FlowSpec

Frédéric Gabut-Deloraine

NEO TELECOMS

FRnOG - 2 décembre 2011

Introduction

Dissemination of Flow Specification Rules

- D(D)oS filtering
- Regular use
- Easy to disseminate

Agenda

Background

- Forwarding traffic
- DDoS Mitigation
  - Remotely Triggered Black Hole
  - Policy/Source Based Routing

How the authors see it

- RFC definitions
- Theoretical network design

Back to real life

- Real network design
- Features
  - DDoS Mitigation
  - Traffic Interception

Background

Forwarding traffic

- A *switch* forwards traffic according to the *MAC destination address*,
- An IP *router* forwards traffic according to the *IP destination address*,
- A *firewall* forwards, shapes, discards, etc. according to a n-tuple (IP src / dst address, L4-L7 headers).

Good news : new routers have firewall features !

Demonstration architecture

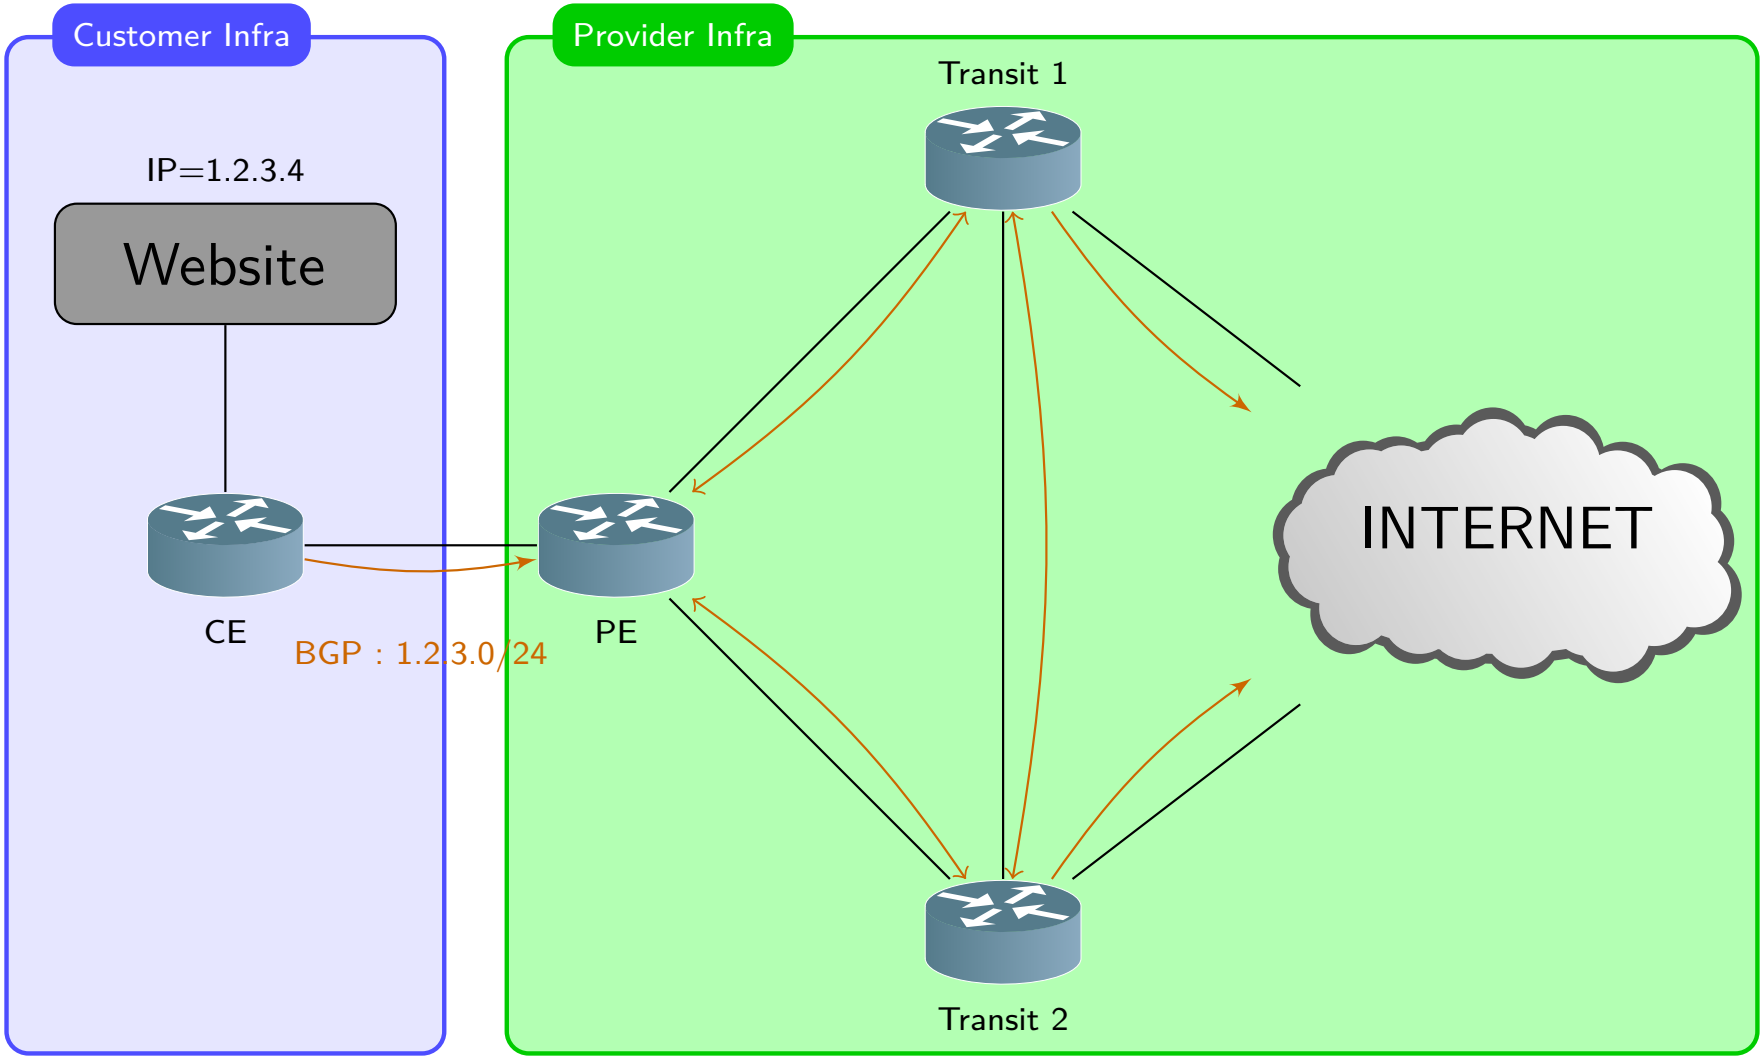

