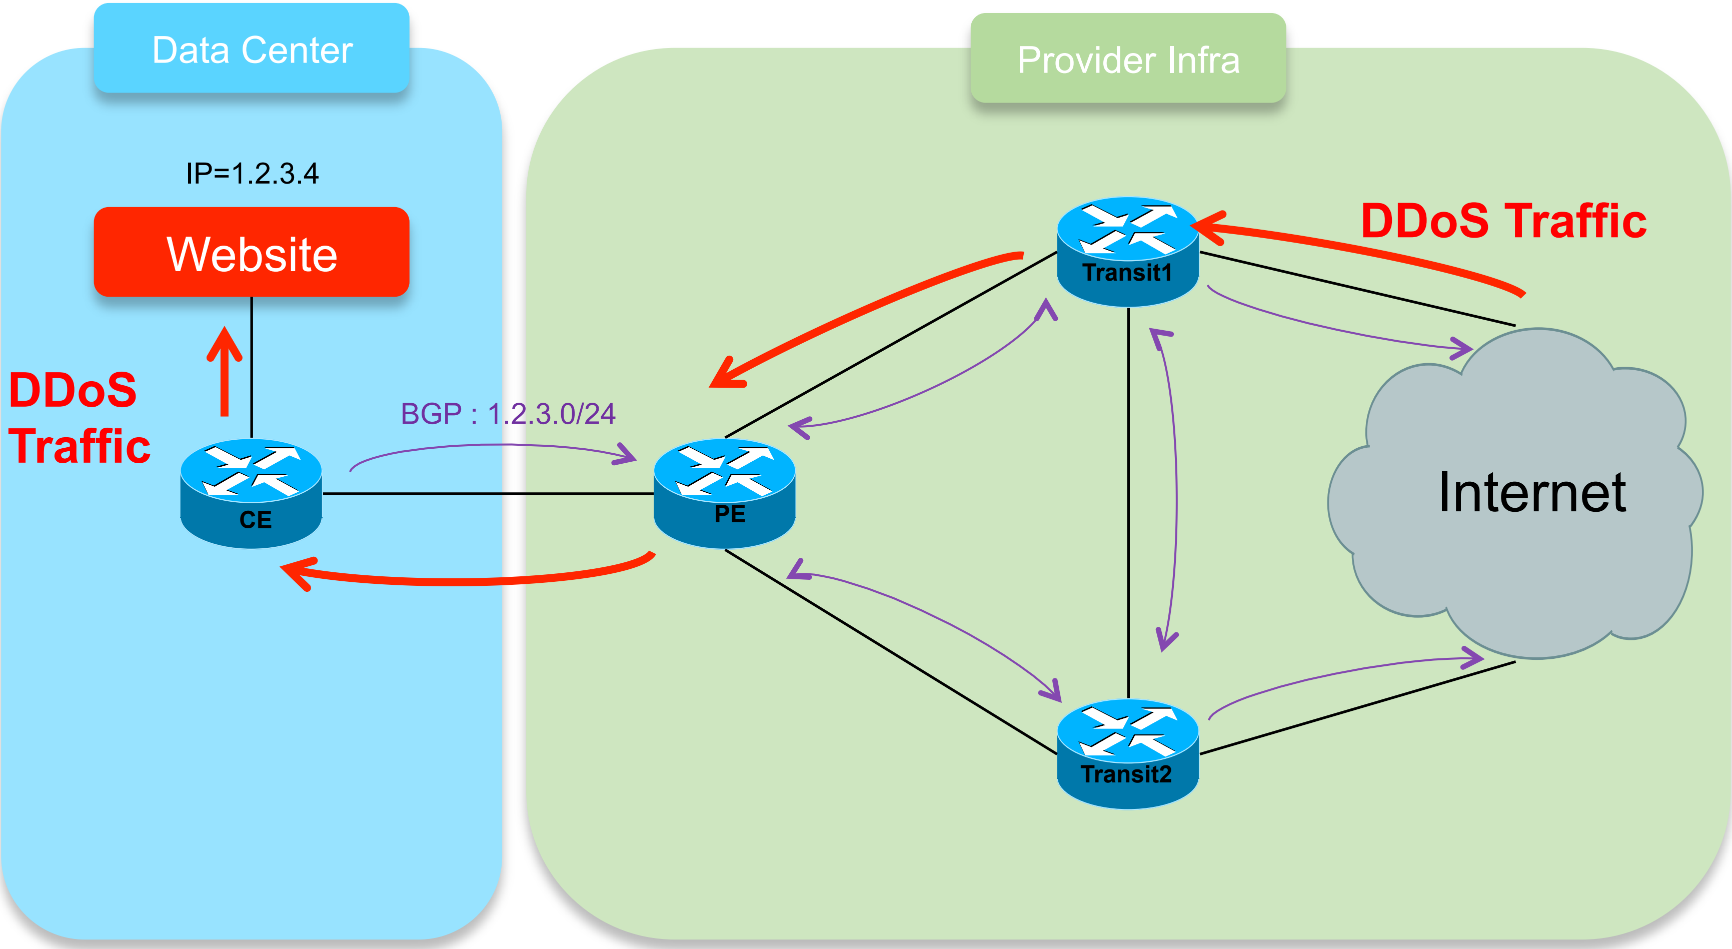

DDoD Mitigation Solutions

DDoS Overview

- 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

Solution: Remotely Triggered Black Hole

It is time to use the blackhole community given by the provider (i.e. 64500:666)

