

BGP Multihoming & Failover using VRRP

Hani Rahrouh

hr@wirelessnetware.ca

First Canadian MUM

October 19th, 2015

Montreal, Canada

About me

- Hani Rahrouh
- MikroTik Certified since 2008
- MikroTik Consultant
- MikroTik Certified Trainer Since 2011 @ www.wirelessnetware.ca

<http://www.mikrotik.com/training/partners/northamerica/canada>

- MikroTik Distributor @ www.netwire.ca

Come to our free MikroTik workshops and Events

Come join us for free workshop.

We know MikroTik RouterBOARDS and RouterOS, so we can answer all your questions and help you learn more about the things you're interested in.

www.wirelessnetware.ca

Overview

Fully redundant fault tolerant internet connectivity


- BGP Multi-homing
 - Why Multi-home?
 - How to Multi-home?
- Dynamic Failover using VRRP
- Manual Failover
- Monitoring

Everything about reliability and redundancy



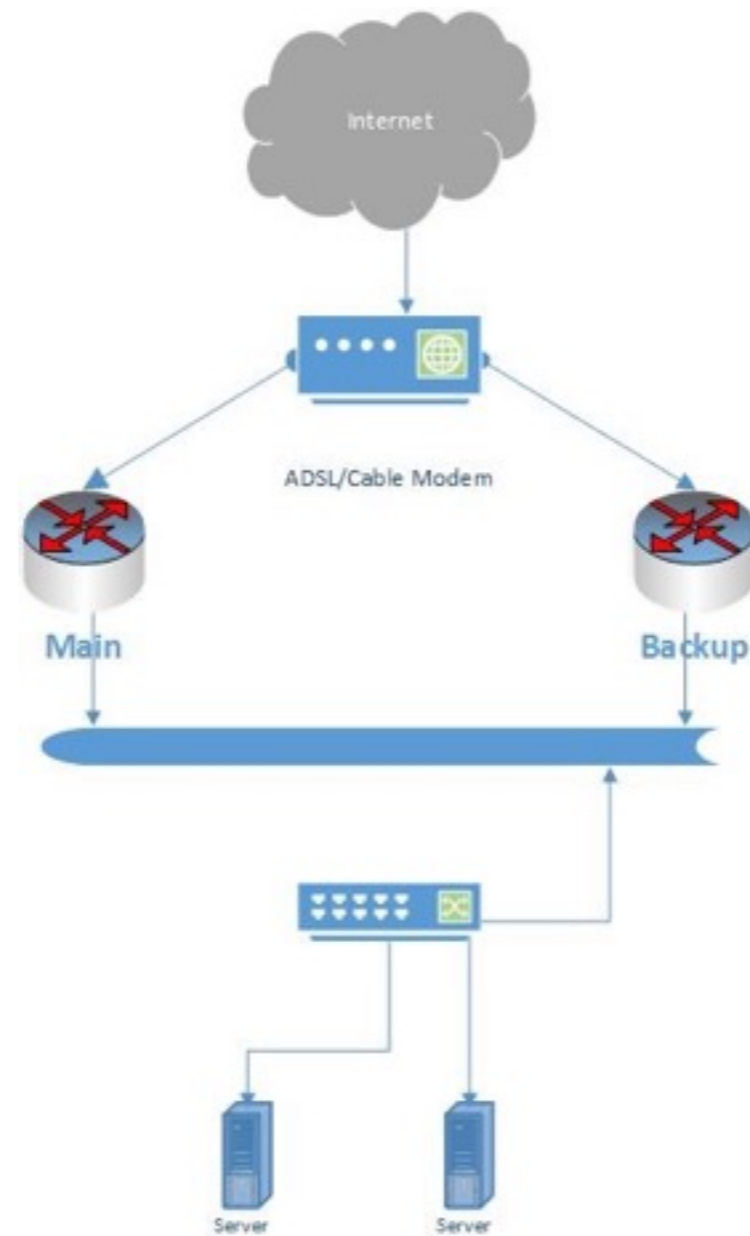
Two of everything

VRRP



Redundant Router(GW)

Virtual Router Redundancy Protocol (VRRP)



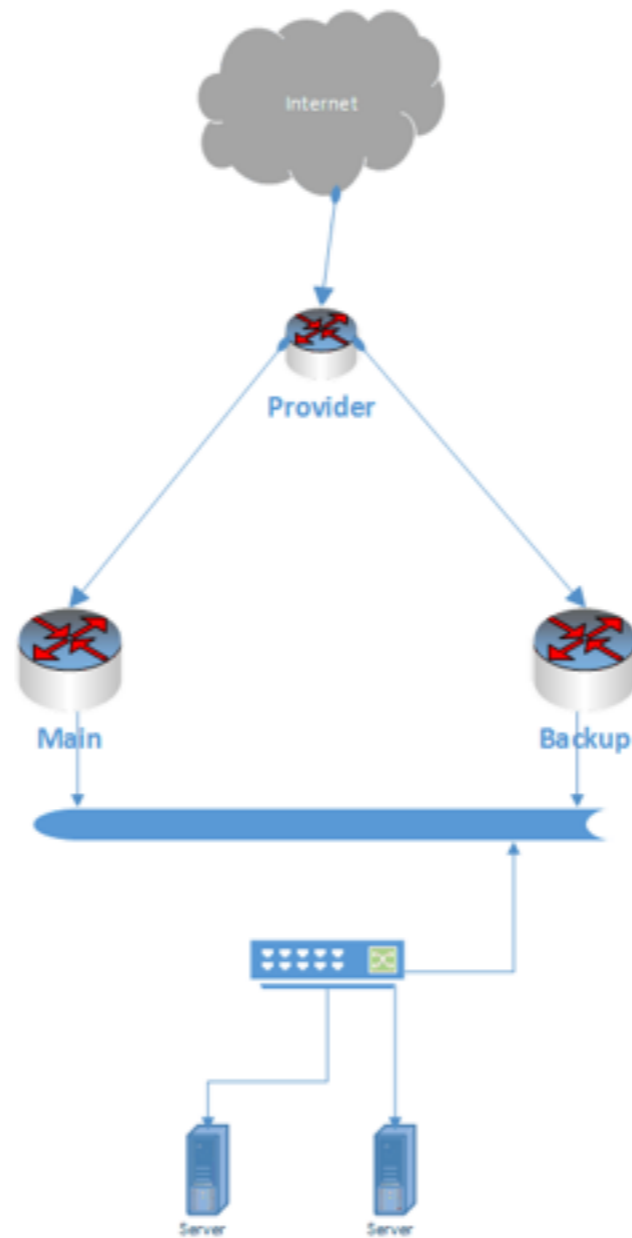
Gateway Failed

Redundant gateway!

Is **VRRP** a good idea when
we have connection-tracking
enable on the routers!

Problems

- **ARP**
Traffic originated by a VRRP cluster uses a Virtual MAC address of the form 00-00-5e-00-01-<VRID>.
- **Connection-Tracking**
- **Gateway failed**



Redundant Gateway

BGP Multi-Homing

Why Multi-home and Why would I want it?

It's all about:

- Redundancy "Fail-over"
- Diversity
- Reliability

Why Multi-home?

- Redundancy "Fail-over"

One connection to internet means the network is depend on:


- Local Router (Configuration, RouterOS and RouterBOARD)
- WAN Media (physical failure, carrier failure)

Why Multi-home?

- Reliability
 - Business critical applications demand continuously availability
 - Lack of redundancy implies lack of reliability implies loss of revenue

Problems

- Connection tracking is unable to keep valid track of connections with multi-homed BGP
- ARP Table Refresh periods
- Client gateway failure
- Downtime
- Hardware failure



