MikroTikSA



NAPAFRICA

TERACO™
AFRICA'S DATA CENTRE

Tech Day Training

Online Edition

MikroTikSA

Housekeeping and Schedule

- Course materials
- Restrooms and smoking area locations

09:10
- BGP Route Filtering (45 minutes)
- Route Filtering Live Demo (15 minutes)
- Route Filtering Q&A (10 minutes)
- *Break (10 minutes)*

10:30
- BGP Communities (45 minutes)
- BGP Communities Live Demo (15 minutes)
- BGP Communities Q&A (15 minutes)
- *Break (15 minutes)*

12:00
- Troubleshooting Tools and Utilities (30 minutes)
- Multihoming / Load balancing (20 minutes)
- API simulation / demonstration (10 minutes)

About MikroTik SA

- Independent Networking Specialist company
- Not owned by / affiliated to MikroTik Latvia
- Official training and support partner for MikroTik
- Specialist in all forms of wireless and wired networking technologies
- Offers high speed PTP links, carrier independent backbone services, high availability SLA's
- David Savage
  - Is a MikroTik Certified Trainer and consultant for the last 12 years
  - Has over 30 years experience in the IT field
  - Teaches general networking and MikroTik RouterOS

MikroTik SA

Lockdown Services

MikroTikSA offers a number of services during lockdown

Group Training

We will shortly be releasing a number of online live training classes for those wanting to begin or continue on their MikroTik Training career. This is intended as an adjunct to instructor led training post lockdown. These will be offered at a number of different levels to suit any skill type.

HOT Labs

Hands On Training Labs are designed to target specific areas of MikroTik configuration which may not ordinarily be dealt with in traditional training. This includes subjects such as scripting, API, Advanced Firewall and a number of other advanced subjects.

©2006-2013 MikroTikSA Pty Ltd

Lockdown Services

MikroTikSA offers a number of services during lockdown

One-on-One

These are sessions for a single company who wants to target a specific subject or set of subjects. Sessions may be booked in 1 hour or 1/2 day segments.

Network Health Check Subscription

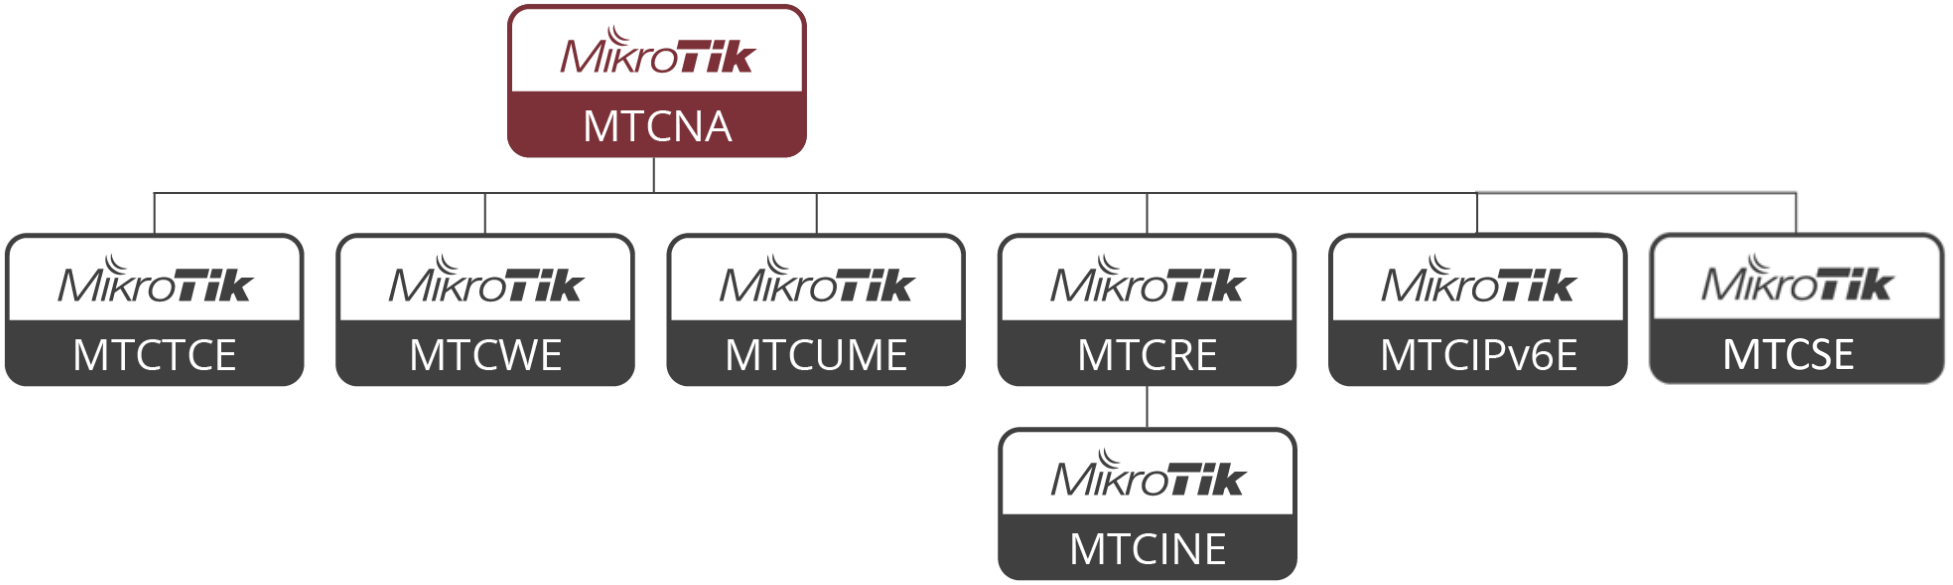
Similar to One-on-One, these interventions are designed to be ongoing as a regular monthly session. During this time we go through your network and confirm that it has been optimized in the most efficient way. Subscription Services are more cost effective than single sessions.

For more information contact us on training@mikrotiksa.com or discuss in person call / WhatsApp David on 083 456 2002

©2006-2013 MikroTikSA Pty Ltd

MikroTik Certification Path

- The MikroTik Certified Network Associate is the foundation course towards any engineer certification
- Once complete you can attend any engineer certification
- The MTCINE requires the MTCRE certification



©2006-2013 MikroTikSA Pty Ltd

MikroTik Certified Training Programs

- **MTCNA** - MikroTik Certified Network Associate (view outline)
- **MTCRE** - MikroTik Certified Routing Engineer (view outline)
- **MTCWE** - MikroTik Certified Wireless Engineer (view outline)
- **MTCTCE** - MikroTik Certified Traffic Control Engineer (view outline)
- **MTCUME** - MikroTik Certified User Management Engineer (view outline)
- **MTCINE** - MikroTik Certified Inter-Networking Engineer (view outline)
- **MTCIPv6E** - MikroTik Certified IPv6 Engineer (view outline)
- **MTCSE** - MikroTik Certified Security Engineer (view outline)

©2006-2013 MikroTikSA Pty Ltd

MikroTik Hardware

- MikroTik manufacture a range of devices for switching and routing
- Devices are based on different hardware platforms
- RouterOS software is consistent across all hardware devices
- Understanding which hardware is appropriate to your application is critical for long term network performance
- Broadly speaking, devices are broken up into switching vs routing roles
- However, it is possible to simulate any kind of network device role with enough configuration
- However – just because you can do something does not mean you should…

©2006-2013 MikroTikSA Pty Ltd

CPE Devices

- Consumer level devices designed to be used at client premises
- Usually features 3-5 ether ports and 1-2 wireless interfaces, some also have SFP ports
- Cost effective means of providing a wireless and ethernet network for small setups
- Can be extended wirelessly or with ether for larger setups



HAP AC2



CAP AC



RB4011



WAP AC

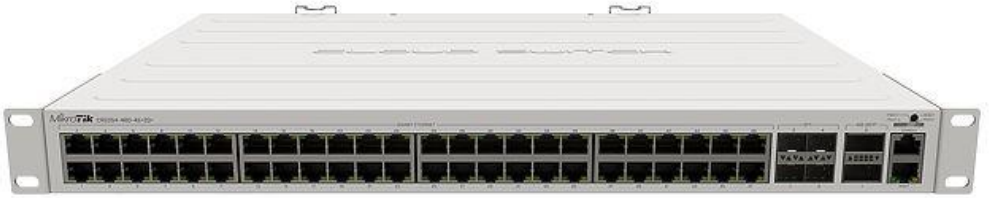
©2006-2013 MikroTikSA Pty Ltd

Switches

- Most multi-port devices have a switch chip and can be used for hardware level switching and VLANs where required
- Care must be taken to ensure the router does not resort to software based switching configuration
- The CRS Cloud Router Switch series comprises devices designed for switching to suit virtually any speed requirement up to 40gbps per port
- They have very limited routing capabilities and are not designed to be used in this role



CRS328-24P-4S+RM

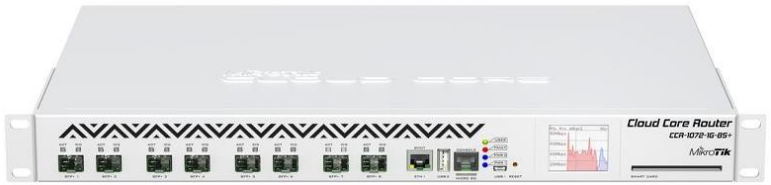


CRS354-48G-4S+2Q+RM

©2006-2013 MikroTikSA Pty Ltd

Routers

- The CCR Cloud Core Routers are dedicated router devices
- They are based on the Tilera multicore CPU platform supporting between 9 – 72 CPU cores
- They usually feature multiple gigabit and 10g interfaces
- They do not have switch chips and as such are not recommended for switching or bridged VLAN configurations
- At the lower end, the RB1100AX4, RB4011/3011 and HEX series routers provide cost effective routing and switching functions in one device



CCR1072-1G-8S+



RB1100AHX4-Dude



HEX-S

©2006-2013 MikroTikSA Pty Ltd

Cloud Hosted Router

- Cloud Hosted Router (CHR) is a RouterOS version intended for running as a virtual machine
- It supports the x86 64-bit architecture and can be used on most of the popular hypervisors such as VMWare, Hyper-V, VirtualBox, KVM and others
- CHR has full RouterOS features enabled by default but has a different licensing model than other RouterOS versions
- Minimal requirements:
  - Package version: RouterOS v6.34 or newer
  - Host CPU: 64-bit with virtualization support
  - RAM: 128MB or more
  - Disk: 128MB disk space for the CHR virtual hard drive (Max: 16GB)

©2006-2013 MikroTikSA Pty Ltd

CHR

- **CHR has been tested on the following platforms:**
  - VirtualBox 5 on Linux and OS X
  - VMWare Fusion 7 and 8 on OS X
  - VMWare ESXi 6.5
  - Qemu 2.4.0.1 on Linux and OS X
  - Hyper-V on Windows Server 2008r2, 2012 and Windows 10
  - Xen Project 4.6.5
  - Xen Server 7.1

License	Speed limit	Price
Free	1Mbit	FREE
P1	1Gbit	$45
P10	10Gbit	$95
P-Unlimited	Unlimited	$250

©2006-2013 MikroTikSA Pty Ltd

CHR License

- **Paid licenses**
- **P1 -** perpetual-1 license level allows CHR to run indefinitely. It is limited to 1Gbps upload per interface.
- **P10 -** perpetual-10 license level allows CHR to run indefinitely. It is limited to 10Gbps upload per interface
- **p-unlimited -** The perpetual-unlimited license level allows CHR to run indefinitely. It has no enforced limitations.
- **Free -** The *free* license level allows CHR to run indefinitely. It is limited to 1Mbps upload per interface. This is the default option
- CHR is available as a pre-installed package on Amazon EC2
- It is eligible for the Free Tier service (750 hours free for the first year)
- Trainer will demonstrate spinning up a new instance

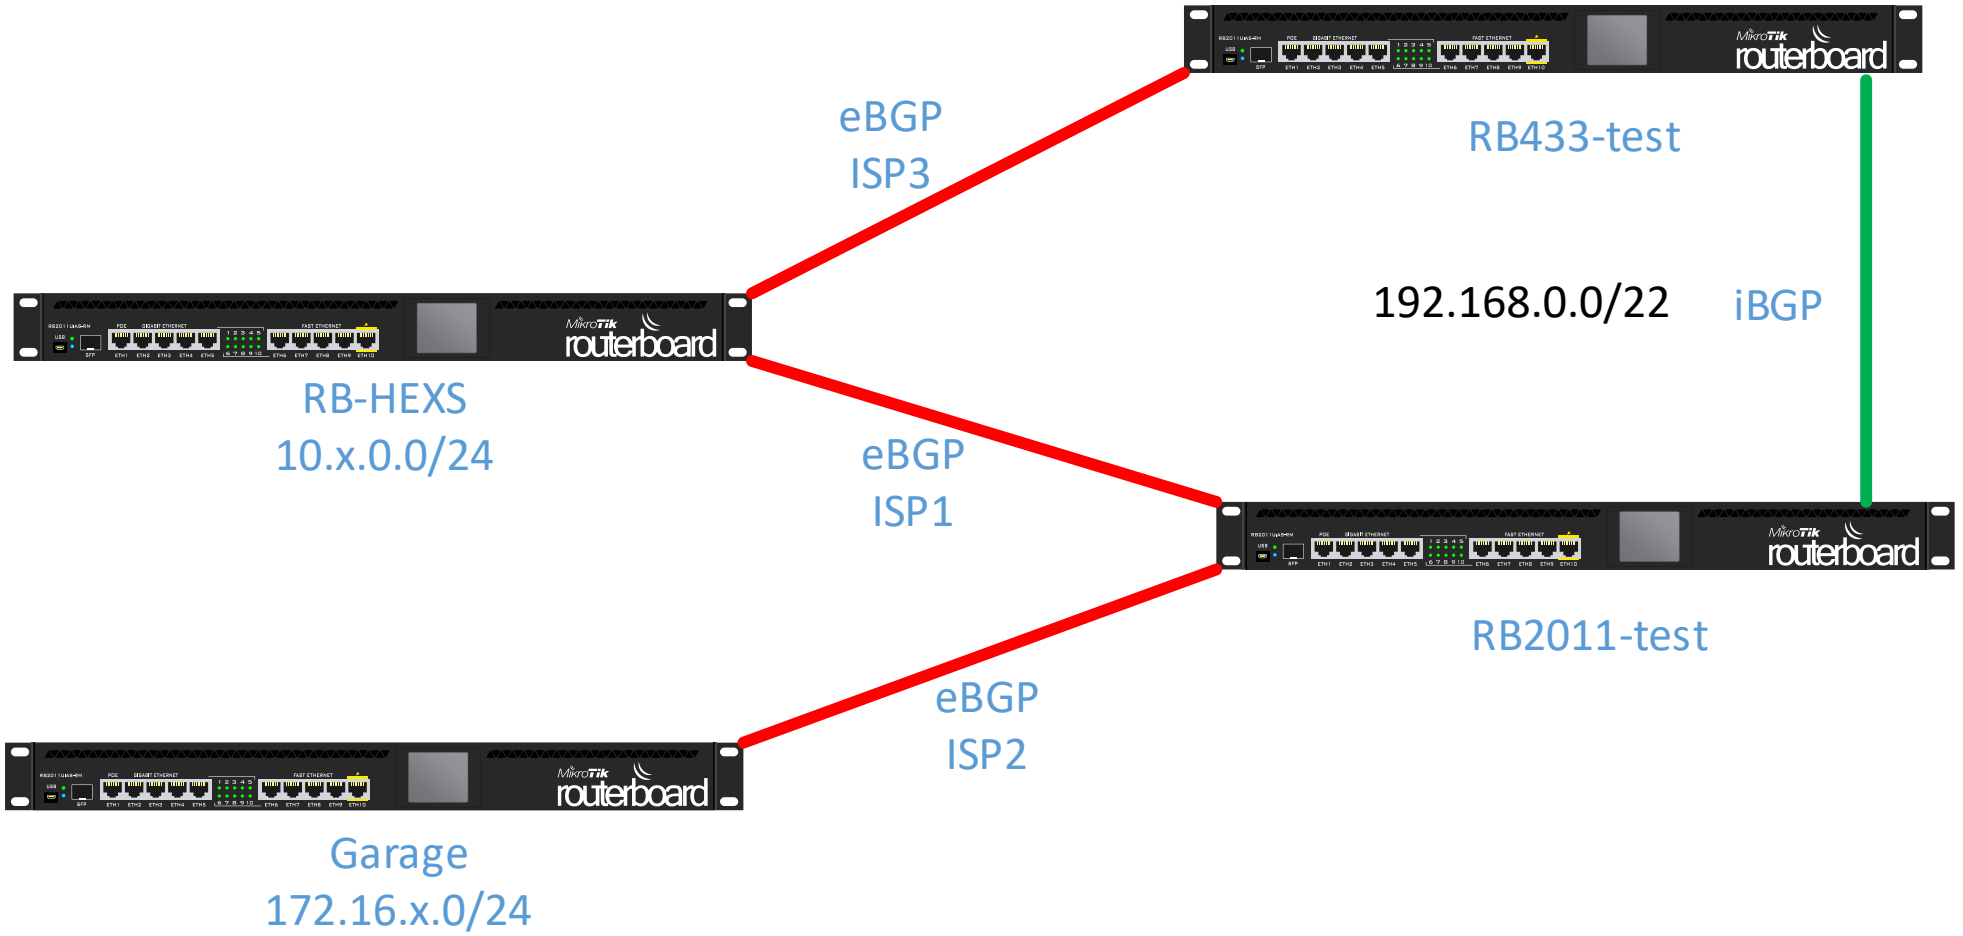
©2006-2013 MikroTikSA Pty Ltd

Interior and Exterior Peering

- Aside from building peering relationships and transporting data between two BGP speakers, two types of peering relationships exist within BGP: *interior peering* and *exterior peering*
- BGP sessions between peers within a single autonomous system are referred to as interior BGP, or iBGP
- BGP running between peers in different autonomous system are referred to as exterior BGP, or eBGP sessions
- Setting up iBGP or eBGP peers is the same process, the only difference is in the AS Number that is specified
  - iBGP occurs between 2 peers with the same AS Number
  - eBGP ASN's are different on each end of the peer
  - Routers will automatically determine if iBGP or eBGP should be used

©2006-2013 MikroTikSA Pty Ltd

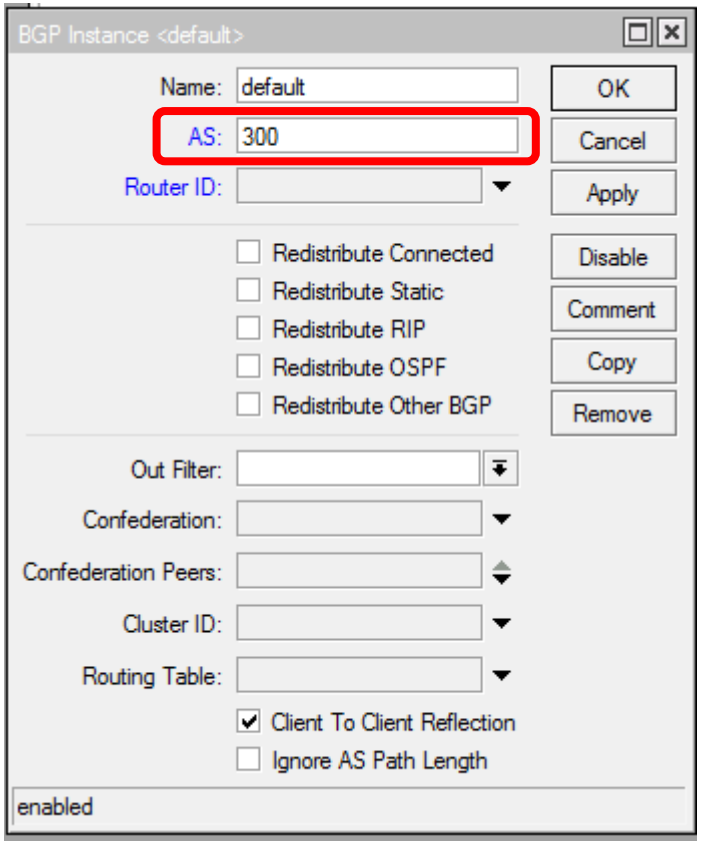
Lab Setup



©2006-2013 MikroTikSA Pty Ltd

Enabling eBGP

- Specify Router ID and AS Number in Routing → BGP → Instance → Default
  - If router-id is not specified, The lowest IP on the Router is Selected automatically as the Router-ID
  - For iBGP especially it is recommended to set a Router ID – usually determined from loopback
  - When using a loopback IGP (OSPF / Static) is required to keep track of loopback IP's



