# Using MikroTik routers for BGP transit and IX points

Juan Miguel Gallardo, MikroTik Trainer and Consultant.
Lisbon, on September 20, 2019.

CODISA
TELCO SOLUTIONS

# ENGINEERING AND PROJECTS

The engine for your ideas


CODISA
TELCO SOLUTIONS

# GLOBAL SUPPORT FOR COMMUNICATION NETWORKS

The best technical support for ISP and Industries.

- PROACTIVE SUPPORT.
- MULTI BRAND SUPPORT.
- CERTIFIED SUPPORT TECHNICIAMS.
- TRANSPARENCY FOR INCIDENTS AND CONFIGURATIONS.

CODISA
TELCO SOLUTIONS

# MIKROTIK TRAINING COURSES

A singular training.

# DEDICATED IP TRANSIT FOR ISP

And others

- Direct circuits.
- Virtual tunnels.
- Backup sceneries.

CODISA
TELCO SOLUTIONS

# TRANSIT AND IX NETWORK

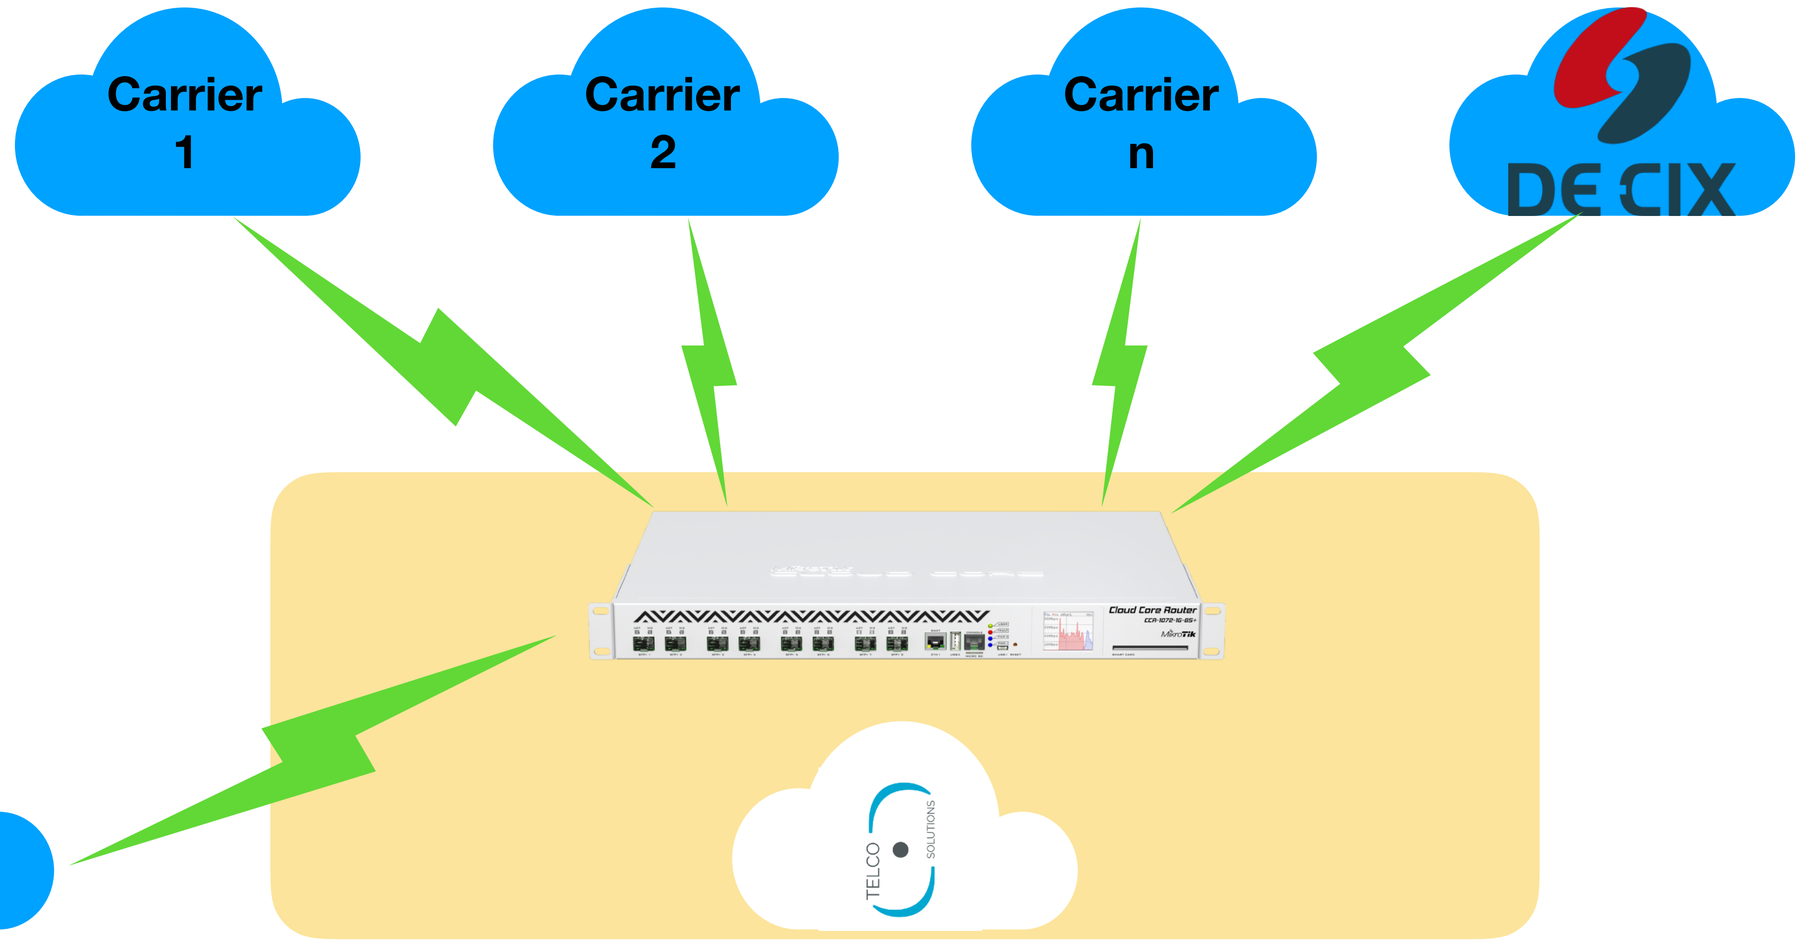


**Full Transit**
**IX Prefixes**
**Default route**

**How do we do it?**

CODISA
TELCO SOLUTIONS

# OWN NETWORKS

ASN 65501

- ASN <=> OWN DOMAIN ==> 65501 (example).
- eBGP <=> Border Gateway Protocol with other ASNs.
- Own networks <=> 10.100.0.0/22, 10.200.0.0/22.
- BGP peers:
  - Transit peer 1: 65510
  - Transit peer 2: 65520
  - DE-CIX route server 1: 48793
  - Customer 1: 65530 <==> 10.200.172.0/22

**We will use private ASN/IPv4 prefixes for this presentation.**
**The shown filters are a very simply configuration for didactic purposes. In real environment, we will need a complex filter configuration to avoid network problems:**
**Own prefixes filtering, bogons filtering, and so on.**

CODISA
TELCO SOLUTIONS

# IMPORT ROUTES ==> OUTGOING TRAFFIC

ASN 65501

- Transit peers: default outgoing traffic when no other preferred.
- Peering: Preferred outgoing traffic.
  - Lower latency.
  - Lower cost.

How to modulate the preference for incoming routes?

- LOCAL_PREF
- SHORTEST AS_PATH
- MED
- OLDEST PATH vs YOUNGER PATH

FILTERS

CODISA
TELCO SOLUTIONS

# IMPORT ROUTES ==> OUTGOING TRAFFIC

ASN 65501

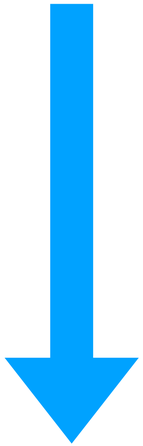

- **LOCAL_PREF:** internal attribute assigned into our network domain.
  - Higher values, preferred routes.
  - Will propagate along our network domain (iBGP), but will not propagate for external peers (eBGP).
- **MED:** Multi Exit Discriminator, can be learned from BGP neighboors.
  - Lower values are for preferred networks.
  - Can be propagated for eBGP peers if they don’t set their own values.

CODISA
TELCO SOLUTIONS

# IMPORT ROUTES ==> OUTGOING TRAFFIC

ASN 65501

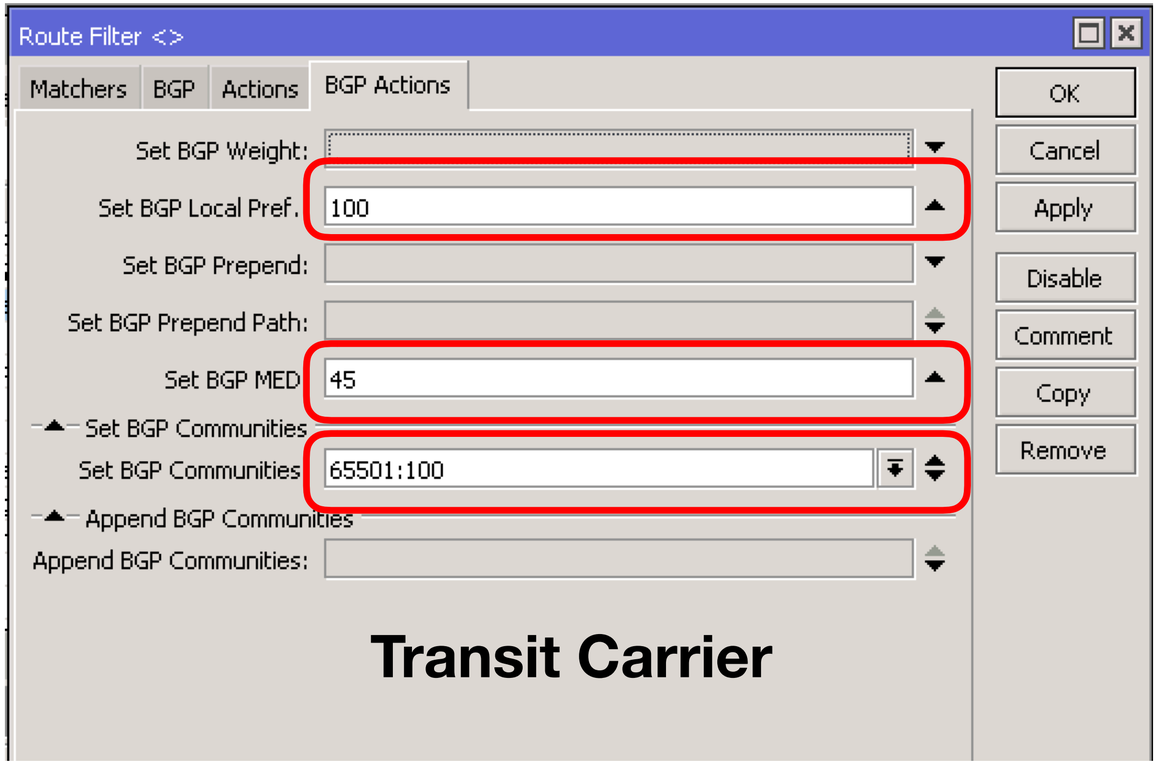


- Local Pref: higher for neutral IX
- BGP MED: lower for neutral IX

- Our outgoing traffic will prefer the IX door.

