
MIKROTIK
professionals conference®

RouterOS v6 to v7 – *The* Routing guide!

According to …

Barry Higgins

Key to symbols used

If you see this symbol -  please laugh!

Presenter bio…

Name: Barry Higgins.

DoB: Before Arpnet

Country of Origin: England (UK)

Trainer: TR0372 (Feb 2016)

Likes: yes please

Dislikes: Top posting!

Hobbies: Hiking

Shameless plug #1 - Consultancy

MikroTik Consultant since 2015

- Consultancy available – Many aspects covered including connection of MikroTik → ‘other brands’
- Find details at https://www.allness.net

Shameless plug #2 - Training

MikroTiK Trainer : Since 2016

- Public sessions held in the UK and Ireland

- Private sessions can be arranged!

- All countries covered so long as you want it in English!
  To date I have trained in Germany, Austria, Norway, Latvia, Italy (Ron & Lorenzo were there), Wales, Scotland, Ireland and now Czechia (Ron and Lorenzo followed me here too!).

Shameless plug #2 – Training – the BUS!

Shameless plug #3 cont...

Shameless plug #3 cont…

Shameless plug #3 cont...

Shameless plug #3 – Riga BootCamp

Riga Bootcamp summer 2024…

21-29 June 2024

Shameless plug #3 – Riga BootCamp

Riga Bootcamp summer 2024…

https://www.mikrotik.camp/

21-29 June 2024

Onward with v7 we go...

And so version 7 continues to roll out week after week with new updates, features and even undocumented ones (bugs)!

Love it or hate it, there is no choice with the roll out of newer routerboard models requiring v7 due to Arm(64) architecture.

V6 is no more. Time to move on!

Routing has changed

There have been lots of changes in how routing is handled in version 7.

A newer linux kernel has meant that how routing was being handled in the ‘backend’ now requires a new ‘frontend’ to handle the changes.

A quick guide to the changes

Buckle up its going to be a fast ride!

(I hear Lorenzo shouting <NEXT> already) 

Winbox → Routing

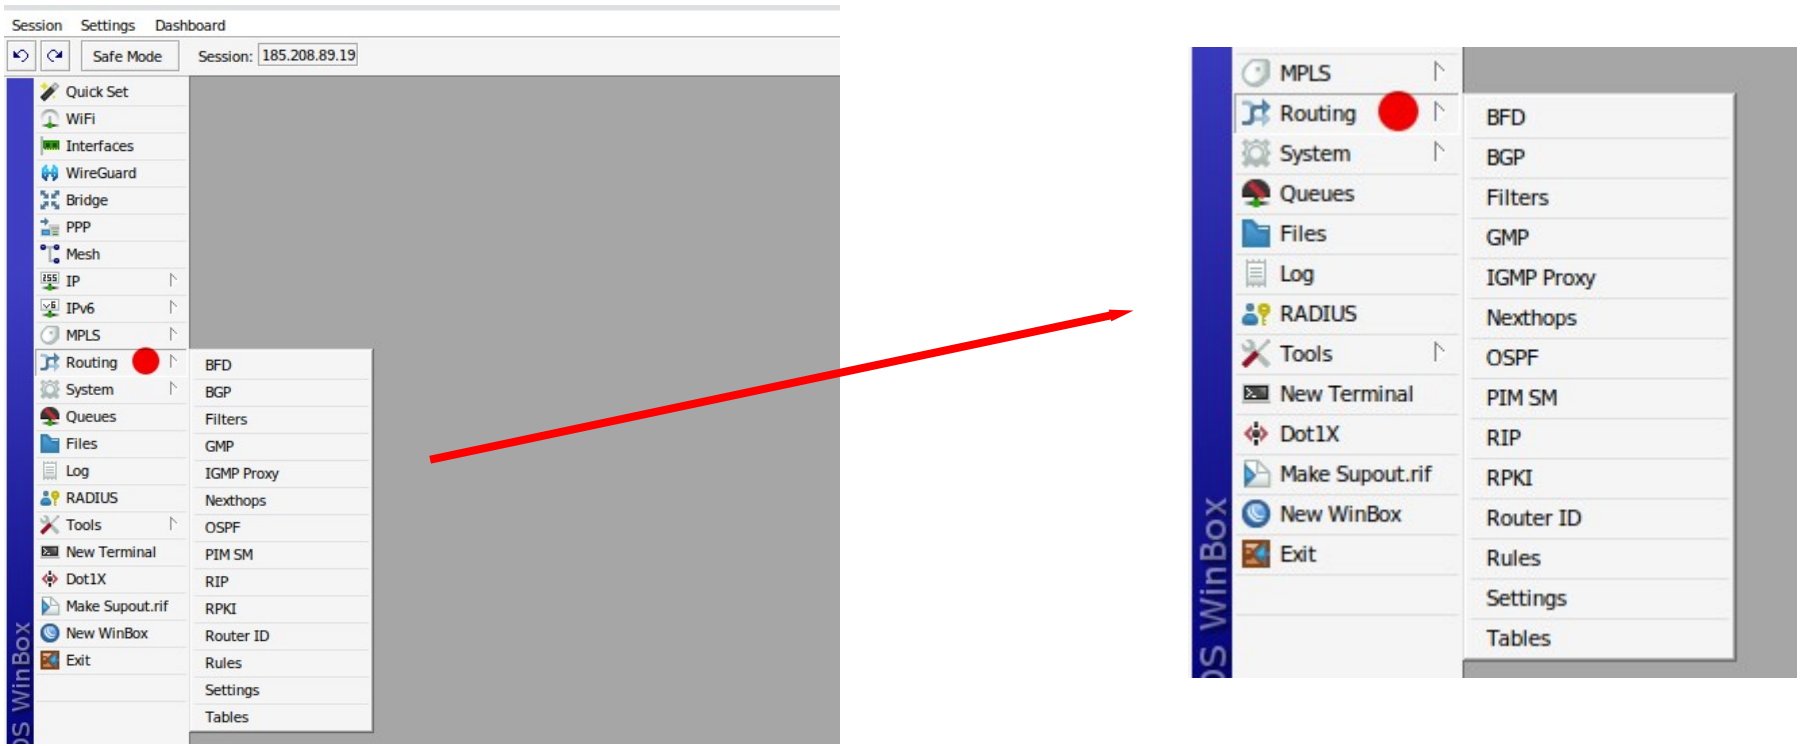

BFD

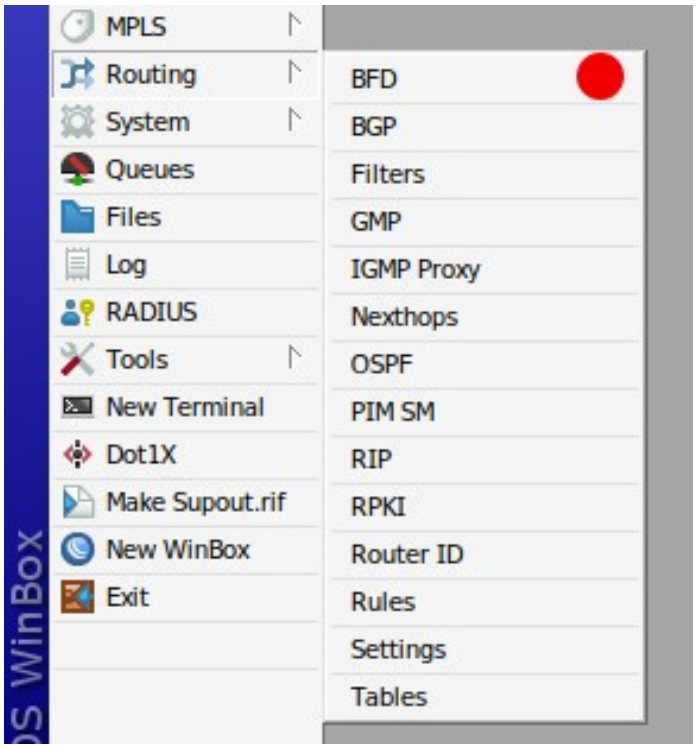

BFD

Bidirectional Forwarding Detection

- A protocol to monitor link failure.
- Reintroduced into RouterOS in v7.10
- v7 now includes RFC 5883 in addition to 5880, 5881 & 5882

