Selective BlackHoling DDoS Damage Control

François Contat

ANSSI

Agence nationale de la sécurité des systèmes d'information

http://www.ssi.gouv.fr/en

October 17th, 2014

DDoS

DDoS

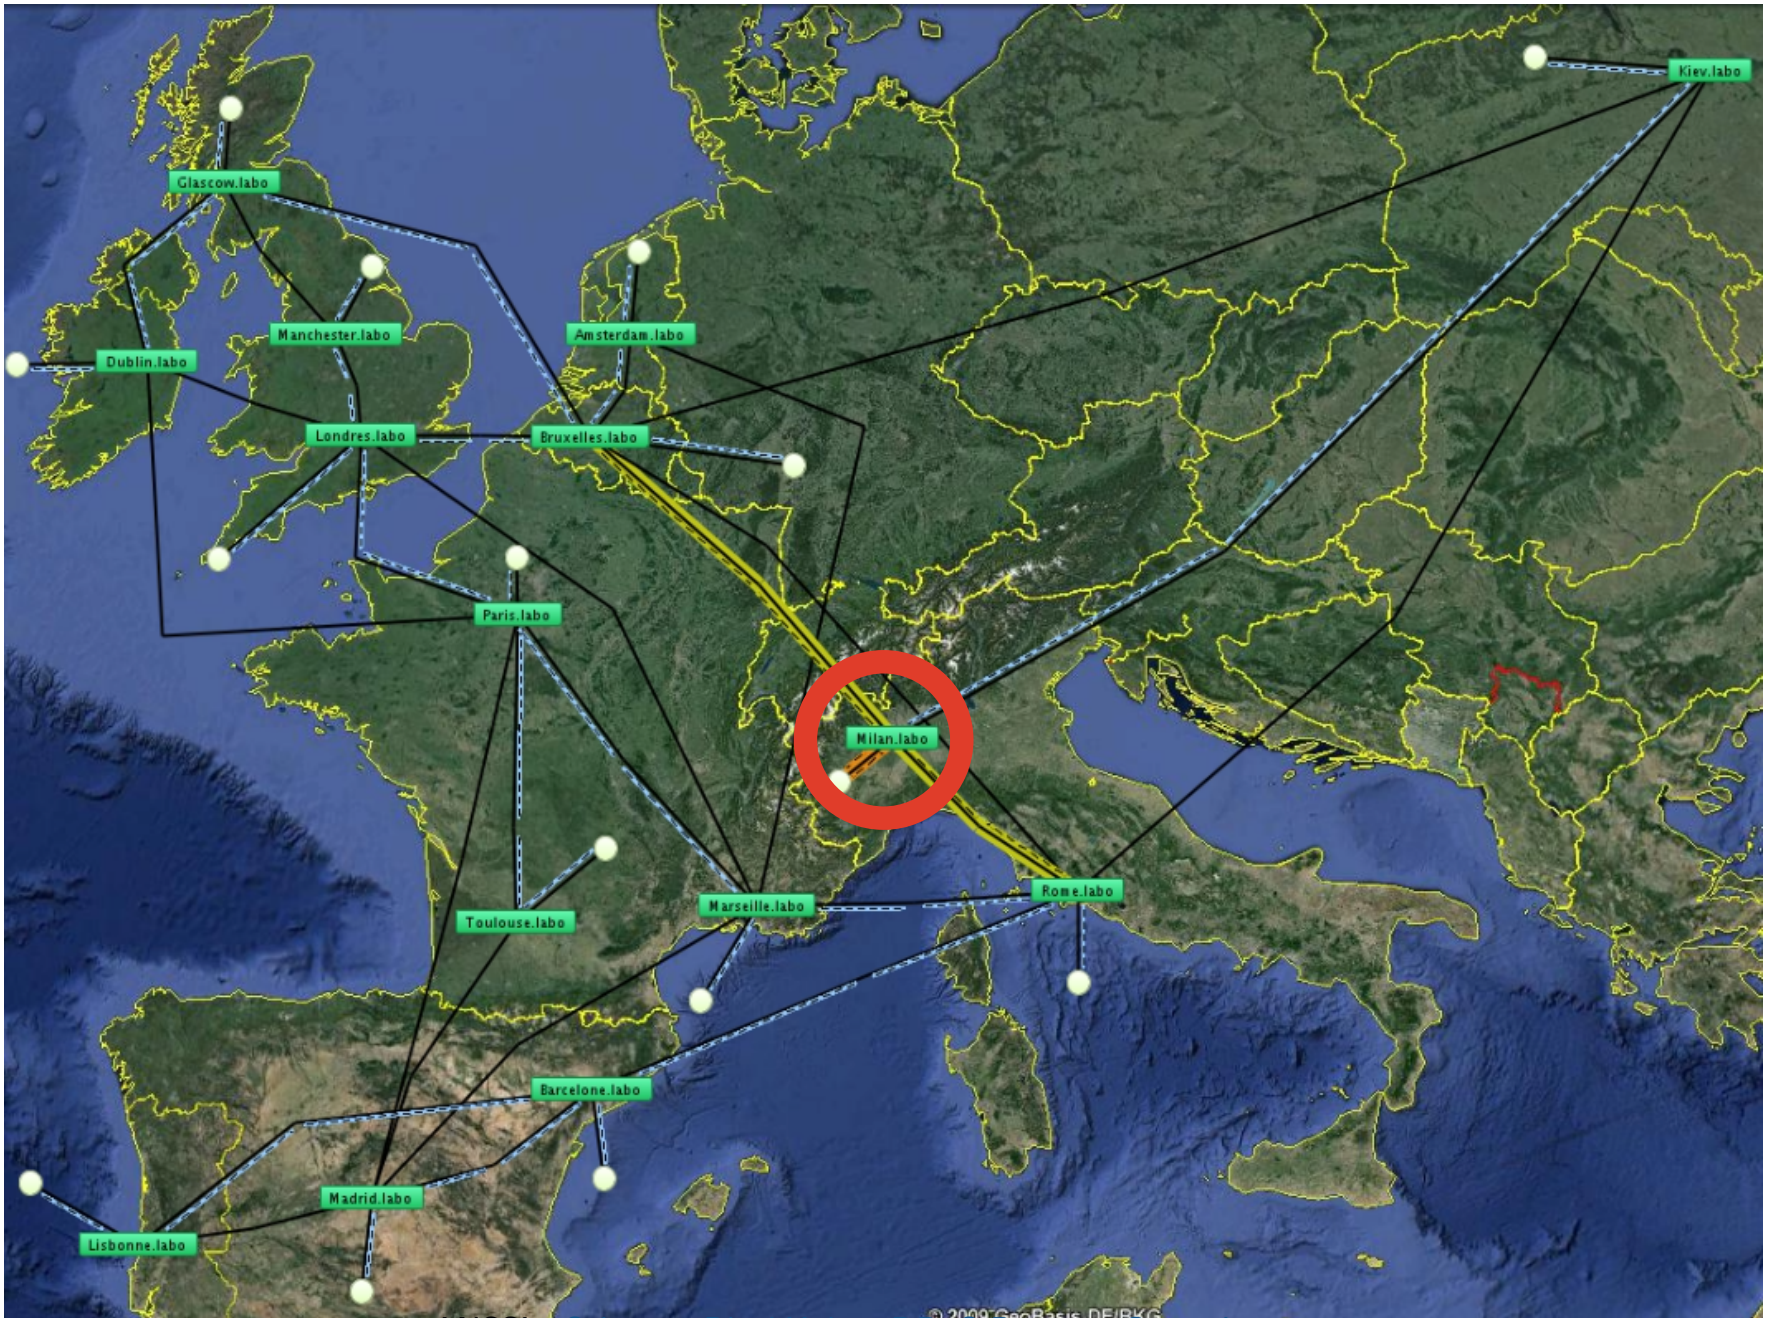
Kiev.labo
Glascow.labo
Manchester.labo
Amsterdam.labo
Dublin.labo
Londres.labo
Bruxelles.labo
Paris.labo
Milan.labo
Rome.labo
Marseille.labo
Toulouse.labo
Barcelone.labo
Madrid.labo
Lisbonne.labo

DDoS Mitigation Solutions - RTBH

RTBH - Remote Triggered BlackHoling

- Null route traffic

DDoS Mitigation Solutions - RTBH

RTBH - Remote Triggered BlackHoling

- Null route traffic
- Trigger the action on a distant provider

DDoS Mitigation Solutions - RTBH

RTBH - Remote Triggered BlackHoling

- Null route traffic
- Trigger the action on a distant provider
- BGP community = transitive attribute added to the BGP route update