Fully redundant fault tolerant internet connectivity

Provider



Core



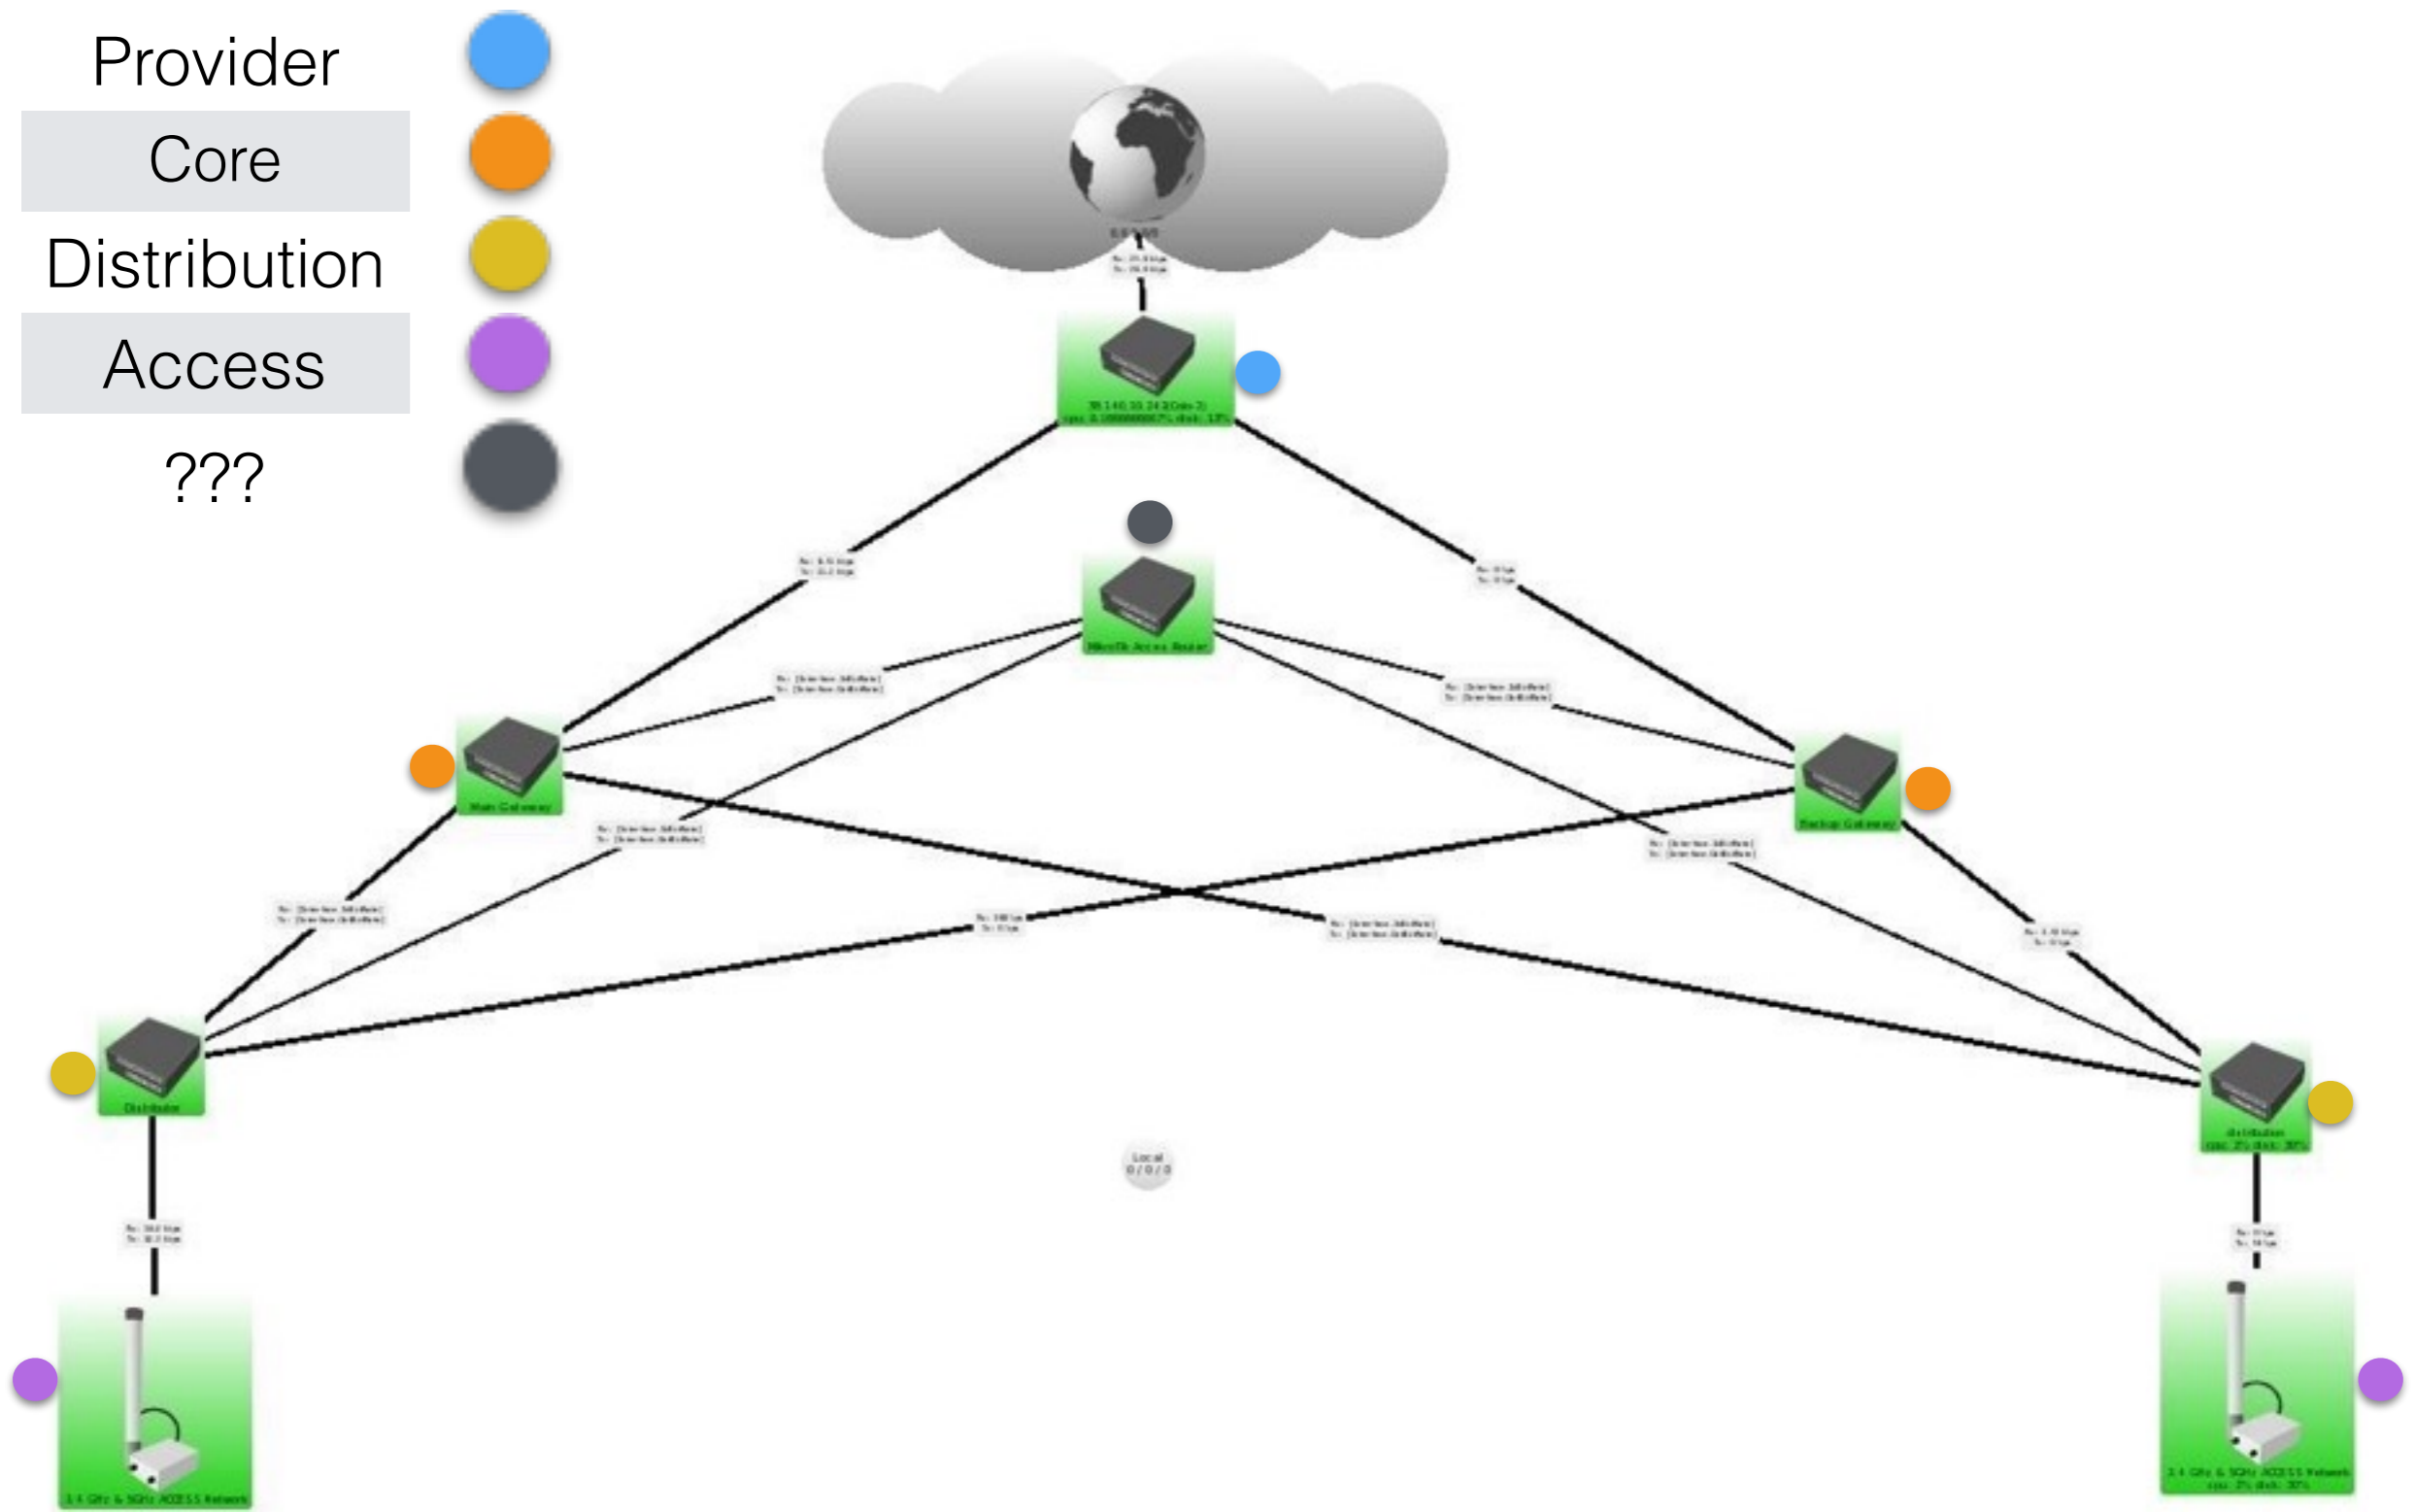
Distribution



Access



???




Multi-homing Definition

- Multi-link “ External” to local network
 - **Two or more link to the same ISP**
 - Multiple Interfaces, Single IP address per interface
 - Two or more link to different ISPs

Best path selection

- Lowest MED (default 0)
- Prefer the route with lowest router ID or ORIGINATOR_ID




Two or more link to the same ISP

- **Multiple Links, Single IP address (space)**
- Multiple Links, Multiple IP address (spaces)

How to Multi-home