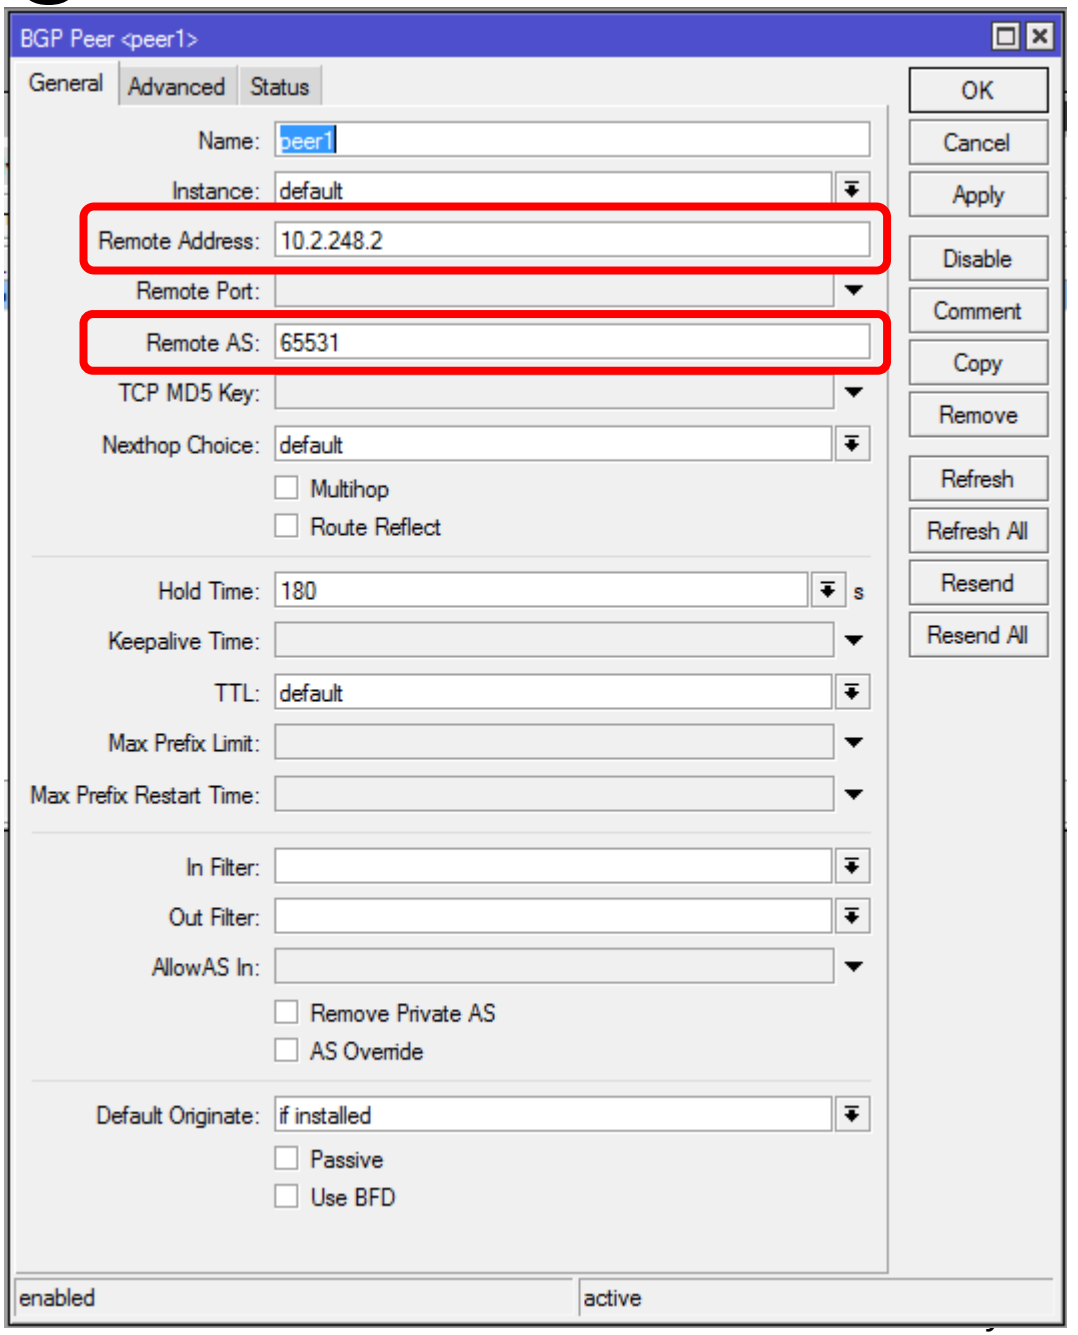
©2006-2013 MikroTikSA Pty Ltd

Defining Peers

- Add a peer in Routing → BGP → Peer
  - Specify the Remote Address and Remote AS of the peer
- To Verify BGP connectivity. Use The Command

```
/routing bgp-peer print
```

- Any state other than established indicates that routers have not become neighbours (use print status for more details)



©2006-2013 MikroTikSA Pty Ltd



- Peer Setup

BGP Routes and Formatting Data

- By default BGP will not advertise routes to peers until they are explicitly defined
- A BGP route is defined as a unit of information that pairs a set of destinations with the attributes of a path to those destinations
- The set of destinations is referred to by BGP as the Network Layer Reachability Information (NLRI), and is a set of systems whose IP addresses are represented by one IP prefix (aggregate / supernet)
- BGP uses update messages to advertise new routing information, withdraw previously advertised routes, or both
- New routing information includes a set of BGP attributes and one or more prefixes with which those attributes are associated

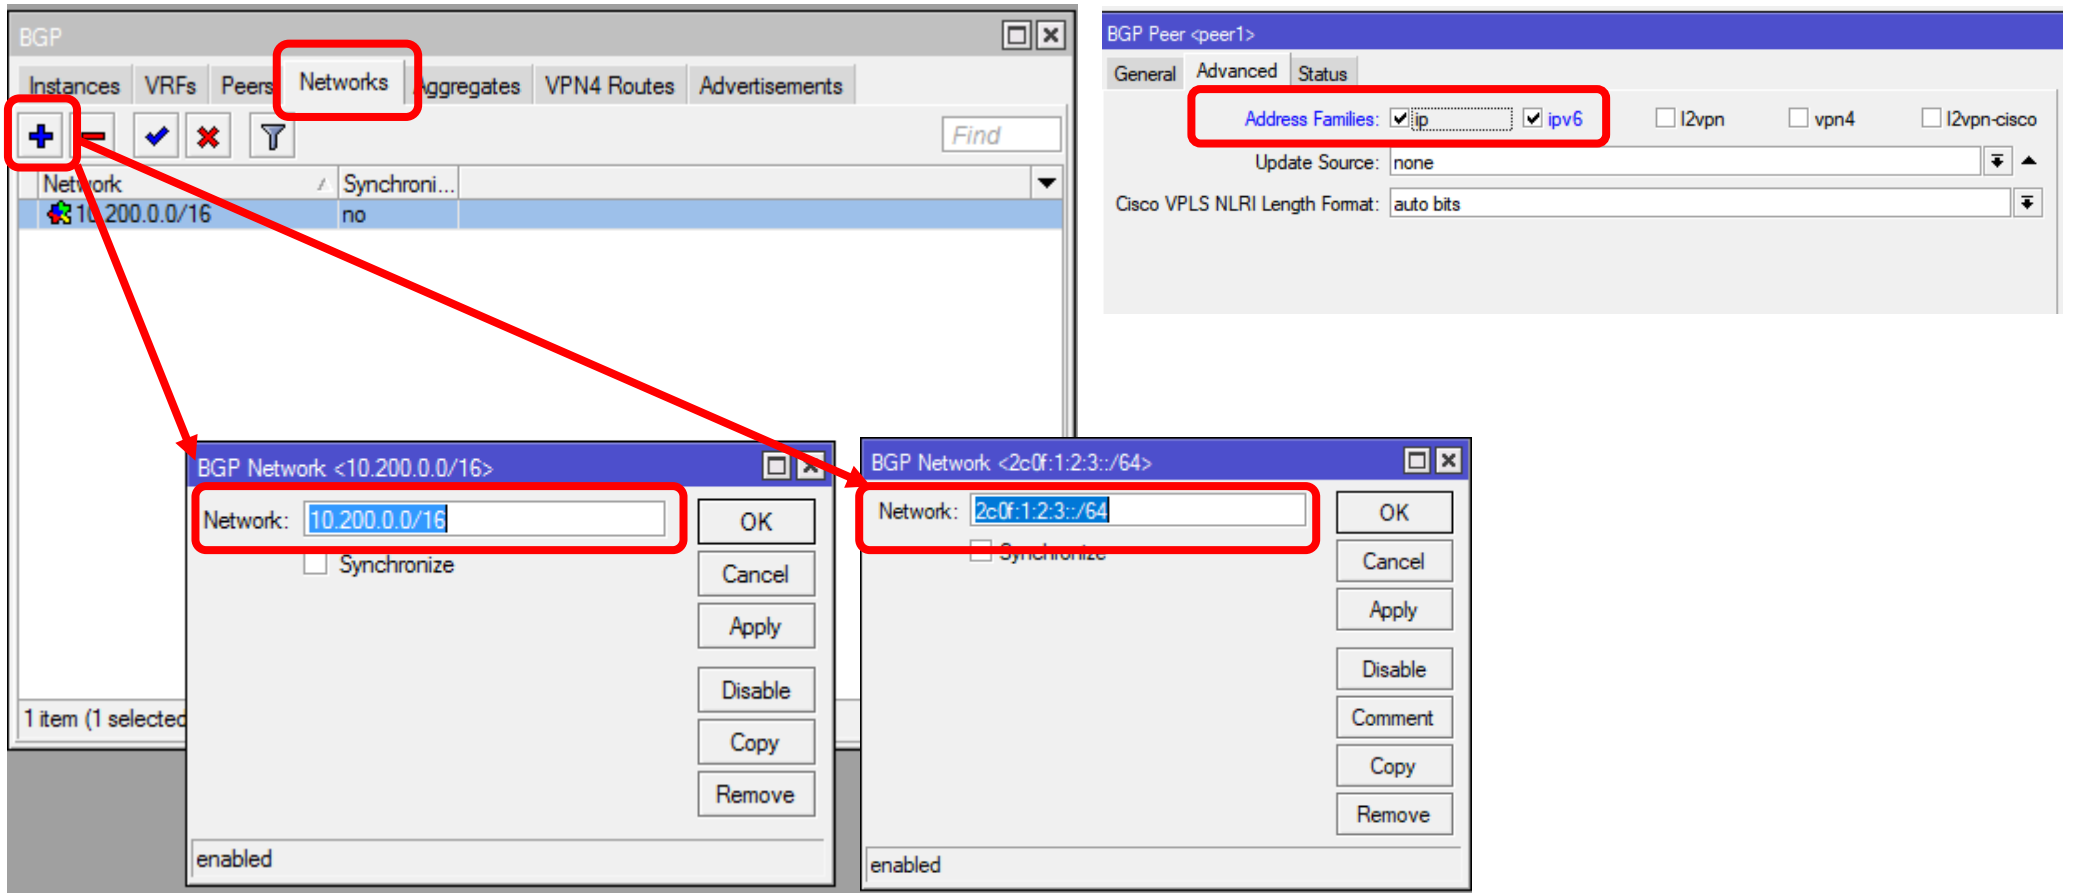
©2006-2013 MikroTikSA Pty Ltd

Networks

- Indicates what networks BGP should originate from the router.
- Both IPv4 and IPv6 networks are natively supported
  - However the peer must be setup to advertise (advanced tab)
- If Synchronise is enabled the network is advertised only if a corresponding route is present in routing table
- Synchronization should be turned off if:
  - Your AS does not provide transit service
  - All the transit routers run BGP
- Disabling sync allows BGP to converge faster.
- Sync can be dangerous if routes are flapping a lot.
- Configurable from **`/routing bgp network`**

©2006-2013 MikroTikSA Pty Ltd

BGP Networks



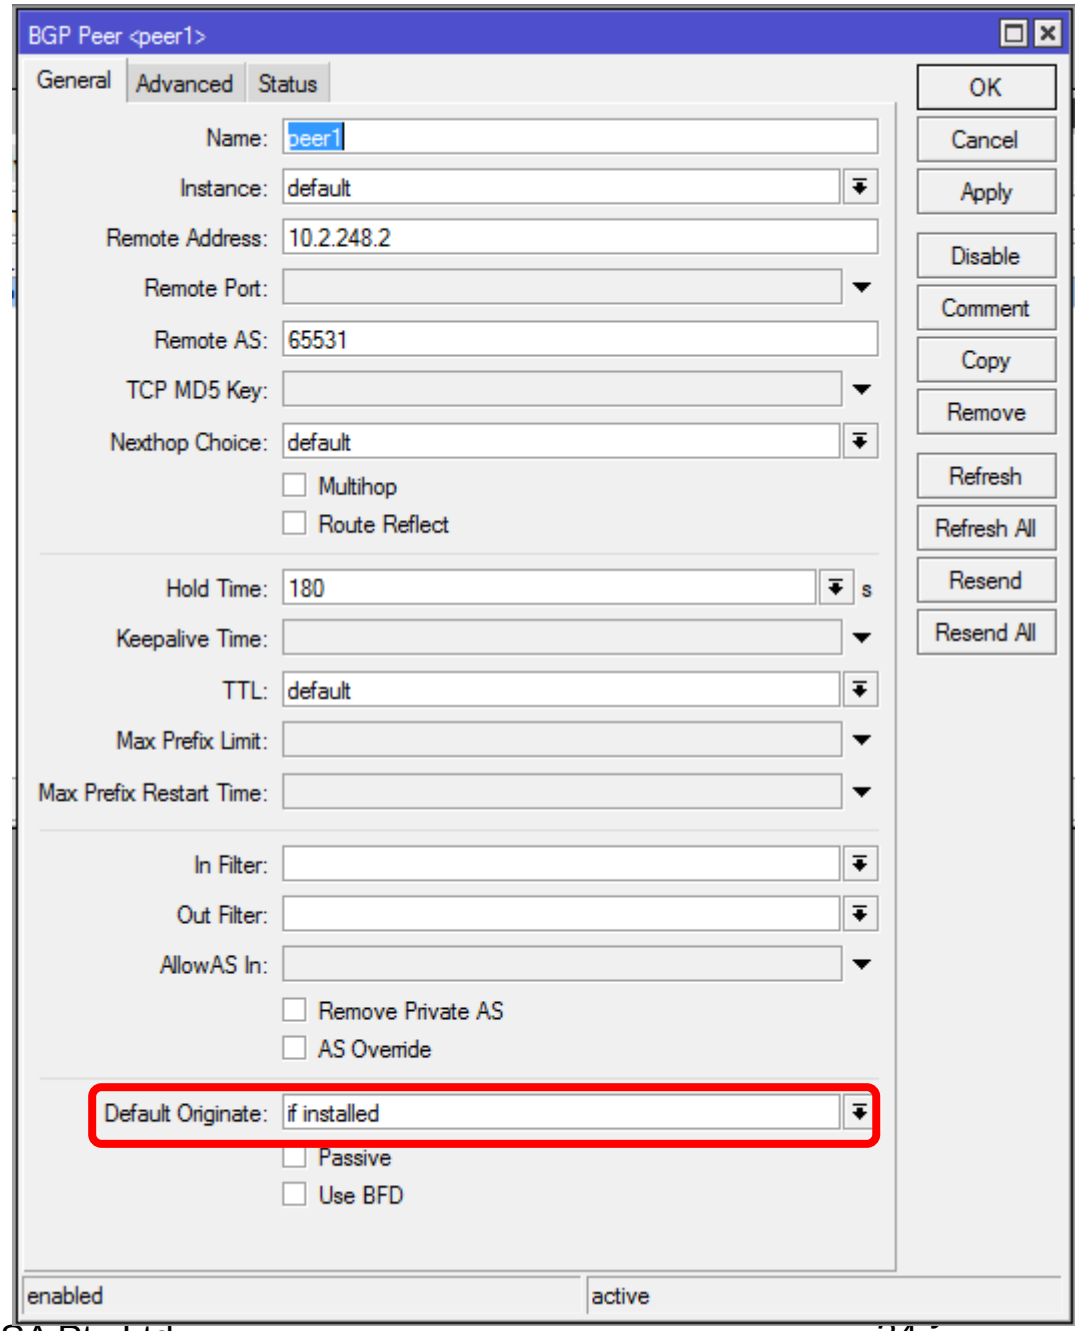
©2006-2013 MikroTikSA Pty Ltd



- Advertise Networks

Default Route

- In most cases a full routing table is not required
- The requirement is to advertise your public range out
- If you do not have multihomed international links you can ask for a default route instead
- Upstream ISP can originate a default route across your peering link



©2006-2013 MikroTikSA Pty Ltd

Default Route

- Default Originate: Send a default route to downstream peer
  - Never – do not send
  - If installed – only send if originator has default route
  - Always – always originate a default to the downstream
- BGP will not advertise the default route between peers transitively (i.e. it is not automatically distributed through the network as it would be with OSPF)
- For iBGP (internal) peers it is recommended to advertise as if-installed
- For eBGP advertised method will depend on inter AS setup (usually advertise always)
  - This is only a consideration when acting as a transit or upstream provider
  - For a single ISP it will not generally be used

©2006-2013 MikroTikSA Pty Ltd

LAB

- Default Routes
  - eBGP advertising – single upstream

Other Peer Options

- TCP MD5 Key: Key used to encrypt comms between peers
- Multihop: Used when eBGP peer is not directly connected (possible security risk – use with MD5 Key only)
  - Note that iBGP uses multihop by default (does not need enabling)
- Hold-time: Specifies the BGP Hold Time value to use when negotiating with peers.
  - According to the BGP specification, if router does not receive successive KEEPALIVE and/or UPDATE and/or NOTIFICATION messages within the period specified in the Hold Time field of the OPEN message, then the BGP connection to the peer will be closed.
  - The minimal hold-time value of both peers will be actually used
- Max-prefix-limit: Maximum number of prefixes to accept from a specific peer.
  - When this limit is exceeded, TCP connection between peers is closed.

©2006-2013 MikroTikSA Pty Ltd

Other Peer Options

- Max-prefix-restart-time: Minimum time interval after which peers can re-establish BGP session when max-prefix-limit is reached
- Allow AS in: How many times to allow own AS number in AS-PATH, before discarding a prefix
- AS Override: If set, then all instances of remote peer's AS number in BGP AS PATH attribute are replaced with local AS number before sending route update to that peer
  - Happens before routing filters and prepending.
- Remove Private AS: If set, then BGP AS-PATH attribute is removed before sending out route update if attribute contains only private AS numbers
  - removal process happens before routing filters are applied and before local AS. Number is prepended to the AS path
  - Currently works only with 16-bit ASNs

©2006-2013 MikroTikSA Pty Ltd

Other Peer Options

- Route Reflect: Used in some iBGP scenarios to reduce the number of full mesh sessions required
- In/Out Filter: For filtering on a per peer basis
- Default Originate: Send a default route to downstream peer
  - Never – do not send
  - If installed – only send if originator has default route
  - Always – always originate a default to the downstream
- Passive: Remote peer must initiate connection
- Use BFD: use Bidirectional Forwarding and Detection
- Multihop: Specifies if the peer is more than 1 hop away, requires IGP setup if enabled (only for eBGP)
- TTL: Time To Live, the hop limit for TCP connection. For example, if 'ttl=255' then only single hop neighbors will be able to establish the connection. This only affects eBGP peers.

