



# Routing Security for lazy kids

RS4LK





# fermi tutti, sentiamo l'Oracolo



# ciao, sono un povero bimbo svogliato, come potrei migliorare la sicurezza nel routing BGP del mio sistema autonomo?

## ↳ PERPLEXITY

Per migliorare la sicurezza nel routing BGP del tuo sistema autonomo, ecco alcuni suggerimenti:

1. Implementa filtri di ingresso e uscita per i prefissi BGP: utilizza filtri di ingresso per limitare i prefissi BGP che accetti dai tuoi peer e filtri di uscita per limitare i prefissi BGP che annunci ai tuoi peer. In questo modo, puoi evitare di accettare o annunciare prefissi BGP non autorizzati o non validi.
2. Utilizza l'autenticazione BGP: utilizza l'autenticazione BGP per verificare l'identità dei tuoi peer BGP e garantire che le informazioni di routing che ricevi siano autentiche e non manipolate.
3. Implementa la crittografia BGP: utilizza la crittografia BGP per proteggere le informazioni di routing che scambiate con i tuoi peer BGP e garantire che non siano intercettate o manipolate da terze parti.
4. Monitora costantemente il tuo sistema autonomo: monitora costantemente il tuo sistema autonomo per rilevare eventuali anomalie o attacchi di routing e agisci prontamente per mitigare i problemi.
5. Partecipa a comunità di sicurezza BGP: partecipa a comunità di sicurezza BGP per rimanere aggiornato sulle ultime minacce e vulnerabilità di routing e condividere informazioni e best practice con altri operatori di rete.

# ho capito, grazie, ma io sono proprio pigro e quello che suggerisci è molto faticoso

## ↳ PERPLEXITY

Capisco che possa sembrare un compito impegnativo, ma la sicurezza del routing BGP è un aspetto critico per la sicurezza della tua rete e dei tuoi dati. Tuttavia, ci sono alcune cose che puoi fare per migliorare la sicurezza del tuo sistema autonomo con un minimo sforzo:

1. Accordo commerciale: il primo passo per migliorare la sicurezza del routing BGP è un accordo commerciale tra organizzazioni che intendono scambiarsi traffico dati su Internet. Questo accordo può includere clausole di sicurezza che limitano l'accesso ai propri prefissi BGP solo ai propri peer autorizzati.
2. Utilizza filtri di ingresso e uscita di base: anche l'implementazione di filtri di ingresso e uscita di base può migliorare significativamente la sicurezza del tuo sistema autonomo. Ad esempio, puoi utilizzare filtri di ingresso per limitare i prefissi BGP che accetti dai tuoi peer e filtri di uscita per limitare i prefissi BGP che annunci ai tuoi peer.
3. Utilizza l'autenticazione BGP: l'autenticazione BGP è un'opzione di sicurezza relativamente semplice da implementare che può aiutare a proteggere il tuo sistema autonomo da attacchi di routing. Puoi utilizzare l'autenticazione BGP per verificare l'identità dei tuoi peer BGP e garantire che le informazioni di routing che ricevi siano autentiche e non manipolate.
4. Monitora costantemente il tuo sistema autonomo: anche il monitoraggio costante del tuo sistema autonomo può aiutare a rilevare eventuali anomalie o attacchi di routing. Puoi utilizzare strumenti di monitoraggio di base per controllare il traffico di routing in ingresso e in uscita dal tuo sistema autonomo.

Anche se questi suggerimenti richiedono meno sforzo rispetto alle soluzioni più avanzate, è importante ricordare che la sicurezza del routing BGP è un aspetto critico della sicurezza della tua rete e dei tuoi dati.



se lo dice lui..



Fortunatamente abbiamo il conforto anche di altre fonti

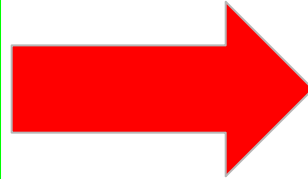
MANRS

Mutually Agreed Norms  
for Routing Security  
by ISOC

BCP38, BCP84,  
BCP194, BCP211  
by IETF

Quindi, come lazy kid, qual è la mia necessità?

Io ti faccio  
vedere la mia  
configurazione



Tu mi dici se  
è conforme a  
MANRS

# Occorre allora uno strumento che:

Riceva la configurazione di un border router (BGP speaking)

Analizzi la configurazione così da comprendere: vendor, transit, peer, customer

Confronti i dati pubblicati sugli IRR con quanto presente nella configurazione

Restituisca segnali di OK, KO, WARNING, INFO

**Costruisca automaticamente un ambiente emulato**



Esegua molteplici test per provare la conformità alle azioni MANRS (anti-spoofing e BGP filtering)



# Verificabilità delle configurazioni, lo stato dell'arte

Batfish: "finds errors and guarantees the correctness of planned or current network configurations. It enables safe and rapid network evolution, without the fear of outages or security breaches."

Bagpipe: "enables an ISP to express its BGP policy in a domain-specific specification language and verify that its router configurations implement this policy."

BGPVerif: "a formal analysis toolkit that enables networking researchers to study and analyze large BGP systems in a sound and automatic fashion."

# Cosa non va generalmente in quegli strumenti?

## Incompletezza

*Perlopiù ricerche abbandonate allo stato embrionale che mancano di alcuni requisiti: non contemplan specifiche caratteristiche del BGP*

## Incomprensione

*I vendor sono molteplici e ciascuno impiega una sintassi (o più di una) che evolve nel tempo*

# Discussione su Batfish

# general <https://www.batfish.org/> 1,031

Tip: Try `🔍` `F` to search this channel

**Harish** 10:20 PM  
Hi Team, Does Batfish enterprise version supports additional protocols like ISIS/MPLS, etc.. ?  
1 reply 8 days ago

Thursday, April 6th

**Matt Slack** 12:55 AM  
joined #general along with Josef.

Friday, April 7th

**Josef** 8:56 AM  
Hello, I'm new to batfish and just added some Cisco IOS Switch configurations.

```
parse_status = bf.q.fileParseStatus().answer().frame()
# View the parse status results
parse_status
```

File_Name	Status	File_Format	Nodes
0 configs/ATBRACSC001.cfg	PARTIALLY_UNRECOGNIZED		
CISCO_IOS	['atbracsc001']		
1 configs/ATKHACSC001.cfg	PARTIALLY_UNRECOGNIZED		
CISCO_IOS	['atkhacsc001']		
2 configs/ATKPACSC001.cfg	PARTIALLY_UNRECOGNIZED		
CISCO_IOS	['atkpacsc001']		

**Thread # general**

**hemanth gade** 1 month ago  
Hey All, wondering if batfish supports ipv6 ?  
1 reply

**Dan Halperin** 1 month ago  
nope, not today. no work planned in this area

Reply...