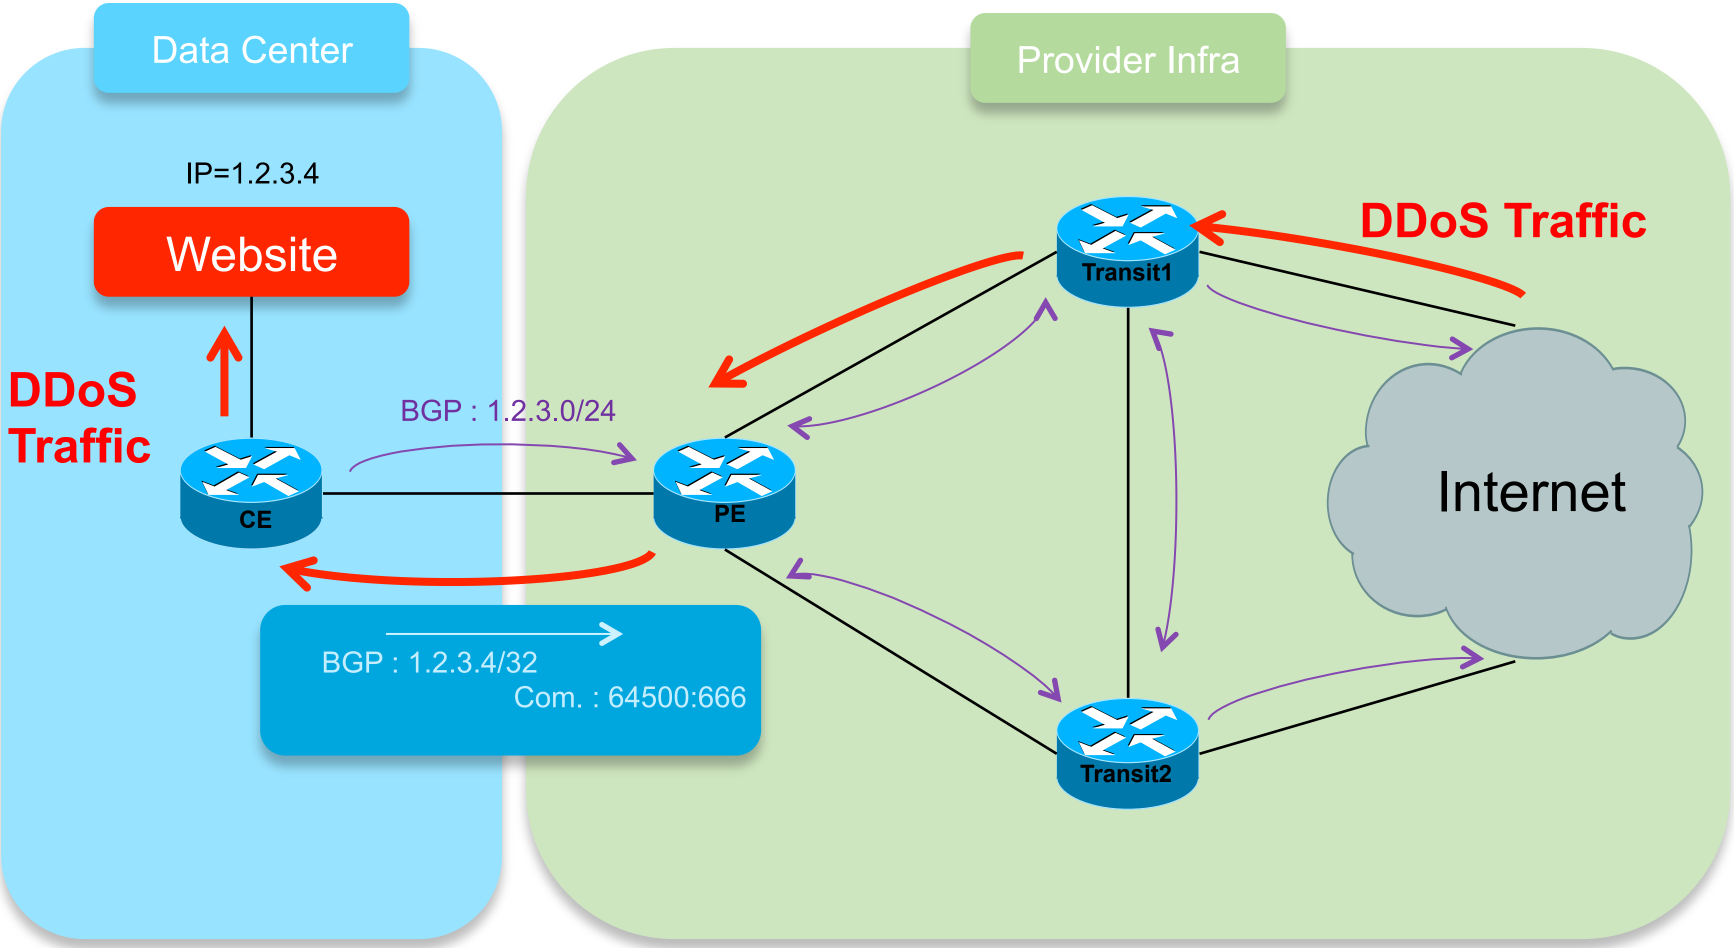

BGP : 1.2.3.4/32
Com. : 64500:666

Solution: Remotely Triggered Black Hole

It is time to use the blackhole community given by the provider (i.e. 64500:666)

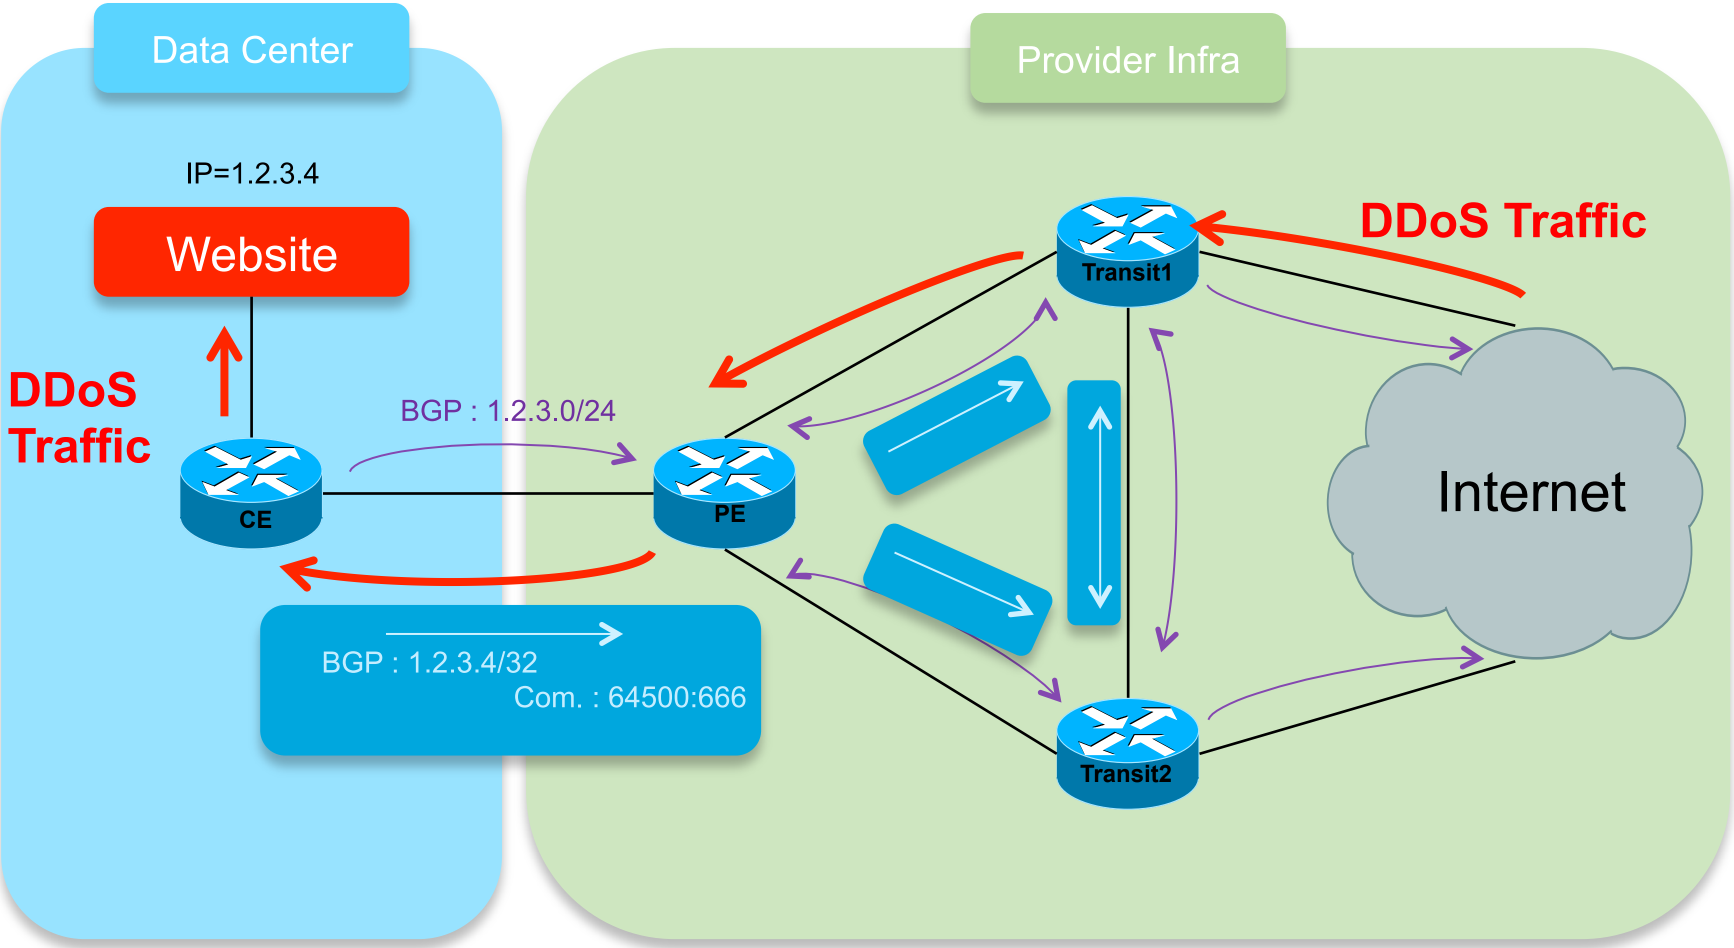

Solution: Remotely Triggered Black Hole

It is time to use the blackhole community given by the provider (i.e. 64500:666)