©2006-2013 MikroTikSA Pty Ltd

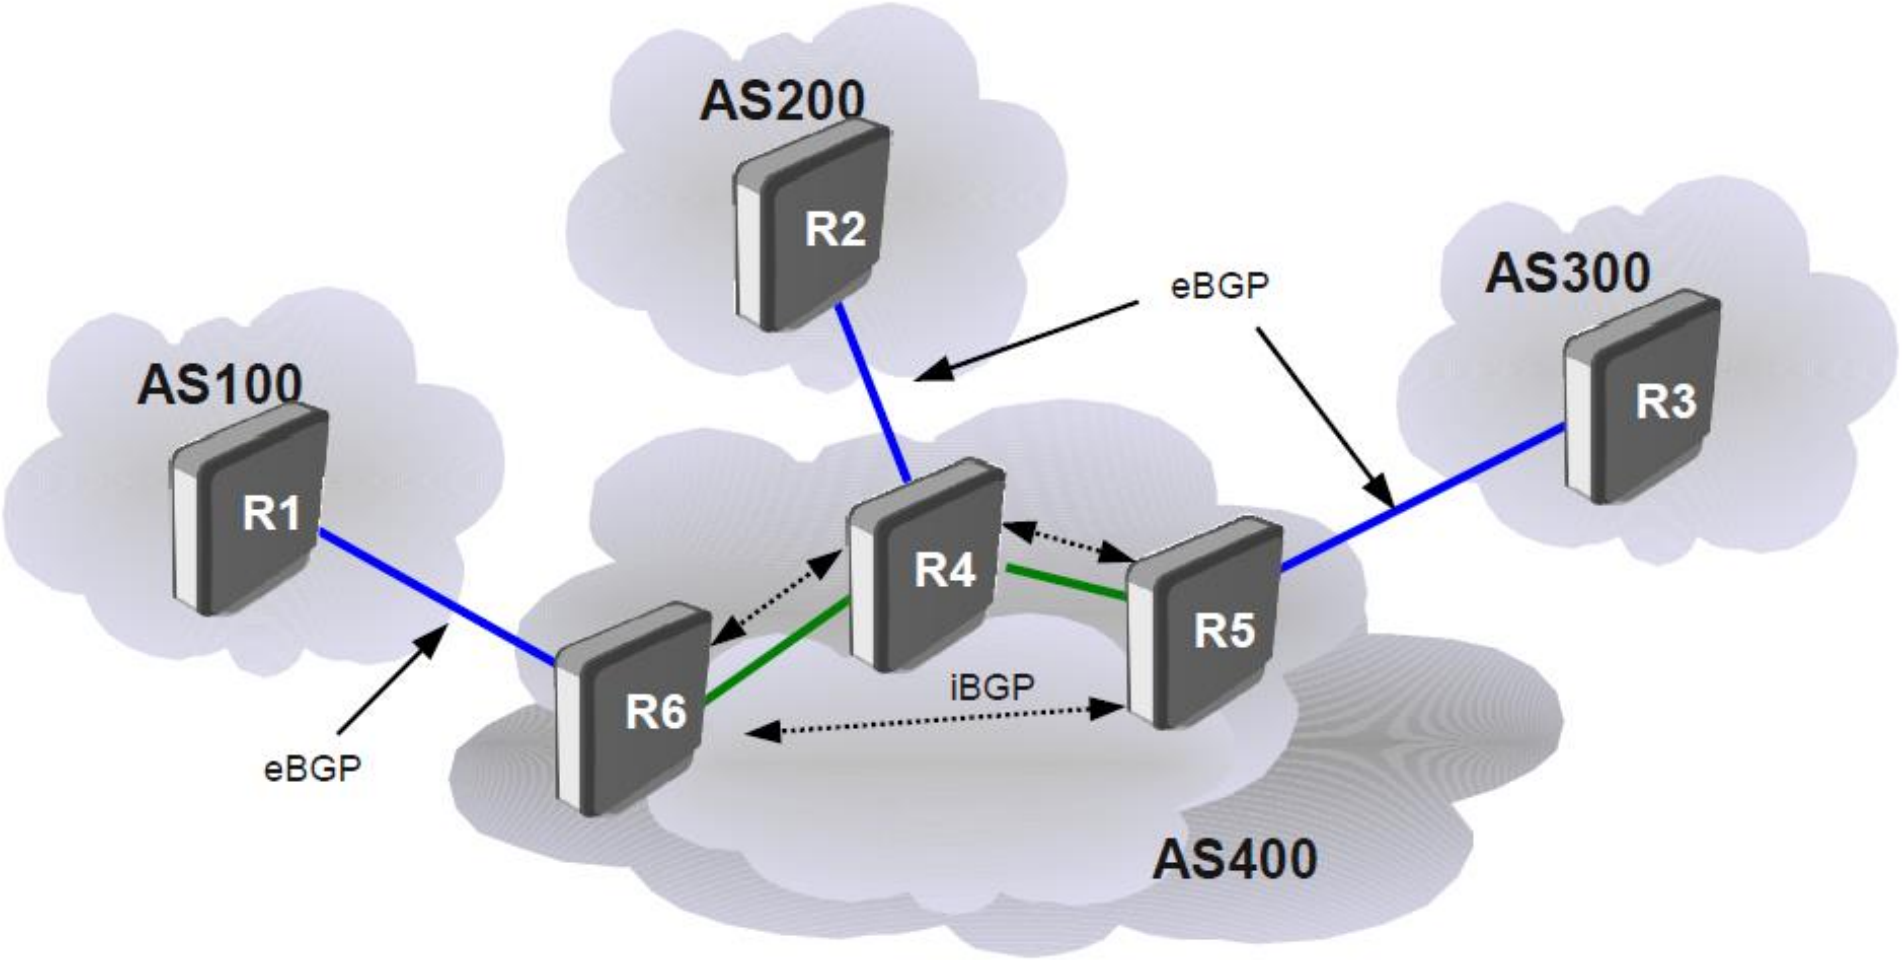
BGP and FW connection tracking

- Connection tracking is unable to keep valid track of connections with multi-homed BGP.
  - Routers may not see all sides of a TCP Conversation
  - Packets related to one connection can travel through different paths
- Do not drop (Forward chain) invalid connections in firewall
- Conntrack can be turned off for better performance (Performance vs Security)

©2006-2013 MikroTikSA Pty Ltd

Interior and Exterior BGP

- **iBGP** – peering between routers inside an AS
- **eBGP** – peering between routers from different ASs



©2006-2013 MikroTikSA Pty Ltd

eBGP

- Usually formed between directly connected peers (AS edge routers)
- eBGP setup from IP on physical Interfaces (not LoopBack)
- eBGP session is run on inter-AS point to point links
- Multi-hop configuration is required if peers are not directly connected (uncommon config)
- Adds AS to advertised prefix's path
- By default Next-hop is changed to self
- BGP session is run on inter-AS point to point links
- Use TTL for added security

©2006-2013 MikroTikSA Pty Ltd

iBGP – Interior BGP

- If you have multiple eBGP routers in your AS they will need to know about each others external routes
- Normal IGP's like OSPF cannot handle route tables the size of the internet
  - You also would not want to advertise your external route table to your internal routers!
- iBGP carries route information between eBGP routers
- Carries some/all Internet prefixes across ISP backbone
- Carries ISP's customer prefixes
- BGP session is run between router **loopback interfaces**
- BGP instance and remote bgp peers have same AS number
  - This is the only difference in setup between eBGP and iBGP

©2006-2013 MikroTikSA Pty Ltd

eBGP Best Practice

- Do not Advertise Prefixes longer than 24 bits (/22 is considered the norm)
  - Filter Prefixes /25-32 bits in length using Routing Filters
- Where Possible only Advertise the whole Prefix assigned by RIR (Aggregation)
- If everyone obeys this rule the BGP Table will be kept small
- Subprefixes of the RIR allocation may be used internally in the ISP Network
  - OSPF or other IGP used to distribute and track sub-prefixes

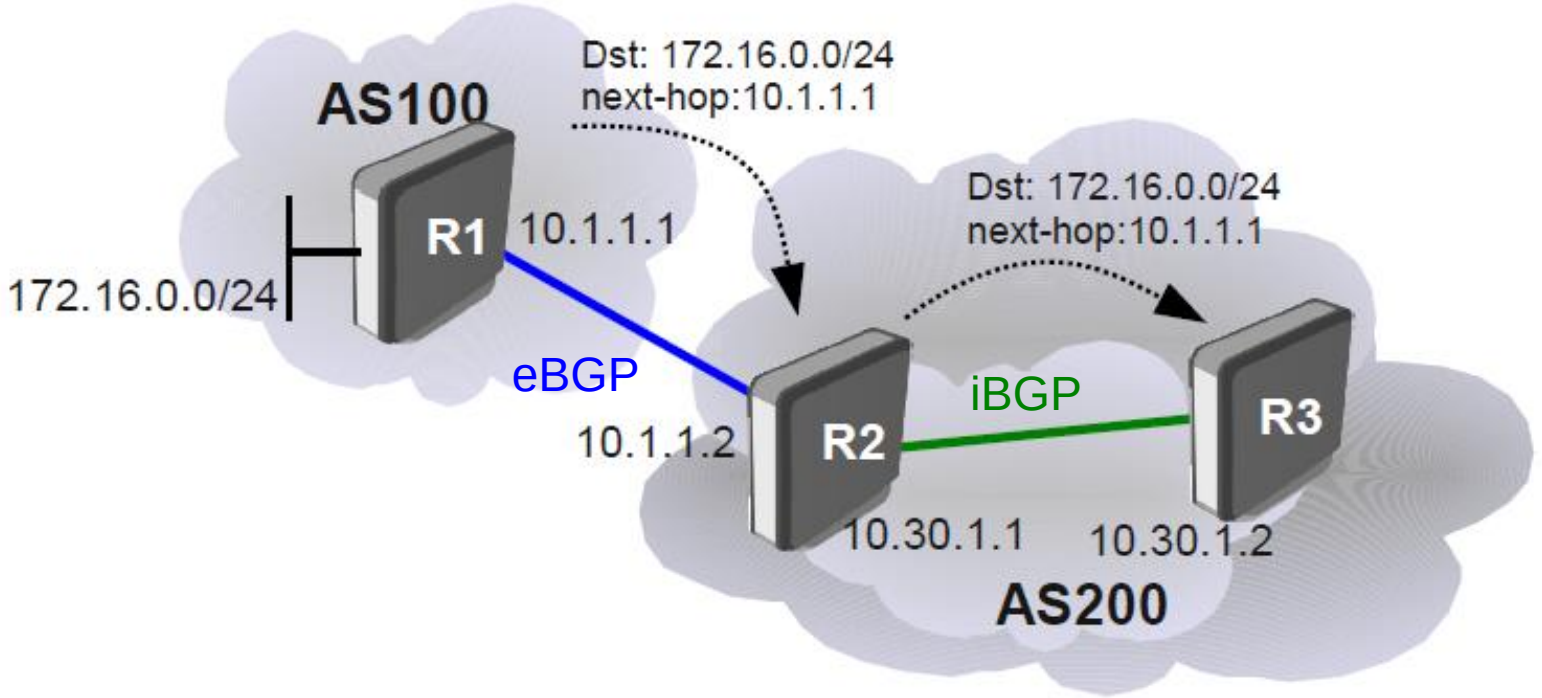
©2006-2013 MikroTikSA Pty Ltd

LAB

- eBGP – second ISP (inx) link
- eBGP – backup router
- iBGP – connect routers

Nexthop

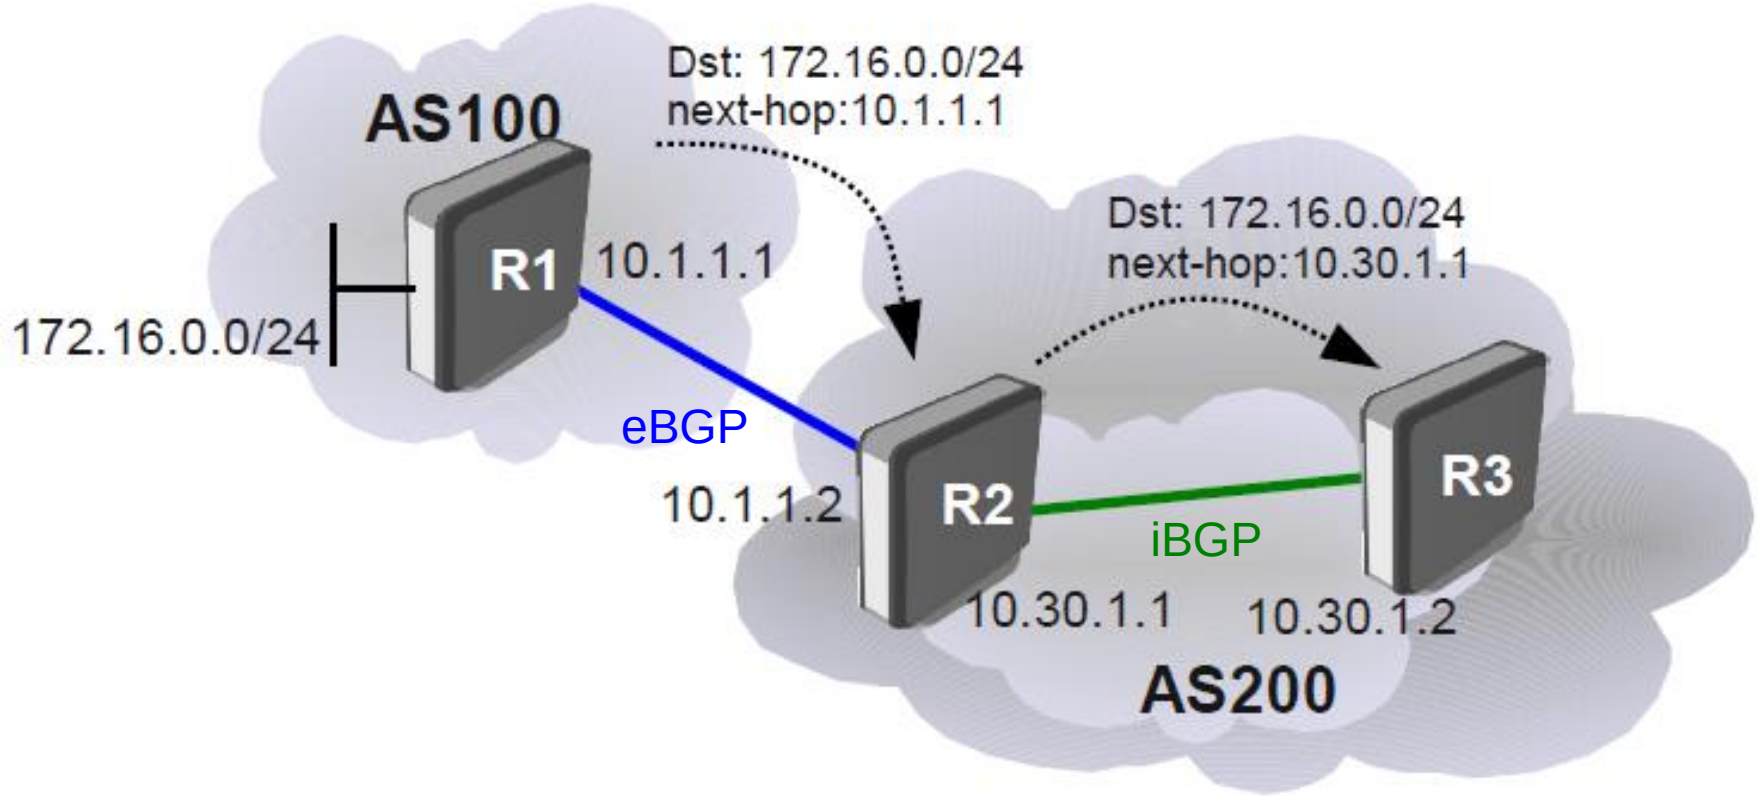
- IP address that is used to reach a certain destination
- For eBGP nexthop is neighbour's IP address
- eBGP advertised nexthop is carried into iBGP (unchanged)
- BGP assigns the first valid path as the current best path. BGP then compares the best path with the next path in the list, until BGP reaches the end of the list of valid paths



©2006-2013 MikroTikSA Pty Ltd

Nexthop self

- Force BGP to use a specific IP as a nexthop
- Can be used to eliminate IGP on directly peered neighbours
- `# config on R2`
- **`/routing bgp peer`** `set peer1 nexthop-choice=force-self`



©2006-2013 MikroTikSA Pty Ltd

LAB

- iBGP – nexthop selection
- iBGP – default route

Route Filter

- Powerful function to modify routes
- Can alter routes entering and leaving router
- Includes default chains for Connected, OSPF routes
- Requires additional chains for BGP peers in/out
- Includes jump and return functions for chain reuse
- Able to match and/or modify any attribute of a route

©2006-2013 MikroTikSA Pty Ltd

Using Route Filters

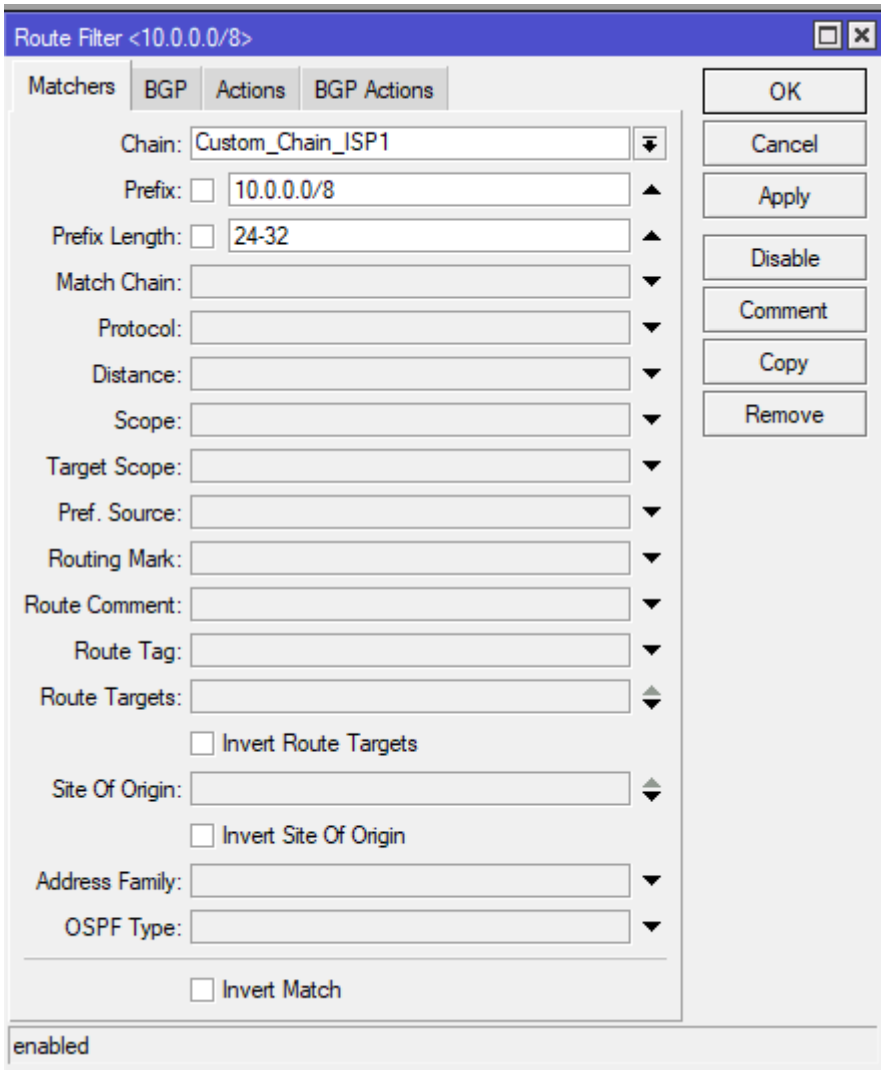
- First 2 tabs define matches for routes
- 2nd tab only to match BGP attributes
- Action defines what to do
- accept - accept the route information
- discard - completely exclude matching prefix from further processing.
- jump - pass control to another filter list by 'jump-target'
- log - log message about this match in system log and continue with the next rule in chain
- passthrough - continue to the next rule in chain (usually used to set BGP attributes)
- reject - reject the routing information for matching prefix.
- return - return to the previous chain from which a jump to the current chain took place

©2006-2013 MikroTikSA Pty Ltd

- Prefix defines the total scope of the match
- Prefix length defines which specific prefix to match within the total scope
- This example will match any route falling within the 10.0.0.0/8 subnet that has a subnet mask of between 24 and 32 bits in length
- These routes will match

  10.0.0.0/24, 10.100.1.0/24, 10.1.2.3
- These will not match

  10.1.0.0/16, 10.0.0.0/8, 10.100.0.0/23



©2006-2013 MikroTikSA Pty Ltd

Prefix Filtering Scenarios

- There are three scenarios for receiving prefixes from other ASNs
  - Customer talking BGP
  - Peer talking BGP
  - Upstream/Transit talking BGP
- Each has different filtering requirements and need to be considered separately

©2006-2013 MikroTikSA Pty Ltd

Receiving Prefixes from Customers

- ISPs should only accept prefixes which have been assigned or allocated to their downstream customer
- If ISP has assigned address space to its customer, then the customer is entitled to announce it back to his ISP
- If the ISP has NOT assigned address space to its customer, then:
  - Check in the five RIR databases to see if this address space really has been assigned to the customer
  - Use the tool: whois –h whois.apnic.net x.x.x.0/24
- Http://RIPE.net/ NCC (has a web based lookup tool for prefixes)
- Http://www.radb.net/ RADb has a database of Public Routed Information and Provides Monitoring services to Network operators.