DDoS Mitigation Appliances

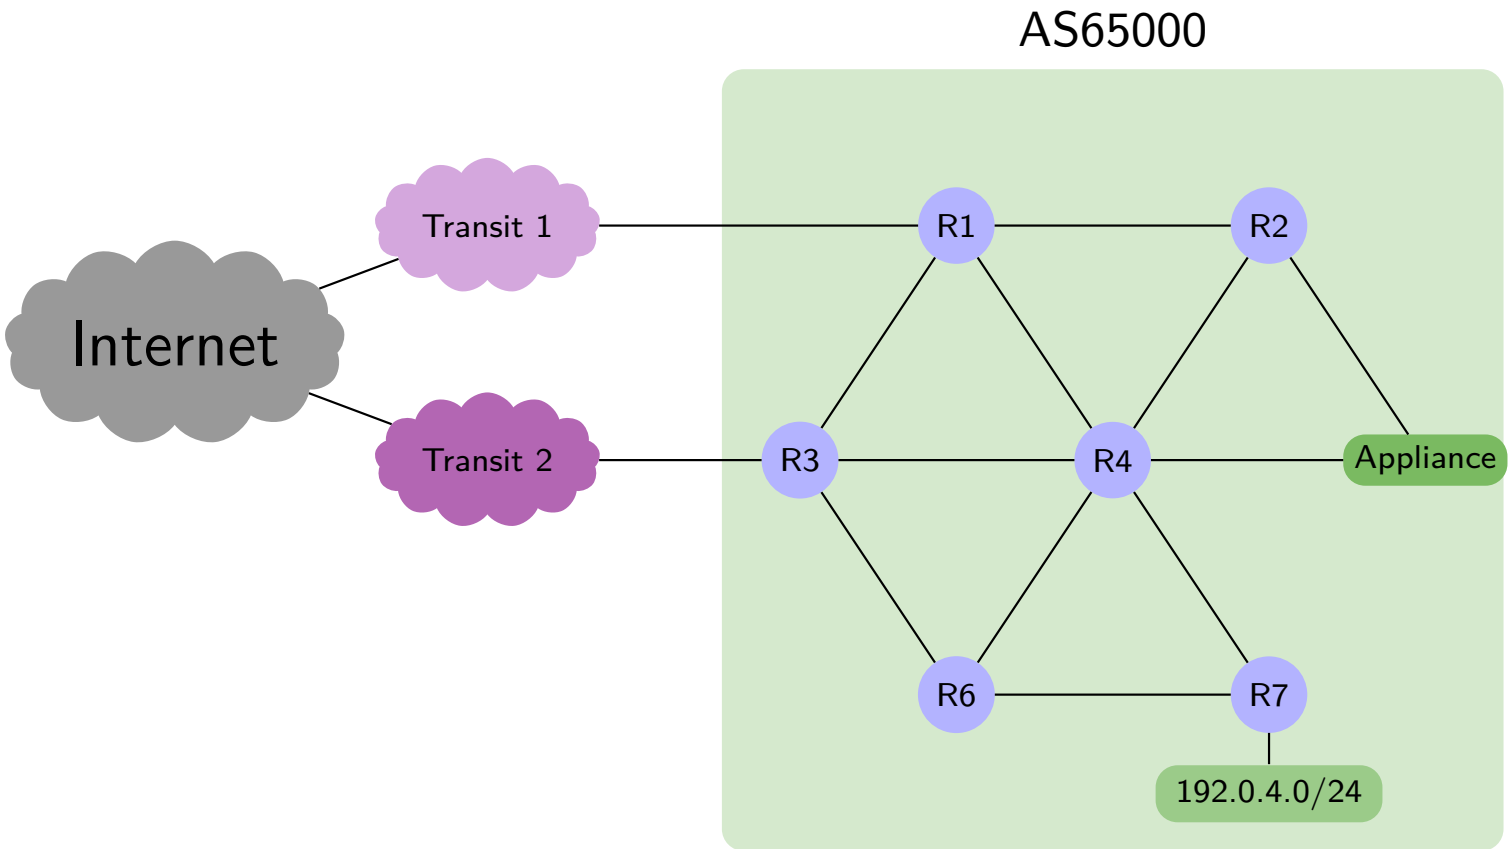

DDoS Mitigation Appliances

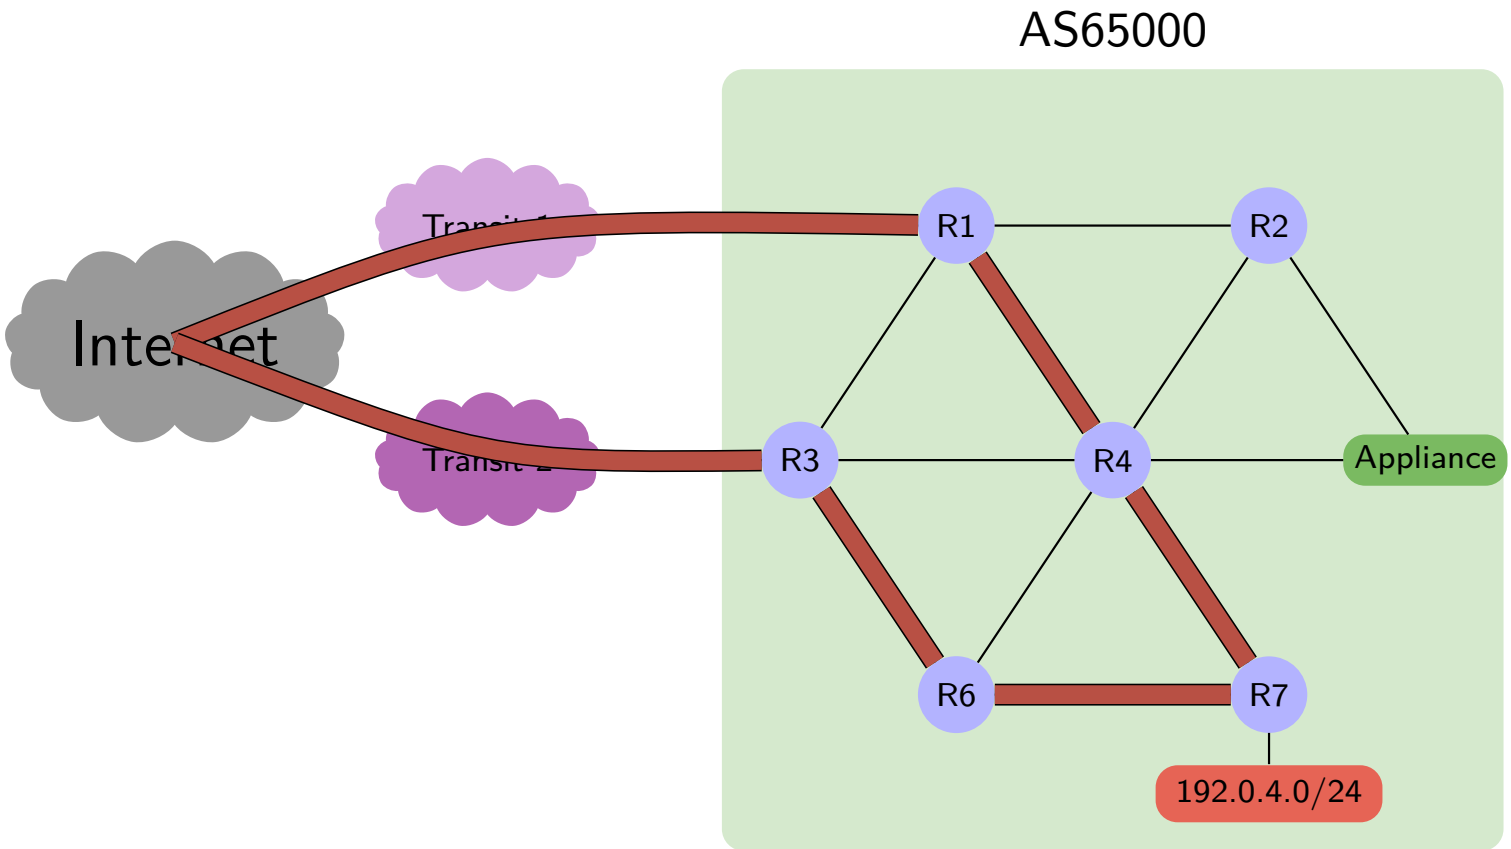

DDoS Mitigation Appliances

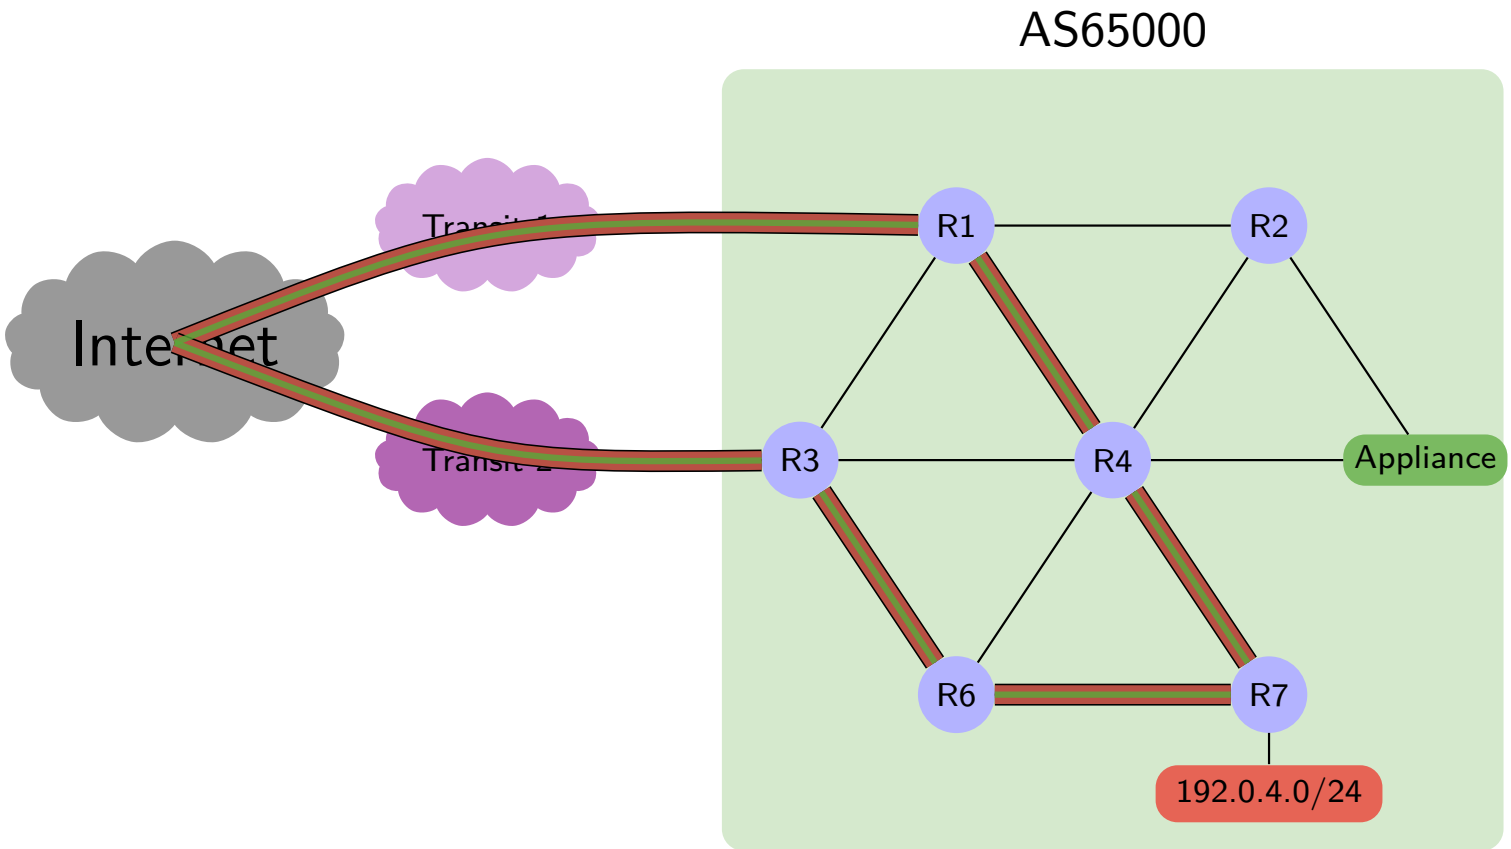

DDoS Mitigation Appliances

- Announce appliance as next-hop inside network

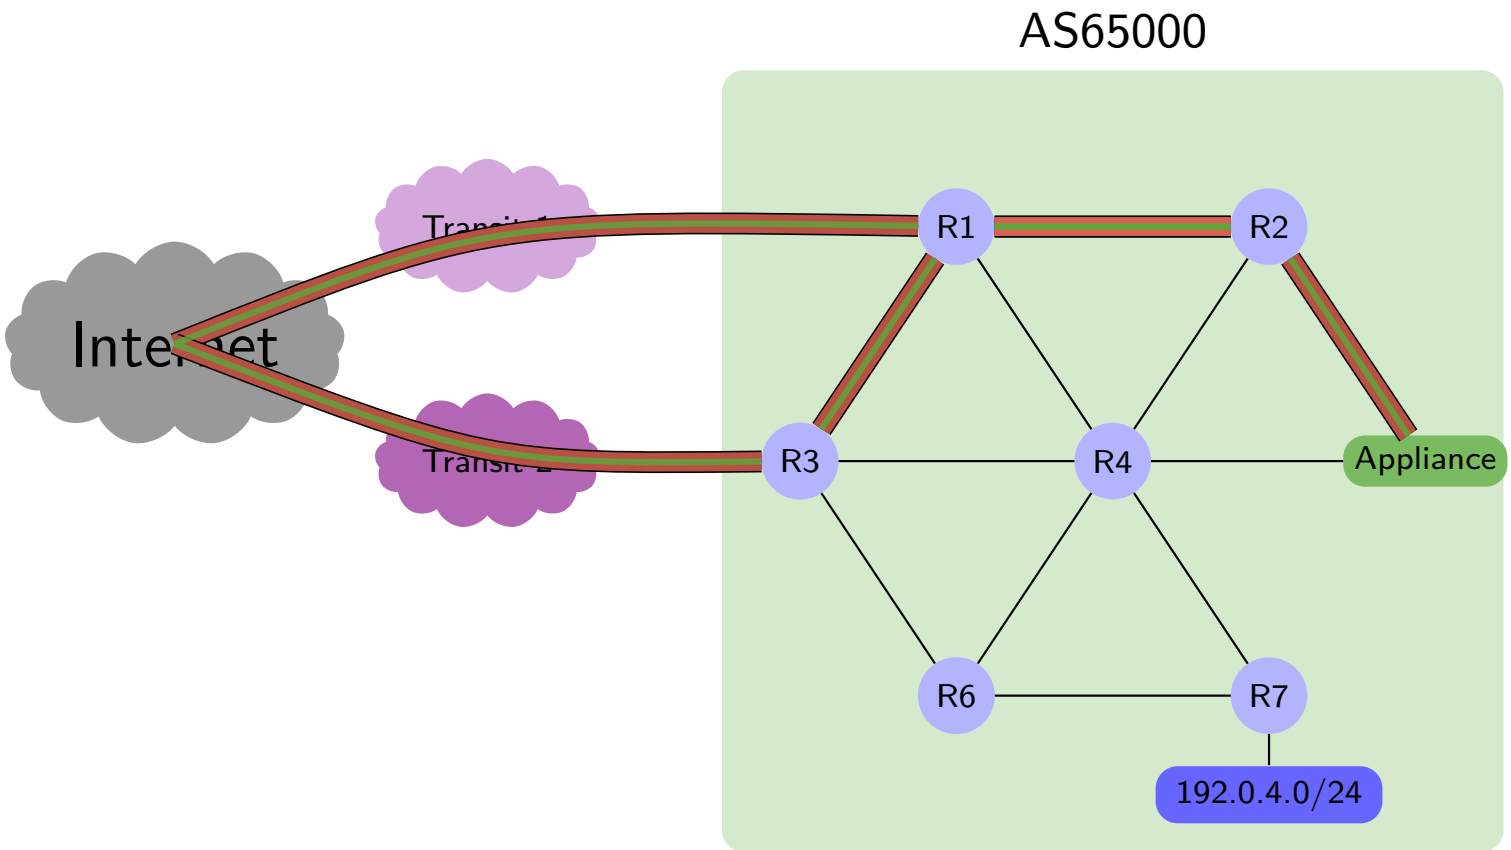