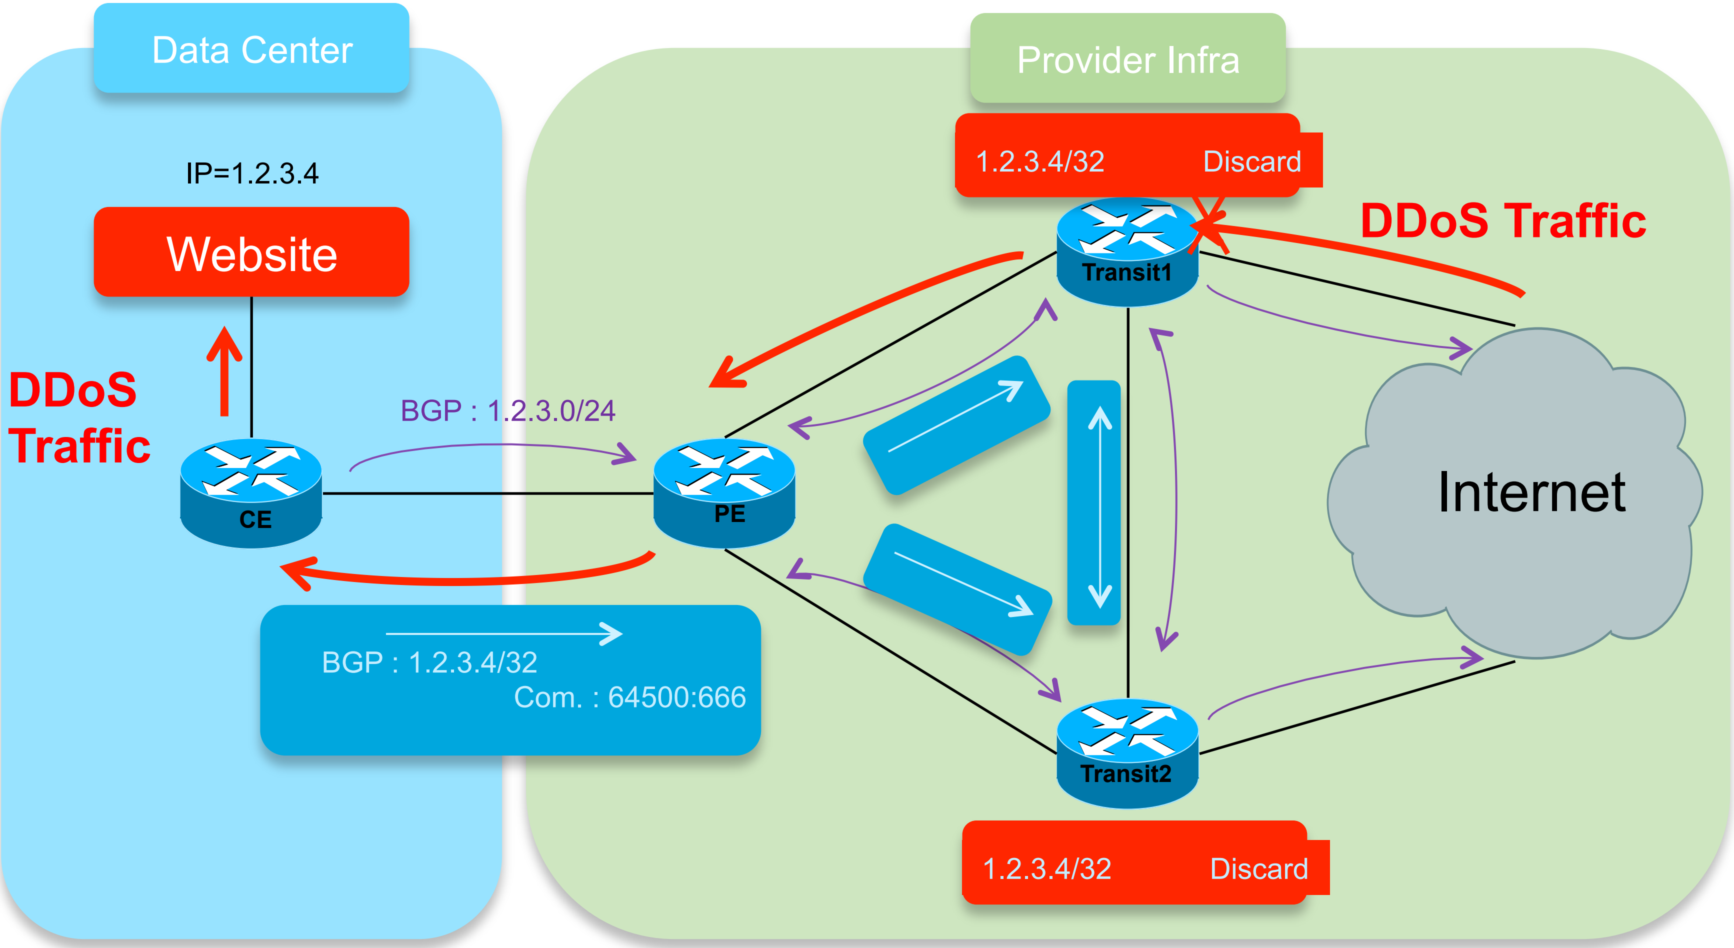

Solution: Remotely Triggered Black Hole

It is time to use the blackhole community given by the provider (i.e. 64500:666)