**Why are we using communities?**

CODISA
TELCO SOLUTIONS

# IMPORT ROUTES ==> OUTGOING TRAFFIC

ASN 65501

**Why are we using communities?**

- We will assign communities over imported routes to ‘mark’ the routes for each provider.
  - It will be useful to provide transit, IX or both routes to our customers, for example.
- In this case:
  - Transit routes will be set with: 65501:100 - 65501:109
  - IX routes will be set with: 65501:110 - 65501:119
- In other cases, we can use communities for:
  - Geo id, router that originates the prefix…
  - To do more complex filters and avoid transit over our network from transit 1 to transit n.
  - Propagate attacked IP address to blackhole servers…

CODISA
TELCO SOLUTIONS

# IMPORT ROUTES ==> OUTGOING TRAFFIC

ASN 65501

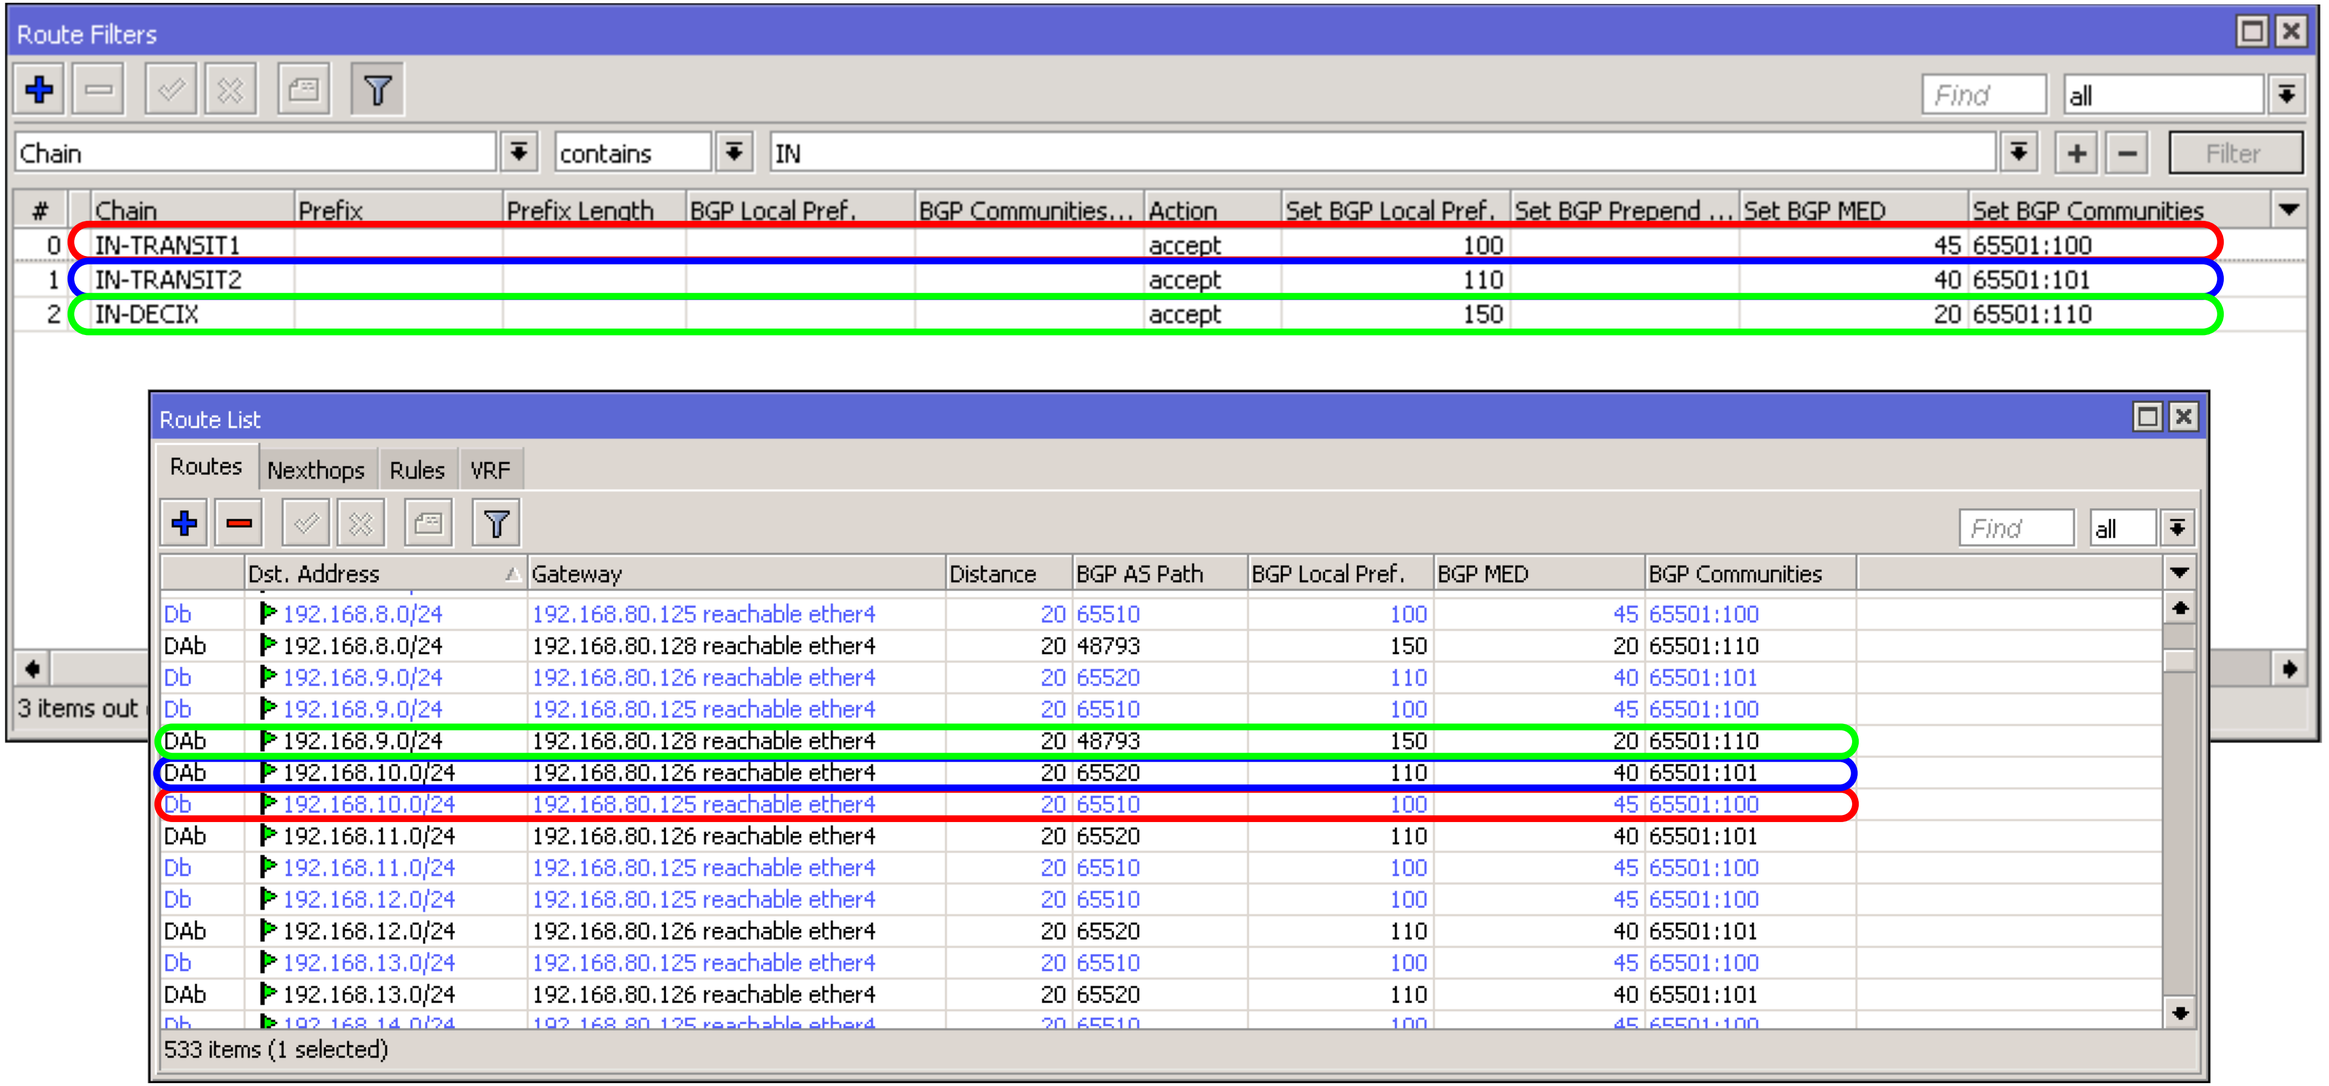


CODISA
TELCO SOLUTIONS

# EXPORT ROUTES ==> INCOMING TRAFFIC

ASN 65501

- Introduce de networks into the BGP world.
- Network size will be used to define if we want to split the aggregate network or not.
  - Advantage: traffic control
  - Disadvantage: more routes in the world.
- The final control will be made by routing filters.
- Optionally, we can create blackhole routes in our routing table.