©2006-2013 MikroTikSA Pty Ltd

Receiving Prefixes from Peer (ISPs)

- A peer is an ISP with whom you agree to exchange prefixes that you originate into the Internet routing table
  - Could also be an aggregation node such as an INX
  - Prefixes you accept from a peer are only those they have indicated they will announce
  - Prefixes you announce to your peer are only those you have indicated you will announce
- Agreeing what each will announce to the other:
  - Exchange of e-mail documentation as part of the peering agreement, and then ongoing updates

    OR
  - Use of the Internet Routing Registry and configuration tools such as the IRRToolSet https://github.com/irrtoolset/irrtoolset

©2006-2013 MikroTikSA Pty Ltd

Receiving Prefixes: Upstream/Transit Provider

- Upstream/Transit Provider is an ISP who you pay to give you transit to the whole Internet
- Receiving prefixes from them is not desirable unless required for multihoming or traffic engineering
- Ask upstream/transit provider to either:
  - originate a default-route (the usual option)

  Or

  - announce one prefix you can use as default

©2006-2013 MikroTikSA Pty Ltd

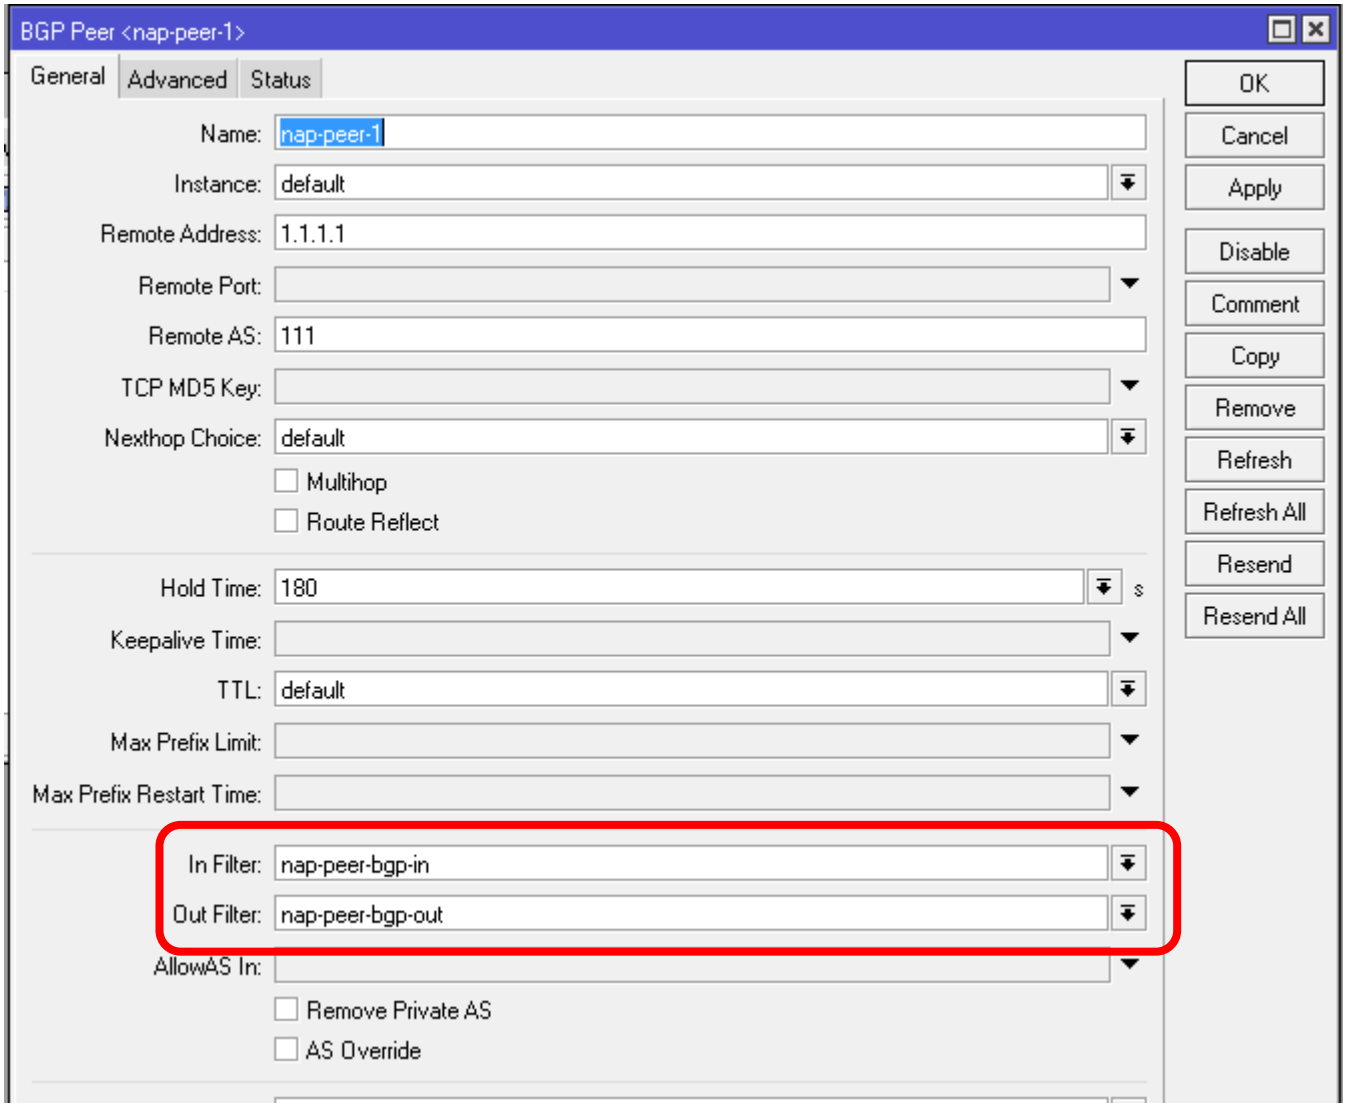
Receiving Prefixes:Upstream/Transit Provider

- If necessary to receive prefixes from any provider, care is required don't accept RFC1918 etc prefixes (i.e. internal networks) ftp://ftp.rfc-editor.org/in-notes/rfc3330.txt
  - don't accept your own prefixes
  - don't accept default (unless you need it)
  - ftp://ftp.rfc-editor.org/in-notes/rfc3330.txt (expanded list)
- Check CYMRU's list of "bogons"
  - http://www.cymru.com/Documents/bogon-list.html
- Or get a BGP feed from the Bogon Route Server
  - http://www.cymru.com/BGP/bogon-rs.html

©2006-2013 MikroTikSA Pty Ltd

Route Filter Peer Setup

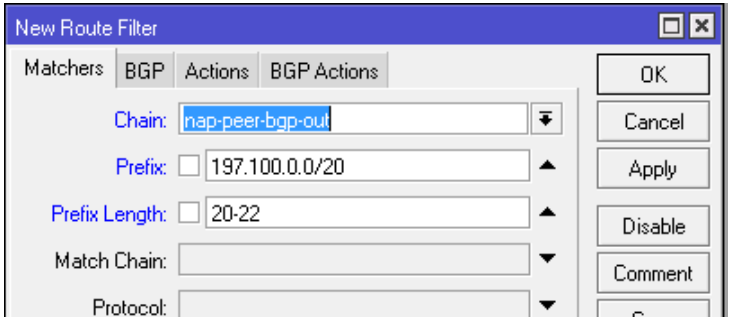
- Specify in and out filters on a per-peer basis

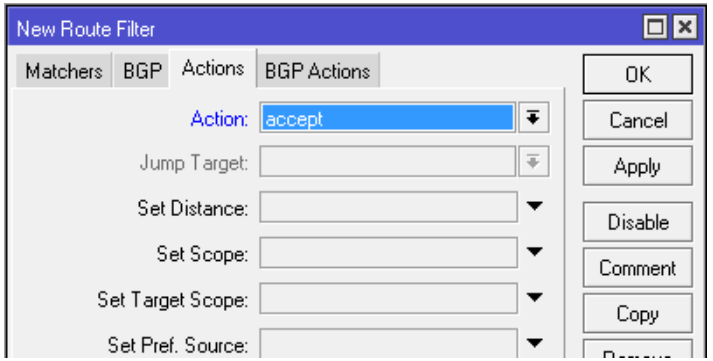


©2006-2013 MikroTikSA Pty Ltd

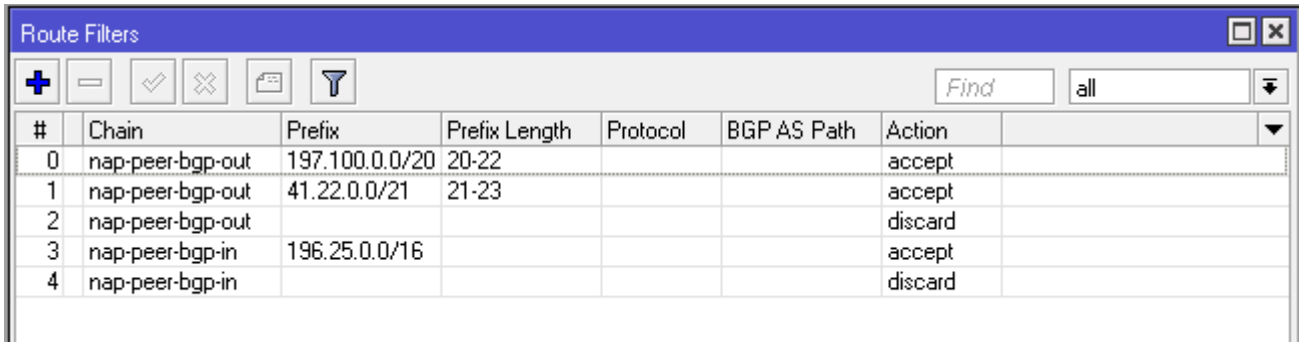
Route Filter

- OUT filter
  - Allow only your aggregated ranges within allowable prefix lengths
  - Drop everything else





- IN filter
  - Allow only prefixes and prefix lengths that peer has specified and you have confirmed with RIR lookup
  - Drop everything else
  - Including your own ranges



#	Chain	Prefix	Prefix Length	Protocol	BGP AS Path	Action
0	nap-peer-bgp-out	197.100.0.0/20	20-22			accept
1	nap-peer-bgp-out	41.22.0.0/21	21-23			accept
2	nap-peer-bgp-out					discard
3	nap-peer-bgp-in	196.25.0.0/16				accept
4	nap-peer-bgp-in					discard

©2006-2013 MikroTikSA Pty Ltd

LAB

- Route Filters – out filter
  - Only advertise own range
  - Drop everything else

BGP Attributes

- BGP attributes are an array of information carried in a BGP update capable of indicating anything from path preference to various additional pieces of information about a route, either within an AS or outside an AS
- There are four basic types of attributes:
  - **Well known mandatory attributes -** these attributes must be recognized by all BGP speakers, and must be included in all update messages. Almost all of the attributes impacting the path decision process, are well known mandatory attributes.
  - **Well known discretionary attributes**; these attributes must be recognized by all BGP speakers, and may be carried in updates, but are not required
  - **Optional transitive attributes -** these attributes may be recognized by some BGP speakers, but not all. They should be preserved and advertised to all peers whether or not they are recognized.
  - **Optional non-transitive attributes -** these attributes may be recognized by some BGP speakers, but not all. If an update containing an optional transitive attribute is received, the update should be advertised to peers without the unrecognized attributes.

©2006-2013 MikroTikSA Pty Ltd

Route Filters Reject vs Drop

- Outgoing filters – there is no difference in results,
- Incoming filters-
- **Reject:**
  - occupies more memory
  - Routes are added to main routing table but made inactive (Routes are blue in Routing Table)
  - This Visibility is good for diagnostics, not good if you are running out of memory
  - Often used to “quarantine” routes on new peers before bringing live
- **Drop:**
  - Drops routes silently
  - No routing information stored in Memory
  - No diagnostics capability

©2006-2013 MikroTikSA Pty Ltd

Common Private List

```
/routing filter
add action=discard chain=rfc5735 comment="default route" prefix=0.0.0.0/0
add action=discard chain=rfc5735 comment="Prefix >24" prefix-length=25-32
add action=discard chain=rfc5735 comment="This network" prefix=0.0.0.0/8 prefix-length=8-32
add action=discard chain=rfc5735 comment=Loopback prefix=127.0.0.0/8 prefix-length=8-32
add action=discard chain=rfc5735 comment=Private prefix=10.0.0.0/8 prefix-length=8-32
add action=discard chain=rfc5735 comment="CG Nat" prefix=100.64.0.0/10 prefix-length=10-32
add action=discard chain=rfc5735 comment=Private prefix=172.16.0.0/12 prefix-length=12-32
add action=discard chain=rfc5735 comment=Private prefix=192.168.0.0/16 prefix-length=16-32
add action=discard chain=rfc5735 comment=Private prefix=169.254.0.0/16 prefix-length=16-32
add action=discard chain=rfc5735 comment=ietf prefix=192.0.0.0/24 prefix-length=24-32
add action=discard chain=rfc5735 comment=ietf prefix=192.0.2.0/24 prefix-length=24-32
add action=discard chain=rfc5735 comment=reserved prefix=192.88.99.0/24 prefix-length=24-32
add action=discard chain=rfc5735 comment="benchmark testing" prefix=198.18.0.0/15 prefix-length=15-32
add action=discard chain=rfc5735 comment="benchmark testing" prefix=198.51.100.0/24 prefix-length=24-32
add action=discard chain=rfc5735 comment="benchmark testing" prefix=203.0.113.0/24 prefix-length=24-32
add action=discard chain=rfc5735 comment=Multicast prefix=224.0.0.0/3 prefix-length=3-32
add action=return chain=rfc5735 comment="Return from chain"
```


©2006-2013 MikroTikSA Pty Ltd

Common Private List

/routing filter

```
add action=discard chain=ipv6martian comment="IPv6 Martian - Default route" prefix=::/0
add action=discard chain=ipv6martian comment="IPv6 Martian - IPv4-compatible IPv6 address  deprecated by RFC4291" prefix=::/96
add action=discard chain=ipv6martian comment="IPv6 Martian - Unspecified address" prefix=::/128
add action=discard chain=ipv6martian comment="IPv6 Martian - Local host loopback address" prefix=::1/128
add action=discard chain=ipv6martian comment="IPv6 Martian - IPv4-mapped addresses" prefix=::ffff:0.0.0.0/96
add action=discard chain=ipv6martian comment="IPv6 Martian - Compatible address (IPv4 format)" prefix=::224.0.0.0/100
add action=discard chain=ipv6martian comment="IPv6 Martian - Compatible address (IPv4 format)" prefix=::127.0.0.0/104
add action=discard chain=ipv6martian comment="IPv6 Martian - Compatible address (IPv4 format)" prefix=::/104
add action=discard chain=ipv6martian comment="IPv6 Martian - Compatible address (IPv4 format)" prefix=::255.0.0.0/104
add action=discard chain=ipv6martian comment="IPv6 Martian - Pool used for unspecified loopback and embedded IPv4 addresses" prefix=::/8
add action=discard chain=ipv6martian comment="IPv6 Martian - OSI NSAP-mapped prefix set (RFC4548)  deprecated by RFC4048" prefix=200::/7
add action=discard chain=ipv6martian comment="IPv6 Martian - Former 6bone now decommissioned" prefix=3ffe::/16
add action=discard chain=ipv6martian comment="IPv6 Martian - Reserved by IANA for special purposes and documentation" prefix=2001:db8::/32
add action=discard chain=ipv6martian comment="IPv6 Martian - Invalid 6to4 packets (IPv4 multicast)" prefix=2002:e000::/20
add action=discard chain=ipv6martian comment="IPv6 Martian - Invalid 6to4 packets (IPv4 loopback)" prefix=2002:7f00::/24
add action=discard chain=ipv6martian comment="IPv6 Martian - Invalid 6to4 packets (IPv4 default)" prefix=2002::/24
add action=discard chain=ipv6martian comment="IPv6 Martian - Invalid 6to4 packets" prefix=2002:ff00::/24
add action=discard chain=ipv6martian comment="IPv6 Martian - Invalid 6to4 packets (IPv4 private 10.0.0.0/8 network)" prefix=2002:a00::/24
add action=discard chain=ipv6martian comment="IPv6 Martian - Invalid 6to4 packets (IPv4 private 172.16.0.0/12 network)" prefix=2002:ac10::/28
add action=discard chain=ipv6martian comment="IPv6 Martian - Invalid 6to4 packets (IPv4 private 192.168.0.0/16 network)" prefix=2002:c0a8::/32
add action=discard chain=ipv6martian comment="IPv6 Martian - Unicast Unique Local Addresses (ULA)  RFC 4193" prefix=fc00::/7
add action=discard chain=ipv6martian comment="IPv6 Martian - Link-local Unicast" prefix=fe80::/10
add action=discard chain=ipv6martian comment="IPv6 Martian - Site-local Unicast  deprecated by RFC 3879 (replaced by ULA)" prefix=fec0::/10
add action=discard chain=ipv6martian comment="IPv6 Martian - Multicast" prefix=ff00::/8
add action=return chain=ipv6martian
```


©2006-2013 MikroTikSA Pty Ltd

LAB

- Bogon Filter Demo

Alternatives to Static Filtering

- To quote Andrew (Teraco circa February 2020)

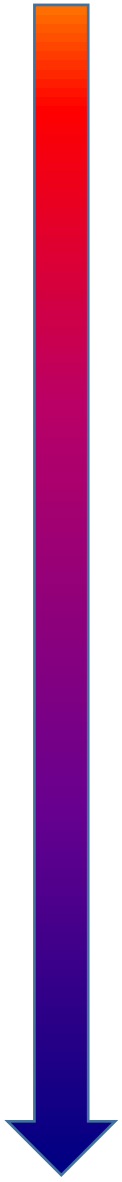
“Static filtering will melt the polar ice caps and lead to the apocalypse”

- An excellent alternative to static filtering is using the RPKI
- Resource Public Key Infrastructure (RPKI) is a cryptographic method of signing records that associate a BGP route announcement with the correct originating AS number.
- RPKI is defined in RFC6480 (An Infrastructure to Support Secure Internet Routing).
- Unfortunately it is not supported in RouterOS (yet…)
- It has been promised that RPKI validation will happen in RouterOS7
- I’ll remind Andrew that ROS7 will restore the ice caps and the hole in the ozone layer

Further reading here: https://blog.cloudflare.com/rpki/

©2006-2013 MikroTikSA Pty Ltd

Route Selection



Routing Table Lookup

Firewall Mangle / Route Rule/Mark match / VRF

Longest Prefix Match

Most specific route to destination

Route Distance

Administrative Distance by Route

BGP Route Selection

AS Path Length, Weight, Localpref etc.

Copyright

BGP decision algorithm

- BGP uses single best path to reach the destination
- BGP always propagates the best path to the neighbours
- Best Path Algorithm runs within **Each Single BGP Instance**
- Different prefix attributes are used to determine best path, like weight, next-hop, as-path, localpref etc.
- If there are no specific settings that can affect the outcome, BGP Best Path Selection Algorithm determines the best route by selecting the shortest path to the destination

©2006-2013 MikroTikSA Pty Ltd

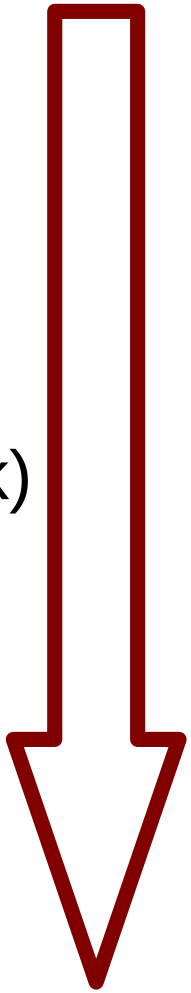
BGP decision algorithm

- For any path to be considered valid it has to meet these requirements:
  - Next-hop IP address of that path is reachable
  - The local AS number is not part of the AS_PATH (basic loop prevention)
  - If BGP synchronization is enabled, the candidate prefix is in the IGP routing table
    - If using OSPF, router-ID must match for the OSPF and BGP process
  - No BGP polices are filtering the candidate prefix.