BGP

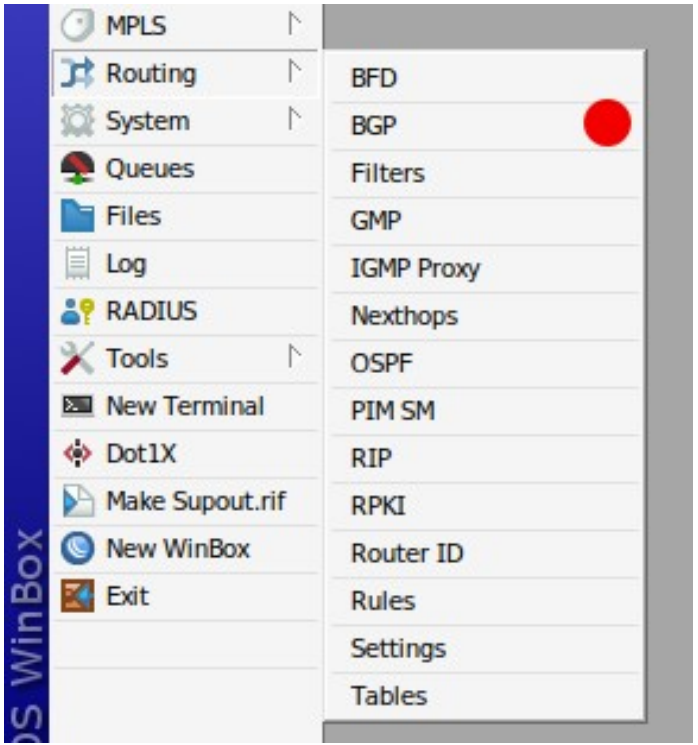

BGP

Border Gateway Protocol

(Or as Sky TV experts call it - ‘Bridging Gap Protocol’ !!)

- 16:00 – 16:45 - Guilherme Ramirez

Filters

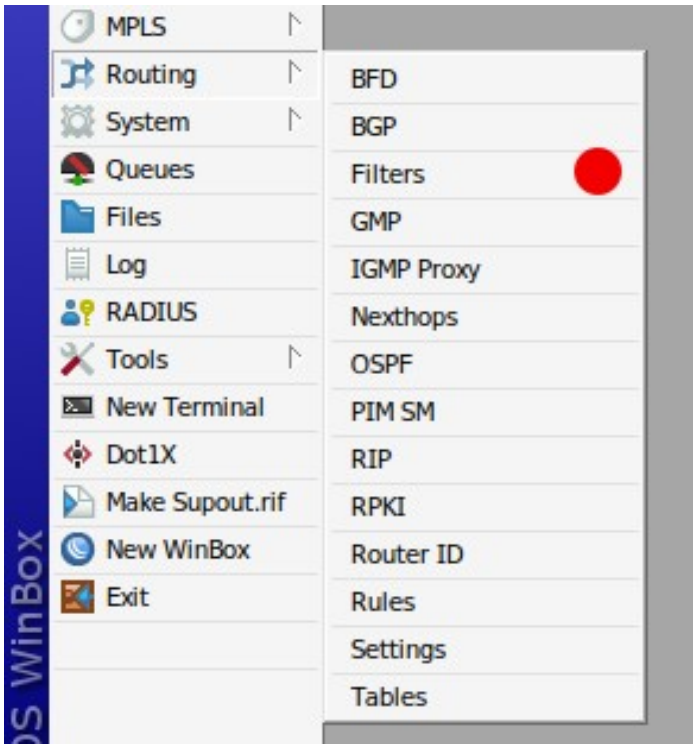

Filters

Filter creation has changed. No more ticking of boxes and selection options.

Filters are now created using a type of scripting using the formula of -

“if ([MATCH]) then { [ACTION] } else { [ACTION] }”

Filters – syntax example

‘( )’ encloses the matching criteria

‘{ }’ encloses the actions should a match occur.

e.g.

```
if ( dst in 10.0.0.0/8 || 192.168.0.0/16 ) {reject} else {set distance 5; accept}
```

Filters – filter window

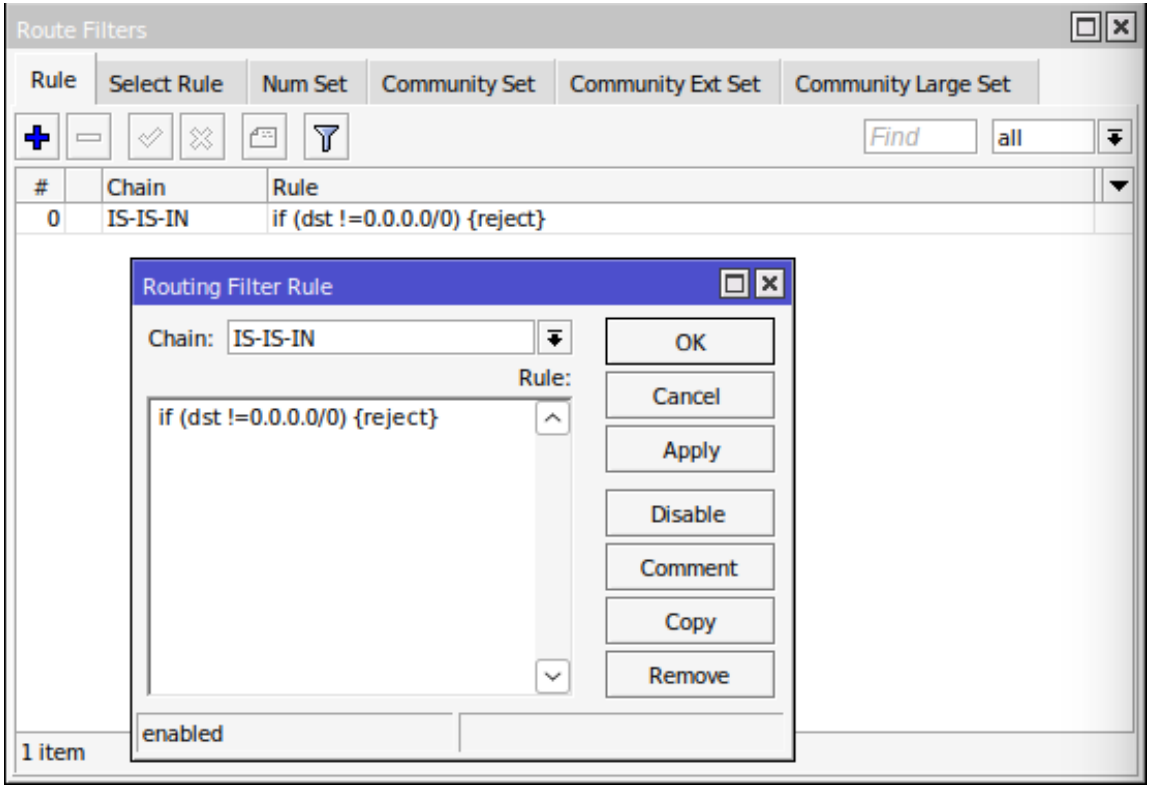

Filters

GMP

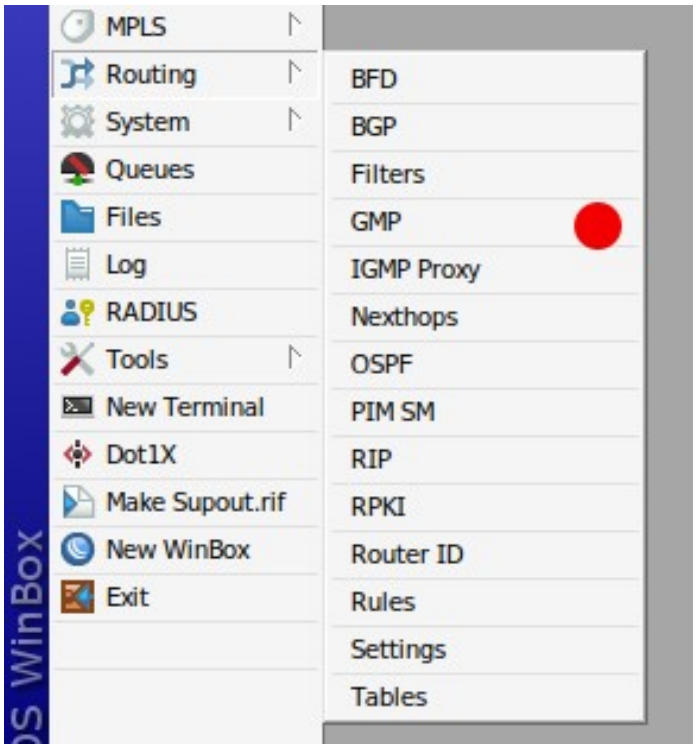

GMP

Group Management Protocol

Multicasting – no time to cover this

<NEXT> !