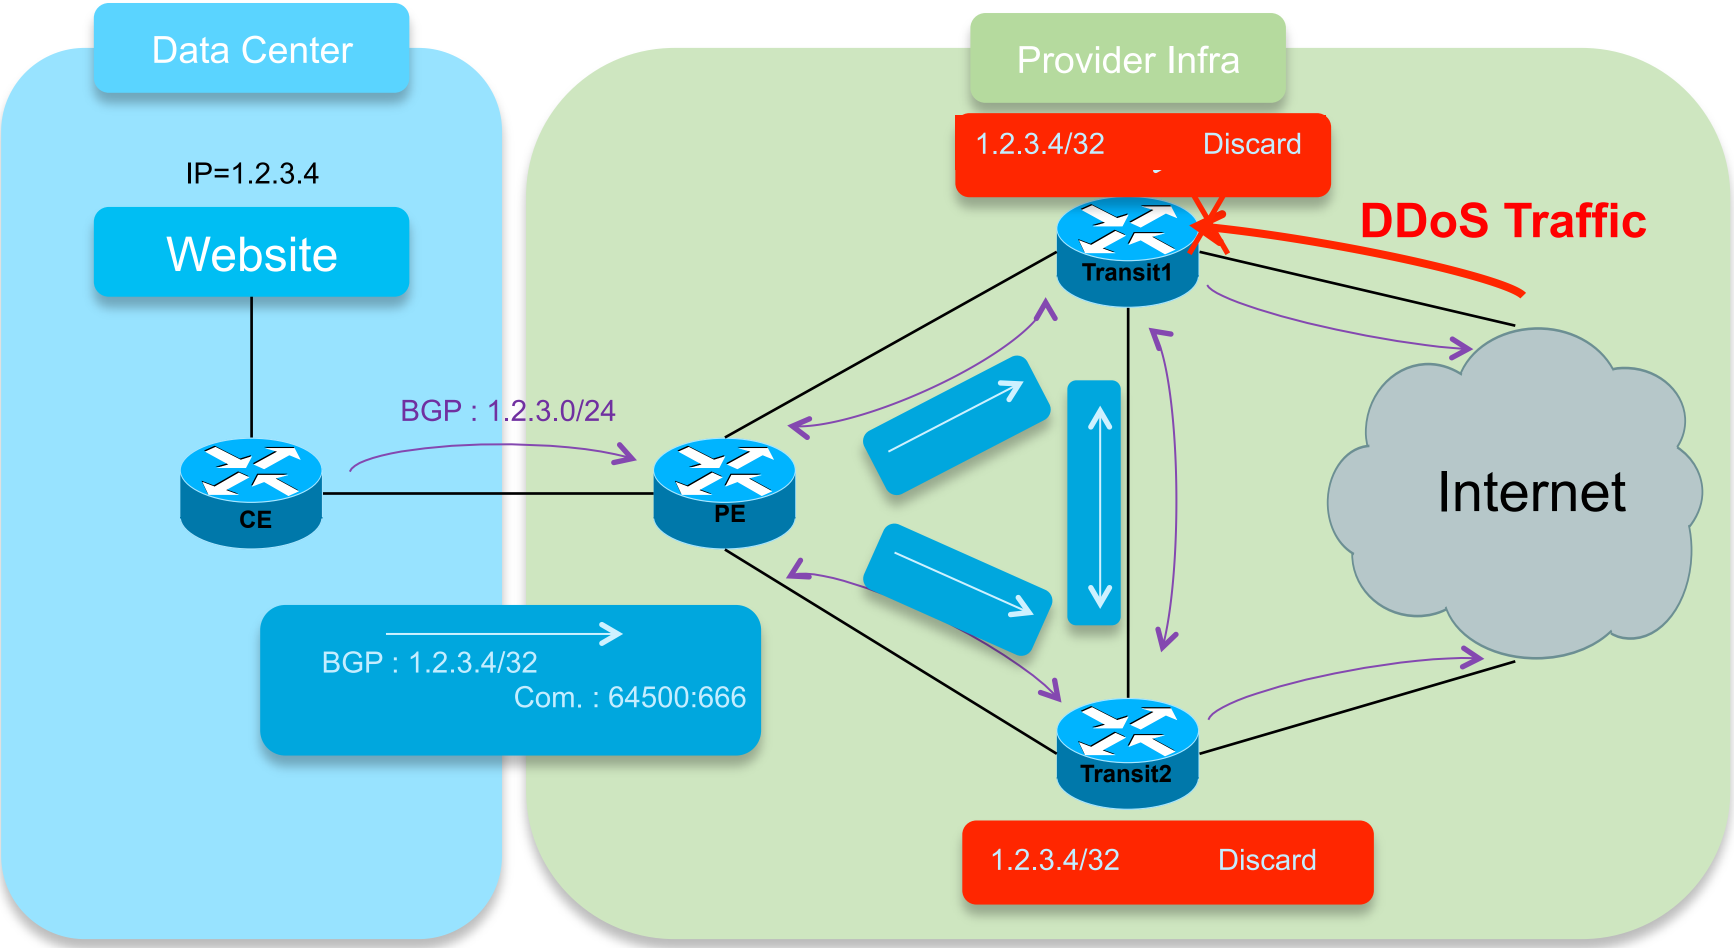

Solution: Remotely Triggered Black Hole

- 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….

Solution: 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?

Solution: Policy Based Routing

- 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…

Scalable DDoS Mitigation

Flowspec as an 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

Dissemination of Flow Specification Rules (RFC5575)

- Why using BGP?

  Simple to extend by adding a new NLRI with MP_REACH_NLRI and MP_UNREACH_NLRI

  Networkwide loopfree point-to-multipoint path is already setup

  Already used for every other kind of technology (IPv4, IPv6, VPN, Multicast, Labels, etc…)

  Inter-domain support

  Networking engineers and architects understand perfectly BGP

- Capability to send via a BGP Address Family

  Match criteria

  Action criteria

Dissemination of Flow Specification Rules (RFC5575)

New NLRI defined (AFI=1, SAFI=133)

1. Destination IP Address (1 component)
2. Source IP Address (1 component)
3. IP Protocol (+1 component)
4. Port (+1 component)
5. Destination port (+1 component)
6. Source Port (+1 component)
7. ICMP Type
8. ICMP Code
9. TCP Flags
10. Packet length
11. DSCP
12. Fragment

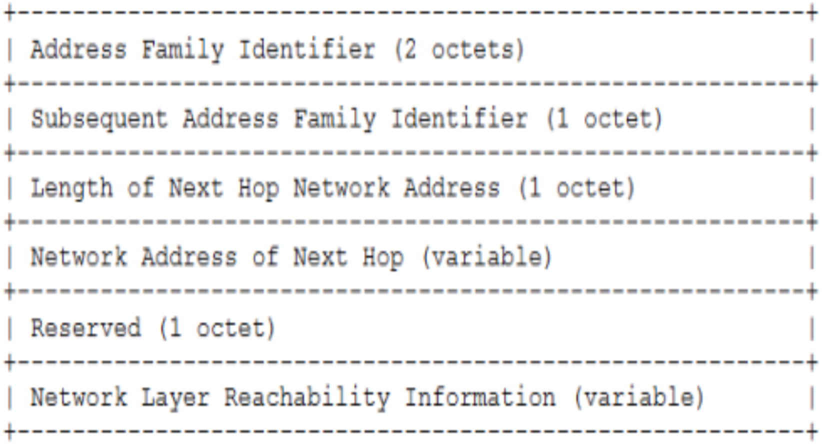

```
+---------------------------------------------------------+
| Address Family Identifier (2 octets)                    |
+---------------------------------------------------------+
| Subsequent Address Family Identifier (1 octet)          |
+---------------------------------------------------------+
| Length of Next Hop Network Address (1 octet)            |
+---------------------------------------------------------+
| Network Address of Next Hop (variable)                  |
+---------------------------------------------------------+
| Reserved (1 octet)                                      |
+---------------------------------------------------------+
| Network Layer Reachability Information (variable)       |
+---------------------------------------------------------+
```

The MP_REACH_NLRI – RFC 4760

Notice from the RFC: “Flow specification components must follow strict type ordering. A given component type may or may not be present in the specification, but if present, it MUST precede any component of higher numeric type value.”

BGP Flowspec Traffic Actions

- Flowspec Traffic Actions

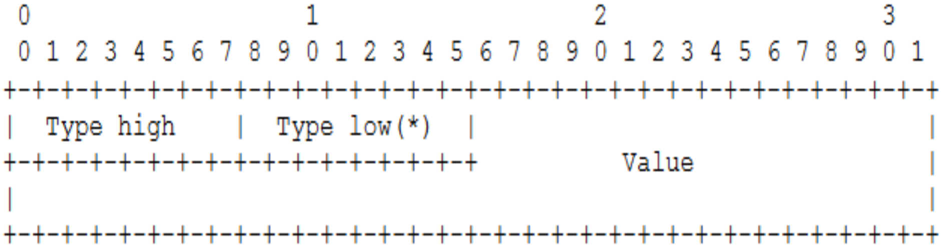

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Type high    |  Type low(*)  |                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+          Value                |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- RFC5575 Flowspec available actions

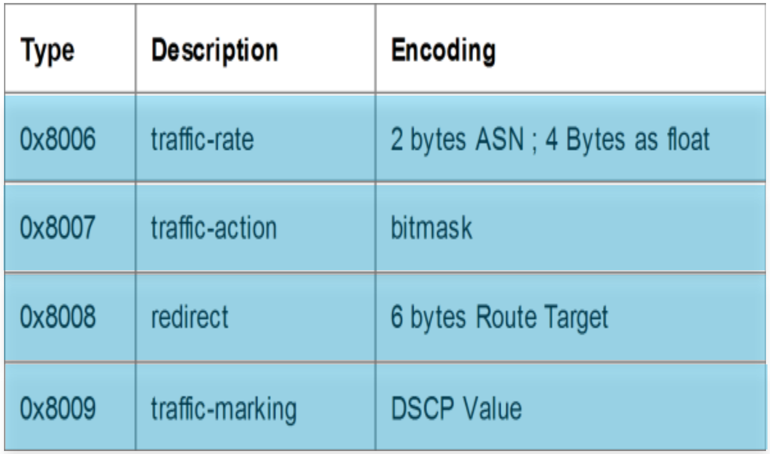

Type	Description	Encoding
0x8006	traffic-rate	2 bytes ASN ; 4 Bytes as float
0x8007	traffic-action	bitmask
0x8008	redirect	6 bytes Route Target
0x8009	traffic-marking	DSCP Value

RFC5575 Architecture

It is time to use the blackhole community given by the provider (i.e. 64500:666)