Demonstration architecture

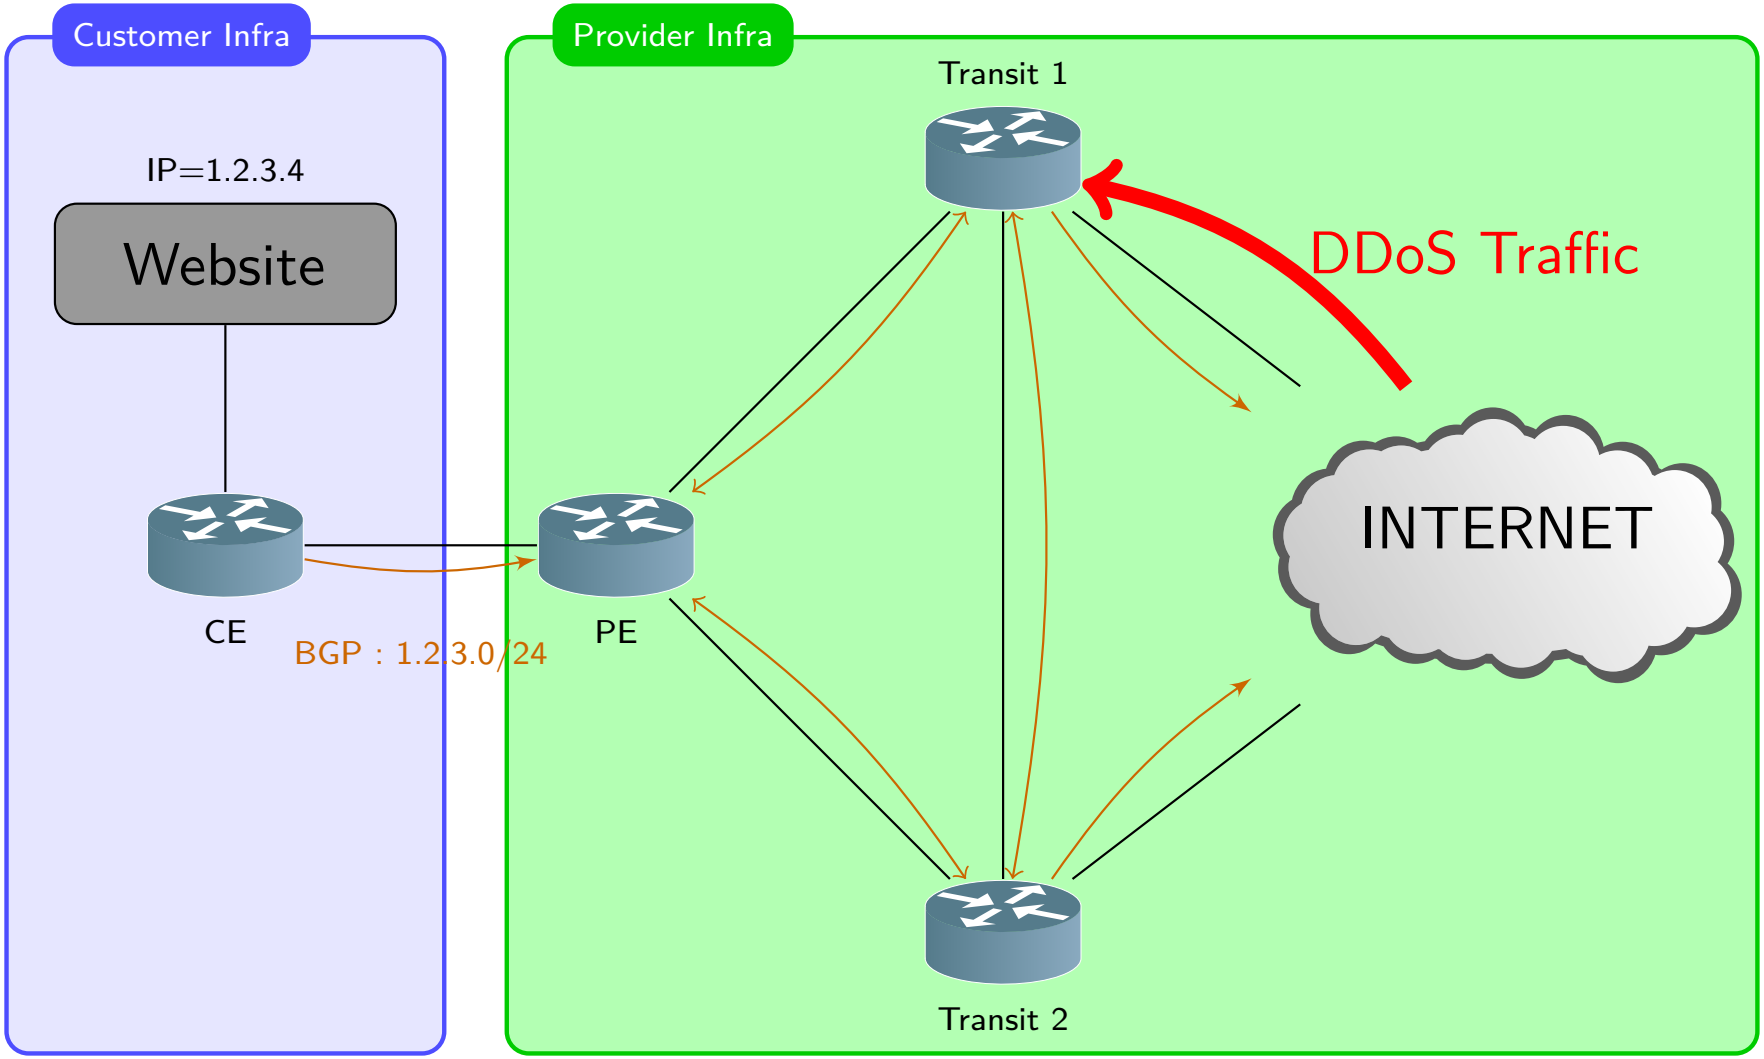

Demonstration architecture

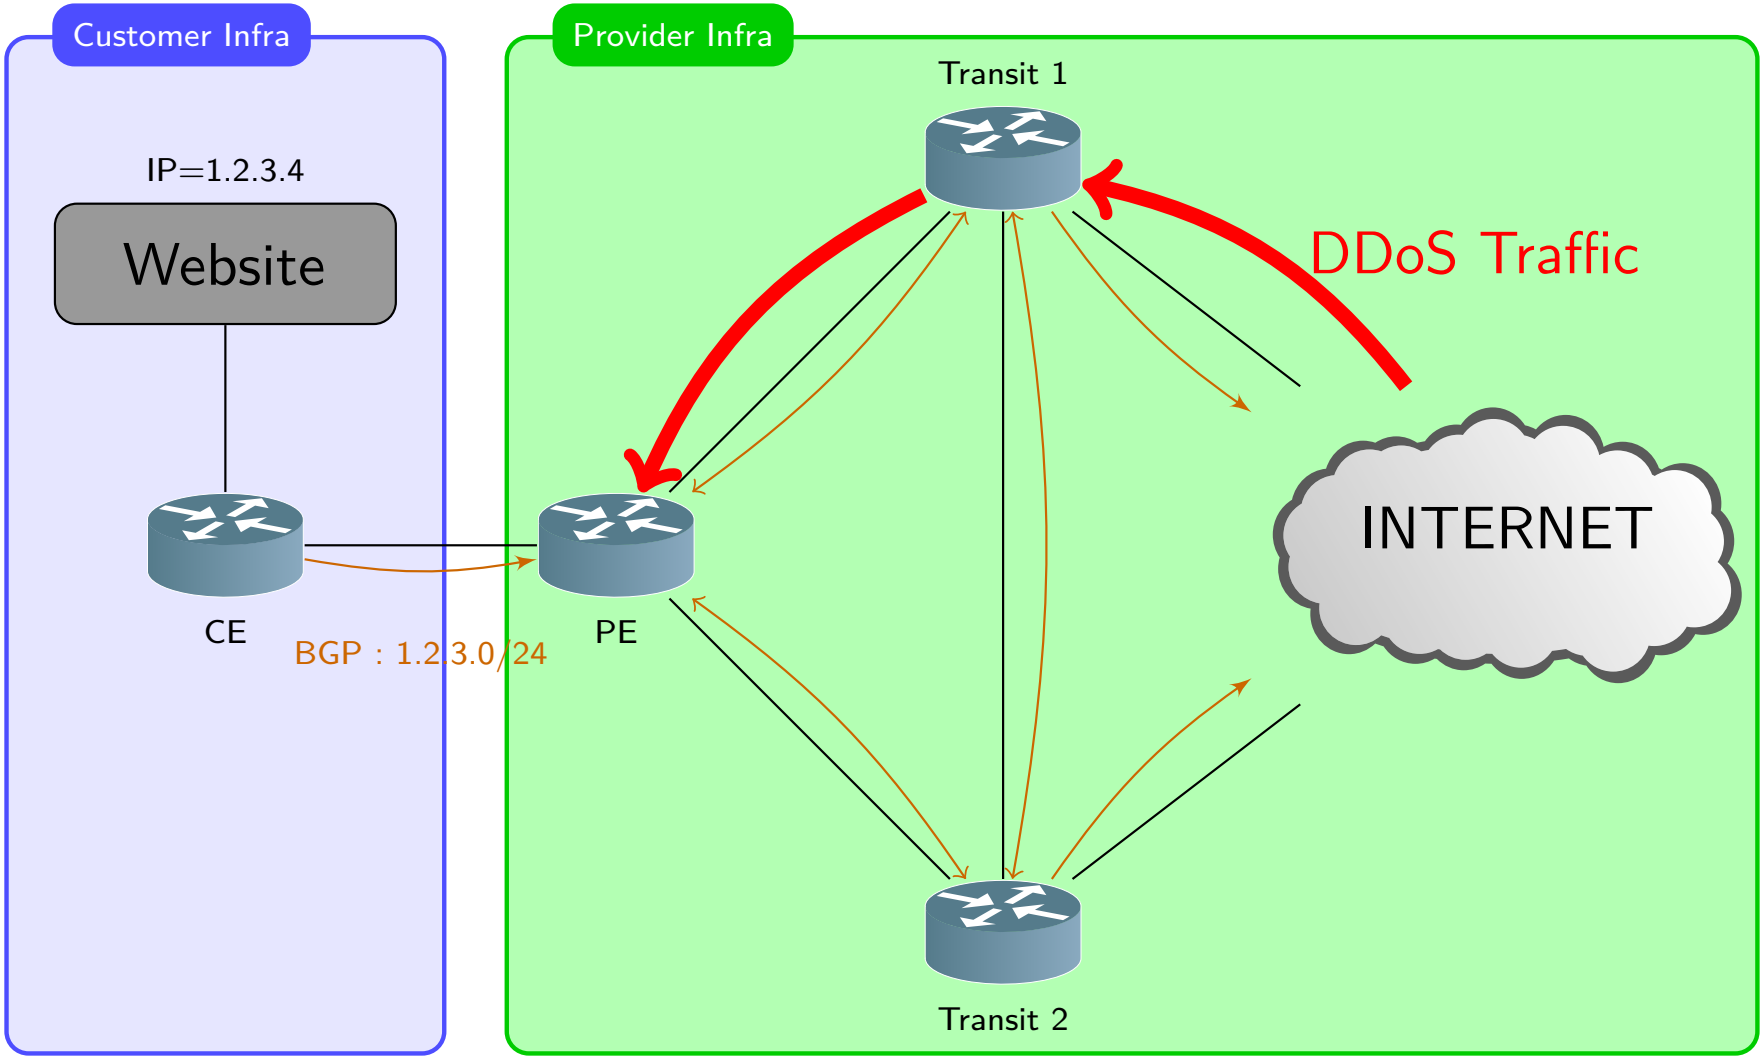

Demonstration architecture

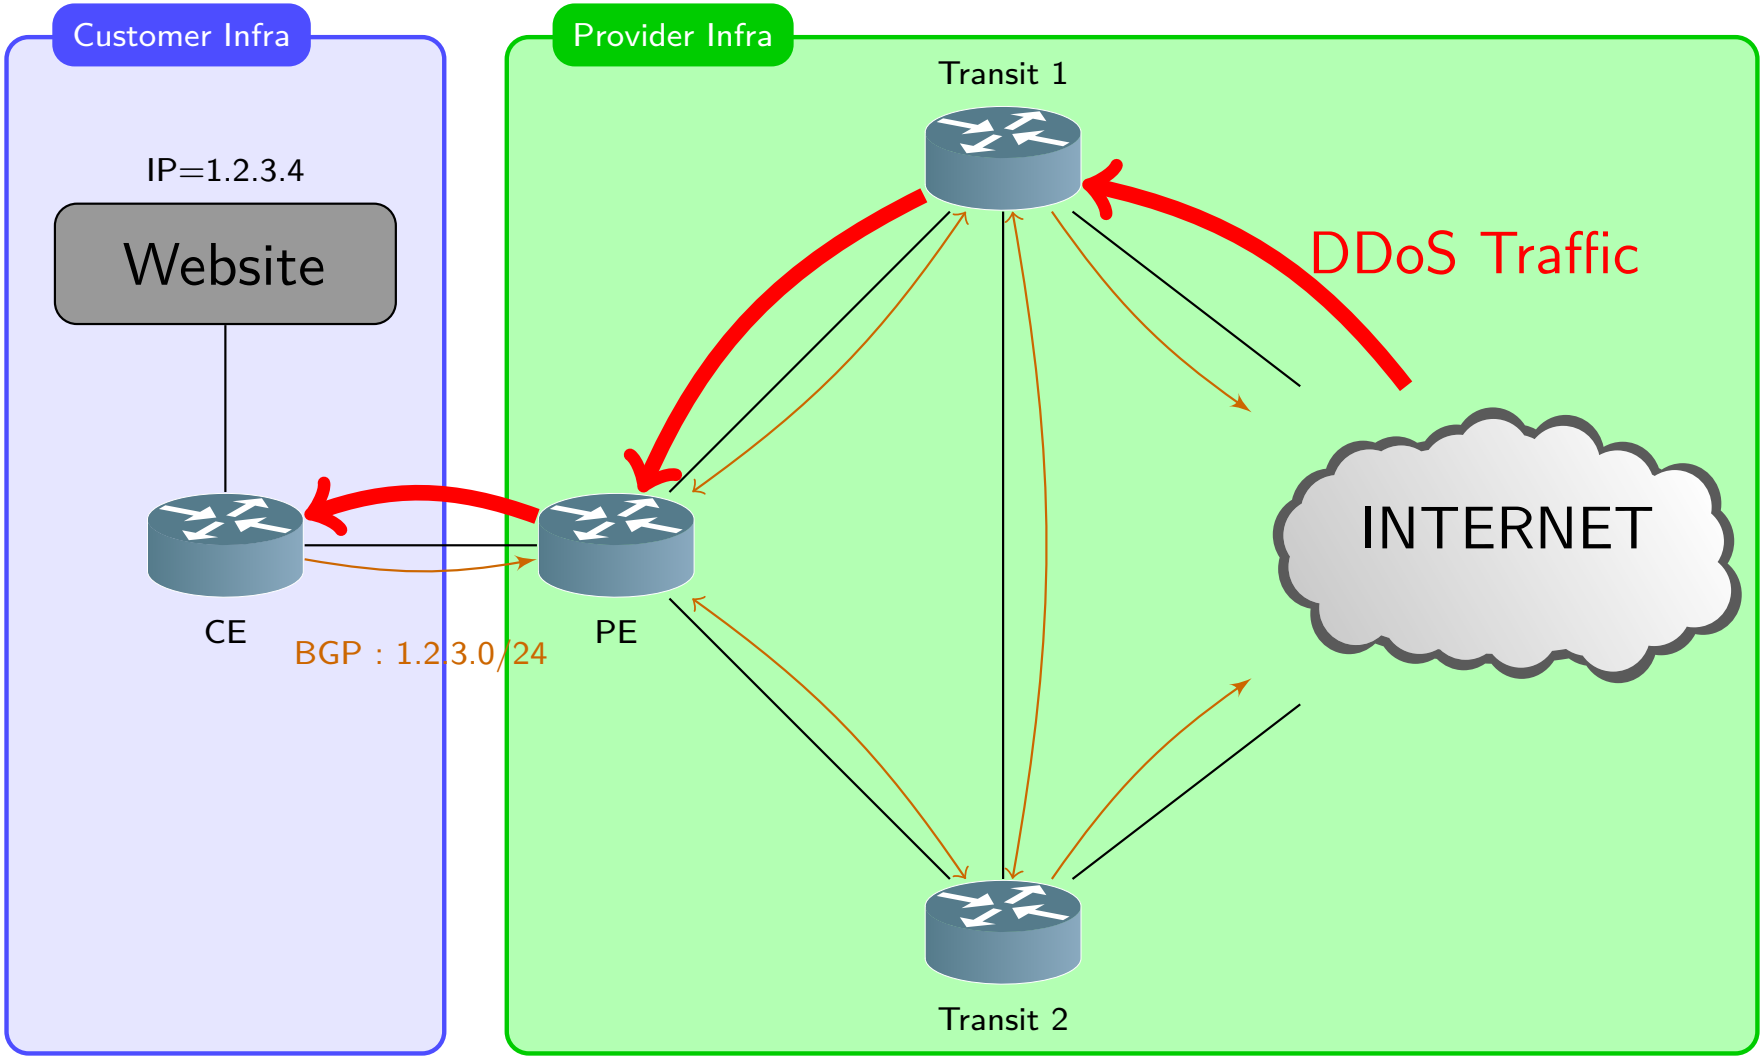

Demonstration architecture

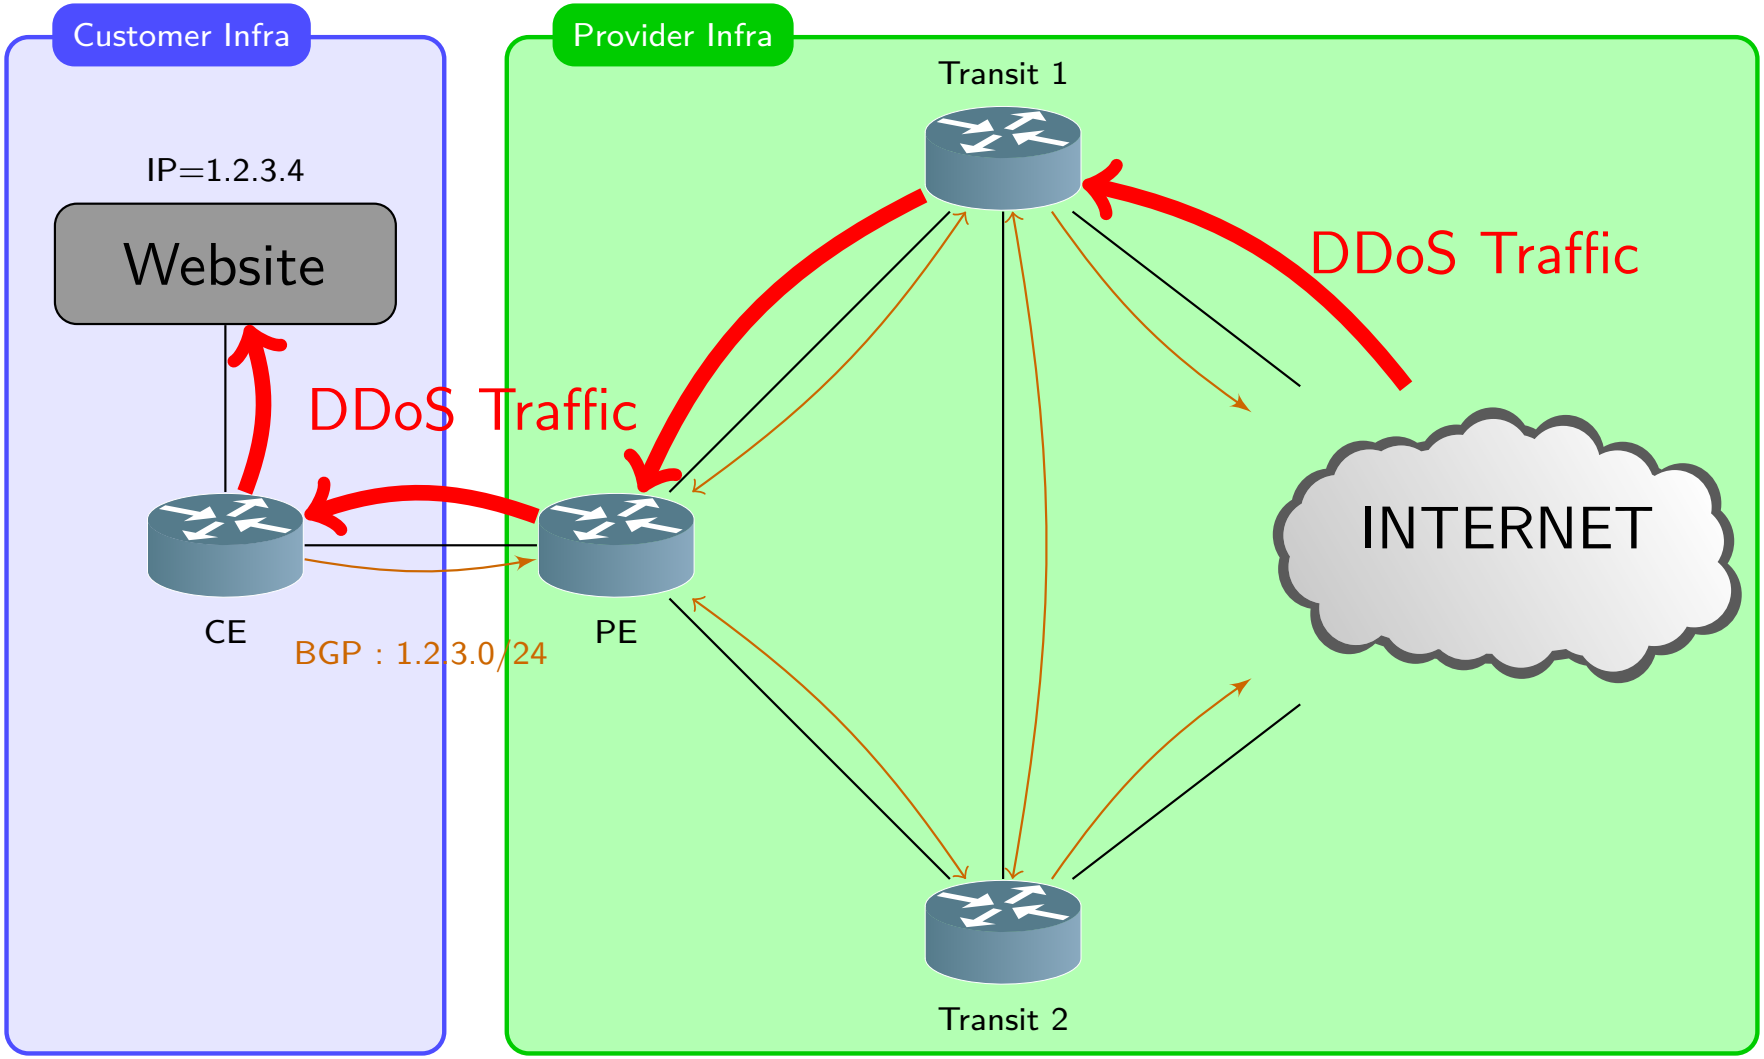

Demonstration architecture

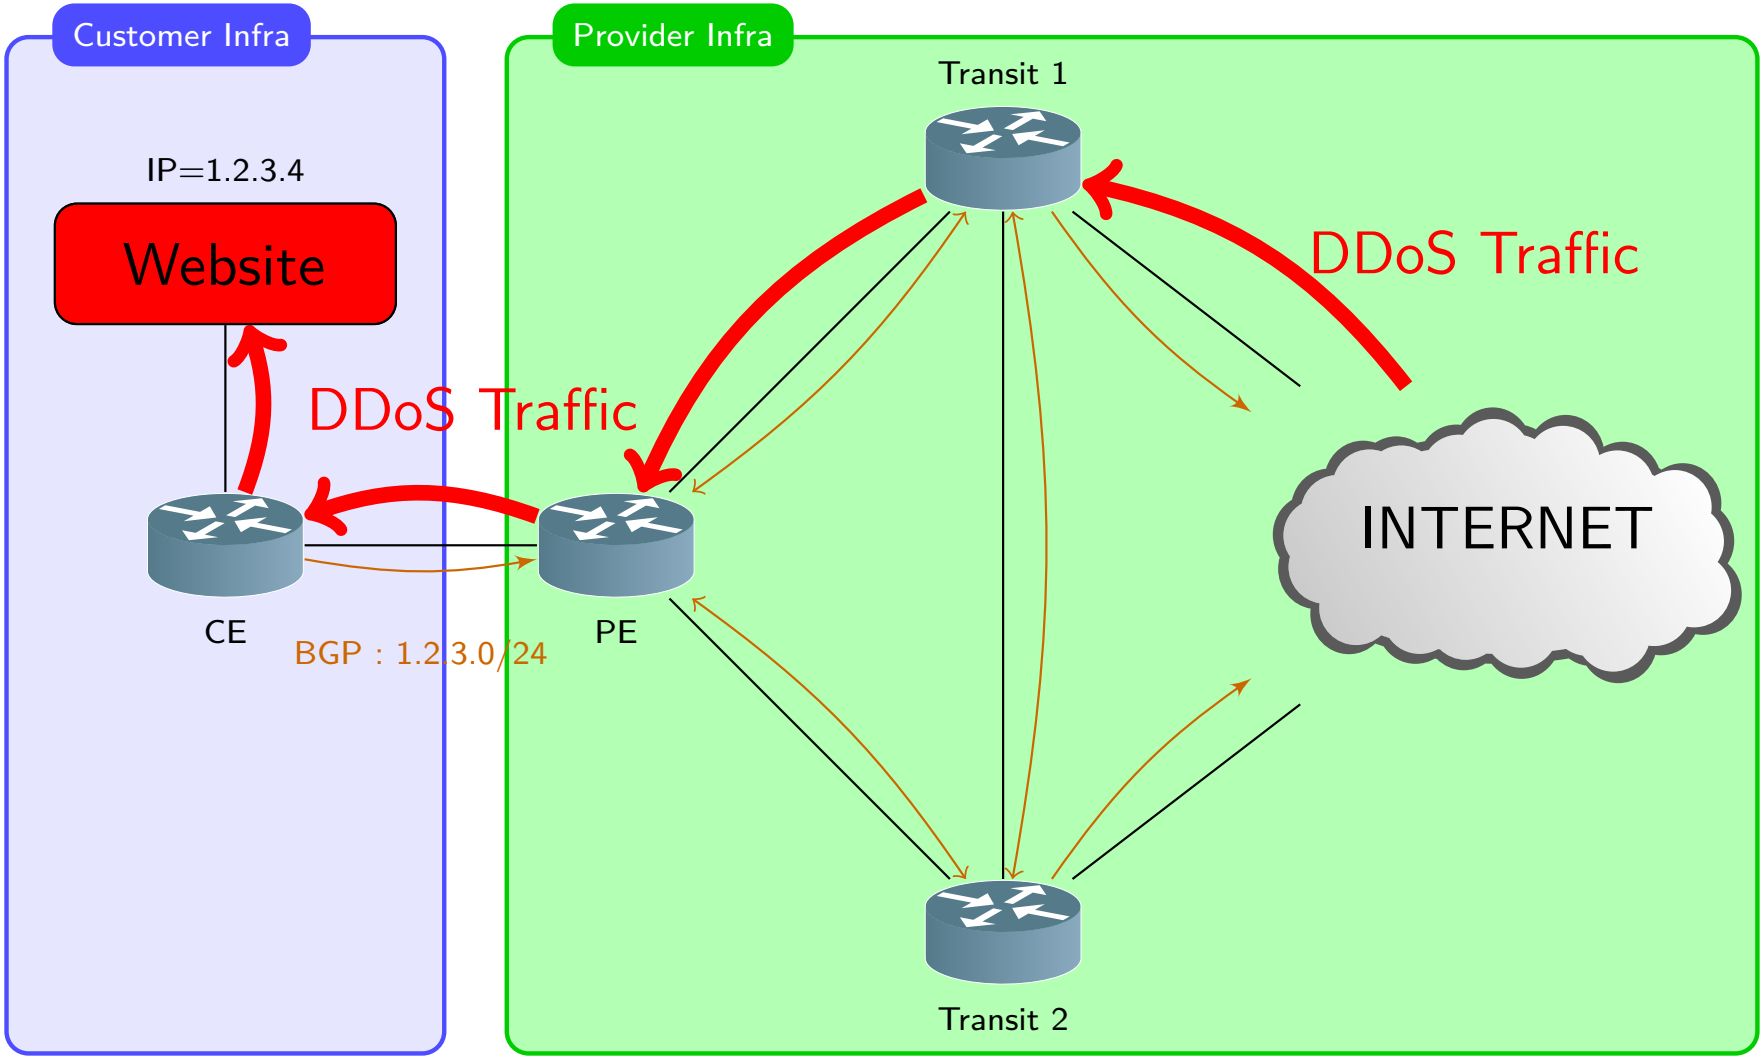

Remotely Triggered Black Hole

Time to
use the blackhole community given by the provider (here 64500:666) !

