

Using MikroTik routers for BGP transit and IX points

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



ENGINEERING AND PROJECTS

The engine for your ideas

- 
- QUALITY.
 - CUSTOMER DEFENSE.
 - SINGULAR PROJECTS.
 - WHITE BRAND FOR COLABORATORS

GLOBAL SUPPORT FOR COMMUNICATION NETWORKS

The best technical support for ISP and Industries.


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



MIKROTIK TRAINING COURSES

A singular training.

- MIKROTIK CERTIFIED EXAMS.
- REAL LABS.
- OWN HANDBOOKS.
- SCHEDULED AND ON DEMAND.
- BASED ON EXPERIENCE.




DEDICATED IP TRANSIT FOR ISP

And others



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

TRANSIT AND IX NETWORK



**Full Transit
IX Prefixes
Default route**

How do we do it?

OWN NETWORKS

ASN 65501

- ASN $\leq \geq$ OWN DOMAIN $\Rightarrow$ 65501 (example).
- eBGP $\leq \geq$ Border Gateway Protocol with other ASNs.
- Own networks $\leq \geq$ 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 $\leq \Rightarrow$ 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.

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

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 neighbors.
 - Lower values are for preferred networks.
 - Can be propagated for eBGP peers if they don't set their own values.

IMPORT ROUTES ==> OUTGOING TRAFFIC

ASN 65501

Route Filter <>

Matchers BGP Actions BGP Actions

Set BGP Weight:

Set BGP Local Pref.:

Set BGP Prepend:

Set BGP Prepend Path:

Set BGP MED:

Set BGP Communities:

Append BGP Communities:

OK Cancel Apply Disable Comment Copy Remove

Transit Carrier

- 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?

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...

IMPORT ROUTES ==> OUTGOING TRAFFIC

ASN 65501

Route Filters

Chain: IN

#	Chain	Prefix	Prefix Length	BGP Local Pref.	BGP Communities...	Action	Set BGP Local Pref.	Set BGP Prepend ...	Set BGP MED	Set BGP Communities
0	IN-TRANSIT1					accept	100		45	65501:100
1	IN-TRANSIT2					accept	110		40	65501:101
2	IN-DECIX					accept	150		20	65501:110

Route List

Routes


	Dst. Address	Gateway	Distance	BGP AS Path	BGP Local Pref.	BGP MED	BGP Communities
Db	192.168.8.0/24	192.168.80.125 reachable ether4	20	65510	100	45	65501:100
DAb	192.168.8.0/24	192.168.80.128 reachable ether4	20	48793	150	20	65501:110
Db	192.168.9.0/24	192.168.80.126 reachable ether4	20	65520	110	40	65501:101
Db	192.168.9.0/24	192.168.80.125 reachable ether4	20	65510	100	45	65501:100
DAb	192.168.9.0/24	192.168.80.128 reachable ether4	20	48793	150	20	65501:110
DAb	192.168.10.0/24	192.168.80.126 reachable ether4	20	65520	110	40	65501:101
Db	192.168.10.0/24	192.168.80.125 reachable ether4	20	65510	100	45	65501:100
DAb	192.168.11.0/24	192.168.80.126 reachable ether4	20	65520	110	40	65501:101
Db	192.168.11.0/24	192.168.80.125 reachable ether4	20	65510	100	45	65501:100
Db	192.168.12.0/24	192.168.80.125 reachable ether4	20	65510	100	45	65501:100
DAb	192.168.12.0/24	192.168.80.126 reachable ether4	20	65520	110	40	65501:101
Db	192.168.13.0/24	192.168.80.125 reachable ether4	20	65510	100	45	65501:100
DAb	192.168.13.0/24	192.168.80.126 reachable ether4	20	65520	110	40	65501:101
Db	192.168.14.0/24	192.168.80.125 reachable ether4	20	65510	100	45	65501:100

533 items (1 selected)

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.




The screenshot shows a BGP configuration window with tabs for Instances, VRFs, Peers, Networks, Aggregates, VPN4 Routes, and Advertisements. The Networks tab is active, displaying a table with columns for Network, Synchronizability, and a third column. The first four rows are highlighted with a red box, indicating they are selected or being edited.

Network	Synchronizability	
10.100.0.0/22	no	
10.100.0.0/24	no	
10.100.1.0/24	no	
10.100.2.0/24	no	
10.100.3.0/24	no	
10.200.0.0/22	no	

6 items (2 selected)

EXPORT ROUTES ==> INCOMING TRAFFIC

ASN 65501



- Attributes aggregation.
- Avoid looping.

EXPORT ROUTES ==> INCOMING TRAFFIC

ASN 65501

Route Filters

Find all

Chain contains OUT

#	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

EXPORT ROUTES ==> INCOMING TRAFFIC

TRANSIT 1 POINT OF VIEW

Route List

Routes | Nexthops | Rules | VRF

+ - ✓ ✗ [icon] Find all [dropdown]

BGP [dropdown] is [dropdown] yes [dropdown] + - Filter

	Dst. Address ▲	Gateway	Distance	Routing Mark	Pref. Source	BGP A	Path	
DAb	▶ 10.100.0.0/22	192.168.80.108 reachable ether1	20			65501	65501,65501	
DAb	▶ 10.100.0.0/24	192.168.80.108 reachable ether1	20			65501	65501,65501	
DAb	▶ 10.100.1.0/24	192.168.80.108 reachable ether1	20			65501	65501,65501	
DAb	▶ 10.100.2.0/24	192.168.80.108 reachable ether1	20			65501	65501,65501	
DAb	▶ 10.100.3.0/24	192.168.80.108 reachable ether1	20			65501	65501,65501	
DAb	▶ 10.200.0.0/22	192.168.80.108 reachable ether1	20			65501	65501,65501	