IGMP Proxy

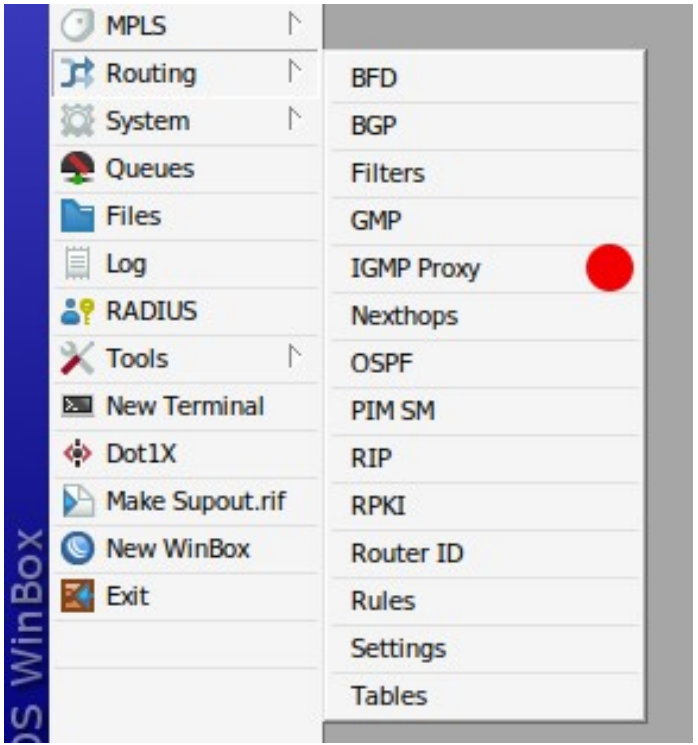

IGMP Proxy

Internet Group Management Protocol

Multicasting – no time to cover this either

<NEXT> !