©2006-2013 MikroTikSA Pty Ltd

Best path selection

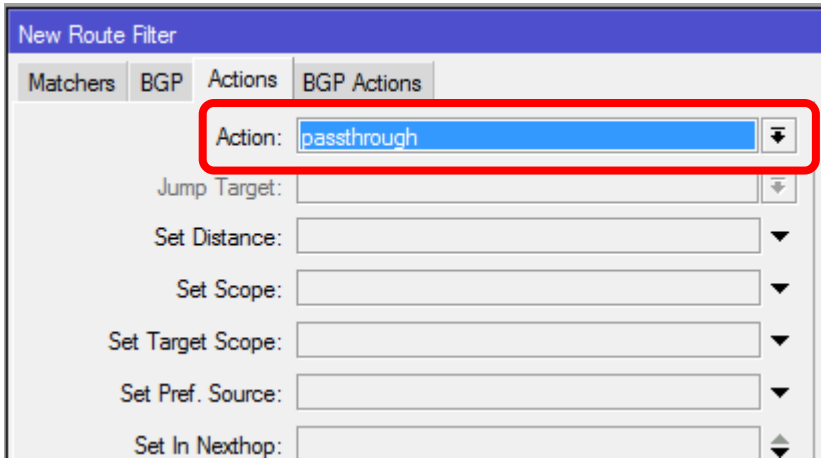
1. Next-hop validation --
2. Highest WEIGHT (default 0)
3. Highest LOCAL-PREF – iBGP Attribute
4. Shortest AS-PATH – eBGP Attribute
5. Prefer locally sourced route (via aggregate or BGP network)
6. Lowest origin type (EGP, IGP, Incomplete in that order)
7. Lowest MED – eBGP Attribute
8. Prefer eBGP over iBGP
9. Prefer the route with lowest router ID or ORIGINATOR_ID
10. Shortest route reflection cluster (default 0)
11. Prefer the path that comes from the lowest neighbour address

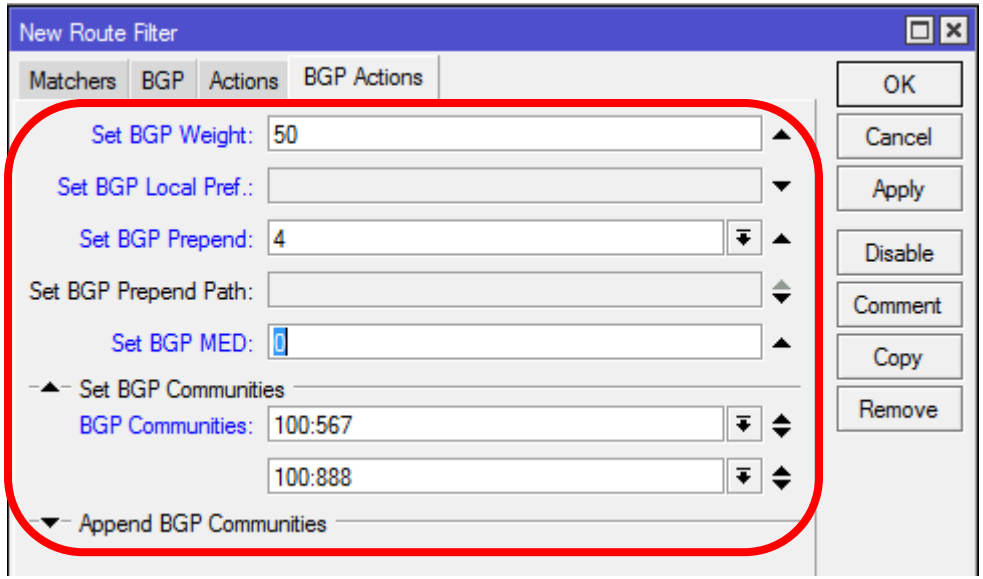


©2006-2013 MikroTikSA Pty Ltd

Modifying Best Path Selection

- Modifying of BGP Attributes is mostly done via Routing Filters
- Create separate bgp-in and bgp-out chains for each peer
- Apply rules as action = **passthrough** or **accept** with BGP modification
- Passthrough is used when modifying attributes in multiple rules
- Accept is used when accepting ranges before a drop





©2006-2013 MikroTikSA Pty Ltd

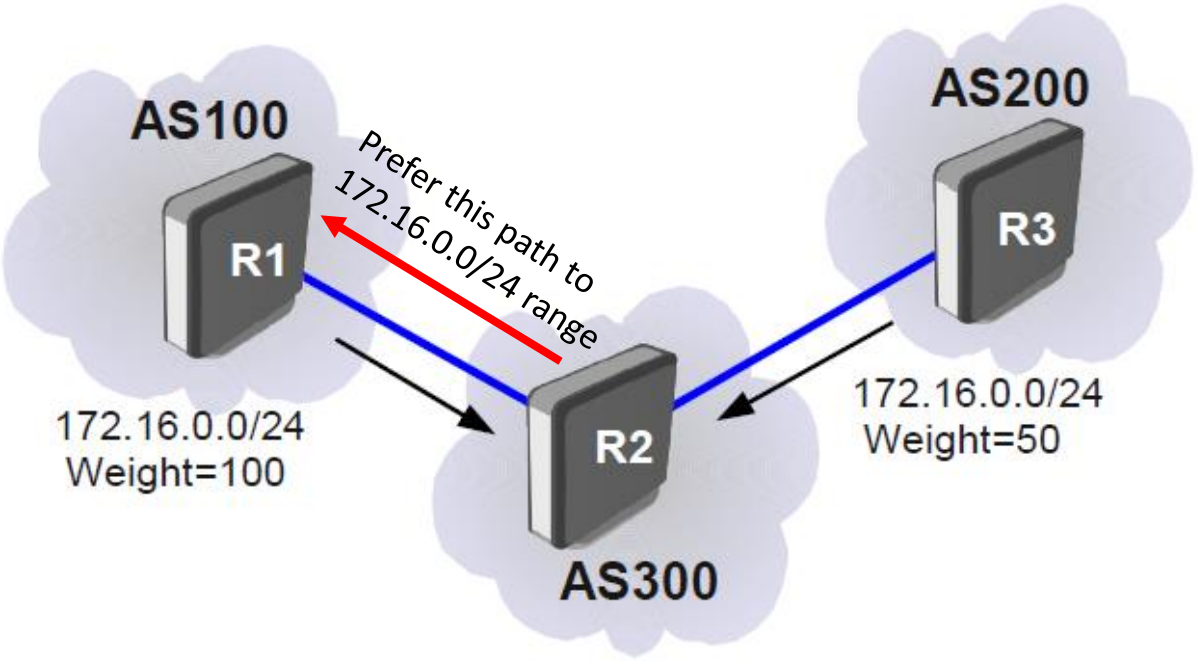
Routes in/out of AS

- Filtering incoming routes will change, how we see the external world, thus influencing how we send traffic
- Filtering outgoing routes will change how the world see us, thus influencing how we receive traffic.
- To modify how traffic LEAVES your network
  - Weight – used on a single router
  - Local-pref – used throughout AS
- To modify how traffic ENTERS your network
  - AS Path Length
  - Longest prefix match (not a BGP attribute)

©2006-2013 MikroTikSA Pty Ltd

Weight

- Used where the is a single router with multiple paths out the AS
- Weight is assigned **locally to the router**
- Prefix without assigned weight has a default value of 0
- Route with higher weight is preferred
- R2 will use Route to R1 for the address 172.16.0.0/24



©2006-2013 MikroTikSA Pty Ltd

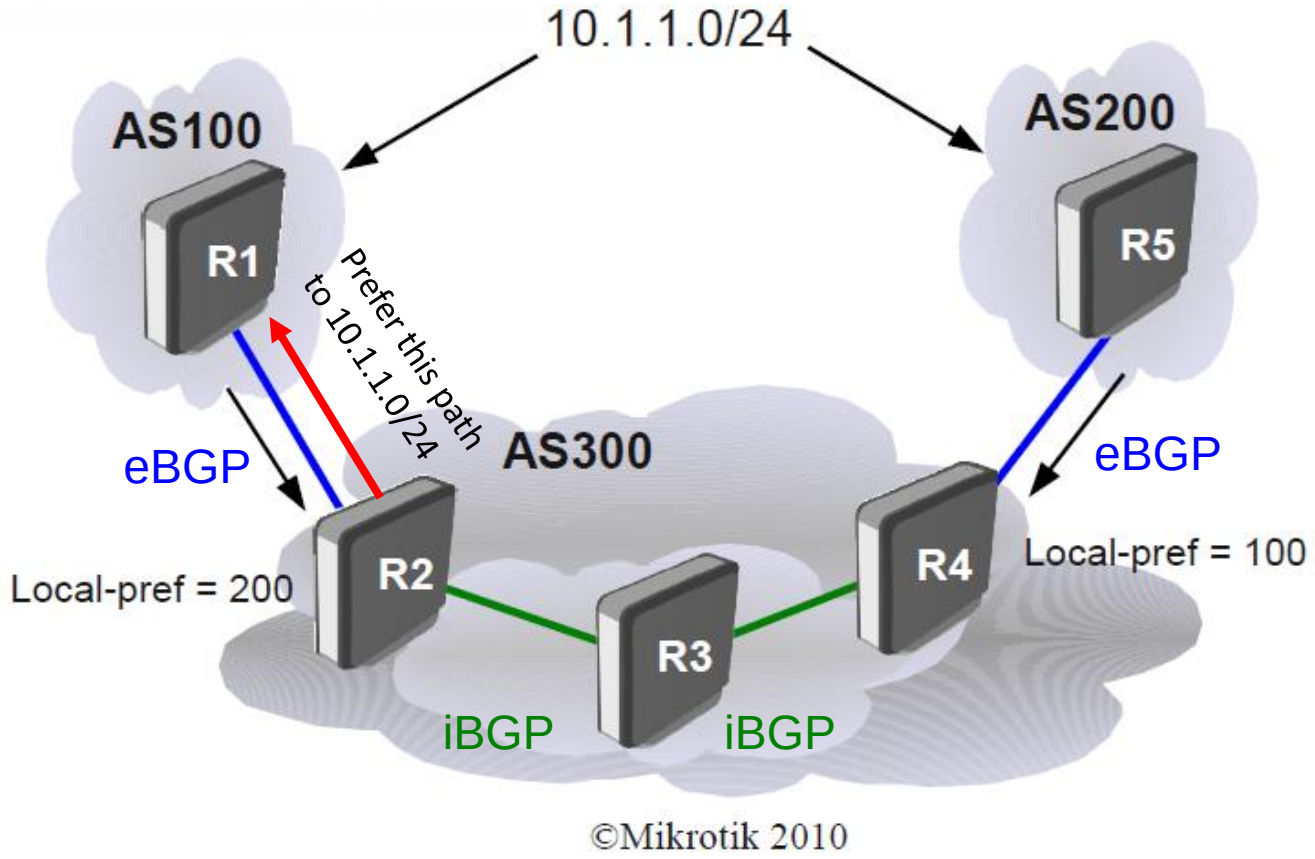
Local Pref

- Local pref is the second BGP attribute.
- You can use local pref to choose the outbound external BGP path.
- Local preference is sent to all internal BGP routers in your autonomous system – it is an iBGP attribute and therefore not exchanged between eBGP routers.
- Local preference is a well-known and discretionary BGP attribute.
- Default value is 100.
- The path with the highest local preference is preferred

©2006-2013 MikroTikSA Pty Ltd

Local Preference

- Used where there are multiple routers running eBGP with paths out the AS
- Indicates which is the **preferred path** to exit AS
- Can be used to replace weight in almost any scenario



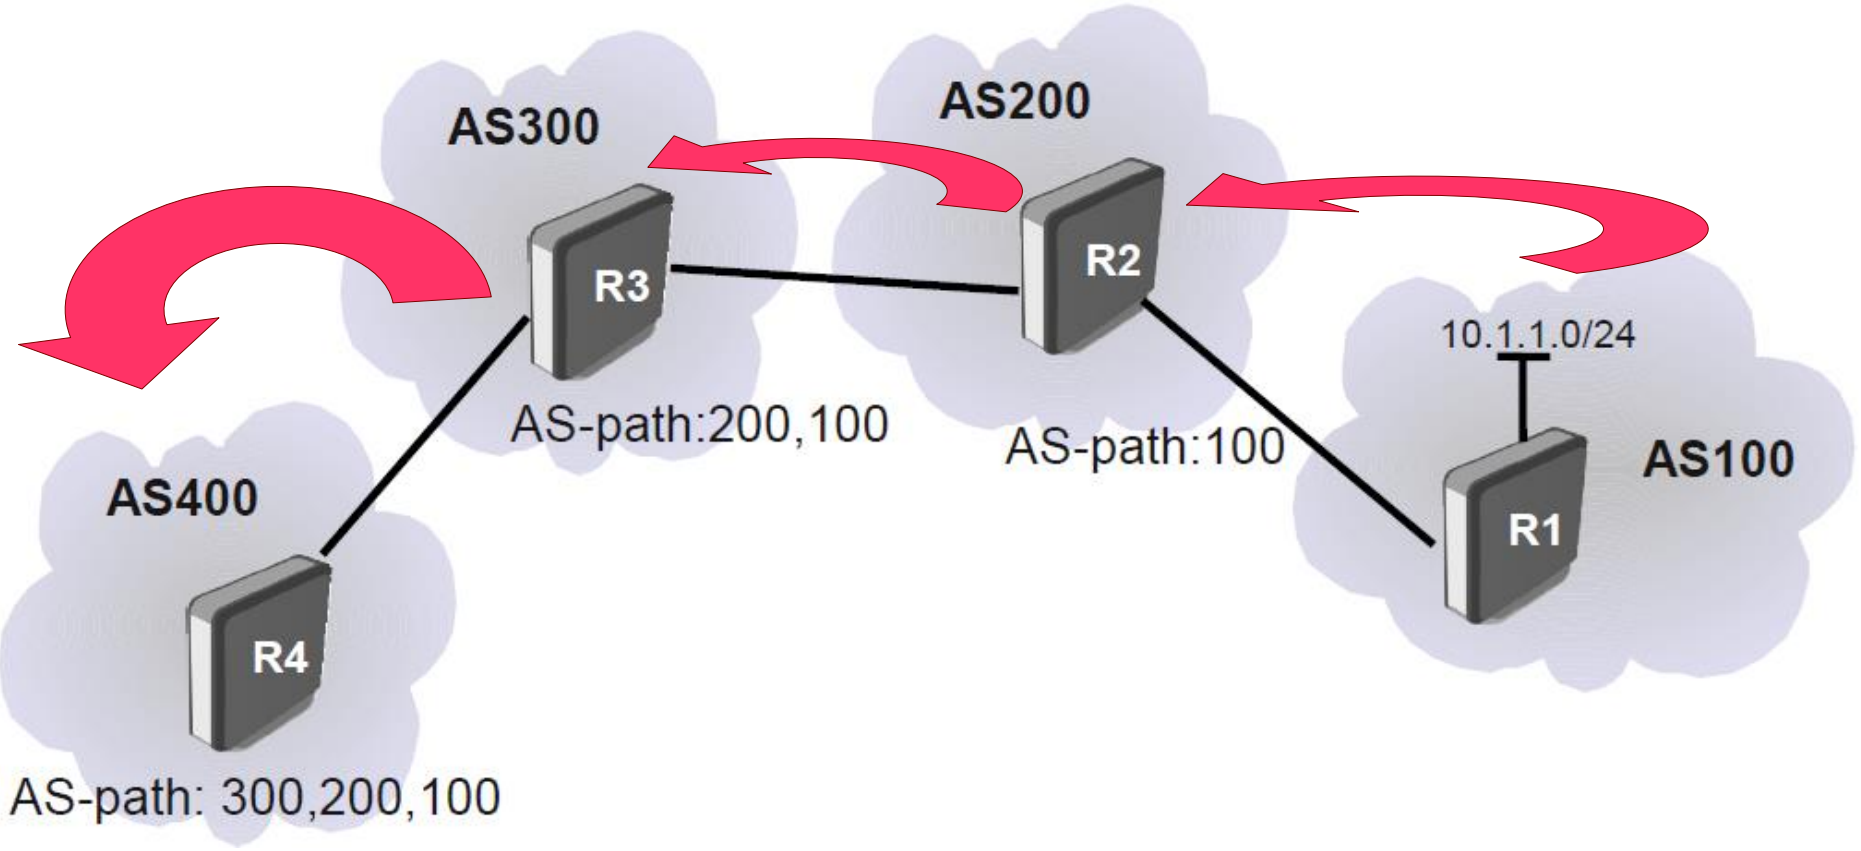
©2006-2013 MikroTikSA Pty Ltd



- BGP Attributes
  - Local-pref

AS Path

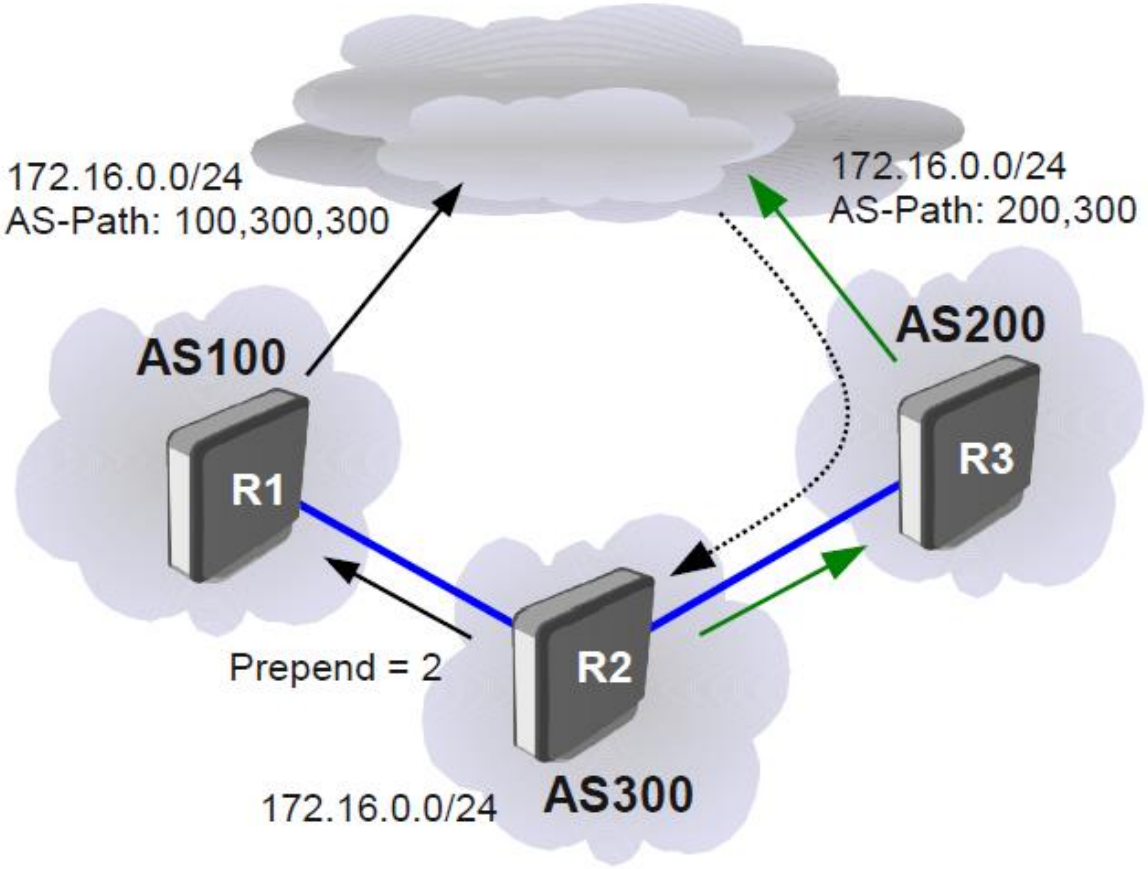
- List of AS numbers that an update has traversed
- Prefer the path with the least number of AS's



©2006-2013 MikroTikSA Pty Ltd

AS-Path Prepend

- AS-Path manipulations can be used to influence best path selection on upstream routers
  - Used to control how bandwidth enters your network
- Upstream ISP thinks one path is longer than another
- Can prepend your Path a number of times



©2006-2013 MikroTikSA Pty Ltd

LAB

- BGP Attributes
  - AS Path Prepend
- Longest prefix match / splitting prefixes

Community

- BGP communities provide additional capability for tagging routes and for modifying BGP routing policy on upstream and downstream routers.
- BGP communities can be appended, removed, or modified selectively on each attribute as the route travels from router to router.
- BGP communities are an optional transitive BGP attribute that can traverse from AS to AS
- The private BGP community pattern could vary from organization to organization, do not need to be registered, and could signify geographic locations for one AS while signifying a method of route advertisement in another AS.

©2006-2013 MikroTikSA Pty Ltd

Community

- Communities can be used to mark a set of prefixes that share a common property.
- Upstream providers can use these marks to apply a common routing policy such as filtering or assigning a specific local preference.
- As a service provider you can make an agreement with your customers on a specific policy to be applied to their prefixes using communities
  - this gives your customers the freedom to change the policy of a prefix just by changing the community attribute value with no support from your side.