- **Basic Multihoming**
 - Multi-home with the same ISP



Basic Multihoming

Multi-homing to the same ISP

(One as backup only)

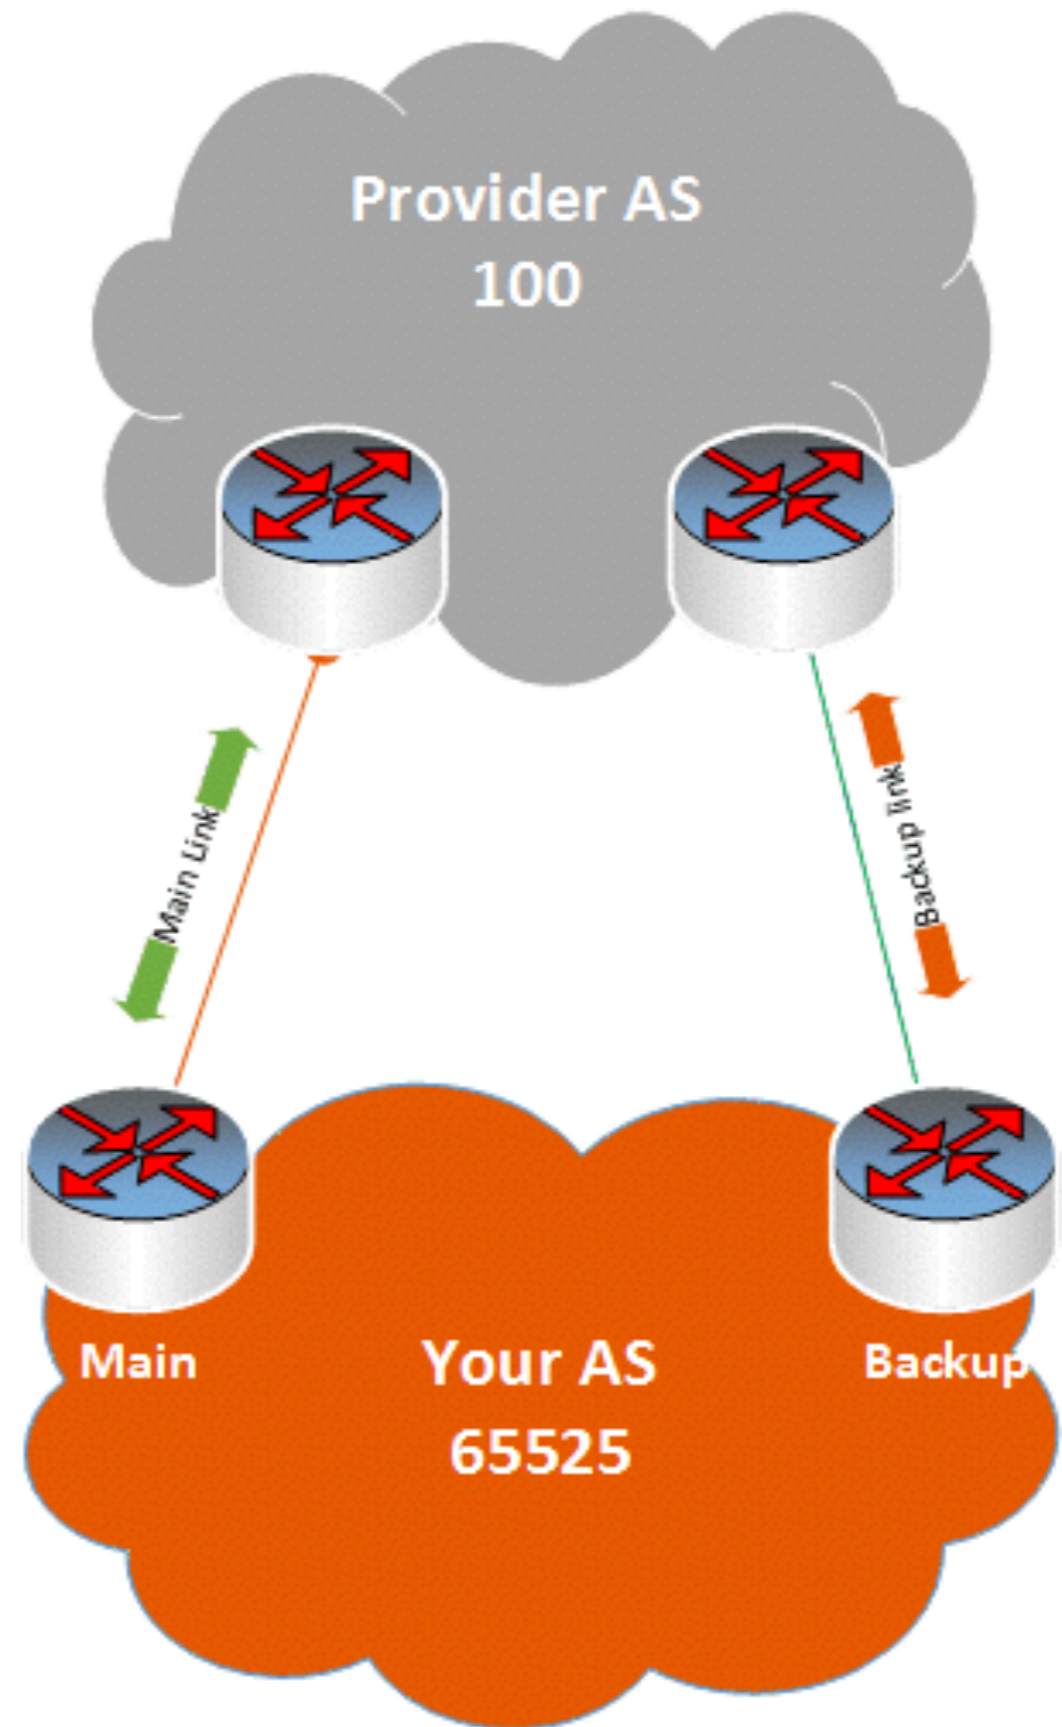
- Use private AS (AS>64511)
- There is no need for public ASN
- AS100 “the provider AS” removes private AS and any customer sub-prefixes from internet announcement

Two link to the same ISP

- Announce /24 aggregate on each link
 - Main link network: announce /24
 - Backup link network: announce /24 with increased metric
- When one link fails, the announcement of /24 aggregate via the backup link ensure continued connectivity.