Nexthops

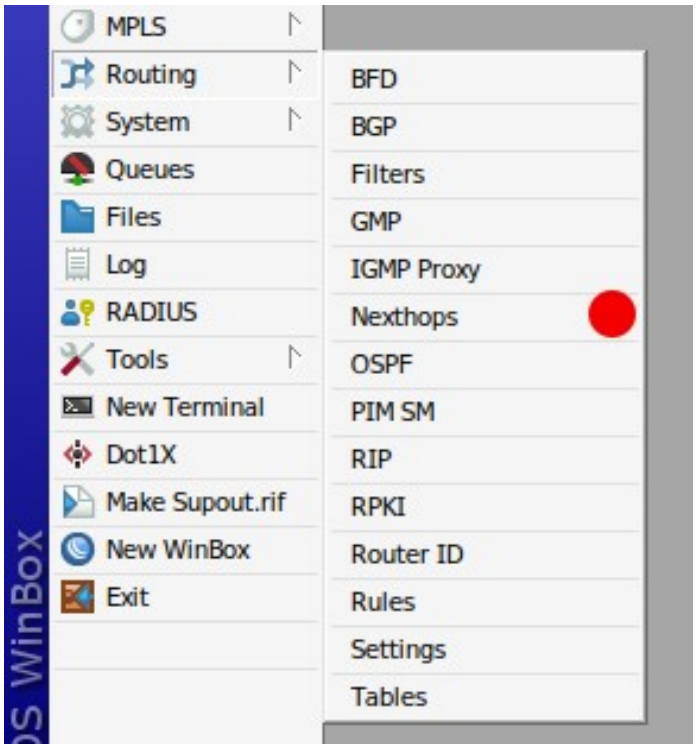

Nexthops

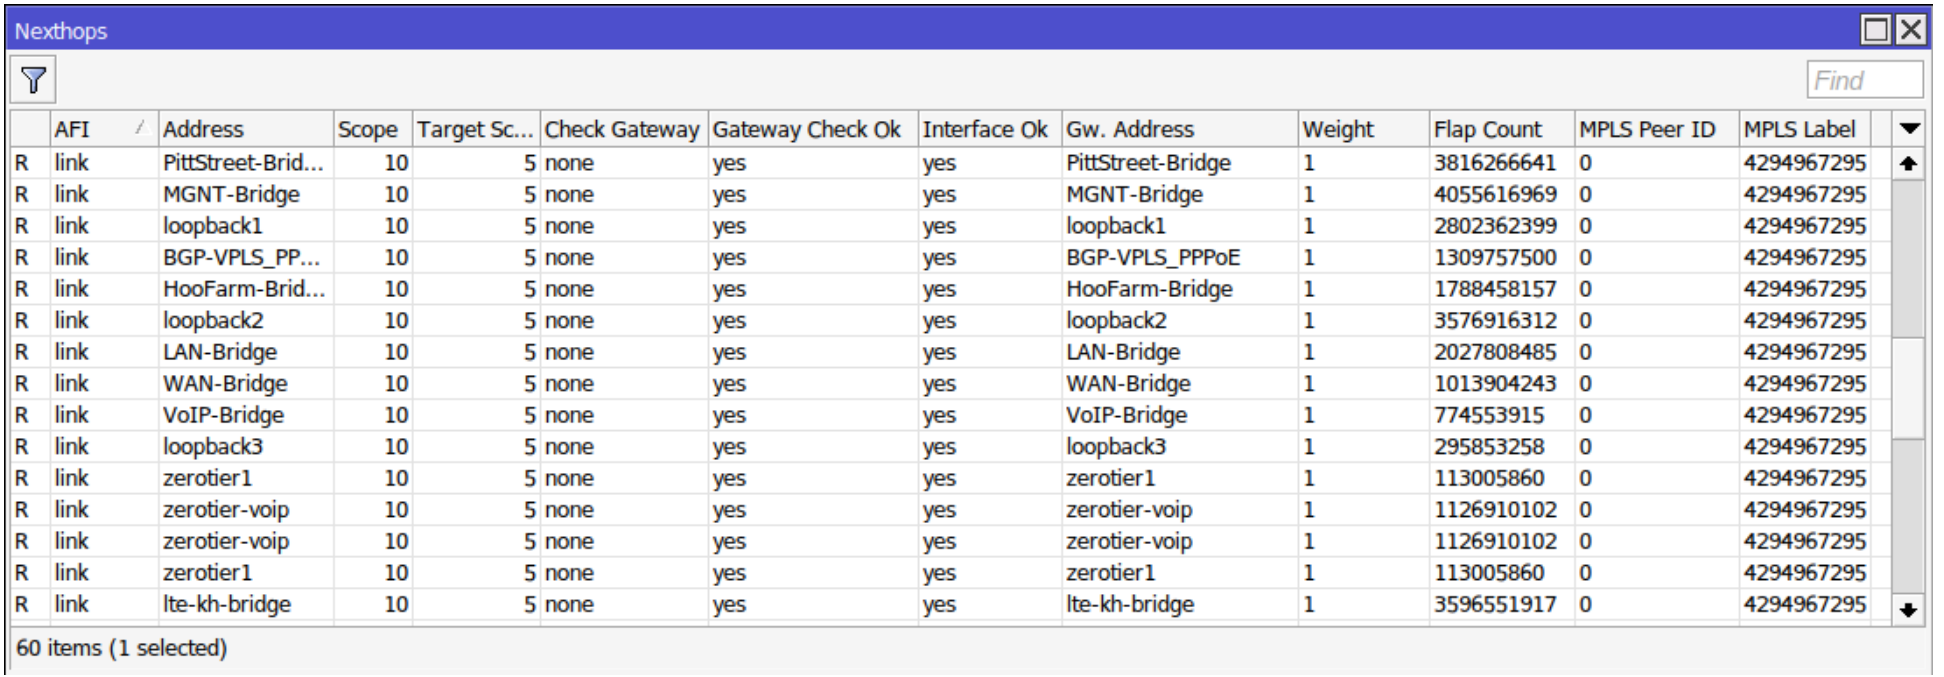

	AFI	Address	Scope	Target Sc...	Check Gateway	Gateway Check Ok	Interface Ok	Gw. Address	Weight	Flap Count	MPLS Peer ID	MPLS Label
R	link	PittStreet-Brid...	10	5	none	yes	yes	PittStreet-Bridge	1	3816266641	0	4294967295
R	link	MGNT-Bridge	10	5	none	yes	yes	MGNT-Bridge	1	4055616969	0	4294967295
R	link	loopback1	10	5	none	yes	yes	loopback1	1	2802362399	0	4294967295
R	link	BGP-VPLS_PP...	10	5	none	yes	yes	BGP-VPLS_PPPoE	1	1309757500	0	4294967295
R	link	HooFarm-Brid...	10	5	none	yes	yes	HooFarm-Bridge	1	1788458157	0	4294967295
R	link	loopback2	10	5	none	yes	yes	loopback2	1	3576916312	0	4294967295
R	link	LAN-Bridge	10	5	none	yes	yes	LAN-Bridge	1	2027808485	0	4294967295
R	link	WAN-Bridge	10	5	none	yes	yes	WAN-Bridge	1	1013904243	0	4294967295
R	link	VoIP-Bridge	10	5	none	yes	yes	VoIP-Bridge	1	774553915	0	4294967295
R	link	loopback3	10	5	none	yes	yes	loopback3	1	295853258	0	4294967295
R	link	zerotier1	10	5	none	yes	yes	zerotier1	1	113005860	0	4294967295
R	link	zerotier-voip	10	5	none	yes	yes	zerotier-voip	1	1126910102	0	4294967295
R	link	zerotier-voip	10	5	none	yes	yes	zerotier-voip	1	1126910102	0	4294967295
R	link	zerotier1	10	5	none	yes	yes	zerotier1	1	113005860	0	4294967295
R	link	lte-kh-bridge	10	5	none	yes	yes	lte-kh-bridge	1	3596551917	0	4294967295

60 items (1 selected)

Nexthops

A new addition to aid visual/informative details of valid nexthops

OSPF

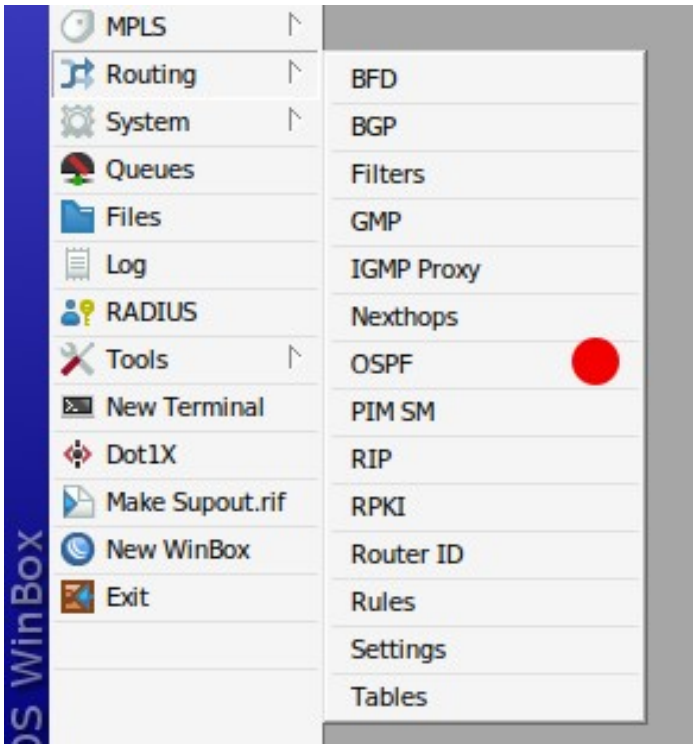

OSPF – Open Shortest Path First

OSPF in v7… where the &*$^ do we start?!!

Without doubt, along with BGP, it is a complete change in how configuration is constructed compared to v6.

- Originally OSPF in v7 was CLI only.
- Thankfully it is now configurable in winbox.

OSPF

The first major change to notice is that OSPFv2 and OSPFv3 are now configured in the same place.
First define which OSPF version to use in the ‘Instance’

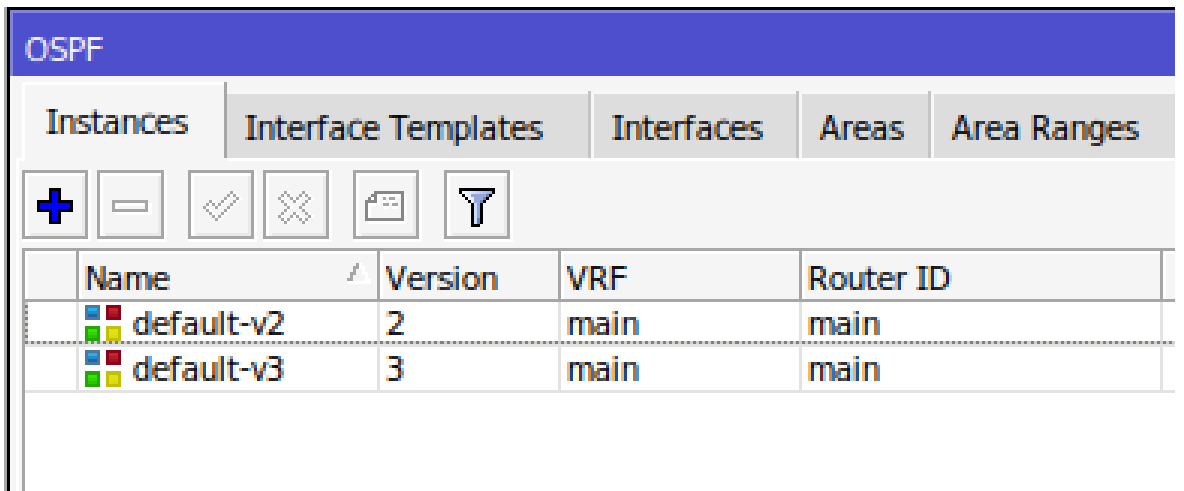

OSPF

Next you need to configure an area as there is no default ‘backbone’ area.

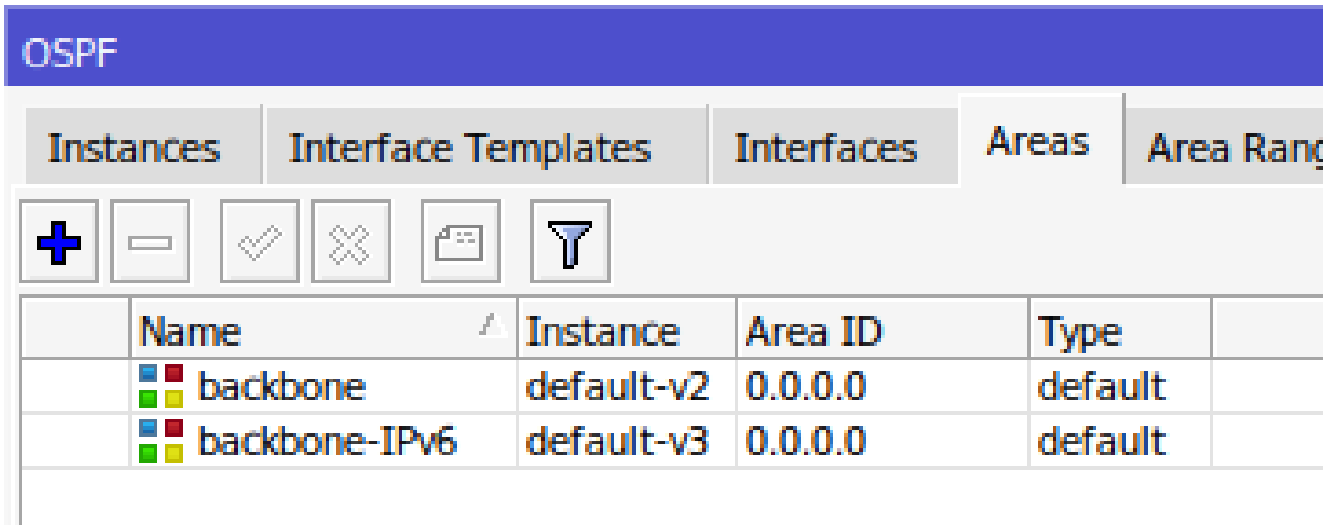

OSPF

You are now at a configuration point similar to how v6 would be presented ‘out of the box’.