©2006-2013 MikroTikSA Pty Ltd

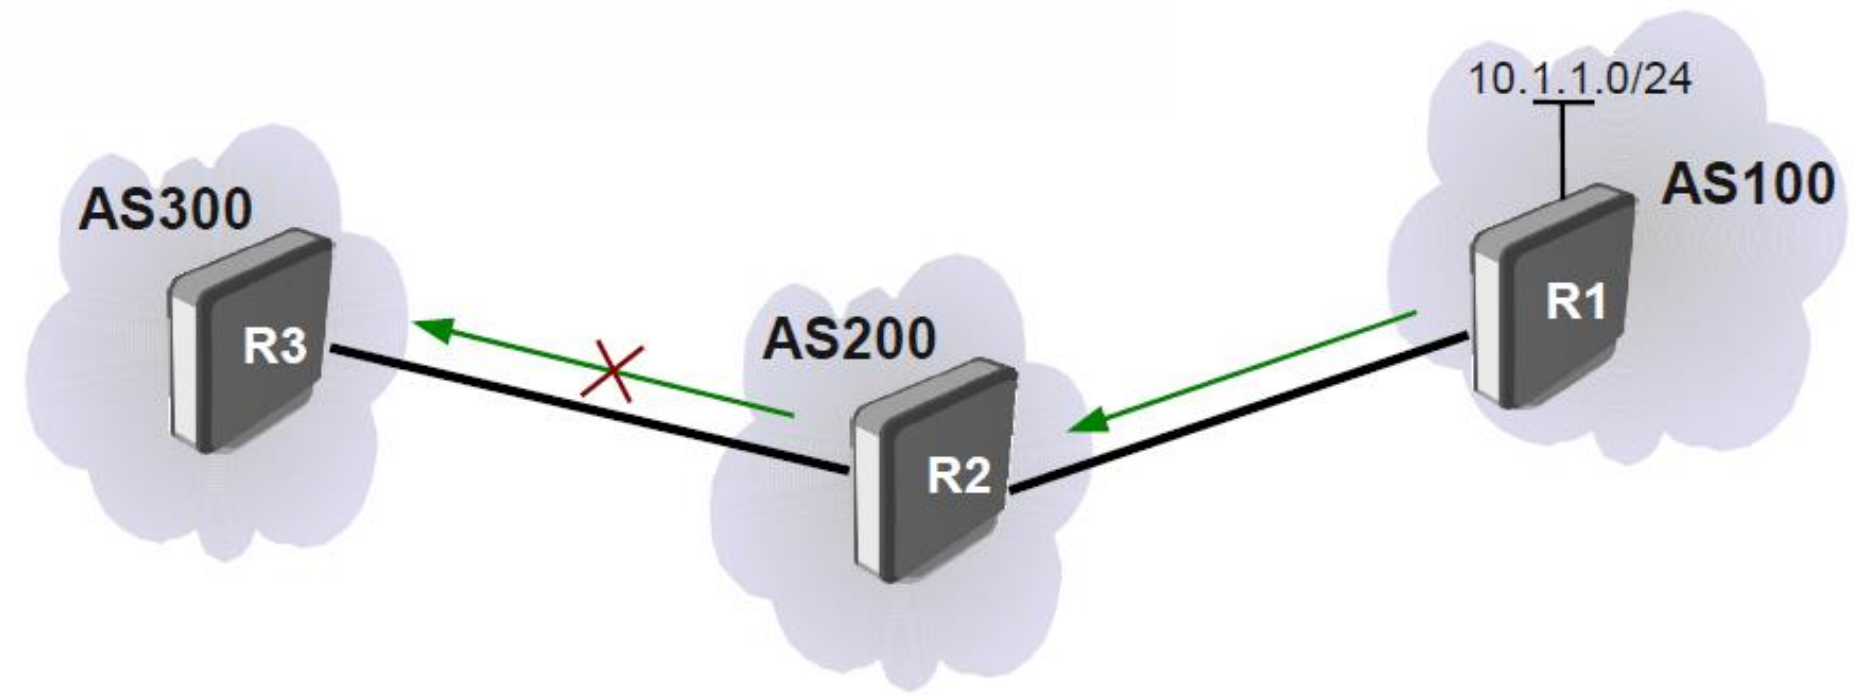
Community

- Attribute that groups destinations
- Filters can be easily applied to whole group
- Default groups:
  - No-export – do not advertise to eBGP peer
  - No-advertise – do not advertise to any peer
  - Local-as – do not send outside local AS (in non-confederation network the same as no-export)

©2006-2013 MikroTikSA Pty Ltd

BGP Community

- Assume that you don't want R2 to propagate routes learned from R1
- **# config on R1**
- **/routing filter** add chain=bgp-out action=passthrough set-bgp-communities=no-export



©2006-2013 MikroTikSA Pty Ltd



- BGP Communities
  - No export

BGP Community

- 32-bit value written in format “xx:yy”
- Private BGP communities follow the convention that the first 16-bits represent the AS of the community origination, and the second 16-bits represent a pattern defined by the originating AS
  - This only applies to ASN below 65535 (old 16 bit format)
- Gives customer more policy control
- Simplifies upstream configuration
- Can be used by ISPs for:
  - AS prepending options
  - Geographic restrictions
  - Blackholing, etc.

©2006-2013 MikroTikSA Pty Ltd

ISP

```
aut-num: AS2588
as-name: LatnetServiss-AS
descr: LATNET ISP
member-of: AS-LATVIA
remarks: +-----------------------------------------
remarks: |
remarks: | x=0 Announce as is
remarks: | x=1 Prepend +1
remarks: | x=2 Prepend +2
remarks: | x=3 Prepend +3
remarks: | x=4 Prepend +4
remarks: | x=5 Prepend +5
remarks: |
remarks: | 2588:400 Latvian Nets
remarks: | 2588:500 Announce to LIX (Latvian Internet Exchange)
remarks: | 2588:666 Don't announce (blackhole)
remarks: | 2588:70x Announce to uplinks with $x prepend
remarks: | 2588:900 Recieved from LIX (Latvian Internet Exchange)
remarks: |
remarks: | For more information please use the email address
remarks: | iproute (at) latnet (dot) lv
remarks: +-----------------------------------------
```

BGP Community Definitions

©2006-2013 MikroTikSA Pty Ltd

Cogent BGP Communities Example

- Local Preference
- All customer routes announced to Cogent will have a local pref of 130.
- The customer can control the local preference for their announcements by using a community string that is passed to Cogent in the BGP session. The following table lists the community strings and the corresponding local preference that will be set when they are used.

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

©2006-2013 MikroTikSA Pty Ltd

Cogent BGP Communities

- These communities affect what actions Cogent will take on the routes before sending them on to Cogent‘s peering partners.

Community String	Effect
174:3000	Do not announce (equivalent to 174:991)
174:3001	Prepend 174 1 time.
174:3002	Prepend 174 2 times.
174:3003	Prepend 174 3 times

©2006-2013 MikroTikSA Pty Ltd

Communities Received from Cogent

- Routes announced to customers by Cogent will have one of the following communities associated with them:

174:22014	Sweden
174:22015	Norway
174:22016	Czech Republic
174:22017	Slovakia
174:22018	Hungaria
174:22019	Ireland
174:22020	Romania
174:22021	Croatia (locally Hrvatska)
174:22022	Slovenia
174:22023	Bulgaria
174:22024	Finland
174:22025	Estonia
174:22026	Ukraine
174:22027	Mexico

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 an EU internal or customer route.
174:22001	Austria
174:22002	Belgium
174:22003	Canada
174:22004	Switzerland
174:22005	Germany
174:22006	Denmark
174:22007	Spain
174:22008	France
174:22009	Italy
174:22010	Netherlands
174:22011	Portugal
174:22012	United Kingdom
174:22013	United States

©2006-2013 MikroTikSA Pty Ltd

Communities In a nutshell

- Route Advertiser and Route Receiver (ISP Admins) discuss policies and exchange useful information e.g. meaning of Policies etc.
- Route Advertiser (BGP out) sets communities according to a design / policy
- Various Communities are set and sent out with various routes...
- Route Receiver Admin sets Router Receiver to look for set communities in routes and implement policy based on the community.
- Now each ISP is implementing / continuing a policy as agreed with their peer

©2006-2013 MikroTikSA Pty Ltd



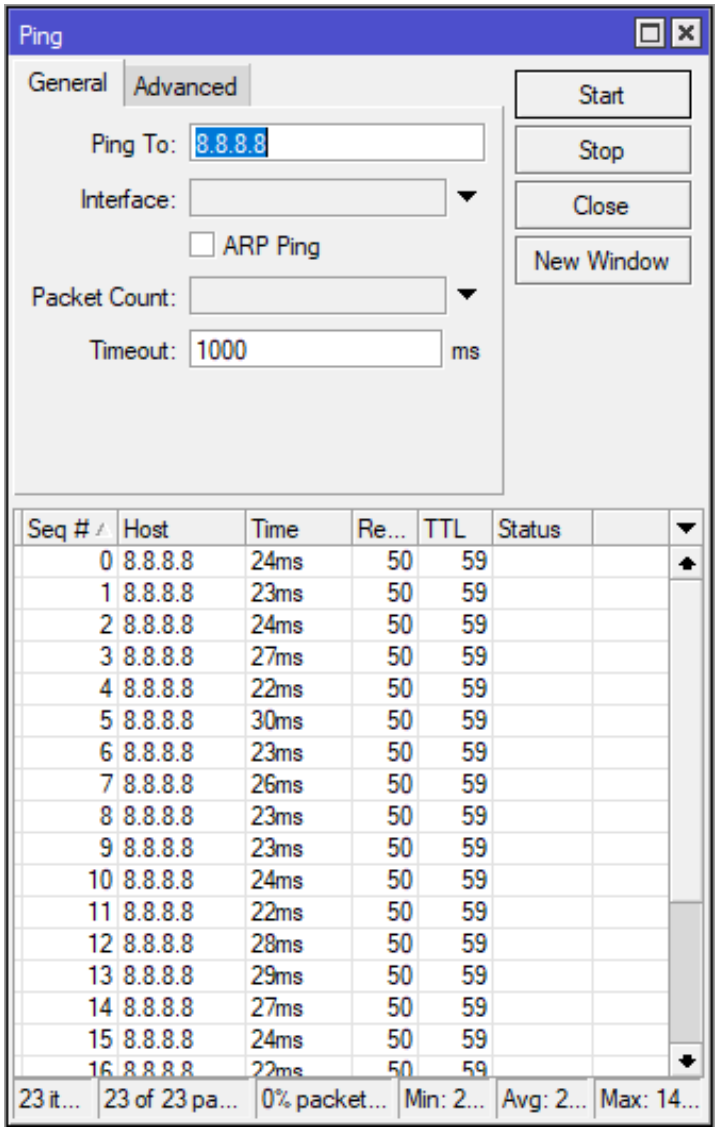
- BGP Communities
  - Set local-pref at ISP
  - Set AS Path Prepend at ISP

Troubleshooting Tools

Checking IP Devices with PING

- Ping is a basic network troubleshooting tool
  - It is based around ICMP – Internet Control Message Protocol which along with Traceroute and Path MTU Discovery form a basis for network troubleshooting
- It shows whether or not a host IP address is contactable
  - Just because a host cannot be pinged does not mean it is not there, a firewall might be preventing ICMP traffic
- It is available from most IP devices
  - From Winbox: Tools → Ping
  - From New Terminal: ping [host_ip or DNS_name]
  - From Windows Command Prompt:

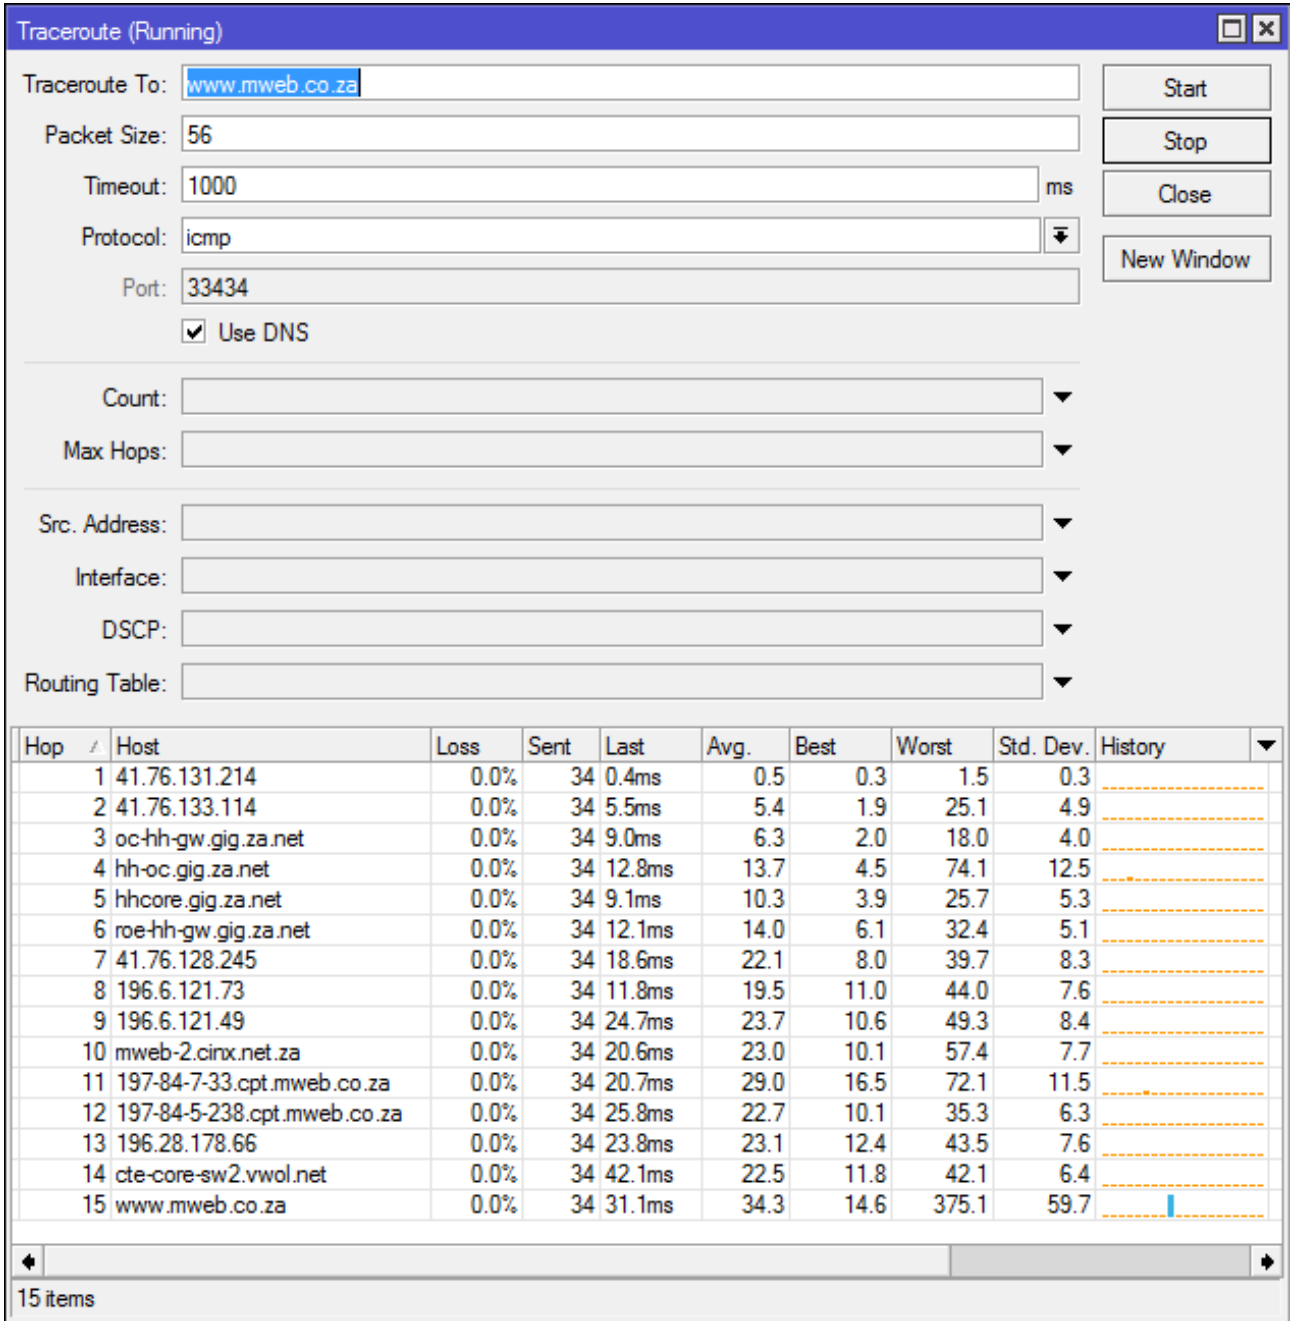
    Start → Run → cmd.exe



MikroTik SA

Traceroute

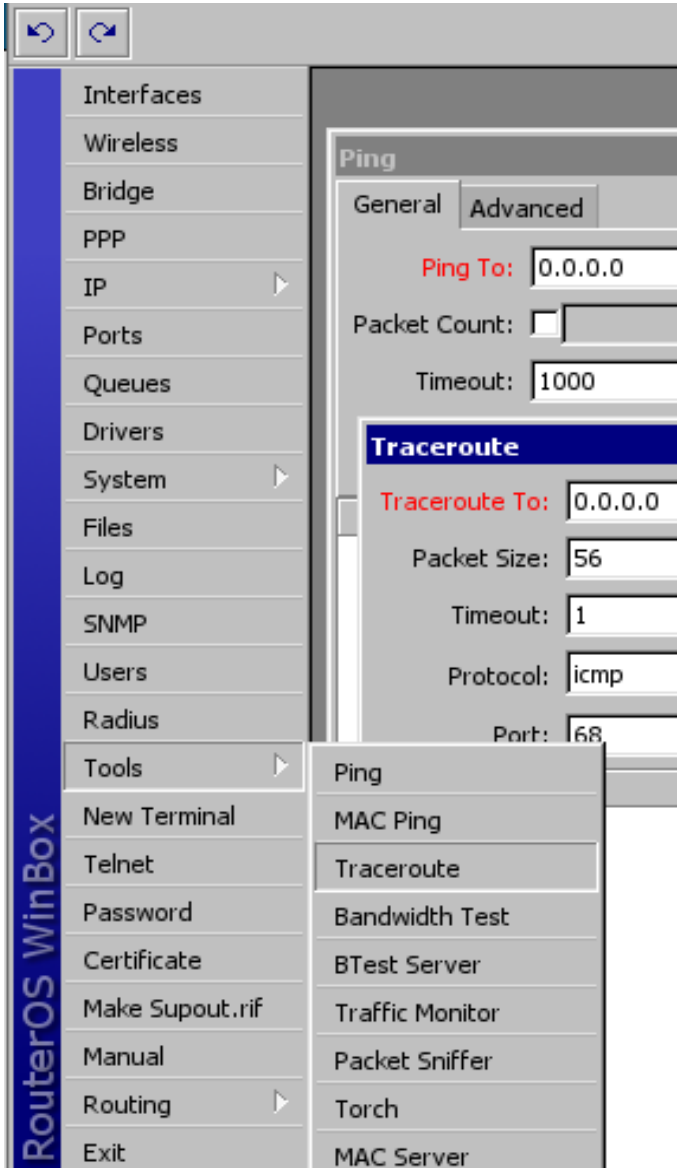
- Traceroute is a utility that records the route (the specific gateway computers at each hop) through the Internet between your computer and a specified destination computer
- It also calculates and displays the amount of time each hop took
- Traceroute is a handy tool both for understanding where problems are in the Internet network and for getting a detailed sense of the Internet itself
- RouterOS traceroute is of the mtr (Multiple TraceRoute) type
  - It will traceroute continuously to show trends in latency
  - Useful for checking latency spikes, jitter, dropouts



Traceroute

- Use DNS – reverse lookup of IP's
- Count – iterations of MTR
- Max Hops – TTL
- Source Address – useful to troubleshoot routing/NAT
- Interface – outgoing interface
- DSCP – QOS pointer
- Routing table – use custom route marks

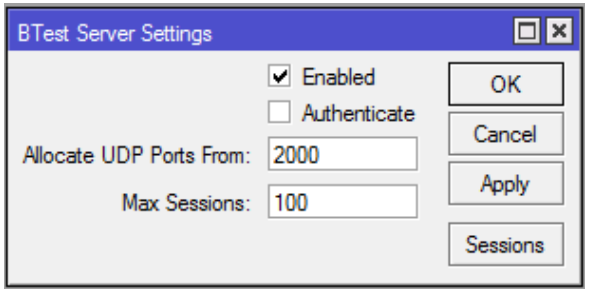
Network Management Tools



- Use ping to check the presence of systems on the network by IP address
- Traceroute can be used to track down routing problems on the network
- Bandwidth test can be used to check the performance of routers and connections
- IP scan can be used to check for multiple network hosts
- All utilities can be found in the Tools menu

