

# RouterOS: Peering the World

by Lorenzo Busatti

MUM CZECH REPUBLIC ON MARCH 27 - 28,  
2015

# About me

## Lorenzo Busatti

- Founder of Grifonline S.r.l., Grosseto – ITALY
- ISP since 1997, WISP since 2006
- MikroTik Certified Trainer since 2010 for:  
MTCNA, MTCWE, MTCRE, MTCTCE, MTCUME,  
MTCINE
- Member of RIPE, AMS-IX, MIX-IT

# About me

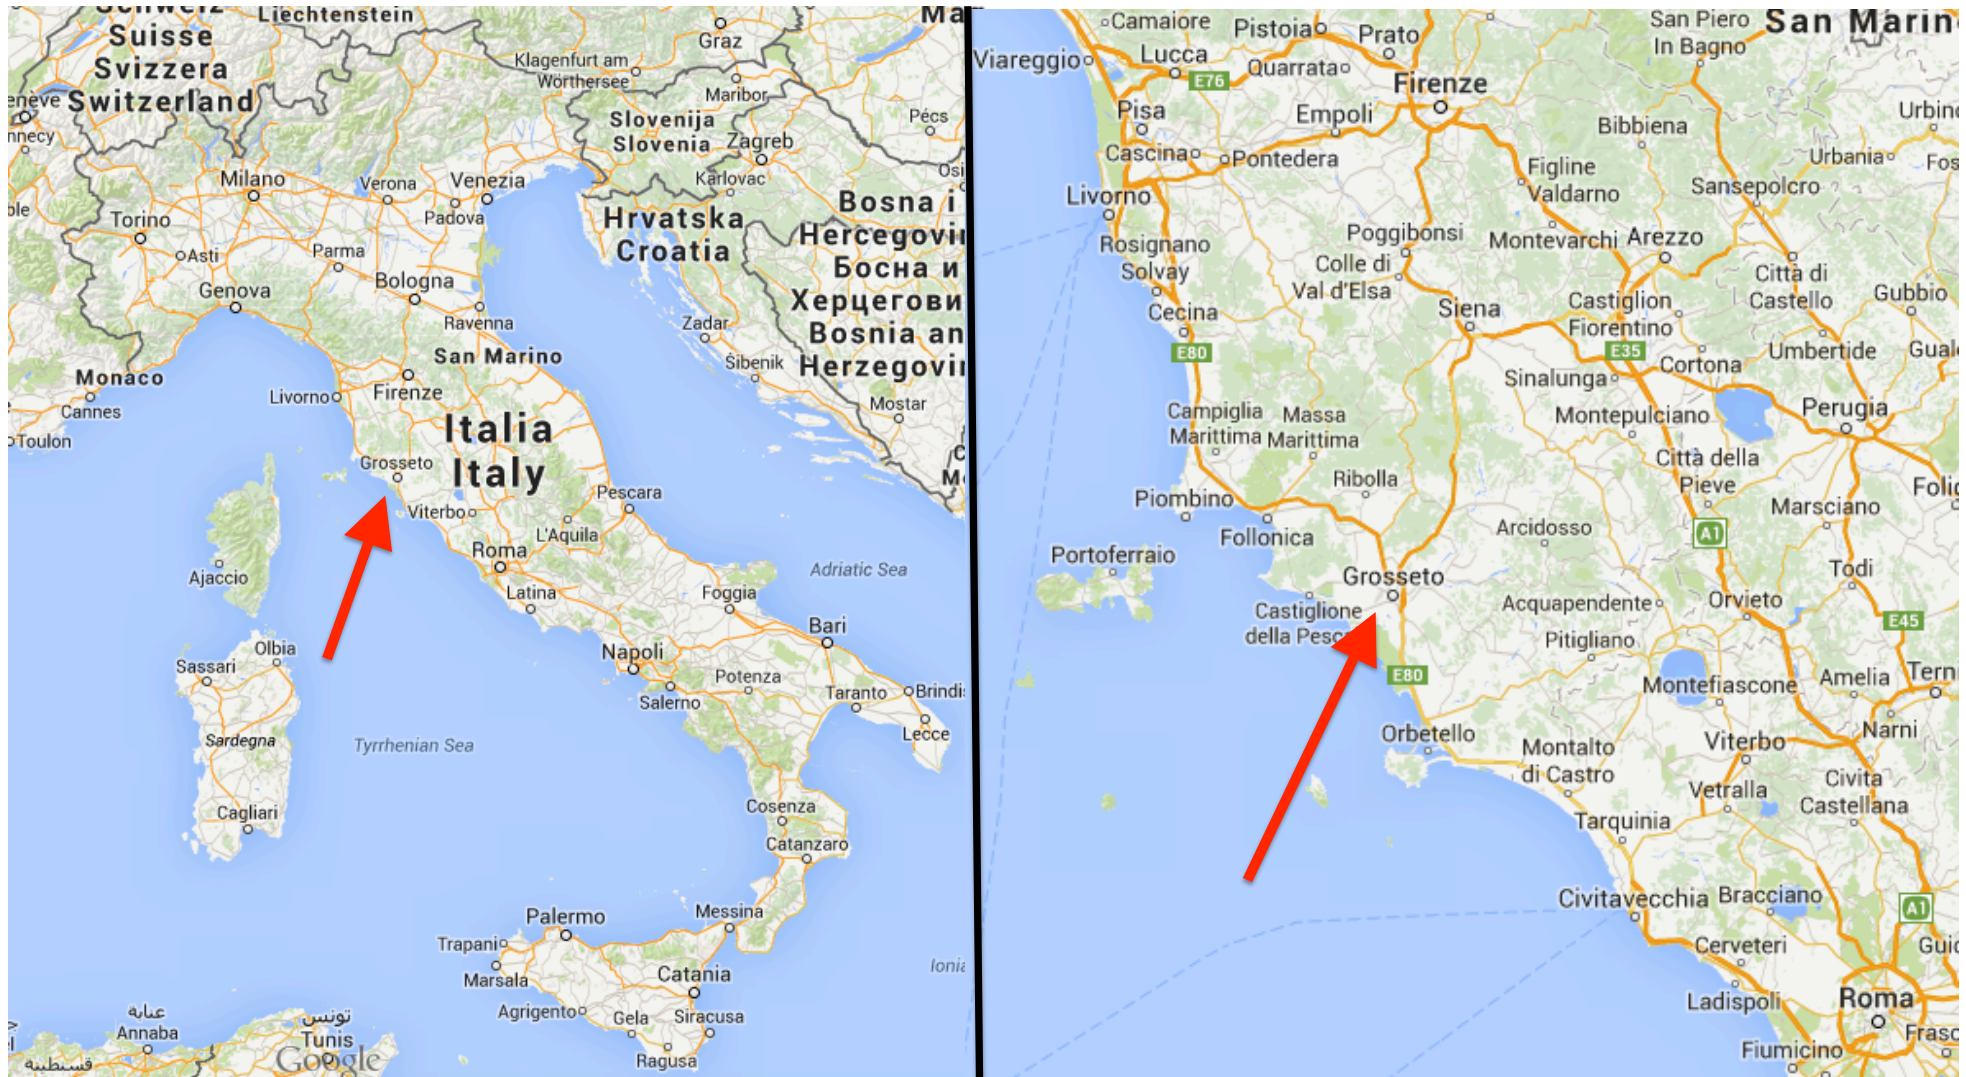
## Lorenzo Busatti

- Founder of Grifonline S.r.l., Grosseto – ITALY
- ISP since 1997, WISP since 2006
- MikroTik Certified Trainer since 2010 for:  
MTCNA, MTCWE, MTCRE, MTCTCE, MTCUME,  
MTCINF
- Member of RIPE, AMS-IX, MIX-IT

I'm a MikroTik *enthusiast*

# I'm a MikroTik *evangelist*

# About me



*Dedicated to Max*

# Presentation Objective


Raise awareness of the peering  
with the big players of the net,  
**using RouterOS.**



# What is the peering?

# Peering and BGP

Commonly used for having more than one ISP



# Wikipedia say:

“In computer networking, peering is a voluntary interconnection of administratively separate Internet networks for the purpose of exchanging traffic between the users of each network.”