Also send to # general

+ 📎 🎤 😊 @ Aa ➤



# Una riflessione sul metodo

Massima astrazione: prendi una routing policy (in RPSLng) e verifica se è conforme alle azioni MANRS

1

Astrazione: prendi una configurazione, leggi le parole, studia la sintassi, comprendi il senso e verifica se è conforme alle azioni MANRS

2

Emulazione: prendi una configurazione, applicala a un router e verifica se è conforme alle azioni MANRS

3

# 1 Massima astrazione

```
aut-num: AS59715
as-name: SBTAP-AS
org: ORG-CdSB1-RIPE
mp-import: afi ipv4.unicast from AS12874 accept any
mp-import: afi ipv6.unicast from AS12874 accept any
mp-import: afi ipv4.unicast from AS6762 accept any
mp-import: afi ipv6.unicast from AS6762 accept any
mp-import: afi ipv4.unicast from AS112 accept AS112
mp-import: afi ipv6.unicast from AS112 accept AS112
mp-export: afi ipv4.unicast to AS12874 announce AS-SBTAP
mp-export: afi ipv6.unicast to AS12874 announce AS-SBTAP
mp-export: afi ipv4.unicast to AS6762 announce AS-SBTAP
mp-export: afi ipv6.unicast to AS6762 announce AS-SBTAP
mp-export: afi ipv4.unicast to AS112 announce any
mp-export: afi ipv6.unicast to AS112 announce any
admin-c: SBT21-RIPE
tech-c: SBT20-RIPE
tech-c: AP7729-RIPE
status: ASSIGNED
mnt-by: RIPE-NCC-END-MNT
mnt-by: SBTAP-MNT
created: 2012-10-04T12:53:46Z
last-modified: 2022-10-24T14:18:26Z
source: RIPE
```



sarà vero?

# 1 Massima astrazione

```
afi ipv4.unicast from AS12874 accept any
afi ipv6.unicast from AS12874 accept any
afi ipv4.unicast from AS6762 accept any
afi ipv6.unicast from AS6762 accept any
afi ipv4.unicast from AS112 accept AS112
afi ipv6.unicast from AS112 accept AS112
afi ipv4.unicast to AS12874 announce AS-SBTAP
afi ipv6.unicast to AS12874 announce AS-SBTAP
afi ipv4.unicast to AS6762 announce AS-SBTAP
afi ipv6.unicast to AS6762 announce AS-SBTAP
afi ipv4.unicast to AS112 announce any
afi ipv6.unicast to AS112 announce any
```

è vero! ma...

V*	185.5.200.0/22	194.85.40.15	0	0	3267	1299	<u>12874</u>	<u>59715</u>	i
V*		193.0.0.56		0	3333	<u>6762</u>	<u>59715</u>	i	
V*		12.0.1.63		0	7018	<u>6762</u>	<u>59715</u>	i	
V*		212.66.96.126		0	20912	<u>12874</u>	<u>59715</u>	i	

dov'è l'AS112?

# 2 Astrazione

Batfish capisce mpBGP?

```
set protocols bgp group NLNOG_V6 type external
set protocols bgp group NLNOG_V6 multihop ttl 30
set protocols bgp group NLNOG_V6 local-address 2a02:cdc5:9715:0:185:5:200:255
set protocols bgp group NLNOG_V6 import DENY_ALL_PREFIX
set protocols bgp group NLNOG_V6 family inet6 unicast
set protocols bgp group NLNOG_V6 export ALLOW_BGP_PREFIX
set protocols bgp group NLNOG_V6 peer-as 199036
set protocols bgp group NLNOG_V6 neighbor 2001:7b8:62b:1:0:d4ff:fe72:7848
```

# 2 Astrazione

```
INFO - Parsing peerings IPv4 addresses.  
INFO - status: ASSIGNED  
INFO - .... no task information  
INFO - status: TERMINATEDNORMALLY  
INFO - .... 2023-04-15 21:20:35.375000+02:00 Begin job.  
WARNING - Peering with AS65332 is in reserved range 64000-131071.  
WARNING - Peering with AS59715 is a iBGP peering.  
WARNING - Peering with AS65332 is in reserved range 64000-131071.  
WARNING - Peering with AS59715 is a iBGP peering.  
WARNING - Peering with AS59715 is a iBGP peering.  
WARNING - Peering with AS64496 is in reserved range 64000-131071.  
WARNING - Peering with AS59715 is a iBGP peering.  
WARNING - Peering with AS64496 is in reserved range 64000-131071.  
INFO - status: ASSIGNED  
INFO - .... no task information  
INFO - status: TERMINATEDNORMALLY  
INFO - .... 2023-04-15 21:20:35.520000+02:00 Begin job.  
INFO - Configuration format is 'FLAT_JUNIPER', loading vendor configuration...  
INFO - Inferring IPv6 interface addresses.  
INFO - Inferring IPv6 routes.  
INFO - Inferring IPv6 BGP peerings.  
WARNING - Peering with AS65332 is in reserved range 64000-131071.  
WARNING - Peering with AS64496 is in reserved range 64000-131071.
```

capisce sessioni BGP

capisce Juniper



# 2 Astrazione

```
- Interface 'xe-0/1/0' remapped into 'ge-0/0/1.0'.
- Interface 'xe-0/1/0.0' remapped into 'ge-0/0/1.0'.
- Interface 'xe-0/1/1' remapped into 'ge-0/0/2.0'.
- Interface 'xe-0/1/1.100' remapped into 'ge-0/0/2.100'.
- Interface 'xe-0/1/1.169' remapped into 'ge-0/0/2.169'.
- Interface 'xe-0/1/2' remapped into 'ge-0/0/3.0'.
- Interface 'xe-0/1/2.0' remapped into 'ge-0/0/3.0'.
- Interface 'xe-0/1/3' remapped into 'ge-0/0/4.0'.
- Resulting interfaces are:
- * lo0 with addresses: []
- * lo0.0 with addresses: []
- * ge-0/0/1.0 with addresses: []
- * ge-0/0/1.0 with addresses: [IPv4Interface('185.5.200.255/23'), IPv6Interface('2a02:cdc6:3074::/64'), IPv6Interface('2a02:cdc6:3074::/64')]
- * ge-0/0/2.0 with addresses: []
- * ge-0/0/2.100 with addresses: [IPv4Interface('185.5.202.2/24')]
- * ge-0/0/2.169 with addresses: [IPv4Interface('169.254.63.74/24')]
- * ge-0/0/3.0 with addresses: []
- * ge-0/0/3.0 with addresses: [IPv4Interface('89.221.32.157/31')]
- * ge-0/0/4.0 with addresses: []
```

capisce le interfacce

e le mappa con IP

# 2 Astrazione

## Initializing the Network and Snapshot

SNAPSHOT\_PATH below can be updated to point to a custom snapshot. More example networks are available in the [networks](#) folder of the Batfish.

```
In [2]: # Initialize a network and snapshot
NETWORK_NAME = "example_network"
SNAPSHOT_NAME = "example_snapshot"

SNAPSHOT_PATH = "networks/example-bgp"

bf.set_network(NETWORK_NAME)
bf.init_snapshot(SNAPSHOT_PATH, name=SNAPSHOT_NAME, overwrite=True)
```

Your snapshot was successfully initialized but Batfish failed to fully recognize some lines in one or more input files. Some unrecognized configuration lines are not uncommon for new networks, and it is often fine to proceed with further analysis. You can help the Batfish developers improve support for your network by running:

```
bf.upload_diagnostics(dry_run=False, contact_info='<optional email address>')
```

to share private, anonymized information. For more information, see the documentation with:

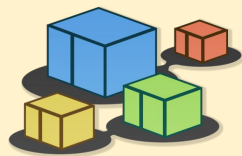
```
help(bf.upload_diagnostics)
```

```
Out[2]: 'example_snapshot'
```

capisce molto, ma non tutto IPv6!