DDoS Mitigation Appliances

- Announce appliance as next-hop inside network
- Appliance cleans the traffic

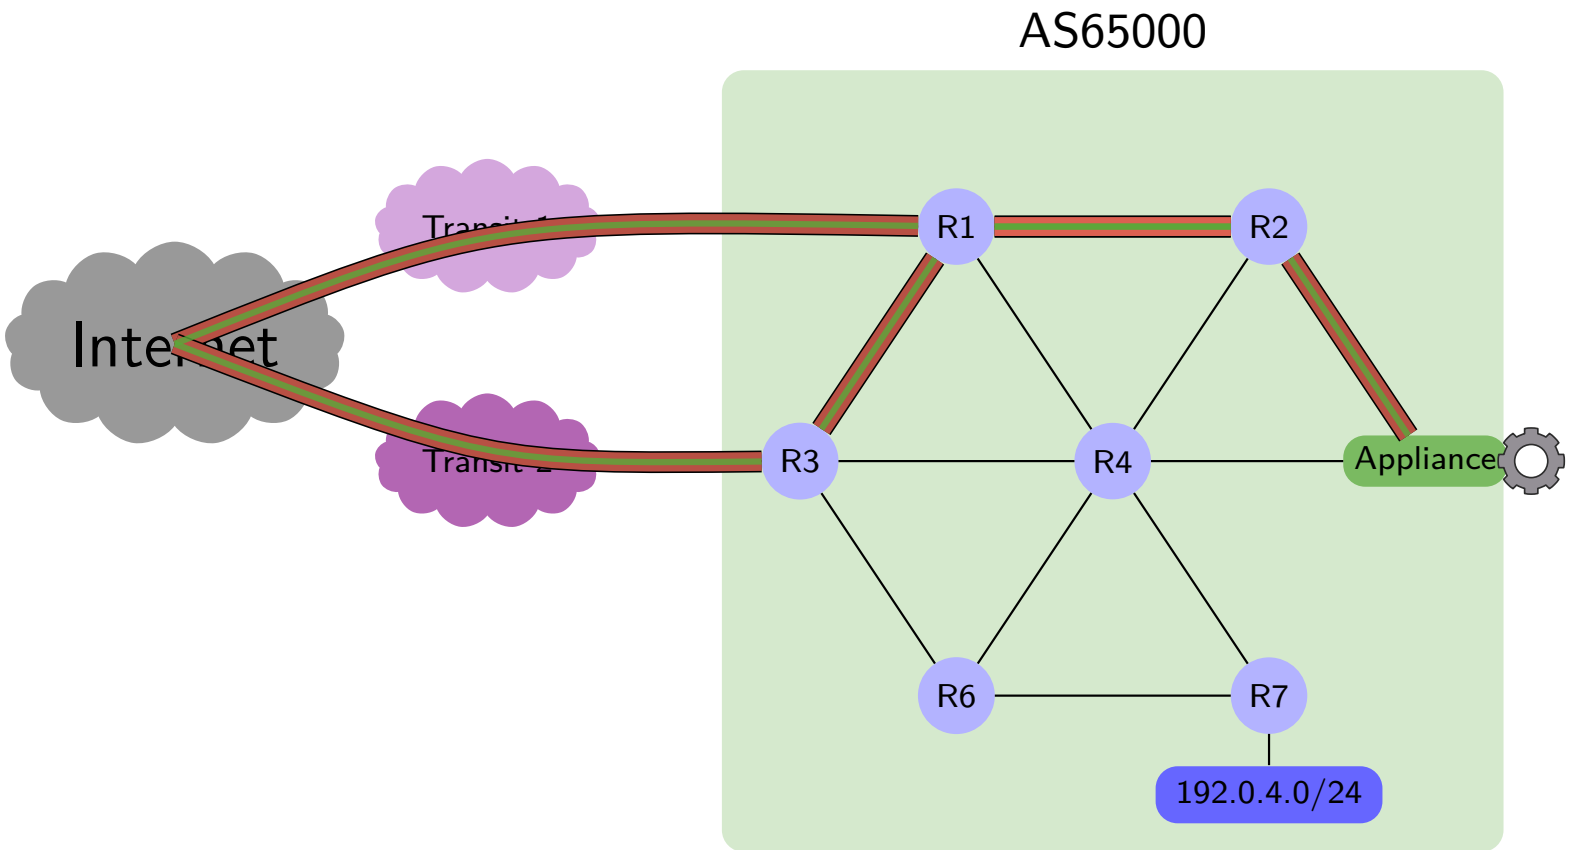

DDoS Mitigation Appliances

- Announce appliance as next-hop inside network
- Appliance cleans the traffic
- Clean traffic flows to the attacked server

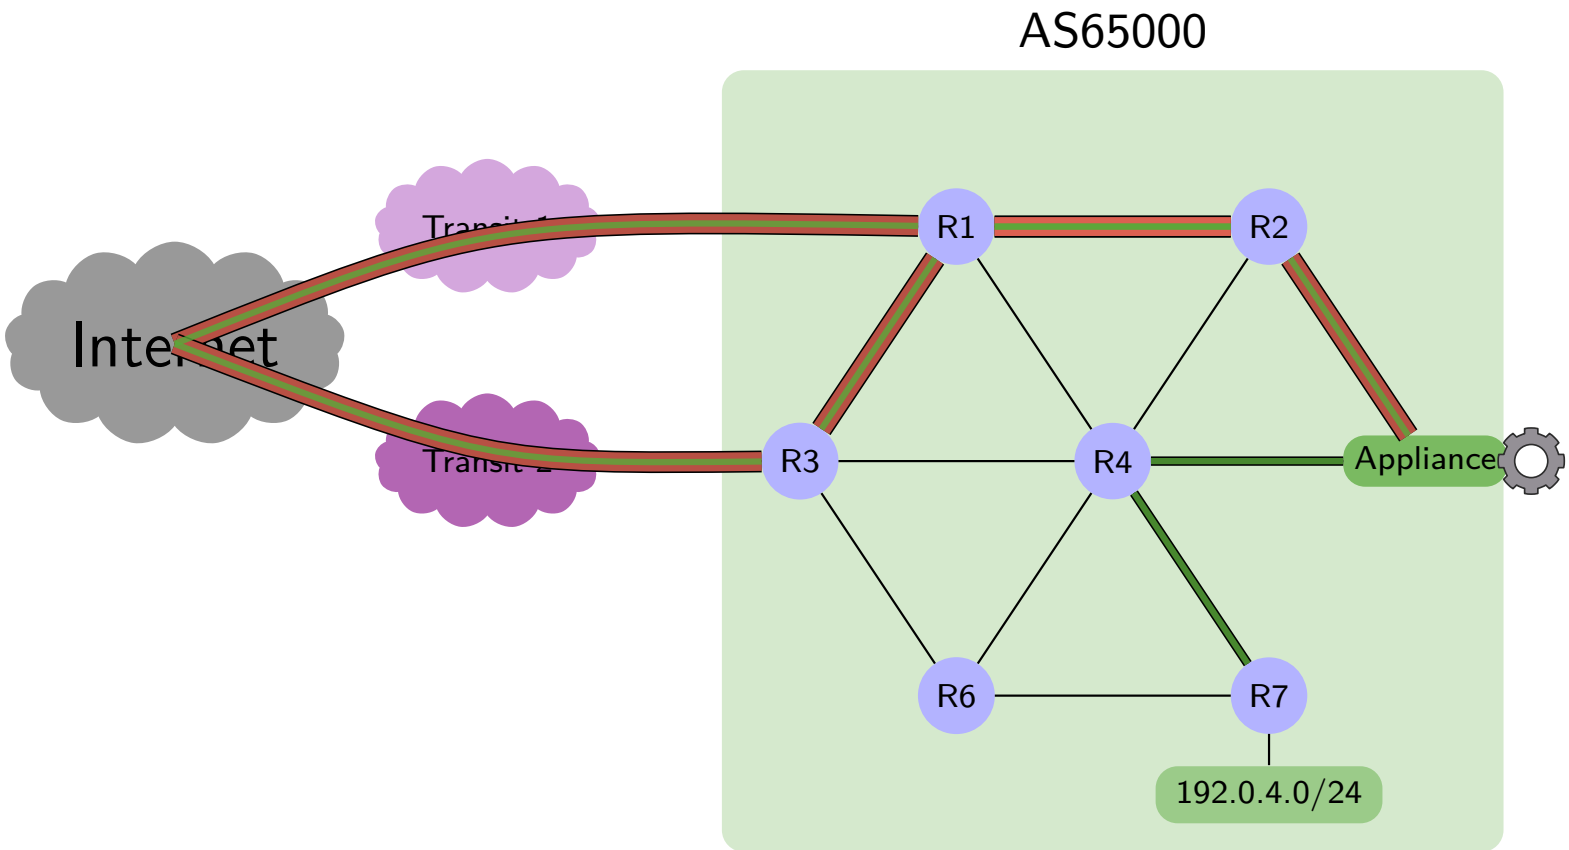