<http://en.wikipedia.org/wiki/Peering>

# Wikipedia say:

And:

“An agreement by two or more networks to peer is instantiated by a physical interconnection of the networks, an exchange of routing information through the Border Gateway Protocol (BGP) routing protocol.”

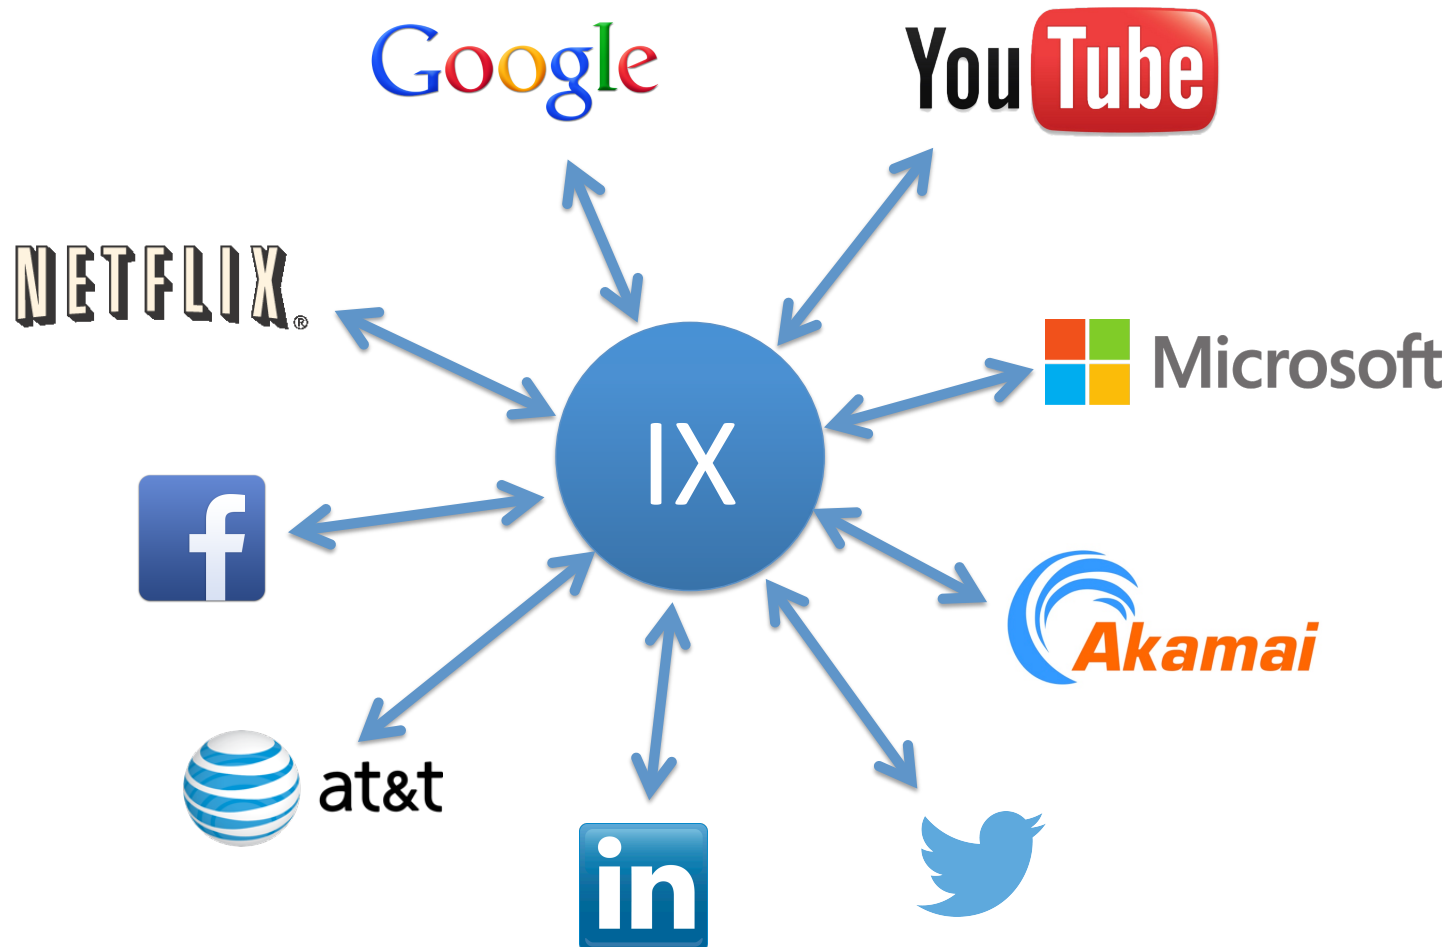
# Internet Exchange

An Internet exchange point (IX or IXP) is a physical infrastructure through which Internet service providers (ISPs) exchange Internet traffic between their networks (autonomous systems).

[http://en.wikipedia.org/wiki/Internet\\_exchange\\_point](http://en.wikipedia.org/wiki/Internet_exchange_point)