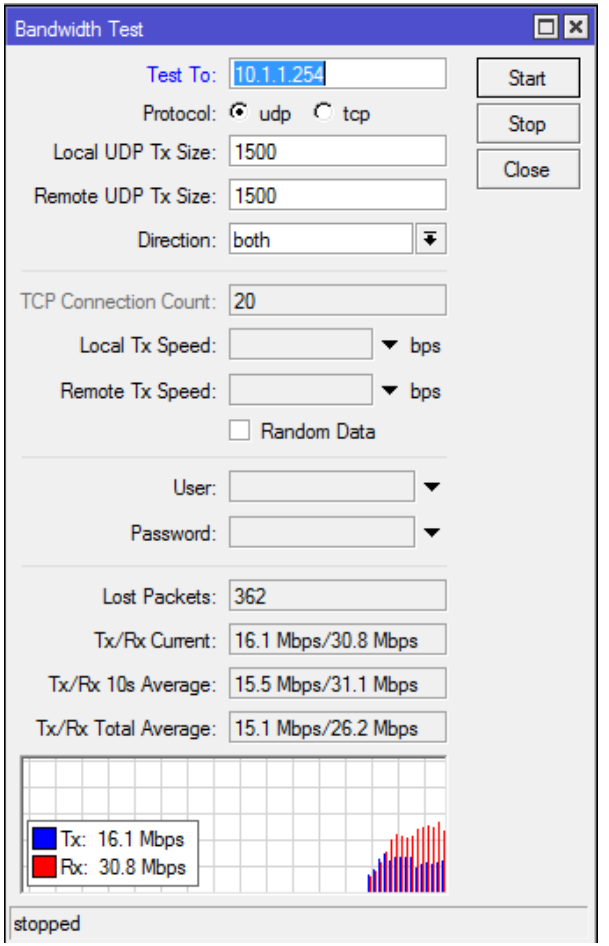
Bandwidth Test Server

- Tools → Btest Server controls whether the service is available
- Enable / disable as required
- Authenticate controls whether a user / password is required or if the server allows testing with no authentication
  - Note since ROS 6.42 authenticated btest is not compatible with previous versions
- Users require at least Winbox, read and test permissions to run test if Authenticate is enabled



Bandwidth Test

- Useful for checking link performance
- Large CPU overhead for running test with TCP
  - Improved multicore performance from 6.44
- UDP tests asynchronous speeds
  - Useful for checking wireless sync rate vs actual performance
- TCP checks “real world” performance
  - Single connection is useful for checking with torch
  - Use multiple connections to simulate normal network activity
  - Random data simulates different protocols and places a high load on the system
- Local and Remote Tx Speed can be used to place a set load on the link
  - This is useful for modelling link performance under different bandwidth conditions



MikroTik SA

Monitoring the Network Traffic

MikroTik RouterOS tools for monitoring the network traffic:

- Interface tx/rx bits/s and packets/s, numbers and graphs
- Torch tool for more detailed traffic report through an interface
- Sniffer for capturing data about Ethernet packets transmitted on a LAN segment connected to an interface
- Firewall logs and connection tracking tables
- Interface traffic graphs
- IP Scanner tool

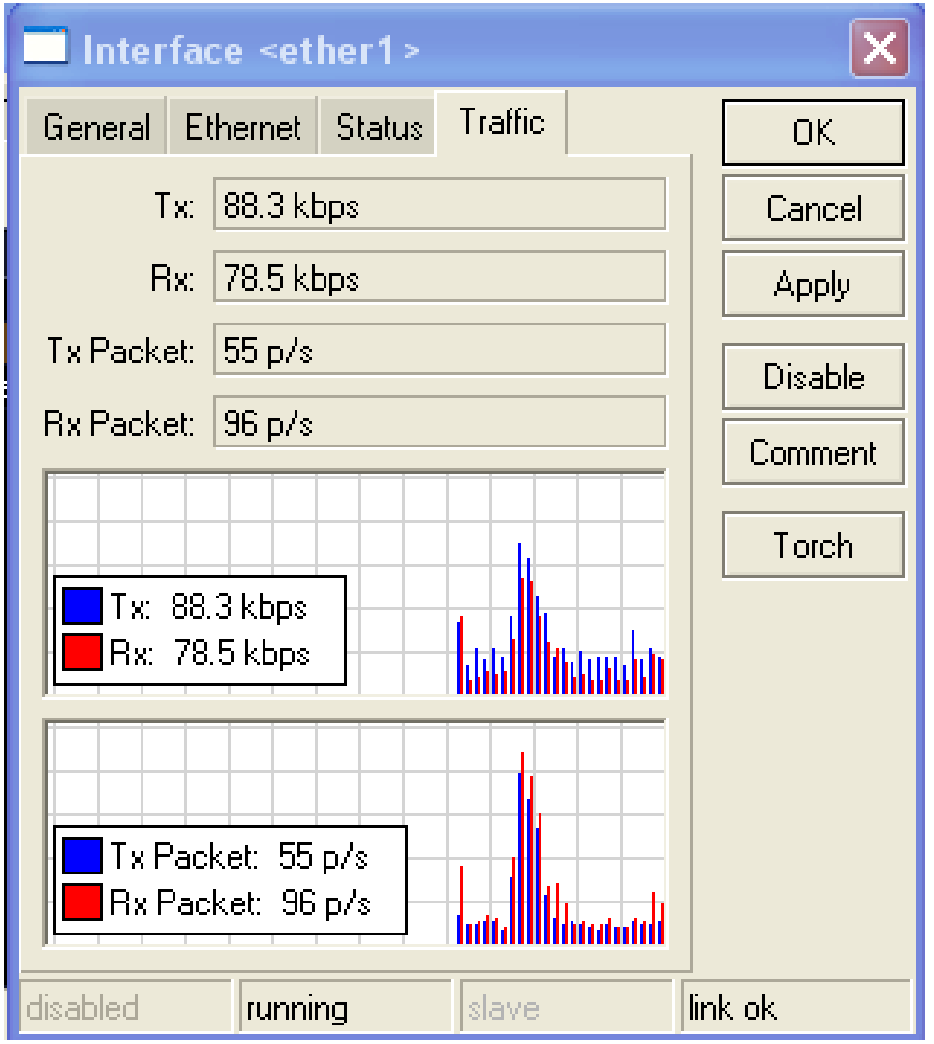
Interface Traffic Monitor

You can open Interfaces in winbox to see tx/rx rates

Open any interface and select the “Traffic” tab to see the graphs and real time speed

Use the “monitor-traffic” command in terminal to get the traffic data per one or more interfaces, for example:

- /interface monitor-traffic ether1
- /interface monitor-traffic ether1,ether2,ether3

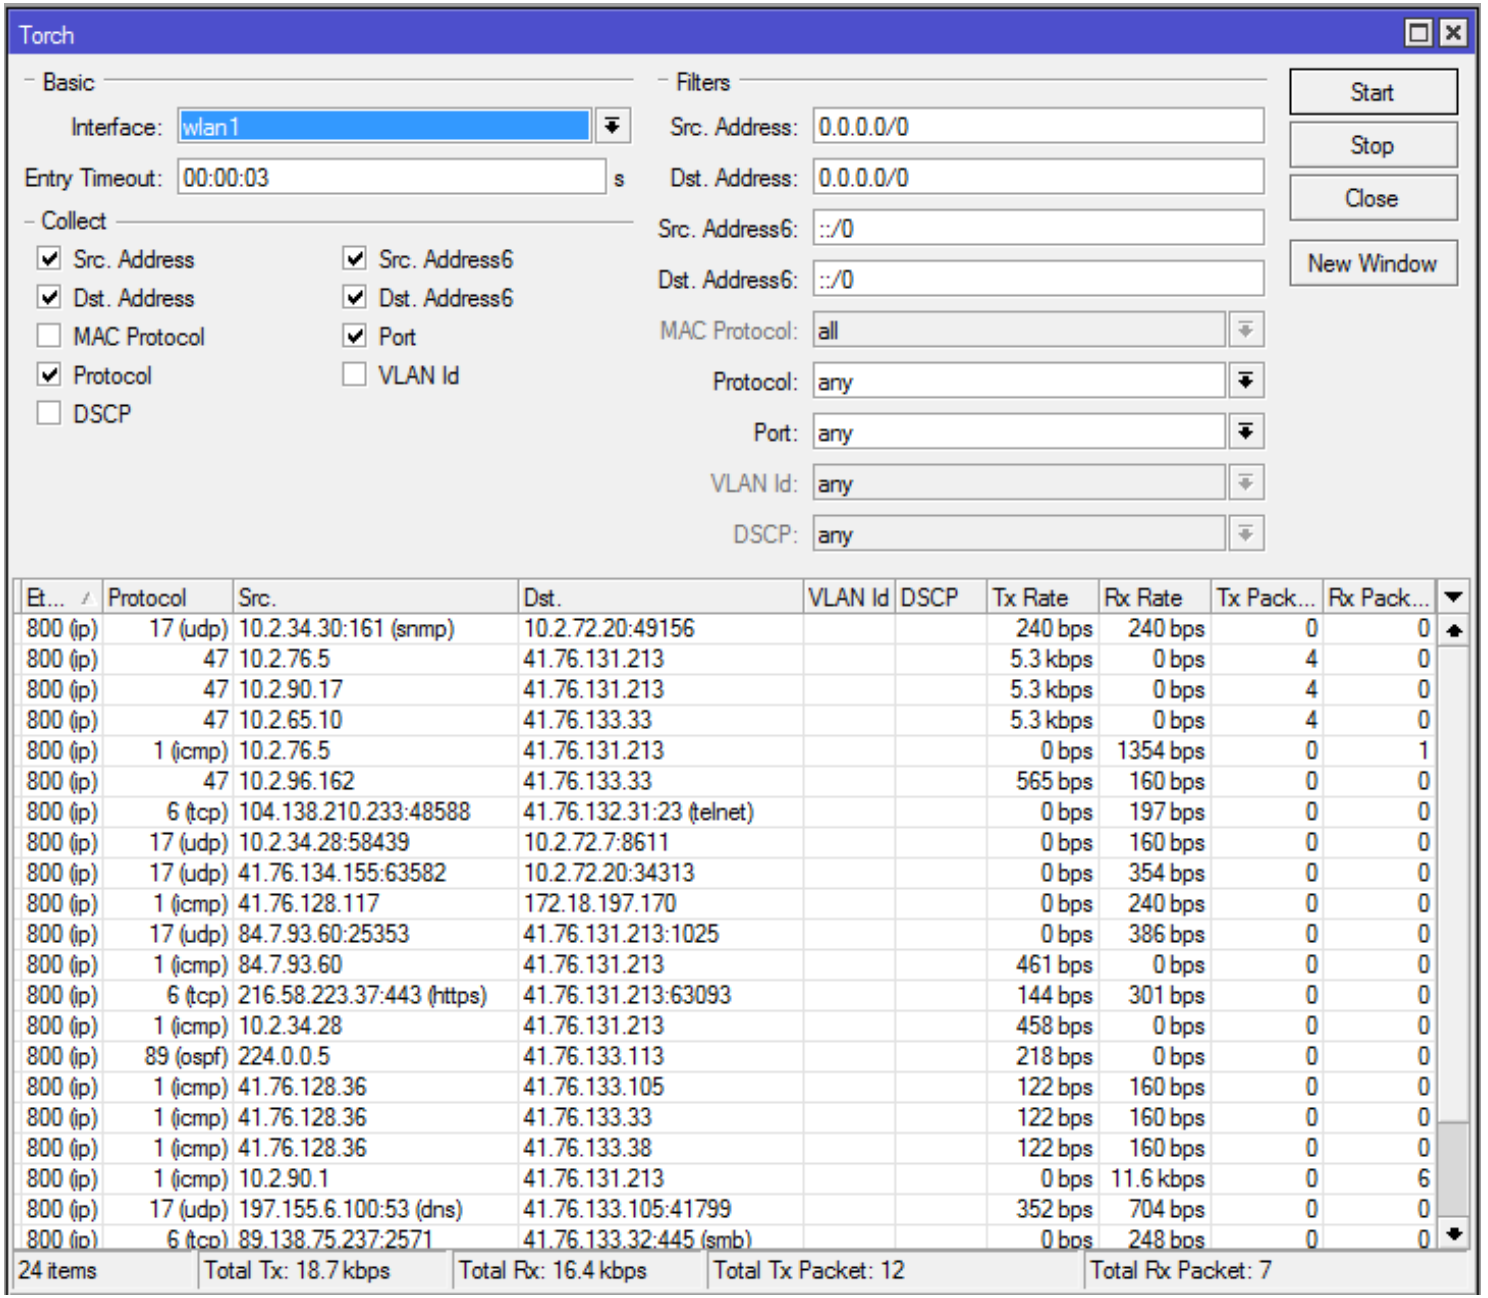


Torch Tool

Torch tool offers a detailed actual traffic report for an interface

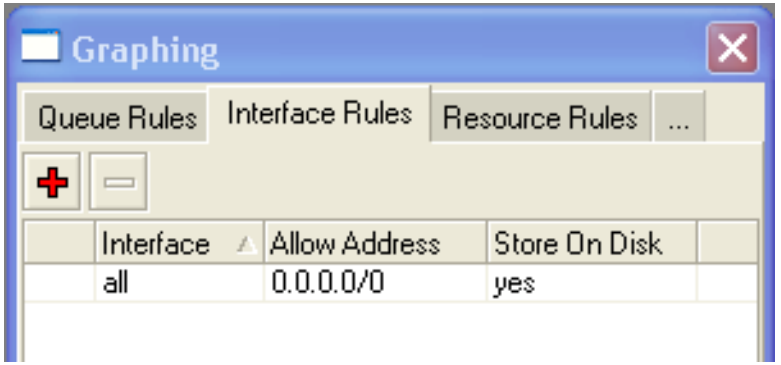
It's easiest to use torch in winbox:

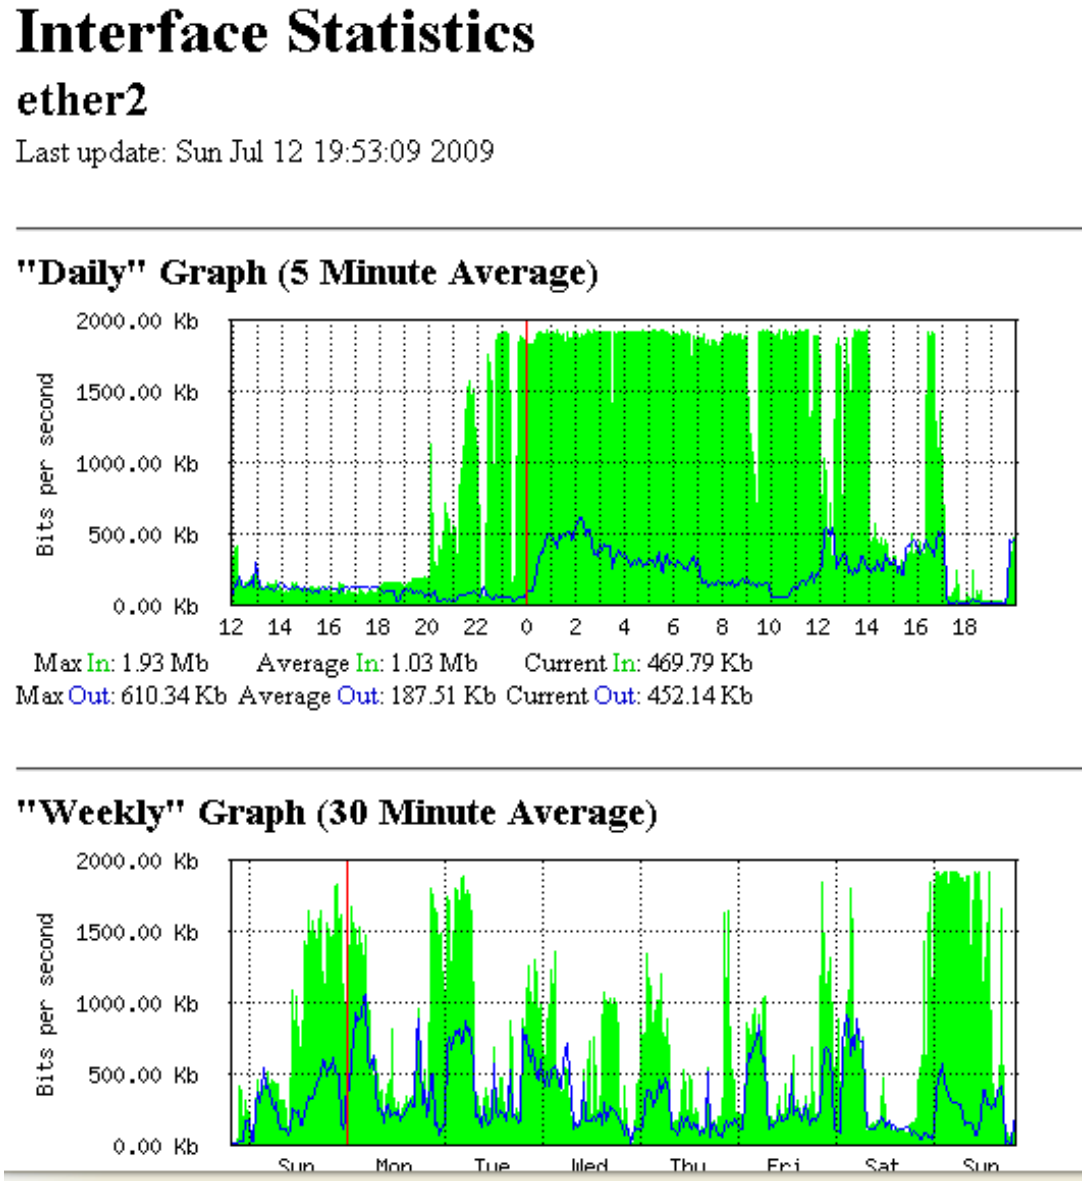
- Go to “Tools” > “Torch”
- Select an interface to monitor and click “Start”
- Use “Stop” and “Start” to freeze/continue
- Refine the display by selecting protocol and port
- Double-click on specific IP address to fill in the Src. Or Dst. Address field (0.0.0.0/0 is for any address)
- Can also identify VLAN tag and DSCP (QOS) markers


Torch
Basic
Interface: wlan1
Entry Timeout: 00:00:03 s
Collect
Src. Address
Dst. Address
MAC Protocol
Protocol
DSCP
Src. Address6
Dst. Address6
Port
VLAN Id
Filters
Src. Address: 0.0.0.0/0
Dst. Address: 0.0.0.0/0
Src. Address6: ::/0
Dst. Address6: ::/0
MAC Protocol: all
Protocol: any
Port: any
VLAN Id: any
DSCP: any
Start
Stop
Close
New Window
Et... / Protocol Src. Dst. VLAN Id DSCP Tx Rate Rx Rate Tx Pack... Rx Pack...
800 (ip) 17 (udp) 10.2.34.30:161 (snmp) 10.2.72.20:49156 240 bps 240 bps 0 0
800 (ip) 47 10.2.76.5 41.76.131.213 5.3 kbps 0 bps 4 0
800 (ip) 47 10.2.90.17 41.76.131.213 5.3 kbps 0 bps 4 0
800 (ip) 47 10.2.65.10 41.76.133.33 5.3 kbps 0 bps 4 0
800 (ip) 1 (icmp) 10.2.76.5 41.76.131.213 0 bps 1354 bps 0 1
800 (ip) 47 10.2.96.162 41.76.133.33 565 bps 160 bps 0 0
800 (ip) 6 (tcp) 104.138.210.233:48588 41.76.132.31:23 (telnet) 0 bps 197 bps 0 0
800 (ip) 17 (udp) 10.2.34.28:58439 10.2.72.7:8611 0 bps 160 bps 0 0
800 (ip) 17 (udp) 41.76.134.155:63582 10.2.72.20:34313 0 bps 354 bps 0 0
800 (ip) 1 (icmp) 41.76.128.117 172.18.197.170 0 bps 240 bps 0 0
800 (ip) 17 (udp) 84.7.93.60:25353 41.76.131.213:1025 0 bps 386 bps 0 0
800 (ip) 1 (icmp) 84.7.93.60 41.76.131.213 461 bps 0 bps 0 0
800 (ip) 6 (tcp) 216.58.223.37:443 (https) 41.76.131.213:63093 144 bps 301 bps 0 0
800 (ip) 1 (icmp) 10.2.34.28 41.76.131.213 458 bps 0 bps 0 0
800 (ip) 89 (ospf) 224.0.0.5 41.76.133.113 218 bps 0 bps 0 0
800 (ip) 1 (icmp) 41.76.128.36 41.76.133.105 122 bps 160 bps 0 0
800 (ip) 1 (icmp) 41.76.128.36 41.76.133.33 122 bps 160 bps 0 0
800 (ip) 1 (icmp) 41.76.128.36 41.76.133.38 122 bps 160 bps 0 0
800 (ip) 1 (icmp) 10.2.90.1 41.76.131.213 0 bps 11.6 kbps 0 6
800 (ip) 17 (udp) 197.155.6.100:53 (dns) 41.76.133.105:41799 352 bps 704 bps 0 0
800 (ip) 6 (tcp) 89.138.75.237:2571 41.76.133.32:445 (smb) 0 bps 248 bps 0 0
24 items
Total Tx: 18.7 kbps
Total Rx: 16.4 kbps
Total Tx Packet: 12
Total Rx Packet: 7