# 3 Emulazione

Studi scientifici hanno dimostrato che



## Kathará

allevia il mal di testa!

Kathará is an open source container-based network emulation system for showing interactive demos/lessons, testing production networks in a sandbox environment, or developing new network protocols.

# 3 Emulazione



Chiediamo a kathara di prendere tutto ciò che viene digerito da Batfish e di creare uno scenario, cioè una piccola Internet con tutti gli elementi essenziali della configurazione:

- **peer** ( $\forall p \in AS1\_SET : \exists rp,2 \in T2; \forall q \in AS2\_SET : \exists rq,1 \in T1$ )
- **transit** ( $\forall p \in FIRT : \exists rp,1 \in T1 \wedge \exists rp,2 \in T2$ )
- **customer** ( $\forall p \in AS2\_SET : \exists rp,2 \in T2 \wedge \exists rp,1 \in T1$ )

**peer** ( $\forall p \in AS1\_SET : \exists rp,2 \in T2 ; \forall q \in AS2\_SET : \exists rq,1 \in T1$ )  
**transit** ( $\forall p \in FIRT : \exists rp,1 \in T1 \wedge \exists rp,2 \in T2$ )  
**customer** ( $\forall p \in AS2\_SET : \exists rp,2 \in T2 \wedge \exists rp,1 \in T1$ )

Non ci son problemi: tu mi dici quello che devo fare e io lo faccio!

Kathará

# 3 Emulazione

```
- Creating Kathara network scenario from configuration...
- Created device under test with name 'as_59715'.
- Searching peering device with name 'as_112'.
- Device 'as_112' not found, creating...
- Device 'as_112' created and directly connect
- Searching peering device with name 'as_
- Device 'as_112' already created.
- Searching peering device with name 'as_676
- Device 'as_6762' not found, creating...
- Last connected 'as_59715' iface idx=2, next peering if
- Created dummy collision domain for 'as_59715' on idx
- Device 'as_6762' created and directly connected with device 'as_59715' on
- Peering with AS6762 is provider, add client.
- Added client 'as_6762_client' on collision domain 'AAAD'.
- Searching peering device with name 'as_6762'.
- Device 'as_6762' already created.
- Added client 'as_59715_client' on iface idx=5 on collision domain 'AAAE'.
```

creazione dello scenario

# 3 Emulazione

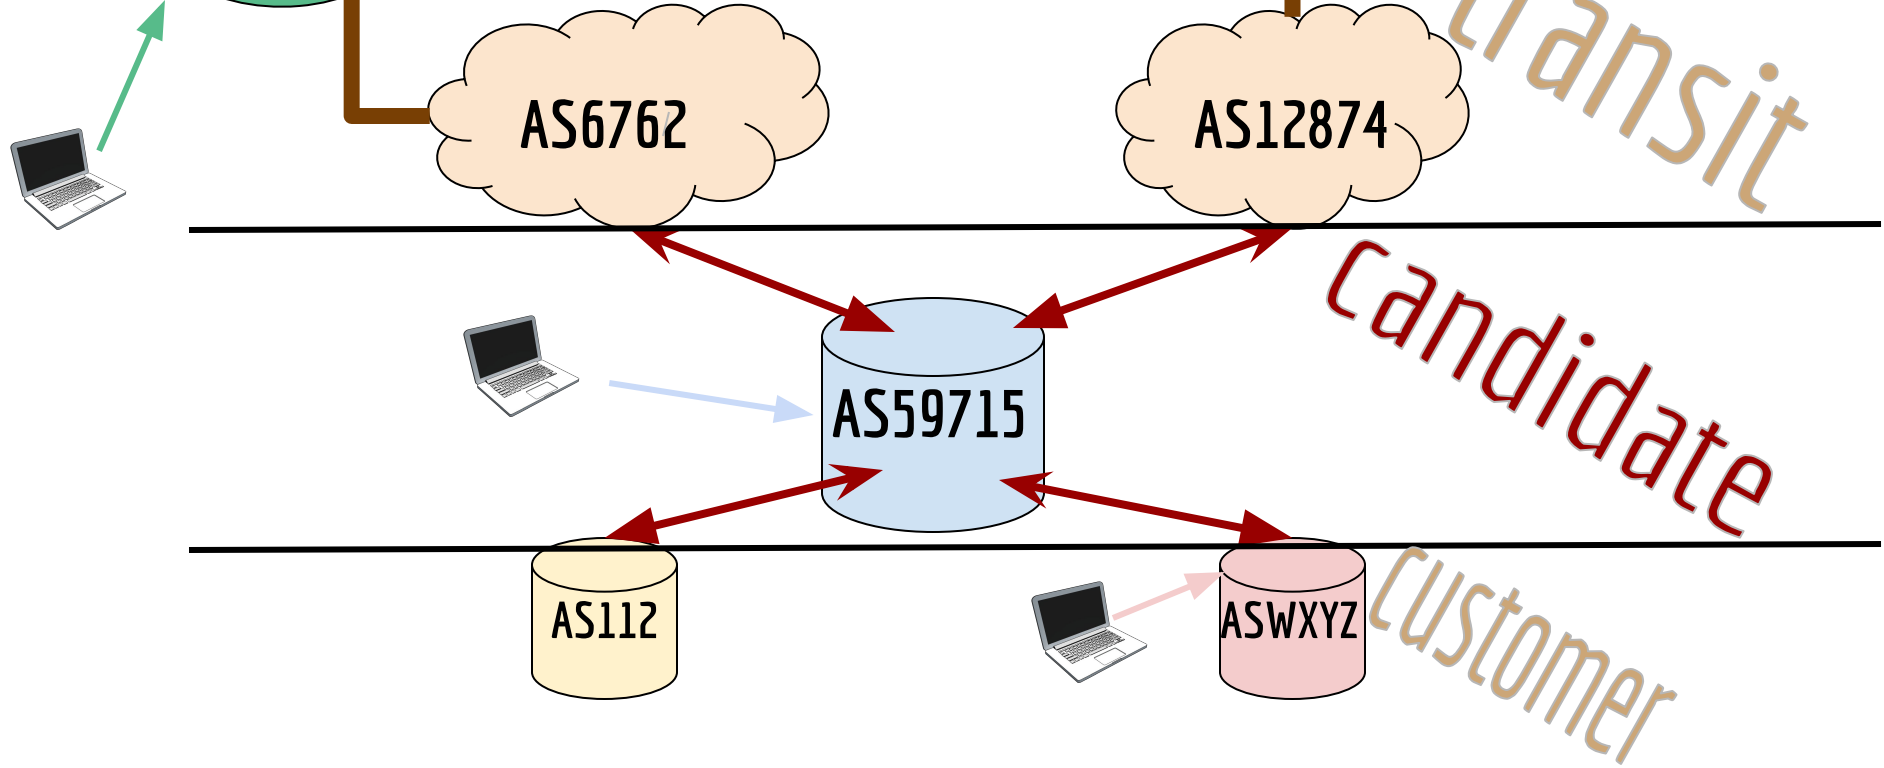


## configurazione delle sessioni BGP

```
- Configuring BGP in peering device 'as_112' ...
- Adding neighbor to 'as_112' on interface IP=185.5.201.241 toward
- Adding neighbor to 'as_112' on interface IP=2a02:cdc5:9715:0:185
- Configuring BGP in peering device 'as_6762' ...
- Adding neighbor to 'as_6762' on interface IP=89.221.32.156 toward
- Adding neighbor to 'as_6762' on interface IP=2001:41a8:60:2::79
- AS6762 is a provider, announcing 0.0.0.0/0 on IPv4...
- AS6762 is a provider, announcing ::/0 on IPv6...
```