DDoS Mitigation Appliances

- Announce appliance as next-hop inside network
- Appliance cleans the traffic
- Clean traffic flows to the attacked server
- Can be inside network or done by direct Transit provider

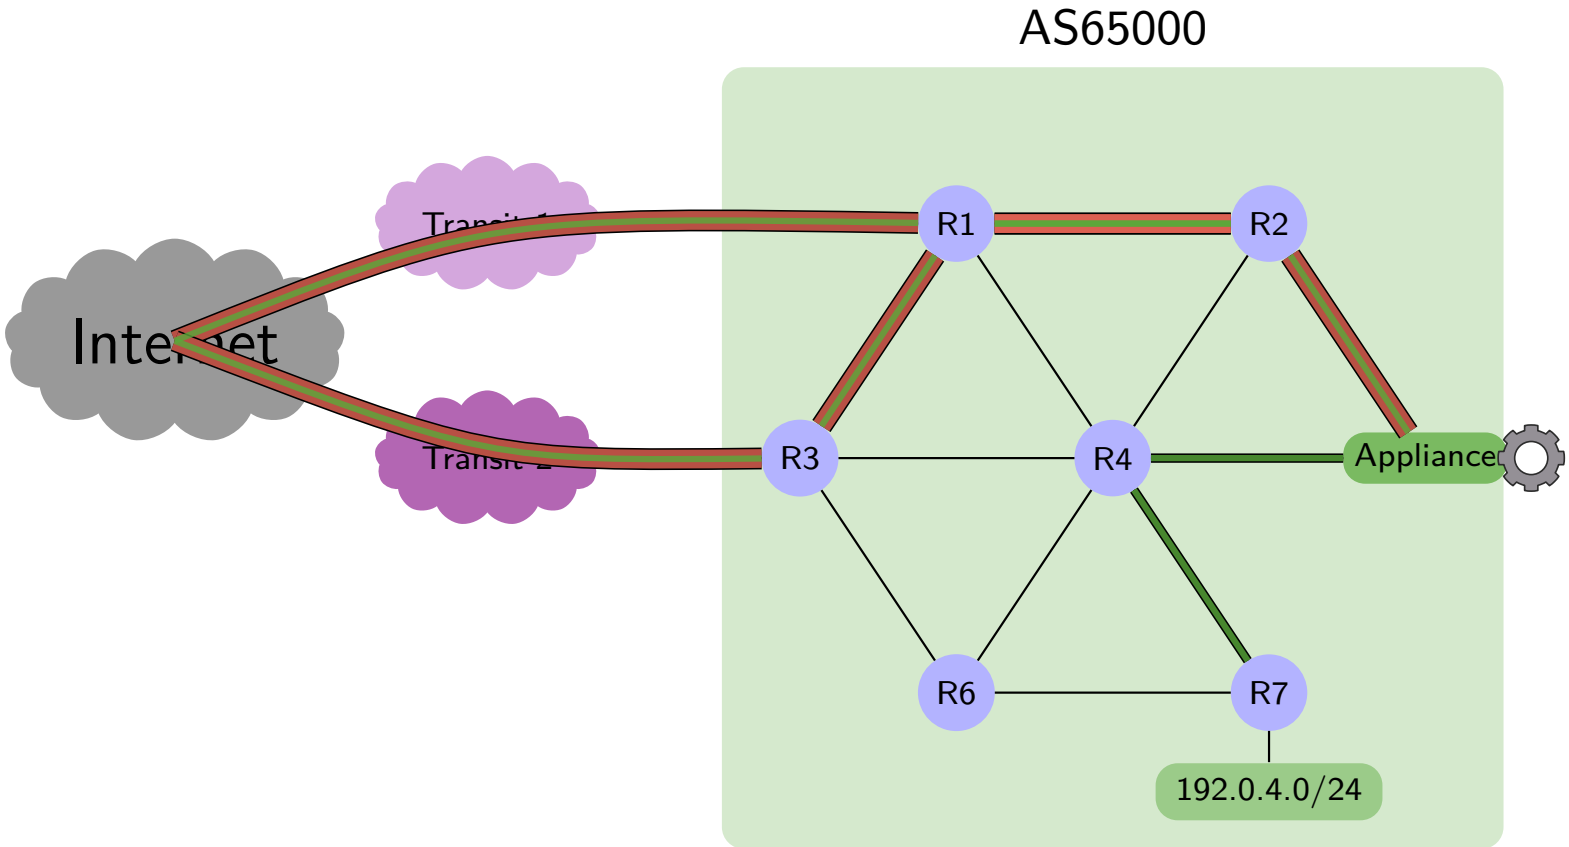

Selective Blackholing
DDoS Damage Control

Selective Blackholing: DDoS Damage Control

Why?