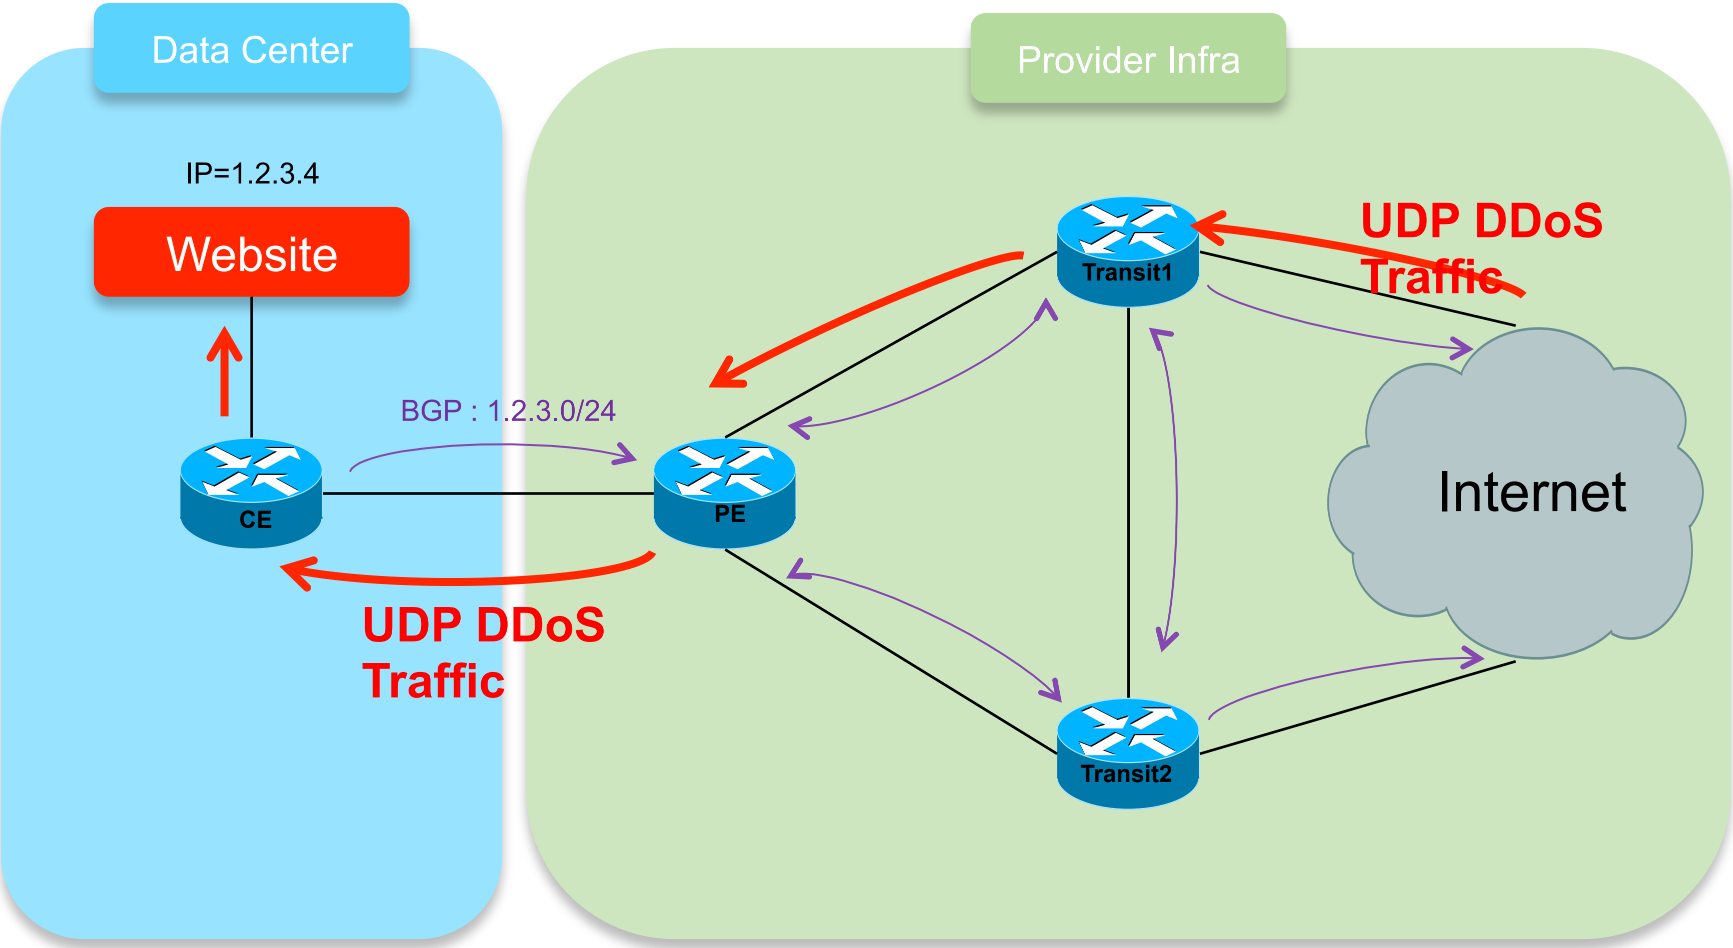

RFC5575 Architecture

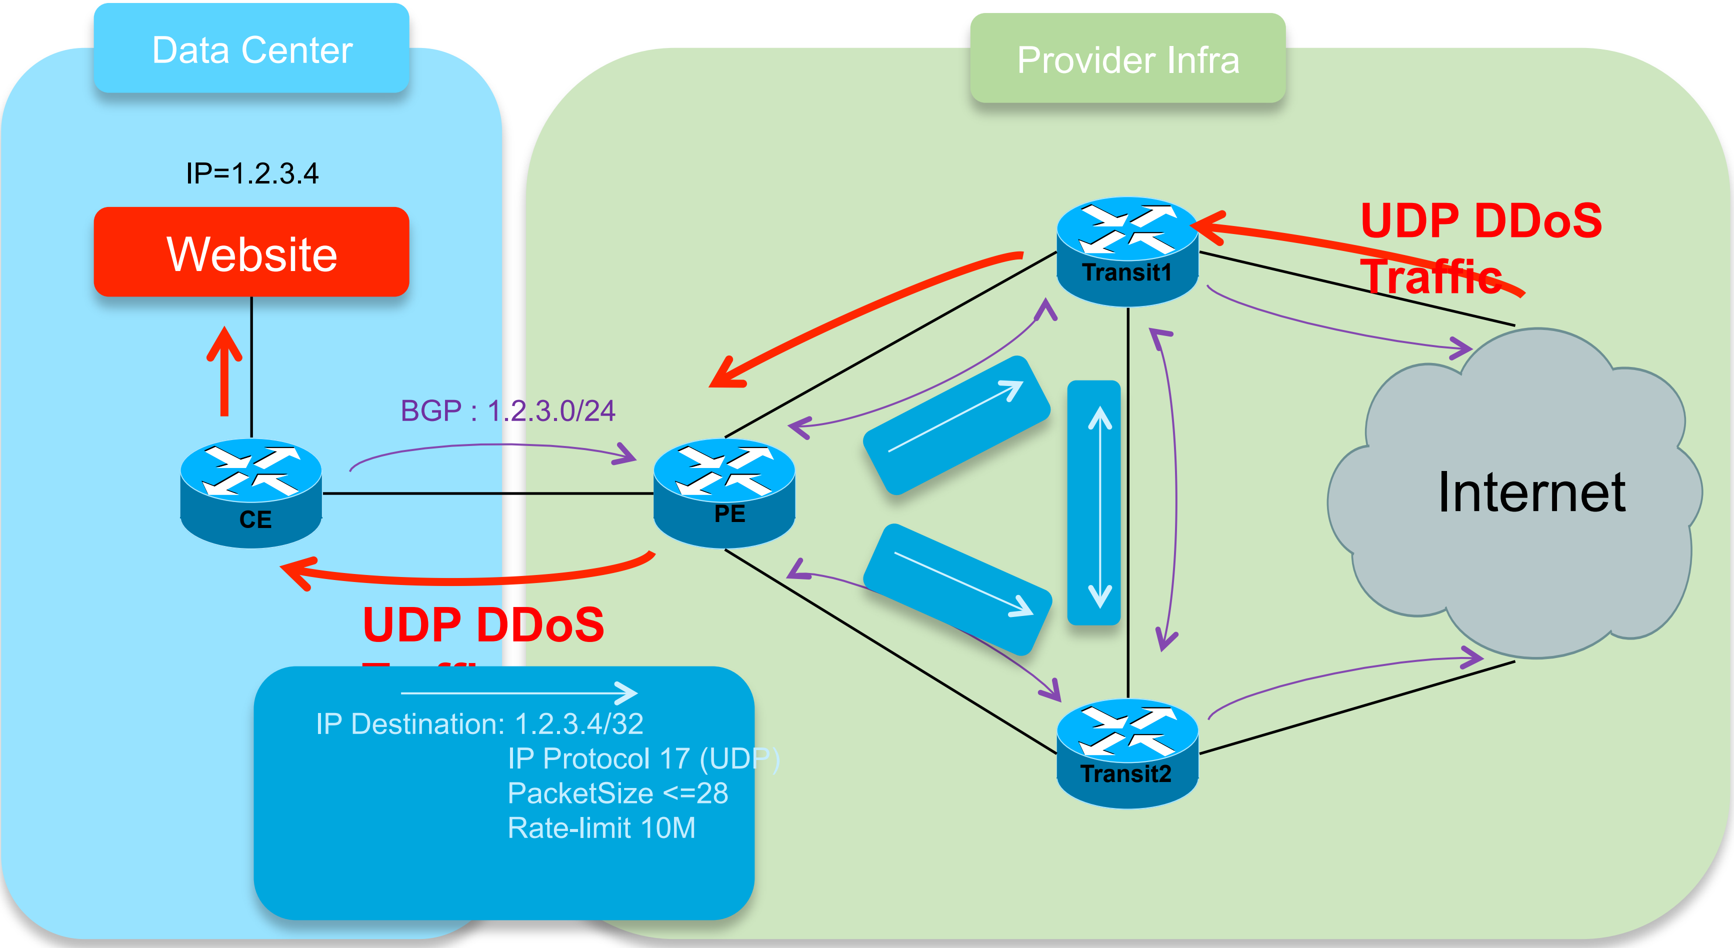

IP Destination: 1.2.3.4/32
IP Protocol 17 (UDP)
PacketSize <=28
Rate-limit 10M

RFC5575 Architecture

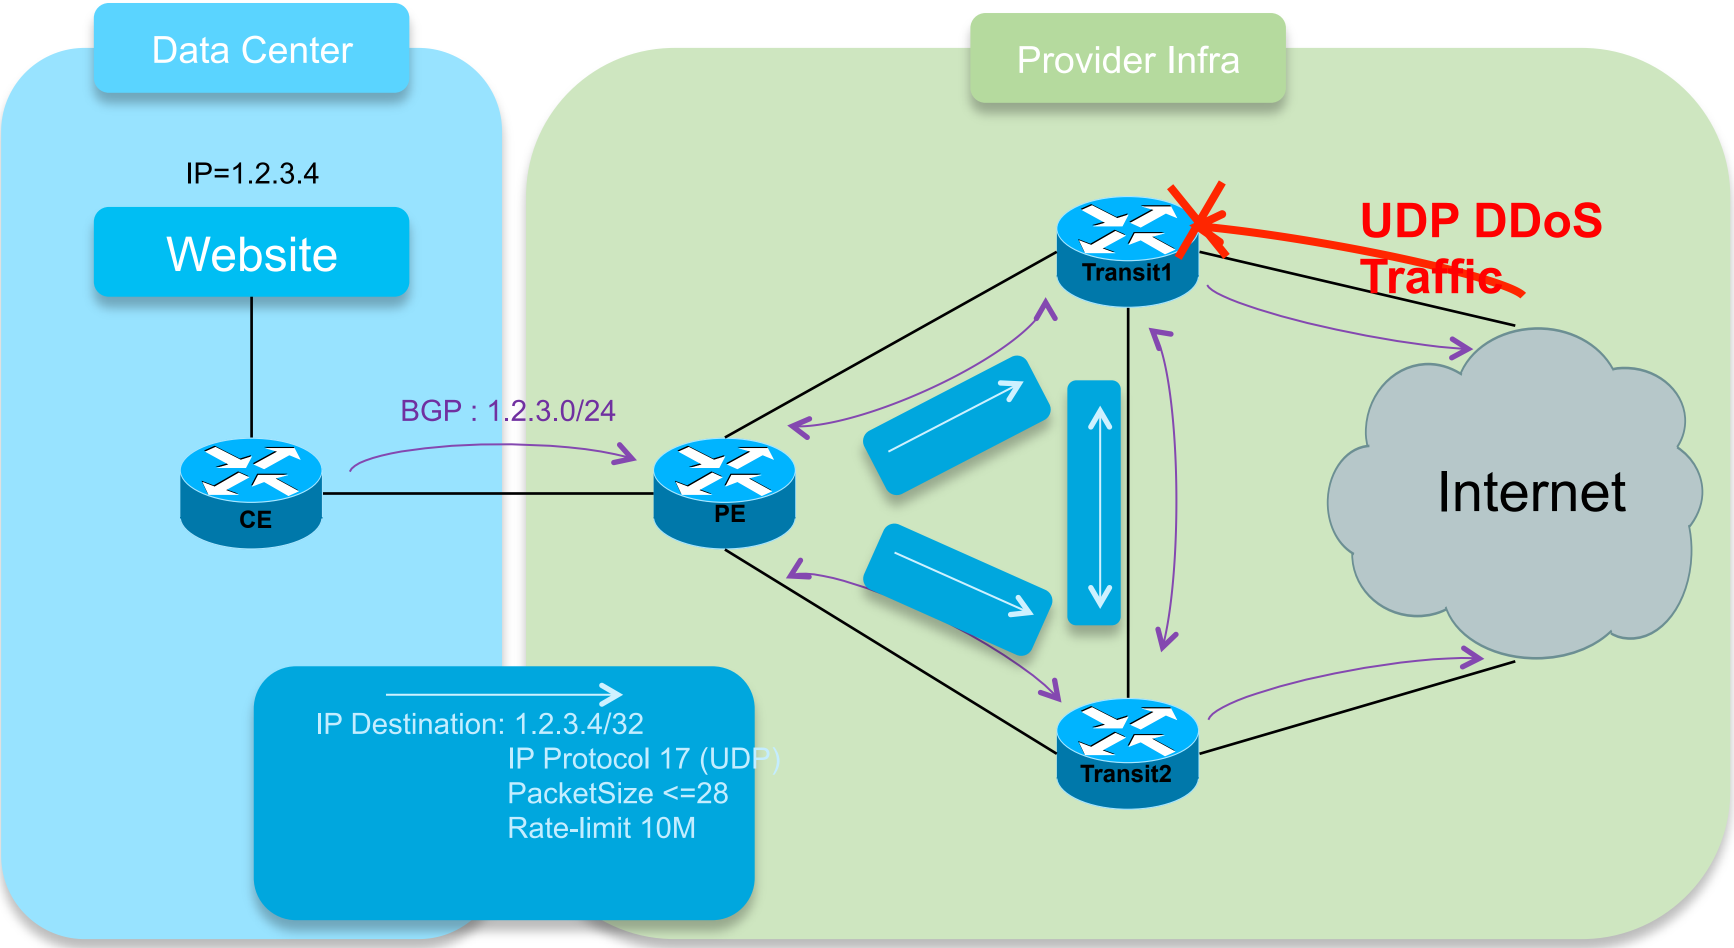

IP Destination: 1.2.3.4/32
IP Protocol 17 (UDP)
PacketSize <=28
Rate-limit 10M

RFC5575 Architecture

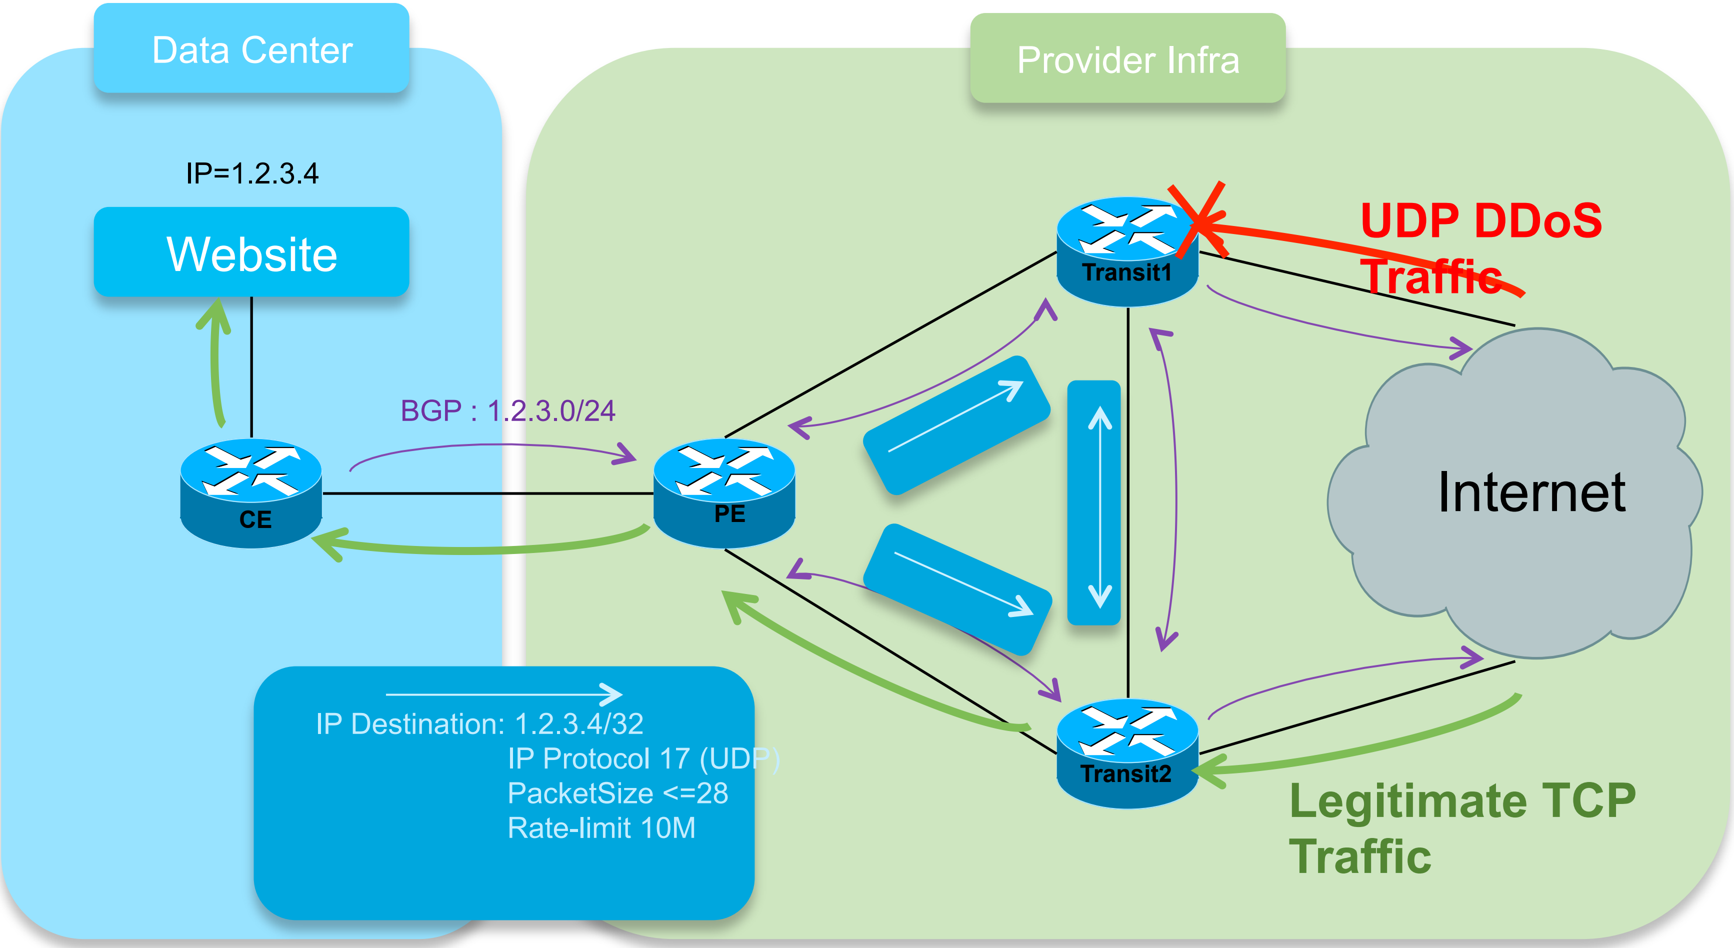

IP Destination: 1.2.3.4/32
IP Protocol 17 (UDP)
PacketSize <=28
Rate-limit 10M

Real life architecture

- In reality this architecture is not deployed
  - Service Provider DO NOT trust the Customer
  - It requires new BGP AFI/SAFI combination to be deployed between Customer and Service provider
  - Both these result in Flowspec not being deployed between Customer and service provider
- What is done instead?
  - SP utilize a central Flowspec speaker(s)
  - Have it BGP meshed within the Service Provider routers
  - Only the central Flowspec speaker is allowed to distribute Flowspec rules
  - Central Flowspec speaker is considered "trusted" by the network
  - Central Flowspec speaker is managed by the service provider