# Internet Exchange

Carriers, ISPs and Content Provider meet together

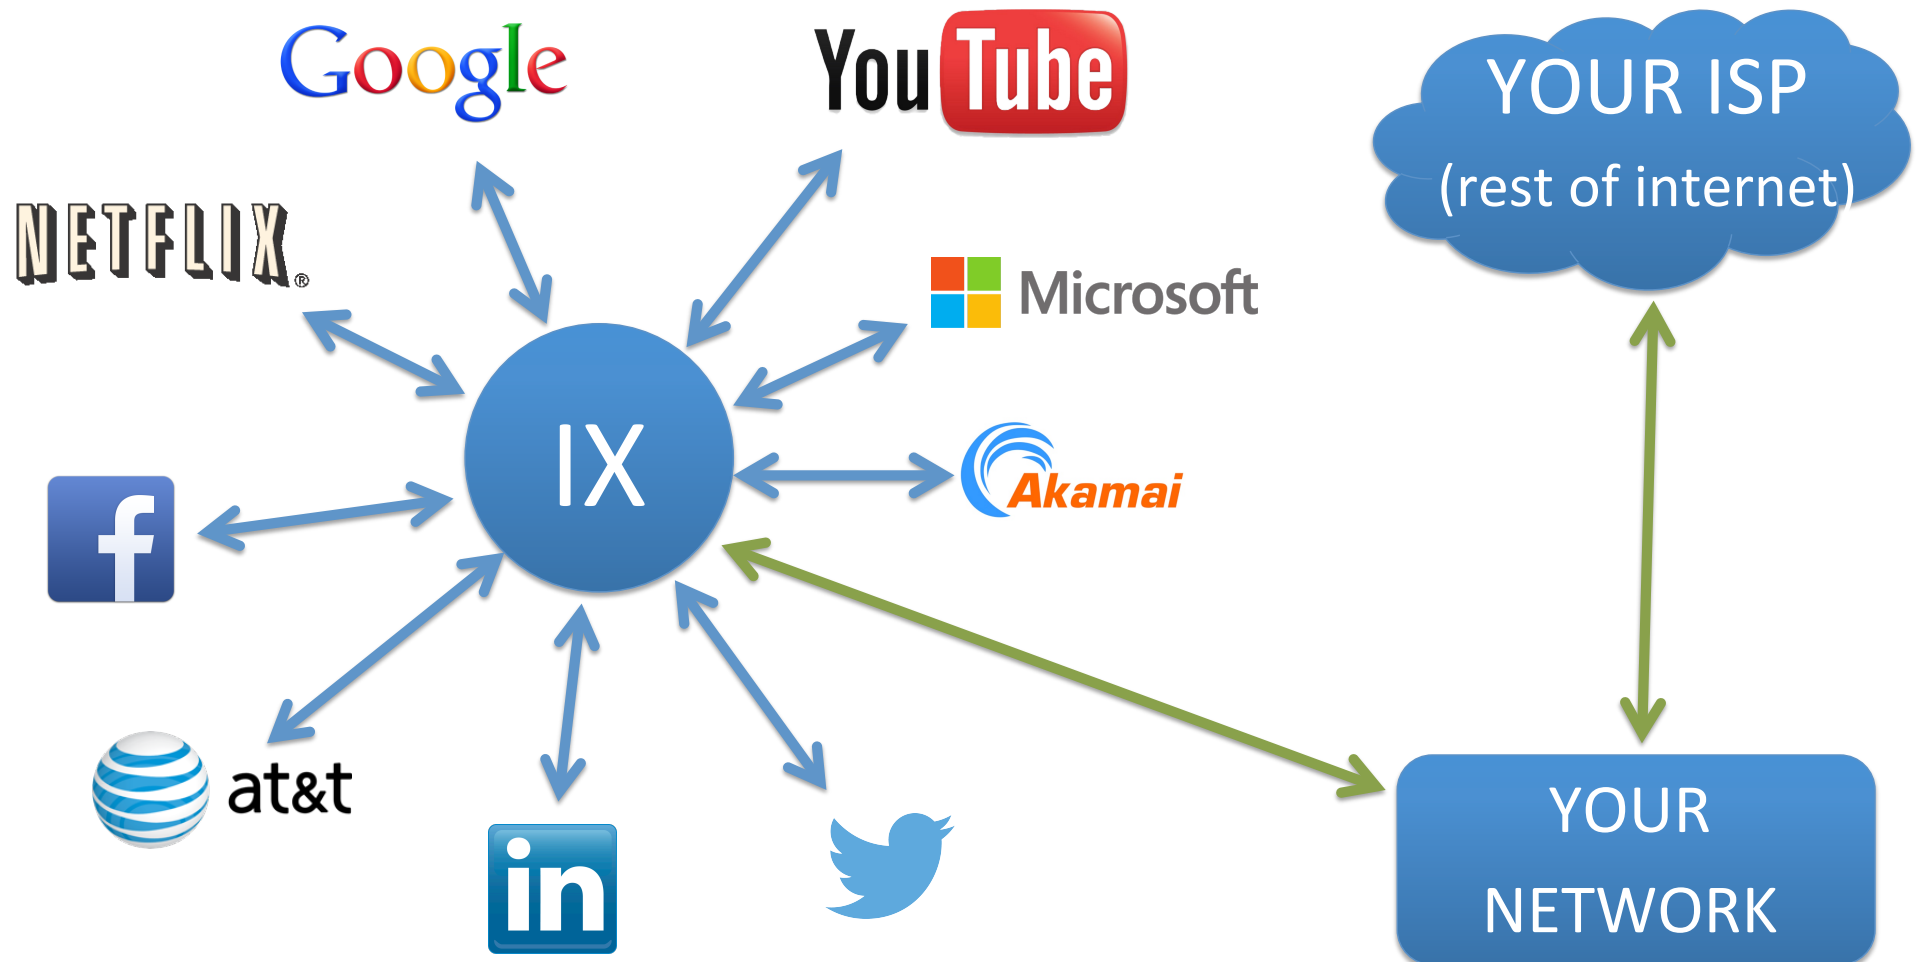


# Internet Exchange (Cont.)

IXPs reduce the portion of an ISP's traffic which must be delivered via their upstream transit providers, thereby reducing the average per-bit delivery cost of their service. Furthermore, the increased number of paths learned through the IXP improves routing efficiency and fault-tolerance.

# Internet Exchange (Cont.)

You can connect to an Internet Exchange





# What I need?

- ✓ A transport up to an IXP, usually a “pseudo-wire” (MPLS L2VPN)
- ✓ An IXP membership
- ✓ To be AS (Autonomous System)
- ✓ A router with BGP capabilities
- ✓ BGP knowledge

# BGP

- **B**order **G**ateway **P**rotocol, the protocol that route the Internet today.
- Designed for Inter-AS routing protocol
- Supported by RouterOS.

# BGP (cont.)

- For every “peer” (i.e.: Google, FB, .....) you need to setup a separate “BGP peer”

# BGP (cont.)

For BGP I'm using RouterOS since version 3.XX.

**I can state: it's working!**

# Public & Private Peering

Each entity that participates in a IXP define if:

- Participate in the "public peering" (Route Server) [many-to-many]
- Only do private “peering” [one-to-one]