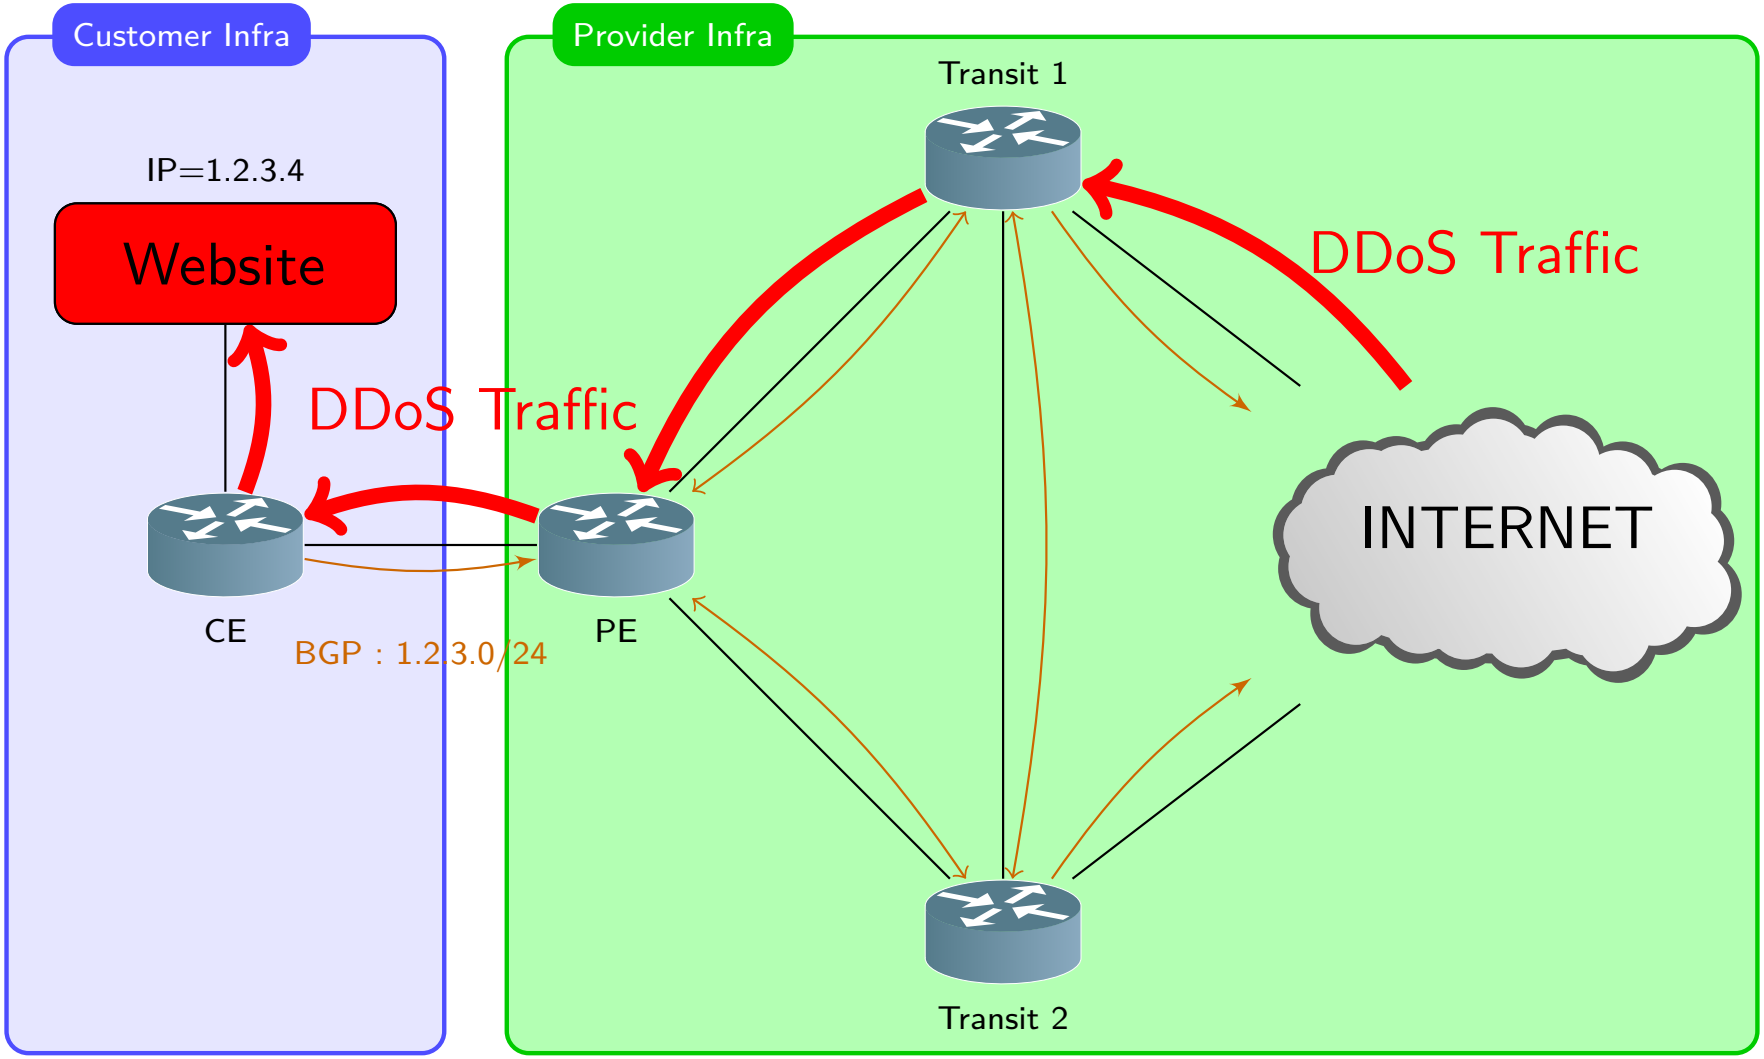

Remotely Triggered Black Hole

Time to
use the blackhole community given by the provider (here 64500:666) !

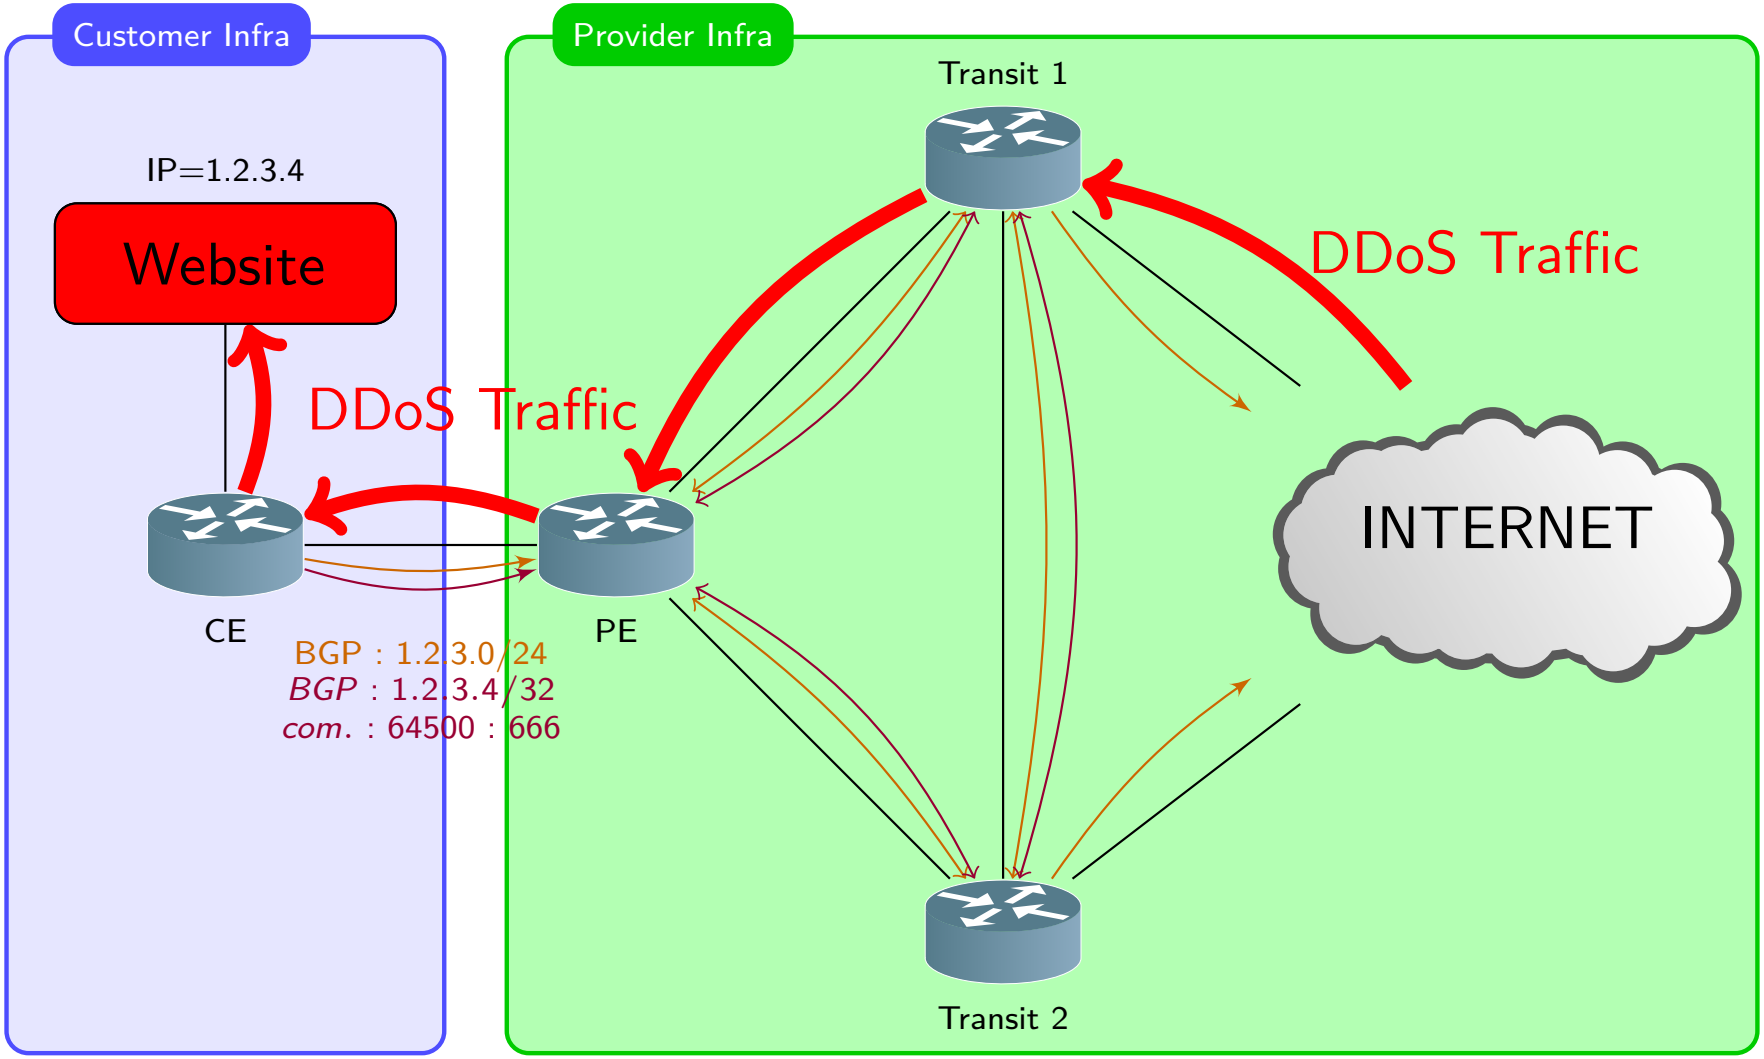

Remotely Triggered Black Hole

Time to

use the blackhole community given by the provider (here 64500:666) !