RFC5575 Architecture

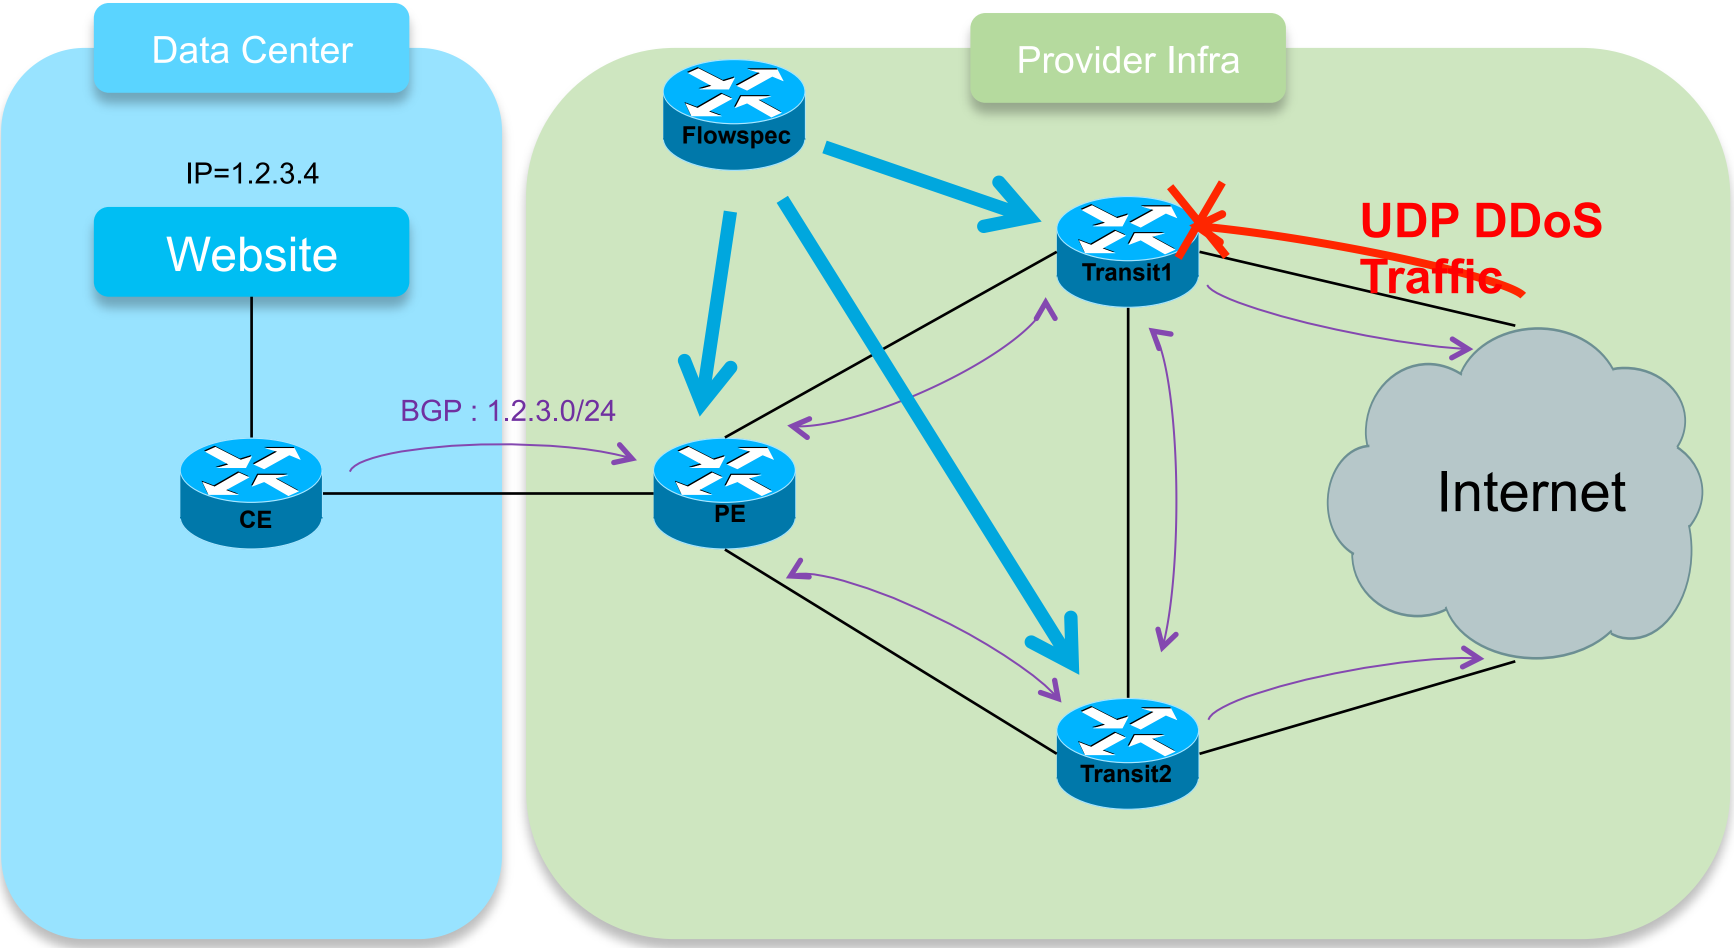

RFC5575 Architecture

Rules inserted by:
- CLI
- Customer Portal
- Workflow
- etc

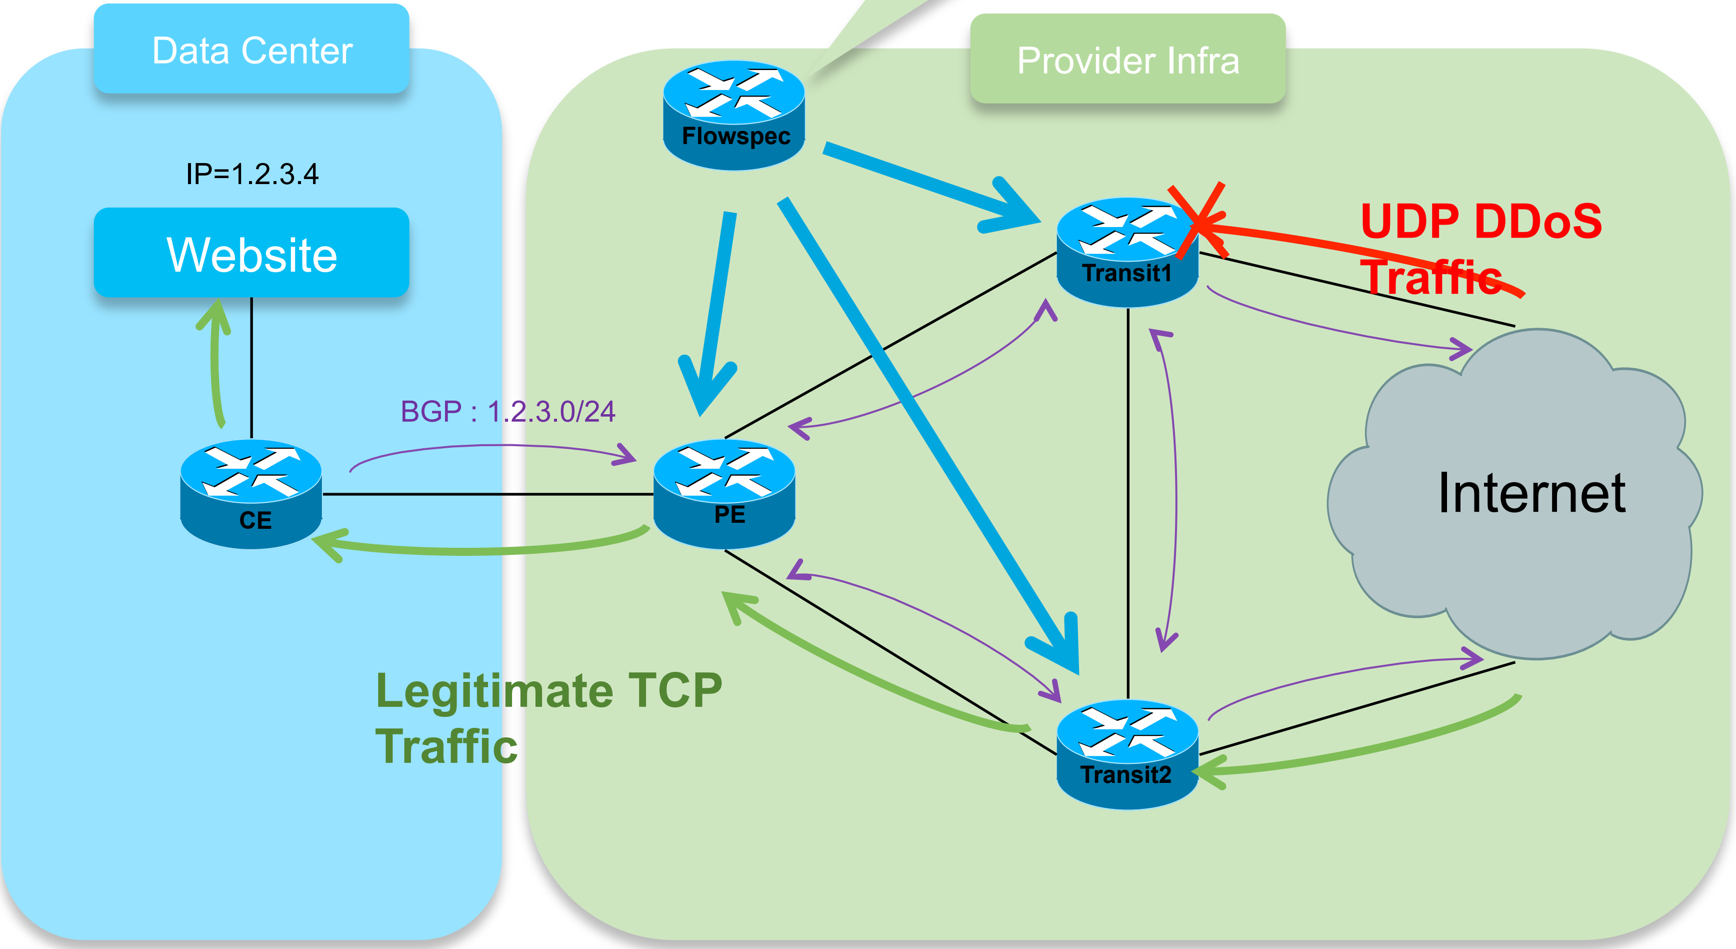

Some thoughts about traffic redirection

- Traffic-rate, traffic-marking are useful for simple attacks, but….
- Traffic-redirect

  Lets you redirect traffic in a VRF (by specifying the VPN RT value)

  Allows to change dynamically the path of a flow without injecting additional BGP routes

- Great too to clean DDoS traffic with a DPI probe

Thank you.

cisco