In v6 if you wanted OSPF to start running you would add an active subnet running on your router

OSPF – v6 ‘Networks’

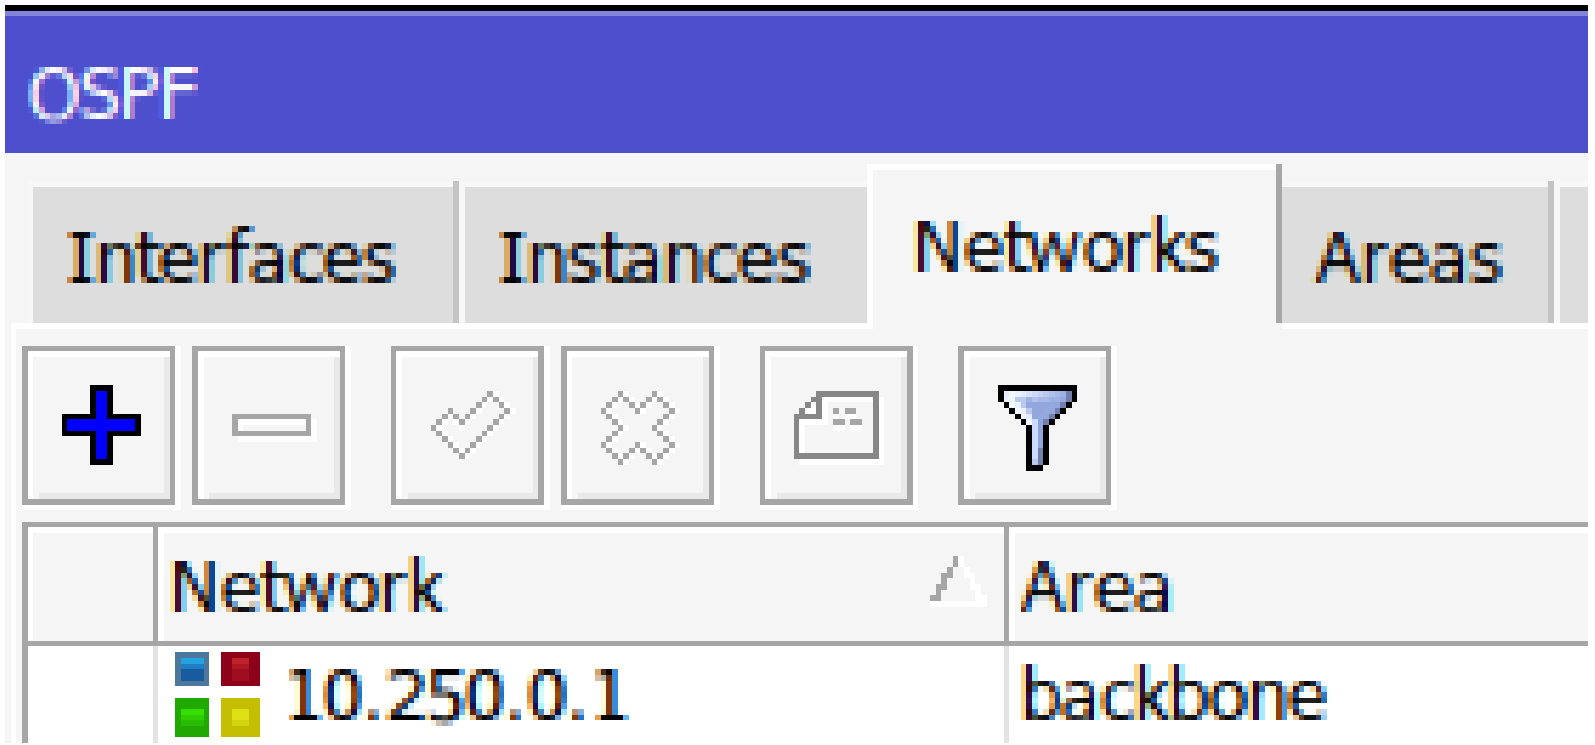

OSPF – v6 Instances (running = yes!)

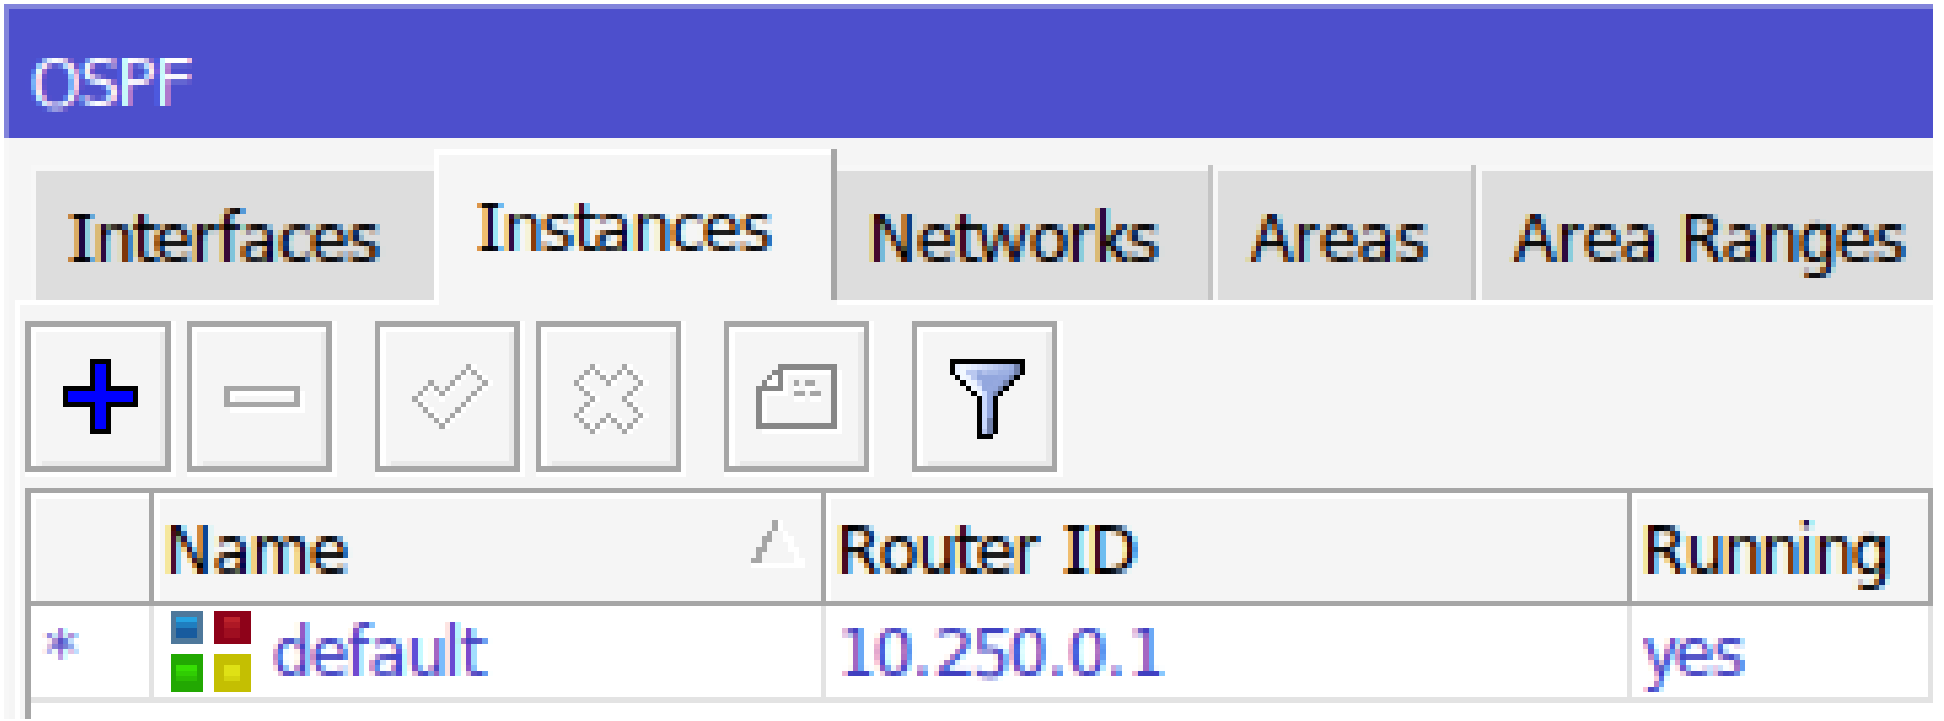

OSPF

With v7 you need to add an ‘interface template’ which has combined both the v6 ‘interface’ and ‘network’ tab as one.