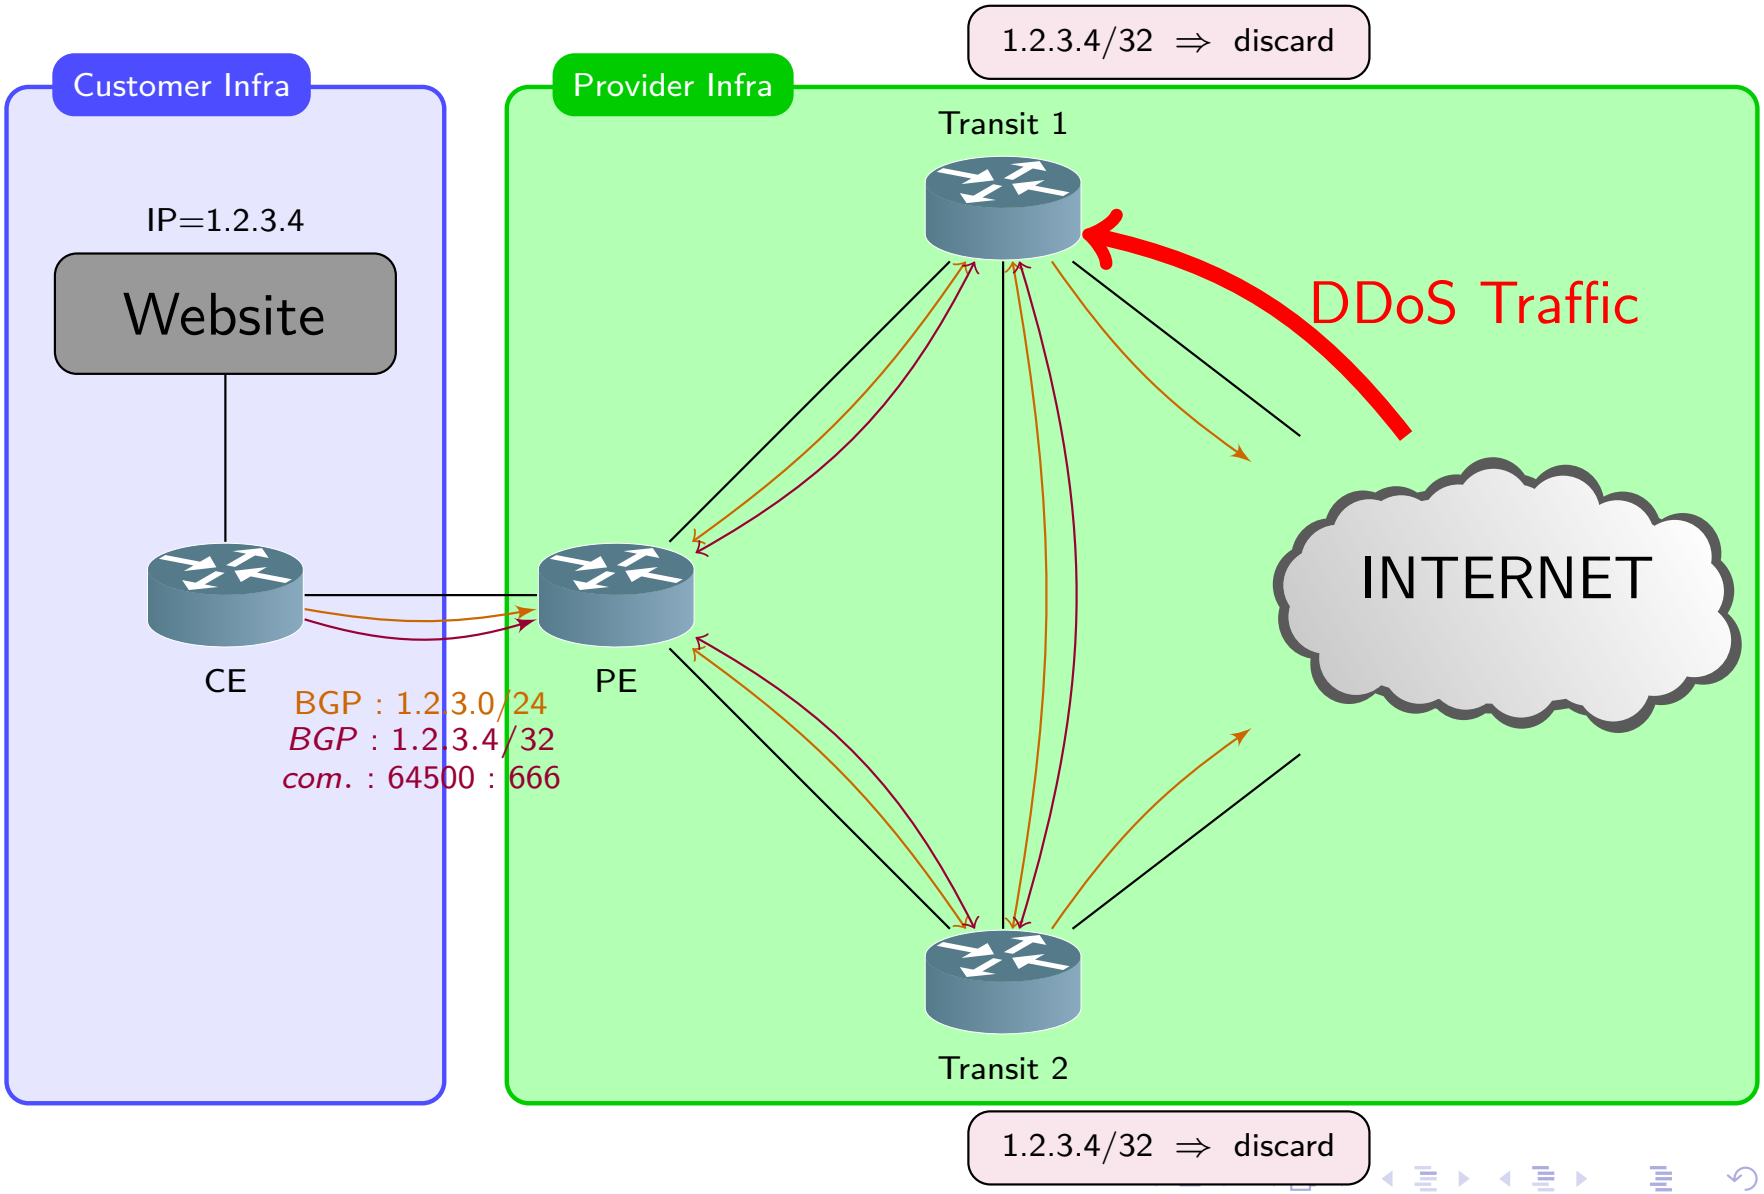

Remotely Triggered Black Hole

Great, my website is 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...

Policy Based Routing

Definition of conditions

- Source / Destination address
- Protocol
- Packet Size, ...

Definition of actions

- Log
- Discard
- Rate-Limit, ...

Sounds nice...

Policy Based Routing

Configuration on the interface

```
xe-0/2/0 {
    description "Transit Interface";
    unit 0 {
        family inet {
            filter {
                input my-ddos-filter;
            }
            address 6.7.8.9/30;
        }
    }
}
```

Policy Based Routing

Configuration of the policy

```
term apnic-udp {
    from {
        source-address {
            le-prefixe-APNIC/24;
        }
        destination-address {
         1.2.3.4/32;
        }
        protocol udp;
    }
    then {
        count my-ddos-filter-counter;
        policer 5m-bw-limit;
    }
}
term everything {
    then accept;
}
```

Policy Based Routing

Interesting feature

- Done in hardware for most carrier grade routers
- Can be used to filter traffic very precisely

But...

- I need to call my provider
- I need him to accept to run this on every router of its backbone
- I need to call him to remove the rule after !

Okay, it won't happen...

FlowSpec as an alternative

Compares to the other solution, FlowSpec :

- Makes static PBR, dynamic !
- Propagates your PBR rules
- Does not need any new communication channel to spread

How ?

- By using your existing MP-BGP Infrastructure

RFC5575 - Dissemination of Flow Specification Rules

Why BGP ?

- Very easy to add NLRI with MP_REACH_NLRI and MP_UNREACH_NLRI
- Communication channel already setup (full mesh, RR)
- Already used for every kind of NLRI (IPv4,6, VPNv4,6, VPLS sig.)
- Only protocol used with transit customers
- Net Eng. already know perfectly BGP !

RFC5575 - Dissemination of Flow Specification Rules

Defines new NLRI (AFI=1, SAFI=133)

Components

1. Source Prefixe (unique)
2. Destination Prefix (unique)
3. IP Protocol (multiple)
4. Port (multiple)
5. Destination Port (multiple)
6. Source Port (multiple)
7. ICMP Type
8. ICMP code
9. TCP Flags
10. Packet length
11. DSCP
12. Fragment

RFC5575 - Traffic filtering actions

Actions are defined in extended communities

type	extended community	encoding
0x8006	traffic-rate	2-byte asn, 4-byte float
0x8007	traffic-action	bitmask
0x8008	redirect	6-byte Route Target
0x8009	traffic-marking	DSCP value

RFC5575 - Principles

How the RFC describes the architecture :

- Your customer already announces you its own prefixes (family inet)
- He advertises inetflow NLRI if the destination address matches (or is more specific) its own announced prefixes. (*validation principle*)
- iBGP propagates the information all over your backbone.

Issues with this

- Your customer needs to support this new family
- It's sometimes hard to setup a simple inet eBGP session. . . Forget about inetflow. . .
- Determining the policy is often complex