Two link to the same ISP

Main Router Configuration



Two link to the same ISP

The image shows two overlapping windows from a network configuration interface. The left window is titled "BGP Instance <bgp-main>" and contains the following fields: Name: bgp-main, AS: 65521, Router ID: 1.1.1.1. Below these are several checkboxes for redistribution: Redistribute Connected, Redistribute Static, Redistribute RIP, Redistribute OSPF, and Redistribute Other BGP. The right window is titled "BGP Peer <peer-main>" and contains the following fields: Name: peer-main, Instance: bgp-main, Remote Address: <Provider router-ip-address>, Remote Port: (dropdown), Remote AS: <provider-ASN>, TCP MD5 Key: (dropdown), Nexthop Choice: default, with checkboxes for Multihop and Route Reflect. At the bottom, there are fields for Hold Time: 180 s and Keepalive Time: (dropdown). Three red circles with numbers 1, 2, and 3 are overlaid on the image. Circle 1 is around the "Routing" menu item in the top-left corner. Circle 2 is around the "Name" field in the BGP Instance window. Circle 3 is around the "Name" field in the BGP Peer window.

Routing

System

BFD

BGP

BGP

Instances VRFs Peers Networks Aggregates VPN4 Routes Ad

BGP Instance <bgp-main>

Name: bgp-main

AS: 65521

Router ID: 1.1.1.1

Redistribute Connected

Redistribute Static

Redistribute RIP

Redistribute OSPF

Redistribute Other BGP

BGP Peer <peer-main>

Name: peer-main

Instance: bgp-main

Remote Address: <Provider router-ip-address>

Remote Port: (dropdown)

Remote AS: <provider-ASN>

TCP MD5 Key: (dropdown)

Nexthop Choice: default

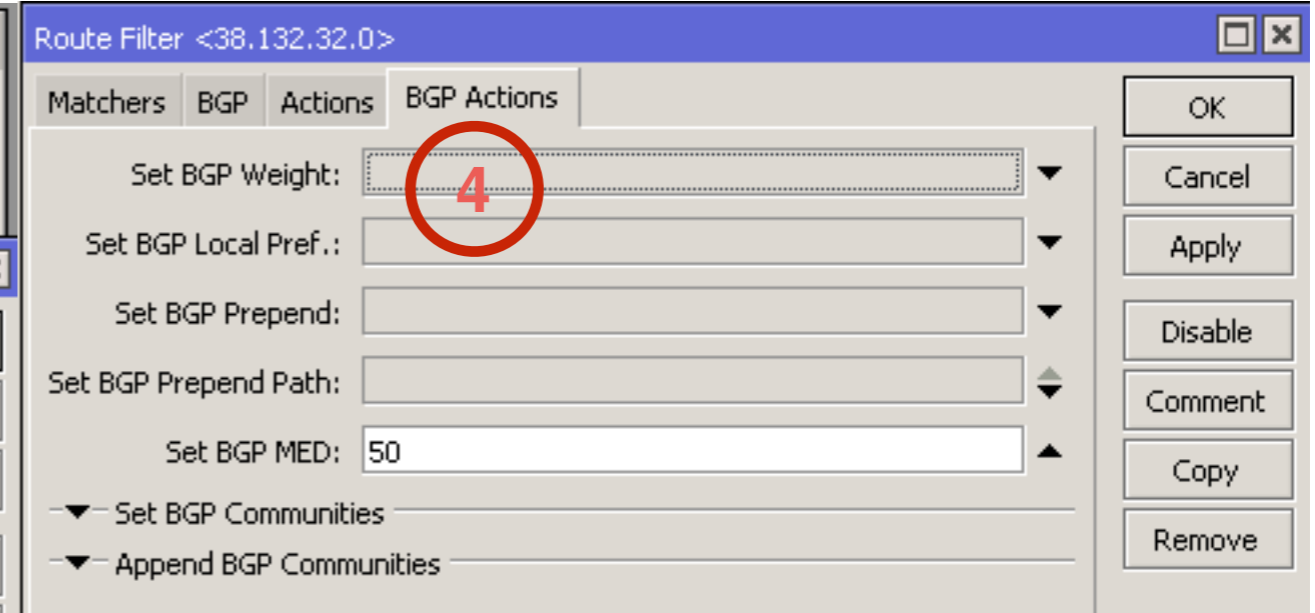

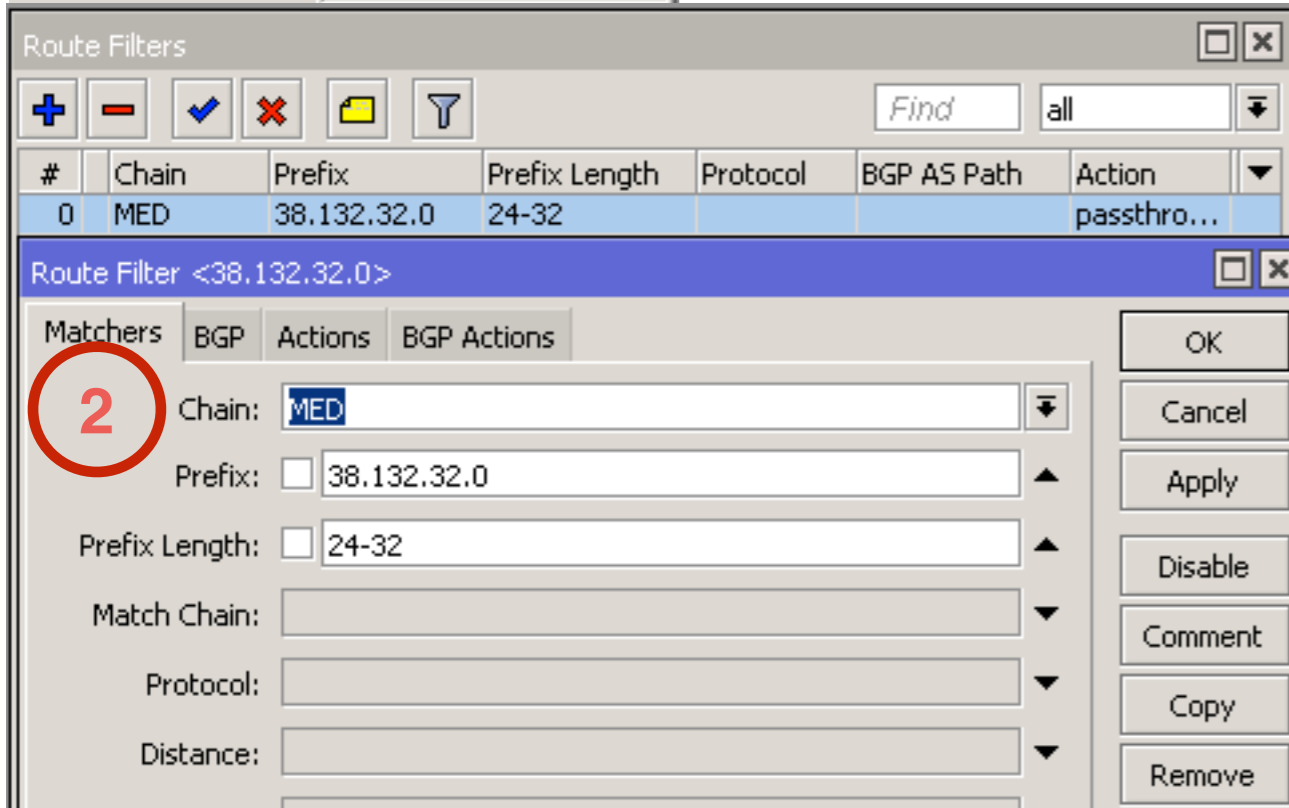

Multihop

Route Reflect

Hold Time: 180 s

Keepalive Time: (dropdown)

Two link to the same ISP



- Lower metric is preferred
- Exchanged between AS and used to make decision inside that AS, not passed to third AS.
- Ignored if received from different ASs

