About me

- **Farzad Heydari Goojani**
  - Training, Support & Consultant
  - Over 5000 hours teaching networks
  - 15 years experience in IT
  - Achieved many international certificates (MTCNA,MTCUME,CCNA,CCNP,CCIE written,LPIC1,ITIL,HP-UX)

MikroTik User Meeting

Heydari.Farzad@gmail.com

Agenda

- **BGP Introduction**
- **eBGP peering on RouterOS**
- **Issues When Redundancy Exists Between eBGP Neighbors**
- **BGP peering with loopback interface**

Heydari.Farzad@gmail.com

BGP

Border Gateway Protocol

Heydari.Farzad@gmail.com

Public IP Address Assignment

- **Step 1**
- ICANN and IANA group public IPv4 addresses by major geographic region.
- **Step 2**
- IANA allocates those address ranges to Regional Internet Registries (RIR).

Heydari.Farzad@gmail.com

Public IP Address Assignment

Regional Internet Registries (RIR)

Heydari.Farzad@gmail.com

MikroTik User Meeting

Public IP Address Assignment

Regional Internet Registries (RIR)

Heydari.Farzad@gmail.com

MikroTik User Meeting

Public IP Address Assignment

Regional Internet Registries (RIR)

Heydari.Farzad@gmail.com

Public IP Address Assignment

Regional Internet Registries (RIR)

Heydari.Farzad@gmail.com

MikroTik User Meeting

Public IP Address Assignment

Regional Internet Registries (RIR)

Heydari.Farzad@gmail.com

MikroTik User Meeting

Introduction to BGP

Differences between BGP & IGPS

- BGP does not require neighbors to be attached to the same subnet.

  BGP routers use a TCP connection (port 179) between the routers to pass BGP messages.

- Instead of choosing the best route just by using an integer metric, BGP uses a more complex process, using a variety of information, called *BGP path attributes*, which are exchanged in BGP routing updates much like IGP metric information.

Heydari.Farzad@gmail.com

Introduction to BGP

AS Numbers

The integer *BGP AS Numbers* uniquely identifies one organization that considers itself autonomous from other organizations.

Each company network connects to the Internet can be considered to be an autonomous system and can be assigned a *BGP ASN*.

(IANA/ICANN also assigns globally unique ASNs.)
Additionally, each ISP has an ASN, or possibly several, depending on the size of the ISP.

Heydari.Farzad@gmail.com

Introduction to BGP

BGP ASNs and the AS SEQ Path Attribute

Note

By default, if no BGP PAs have been explicitly set, BGP routers use the BGP AS_PATH (autonomous system path) PA when choosing the best route among many competing routes.

Heydari.Farzad@gmail.com

Introduction to BGP

AS Path

When a router uses BGP to advertise a route, the prefix/length is associated with a set of PAs, including the AS_Path.

The AS_Path says:

"*If you use this path (route), the path will go through this list of ASNs*."

Heydari.Farzad@gmail.com

Introduction to BGP

AS Path

BGP uses the AS_Path to:

- Choose the best route for a prefix based on the shortest AS_Path (fewest number of ASNs listed).

- Prevent routing loops.

Heydari.Farzad@gmail.com

Introduction to BGP

AS Path

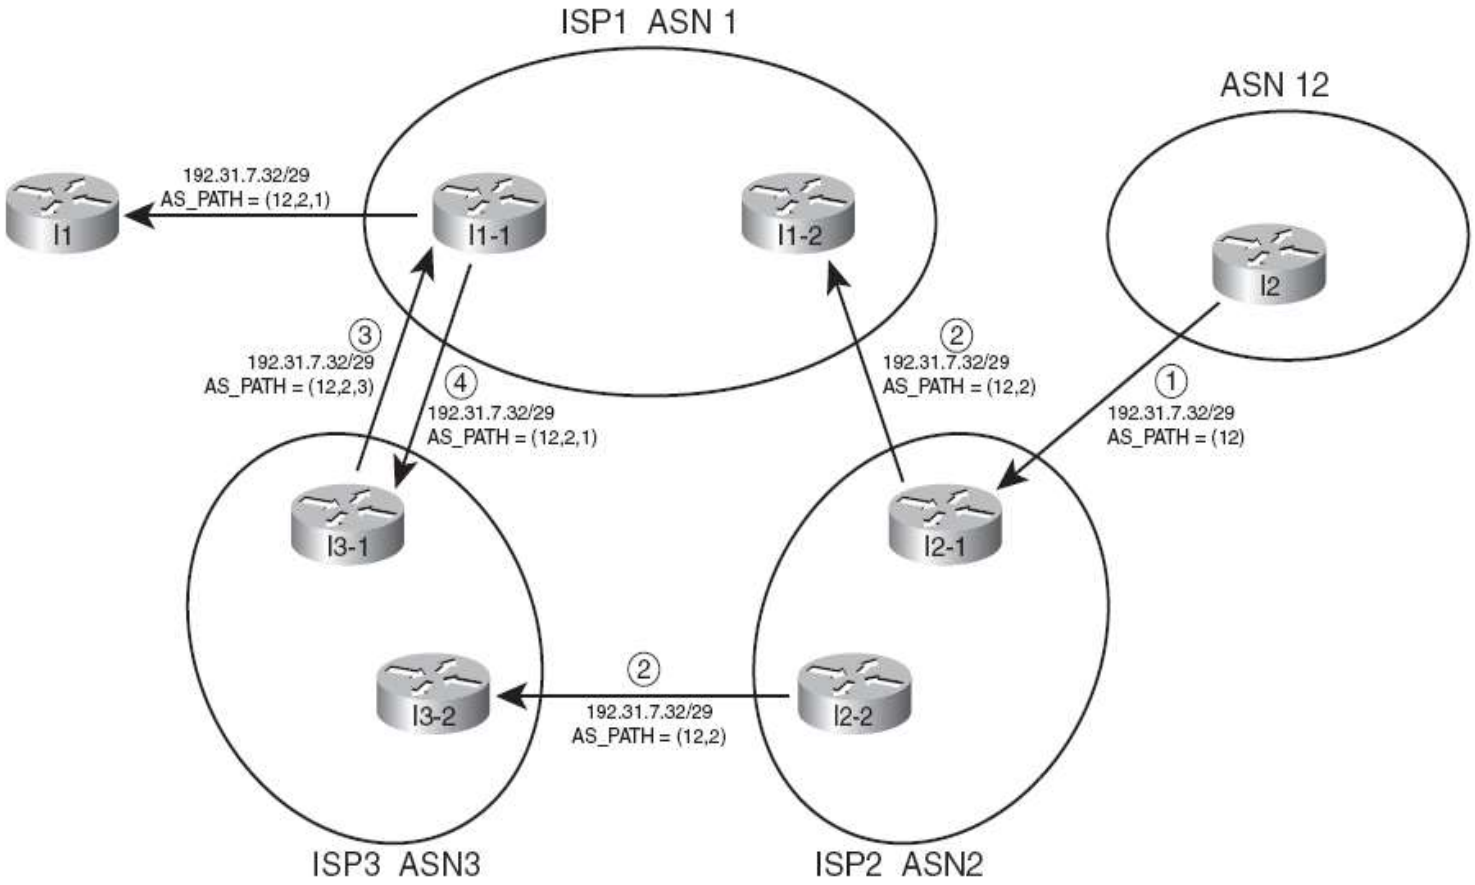

Heydari.Farzad@gmail.com

Introduction to BGP

Preventing routing loops in BGP

The *ASNs* listed in the *AS_Path* will do that.

When a BGP router receives an update, and a route advertisement lists an *AS_Path* with its own ASN, the router ignores that route.

Heydari.Farzad@gmail.com

Introduction to BGP

Preventing routing loops in BGP

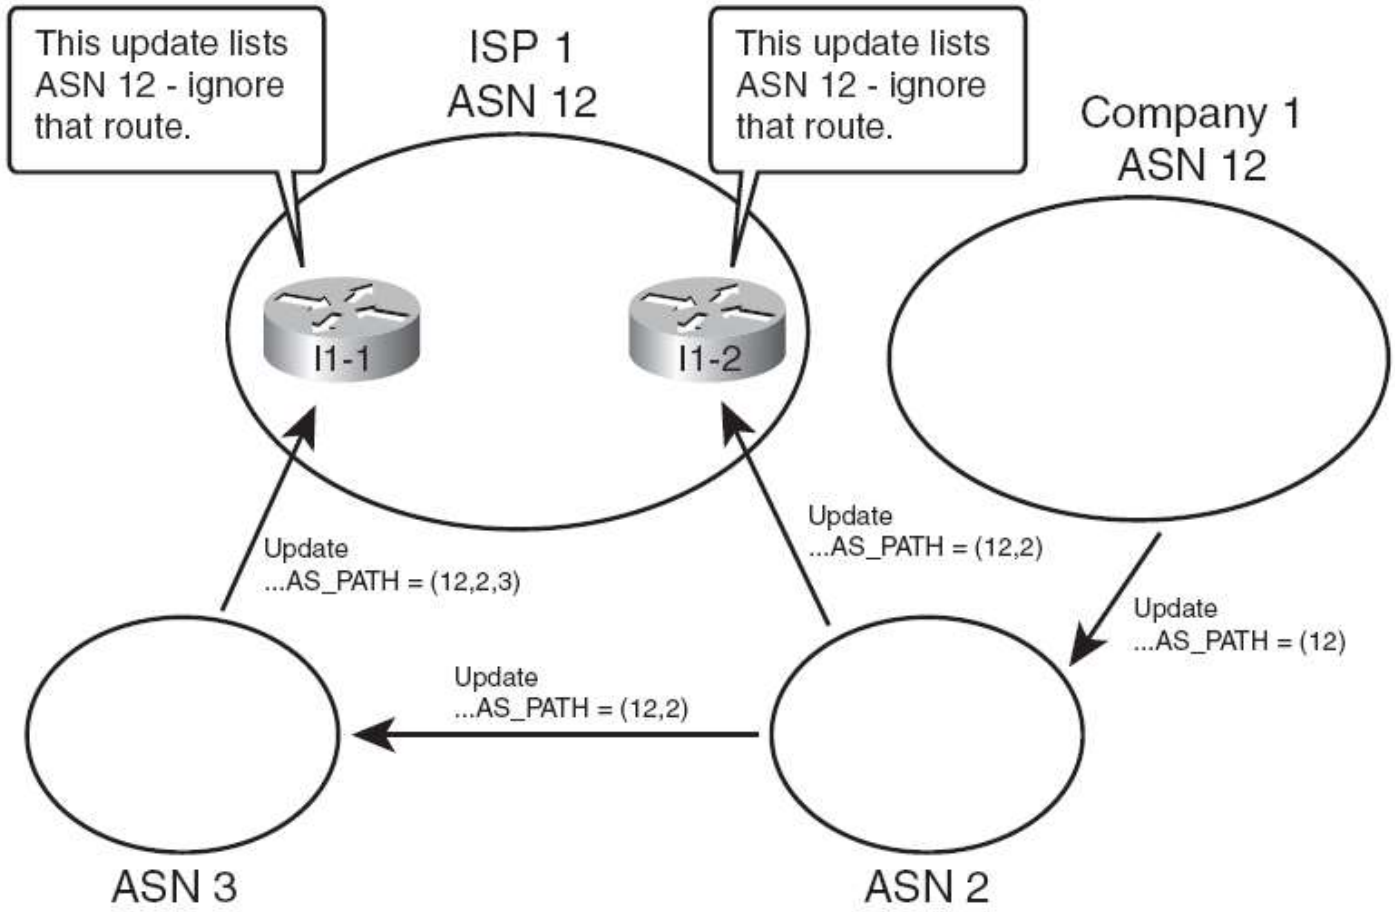

Heydari.Farzad@gmail.com

Introduction to BGP

IANA administers the assignment of unique values of *ASNs*

If *ASNs* are duplicated, the BGP loop prevention process can actually prevent parts of the Internet from learning about a route.

Heydari.Farzad@gmail.com

Introduction to BGP

Internal and External BGP

BGP defines two classes of neighbors (peers):

Internal BGP (iBGP)

BGP neighbor is in the same ASN.

When advertising , a BGP router does not updates the AS_Path PA.

External BGP (eBGP)

BGP neighbor is in the different ASN.

When advertising , a BGP router updates the AS_Path PA.