RFC5575 - Architecture

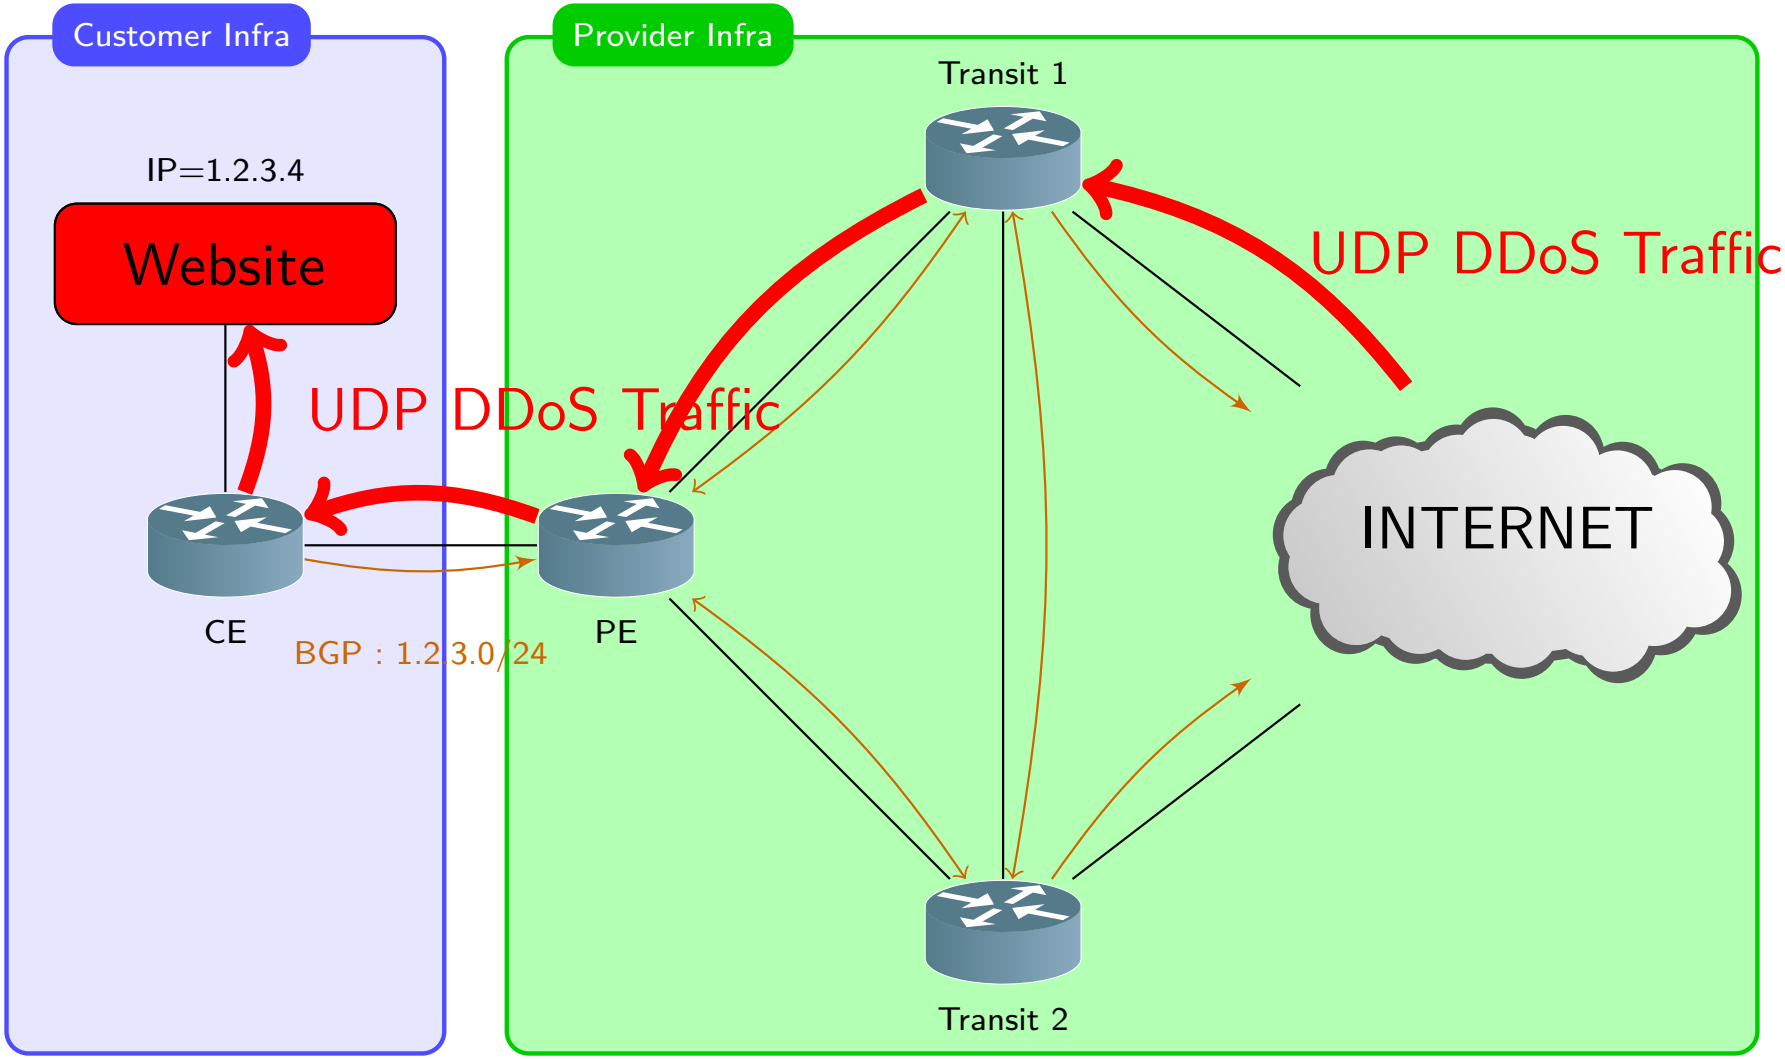

RFC5575 - Architecture

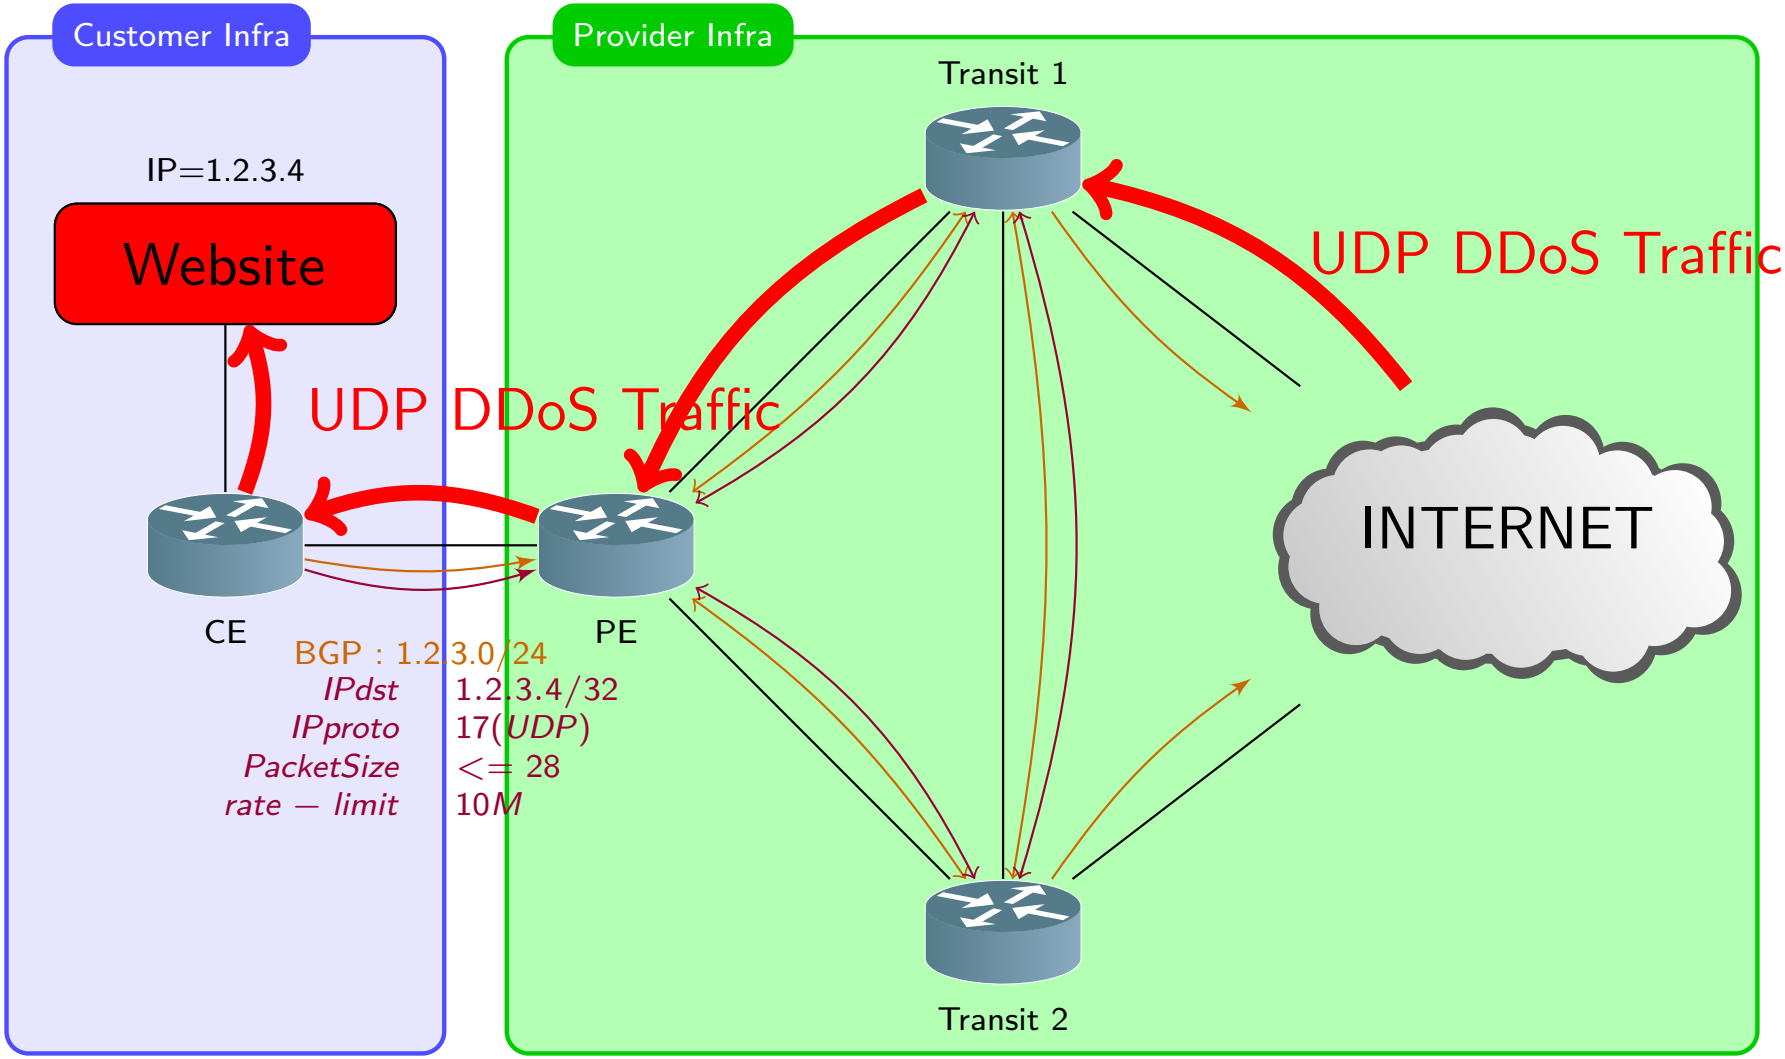

RFC5575 - Architecture

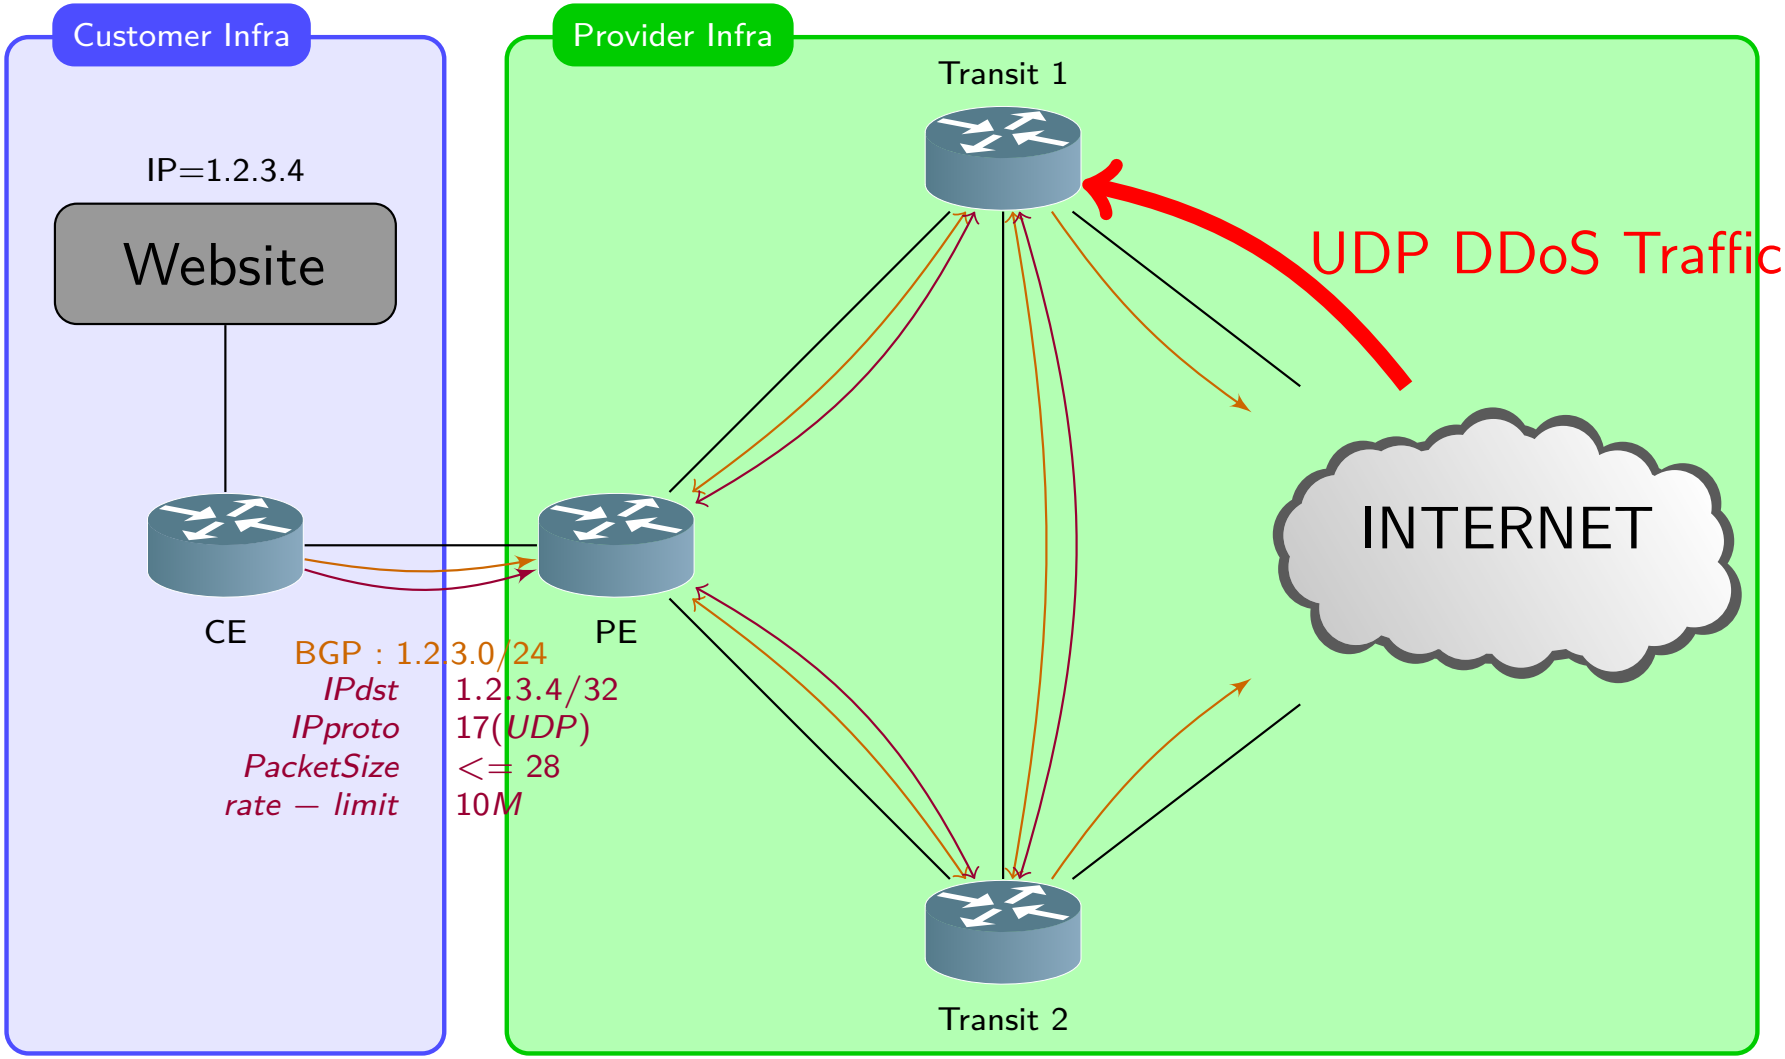

RFC5575 - Architecture

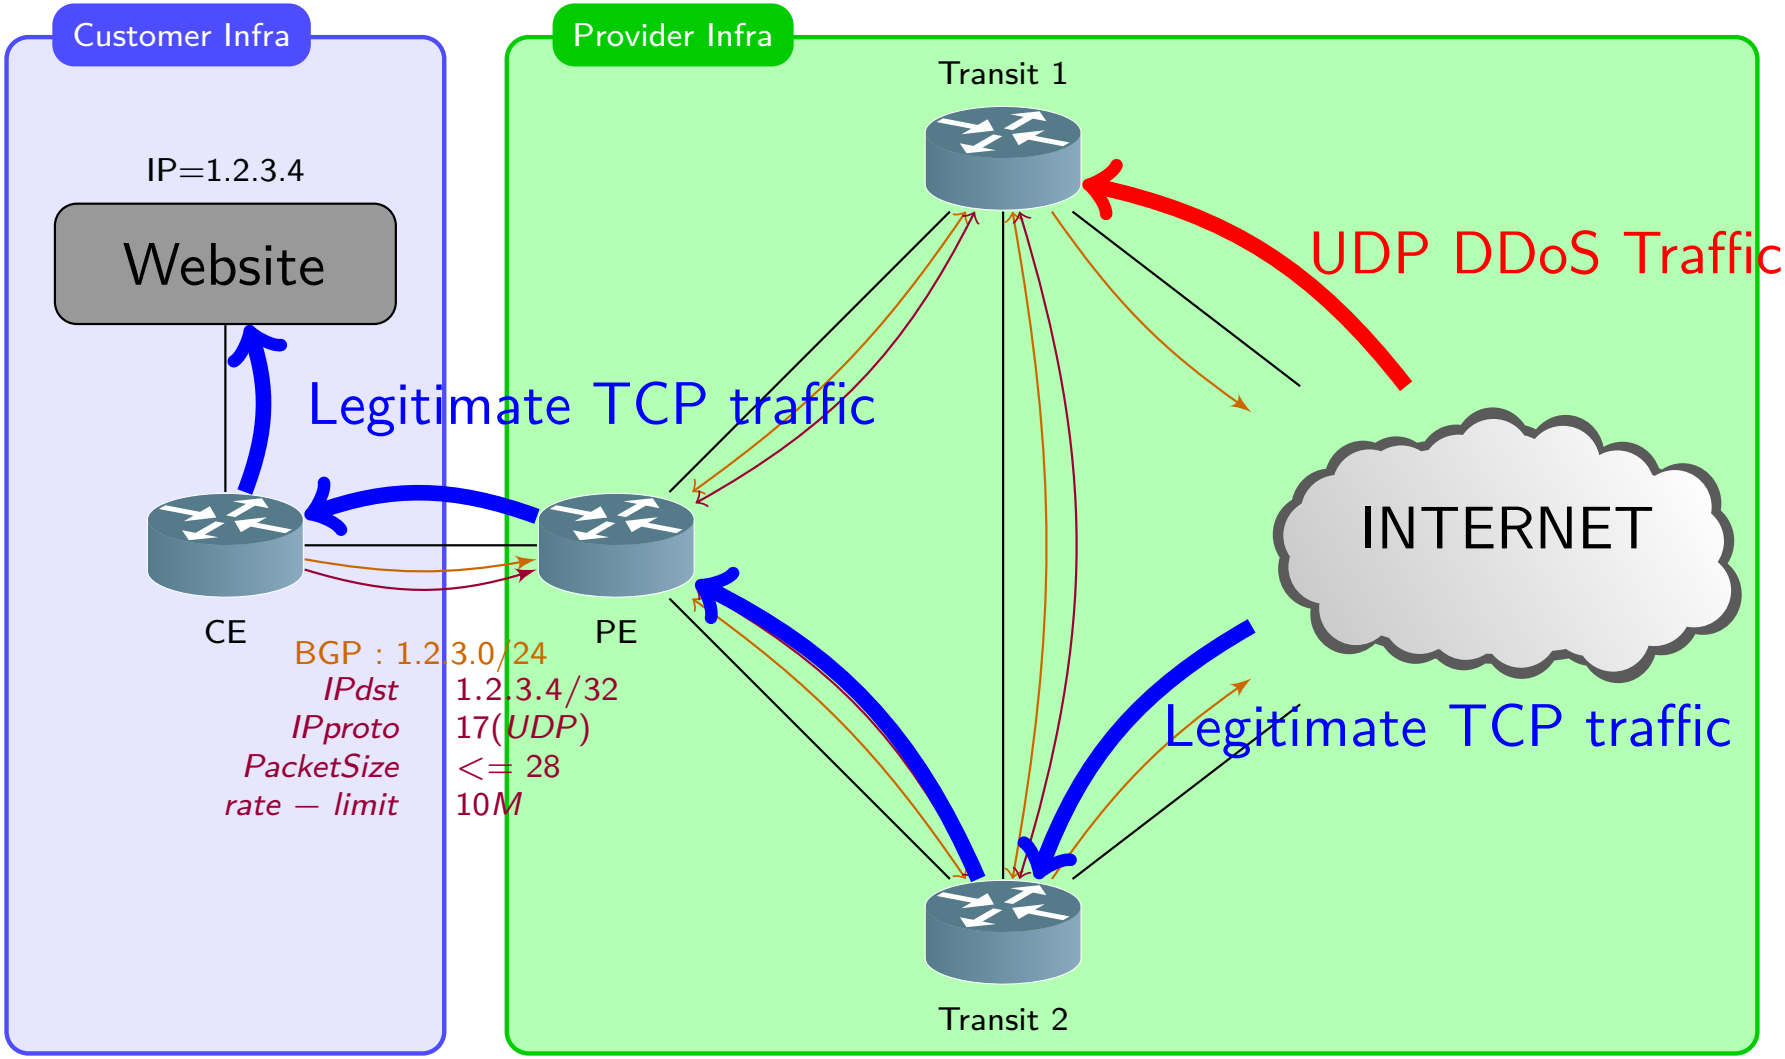

Real life architecture

This architecture is not deployed...

- You DO NOT trust your customer
- We have enough BGP related bugs (Cisco, Juniper and Redback yesterday)
- So, we won't enable inetflow on eBGP sessions...

What do we do instead ?

- Centralized inetflow speaker
- Meshed with core routers
- Only one peer allowed to announce inetflow
- Considered "trusted" by the network (*no-validate*)