# Public & Private Peering

Check the IXs member list:

| ↕ Organisation             | ↕ URL   | ↕ AS   | ↕ Policy     | ↕ Routerserver | ↕ Location          | ↕ Vlan |
|----------------------------|---|--------|--------------|----------------|---------------------|--------|
| GoDaddy.com, Inc           | ▼ <a href="http://www.godaddy.com/">http://www.godaddy.com/</a>             | 26496  | open         | Yes            | Telecity 5          | ISP    |
| Golden Telecom Limited     | ▼ <a href="http://www.goldentelecom.com/">http://www.goldentelecom.com/</a> | 3216   | case-by-case | No             | Equinix AM1/2       | ISP    |
| Google Ireland Limited     | ▼ <a href="http://www.google.com">http://www.google.com</a>                 | 15169  | case-by-case | Yes            | TeleCity 2, Vancis  | ISP    |
| GooMobile srl              | ▼ <a href="http://www.goomobile.it">http://www.goomobile.it</a>             | 198291 | open         | Yes            | Lepida Spa          | ISP    |
| Gopas Solutions GmbH (new) | ▼ <a href="http://www.gopas.de/">http://www.gopas.de/</a>                   | 13157  | open         | Yes            | Peering GmbH / ECIX | ISP    |
| Goscomb Technologies Ltd   | ▼ <a href="http://www.goscomb.net/">http://www.goscomb.net/</a>             | 39326  | open         | No             | NIKHEF              | ISP    |
| Green.ch AG                | ▼ <a href="http://www.green.ch">http://www.green.ch</a>                     | 1836   | open         | No             | Vancis              | ISP    |
| Greenhost                  | ▼ <a href="https://greenhost.nl/">https://greenhost.nl/</a>                 | 47172  | open         | Yes            | NL-IX               | ISP    |
| Grifonline S.r.l.          | ▼ <a href="http://www.linkwave.it">http://www.linkwave.it</a>               | 43942  | open         | Yes            | Vancis              | ISP    |
| GTT Communications, Inc.   | ▼ <a href="http://www.gtt.net/">http://www.gtt.net/</a>                     | 3257   | case-by-case | No             | TeleCity 2          | ISP    |

# Public & Private Peering

Check the IXs member list:

| ↕ Organisation              | ↕ URL   | ↕ AS   | ↕ Policy     | ↕ Routeserver |
|-----------------------------|---|--------|--------------|---------------|
| 1&1 Internet AG ▼           | <a href="http://www.1und1.de">http://www.1und1.de</a>                 | 8560   | open         | No            |
| 2AT B.V. ▼                  | <a href="http://www.2at.nl/">http://www.2at.nl/</a>                   | 197219 | case-by-case | Yes           |
| A2B Internet BV ▼           | <a href="http://www.a2b-internet.com">http://www.a2b-internet.com</a> | 51088  | open         | Yes           |
| AC Webconnecting BV ▼       | <a href="http://acwebconnecting.com">http://acwebconnecting.com</a>   | 47836  | open         | Yes           |
| Adamo Telecom Iberia S.A. ▼ | <a href="http://www.adamo.es/">http://www.adamo.es/</a>               | 35699  | open         | No            |
| Adeli ▼                     | <a href="http://www.adeli.fr/">http://www.adeli.fr/</a>               | 43142  | open         | Yes           |
| Adobe Systems ▼             | <a href="http://www.adobe.com">http://www.adobe.com</a>               | 15224  | case-by-case | No            |

# Public Peering

If you participate at the “Route Server” with policy “open” you’ll need just one peer and you’ll be automatically peering with the others in the “Route Server” with policy “open” .



# Private Peering

You can write to a member for request to peer, or you can receive a request by the others.

In most cases the peering is free, but this isn't a rule!

# Public & Private Peering

In the real life probably you'll have a public peering and some private peerings, they don't exclude each other.

# Traceroute

If you're just buying internet from an ISP, your customers should obtain something like this:

```
Lorenzos-MacBook-Pro:~ Lorenzo$ traceroute www.facebook.com
```

```
traceroute to star.cl0r.facebook.com (31.13.73.193), 64 hops max, 52 byte packets
```

```
1  192.168.176.1 (192.168.176.1)  0.969 ms  0.685 ms  0.775 ms

2  192.168.25.1 (192.168.25.1)  1.500 ms

   177.184.136.65 (177.184.136.65)  2.140 ms  2.137 ms

3  177.184.137.89 (177.184.137.89)  3.972 ms *  2.482 ms

4  * gvt-be-5-90.rd03.flg.gvt.net.br (179.184.83.81)  17.038 ms *

5  gvt-te-0-0-0-4.rc01.flg.gvt.net.br (179.185.131.57)  14.076 ms

   gvt-te-0-7-0-3.rc01.flg.gvt.net.br (189.59.251.157)  9.542 ms

   gvt-te-0-2-0-4.rc01.flg.gvt.net.br (179.185.131.61)  6.370 ms

6  186.211.241.1 (186.211.241.1)  88.708 ms  103.148 ms  114.605 ms

7  nota.br01.mia1.tfbnw.net (198.32.124.200)  89.564 ms  87.921 ms  103.844 ms

8  * * *

9  * * *

10 edge-star-shv-13-mia1.facebook.com (31.13.73.193)  91.461 ms  88.426 ms  90.890 ms
```

# Traceroute

.. And like this:

(Google.com  
15 hops)

```
traceroute to www.google.com.br (74.125.225.24), 64 hops max, 52 byte packets

 1  192.168.176.1 (192.168.176.1)  15.849 ms  1.824 ms  0.937 ms

 2  177.184.136.65 (177.184.136.65)  2.634 ms  5.032 ms  1.962 ms

 3  177.184.137.89 (177.184.137.89)  2.959 ms * *

 4  core001.plaz.wirelink.com.br.45.60.187.in-addr.arpa (187.60.45.1)  5.646 ms *  3.658 ms

 5  xe-0-3-1.csrl.fcsl.for.gblx.net (159.63.52.129)  18.874 ms  2.142 ms  2.037 ms

 6  po3-40g.asr1.grul.gblx.net (67.16.139.166)  419.514 ms

    po5-30g.asr1.grul.gblx.net (67.16.130.58)  326.170 ms

    gvt-te-0-6-0-11.rc01.spo.gvt.net.br (179.184.72.253)  48.695 ms

 7  google-1.ar5.grul.gblx.net (64.208.110.102)  51.403 ms  43.347 ms

    gvt-te-0-4-0-2.rt01.spo.gvt.net.br (187.115.215.58)  48.099 ms

 8  72.14.198.181 (72.14.198.181)  47.343 ms

    216.239.51.228 (216.239.51.228)  43.387 ms

    72.14.198.181 (72.14.198.181)  49.158 ms

 9  209.85.245.53 (209.85.245.53)  58.193 ms  76.383 ms  52.908 ms

10  209.85.245.54 (209.85.245.54)  164.945 ms  161.104 ms  167.806 ms

11  209.85.245.54 (209.85.245.54)  155.576 ms

    64.233.175.206 (64.233.175.206)  178.308 ms

    209.85.252.96 (209.85.252.96)  174.168 ms

12  72.14.239.91 (72.14.239.91)  195.022 ms

    72.14.233.196 (72.14.233.196)  196.196 ms *

13  72.14.232.163 (72.14.232.163)  197.179 ms

    209.85.254.239 (209.85.254.239)  197.063 ms

    72.14.232.163 (72.14.232.163)  195.611 ms

14  72.14.237.109 (72.14.237.109)  197.291 ms  196.044 ms

    72.14.232.163 (72.14.232.163)  194.559 ms

15  72.14.237.109 (72.14.237.109)  189.895 ms

    ord08s12-in-f24.1e100.net (74.125.225.24)  196.617 ms  251.374 ms
```

# Traceroute

With direct peering **your customer** should obtain something like this:

traceroute to [www.netflix.com](http://www.netflix.com)

- 1 The ip address of your customer router
- 2 The ip address of your PPPoE server or something else
- 3 The ip address of your BGP router
- 4 The ip address of the destination BGP router
- 5 The ip address of the destination server.