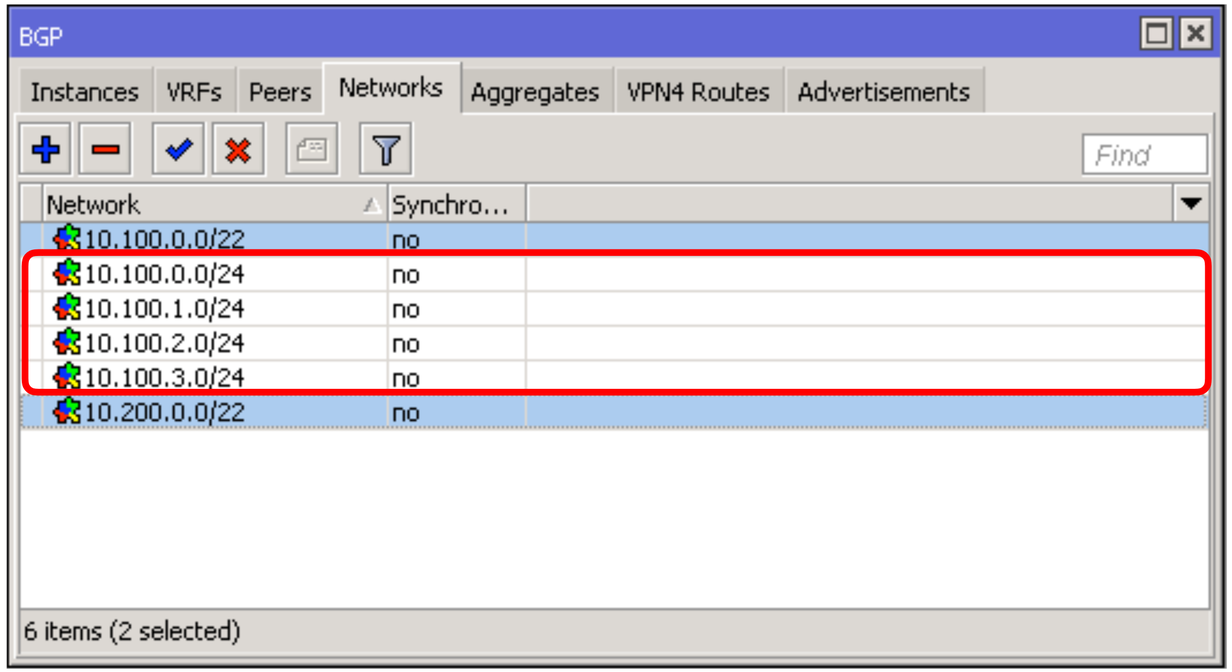


CODISA
TELCO SOLUTIONS

# EXPORT ROUTES ==> INCOMING TRAFFIC

ASN 65501

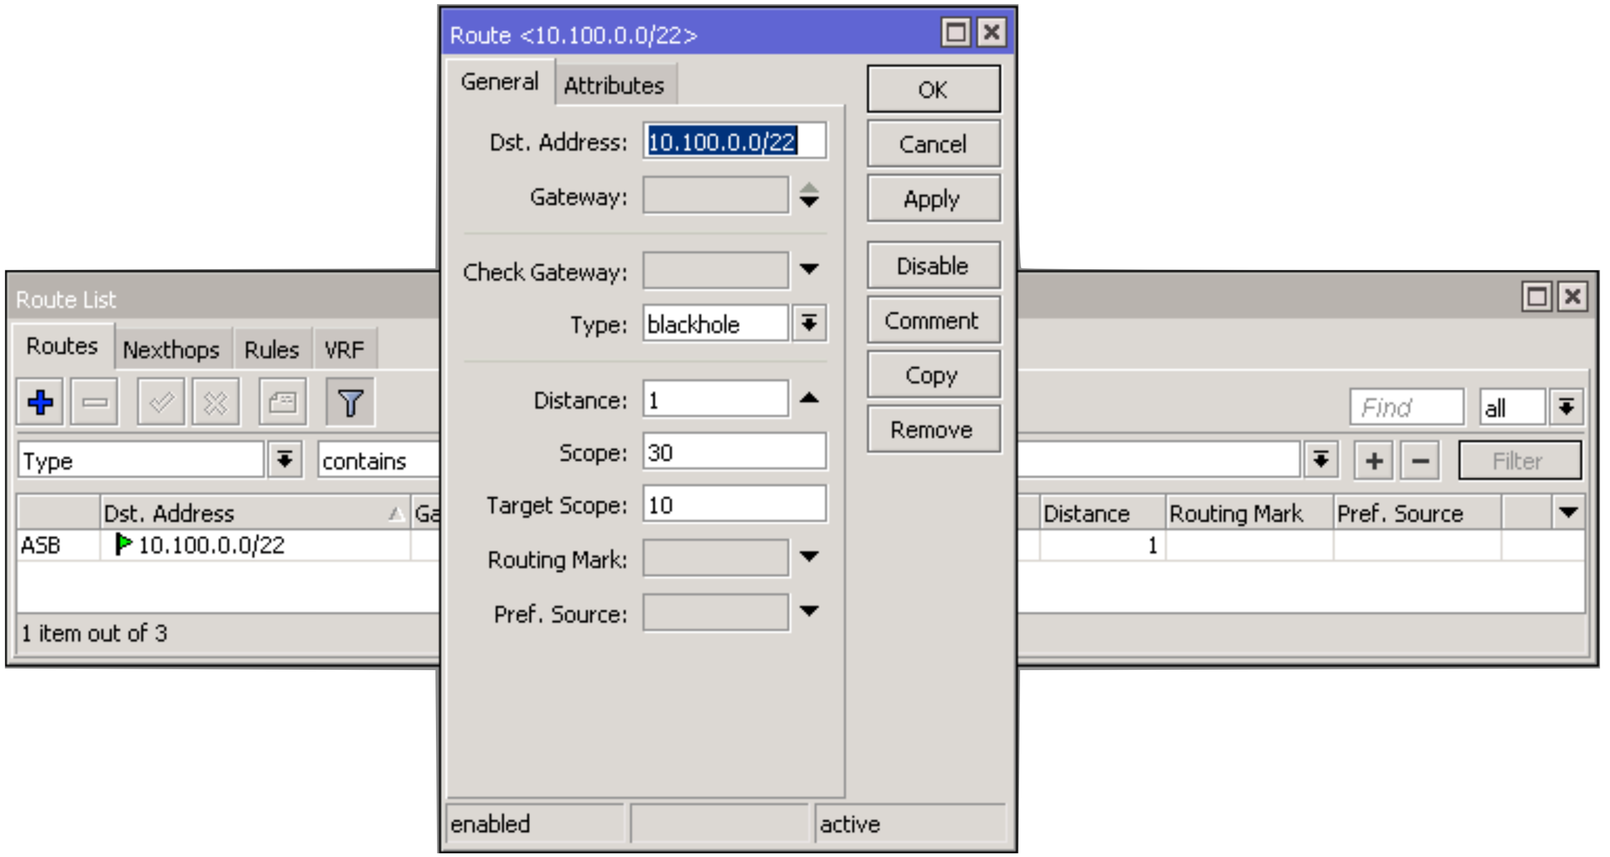


- Attributes aggregation.
- Avoid looping.

CODISA
TELCO SOLUTIONS

# EXPORT ROUTES ==> INCOMING TRAFFIC

ASN 65501

Route Filters

Find | all

Chain | contains | OUT | Filter

| # | Chain | Prefix | Prefix Length | BGP Local Pref. | BGP Communities... | Action | Set BGP Local Pref. | Set BGP Prepend ... | Set BGP MED | Set BGP Communities |
|---|---|---|---|---|---|---|---|---|---|---|
| 3 | OUT-TRANSIT-1 | 10.100.0.0/22 | 22-24 | | | accept | | 65501, 65501 | | |
| 4 | OUT-TRANSIT-1 | 10.200.0.0/22 | 22-24 | | | accept | | 65501, 65501 | | |
| 5 | OUT-TRANSIT-1 | | | | | discard | | | | |
| 6 | OUT-TRANSIT-2 | 10.100.0.0/22 | 22-24 | | | accept | | 65501 | | |
| 7 | OUT-TRANSIT-2 | 10.200.0.0/22 | 22-24 | | | accept | | 65501 | | |
| 8 | OUT-TRANSIT-2 | | | | | discard | | | | |
| 9 | OUT-DECIX | 10.100.0.0/22 | 22-24 | | | accept | | | | |
| 10 | OUT-DECIX | 10.200.0.0/22 | 22-24 | | | accept | | | | |
| 11 | OUT-DECIX | | | | | discard | | | | |