Heydari.Farzad@gmail.com

Introduction to BGP

Internal and External BGP

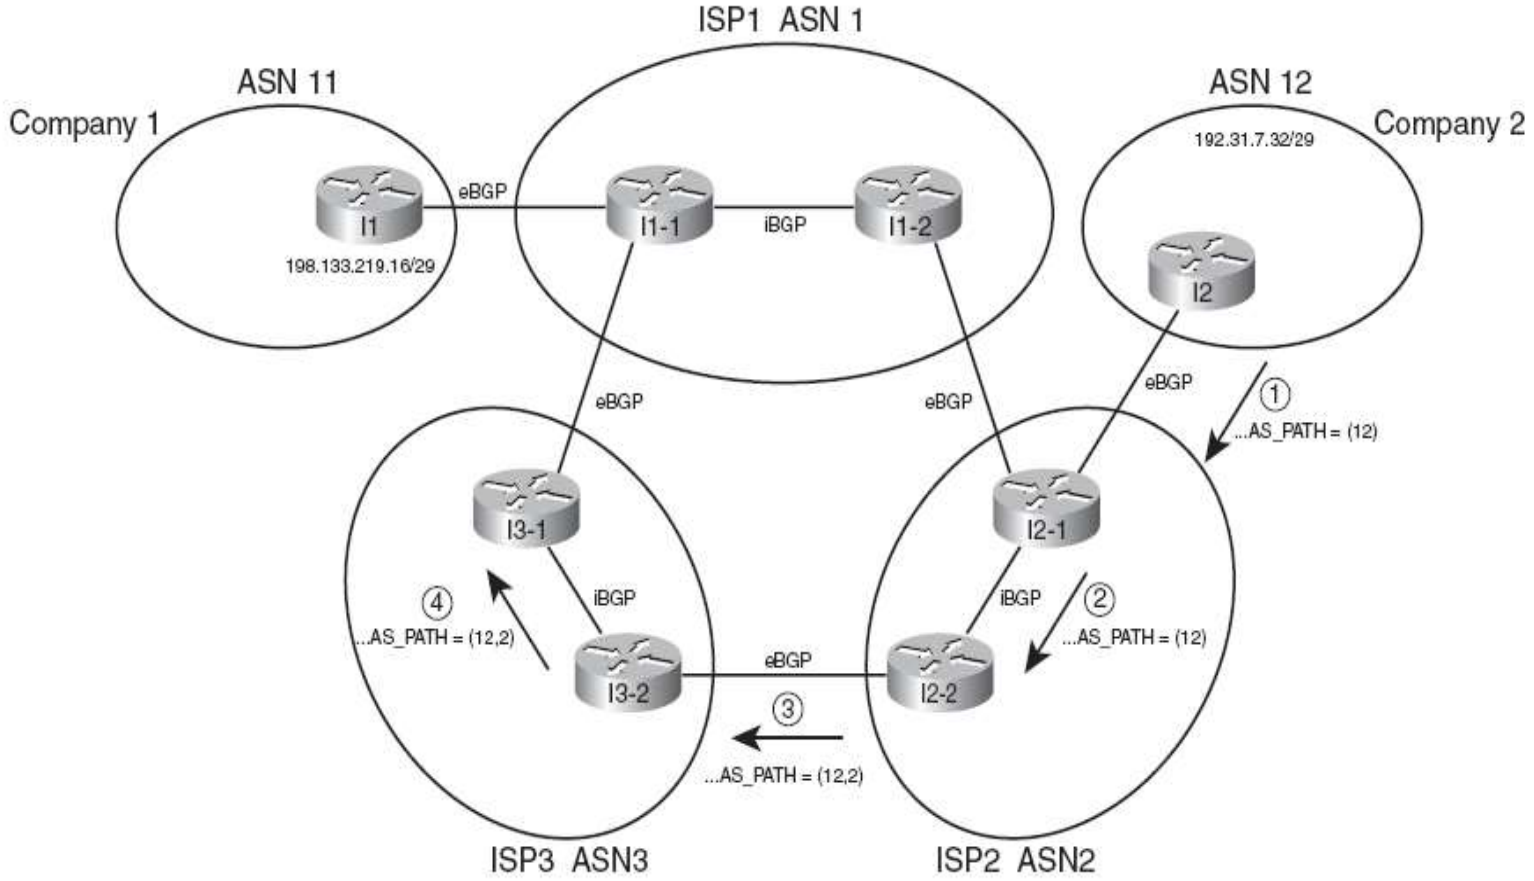

Heydari.Farzad@gmail.com

WIKIPEDIA

BGP table growth on the Internet.

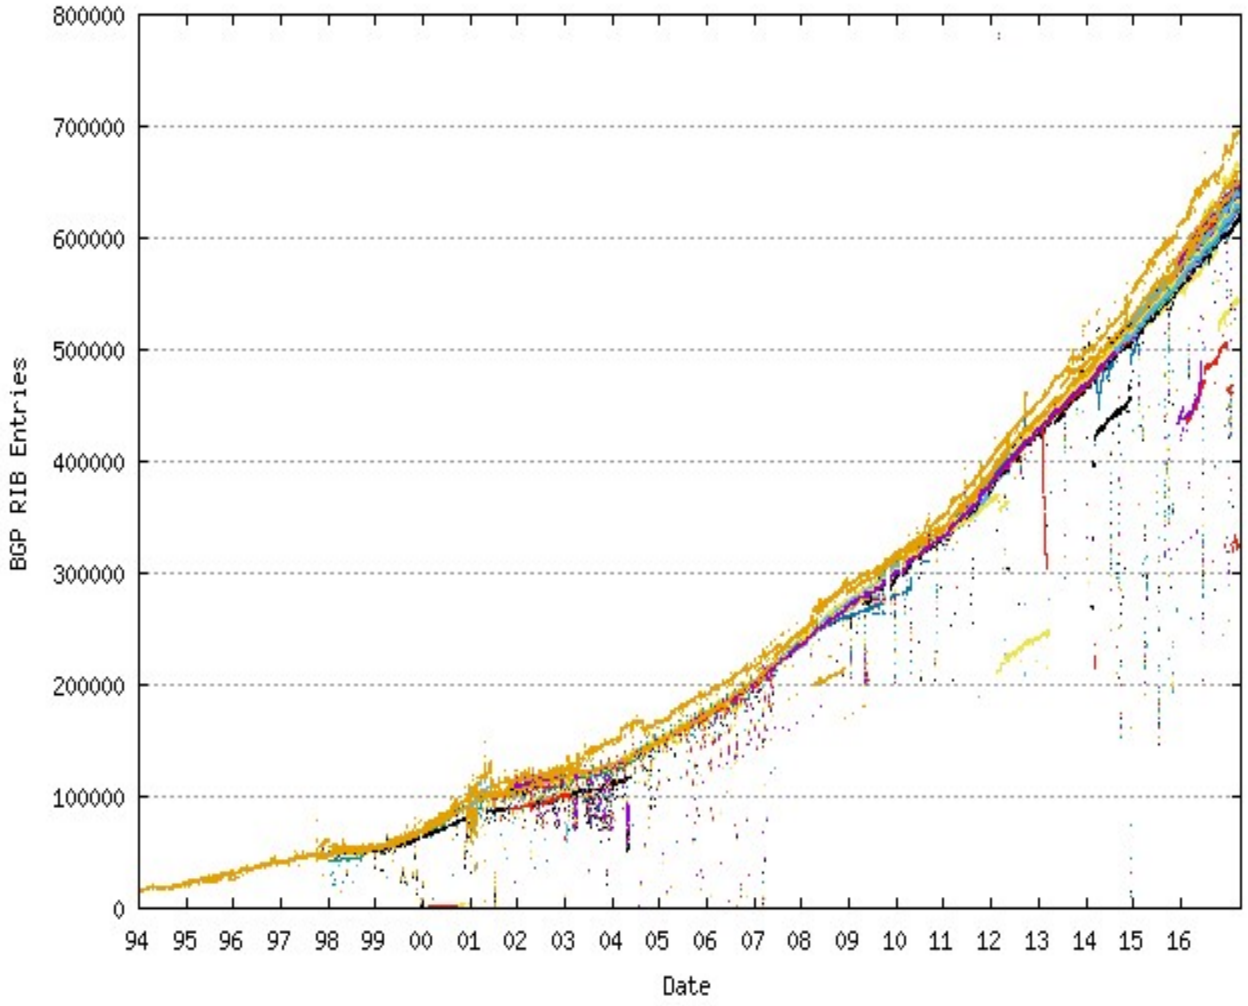

Heydari.Farzad@gmail.com

MikroTik User Meeting

Part II
External BGP

Heydari.Farzad@gmail.com

External BGP

- BGP first forms a neighbor relationship with peers.
- BGP then learns information from its neighbors, placing that information in the BGP table.
- Finally, BGP analyzes the BGP table to choose the best working route for each prefix in the BGP table, placing those routes into the IP routing table.

Heydari.Farzad@gmail.com

Requirements for Forming eBGP Neighborships

- Each router must be part of a TCP connection with the other router, with the remote router's IP address used in that TCP connection matching what the local router configures in a BGP neighbor remote-as .

BGP peer first form a TCP connection; later, BGP messages flow over that connection, which allows BGP routers to know when the messages arrived at the neighbor, and when they did not.

Heydari.Farzad@gmail.com

BGP peer with single link

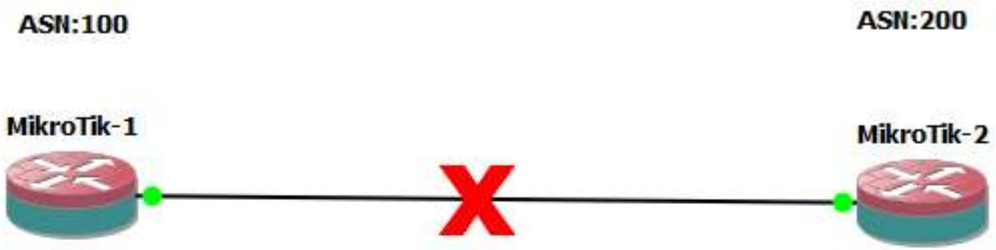

Heydari.Farzad@gmail.com

Redundancy Between BGP

Heydari.Farzad@gmail.com

Issues When Redundancy Exists Between eBGP Neighbors

Solutions

1. Configure two peers on each router,
one for each of the neighbor's interface IP addresses.

If one link fails, the other neighborship can remain up and working.

However, both neighborships exchange BGP routes, consuming bandwidth and more memory in the BGP table.

Heydari.Farzad@gmail.com

Solutions

2. use bridge interfaces(loopback) as the TCP connection endpoints

The two routers each configure a loopback interface and IP address, and use those loopback IP addresses as the source of their single BGP TCP connection.

If one of the multiple links fails, the loopback interface does not fail.