Two link to the same ISP

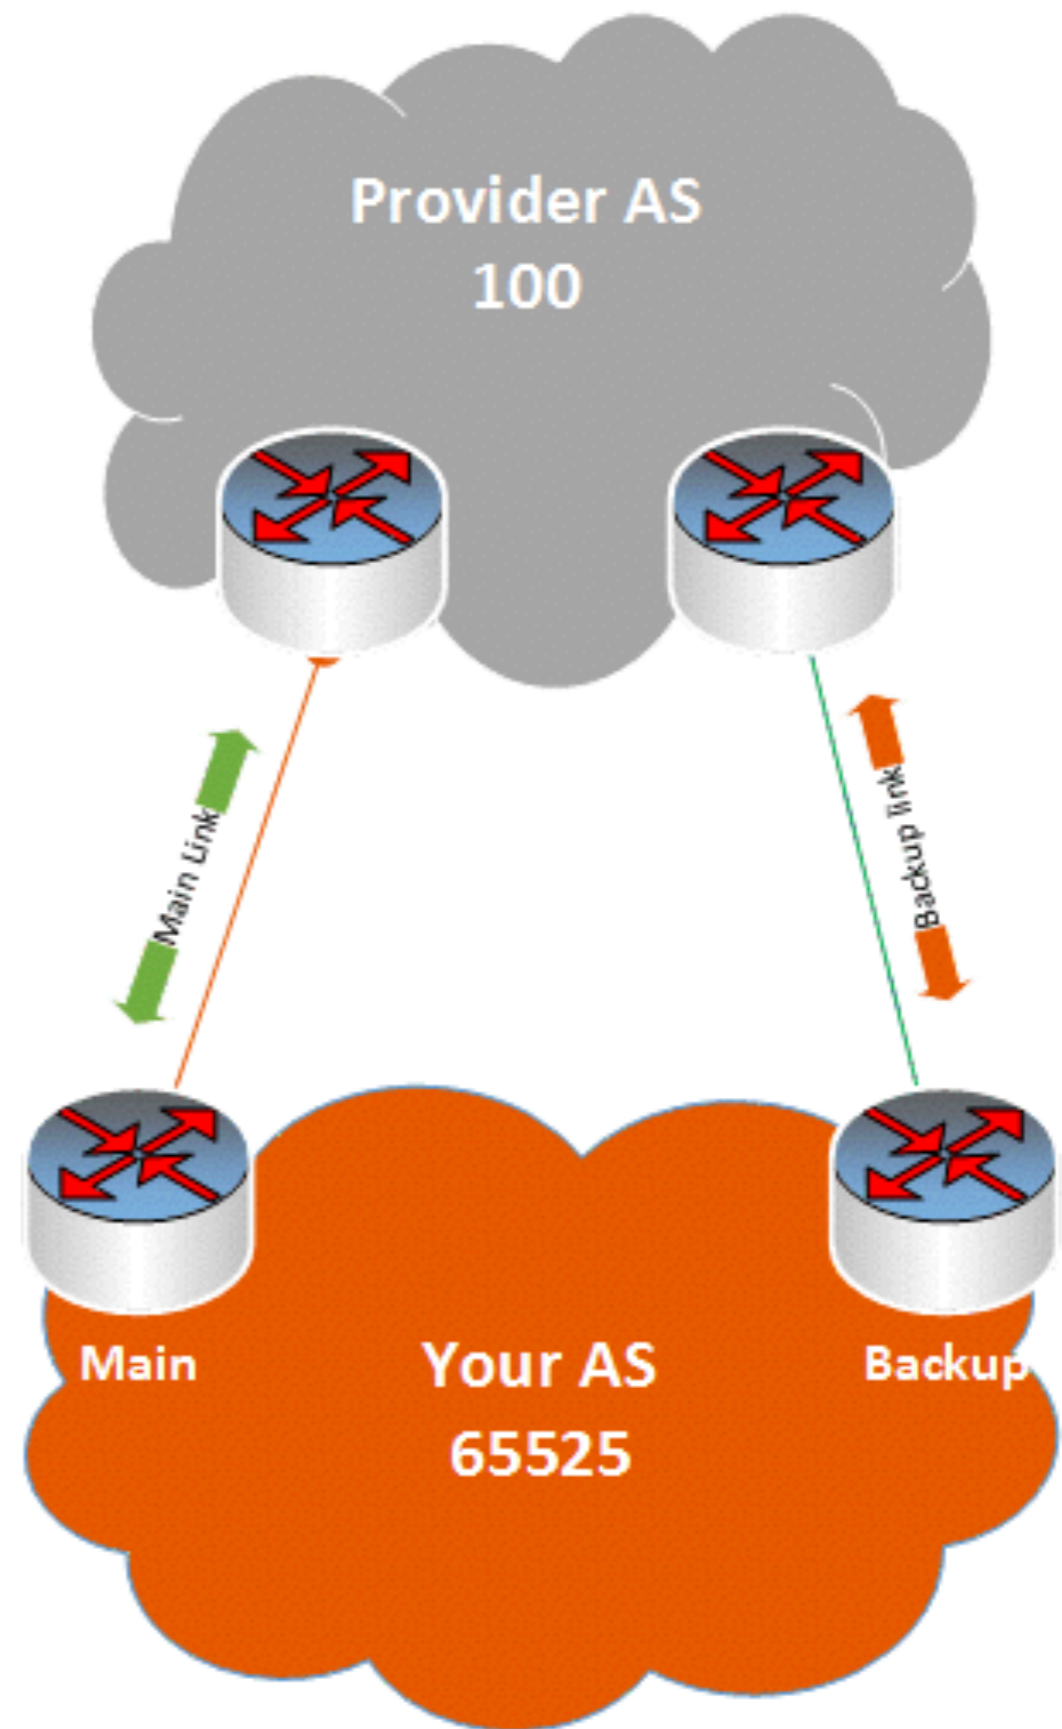
Used to hint an external neighbour about path preference into an AS

The image shows a network configuration interface with the following elements:

- Routing Menu:** A red circle labeled '1' highlights the 'Routing' menu.
- BGP Tab:** A red circle labeled '2' highlights the 'Peers' tab in the BGP configuration area.
- BGP Peer Configuration:** The 'BGP Peer <peer-main>' window is open, showing the 'General' tab. The 'Out Filter' field is set to 'MED', which is circled in red with a '3'. Other fields include Name (peer-main), Instance (bgp-main), Remote Address (<Provider IP Address>), Remote AS (<Provider ASN>), TCP MD5 Key, Nexthop Choice (default), Hold Time (180 s), Keepalive Time, TTL (default), Max Prefix Limit, Max Prefix Restart Time, AllowAS In, and Default Originate (never).

Two link to the same ISP

Backup Router Configuration



Two link to the same ISP

The image shows a screenshot of a network configuration interface with three red circles highlighting specific elements:


- 1**: A red circle around the "Routing" menu item in the top-left navigation pane.
- 2**: A red circle around the "Name" field in the "BGP Instance <bgp-backup>" configuration dialog, which contains the text "bgp-backup".
- 3**: A red circle around the "Name" field in the "BGP Peer <peer-colo>" configuration dialog, which contains the text "peer-colo".

The interface also shows a table of BGP instances and a list of BGP peers.