# Traceroute

Your customers will start to think you're a real ISP, connected directly to the Big internet players and you're not just "reselling" other ISPs.

# The Internet Exchange

# IXP in the world

Many Internet Exchange in the continents. The largest ones in the world are DE-CIX in Frankfurt, AMS-IX in Amsterdam, LINX in London, the Moscow Internet Exchange, Equinix Ashburn in Washington D.C., and JPNAP in Tokyo, and so on.

[http://en.wikipedia.org/wiki/List\\_of\\_Internet\\_exchange\\_points](http://en.wikipedia.org/wiki/List_of_Internet_exchange_points)

[http://en.wikipedia.org/wiki/List\\_of\\_Internet\\_exchange\\_points\\_by\\_size](http://en.wikipedia.org/wiki/List_of_Internet_exchange_points_by_size)






# DE-CIX, Frankfurt

1<sup>st</sup> in the world, for traffic exchange.

The last 2 days traffic average:






# DE-CIX, Frankfurt

1<sup>st</sup> in the world, for traffic exchange.

The last 5 years traffic average:






# AMS-IX, Amsterdam

2<sup>nd</sup> in the world, for traffic exchange.

The last 2 days traffic average:






# AMS-IX, Amsterdam

2<sup>nd</sup> in the world, for traffic exchange.

The last year traffic average:





# AMS-IX, Amsterdam

Some numbers:

ASNs

**711**

Ports

**1404**

Peak (Tb/s)

**3.52**


Cur (Tb/s)

**1.724**

Capacity (Tb/s)

**14.4314**

Peering around the globe





# AMS-IX, Amsterdam

2014, November:

ASNs

**685**

Ports

**1324**

Peak (Tb/s)

**3.339**

Cur (Tb/s)

**3.142**

Capacity (Tb/s)

**12.907**

2015, March:

ASNs

**711**

Ports

**1404**

Peak (Tb/s)

**3.52**

Cur (Tb/s)


**1.724**

Capacity (Tb/s)

**14.4314**


# MIX-IT, Milan

The biggest in Italy, yesterday:



# MIX-IT, Milan

The biggest in Italy, one year average:





# Why you should know that?

# Things you should know:

- ✓ Internet isn't only your main ISP
- ✓ You can be part of the internet, not just a customer.
- ✓ You can use RouterOS for these tasks!

That's why I'm telling you my experiences.

# RouterOS at the IXPs

# What is RouterOS?

Some years ago I had a training at RIPE in Amsterdam and there I meet an AMS-IX technician, taking this training with me. He discovered I'm a MikroTik Trainer and asked me some informations. He know that RouterOS existed but he never used one. I showed him some RouterOS functionalities.

# What is RouterOS?

I said him that I'll use RouterOS for peering at the AMS-IX. He replied me I'll be the first one.

In every IXP are used **only Cisco, Juniper and Huawei**.

In the IXPs websites exist setup guides only for Cisco and Juniper.

# What is RouterOS?

Some months later I started to peer at AMS-IX using RouterOS. In few days I had dozens of running peers!

This technician was very impressed by RouterOS, I never called for an help (and they're used to provide it).

# Epilogue

Years later this technician meet my colleague and friend **Tom**, an Irish Trainer. Tom started to talk about RouterOS. The technician stopped him and said: I know a very good technician that is using RouterOS at AMS-IX with success, if you need help I can recommend him. (me!!!)

# Epilogue

The world is really small and RouterOS is now famous in the bigger Internet Exchange of the world! :-)



# Epilogue

Lots of Cisco stuff ....

.....

.....

## 12. Mikrotik Configuration Hints

By default Mikrotik routers have their own proprietary Mikrotik Discovery Protocol and CDP enabled. To turn these discovery protocols off, in the Web UI go to *IP > Neighbors > Discovery Interfaces* and disable the protocols on the AMS-IX-facing interface.

.....

Lots of Juniper stuff....

| Name   | Instance | Remote Address | Remote AS  | TTL | Remote ID | Prefix Co... | State       |
|--|----------|----------------|------------|-----|-----------|--------------|-------------|
| MIX RouteServer1 IPv4                        |          |                | 61968 rr 1 |     |           | 1983         | established |
| MIX RouteServer1 IPv6                        |          |                | 61968 rr 1 |     |           | 125          | established |
| MIX RouteServer2 IPv4                        |          |                | 61968 rr 1 |     |           | 4588         | established |
| MIX RouteServer2 IPv6                        |          |                | 61968 rr 1 |     |           | 345          | established |
| MIX ServerTest IPv4                          |          |                | 16004 rr 1 |     |           |              | established |
| MIX ServerTest IPv6                          |          |                | 16004 rr 1 |     |           | 1            | established |
| retelit                                      |          |                | 28716 rr 1 |     |           | 508499       | established |
| retelit IPv6                                 |          |                | 28716 rr 1 |     |           | 19791        | established |
| zzz AMS-IX: Akamai 1 IPv4 [NEW]              |          |                | 20940 rr 1 |     |           | 17           | established |
| zzz AMS-IX: Akamai 1 IPv6                    |          |                | 20940 rr 1 |     |           | 2            | established |
| zzz AMS-IX: Akamai 2 IPv4 [NEW]              |          |                | 20940 rr 1 |     |           | 23           | established |
| zzz AMS-IX: Akamai 2 IPv6                    |          |                | 20940 rr 1 |     |           | 2            | established |
| zzz AMS-IX: BIT IPv4 [NEW]                   |          |                | 12859 rr 1 |     |           | 65           | established |
| zzz AMS-IX: BIT IPv6                         |          |                | 12859 rr 1 |     |           | 29           | established |
| zzz AMS-IX: Cablesystems IPv4 [NEW]          |          |                | 6327 rr 1  |     |           | 1664         | established |
| zzz AMS-IX: Cablesystems IPv6                |          |                | 6327 rr 1  |     |           |              | idle        |
| zzz AMS-IX: Computerline IPv4 [NEW]          |          |                | 41913 rr 1 |     |           | 17           | established |
| zzz AMS-IX: Computerline IPv6                |          |                | 41913 rr 1 |     |           | 4            | established |
| zzz AMS-IX: Facebook 1 IPv4 [NEW]            |          |                | 32934 rr 1 |     |           | 22           | established |
| zzz AMS-IX: Facebook 1 IPv6                  |          |                | 32934 rr 1 |     |           | 8            | established |
| zzz AMS-IX: Facebook 2 IPv4 [NEW]            |          |                | 32934 rr 1 |     |           | 22           | established |
| zzz AMS-IX: Facebook 2 IPv6                  |          |                | 32934 rr 1 |     |           | 8            | established |
| zzz AMS-IX: GoDaddy IPv4 [NEW]               |          |                | 26496 rr 1 |     |           | 595          | established |
| zzz AMS-IX: GoDaddy IPv6                     |          |                | 26496 rr 1 |     |           | 8            | established |
| zzz AMS-IX: HE IPv4 [NEW]                    |          |                | 6939 rr 1  |     |           | 59759        | established |
| zzz AMS-IX: HE IPv6                          |          |                | 6939 rr 1  |     |           | 10380        | established |
| zzz AMS-IX: LinkedIn IPv4 [NEW]              |          |                | 14413 rr 1 |     |           | 4            | established |
| zzz AMS-IX: LinkedIn IPv6                    |          |                | 14413 rr 1 |     |           | 1            | established |
| X zzz AMS-IX: MANDA IPv4                     |          |                | 8365 rr 1  |     |           |              | idle        |
| zzz AMS-IX: MANDA IPv6                       |          |                | 8365 rr 1  |     |           | 13           | established |
| zzz AMS-IX: Microsoft 1 IPv4 [NEW]           |          |                | 8075 rr 1  |     |           | 165          | established |
| zzz AMS-IX: Microsoft 1 IPv6                 |          |                | 8075 rr 1  |     |           | 16           | established |
| zzz AMS-IX: Microsoft 2 IPv4 [NEW]           |          |                | 8075 rr 1  |     |           | 165          | established |
| zzz AMS-IX: Microsoft 2 IPv6                 |          |                | 8075 rr 1  |     |           | 16           | established |
| zzz AMS-IX: NLhosting IPv4 [NEW]             |          |                | 24793 rr 1 |     |           | 1            | established |
| zzz AMS-IX: Netflix IPv4 [NEW]               |          |                | 2906 rr 1  |     |           | 2            | established |
| zzz AMS-IX: Netflix IPv4 [NEW]               |          |                | 2906 rr 1  |     |           | 2            | established |
| zzz AMS-IX: Netflix IPv6                     |          |                | 2906 rr 1  |     |           | 2            | established |
| zzz AMS-IX: Netflix IPv6                     |          |                | 2906 rr 1  |     |           | 2            | established |
| zzz AMS-IX: Open Carrier IPv4 [NEW]          |          |                | 41692 rr 1 |     |           | 63           | established |
| zzz AMS-IX: Open Carrier IPv6                |          |                | 41692 rr 1 |     |           | 20           | established |
| zzz AMS-IX: Packet Clearing House 1 IPv4 ... |          |                | 42 rr 1    |     |           | 48           | established |
| zzz AMS-IX: Packet Clearing House 1 IPv6     |          |                | 42 rr 1    |     |           | 42           | established |
| zzz AMS-IX: Packet Clearing House 2 IPv4 ... |          |                | 3856 rr 1  |     |           | 1            | established |