# Internet

## 3 Emulazione

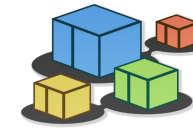




```

set lab@vr-vmx> show configuration | display set
set version 18.2R1.9
set system login user vrnnetlab uid 2000
set system login user vrnnetlab class super-user
set system login user vrnnetlab authentication encrypted-password "$6$EW.Wa.FU$02rxc6GkbIU/7XtLk86fzom6wFmQP1u87YbJQTtHLFGNymEXL/gLG01uyIv.Eeq
25xi0b7pC.DZUcPHQwpl"
set system root-authentication encrypted-password "$6$ZV%kArQl$VHz3zZJ0EJGxUfXbhU4500c6uk.Ck.J.,.uLdbq97PSCkRE0xUyVhPppDlleKd00BfDaYRFmK/NkZsgr
Hhadi1"
set system host-name vr-vmx
set system management-instance
set system services ssh
set system services extension-service request-response grpc clear-text port 57400
set system services extension-service request-response grpc max-connections 4
set system services netconf ssh
set system services netconf rfc-compliant
set system syslog user * any emergency
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set chassis fpc 0 pic 0 number-of-ports 96
set interfaces ge-0/0/1 description DOWNLINK-HUAWEI-10G-WAN
set interfaces ge-0/0/1 unit 0 family inet address 185.5.203.254/24
set interfaces ge-0/0/1 unit 0 family inet address 185.5.201.201/23
set interfaces ge-0/0/1 unit 0 family inet address 185.5.200.255/23
set interfaces ge-0/0/1 unit 0 family inet6 address 2a02:cdc5:9715:0:185:5:200:255/64
set interfaces ge-0/0/1 unit 0 family inet6 address 2a02:cdc5:9715:0:185:5:203:254/64
set interfaces ge-0/0/1 unit 0 family inet6 address 2a02:cdc6:3074::/64
deactivate interfaces ge-0/0/1 unit 0 family inet6 address 2a02:cdc6:3074::/64
set interfaces ge-0/0/1 unit 0 family inet6 address 2a02:cdc6:3074:0:185:5:200:255/64
set interfaces ge-0/0/1 unit 0 family inet6 address 2a02:cdc6:3074:0:185:5:203:254/64
set interfaces ge-0/0/2 description DOWNLINK-HUAWEI-10G-RSPTP
set interfaces ge-0/0/2 vlan-tagging
set interfaces ge-0/0/2 unit 100 vlan-id 100
set interfaces ge-0/0/2 unit 100 family inet address 185.5.202.2/24 preferred
set interfaces ge-0/0/2 unit 100 family inet address 185.5.202.2/24 vrrp-group 10 virtual-address 185.5.202.254
set interfaces ge-0/0/2 unit 100 family inet address 185.5.202.2/24 vrrp-group 10 priority 150
set interfaces ge-0/0/2 unit 100 family inet address 185.5.202.2/24 vrrp-group 10 advertise-interval 1
set interfaces ge-0/0/2 unit 100 family inet address 185.5.202.2/24 vrrp-group 10 accept-data
set interfaces ge-0/0/2 unit 100 family inet address 185.5.202.2/24 vrrp-group 10 authentication-type simple
set interfaces ge-0/0/2 unit 100 family inet address 185.5.202.2/24 vrrp-group 10 authentication-key "$9$0cnlnCuRhr8xdp0gkv87NUJHmT3n/HtuLJHTF3
/00hSrevLwVo"
set interfaces ge-0/0/2 unit 100 family inet6 address 2a02:cdc5:9715:a000:185:5:202:2/64 primary
set interfaces ge-0/0/2 unit 100 family inet6 address 2a02:cdc5:9715:a000:185:5:202:2/64 vrrp-inet6-group 10 virtual-inet6-address 2a02:cdc5:97
15:a000:185:5:202:254
set interfaces ge-0/0/2 unit 100 family inet6 address 2a02:cdc5:9715:a000:185:5:202:2/64 vrrp-inet6-group 10 virtual-link-local-address fe80::1
85:5:202:254
set interfaces ge-0/0/2 unit 100 family inet6 address 2a02:cdc5:9715:a000:185:5:202:2/64 vrrp-inet6-group 10 priority 150
set interfaces ge-0/0/2 unit 100 family inet6 address 2a02:cdc5:9715:a000:185:5:202:2/64 vrrp-inet6-group 10 preempt
set interfaces ge-0/0/2 unit 100 family inet6 address 2a02:cdc5:9715:a000:185:5:202:2/64 vrrp-inet6-group 10 accept-data
set interfaces ge-0/0/2 unit 169 vlan-id 169
set interfaces ge-0/0/2 unit 169 family inet rpf-check mode loose
set interfaces ge-0/0/2 unit 169 family inet address 169.254.63.74/24
set interfaces ge-0/0/2 unit 169 family inet6 rpf-check mode loose
set interfaces ge-0/0/2 unit 169 family inet6 address 2a02:cdc5:9715:169:254:63:074/64
set interfaces ge-0/0/3 description DOWNLINK-HUAWEI-10G-WAN-TRANSIT
set interfaces ge-0/0/3 unit 0 family inet rpf-check mode loose
set interfaces ge-0/0/3 unit 0 family inet filter input FW_SBTAP_IPV4_IN
set interfaces ge-0/0/3 unit 0 family inet filter output FW_SBTAP_IPV4_OUT
set interfaces ge-0/0/3 unit 0 family inet address 69.221.32.157/31
set interfaces ge-0/0/3 unit 0 family inet6 rpf-check mode loose
set interfaces ge-0/0/3 unit 0 family inet6 filter input FW_SBTAP_IPV6_IN
set interfaces ge-0/0/3 unit 0 family inet6 filter output FW_SBTAP_IPV6_OUT
set interfaces ge-0/0/3 unit 0 family inet6 address 2001:41a8:60:2::7a/126
set interfaces fxp0 unit 0 family inet address 10.0.0.15/24
set forwarding-options rpf-loose-mode-discard family inet
set forwarding-options rpf-loose-mode-discard family inet6
set routing-options options syslog level debug
set routing-options rib inet5.0 static route 100::/128 discard
set routing-options rib inet5.0 static route 100::/128 retain
set routing-options rib inet5.0 static route 100::/128 install
---(more 82)---