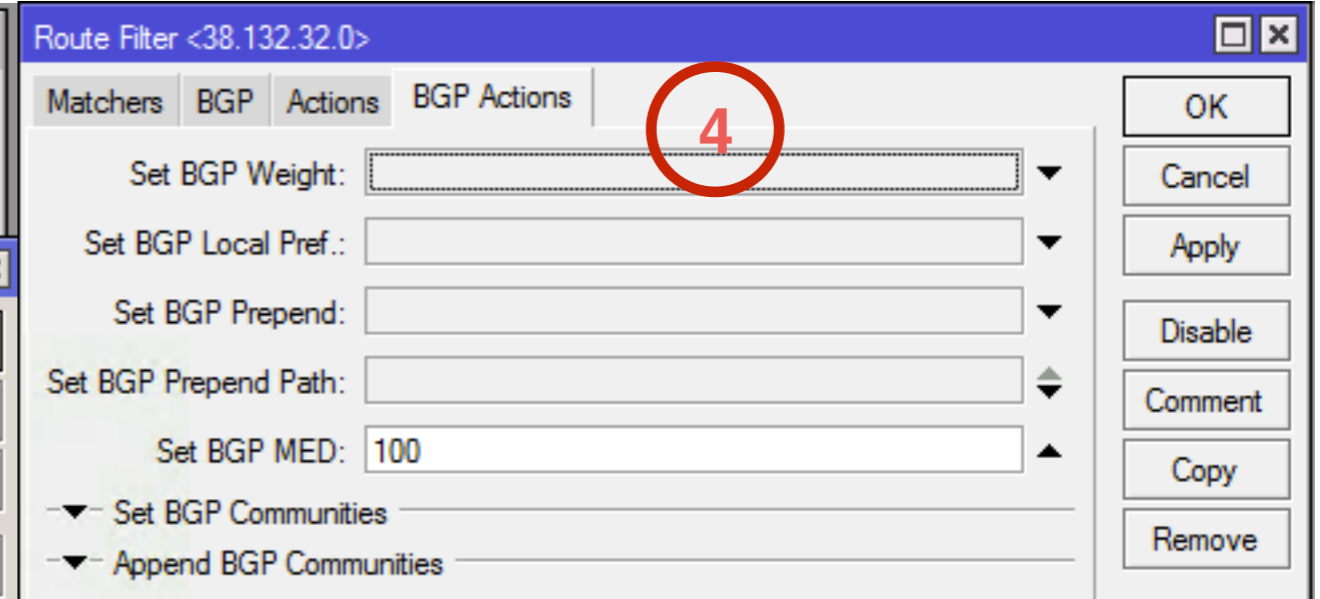
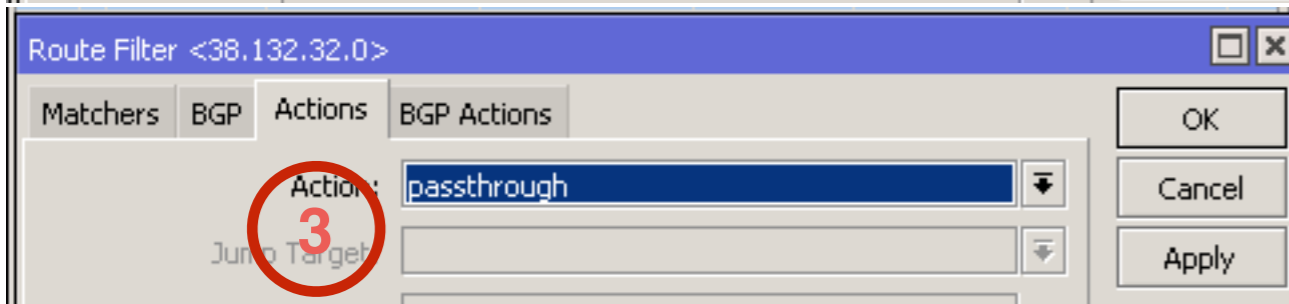
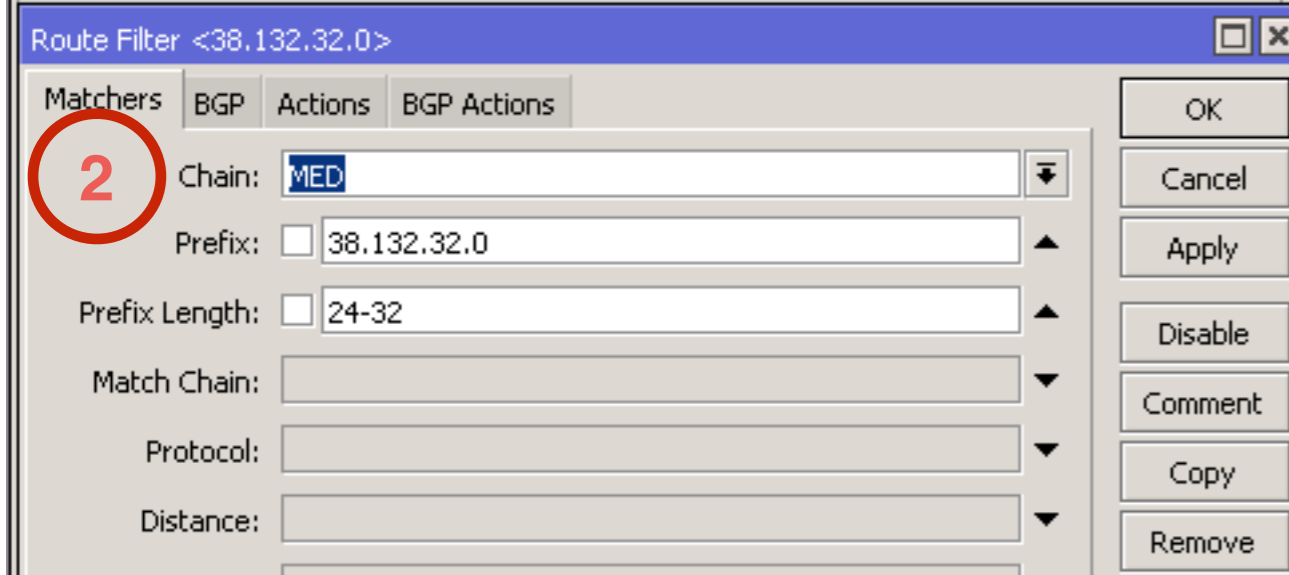
Name	AS	Router ID	Out Filter	Confeder...
bgp-backup	65521	1.1.1.1		

Name
peer-colo

Two link to the same ISP



#	Chain	Prefix	Prefix Length	Protocol	BGP AS Path	Action
0	MED	38.132.32.0	24-32			passthro...



- Lower metric is preferred
- Exchanged between AS and used to make decision inside that AS, not passed to third AS.
- Ignored if received from different ASs

Two link to the same ISP

Used to hint an external neighbour about path preference into an AS

The image shows a network configuration interface with the following elements:

- Routing System:** A sidebar menu with 'Routing' and 'System' options. 'Routing' is circled in red with the number '1'.
- BGP Configuration:** A sub-menu with 'Instances', 'VRFs', 'Peers', and 'Networks'. 'Peers' is circled in red with the number '2'.
- BGP Peer Configuration:** A window titled 'BGP Peer <peer-main>' with the following fields:
 - Name: peer-main
 - Instance: bgp-main
 - Remote Address: <Provider IP Address>
 - Remote Port: (empty)
 - Remote AS: <Provider ASN>
 - TCP MD5 Key: (empty)
 - Nexthop Choice: default
 - Hold Time: 180 s
 - Keepalive Time: (empty)
 - TTL: default
 - Max Prefix Limit: (empty)
 - Max Prefix Restart Time: (empty)
 - In Filter: (empty)
 - Out Filter: MED (circled in red with the number '3')
 - AllowAS In: (empty)
 - Default Originate: never

D- Dynamic
A- Active
b- BGP

MED=50

Main interface


```
[h-rahrouh@24/7Systems] /ip route> print detail where dst-address=38.132.32.0/24
Flags: X - disabled, A - active, D - dynamic, C - connect, S - static, r - rip, b - bgp, o - ospf, m - mme,
B - blackhole, U - unreachable, P - prohibited
0 ADb  dst-address=38.132.32.0/24 gateway=172.28.1.2 gateway-status=172.28.1.2 reachable via bridge-main distance=20 scope=40 target-scope=10
      bgp-as-path="65521" bgp-med=50 bgp-origin=igp received-from=peer2-mum-main
1 Db  dst-address=38.132.32.0/24 gateway=172.29.1.2 gateway-status=172.29.1.2 reachable via bridge-backup distance=20 scope=40 target-scope=10
      bgp-as-path="65521" bgp-med=100 bgp-origin=igp received-from=peer3-mum-backup
[h-rahrouh@24/7Systems] /ip route>
```

D- Dynamic
b- BGP

Backup interface

VRRP

Enable VRRP configuration
on Core layer network



VRRP Setup on main Router

The image shows a sequence of configuration steps for VRRP on a main router. The steps are numbered 1 through 5:

- 1**: In the **Interface List**, the **vrrp-main** interface is selected. The **Interface** column shows **bridge-dist**.
- 2**: The **Interface <vrrp-main>** configuration window is open. The **Interface** dropdown is set to **bridge-dist**. Other settings include **VRID: 50**, **Priority: 254**, **Interval: 1.00 s**, and **Preemption Mode** checked. The **Authentication** is set to **none**, and the **V3 Protocol** is **IPv4**.
- 3**: The **Address List** window is open, showing the **Address List** table with columns **Address**, **Network**, and **Interface**.
- 4**: The **Address <192.168.168.250/24>** configuration window is open. The **Address** is **192.168.168.250/24**, the **Network** is **192.168.168.0**, and the **Interface** is **bridge-dist**.
- 5**: The **Address <192.168.168.254>** configuration window is open. The **Address** is **192.168.168.254**, the **Network** is **192.168.168.254**, and the **Interface** is **vrrp-main**.

VRRP Setup on backup Router

The screenshot illustrates the VRRP configuration process on a backup router. It shows the following steps:

- 1**: Selecting the `vmp-backup` VRRP instance in the `Interface List`.
- 2**: Configuring the VRRP instance for the `bridge-dist` interface, setting the `VRID` to `50`, `Priority` to `250`, and `Interval` to `1.00` seconds. The `Preemption Mode` is checked.
- 3**: Configuring the address `192.168.168.254` on the `vmp-backup` interface.
- 4**: Configuring the address `192.168.168.251/24` on the `bridge-dist` interface.