Heydari.Farzad@gmail.com

eBGP Neighbor Configuration with loopback

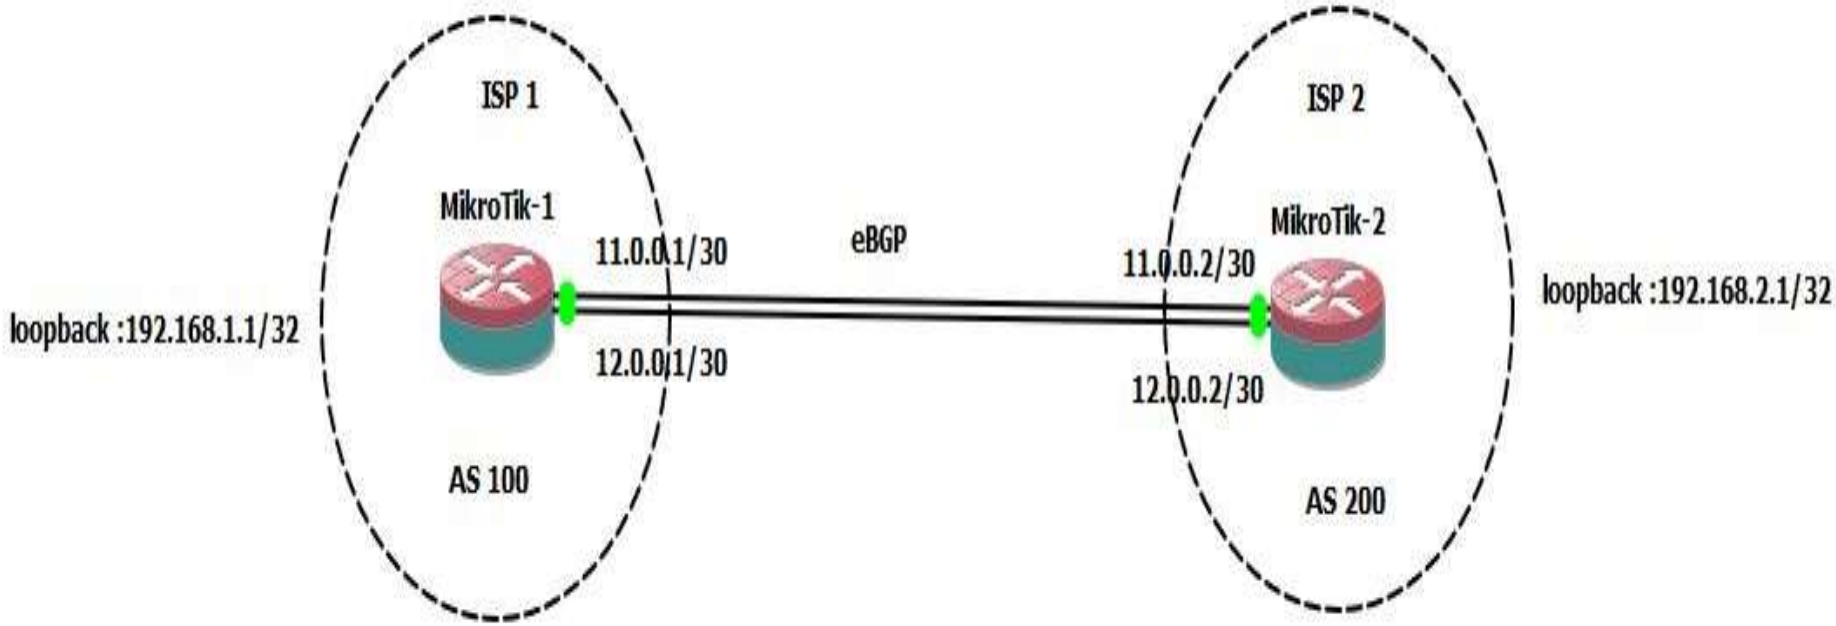

Heydari.Farzad@gmail.com

Step 1

Make a loopback interface.

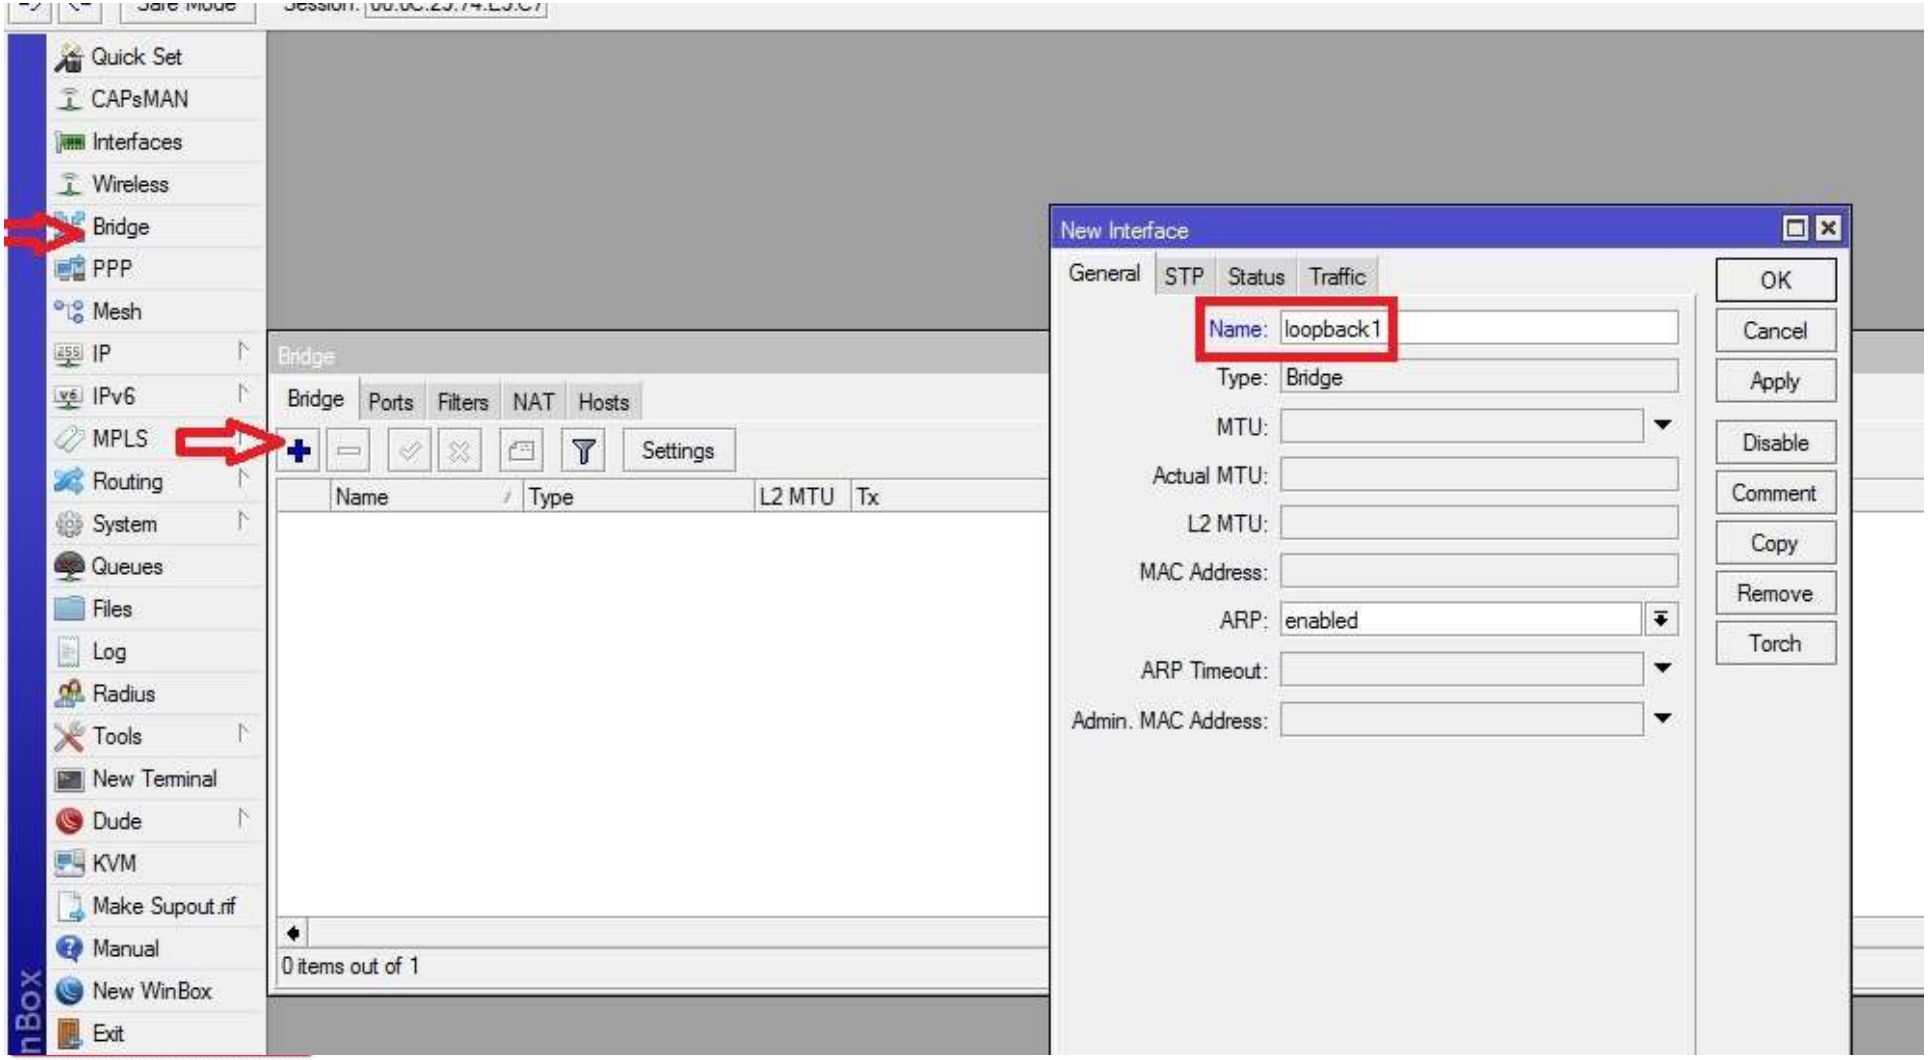

Step 2

Configure an IP address on a loopback interface on each router.