```



# Kathará

## configurazione del candidato caricata nel JunOS su kathara



# 3 Emulazione



## sessioni BGP del JunOS su kathara

```
vrnetlab@vr-vmx> show bgp summary
Groups: 18 Peers: 23 Down peers: 19
Table          Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0
    1          1          0          0          0          0
inet6.0
    1          1          0          0          0          0
Peer           AS           InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn  State|#Active/Received/Accepted/Damped...
38.229.6.20    65332         0         0         0         0         59 Active
38.229.46.20   65332         0         0         0         0         59 Active
38.229.240.100 64496         0         0         0         0         59 Active
38.229.244.100 64496         0         0         0         0         59 Active
64.79.149.244  397601        0         0         0         0         59 Active
89.221.32.156  6762          4         3         0         0         22 1/1/1/0          0/0/0/0
169.254.63.100 59715         0         0         0         0         59 Connect
169.254.63.101 59715         0         0         0         0         59 Connect
178.248.237.29 197068        0         0         0         0         59 Active
185.5.200.243  59715         0         0         0         0         59 Connect
185.5.201.241  112           2         2         0         0         22 0/0/0/0          0/0/0/0
185.5.202.131  59715         0         0         0         0         59 Connect
212.114.120.72 199036        0         0         0         0         59 Connect
2001:7b8:62b:1:0:d4ff:fe72:7848 199036        0         0         0         0         0         59 Active
2001:1898:2400:3000:8000::105 397601        0         0         0         0         0         59 Active
2001:41a8:60:2::79 6762          4         3         0         0         5 Establ
inet6.0: 1/1/1/0
2604:8800:60:240::100 64496         0         0         0         0         59 Active
2604:8800:240::100 64496         0         0         0         0         59 Active
2620:0:6b0:8000::20 65332         0         0         0         0         59 Active
2620:0:6b0:ff00::20 65332         0         0         0         0         59 Connect
2a02:cdc5:9715:0:185:5:201:241 112           2         2         2         0         0         8 Establ
inet6.0: 0/0/0/0
2a02:cdc5:9715:169:169:254:63:100 59715         0         0         0         0         0         59 Connect
2a02:cdc5:9715:169:169:254:63:101 59715         0         0         0         0         0         59 Connect
```

# 3 Emulazione



## router FRR su kathara per emulazione transit

```
as_6762# sh bgp summary
```

```
IPv4 Unicast Summary:
```

```
BGP router identifier 89.221.32.156, local AS number 6762 vrf-id 0
```

```
BGP table version 4
```

```
RIB entries 4, using 768 bytes of memory
```

```
Peers 1, using 21 KiB of memory
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	PfxSnt
89.221.32.157	4	59715	4	5	0	0	0	00:00:21	3	4

```
Total number of neighbors 1
```

```
IPv6 Unicast Summary:
```

```
BGP router identifier 89.221.32.156, local AS number 6762 vrf-id 0
```

```
BGP table version 3
```

```
RIB entries 3, using 576 bytes of memory
```

```
Peers 1, using 21 KiB of memory
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	PfxSnt
2001:41a8:60:2::7a	4	59715	4	5	0	0	0	00:00:03	2	3

```
Total number of neighbors 1
```

```
as_6762# █
```



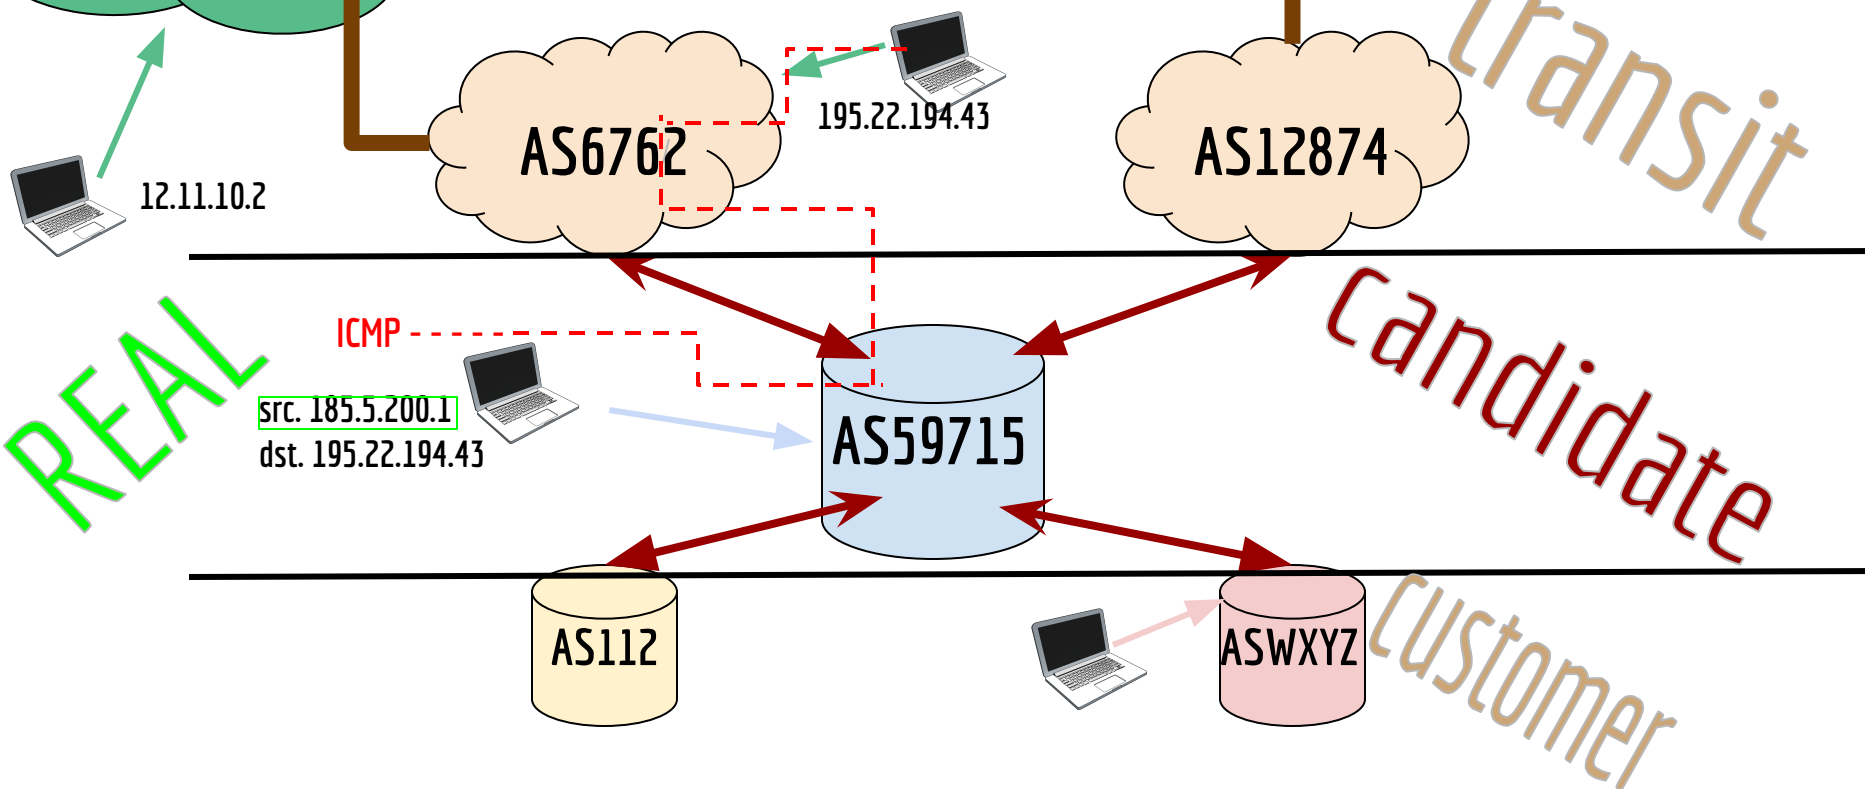
# MANRS

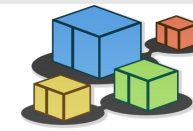
## ACTION 3 for network operators

Blocking traffic from spoofed IP addresses: an operator should implement a system that enables Source Address Validation for its own network and its customers. It is then necessary to include anti-spoofing filters to prevent packets with an incorrect source IP address from entering or leaving the network.