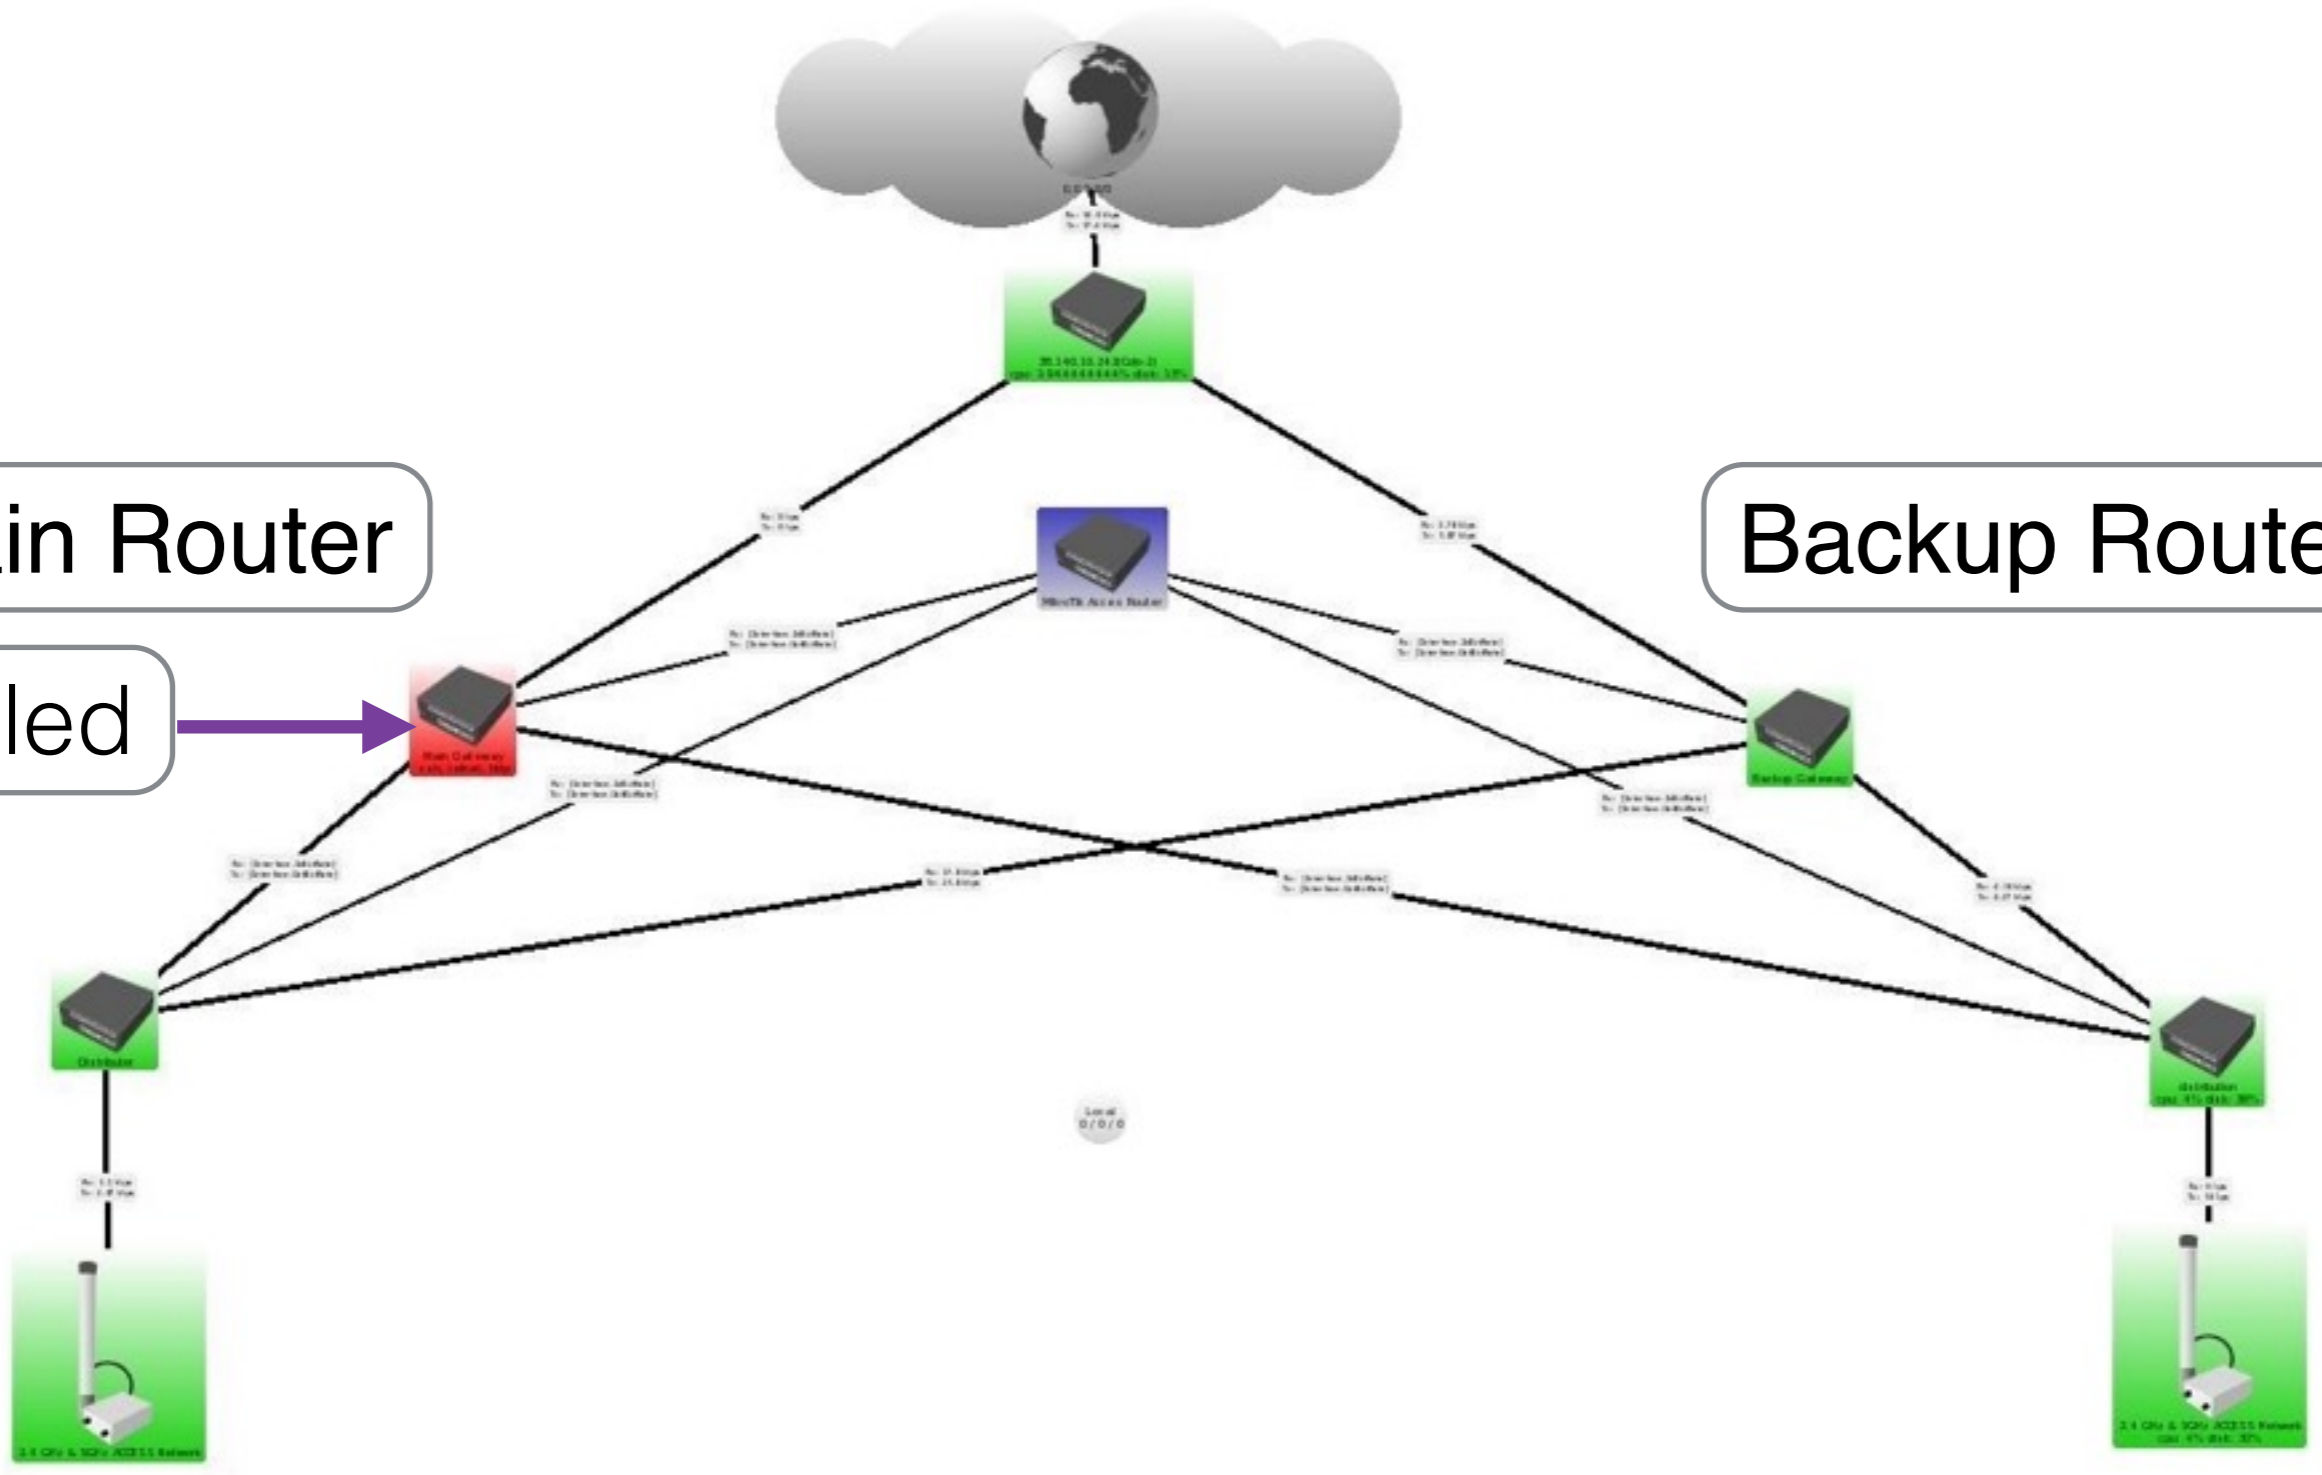
Address	Network	Interface
10.1.1.1	10.1.1.2	ether5
38.132.32.5/30	38.132.32.4	ether10
38.132.32.32/24	38.132.32.0	loopback
172.29.1.2/28	172.29.1.0	ether1
192.168.168.251/24	192.168.168.0	bridge-dist
192.168.168.254	192.168.168.254	vmp-backup