Real life architecture

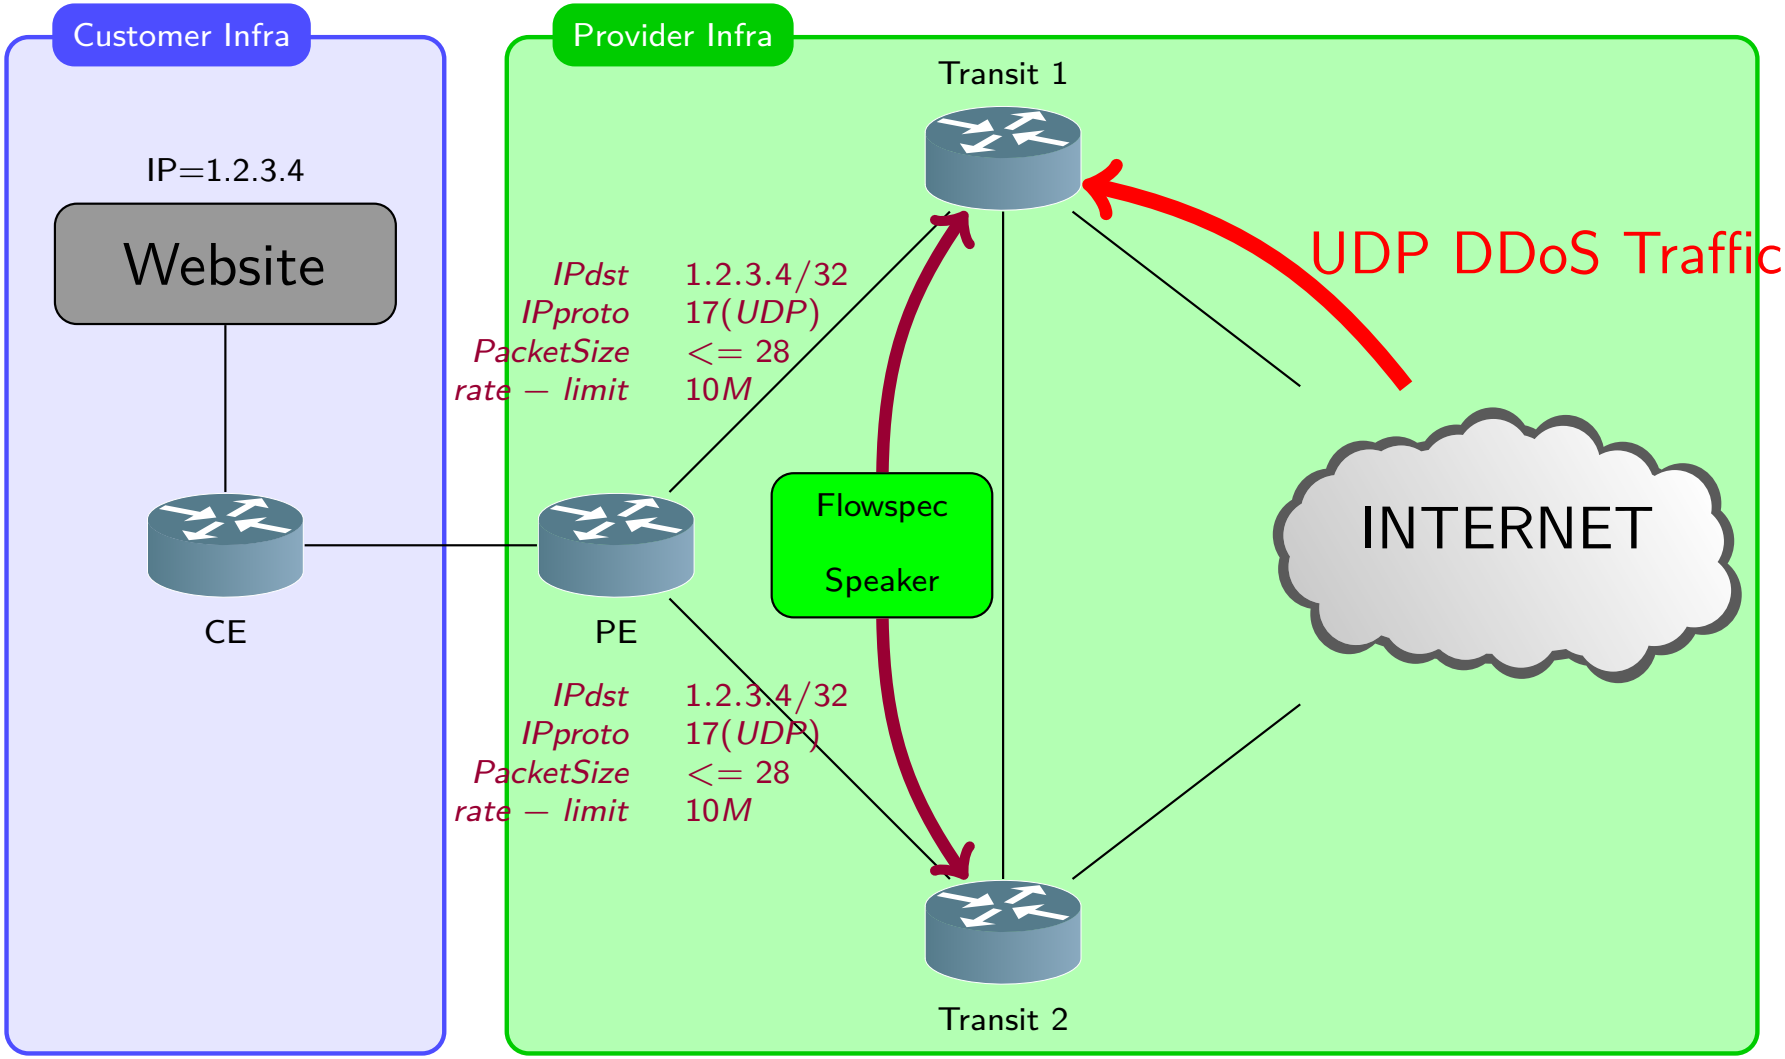

Real life architecture

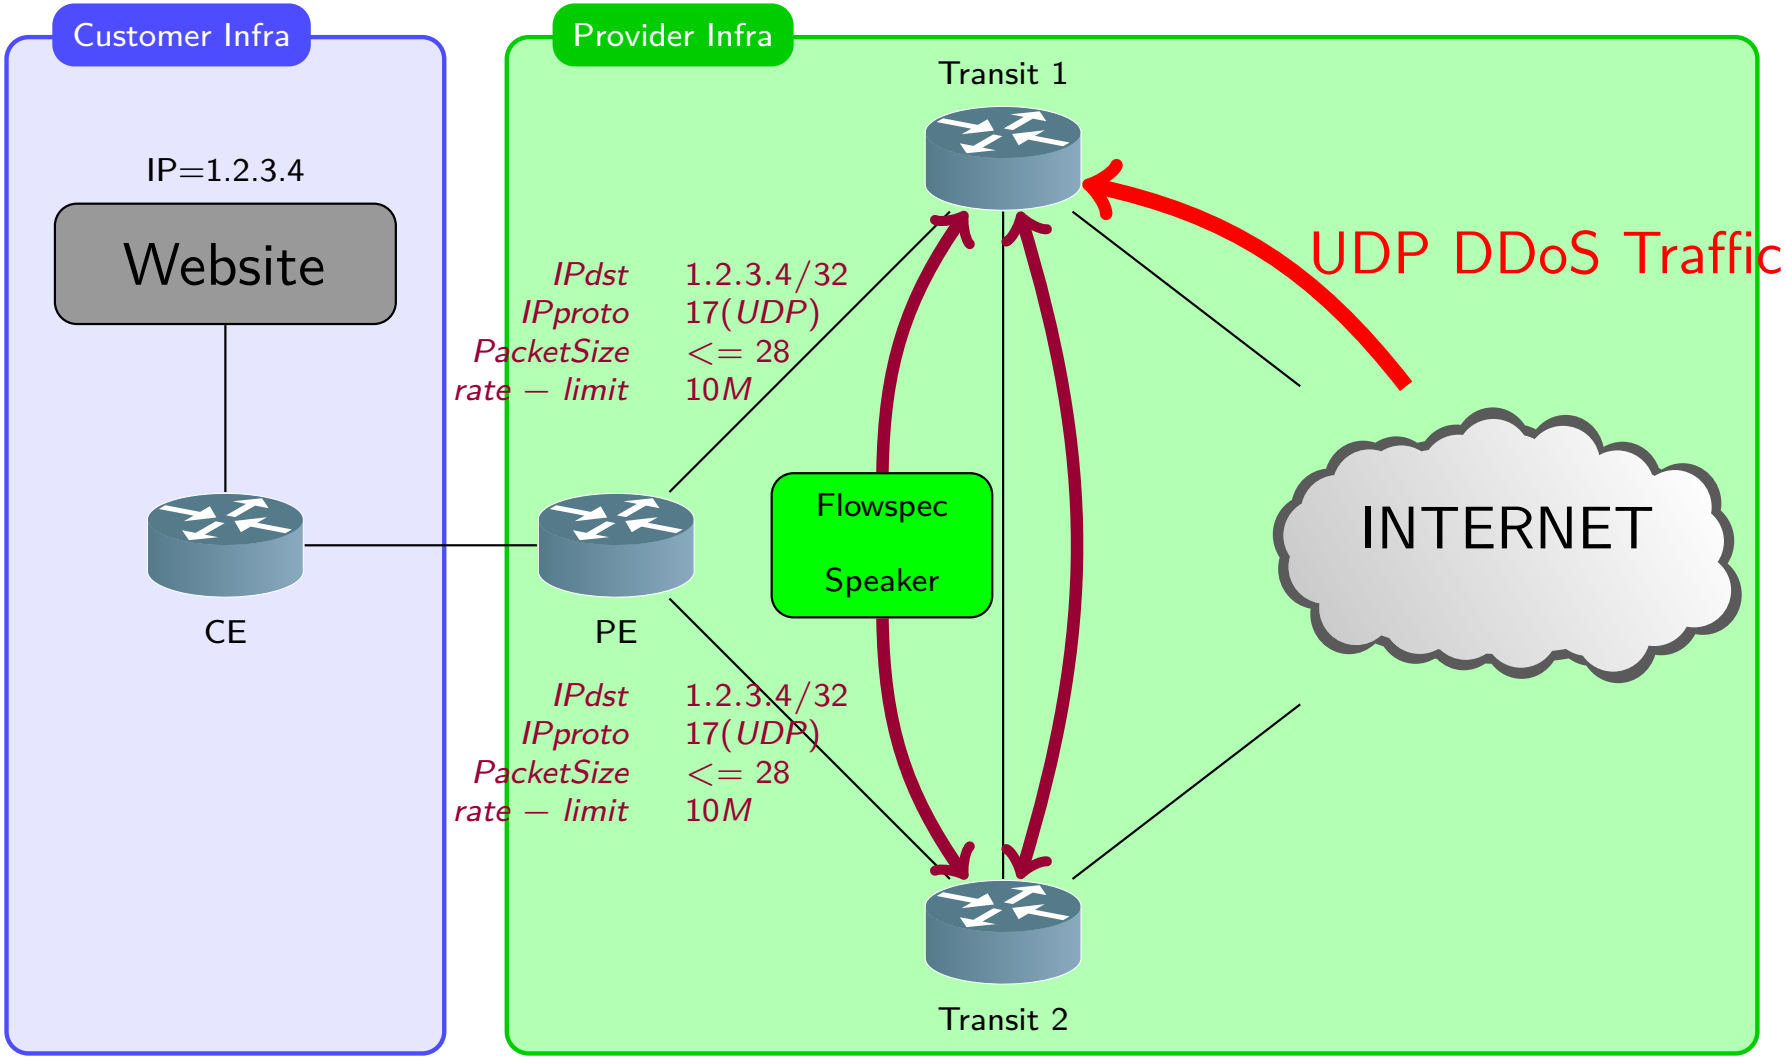

Real life architecture

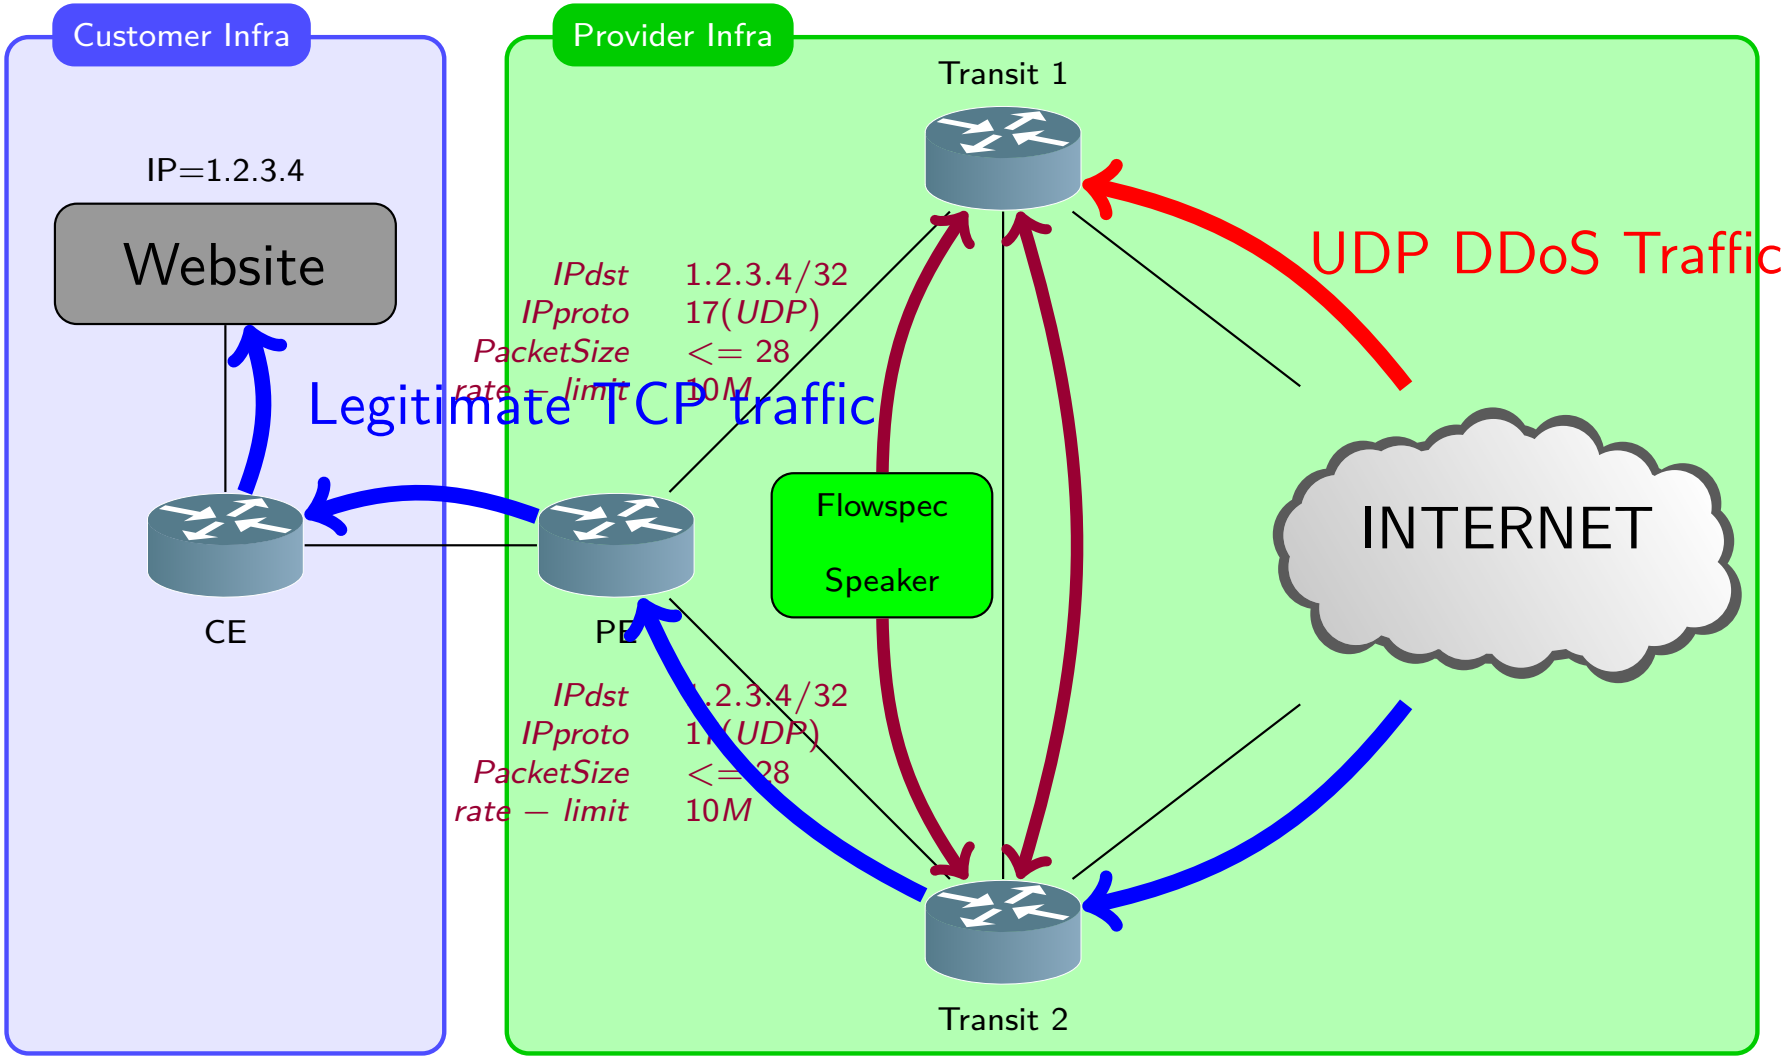

Few words about traffic redirection

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

traffic-redirect

- lets you redirect traffic in a VRF (which import the specified rt)
- lets you change dynamically the path of a flow without injecting BGP more specific routes

Great tool for cleaning DDoS traffic with a DPI probe without interaction with your global forwarding table

Configuration extract

JunOS configuration example

```
lab@lab-mx80> show configuration protocols bgp
group flowspec-src {
    import then-accept;
    family inet {
        unicast;
        flow {
            no-validate then-accept;
        }
    }
    peer-as 8218;
    neighbor 192.168.200.1;
}
```

Command samples

Useful JunOS show commands

```
lab@lab-mx80> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                   1          1          0          0          0          0
inetflow.0               1          1          0          0          0          0
Peer                     AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn State|#Act
192.168.200.1          8218          4          3        0          0          15 Establ
  inet.0: 1/1/1/0
  inetflow.0: 1/1/1/0
```

```
lab@lab-mx80> show route receive-protocol bgp 192.168.200.1 table inetflow.0
inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
* 192.168.200/24,192.168.200.1,proto=17,len>=0&<=28/term:1 (1 entry, 1 announced)
     Accepted
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: traffic-rate:0:9600
```

```
lab@lab-mx80> show firewall filter __flowspec_default_inet__
Filter: __flowspec_default_inet__
Counters:
Name                                                Bytes              Packets
192.168.200/24,192.168.200.1,proto=17,len>=0&<=28  15065136                19464
```

How to play with FlowSpec at home ?

FlowSpec speaker : Arbor Networks or exabgp

exabgp

- BSD License BGP speaker written by T. Mangin
- supports the entire RFC5575
- JunOS like configuration

FlowSpec listener

Get an Alcatel SR or a Juniper MX

exabgp sample configuration

```
neighbor 192.168.200.2 {
        description "mx80";
        router-id 192.168.200.1;
        local-address 192.168.200.1;
        local-as 8218;
        peer-as 8218;
        graceful-restart 5;

        flow {
                route optional-name-of-the-route {
                        match {
                                source 192.168.200.1/32;
                                destination 192.168.200.50/24;
                               protocol udp;
                               packet-length <29;

                        }
                        then {
                                rate-limit 9600;
                        }
                }
        }
        static {
                route 10.0.5.0/24 {
                        next-hop 192.168.200.1;
                        local-preference 10;
                        community [ 0x87654321 ];
                }
        }
}
```

You want to sell protected IP transit !

What do you need ?

- Traffic analyzer to qualify the attacks (netflow parser)
- Flowspec interface to manage the mitigation
- Long list of prefix-list (EU prefixes, APNIC prefixes, . . .)

But that's not enough

- A lot of attacks can be easily qualified (ICMP flood, UDP flood)
- But others can't (e.g. TCP SYN flood)
- You need a cleaning box (TCP SYN Proxy, URL analyze, . . .)
- And that's really expensive !

What's next ?

- Support IPv6 and VPNv6 (draft-raszuk-idr-flow-spec-v6) ☺
- More vendor support (only Alcatel and Juniper routers today)
- Deeper RFC implementation (Juniper does not support the full RFC)
- New features needed (e.g. traffic mirroring)

Any questions ?

Contact information

- Frédéric Gabut-Deloraine
- fgabut@neotelecoms.com
- NEO TELECOMS