# Internet

## 3 Emulazione





```
root@as_59715_client:/# python3 -m scapy
INFO: Can't import matplotlib. Won't be able to plot.
INFO: Can't import PyX. Won't be able to use psdump() or pdfdump().
WARNING: IPython not available. Using standard Python shell instead.
AutoCompletion, History are disabled.
```

```

      aSPY//YASa
    apyyuyCY/////////YCa
  sY/////////YSpCs  scpCY//Pp
aSP auyuyuySCP//Pp  syY//C
AYAsAYYYYYYYYY//Ps  cY//S
  pCCCCY//p  cSSps y//Y
  SPPPP//a  pP//AC//Y
    A//A  cyP//C
  p//Ac  sC//a
  P//Y/Cpc  A//A
  scccccp//pSP//p  p//Y
  sY/////////y caa  S//P
  caYcyayP//Ya  pY//a
  sY/PsY/////////YCc  aC//p
  sc sccaCY//PCyapaayCP//Ys
    spCPY/////////YPSps
      ccaacs

```

```
| Welcome to Scapy
| Version 2.4.5
| https://github.com/secdev/scapy
| Have fun!
| Craft me if you can. — IPv6 layer
```

```
>>> send(IP(src="185.5.200.1", dst="195.22.194.43")/ICMP())
```

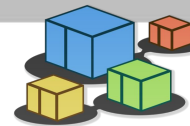
```
* Sent 1 packets.
```

```
>>> █
```

inviamo un pacchetto ICMP

src 185.5.200.1

dst 195.22.194.43



```
root@as_6762_client:~# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
17: eth0@if16: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether 9e:e0:5c:97:6d:ae brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 195.22.194.43/24 scope global eth0
        valid_lft forever preferred_lft forever
root@as_6762_client:~# tcpdump -i eth0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 262144 bytes
10:40:50.290637 IP 185.5.200.1 > 195.22.194.43: ICMP echo request, id 0, seq 0, length 8
10:40:50.290676 IP 195.22.194.43 > 185.5.200.1: ICMP echo reply, id 0, seq 0, length 8
10:40:55.319204 ARP, Request who-has 195.22.194.1 tell 195.22.194.43, length 28
10:40:55.319260 ARP, Request who-has 195.22.194.43 tell 195.22.194.1, length 28
10:40:55.319262 ARP, Reply 195.22.194.43 is-at 9e:e0:5c:97:6d:ae (oui Unknown), length 28
10:40:55.319263 ARP, Reply 195.22.194.1 is-at ee:89:68:ec:3b:bd (oui Unknown), length 28