9 items out of 12

CODISA TELCO SOLUTIONS

# EXPORT ROUTES ==> INCOMING TRAFFIC

TRANSIT 1 POINT OF VIEW

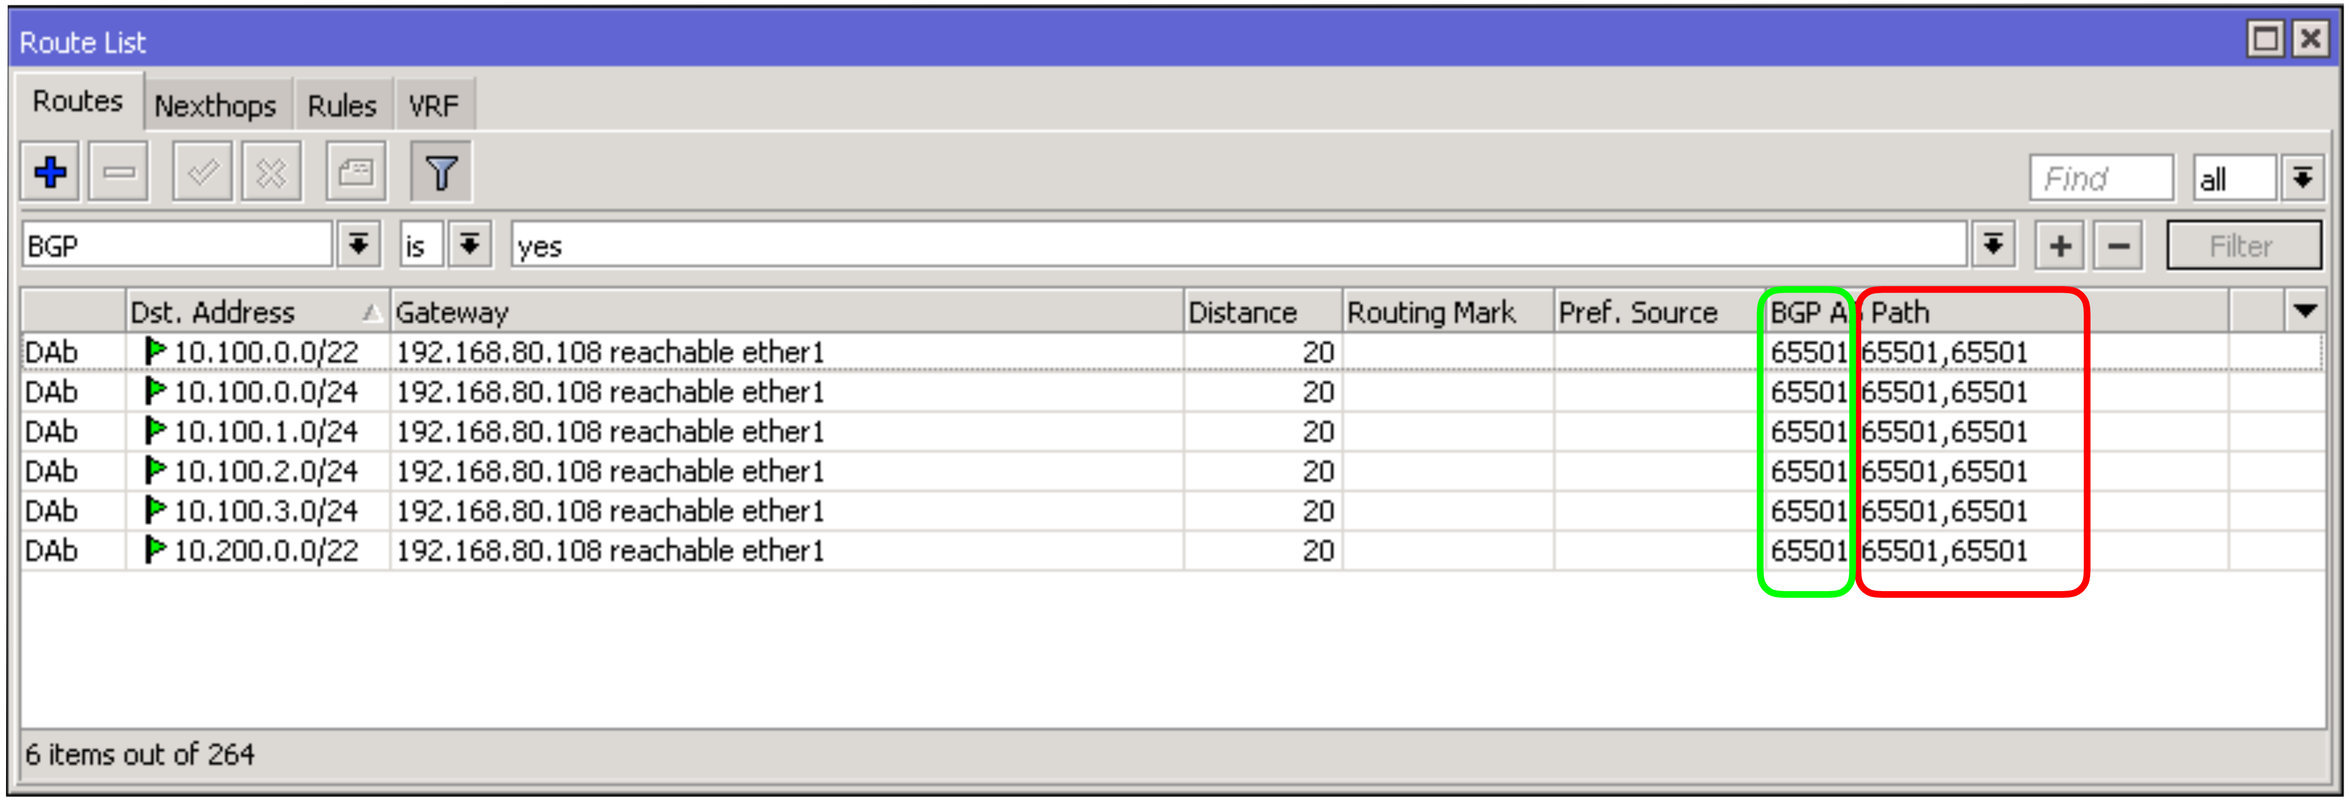


CODISA
TELCO SOLUTIONS

# EXPORT ROUTES ==> INCOMING TRAFFIC

TRANSIT 2 POINT OF VIEW