- RTBH makes the attacked service unavailable
- DDoS mitigation appliances have bandwidth limitations
- DDoS bandwidth $>$ targeted backbone total bandwidth

Selective Blackholing: DDoS Damage Control

Why?

- RTBH makes the attacked service unavailable
- DDoS mitigation appliances have bandwidth limitations
- DDoS bandwidth $>$ targeted backbone total bandwidth

Objectives

- Make the service available to selected hosts based on:
  - Continent
  - Country
- Ex: tax services

Selective Blackholing: DDoS Damage Control

Why?

- RTBH makes the attacked service unavailable
- DDoS mitigation appliances have bandwidth limitations
- DDoS bandwidth $>$ targeted backbone total bandwidth

Objectives

- Make the service available to selected hosts based on:
  - Continent
  - Country
- Ex: tax services

How does it work?

- RTBH + whitelisting for selected hosts

Damage Control vs Mitigation

	Mitigation	Damage control
Unsollicitated traffic	Blocked	Residual
Service availability	Full	Controlled
Bandwidth limit	Licence	None

- Mitigation fits low bandwith DDoS
- Damage control fits high bandwidth DDoS
- Mitigation and Damage control complement one another

Selective Blackholing Features

Limit geographically the scope of an announced prefix

Implemented by transit providers

- Not signature based
- No appliance
- No bandwidth limit
- Cost limited to architecture design (configuration)
- Attacked service reachable by selected hosts
- Residual unsollicitated traffic, but service still available

Customer Setup

Example of communities for customers

Null route	Community
Outside country	65000:blk-ctry
Outside continent	65000:blk-ctnt
RTBH	65000:rtbh

Transit provider setup

Transit provider defines an internal AS number for SBH

- Used for differentiation between communities added by customer and internal ones
- Dedicated AS number not announced on the Internet
- Example: 64666

Transit Provider Setup: Incoming Annonces

For each eBGP router in the backbone, define:

- Local city community
- Local country community
- Local continent community

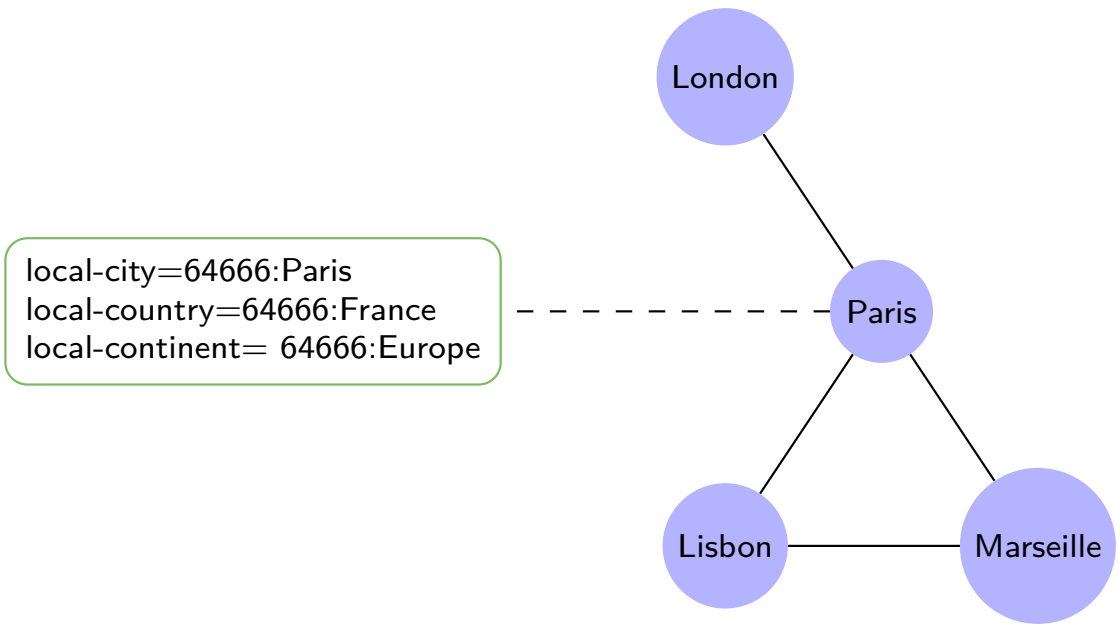

Transit Provider Setup: Incoming Annonces

For each eBGP router in the backbone, define:

- Local city community
- Local country community
- Local continent community

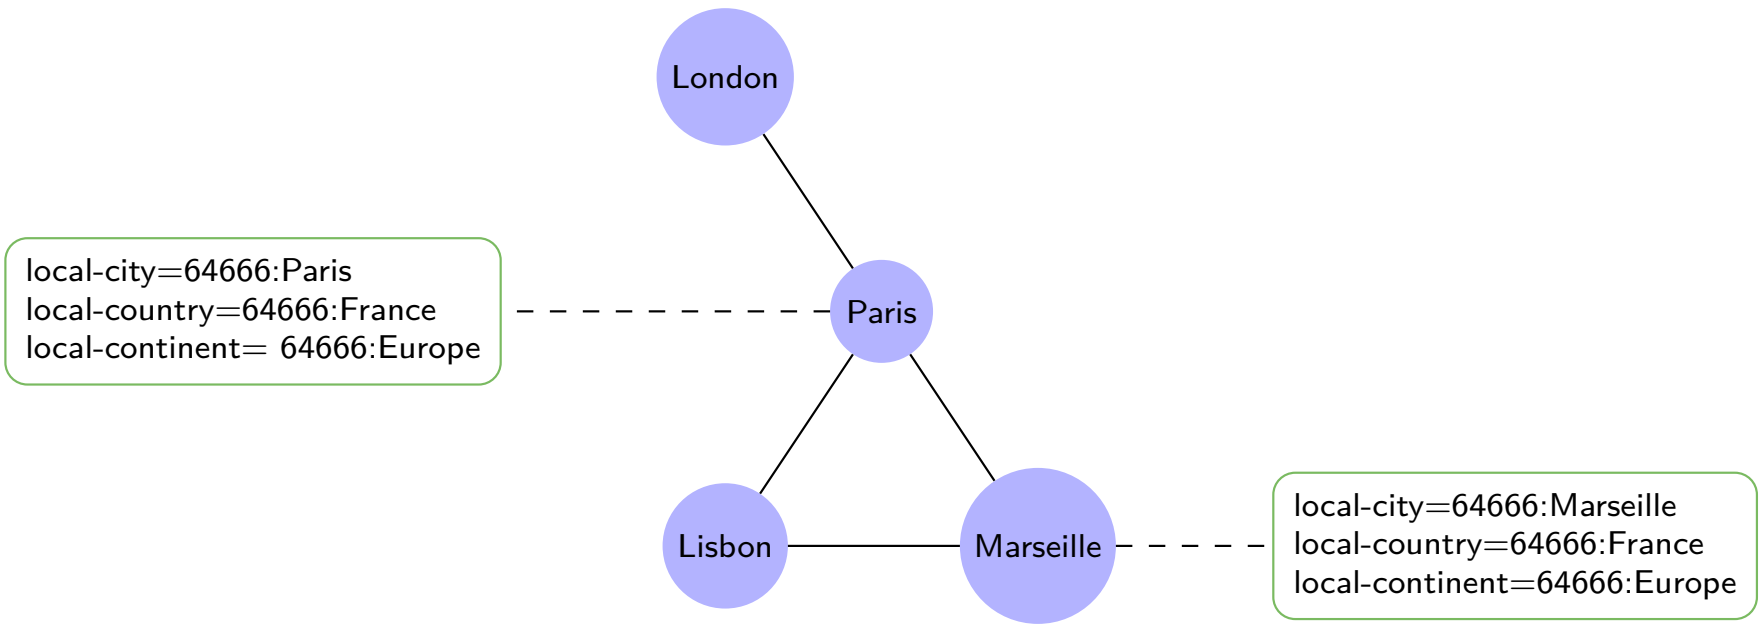

Transit Provider Setup: Incoming Annonces

For each eBGP router in the backbone, define:

- Local city community
- Local country community
- Local continent community

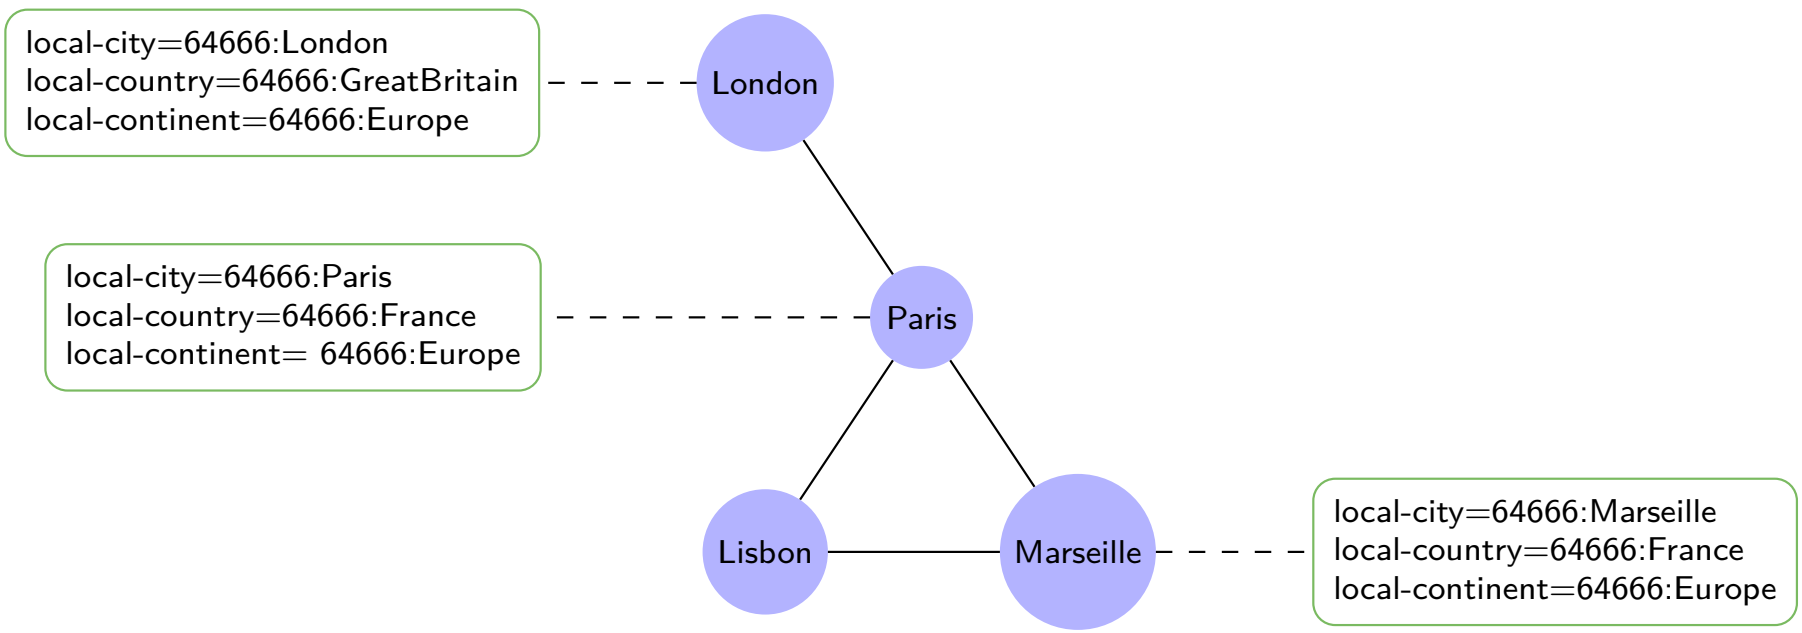

Transit Provider Setup: Incoming Annonces

For each eBGP router in the backbone, define:

- Local city community
- Local country community
- Local continent community

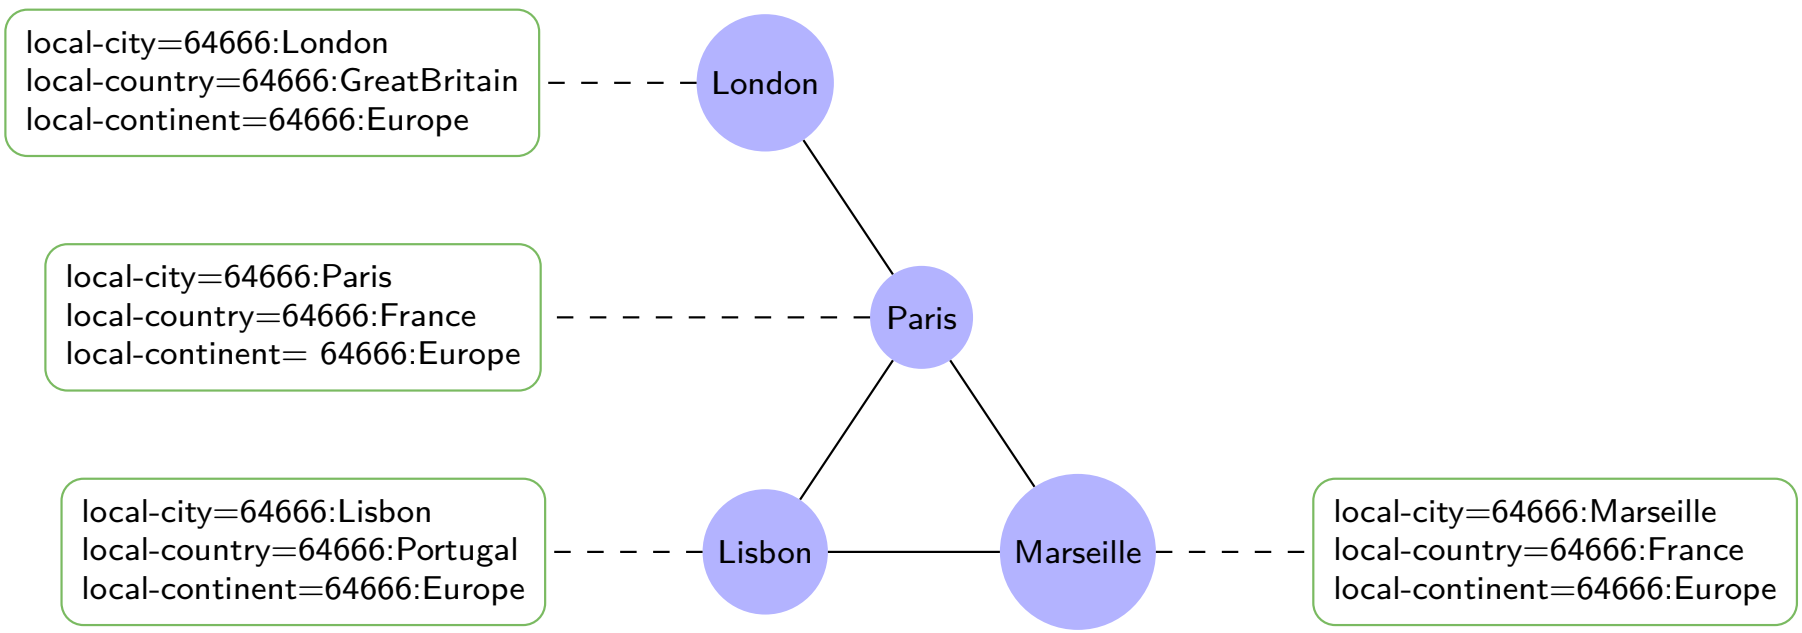

Transit Provider Setup: Incoming Annonces

For each eBGP router in the backbone, define:

- Generic route-map to insert internal communities based on generic customer communities

```
route-map RX:Customer
if customer.community == 65000:blk-ctry then
  add local-country
fi
if customer.community == 65000:blk-ctnt then
  add local-continent
fi
.
.
if customer.community == 65000:rtbh then
  set next-hop Null0
fi
```

Transit Provider Setup: Transfering Announces

Propagation mechanism:

```
route-map TO:Customer
if 65000:blk-ctnt|blk-ctry|blk-1k|blk-2.5k then
# The prefix wants to be in selective blackholing
  if prefix.community == 64666:local-city or
    prefix.community == 64666:local-country or
    prefix.community == 64666:local-continent then
    # The prefix fits the router scope
      forward announce
  else
      set next-hop Null0
  fi
fi
```