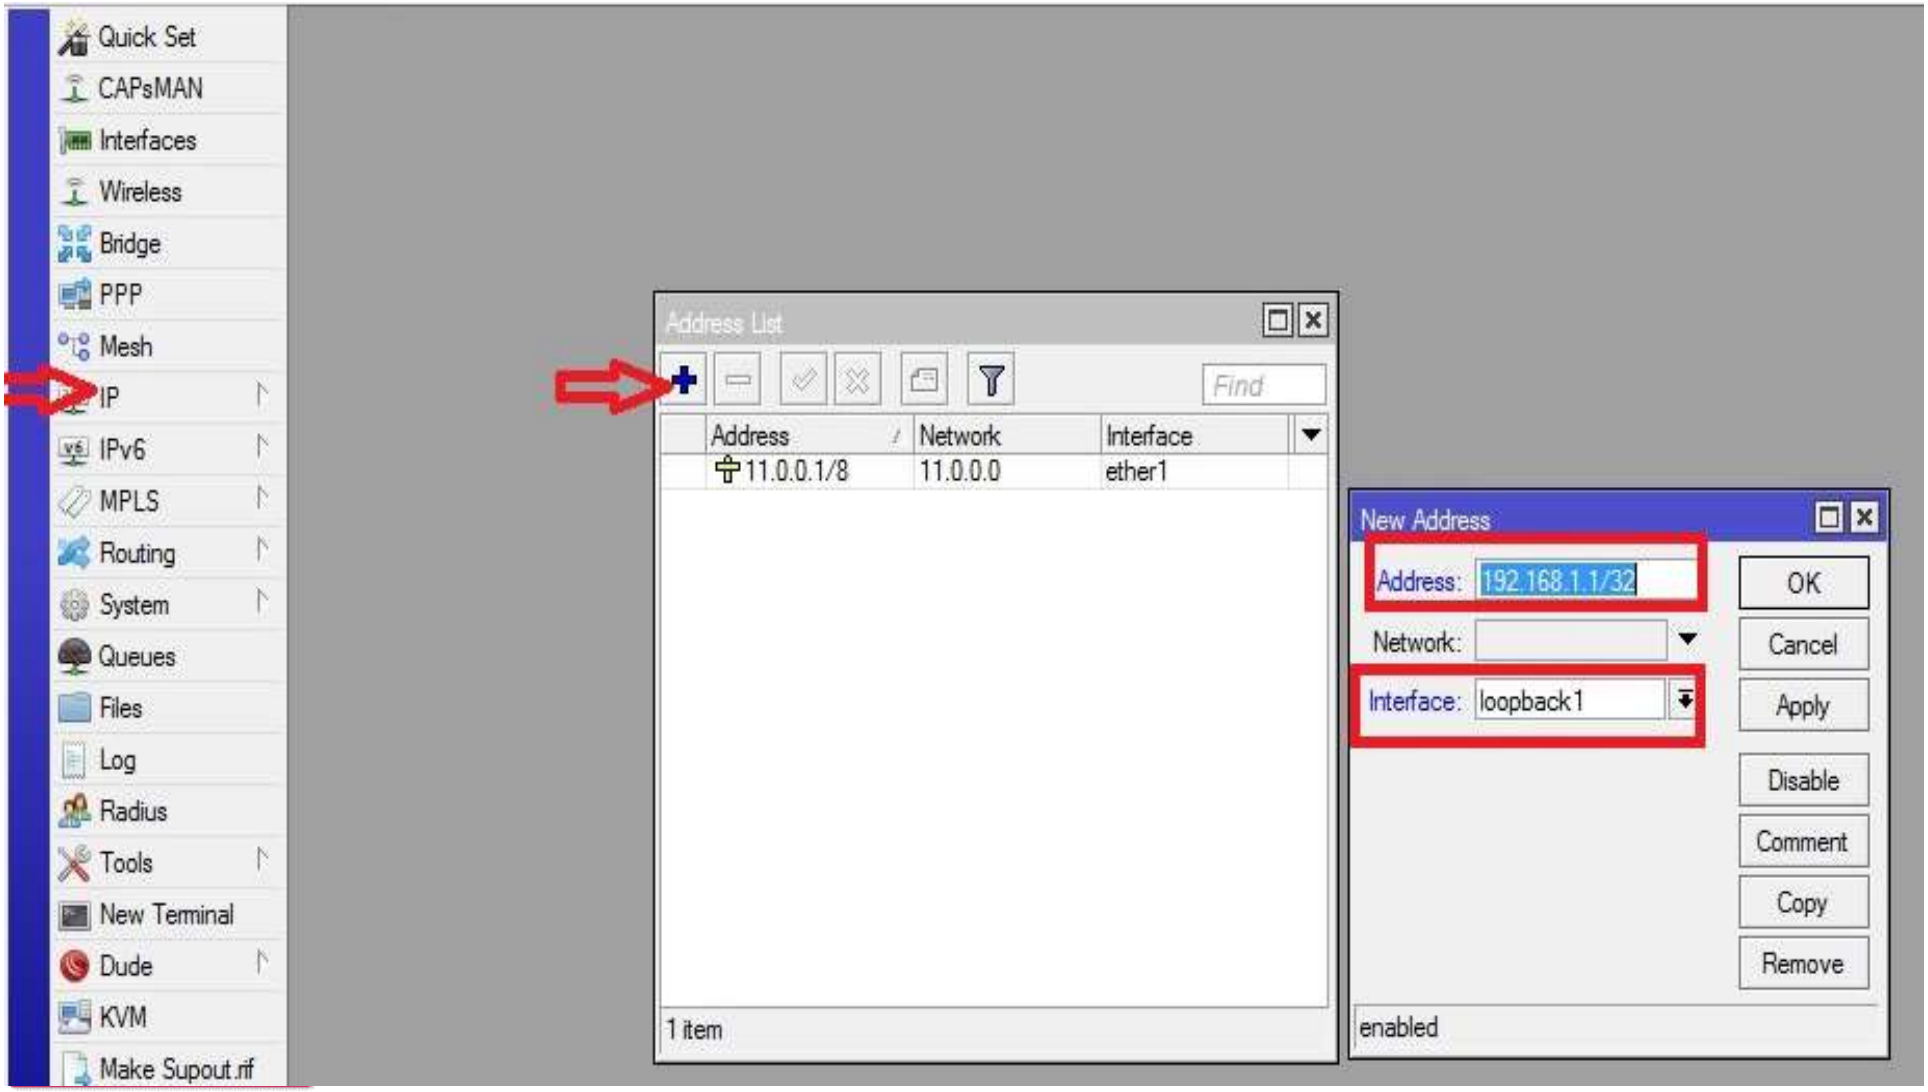

Step3

Configure the BGP on router 1

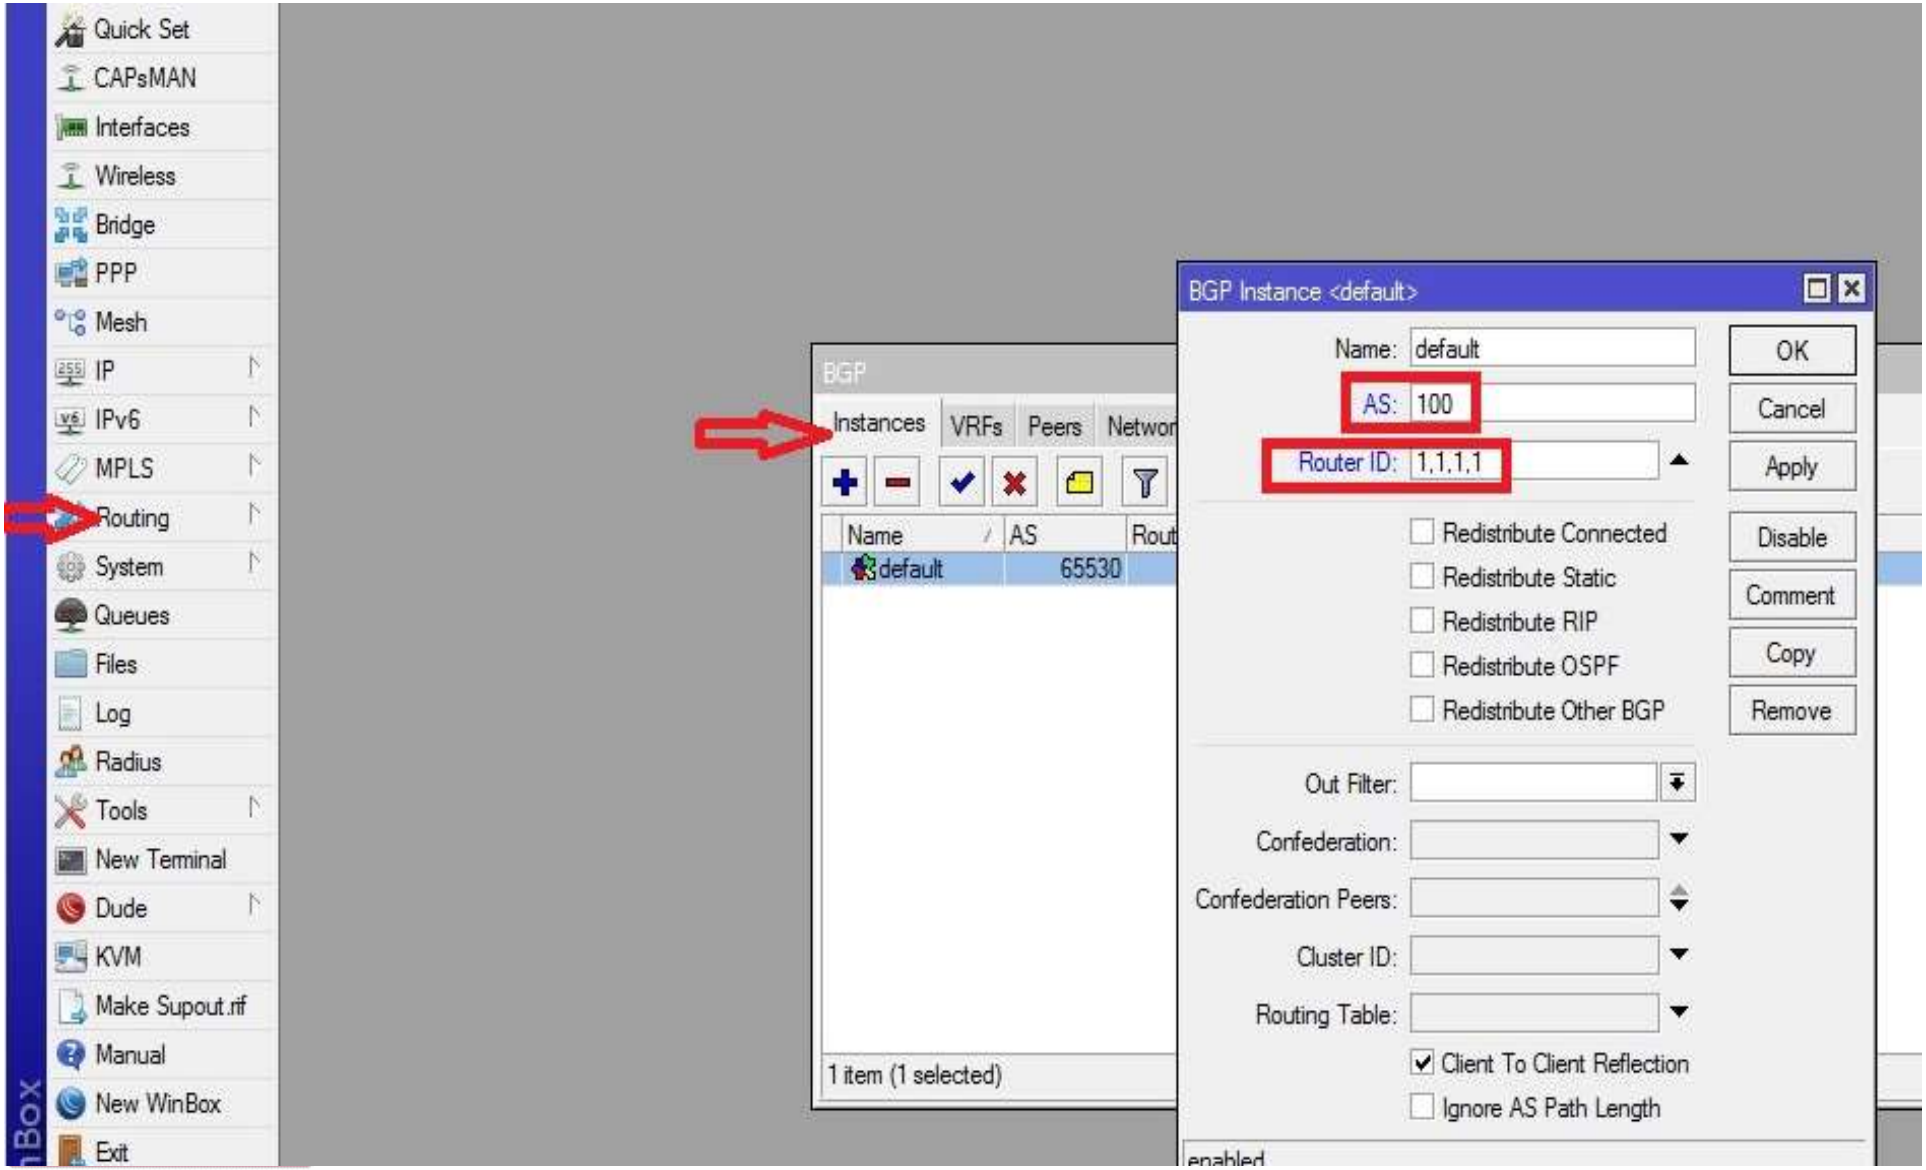

Step3

Configure the BGP on router 2