OSPF – Interfce Template

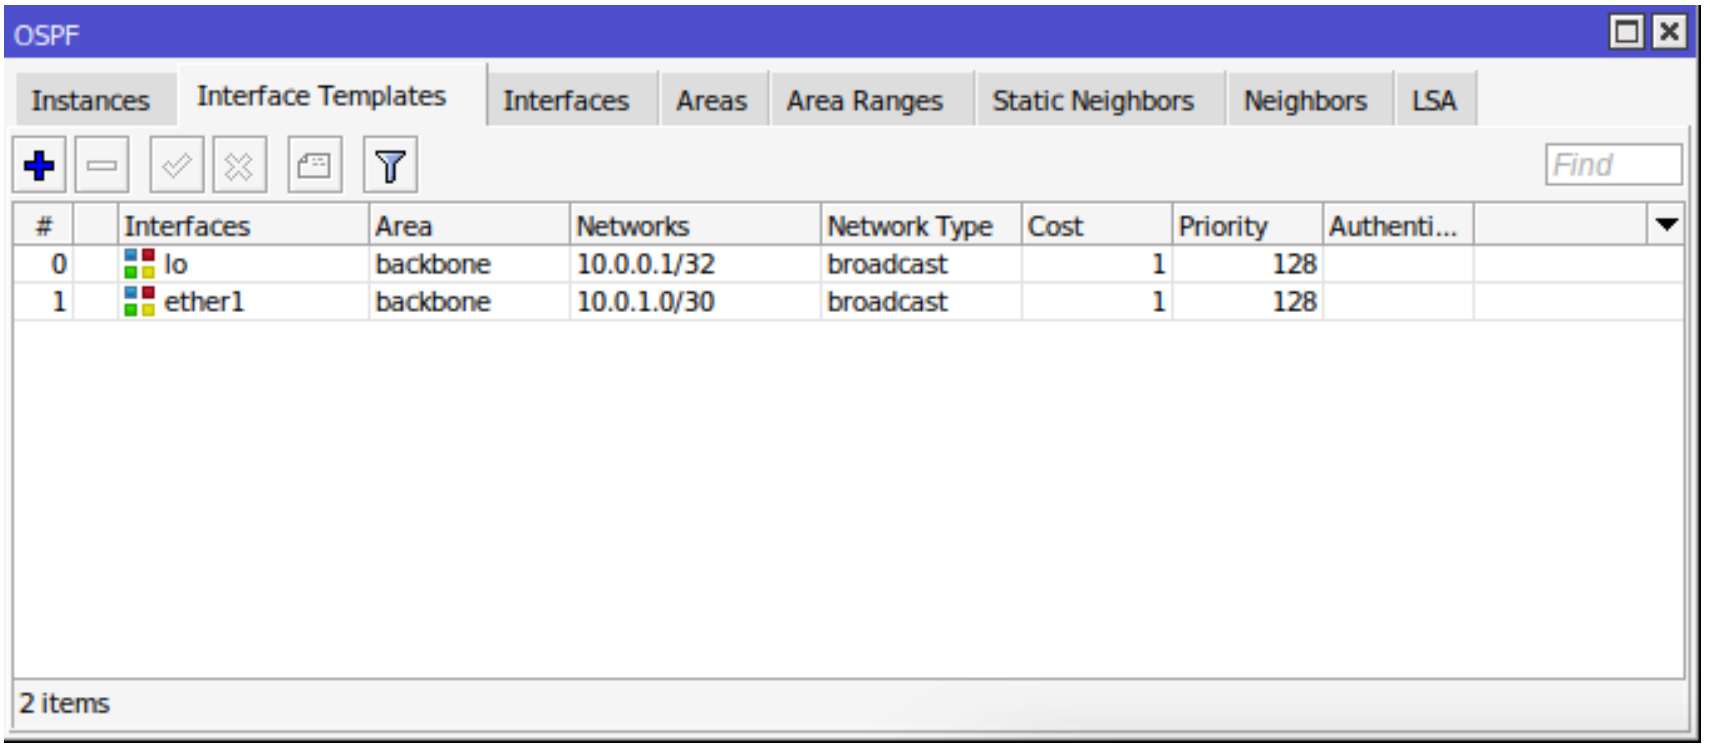

OSPF – Interface Template (top half)

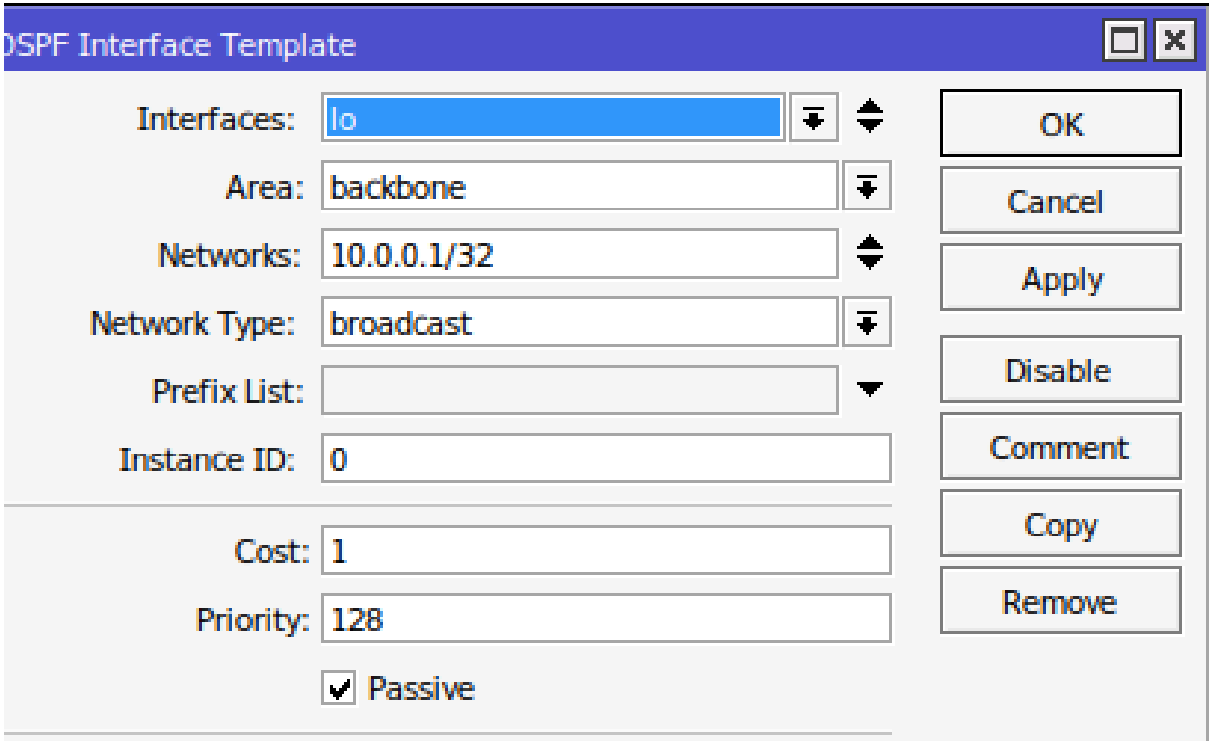

N.B – Cost and priority differ from v6 defaults -
Cost = 10
priority = 1

OSPF – Interface Template (bottom half)