55 items (1 selected)

# RouterOS

My RouterOS BGP router:

```
[admin@BGP] > ip route print count-only  
1088614
```

# Resources

# Resources

For fast BGP knowledge these two MikroTik trainings will help you into this world:

**MTCRE:** MikroTik Certified Routing Engineer

**MTCINE:** MikroTik Certified Inter-  
Networking Engineer

# Resources

For self-study and for see lots of examples and some very interesting case studies use this beautiful presentation made by my colleague Wardner Maia:

<http://mum.mikrotik.com/presentations/HR13/maia.pdf>

# Resources

For BGP security tips you should be here yesterday and follow the beautiful presentation made by my colleague Tom Smith

# Tips & Tricks



# Security

When planning a peer you should consider the use of the MD5 password.

The MAC Address on the interfaces you're using verso the IXP should not change: their router doesn't tolerate this action and will close all the connections.

# Filtering

Use the filters for:

- Avoid to announce private addresses;
- Avoid to receive private addresses;
- Avoid to receive your own addresses;
- Avoid to be used as a “transit”
- Manipulate upload and download preferences

# Neighbor Discovery Protocol

**Disable** the MikroTik Neighbor Discovery protocol (MNDP) on the interfaces you're using verso the IXP: their router doesn't tolerate such broadcast traffic, they will treat you like a threat.

# PeeringDB

PeeringDB is "a freely available web-based database of networks that are interested in peering", which has become the default location for Internet peering data.

[www.peeringdb.com](http://www.peeringdb.com)

# PeeringDB

With PeeringDB you can know where an AS is connected, at which speed, addresses, notes, policies, contacts (for peering purpose only!!), etc.

| Company Information        |  |              |                       |
|----------------------------|--|--------------|-----------------------|
| Company Name               | Yahoo!   |              |                       |
| Also Known As              | Favorite whipping post of the tech media   |              |                       |
| Company Website            | <a href="http://www.yahoo.com/">http://www.yahoo.com/</a>  |              |                       |
| Primary ASN                | 10310  |              |                       |
| IRR Record                 | AS-YAHOO   |              |                       |
| Network Type               | Content  |              |                       |
| Approx Prefixes            | 500  |              |                       |
| Traffic Levels             | 200-300 Gbps   |              |                       |
| Traffic Ratios             | Heavy Outbound   |              |                       |
| Geographic Scope           | Global   |              |                       |
| Looking Glass URL          |  |              |                       |
| Route Server URL           |  |              |                       |
| Notes                      | We require sessions to be built to both IPs for all dual-attached public exchange points. Peers that only establish adjacencies with only one router may be eliminated for non-compliance. |              |                       |
| Protocols Supported        | Unicast IPv4 <input checked="" type="checkbox"/> Multicast <input type="checkbox"/> IPv6 <input checked="" type="checkbox"/>   |              |                       |
| Date Last Updated          | 2014-10-27 15:22:26 UTC  |              |                       |
| Peering Policy Information |  |              |                       |
| Peering Policy URL         |  |              |                       |
| General Policy             | Selective  |              |                       |
| Multiple Locations         | Not Required   |              |                       |
| Ratio Requirement          | No   |              |                       |
| Contract Requirement       | Not Required   |              |                       |
| Contact Information        |  |              |                       |
| Role                       | Contact Name   | Telephone    | E-Mail                |
| Technical                  | Igor Gashinsky   | 917-807-2213 | igor@yahoo-inc.com    |
| Technical                  | Matthew Petach   | 408-349-7231 | mpetach@yahoo-inc.com |
| Policy                     | Peering Committee  |              | peering@yahoo-inc.com |
| Technical                  | NOC  | 866-GET-YNOC | ynoc@yahoo-inc.com    |
| NOC                        | NOC  | 408-349-5555 | ynoc@yahoo-inc.com    |