Example: Limit to local country 65000:blk-ctry

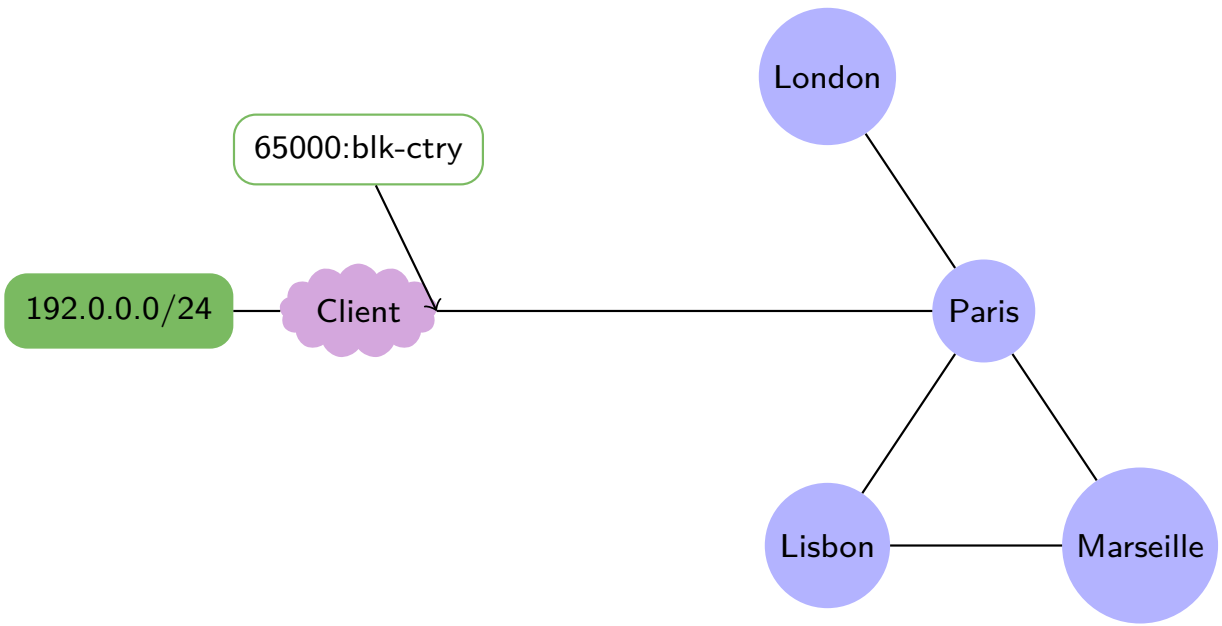

Example: Limit to local country 65000:blk-ctry

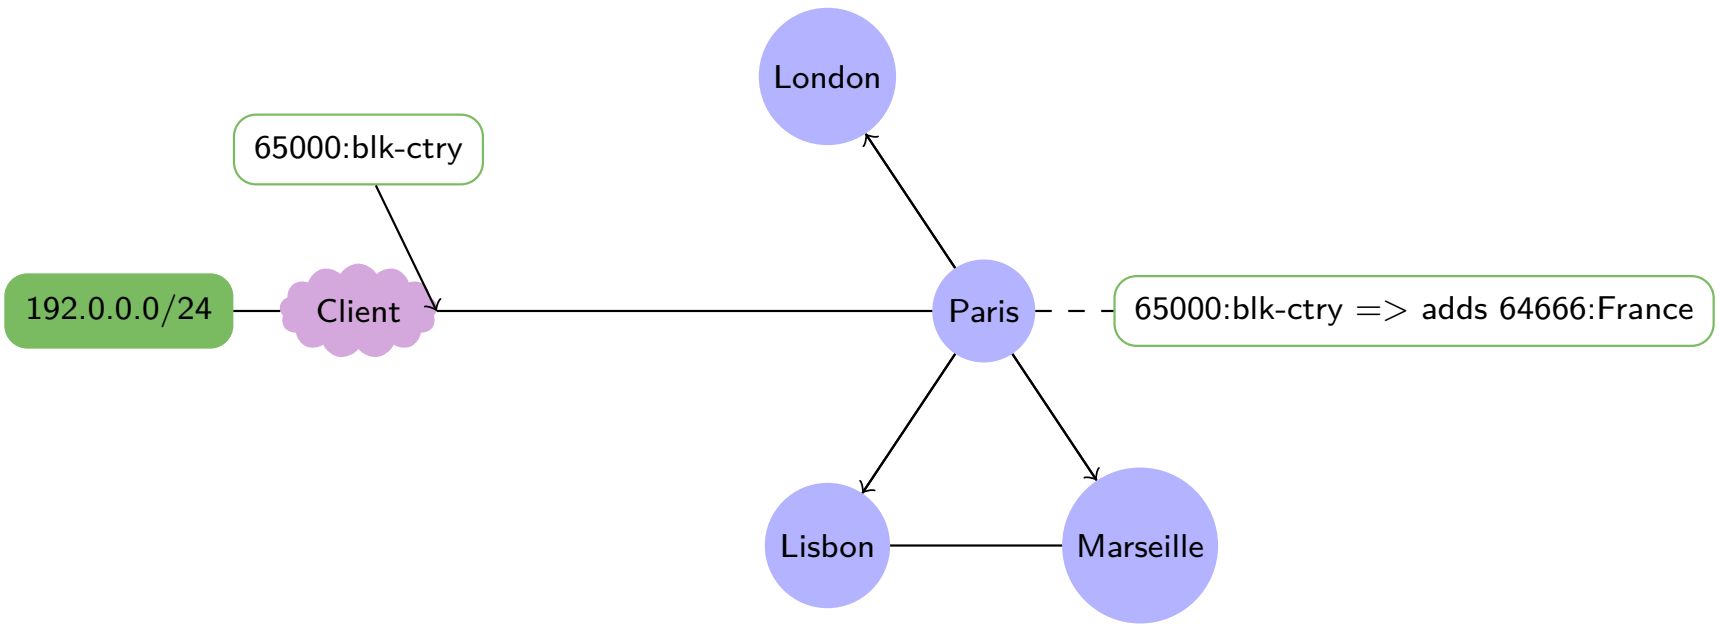

Example: Limit to local country 65000:blk-ctry

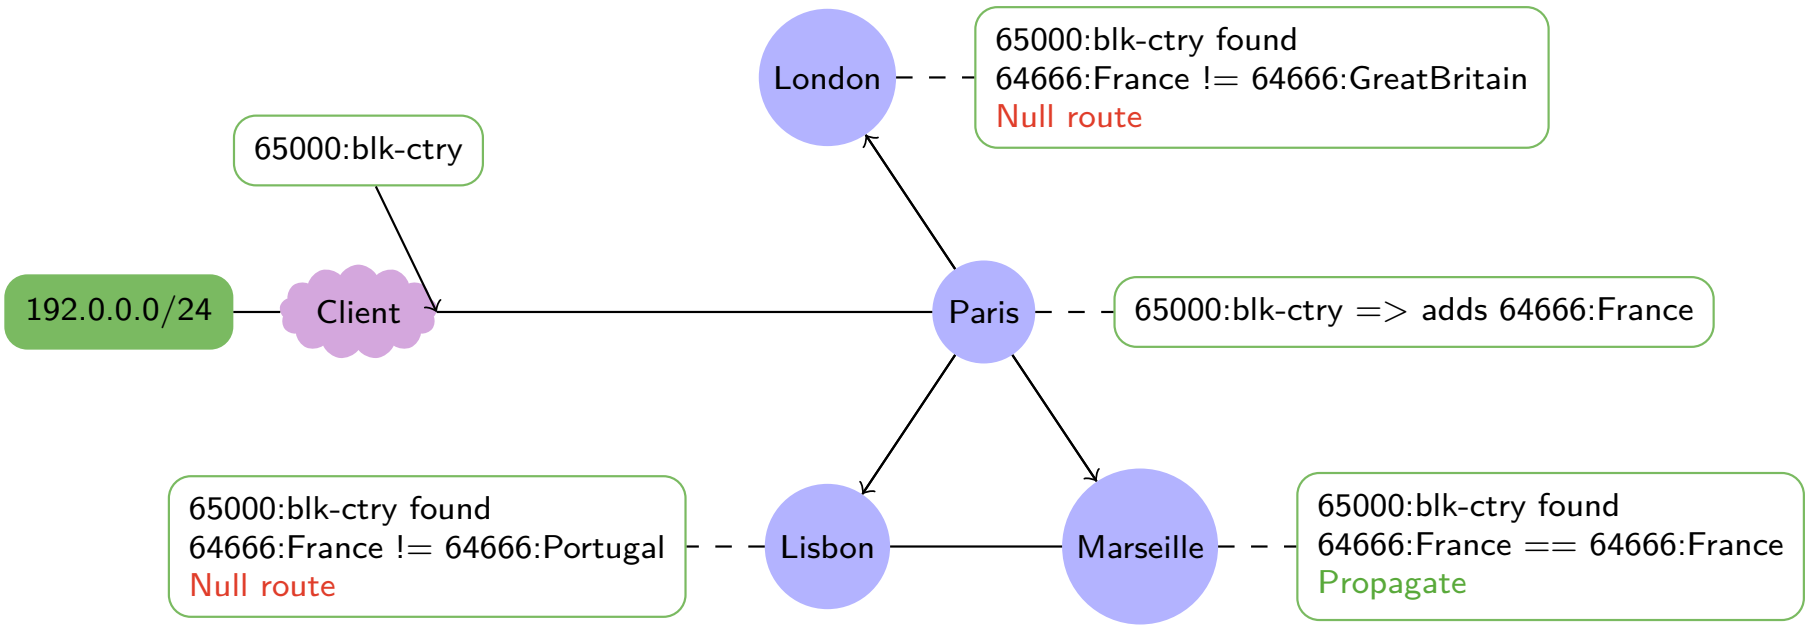

SBH Demo

Proof of concept

- 15 customer LXC routers using OpenBGP
- 16 Cisco routers (backbone)
- Implementation of SBH on all routers
- Iperf to generate DDoS
- Intermapper to graph supervision

SBH Demo

Demo

Radius Selective Blackholing

Null route	Community
Outside country	65000:blk-ctry
Outside continent	65000:blk-ctnt
Outside 1000km	65000:blk-1k
Outside 2500km	65000:blk-2.5k
RTBH	65000:rtbh

Selective Blackholing: Radius Filtering

Methodology:

- Collect latitude and longitude for each eBGP router
- Calculate distance through Harversine formula
- Generate route-map and upload to routers

Drawbacks:

- Renew Haversine calculations
- Recalculate backbone route-maps if you add or remove a router