Graphing

- You can add rules to give you basic graphing for router interfaces
- Use Tools → Graphing to add support for interface graphing
- Use the web interface to view graphs

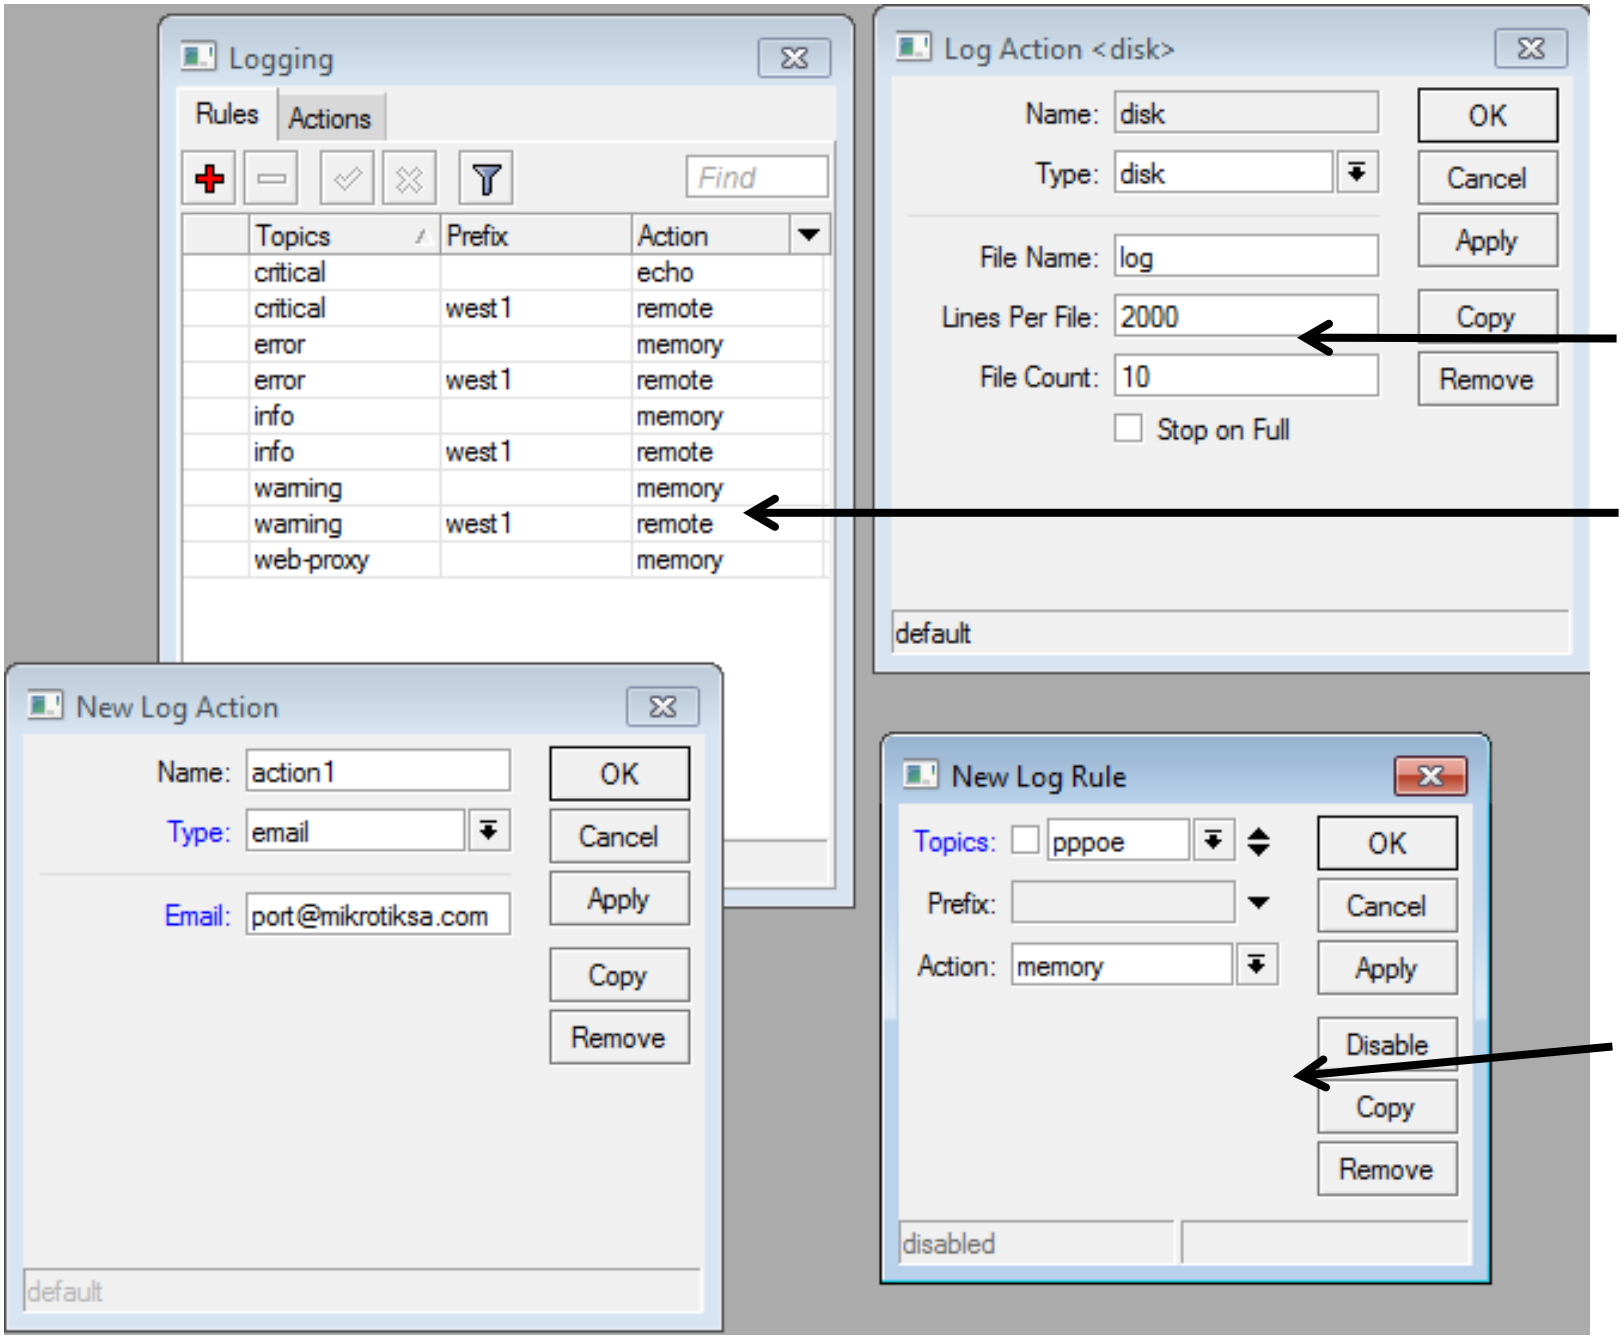




MikroTik SA

Logging

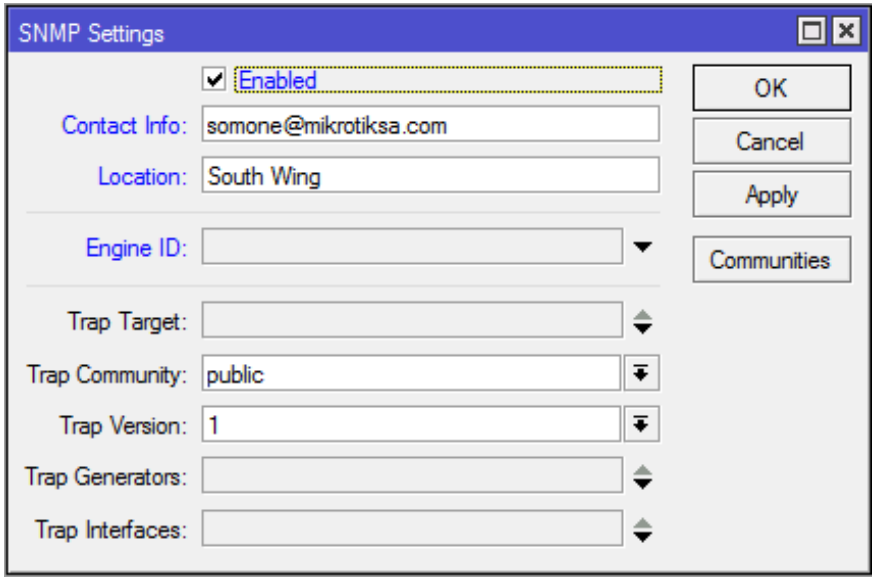
- System → Logging allows you to configure logging options for the router
- You can log to memory, disk, email and remote syslog
- You can log all functions of RouterOS
- All messages stored in the routers local memory can be printed from /log menu.
  - Each entry contains the time and date when event occurred, topics that the message belongs to and the message itself.
- If you have enough memory/diskspace you should increase the number of log entries stored


Logging
Rules
Actions
Find
Topics
Prefix
Action
critical
echo
critical
west1
remote
error
memory
error
west1
remote
info
memory
info
west1
remote
warning
memory
warning
west1
remote
web-proxy
memory
Log Action <disk>
Name: disk
Type: disk
File Name: log
Lines Per File: 2000
File Count: 10
Stop on Full
OK
Cancel
Apply
Copy
Remove
default
New Log Action
Name: action1
Type: email
Email: port@mikrotiksa.com
OK
Cancel
Apply
Copy
Remove
default
New Log Rule
Topics: pppoe
Prefix:
Action: memory
OK
Cancel
Apply
Disable
Copy
Remove
disabled
Increase line and file count
Add multiple rules and locations
Add new rules

MikroTik SA

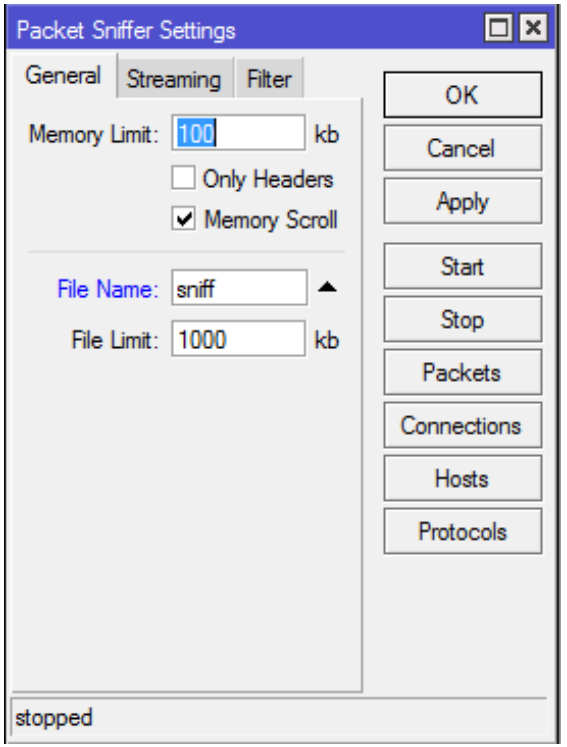
SNMP

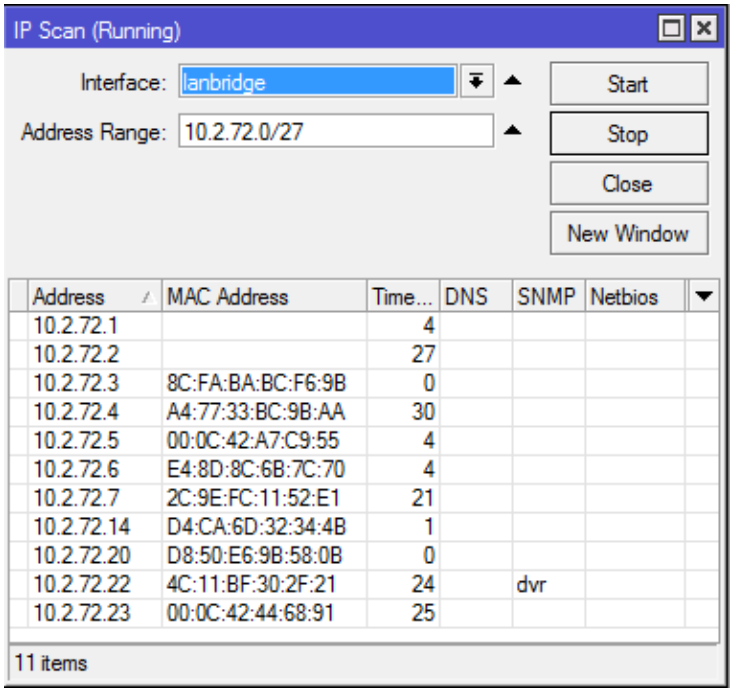
- Simple Network Management Protocol (SNMP) is an Internet-standard protocol for managing devices on IP networks
- SNMP can be used to graph various data with tools such as CACTI, MRTG or The Dude
- Enabled in IP → SNMP



IP Scan and Sniffer

- Tool → Packet Sniffer allows capturing of packets moving through an interface
  - Generally used with applications like Wireshark to examine packet headers
- Tool → IP Scan can scan an IP range to determine which devices are responding
  - Ranges can be local or remote



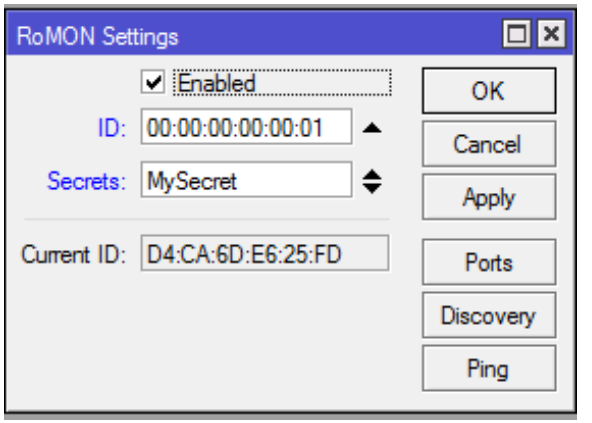


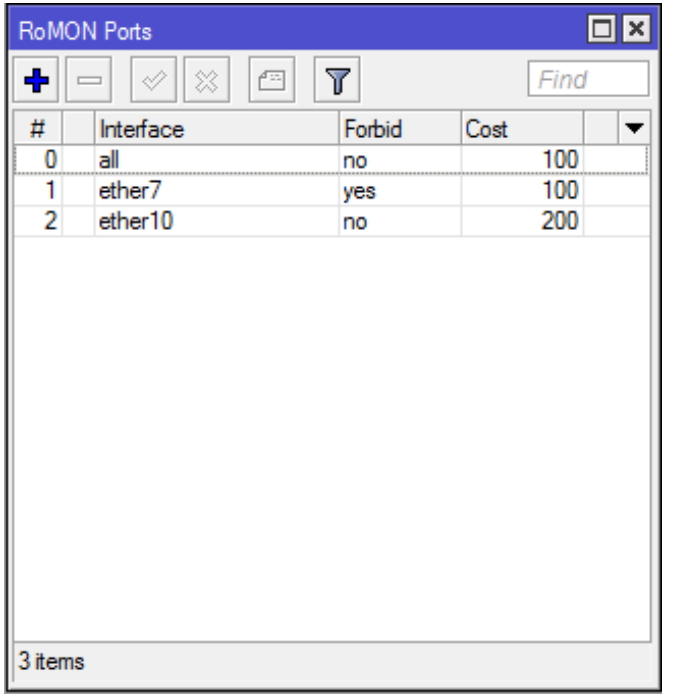
RoMON

- RoMON - Router Management Overlay Network
- RoMON works by establishing an independent MAC layer peer discovery and data forwarding network
- RoMON network operates independently from L2 or L3 forwarding configuration
- Each router on the RoMON network is assigned a RoMON ID
- The RoMON ID can be selected from port MAC address or specified by user
- The RoMON protocol does not provide encryption services except for actual RoMON comms
  - Encryption is provided at "application" level, by e.g. using ssh or by using secure winbox

Configure RoMON

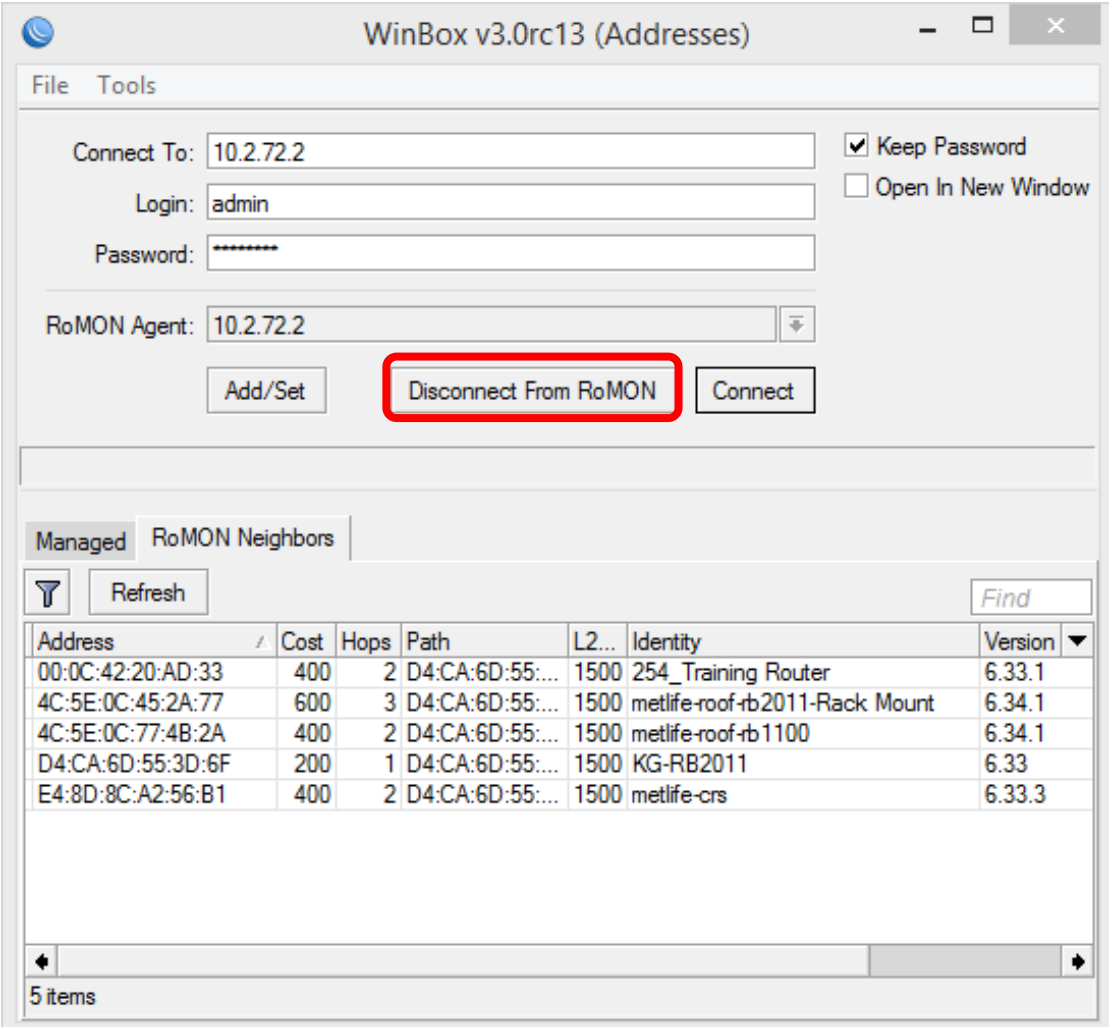
- Tool → RoMON allows the service to be enabled/disabled
- ID can optionally be specified otherwise default is ether1 MAC
- Secrets will encrypt RoMON comms with MD5 / AES – secret must be the same for adjacent ports
- RoMON Ports allows setting up ports individually with costs





Connecting to RoMON

- Winbox V3 must be used
- Select a RoMON enabled router and choose “Connect to RoMON”
- RoMON enabled routers will now be displayed



MikroTik SA

Contact Us

General Training Queries

training@mikrotiksa.com

021 557 6868

David Savage

David@mikrotiksa.com

083 456 2002

©2006-2013 MikroTikSA Pty Ltd