```

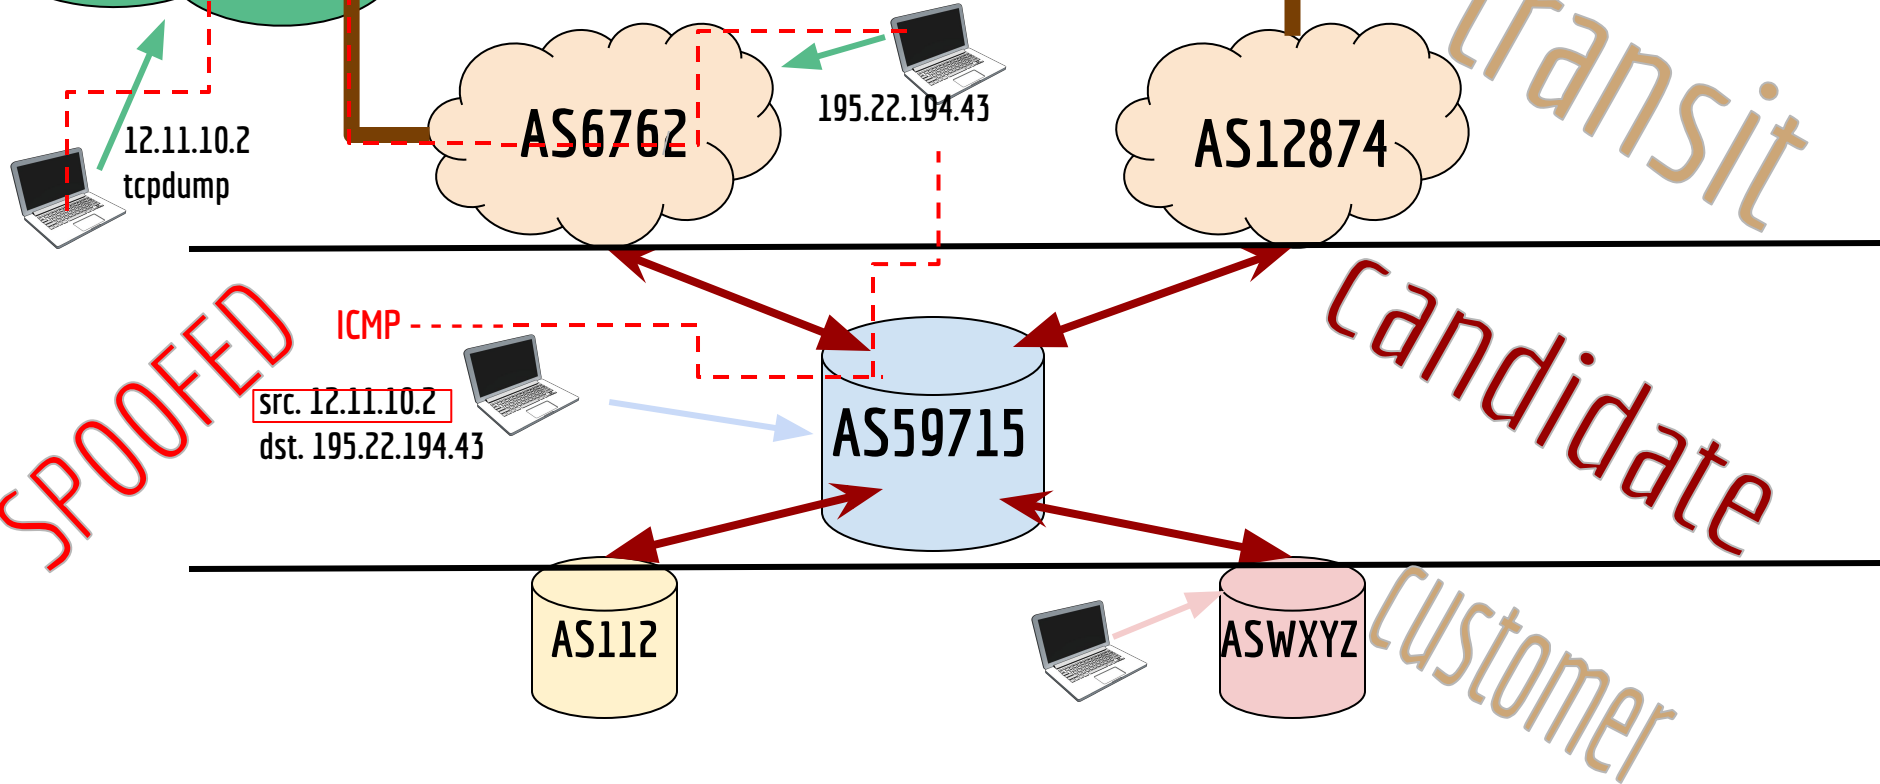
riceviamo il pacchetto ICMP

src 185.5.200.1

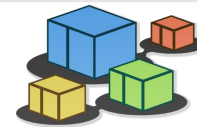
dst 195.22.194.43

# Internet

## 3 Emulazione







```
root@as_59715_client:/# python3 -m scapy
INFO: Can't import matplotlib. Won't be able to plot.
INFO: Can't import PyX. Won't be able to use psdump() or pdfdump().
WARNING: IPython not available. Using standard Python shell instead.
AutoCompletion, History are disabled.
```

```

      aSPY//YASa
    apuuuCY/////////YCa
  sY/////////YSpS  scpCY//Pp
aYp auuuuuSCP//Pp  syY//C
AYAsAYYYYYYYY//Ps  cY//S
  pCCCY//p  cSSps y//Y
  SPPPP//a  pP//AC//Y
    A//A  cyP////C
  p//Ac  sC///a
  P////Cpc  A//A
scccccp///pSP///p  p//Y
sY/////////y caa  S//P
  cayCyayP//Ya  pY//a
  sY/PsY////////C  aC//p
  sc  sccaCY//PCyPaapyCP//YSs
    spCPY/////////YSpS
      ccaacs

```

```

| Welcome to Scapy
| Version 2.4.5
|
| https://github.com/secdev/scapy
|
| Have fun!
|
| We are in France, we say Skappee.
| OK? Merci. — Sebastien Chabal

```

```
>>> send(IP(src="12.11.10.2", dst="195.22.194.43")/ICMP())
```

```
• Sent 1 packets.
```

```
>>> █
```

inviamo un pacchetto ICMP

src 12.11.10.2

dst 195.22.194.43



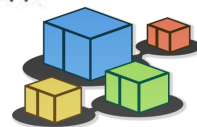
```
root@as_6762_client:~# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
17: eth0@if16: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether 9e:e0:5c:97:6d:ae brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 195.22.194.43/24 scope global eth0
        valid_lft forever preferred_lft forever
root@as_6762_client:~# tcpdump -i eth0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 262144 bytes

```

su 195.22.194.43 non arrivano  
pacchetti ICMP

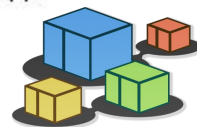


```
root@as_6762_spoof:/# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
25: eth0@if24: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether f6:6a:64:c9:db:2c brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 12.11.10.2/24 scope global eth0
        valid_lft forever preferred_lft forever
root@as_6762_spoof:/# tcpdump -i eth0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 262144 bytes
[]
```



# Kathará

su 12.11.10.2 non arrivano  
pacchetti ICMP



```
root@as_6762_spoof:~# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host localhost
        valid_lft forever preferred_lft forever
    inet 0:0:0:0:0:0:0:0/128 scope host ::1
        valid_lft forever preferred_lft forever
25: eth0@if24: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether f6:ba:64:00:00:db:2c brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 12.11.10.2/24 scope global eth0
        valid_lft forever preferred_lft forever
root@as_6762_spoof:~# tcpdump -i eth0
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 62144
```