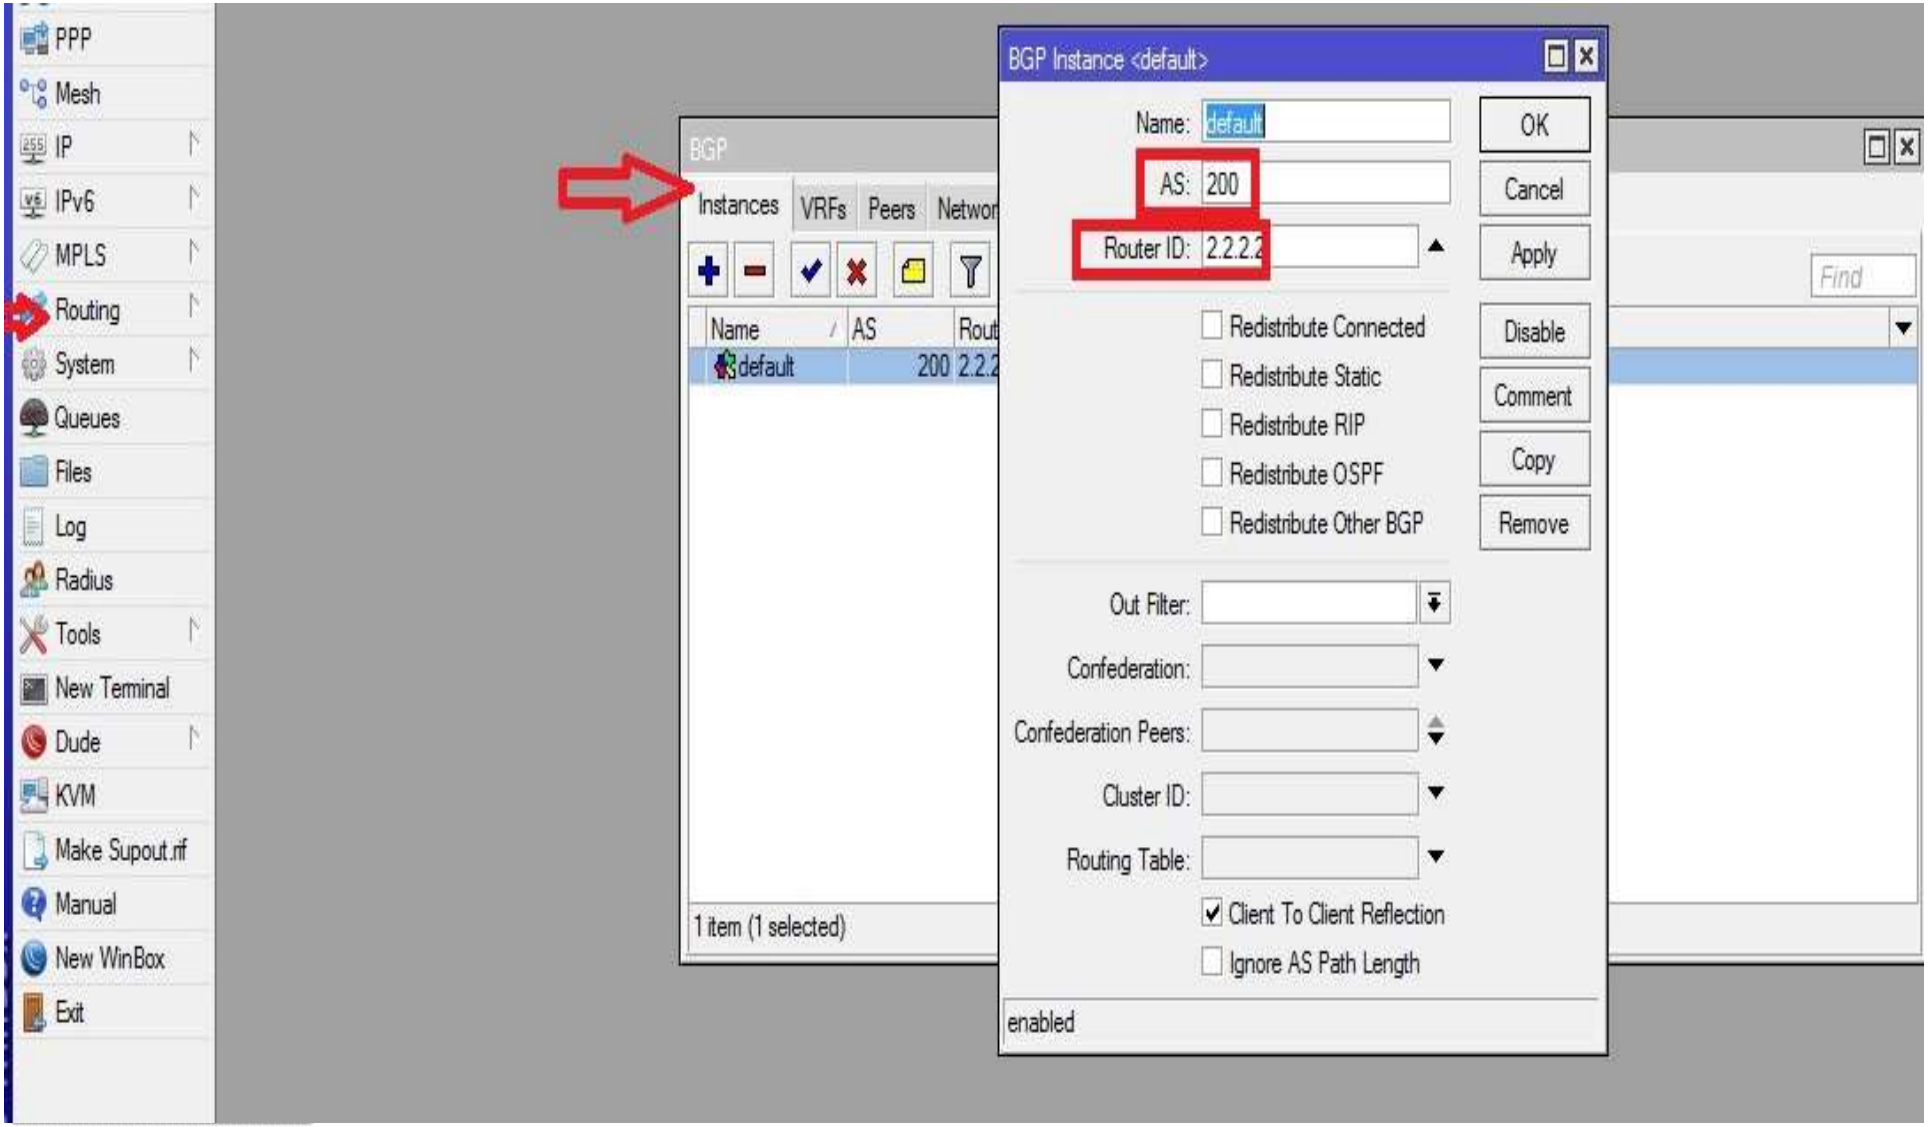

Step 4

Configure the BGP peer on each router to refer to the other router's loopback IP address.