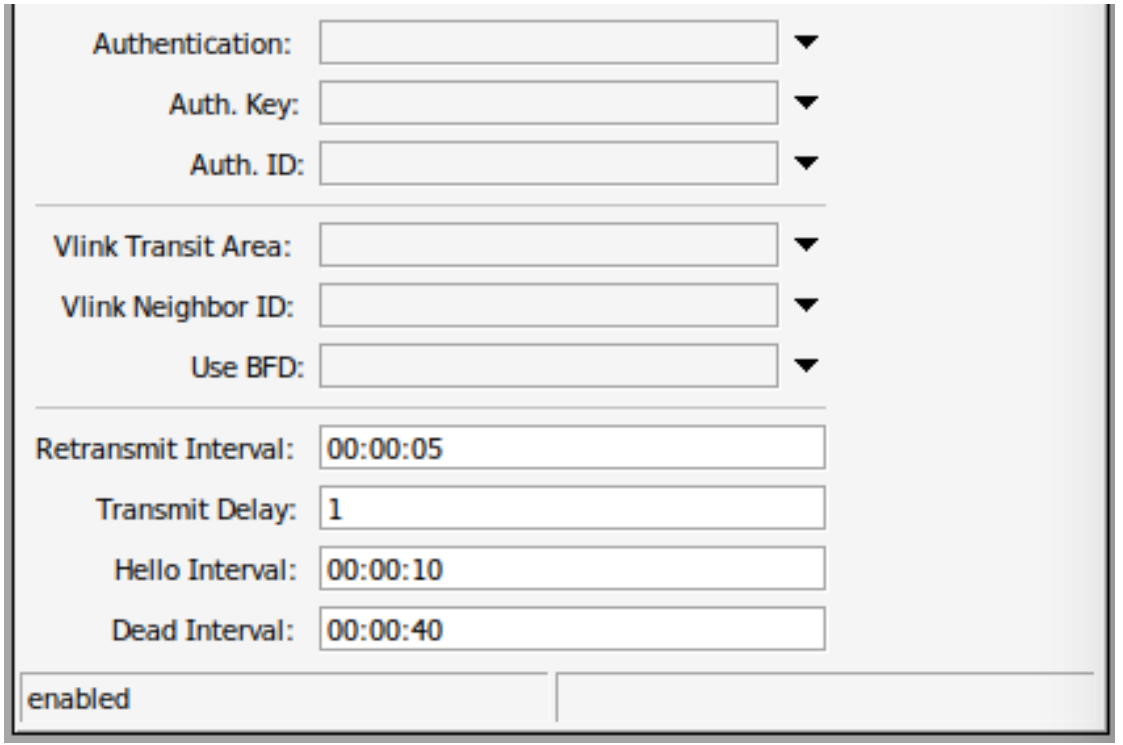

OSPF – Interface Template - Authentication options

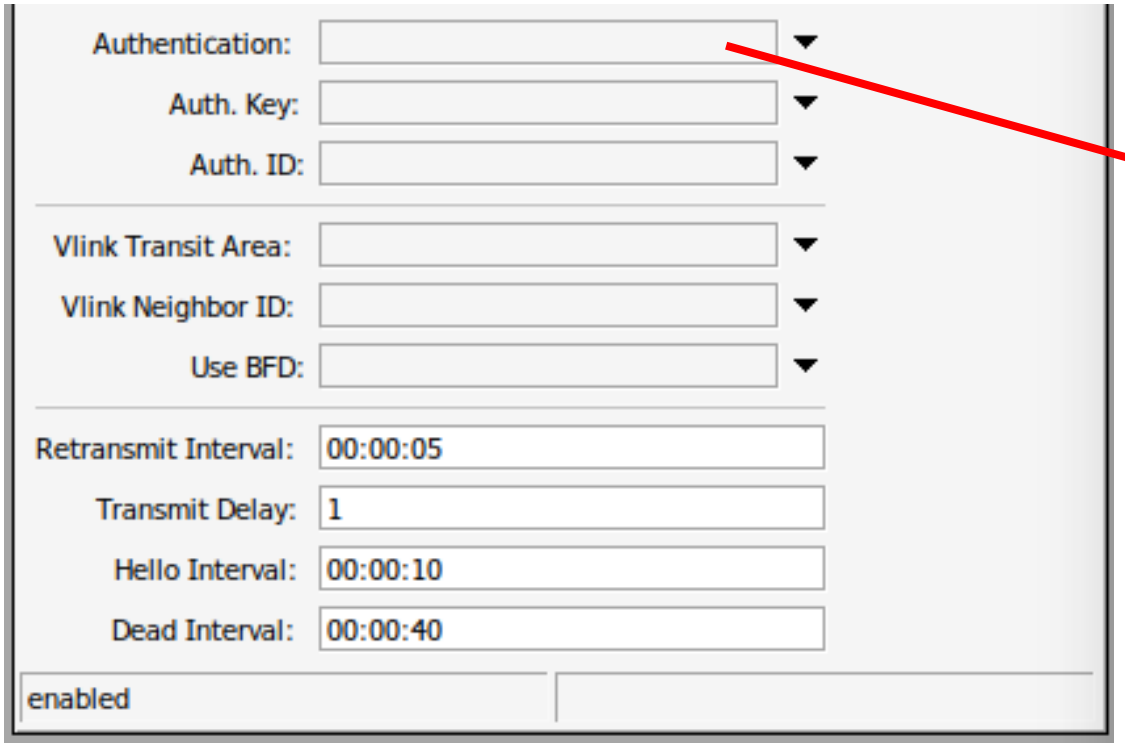

MD5
sha1
sha256
sha384
sha512
Simple

OSPF - Interfaces

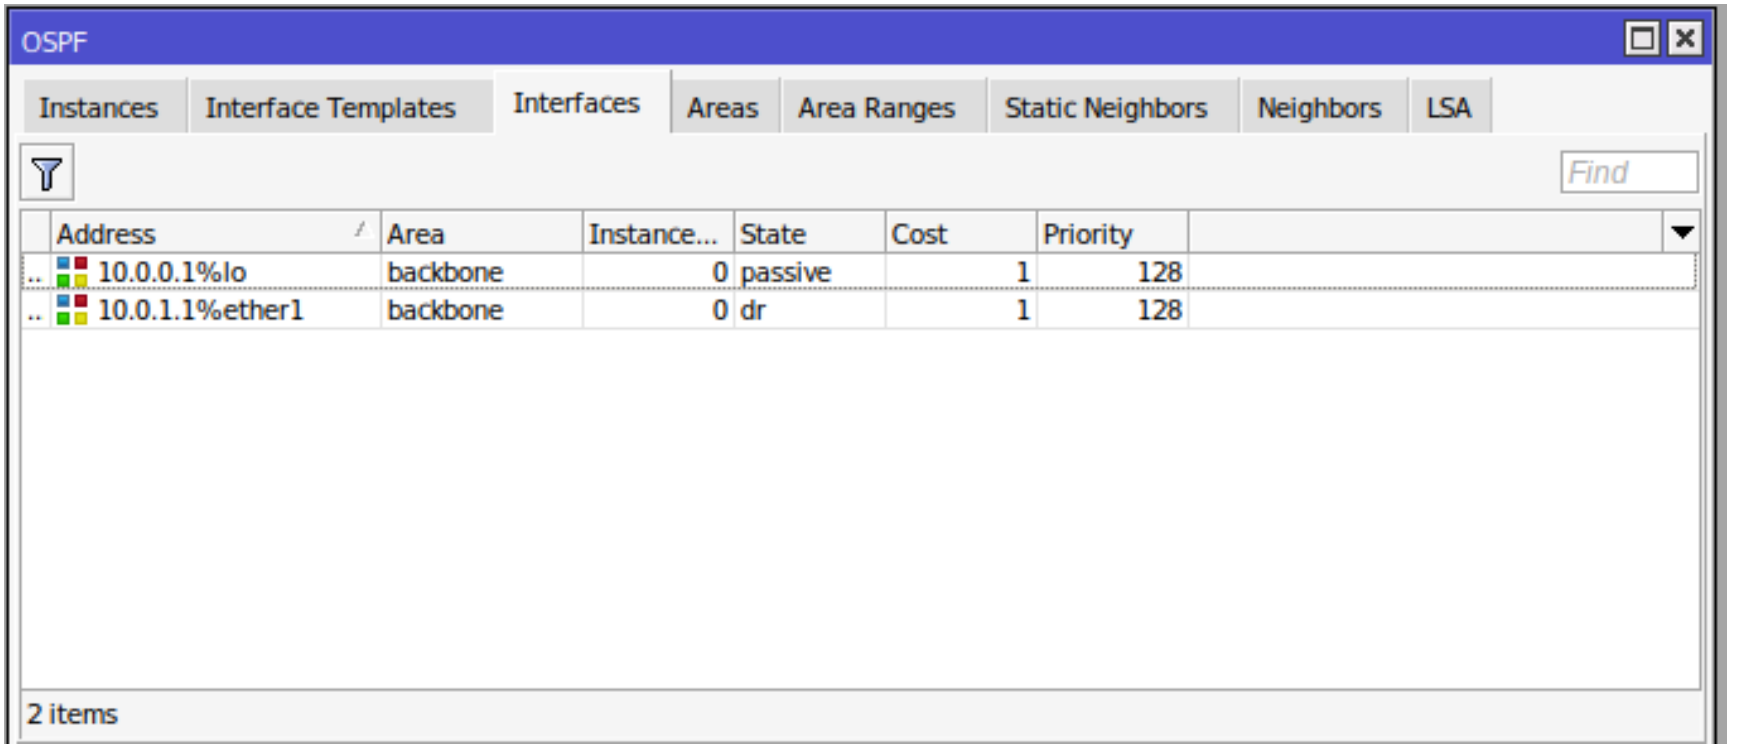

OSPF – routes / costs

The OSPF route table has now moved to IP → Routes

OSPF – routes / costs

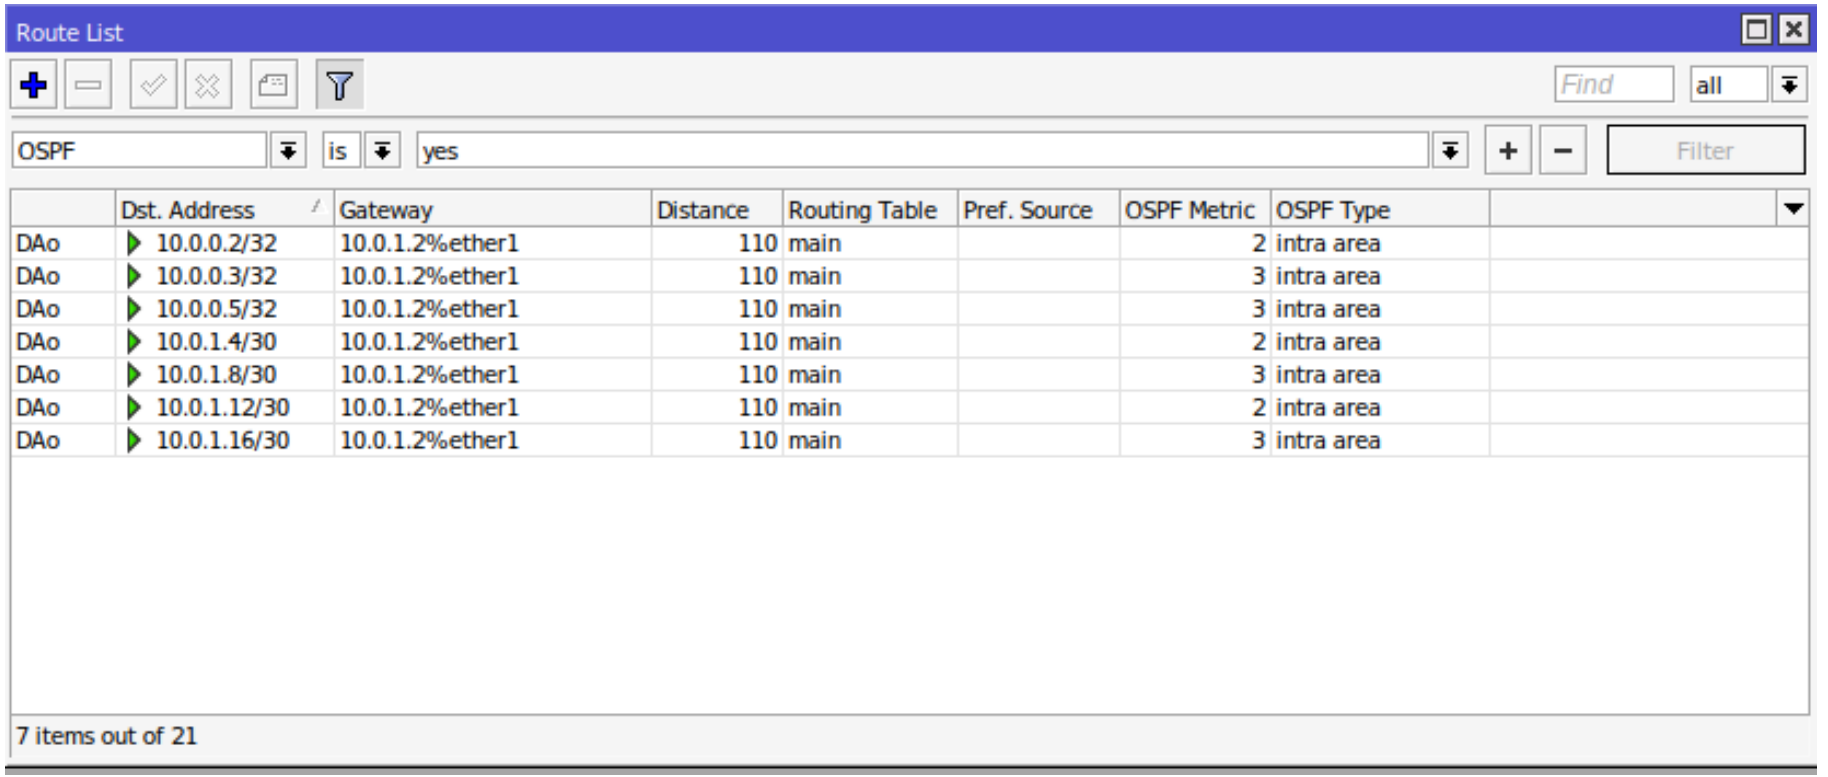

IP Route

IP Route

Before I forget this important little part of routing changes, there is the route table we all know so well.