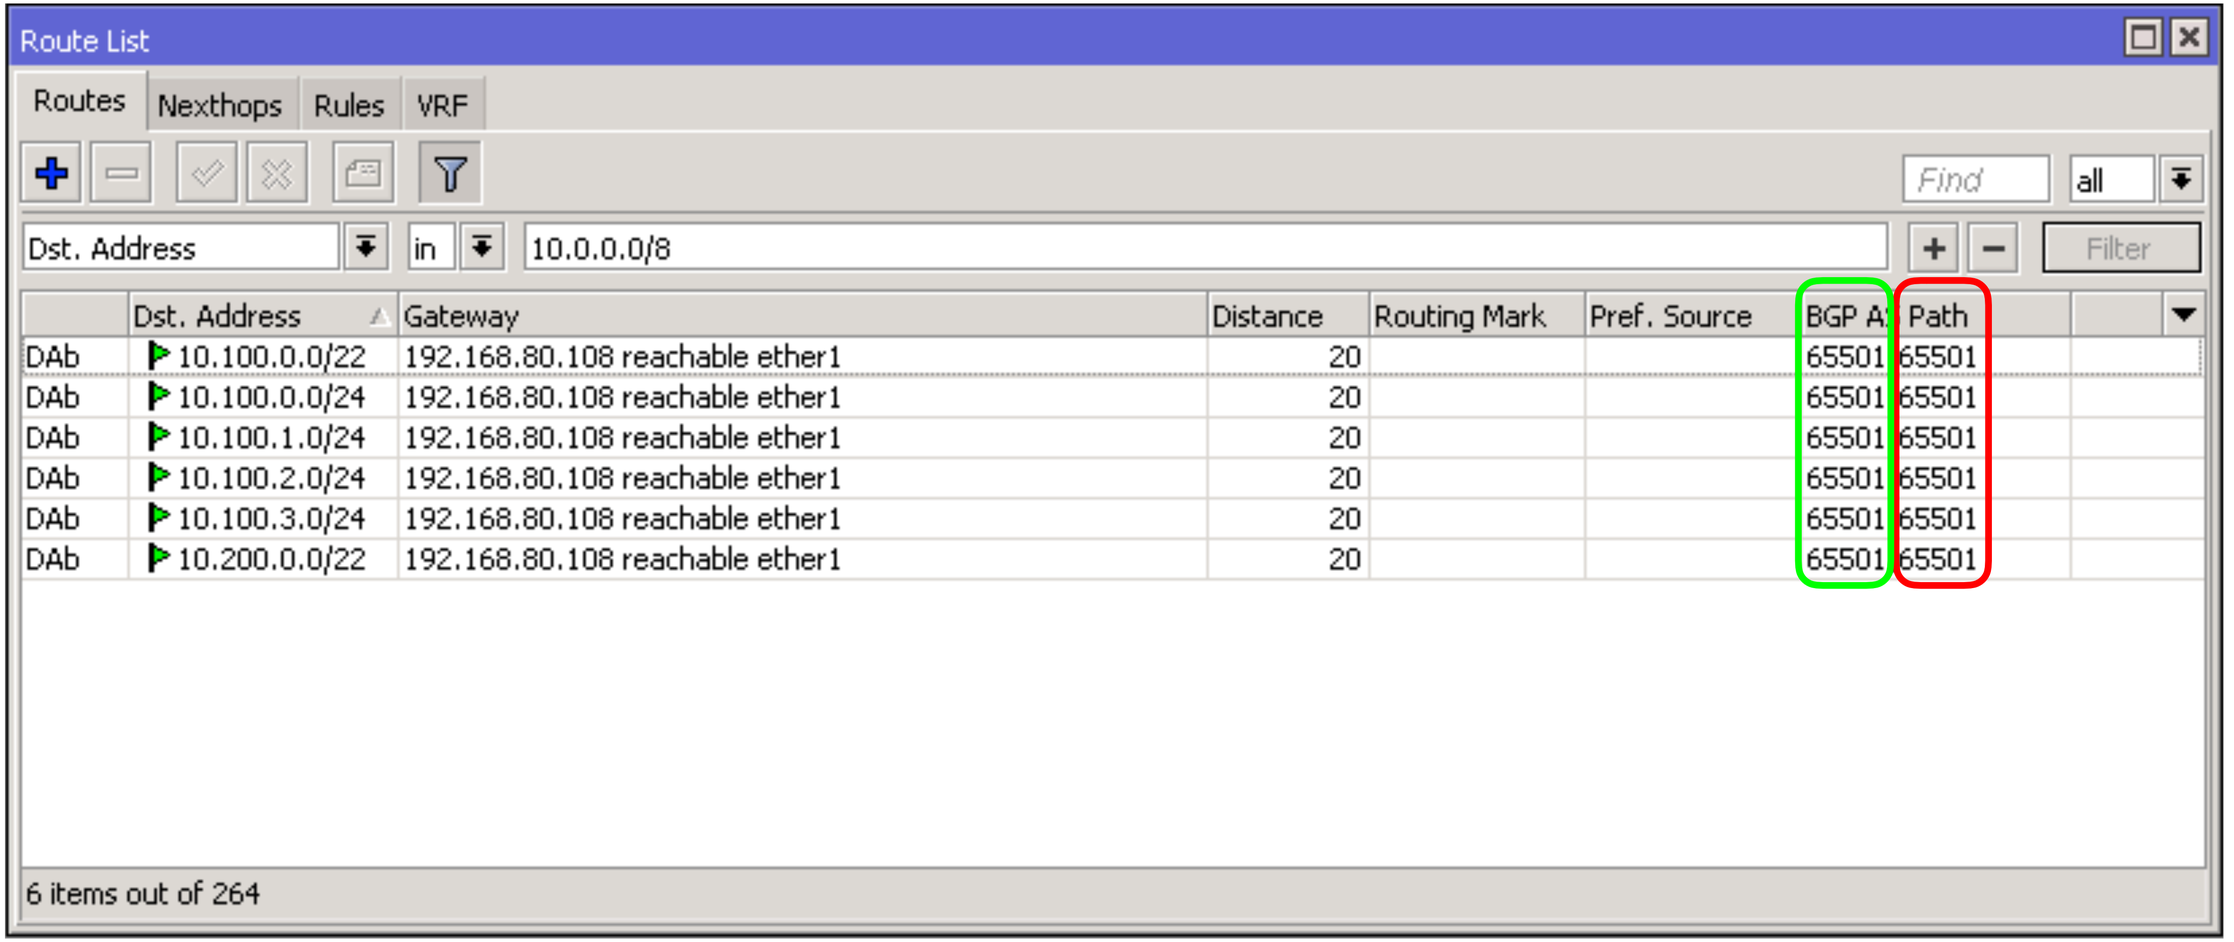


CODISA
TELCO SOLUTIONS

# EXPORT ROUTES ==> INCOMING TRAFFIC

DECIX POINT OF VIEW

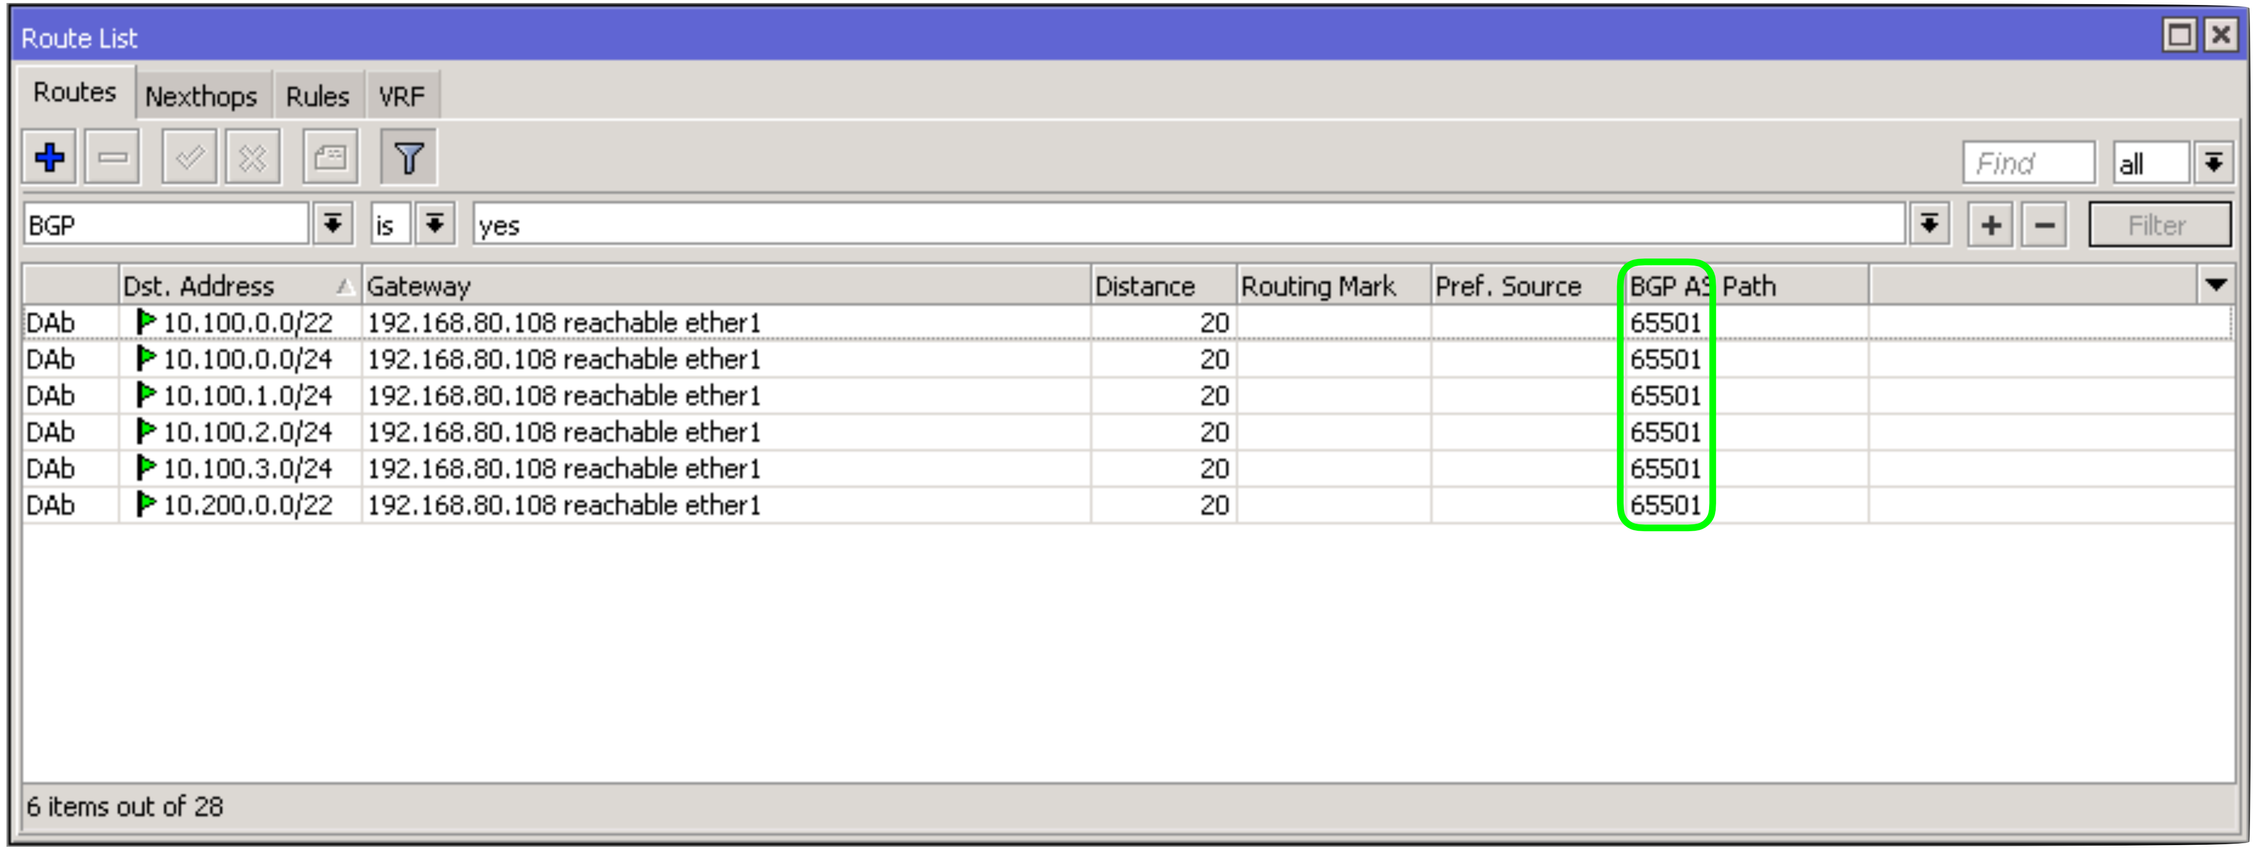


CODISA
TELCO SOLUTIONS

# IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22