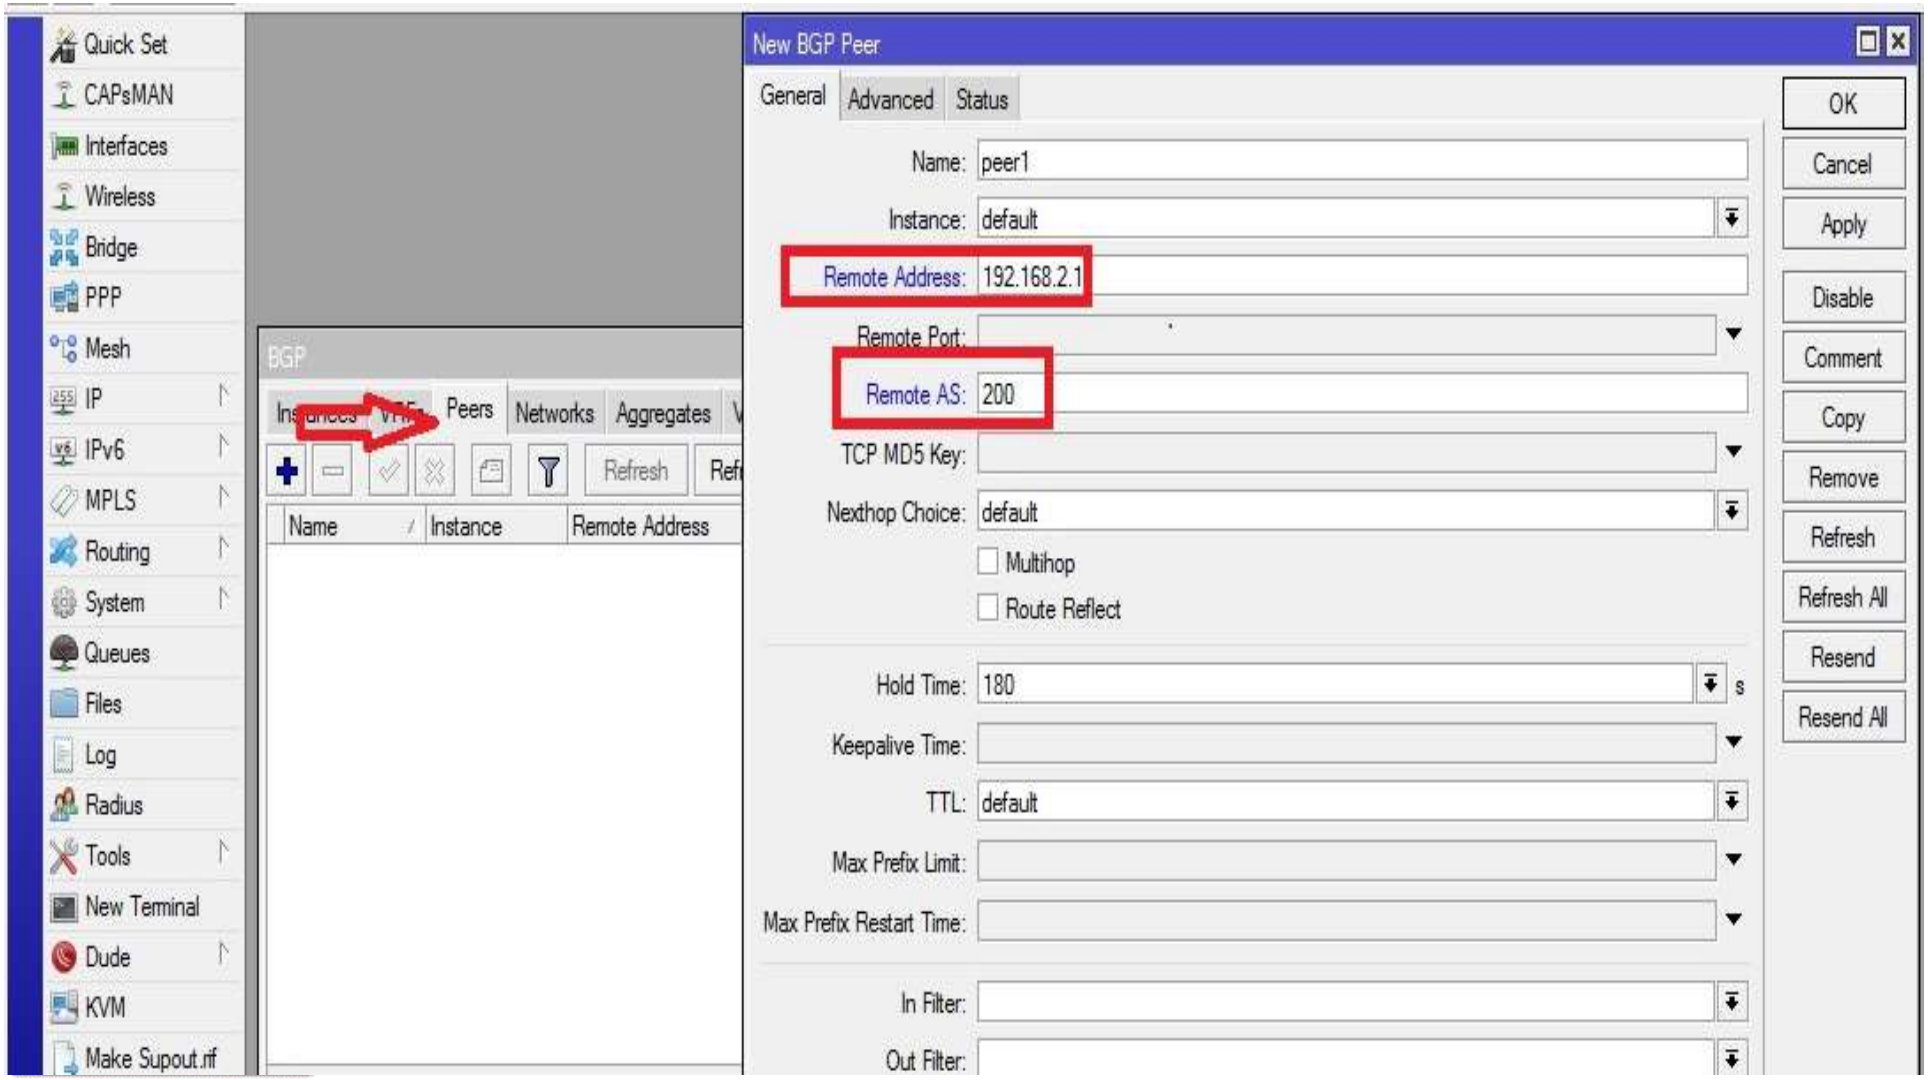

Step 4

Configure the BGP peer on router 2 to refer to the other

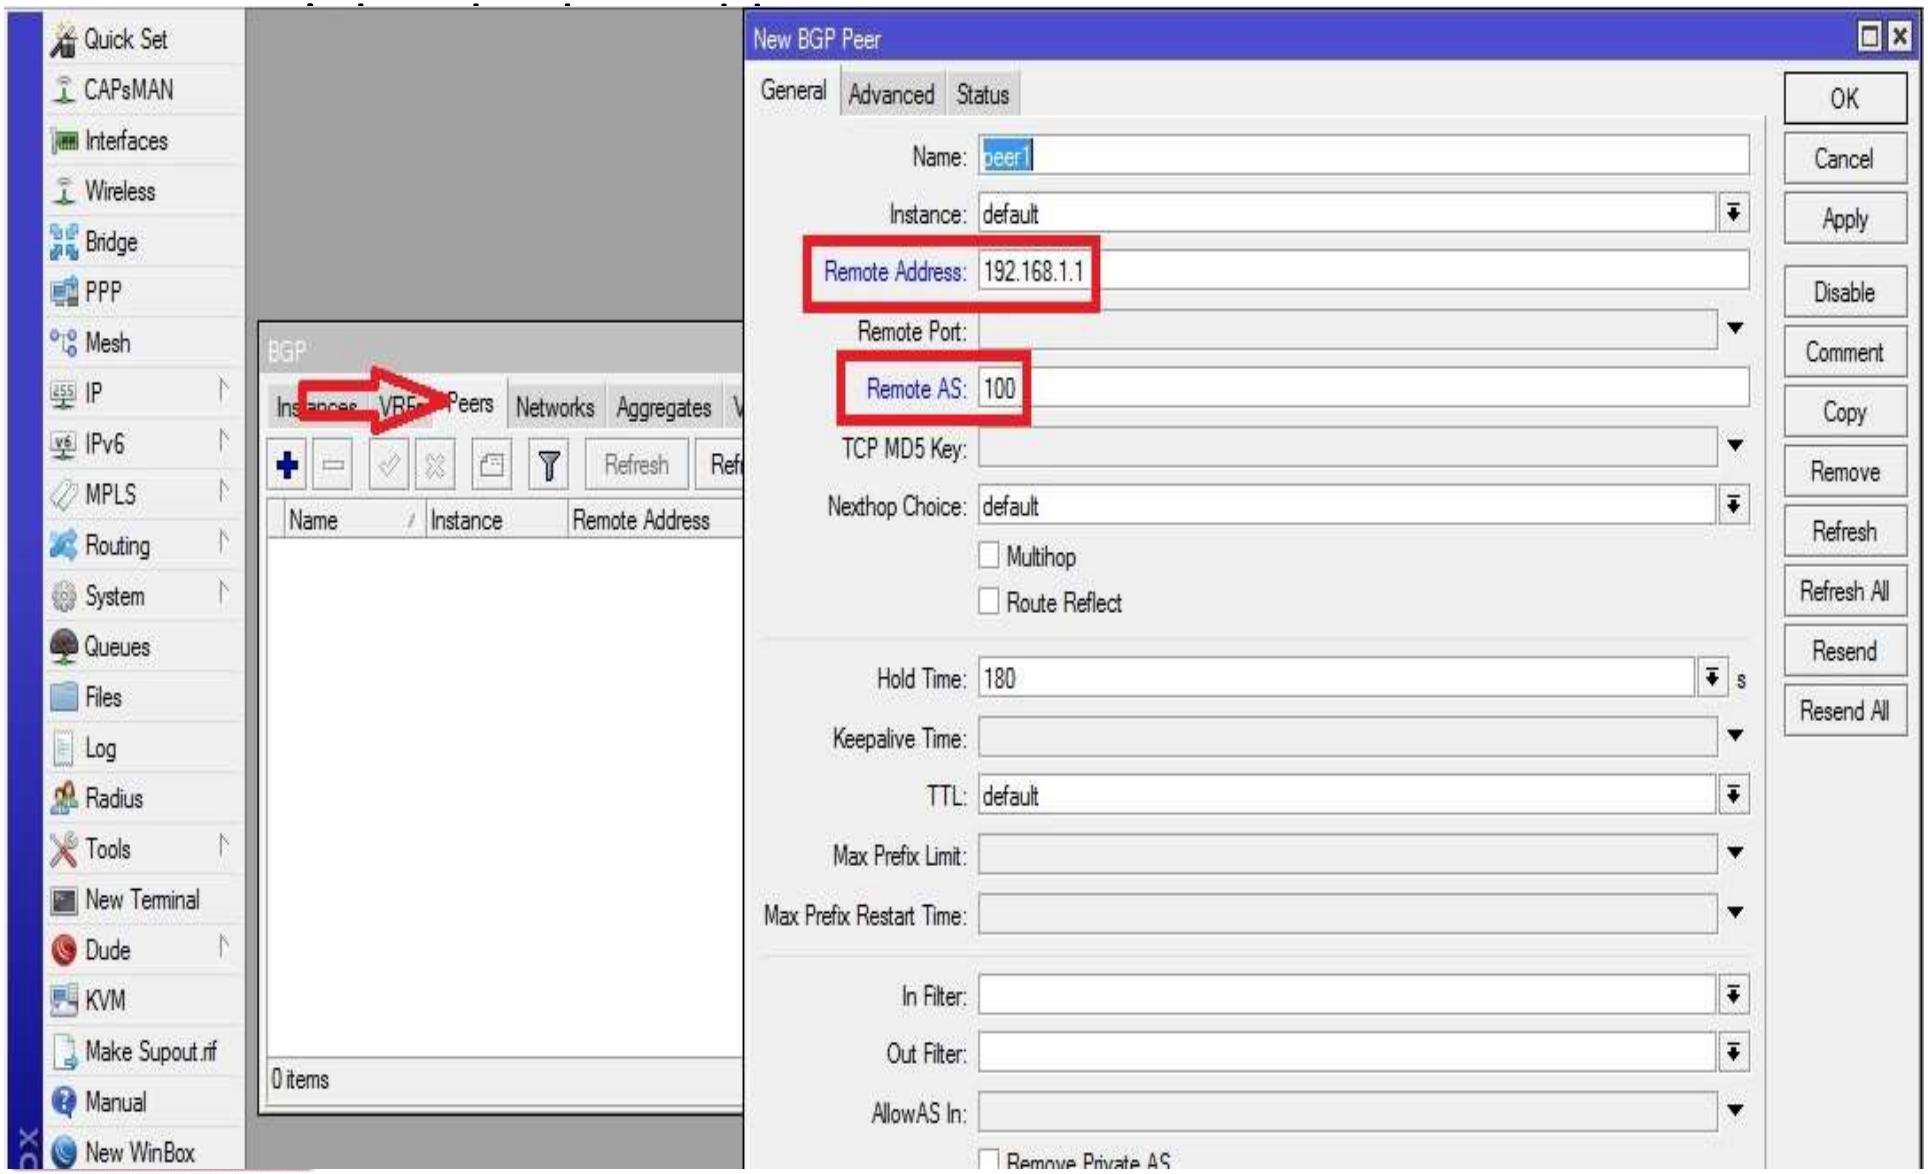

Step 5

Tell BGP on each router to use the loopback IP address as the source IP address.

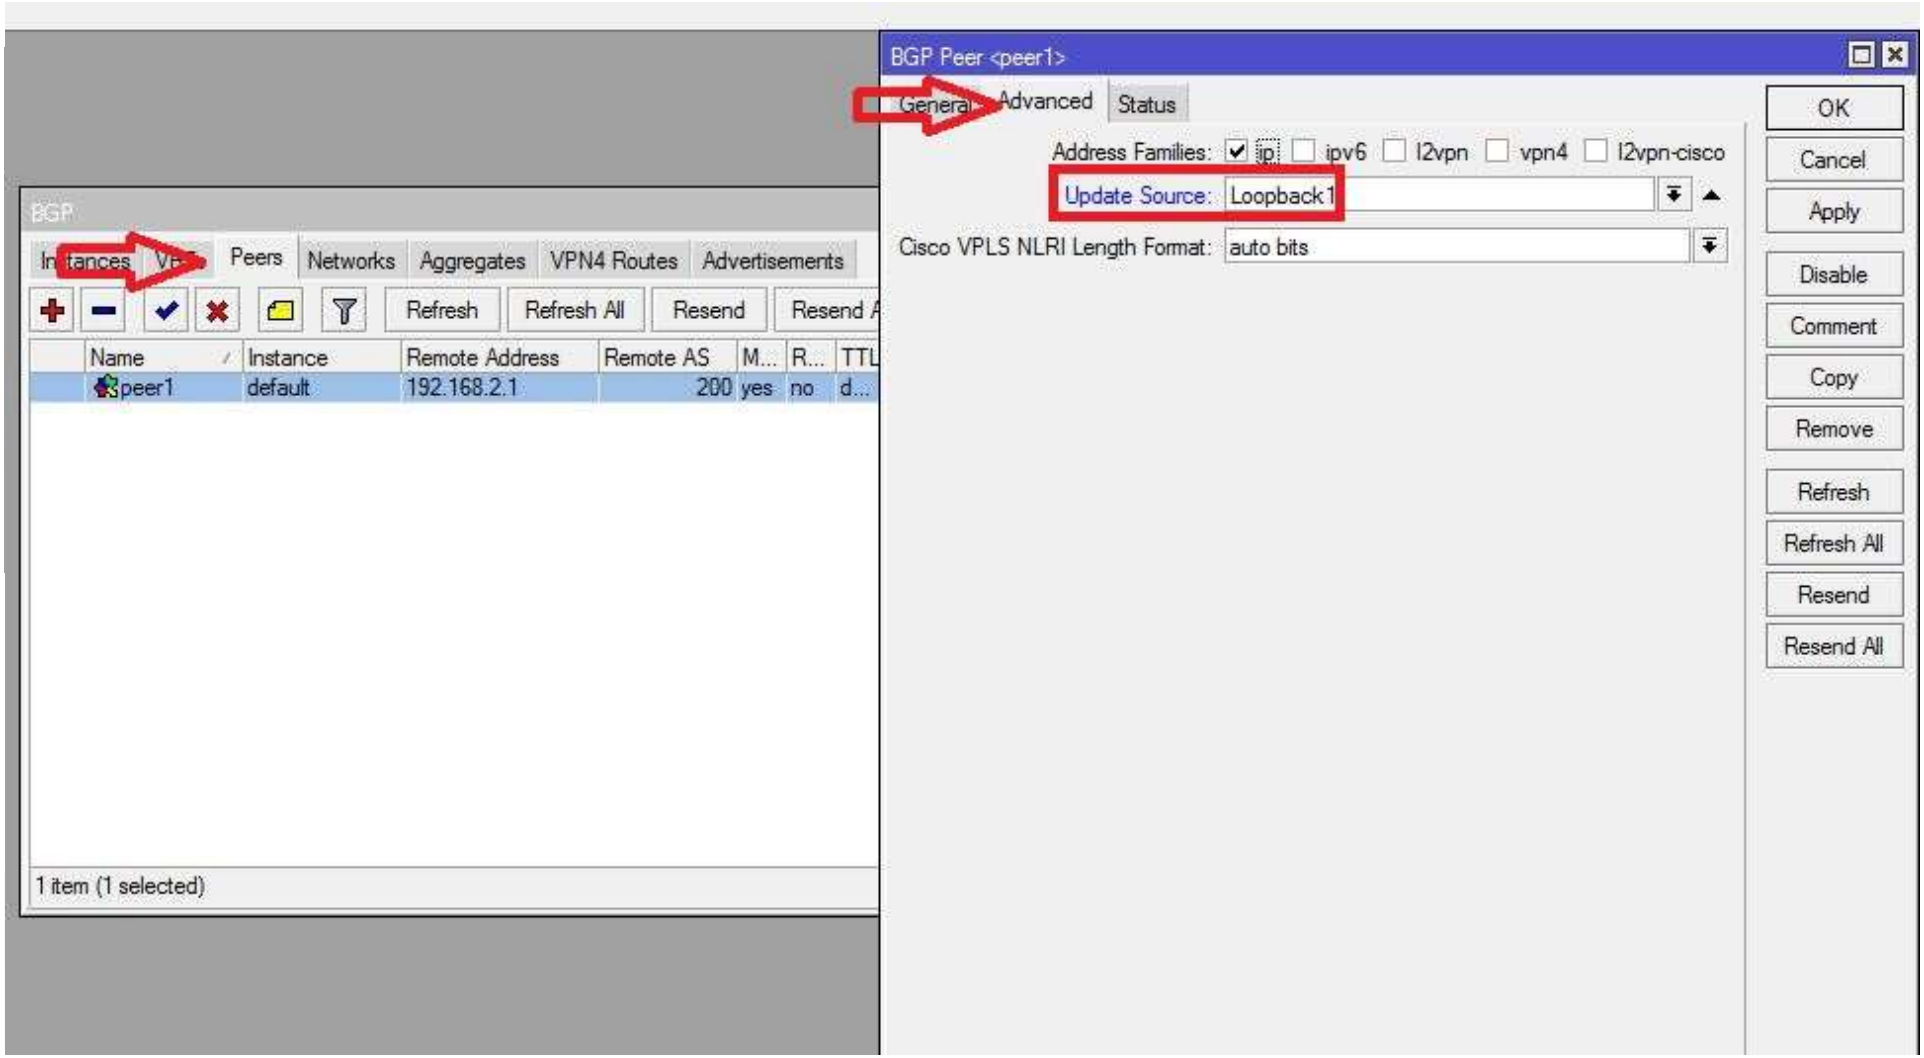

Step 6

Make sure each router has IP routes so that they can forward packets to the loopback interface IP address of the other router.