Example of Radius Selective Blackholing

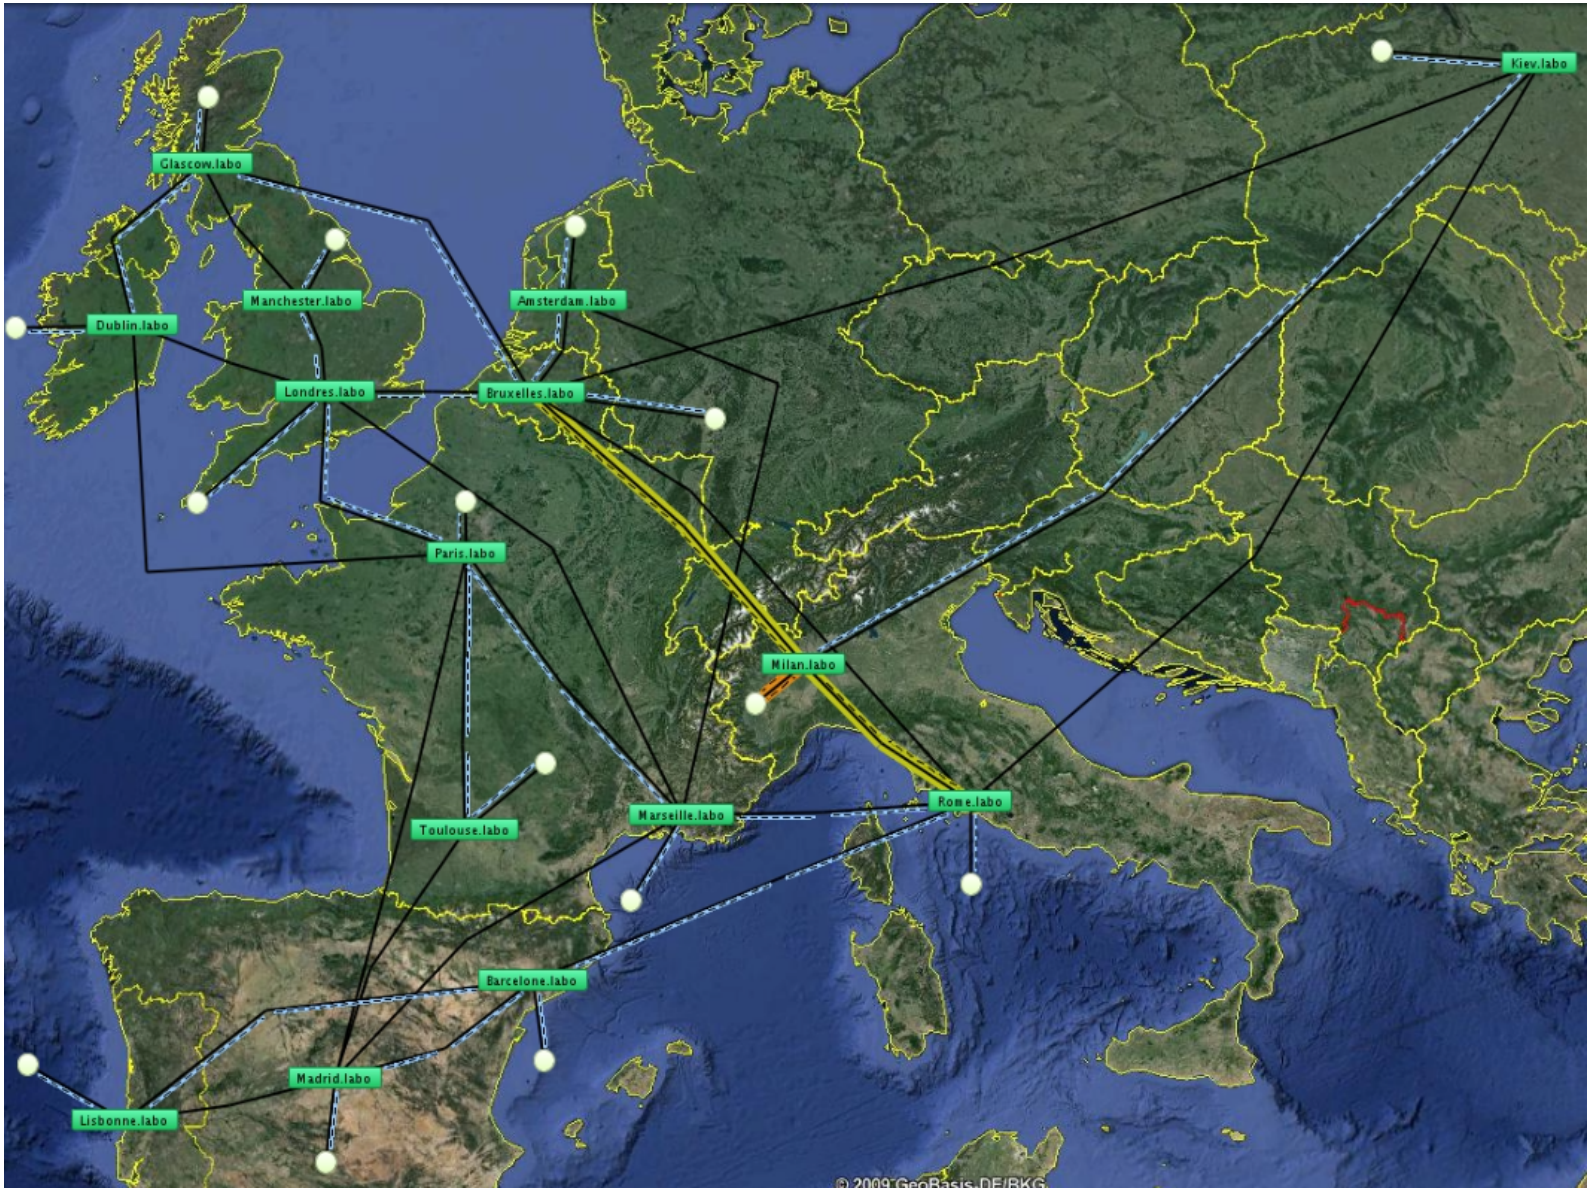

Example of Radius Selective Blackholing

Example of Radius Selective Blackholing

2500km of Kiev

- 64666:250 -> France: Paris+Marseille+Toulouse
- 64666:276 -> UK: Glascow+London+Manchester
- 64666:380 -> Italy: Milan + Rome
- 64666:20008 -> Barcelona
- 64666:56 -> Belgium: Brussels
- 64666:528 -> Netherlands: Amsterdam
- 64666:616 -> Ukraine: Kiev

Implementation

- Populate database
- Customize route-map generation script
- Generate custom values for internal communities:

  Example: blk-ctry = 65000:664, blk-cntnt = 65000:662
  Example: France = 64666:250 (ISO31661), Europe = 64666:2000

- For each eBGP router:

  Generate generic route-map RX:Customer
  Generate route-map TO:Customer

Conclusion

Conclusion

- RTBH successor
- SBH and DDoS mitigation appliances complement one another
- Efficient solution with no appliance or licence cost, only configuration
- POC showing the capabilities of the solution
- Radius implementation filtering implies too much modifications on routers
- Implemented by NLayer (GTT) and Atrato

Questions?

Author: Job Snijders <job[at]instituut[dot]net>

https://ripe68.ripe.net/archives/video/137/
http://mailman.nanog.org/pipermail/nanog/2014-February/064381.html
http://noc.as5580.net/~job/example_community_calculator.py

How the filtering process works?

Mechanisms

- Translate community sent by customer to inner community
- On each border router:

```
if receiv-community == 666 then
  next-hop = /dev/null
else if receiv-community != 660|662|663|664 then
  forward
else if receiv-community == local router continent then
  forward
else if receiv-community == local router country then
  forward
else if receiv-community == local router city then
  forward
else
  set next-hop to /dev/null
```

config example

```
ip community-list standard THIS:METRO permit 64666:20001
ip community-list standard THIS:COUNTRY permit 64666:250
ip community-list standard THIS:CONTINENT permit 64666:2000
! communities which are exposed in public documentation for customers
ip community-list standard OUTSIDE:THIS:CONTINENT:DISCARD permit 65000:660
ip community-list standard OUTSIDE:2500KM:RADIUS:DISCARD permit 65000:662
ip community-list standard OUTSIDE:1000KM:RADIUS:DISCARD permit 65000:663
ip community-list standard OUTSIDE:THIS:COUNTRY:DISCARD permit 65000:664
ip community-list standard GLOBAL:BLACKHOLE permit 65000:666
ip community-list standard SCOPED:ACTION permit 65000:660
ip community-list standard SCOPED:ACTION permit 65000:662
ip community-list standard SCOPED:ACTION permit 65000:663
ip community-list standard SCOPED:ACTION permit 65000:664
ip route 10.0.0.1 255.255.255.255 null0
route-map RX:iBGP permit 100
match community GLOBAL:BLACKHOLE
continue 1100
route-map RX:iBGP permit 100
match community OUTSIDE:THIS:COUNTRY:DISCARD OUTSIDE:THIS:CONTINENT:DISCARD
continue 1100
route-map RX:iBGP permit 110
match community OUTSIDE:1000KM:RADIUS:DISCARD OUTSIDE:2500KM:RADIUS:DISCARD
continue 1100
route-map RX:iBGP permit 1000
route-map RX:iBGP permit 1100
match community THIS:METRO THIS:COUNTRY THIS:CONTINENT
route-map RX:iBGP permit 1101
set community no-export additive
set ip next-hop 10.0.0.1
set local-preference 1000000
router bgp 65000
neighbor 192.168.1.16 route-map RX:iBGP in
```