| Public Peering Exchange Points             |       |                           |          |                          |                                     |                          |
|--|-------|---------------------------|----------|--------------------------|-------------------------------------|--------------------------|
| Exchange Point Name                        | ASN   | IP Address                | Mbit/sec |                          |                                     |                          |
| <a href="#">AMS-IX</a>                     | 10310 | 2001:7f8:1::A501:310:1/64 | 20000    |                          |                                     |                          |
| <a href="#">AMS-IX</a>                     | 10310 | 80.249.209.110            | 20000    |                          |                                     |                          |
| <a href="#">AMS-IX</a>                     | 10310 | 80.249.209.163            | 20000    |                          |                                     |                          |
| <a href="#">AMS-IX</a>                     | 10310 | 2001:7f8:1::A501:310:2/64 | 20000    |                          |                                     |                          |
| <a href="#">BNIIX</a>                      | 10310 | 206.130.61.11             | 1000     |                          |                                     |                          |
| <a href="#">BNIIX</a>                      | 10310 | 2001:504:20:81::11/64     | 1000     |                          |                                     |                          |
| <a href="#">CoreSite - Any2 California</a> | 10310 | 2001:504:13:0:0:0:195/64  | 10000    |                          |                                     |                          |
| <a href="#">CoreSite - Any2 California</a> | 10310 | 206.72.210.195/23         | 10000    |                          |                                     |                          |
| <a href="#">CoreSite - Any2 Denver</a>     | 10310 | 2605:6c00:303:303::25/64  | 10000    |                          |                                     |                          |
| <a href="#">CoreSite - Any2 Denver</a>     | 10310 | 206.51.46.25              | 10000    |                          |                                     |                          |
| <a href="#">DE-CIX Frankfurt</a>           | 10310 | 80.81.193.115             | 10000    |                          |                                     |                          |
| <a href="#">DE-CIX Frankfurt</a>           | 10310 | 2001:7f8::2846:0:1/64     | 10000    |                          |                                     |                          |
| 1 2 3 4 5 of 8 Next > Last >>              |       |                           |          |                          |                                     |                          |
| Private Peering Facilities                 |       |                           |          |                          |                                     |                          |
| Facility Name                              | ASN   | City                      | Country  | SONET                    | Ethr                                | ATM                      |
| <a href="#">CoreSite - DE1</a>             | 10310 | Denver                    | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix Ashburn (DC1-DC11)</a> | 10310 | Ashburn                   | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix Chicago (CH1/CH2)</a>  | 10310 | Chicago                   | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix Dallas (DA1)</a>       | 10310 | Dallas                    | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix Los Angeles (LA1)</a>  | 10310 | Los Angeles               | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix New York (111 8th)</a> | 10310 | New York                  | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix Palo Alto (SV8)</a>    | 10310 | Palo Alto                 | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix San Jose (SV1/5)</a>   | 10310 | San Jose                  | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix Seattle (SE2/3)</a>    | 10310 | Seattle                   | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix Sydney</a>             | 17457 | Mascot (Sydney) NSW       | AU       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix Tokyo (TY1)</a>        | 10310 | Tokyo                     | JP       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| <a href="#">Equinix Vienna, VA (DC7)</a>   | 10310 | Vienna                    | US       | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| 1 2 of 2 Next > Last >>                    |       |                           |          |                          |                                     |                          |

# BGP Community

The BGP Community will let you manage “things” of your peering partner without calling him. You can manage these “things” yourself.

And they’re very useful in the real life.

# BGP Community

You can know the available BGP Community settings from your peer, or querying public DB like RIPE or anothers.

For example let's have a look at Cogent, AS174 . Just for illustration purpose!



# AS174 Community

| BGP Community String | Local Pref | Effect  |
|----------------------|------------|---|
| 174:10               | 10         | Set customer route local preference to 10<br>(below everything-least preferred) |
| 174:70               | 70         | Set customer route local preference to 70<br>(below peers)                      |
| 174:120              | 120        | Set customer route local preference to 120<br>(below customer default)          |
| 174:125              | 125        | Set customer route local preference to 125<br>(below customer default)          |
| 174:135              | 135        | Set customer route local preference to 135<br>(above customer default)          |
| 174:140              | 140        | Set customer route local preference to 140<br>(above customer default)          |

# AS174 Community

| BGP Community String | Effect   |
|----------------------|--|
| 174:970              | Do <u>not send</u> route to NA ( <u>North America</u> ) - not accepted in NA |
| 174:971              | Do not send route to NA peers.   |
| 174:975              | Set local preference to 10 in NA.  |
| 174:980              | Do not send route to EU (Europe) - not accepted in EU                        |
| 174:981              | Do not send route to EU peers.   |
| 174:985              | Set local preference to 10 in EU.  |
| 174:990              | Do not send route to BGP customers, or peers.                                |
| 174:991              | Do not send route to peers.  |

# AS174 Community

| BGP Community String | Effect                   |
|----------------------|--------------------------|
| 174:3000             | Do not <u>announce</u> . |
| 174:3001             | Prepend 174 1 time.      |
| 174:3002             | Prepend 174 2 times.     |
| 174:3003             | Prepend 174 3 times      |

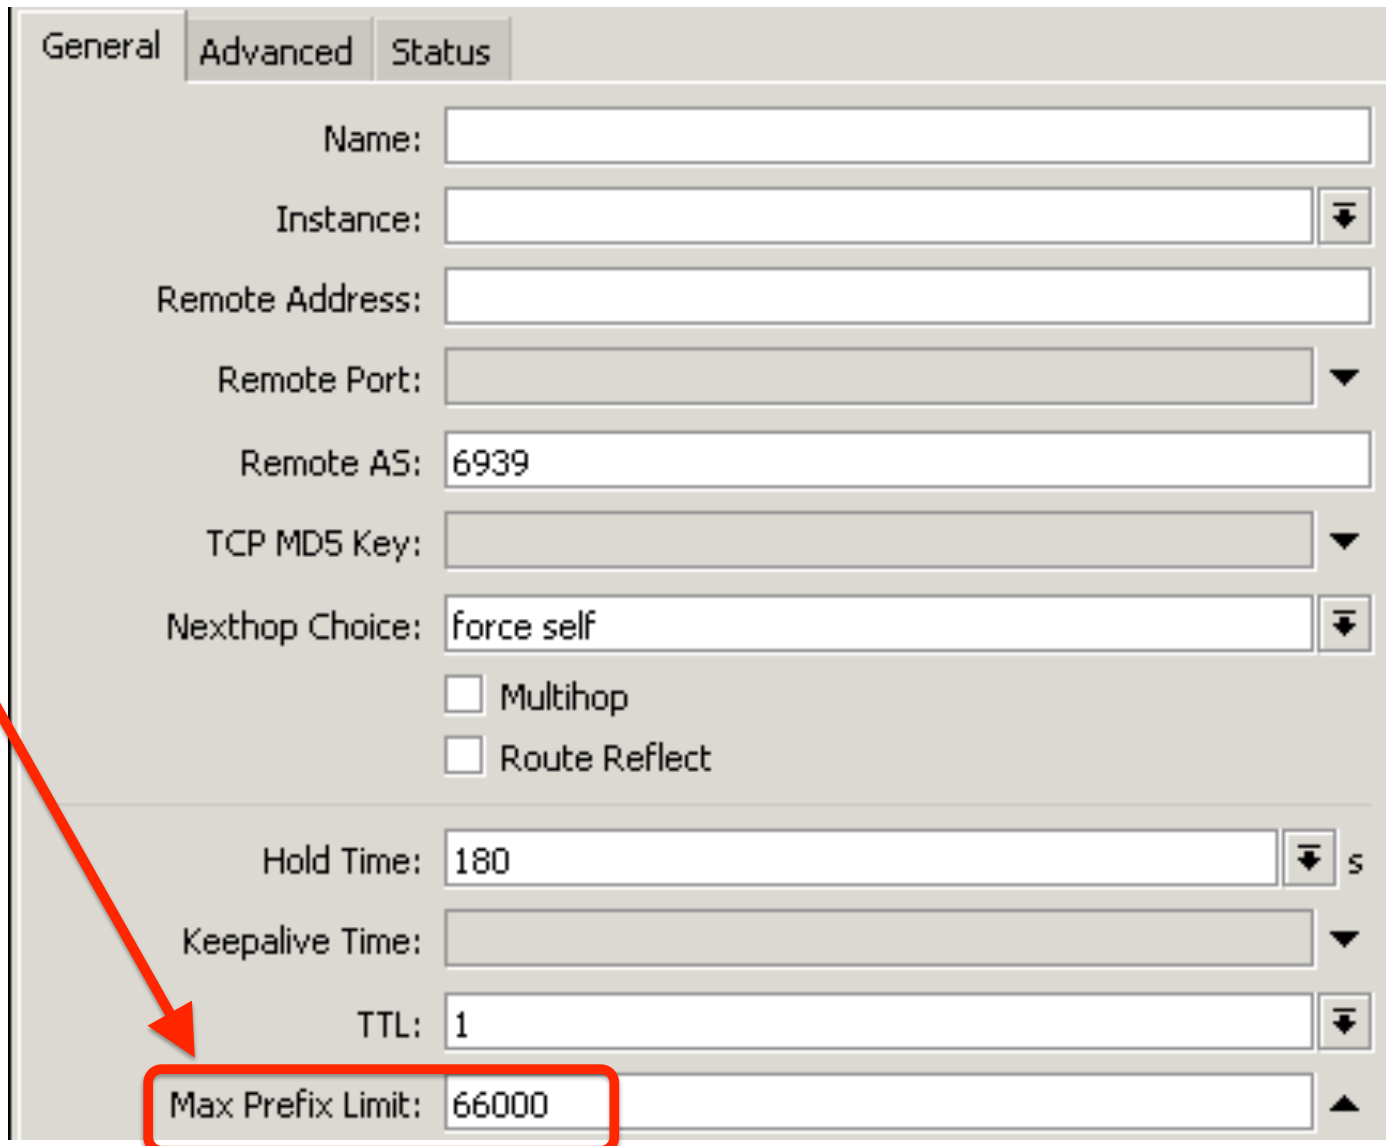
| BGP Community String | Description  |
|----------------------|--|
| 174:21000            | Route is learned from NA (North America) non-customer. |
| 174:21001            | Route is NA internal or customer route.                |
| 174:21100            | Route is learned from EU (Europe) non-customer.        |
| 174:21101            | Route is a EU internal or customer route.              |