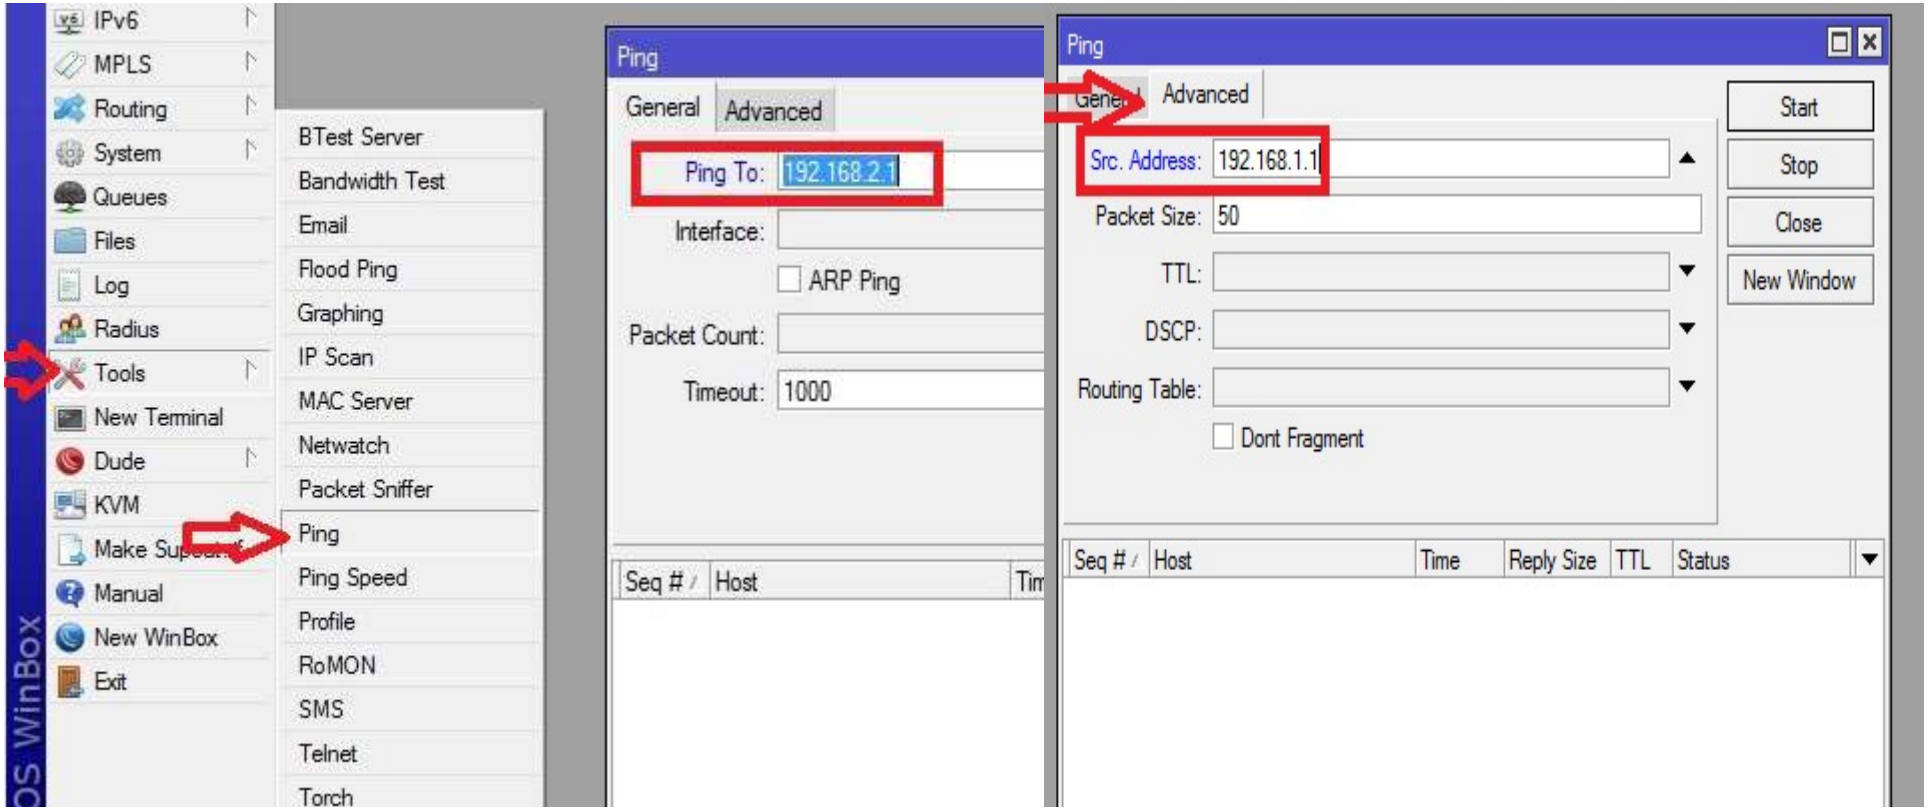

Note

If using static routes, make sure to configure the routes so that all redundant paths would be used.

If using an IGP, make sure the configuration allows the two routers to become IGP neighbors over all redundant links.

Heydari.Farzad@gmail.com

Step 7

Configure eBGP multihop.

By default, when building packets to send to an eBGP peer, RouterOS sets the IP Time-To-Live (TTL) field in the IP header to a value of 1.

With this default action, the eBGP neighborship fails to complete when using loopback interface IP addresses.