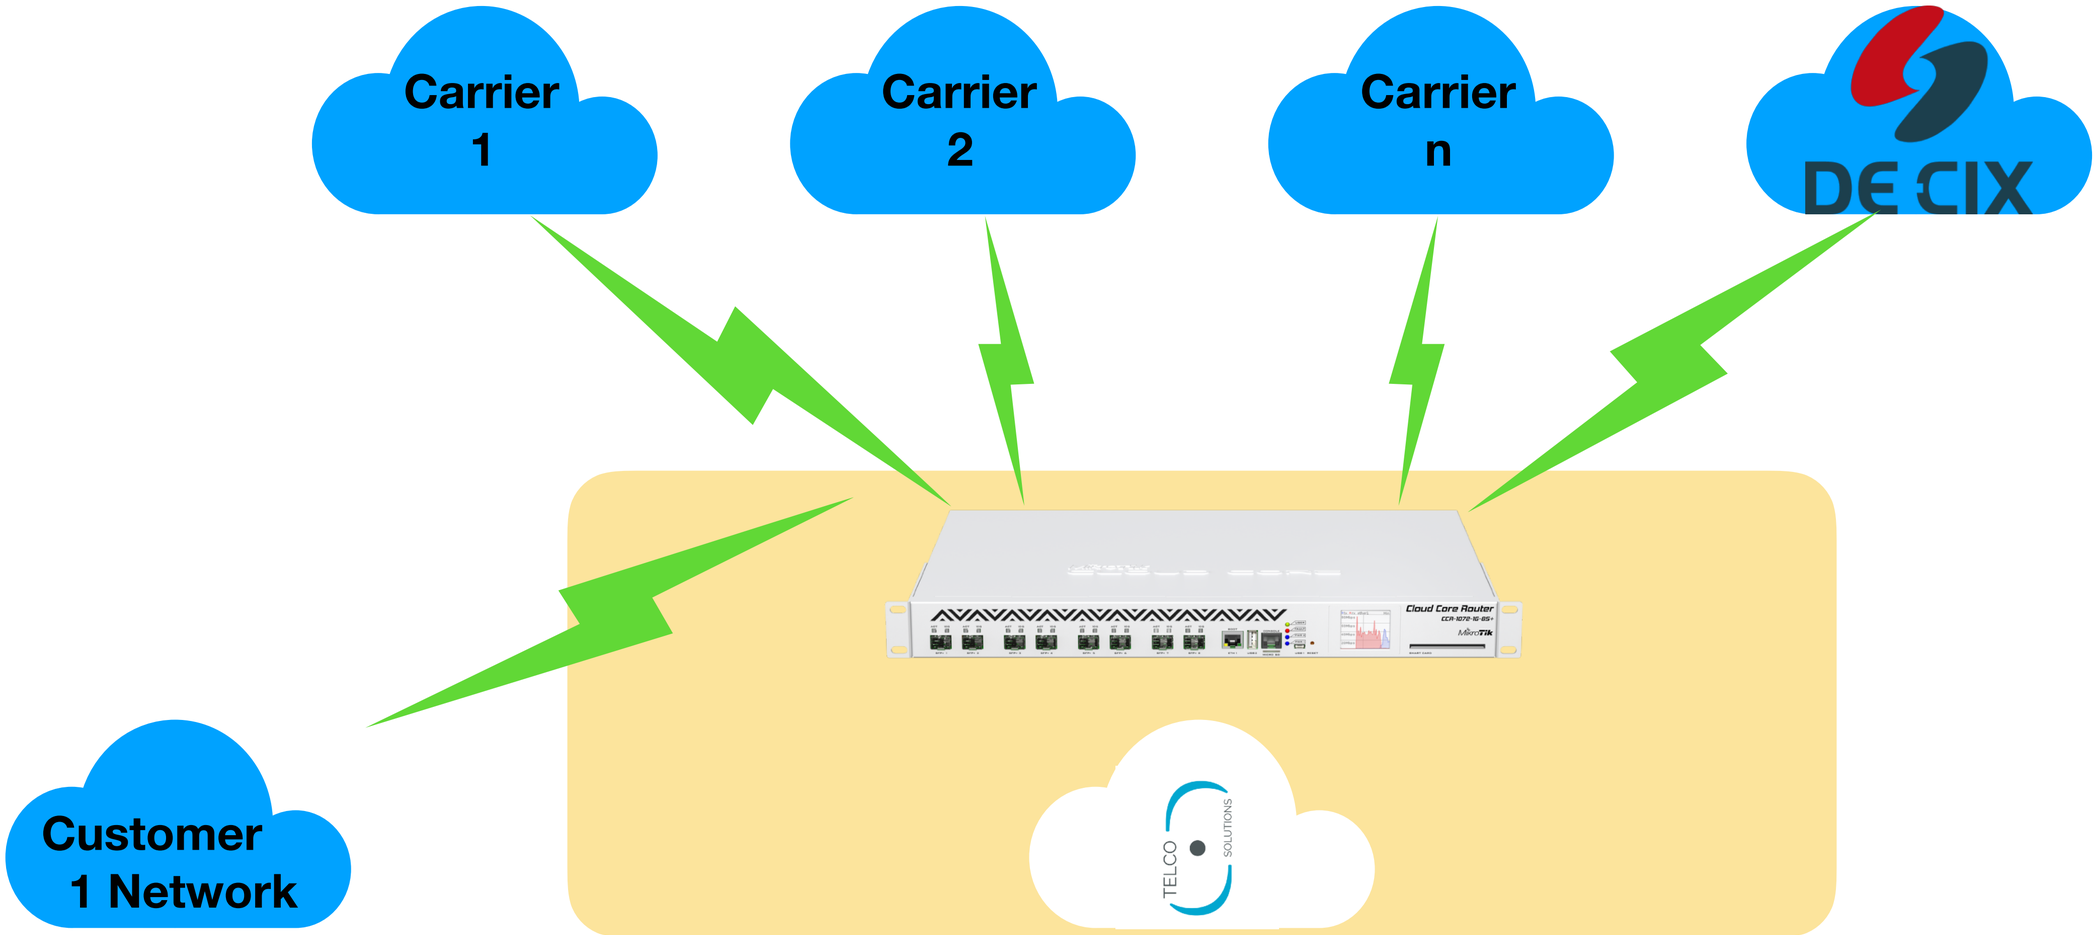


CODISA
TELCO SOLUTIONS

# IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22

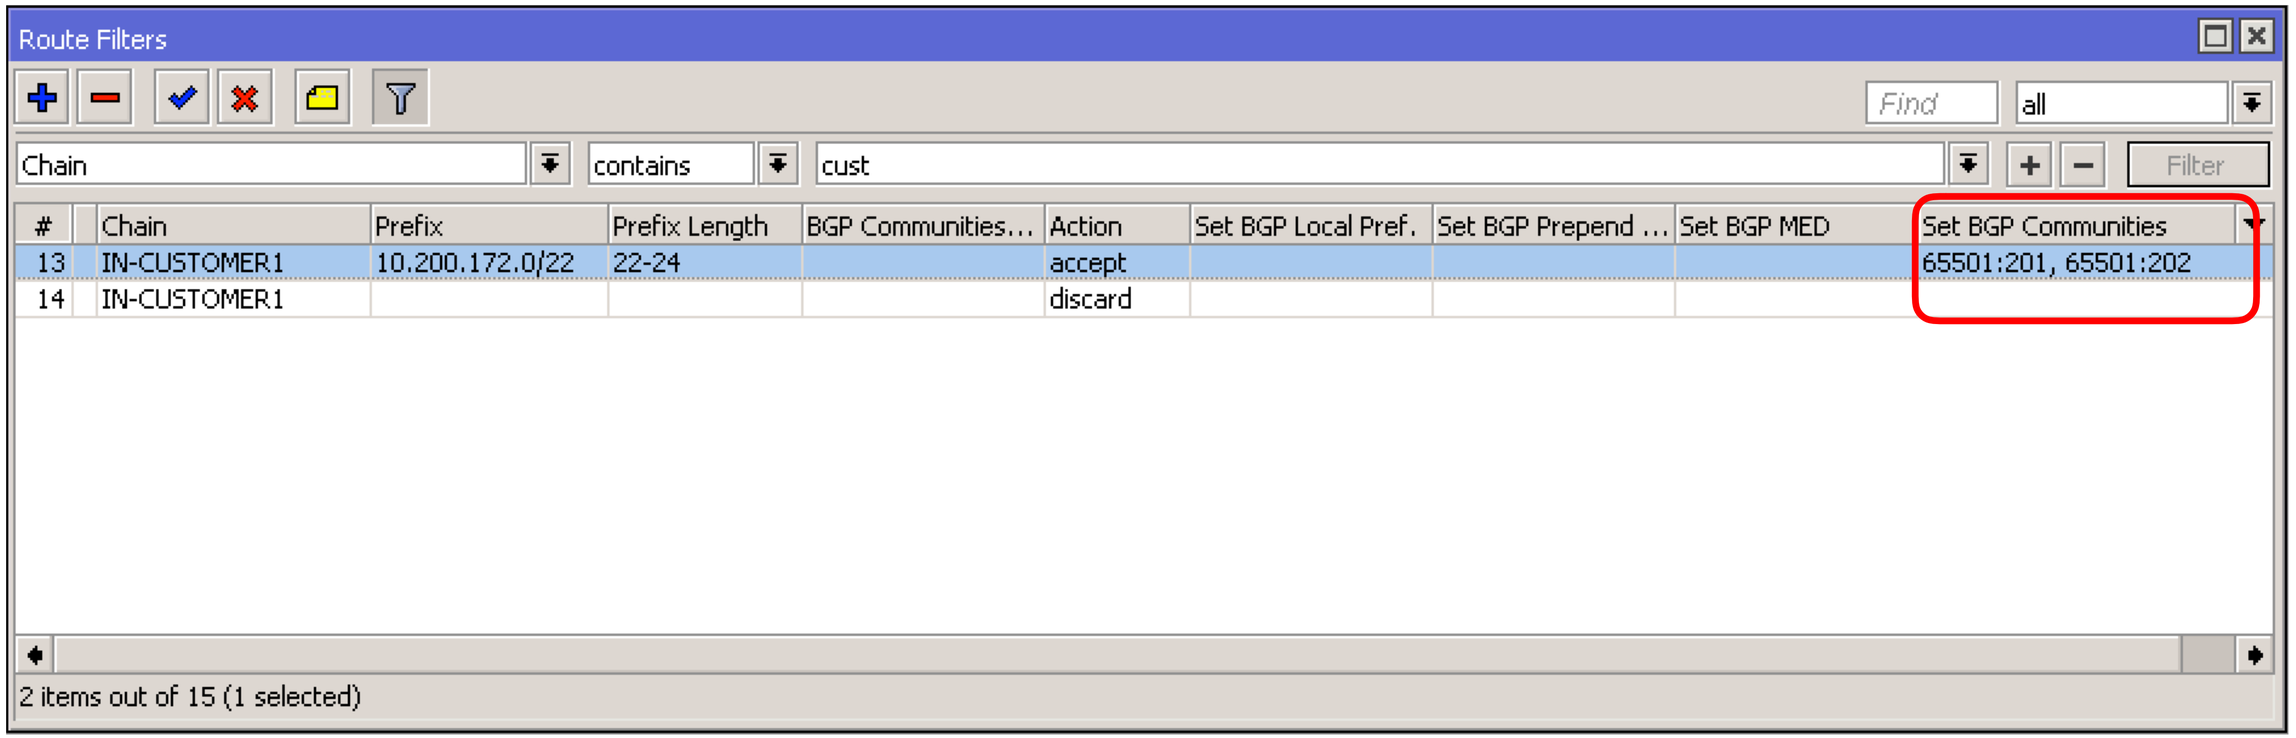


COMMUNITIES:

65501:201—> Announce for transit.

65501:202—> Announce for IX.