u 12.11.10.2 non arrivano  
par netfump



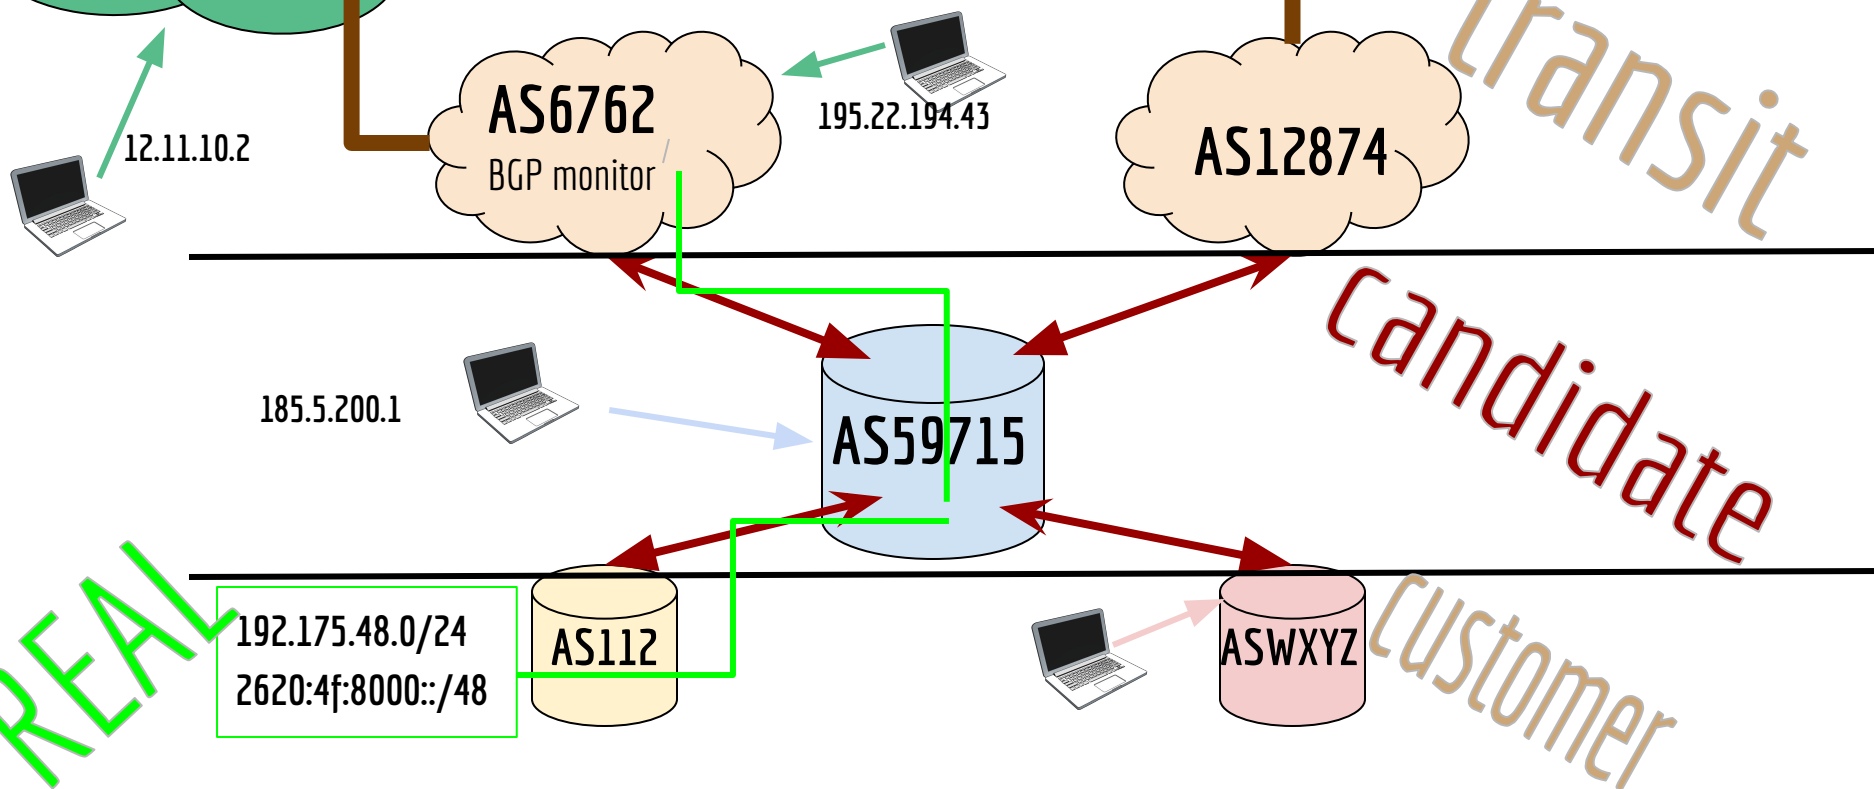
# MANRS

## **ACTION 4 for network operators**

Preventing the propagation of incorrect routing information: it is necessary to implement a system on the basis of which an operator announces to adjacent networks only and exclusively its own prefixes and those of its customers. Likewise, the operator must verify that the prefixes announced by its customers are really their own.

# Internet

## 3 Emulazione

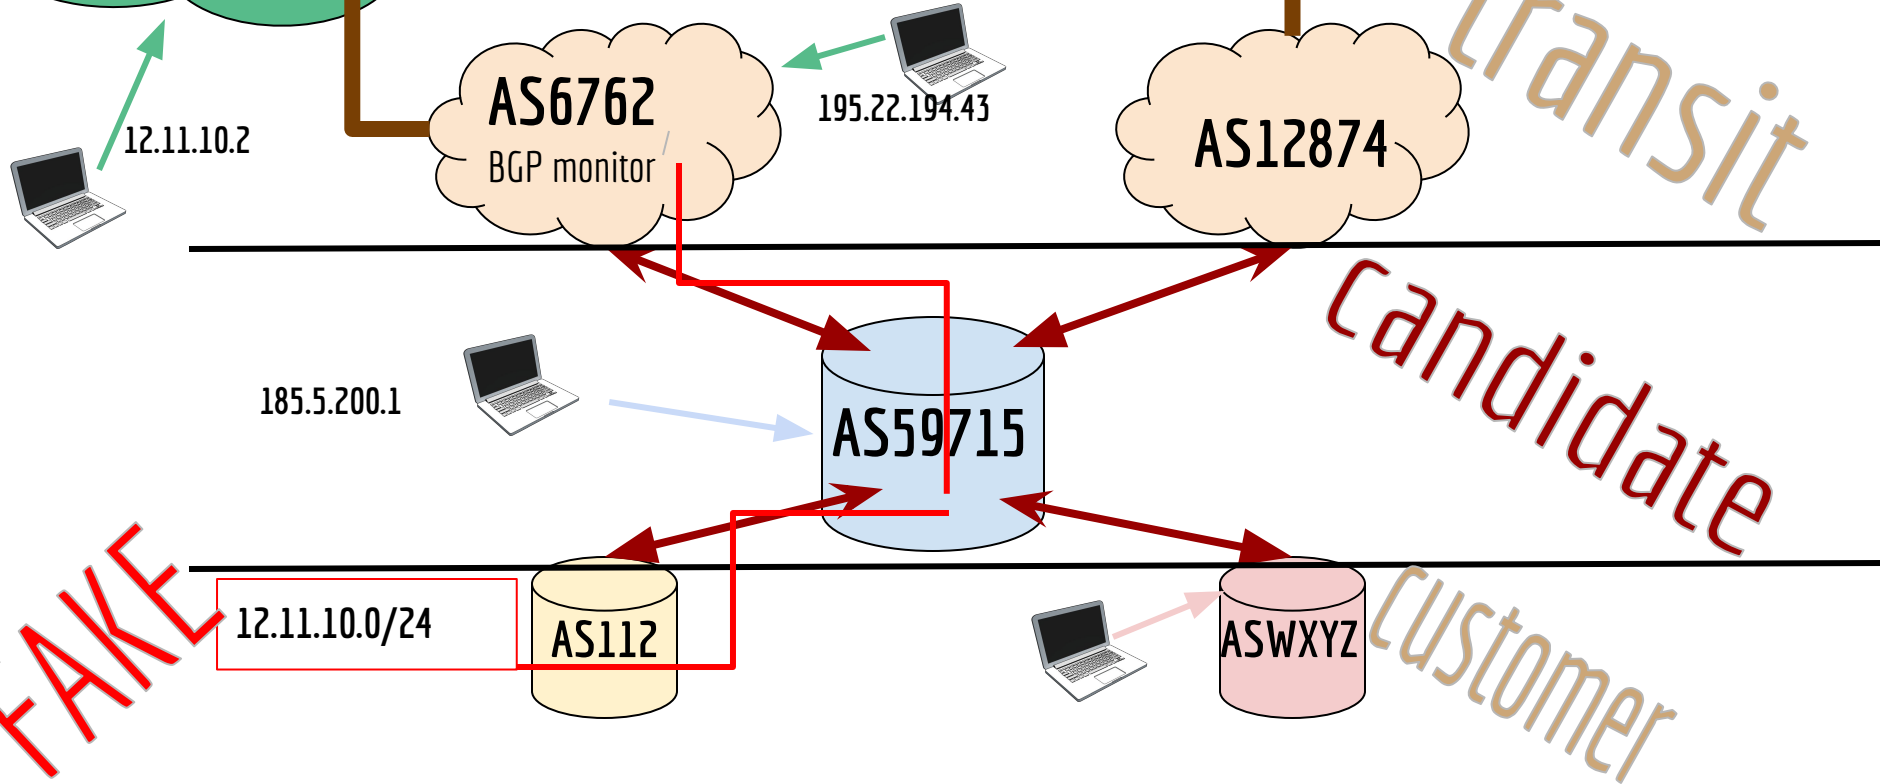


REAL

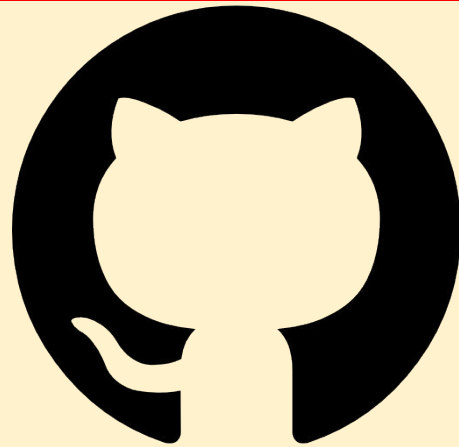


# Internet

## 3 Emulazione



Su GitHub il progetto in fieri:



<https://github.com/Skazza94/rs4lk>

Aiutaci, cerchiamo volontari!



Speciali ringraziamenti vanno:

- al professore **Fabio Fioravanti** per la fiaccola che mantiene sempre accesa durante il mio cammino;
- all'Università degli studi "G. d'Annunzio" di Chieti-Pescara per il sostegno a questa ricerca attraverso la messa a disposizione di amplissima capacità di calcolo;
- al professore **Gianluca Amato** per il prezioso supporto;
- al collega **Mariano Scazzariello** per l'entusiasmo, la professionalità e la continua collaborazione.





RSALK



Domande?