# Prefix limit

For every peer should be a very good idea to set up a prefix limit. It should be limit the damage from receiving too many routes, i.e. in case of some “leaks”.


When according for a peer this information should be exchanged!

# Prefix limit




General Advanced Status


Name:


Instance:  

Remote Address:

Remote Port:  


Remote AS:


TCP MD5 Key:  


Nexthop Choice:  


☐ Multihop

☐ Route Reflect

Hold Time:   s

Keepalive Time:  

TTL:  

**Max Prefix Limit:**  

# Shit Happen

**Subject: ASXXXXX - IPv4 bogon prefixes leak**

Hi AMS-IX peers,

I would like to apologise because between 12:11PM and 12:13PM (London time) we accidentally leaked about 3k IPv4 bogon prefixes.

That was caused by a human mistake and it was immediately corrected.

We took necessary actions to make sure it will not happen again in the future.

If you peer with us at AMS-IX could you please reset the peering session(s), if it is still down:

ASN: XXXXX

IP: XX.XX.XX.XX

Please accept our apology for problems the issue may caused.

# Peer Status

Check the uptime, the prefix count (maybe you'll update the "prefix limit"), and the updates that happened.


| General              | Advanced | Status   |
|----------------------|----------|--|
| Remote ID:           |          |  |
| Local Address:       |          |  |
| Uptime:              |          | 7d 05:52:51  |
| Prefix Count:        |          | 59761  |
| Updates Sent:        |          | 3  |
| Updates Received:    |          | 185 650  |
| Withdrawn Sent:      |          |  |
| Withdrawn Received:  |          | 90 983   |
| Remote Hold Time:    |          | 180 s  |
| Used Hold Time:      |          | 180 s  |
| Used Keepalive Time: |          | 60 s   |
|                      |          | <input checked="" type="checkbox"/> Refresh Capability |
|                      |          | <input checked="" type="checkbox"/> AS4 Capability     |

# Winbox

Do you want to check your routing table using Winbox on your BGP router?



# Winbox



Don't click here!!!

# Winbox

You're surely don't want your router will upload to your winbox client all the "thousand of thousands" routes.

The CLI (Command Line Interface) will do the task better!

# Winbox

```
ip route print count-only
```

```
routing bgp advertisements print <peer_name>
```

```
ip route print detail where <ip_addresses> in  
dst-address and bgp
```

```
ip route print where received-  
from=<peer_name>
```

**[admin@BGP] > ip route print detail where 90.189.192.1 in dst-address and bgp**

Flags: X - disabled, A - active, D - dynamic,

C - connect, S - static, r - rip, b - bgp, o - ospf, m - mme,

B - blackhole, U - unreachable, P - prohibit

- 0 ADb dst-address=90.189.128.0/17 gateway=xxx.yyy.zzz.kkk  
gateway-status=xxx.yyy.zzz.kkk reachable via ether5 distance=20  
scope=40 target-scope=10 bgp-as-path="8928,12389,41440" bgp-med=10  
bgp-communities=8928:10403,8928:10901,8928:10902,  
8928:11003,8928:65103,8928:65152,8928:65193,8928:65194,  
41440:1501  
received-from=Peer1
  
- 1 Db dst-address=90.189.128.0/17 gateway=xxx.yyy.zzz.kkk  
gateway-status=xxx.yyy.zzz.kkk reachable via ether1 distance=20  
scope=40 target-scope=10 bgp-as-path="28716,1299,12389,41440"  
received-from=Peer2
  
- 2 ADb dst-address=90.189.192.0/23 gateway=xxx.yyy.zzz.kkk  
gateway-status=xxx.yyy.zzz.kkk reachable via ether5 distance=20  
scope=40 target-scope=10 bgp-as-path="8928,12389" bgp-med=10  
bgp-communities=8928:10403,8928:10901,8928:10902,  
8928:11003,8928:65103,8928:65152,8928:65193,8928:65194  
received-from=Peer1
  
- 3 Db dst-address=90.189.192.0/23 gateway=xxx.yyy.zzz.kkk  
gateway-status=xxx.yyy.zzz.kkk reachable via ether3 distance=20  
scope=40 target-scope=10 bgp-as-path="6939,12389" bgp-med=1  
received-from=Peer3
  
- 4 Db dst-address=90.189.192.0/23 gateway=xxx.yyy.zzz.kkk  
gateway-status=xxx.yyy.zzz.kkk reachable via ether1 distance=20  
scope=40 target-scope=10 bgp-as-path="28716,6939,12389"  
received-from=Peer2

# Wrap up

- ✓ Expand your network and **become part of Internet.**
- ✓ **Use RouterOS for BGP**, it's working very well.
- ✓ Improve your **BGP knowledge.**
- ✓ And quoting my friend Tom Smith:  
*“Buy bigger routers, bitches!”*



# Thank you!

## Q & A

<http://training.grifonline.it>

[training@grifonline.it](mailto:training@grifonline.it)