CODISA TELCO SOLUTIONS

# IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22

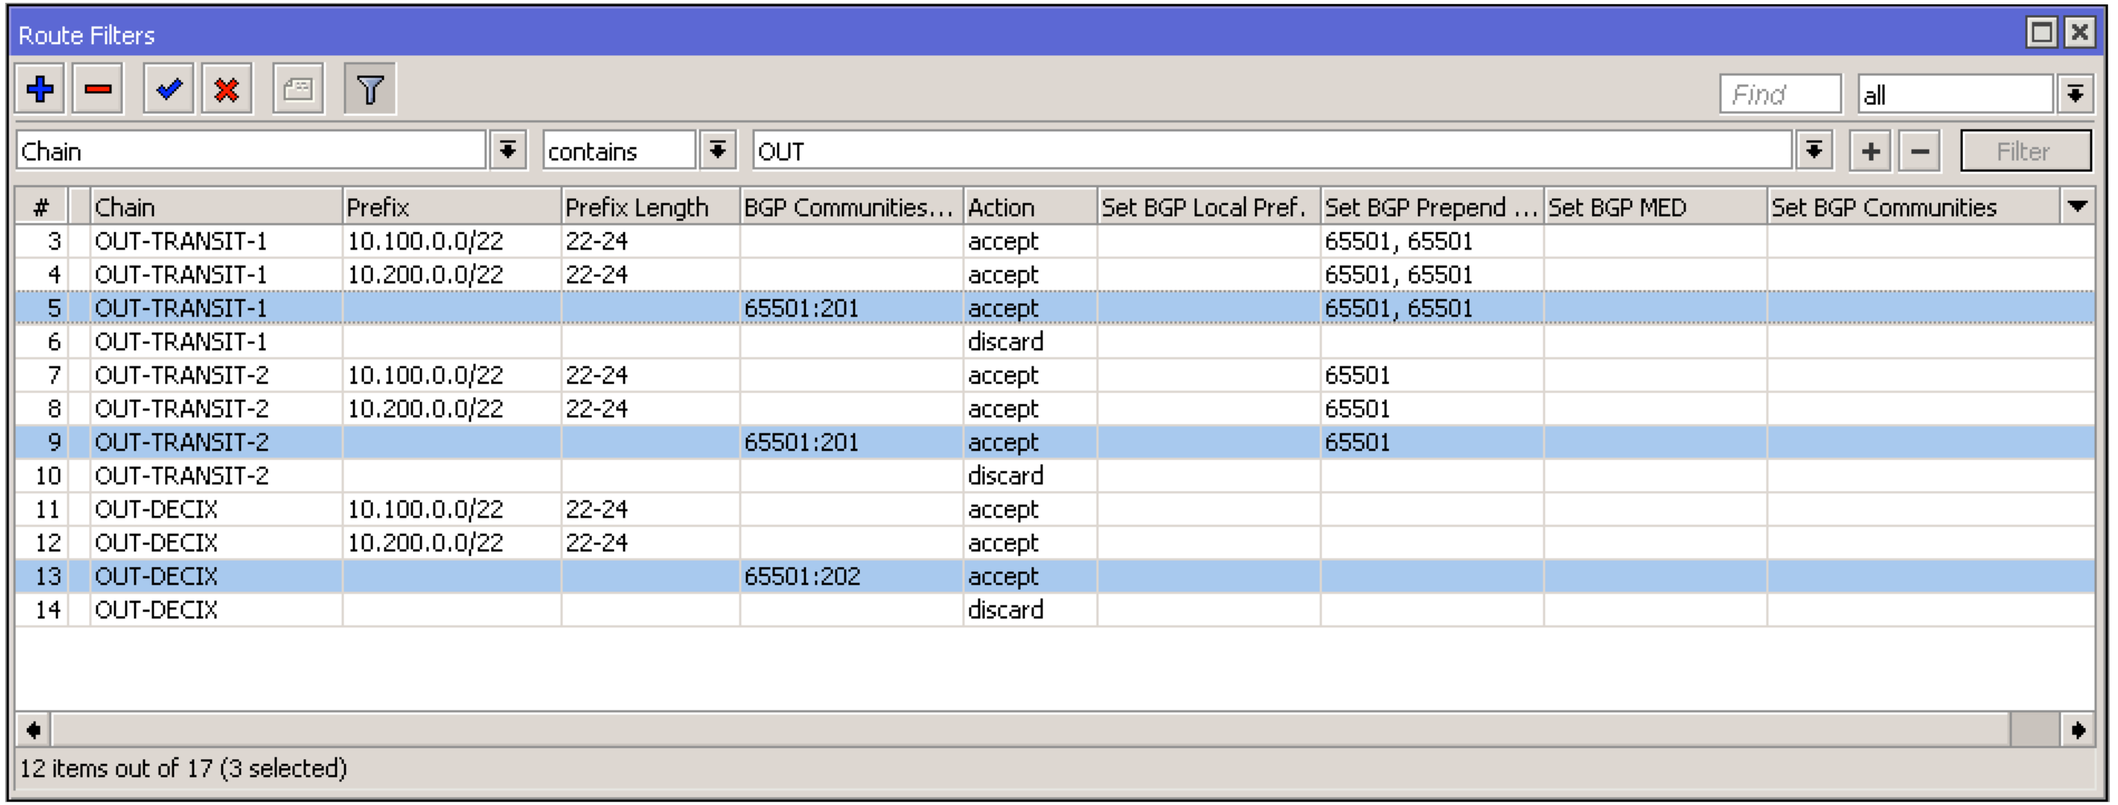


CODISA
TELCO SOLUTIONS

# IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22

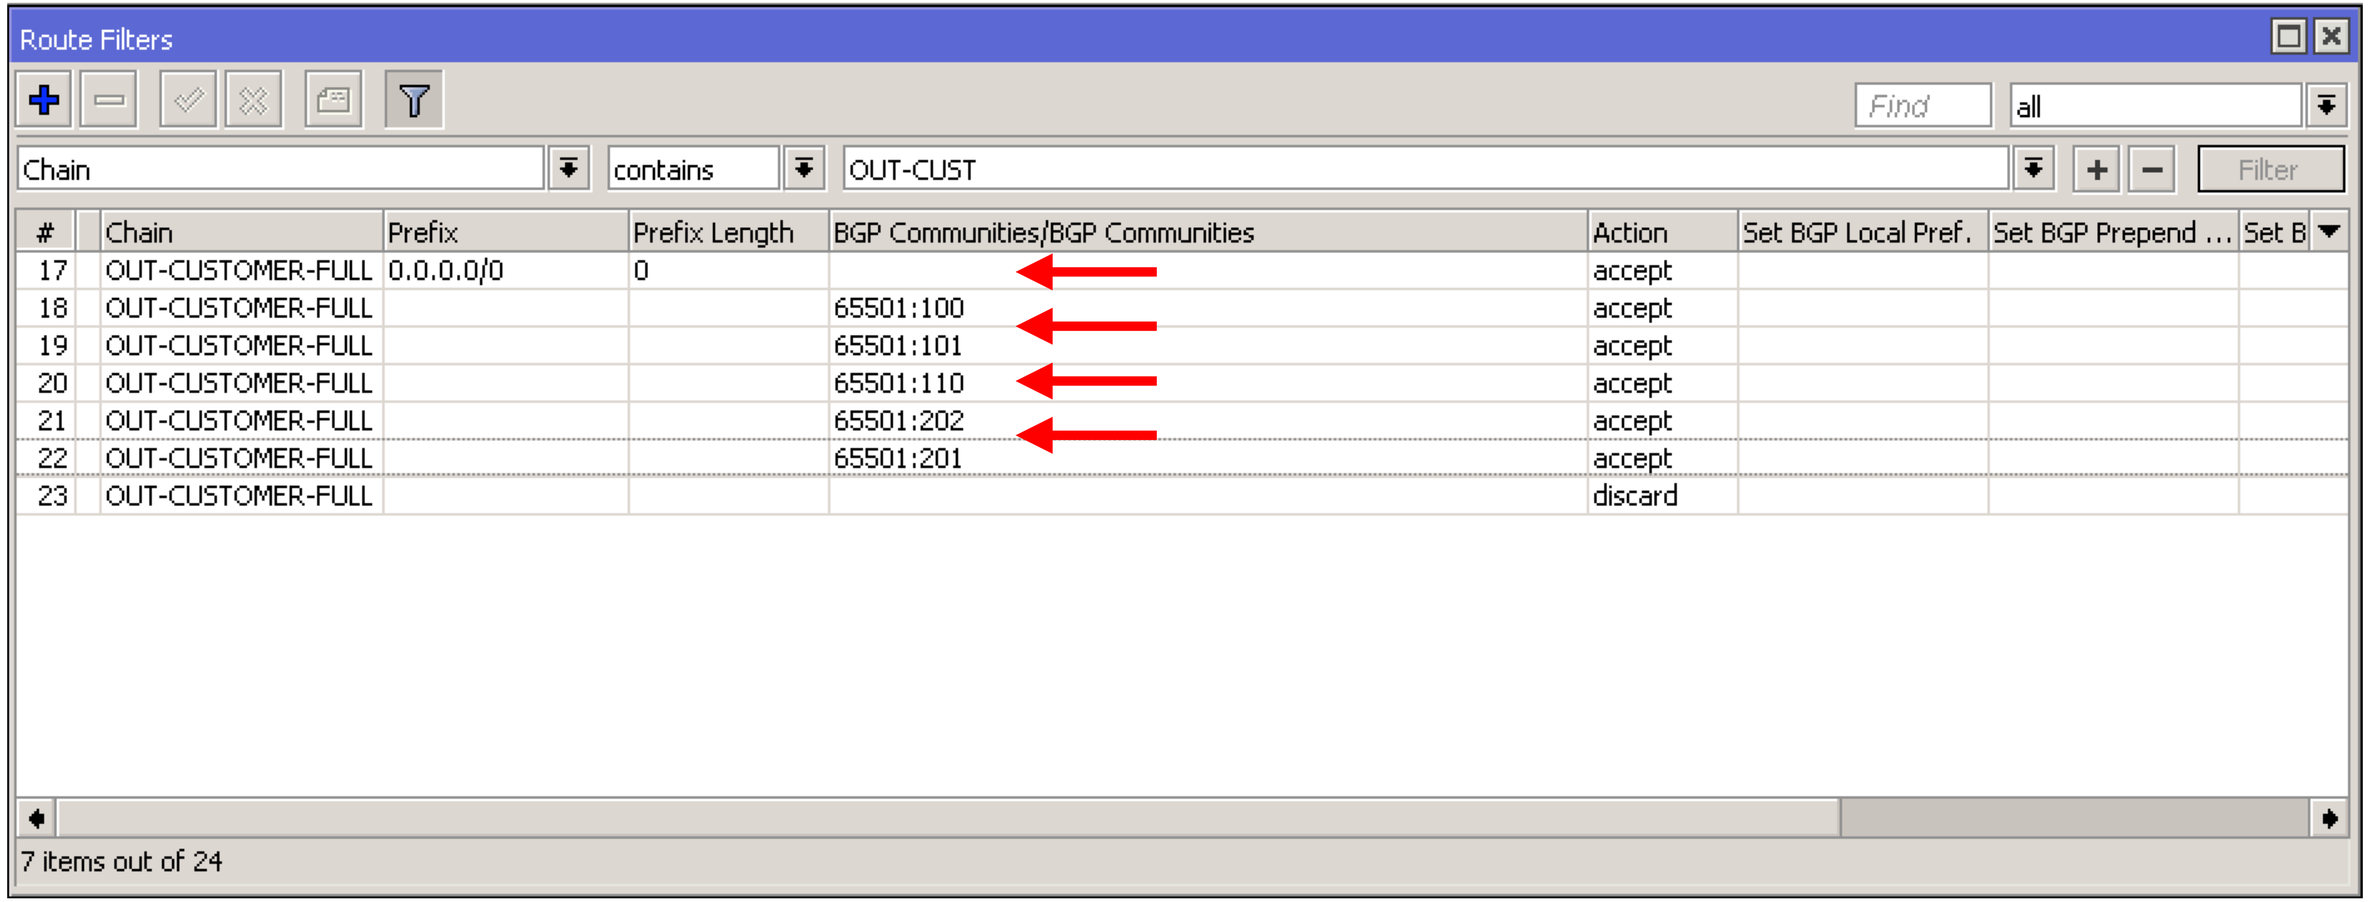


CODISA TELCO SOLUTIONS

# IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22

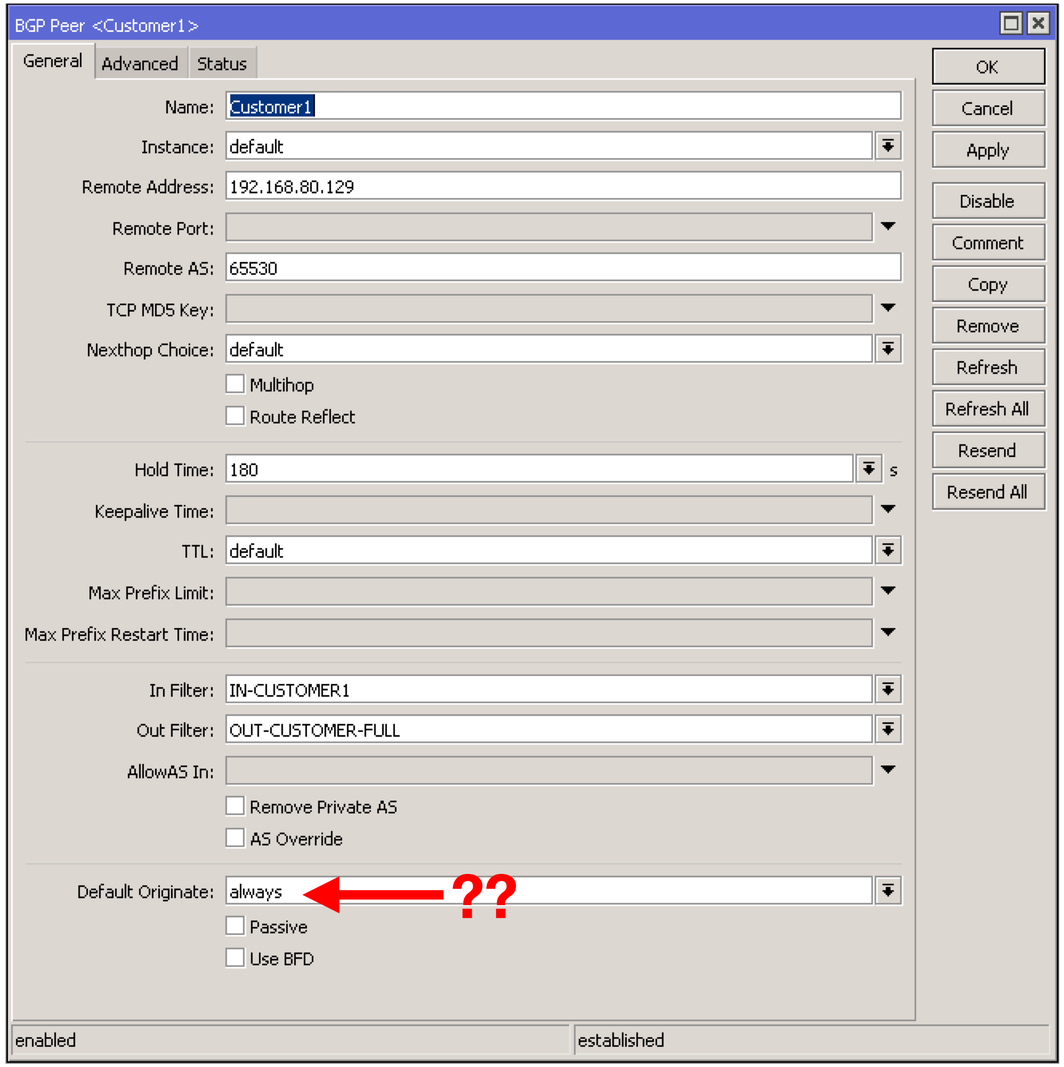


CODISA TELCO SOLUTIONS

# TRANSIT, IX AND CUSTOMERS CONNECTED

IS ANYMORE FOR US?

CODISA
TELCO SOLUTIONS

# OTHER USEFUL USES FOR COMMUNITIES

- Propagate black holing prefixes detected by DDoS detection tools.


We are Fast Netmon Partners, and we can introduce this tool in your network.

CODISA
TELCO SOLUTIONS

# DDoS mitigation

IP: 185.X.Y.Z
Attack uuid: 4cce6e17-b7df-4b69-88c7-718562377d07
Attack severity: middle
Attack type: udp_flood
**Initial attack power: 100029 packets per second**
**Peak attack power: 100029 packets per second**
Attack direction: incoming
Attack protocol: udp
Detection source: automatic
Host network: 185.X.Y.Z/22
Protocol version: IPv4
**Total incoming traffic: 919 mbps**
Total outgoing traffic: 0 mbps
**Total incoming pps: 100029 packets per second**
Total outgoing pps: 92 packets per second
**Incoming udp pps: 99988 packets per second**
Outgoing udp pps: 0 packets per second

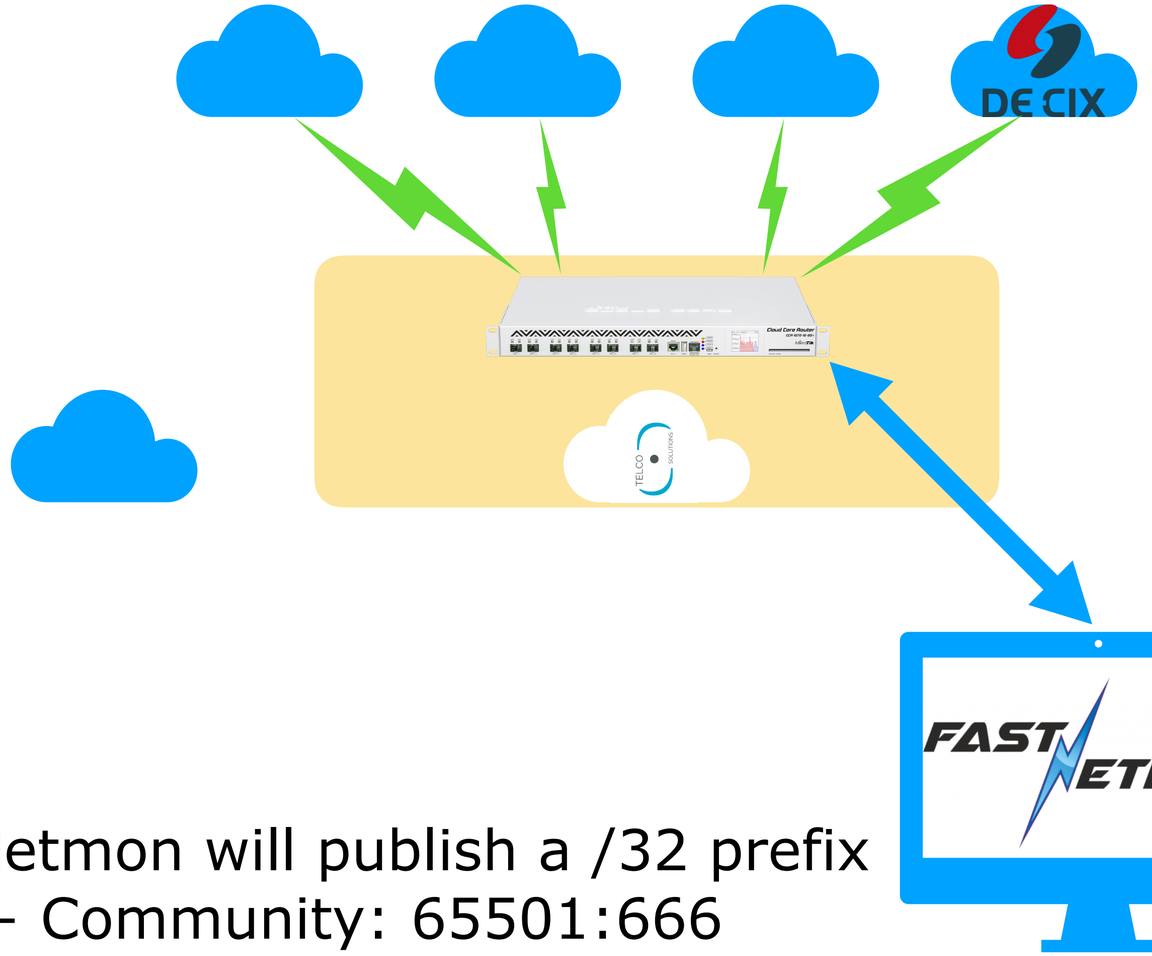


TRAFFIC FLOW Analysis
+
Permanent BGP Session

If Attack...

Fast Netmon will publish a /32 prefix
+ Community: 65501:666

CODISA
TELCO SOLUTIONS