6 items out of 264

EXPORT ROUTES ==> INCOMING TRAFFIC

TRANSIT 2 POINT OF VIEW

Route List

Routes | Nexthops | Rules | VRF

+ - ✓ ✗ [icon] [icon] Find all [icon]

Dst. Address [icon] in [icon] 10.0.0.0/8 + - Filter

	Dst. Address ▲	Gateway	Distance	Routing Mark	Pref. Source	BGP A	Path	
DAb	▶ 10.100.0.0/22	192.168.80.108 reachable ether1	20			65501	65501	
DAb	▶ 10.100.0.0/24	192.168.80.108 reachable ether1	20			65501	65501	
DAb	▶ 10.100.1.0/24	192.168.80.108 reachable ether1	20			65501	65501	
DAb	▶ 10.100.2.0/24	192.168.80.108 reachable ether1	20			65501	65501	
DAb	▶ 10.100.3.0/24	192.168.80.108 reachable ether1	20			65501	65501	
DAb	▶ 10.200.0.0/22	192.168.80.108 reachable ether1	20			65501	65501	

6 items out of 264

EXPORT ROUTES ==> INCOMING TRAFFIC

DECIX POINT OF VIEW

Route List

Routes | Nexthops | Rules | VRF

Find all

BGP is yes


	Dst. Address	Gateway	Distance	Routing Mark	Pref. Source	BGP AS Path	
DAb	10.100.0.0/22	192.168.80.108 reachable ether1	20			65501	
DAb	10.100.0.0/24	192.168.80.108 reachable ether1	20			65501	
DAb	10.100.1.0/24	192.168.80.108 reachable ether1	20			65501	
DAb	10.100.2.0/24	192.168.80.108 reachable ether1	20			65501	
DAb	10.100.3.0/24	192.168.80.108 reachable ether1	20			65501	
DAb	10.200.0.0/22	192.168.80.108 reachable ether1	20			65501	

6 items out of 28

IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22



IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22

Route Filters

Find

Chain contains

#	Chain	Prefix	Prefix Length	BGP Communities...	Action	Set BGP Local Pref.	Set BGP Prepend ...	Set BGP MED	Set BGP Communities
13	IN-CUSTOMER1	10.200.172.0/22	22-24		accept				65501:201, 65501:202
14	IN-CUSTOMER1				discard				

2 items out of 15 (1 selected)

COMMUNITIES:

65501:201—> Announce for transit.

65501:202—> Announce for IX.

IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22

Route Filters

Find

all

Chain

contains

OUT

Filter

#	Chain	Prefix	Prefix Length	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			65501:201	accept		65501, 65501			
6	OUT-TRANSIT-1				discard					
7	OUT-TRANSIT-2	10.100.0.0/22	22-24		accept		65501			
8	OUT-TRANSIT-2	10.200.0.0/22	22-24		accept		65501			
9	OUT-TRANSIT-2			65501:201	accept		65501			
10	OUT-TRANSIT-2				discard					
11	OUT-DECIX	10.100.0.0/22	22-24		accept					
12	OUT-DECIX	10.200.0.0/22	22-24		accept					
13	OUT-DECIX			65501:202	accept					
14	OUT-DECIX				discard					

12 items out of 17 (3 selected)

IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22

Route Filters

Findall

ChaincontainsOUT-CUST

#	Chain	Prefix	Prefix Length	BGP Communities/BGP Communities	Action	Set BGP Local Pref.	Set BGP Prepend ...	Set B
17	OUT-CUSTOMER-FULL	0.0.0.0/0	0		accept			
18	OUT-CUSTOMER-FULL			65501:100	accept			
19	OUT-CUSTOMER-FULL			65501:101	accept			
20	OUT-CUSTOMER-FULL			65501:110	accept			
21	OUT-CUSTOMER-FULL			65501:202	accept			
22	OUT-CUSTOMER-FULL			65501:201	accept			
23	OUT-CUSTOMER-FULL				discard			

7 items out of 24

IMPORT // EXPORT CUSTOMER ROUTES

ASN 65530

PREFIX: 10.200.172.0/22

BGP Peer <Customer1>

General Advanced Status

Name: Customer1

Instance: default

Remote Address: 192.168.80.129

Remote Port:

Remote AS: 65530

TCP MD5 Key:

Nexthop Choice: default

☐ Multihop

☐ Route Reflect

Hold Time: 180 s

Keepalive Time:

TTL: default

Max Prefix Limit:

Max Prefix Restart Time:

In Filter: IN-CUSTOMER1

Out Filter: OUT-CUSTOMER-FULL

AllowAS In:

☐ Remove Private AS

☐ AS Override

Default Originate: always

☐ Passive

☐ Use BFD

OK

Cancel

Apply

Disable

Comment

Copy

Remove

Refresh

Refresh All

Resend

Resend All

enabled established

TRANSIT, IX AND CUSTOMERS CONNECTED

IS ANYMORE FOR US?

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.

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



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

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 20 for worst
Transit prefixes	40 for default Up to 50 for worst

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 2685-378 Prior Velho +351 308 809564



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

**NETWORK
ENGINEERING**

**TECHNICAL
SUPPORT**

TRAINING

**INTERNET
ACCESS**