Main Router

Backup Router

Failed



Problems


- Router crash!
- Interface failure

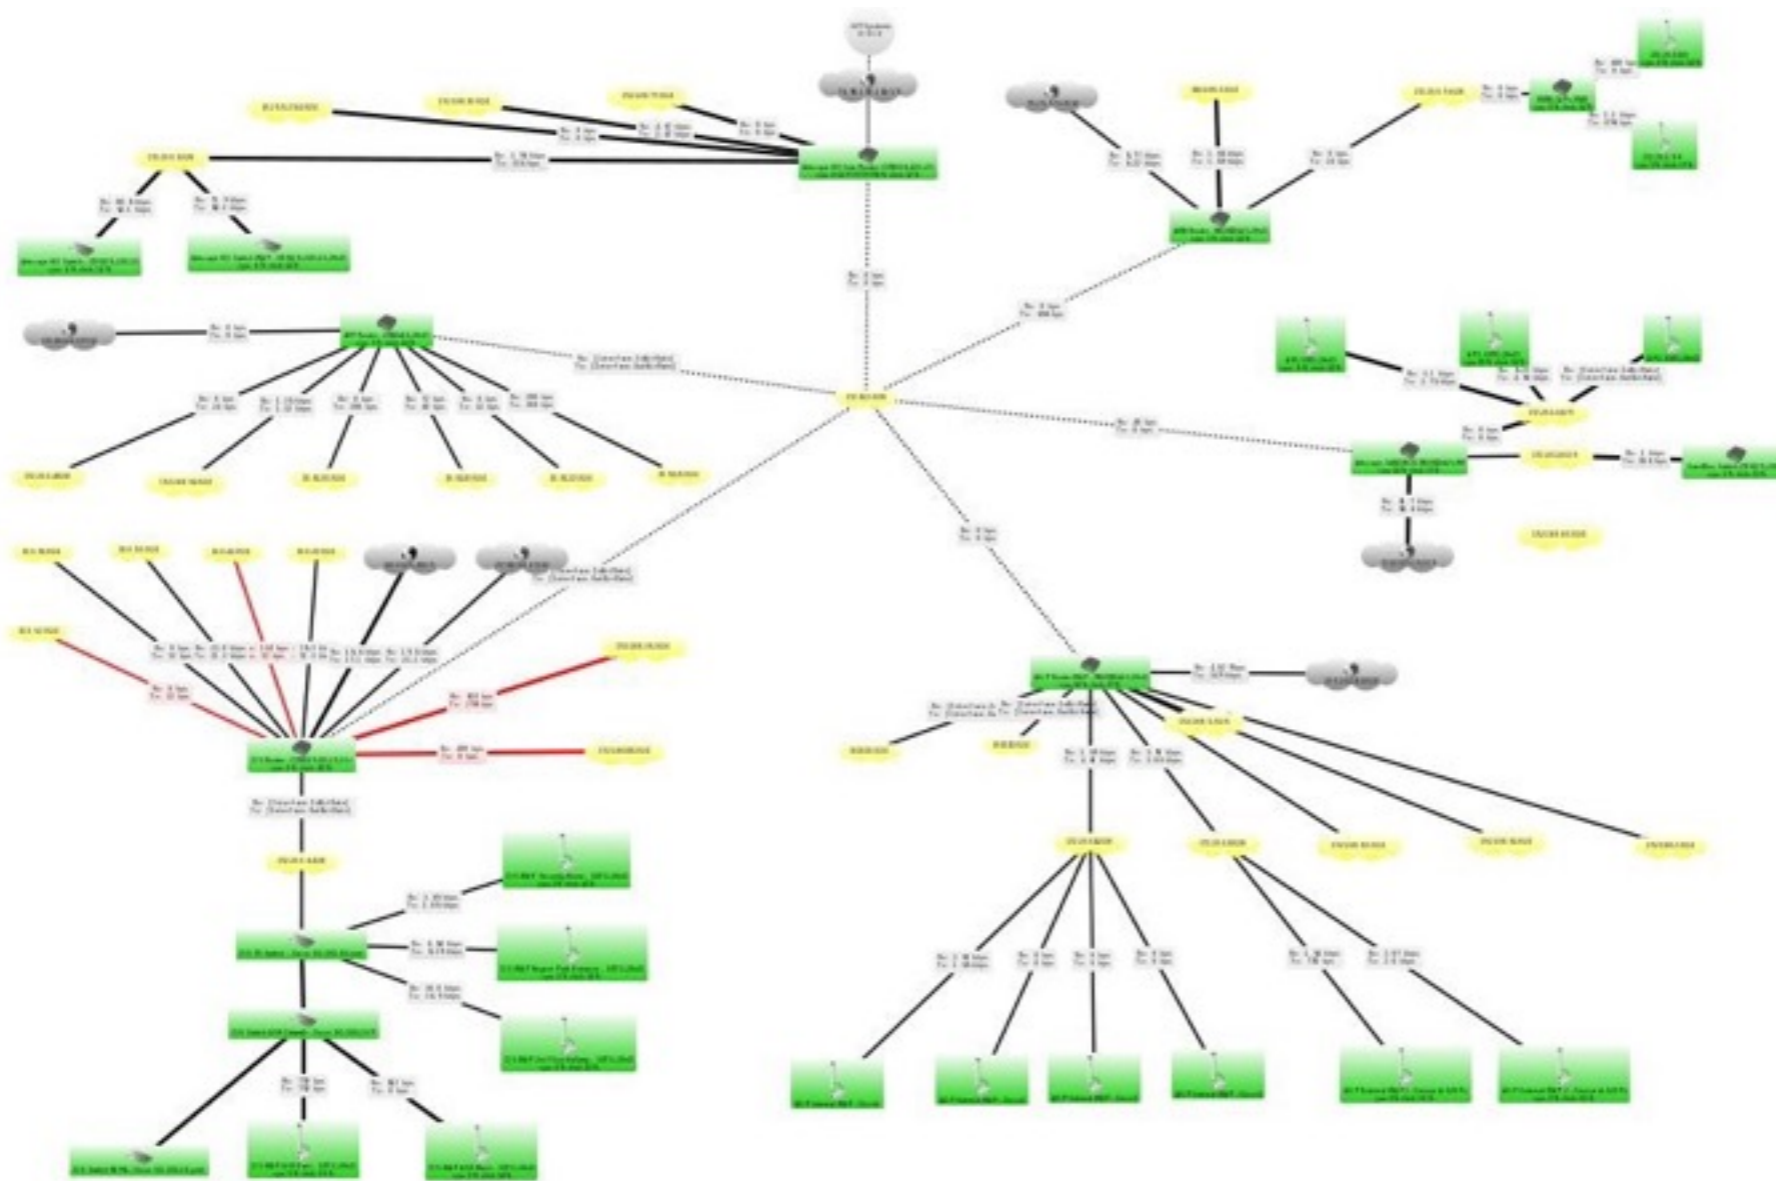
Solutions

Monitoring

- E-mail notification
- SMS

Access Router





The-Dude

The Dude is **free of charge!**

Question?

Enjoy!