# Recomended Values for incoming filters

Outgoing Traffic

| Localpref | |
|---|---|
| Internal | 999 |
| Customer overweight | 200 |
| Customer Default | 190 |
| Customer Underweight | 180 |
| Peering overweight | 140 |
| Peering Default | 130 |
| Peering underweight | 120 |
| Transit Default | 100 |
| Transit underweight | 90 |

| MED (metric) | |
|---|---|
| Internal | 0 |
| Customer prefixes | 0 for default |
| Peering prefixes | 10 for best<br>20 for worst |
| Transit prefixes | 40 for default<br>Up to 50 for worst |

CODISA
TELCO SOLUTIONS

# What about incoming traffic?

- Set the metric of the sent prefixes to zero. It could be OK if the other party has not set it.
- Try to set some AS prepends on the link you do not want to be used. If the other party decides on the basis of localpref, it doesn’t matter how much you enlarge the AS path.
- Be in touch with the other side to try the route definition together.

# Acknowledgments

Thanks to DE-CIX. They allowed us to use their name, logo and peering guides information for this presentation.

https://www.de-cix.net

Ms. Theresa Bobis: theresa.bobis@de-cix.net
Mr. Da Costa: darwin.costa@de-cix.net

| Name | Address |
| --- | --- |
| Equinix (Itconic) | Av. Severiano Falcão 14<br>2685-378 Prior Velho<br>+351 308 809564 |

CODISA
TELCO SOLUTIONS

TELCO SOLUTIONS

924 11 11 28
info@codisats.es
www.codisats.es
Badajoz - Spain

**NETWORK ENGINEERING**

**TECHNICAL SUPPORT**

**TRAINING**

**INTERNET ACCESS**