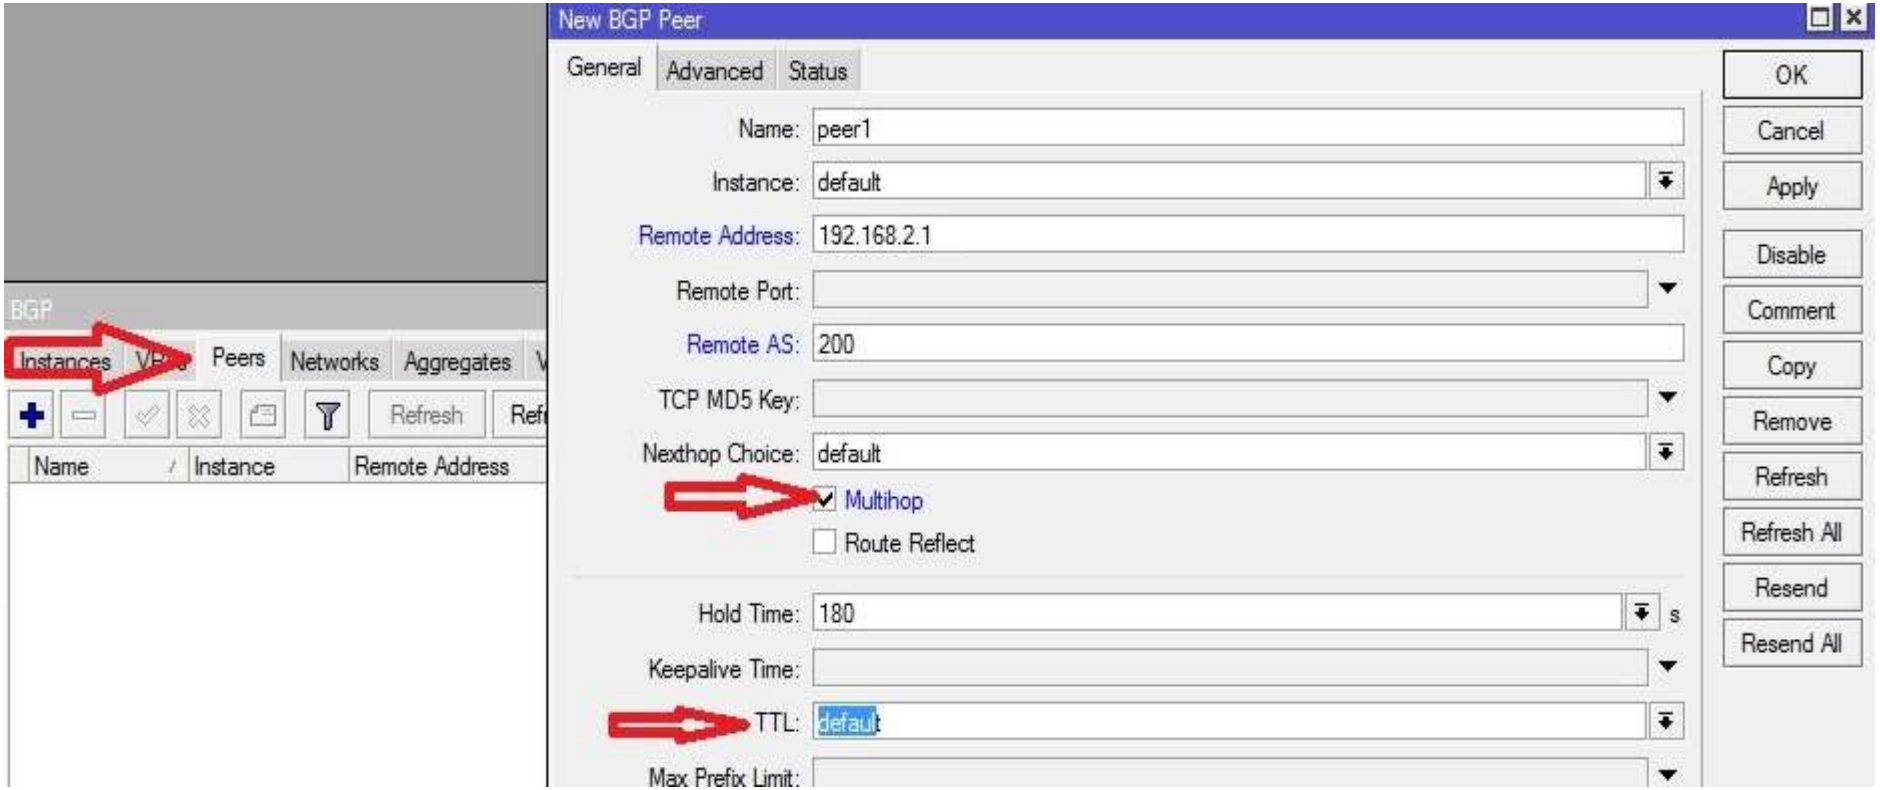

eBGP Neighbor Configuration with loopback

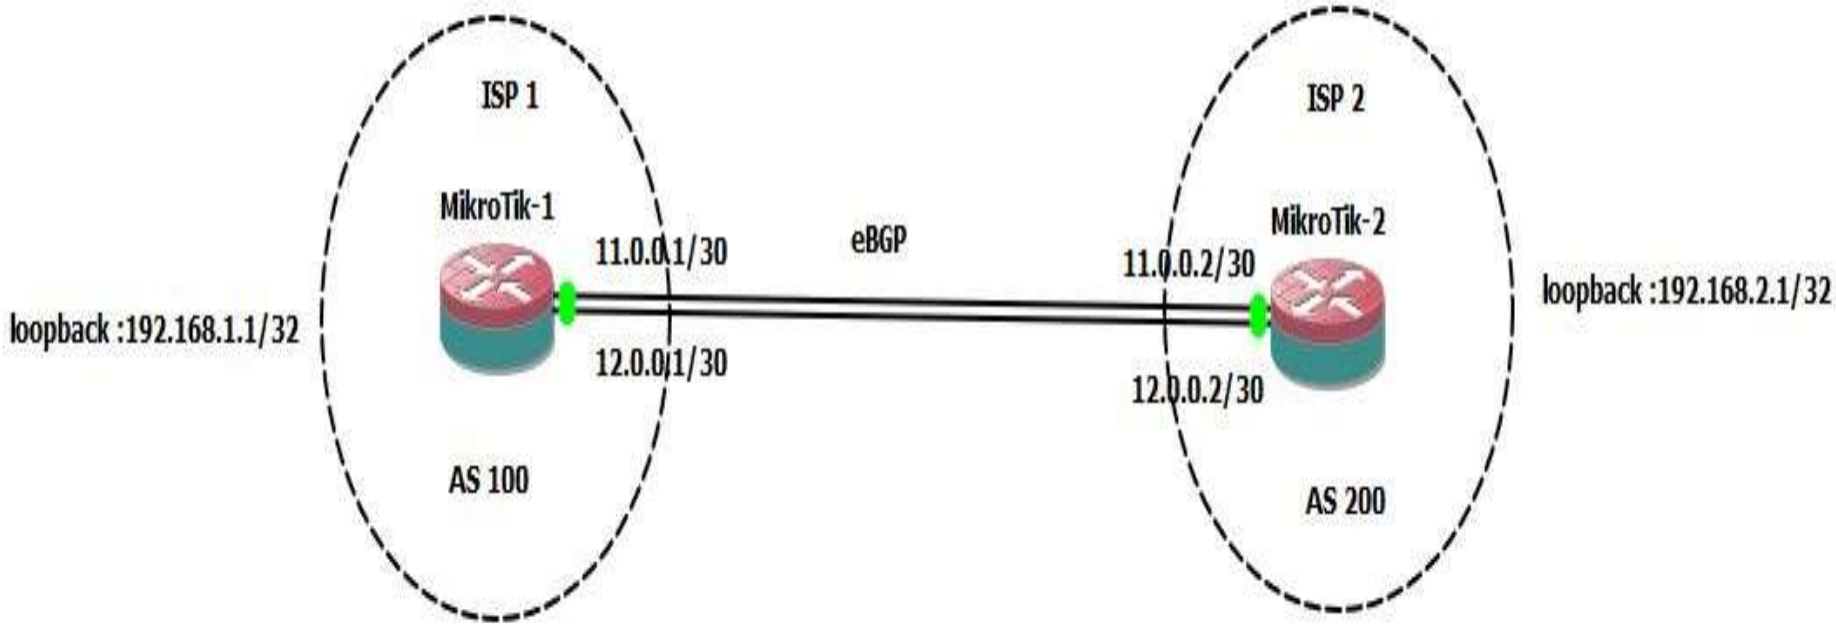

Heydari.Farzad@gmail.com

Verifying eBGP Neighbor Status

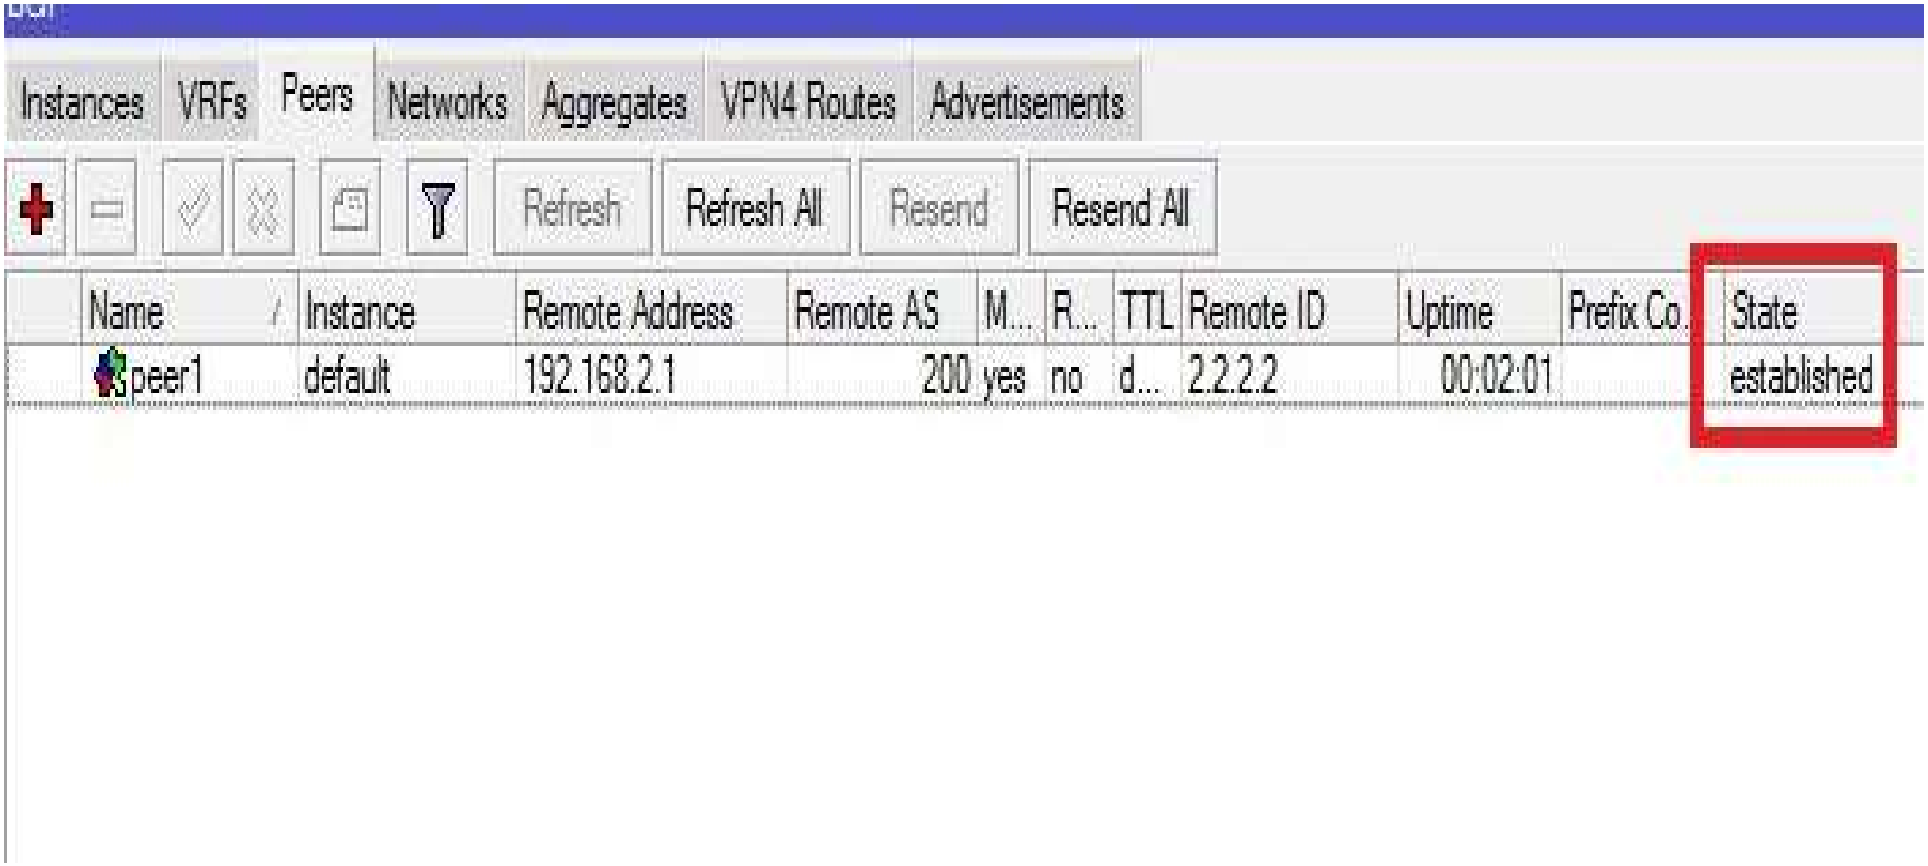

Heydari.Farzad@gmail.com

External BGP

BGP Internals and Verifying eBGP Neighbors

State	Typical Reasons
Idle	The BGP process is either administratively down or awaiting the next retry attempt.
Connect	The BGP process is waiting for the TCP connection to be completed. You cannot determine from this state information whether the TCP connection can complete.
Active	The TCP connection has been completed, but no BGP messages have been sent to the peer yet.
Opensent	The TCP connection exists, and a BGP Open message has been sent to the peer, but the matching Open message has not yet been received from the other router.
Openconfirm	An Open message has been both sent to and received from the other router. The next step is to receive a BGP Keepalive message (to confirm all neighbor-related parameters matched) or BGP Notification message (to learn there is some mismatch in neighbor parameters).
Established	All neighbor parameters match, the neighbor relationship works, and the peers can now exchange Update messages.

Heydari.Farzad@gmail.com

Verifying eBGP peer Status

State

either lists the BGP peer state

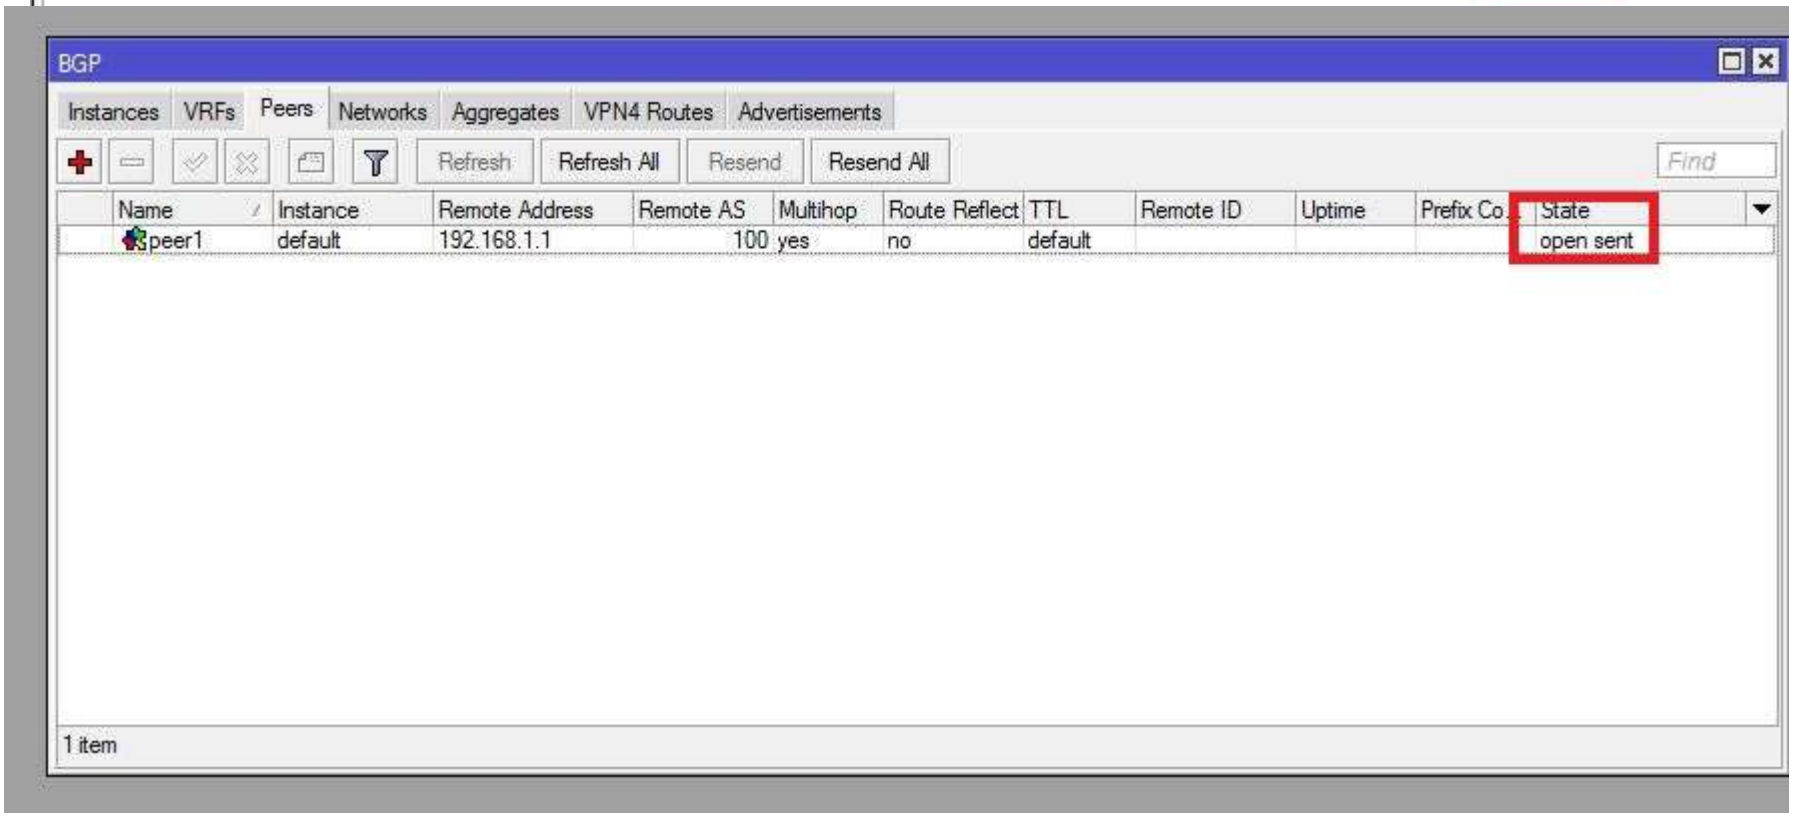

Heydari.Farzad@gmail.com

Heydari.Farzad@gmail.com

Contact: Heydari.Farzad@gmail.com