Found in both v6 and v7 but there changes worth noting when it comes to the ‘flags’

IP Route - Flags

D – Dynamic

X – Disabled

I – Inactive

A – Active

c – Connect(ed)

S – static

r -rip

b – bgp

o – ospf

i – is-is

d - dhcp

v - vpn

m - modem

y - bgp-mpls-vpn

H - hw-offloaded

+ - ECMP

PIM SM

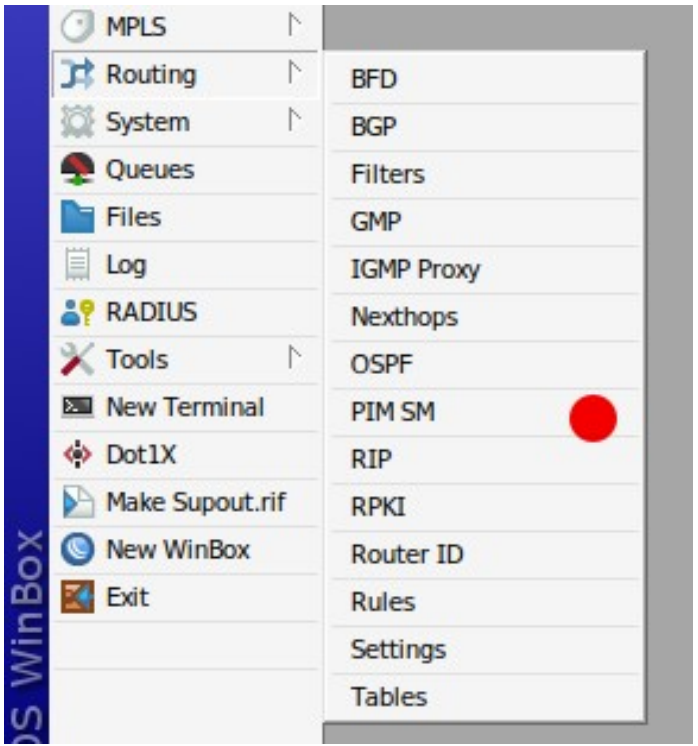

PIM SM

Protocol Independent Multicast

More multicast stuff. No time

<next>!

RIP

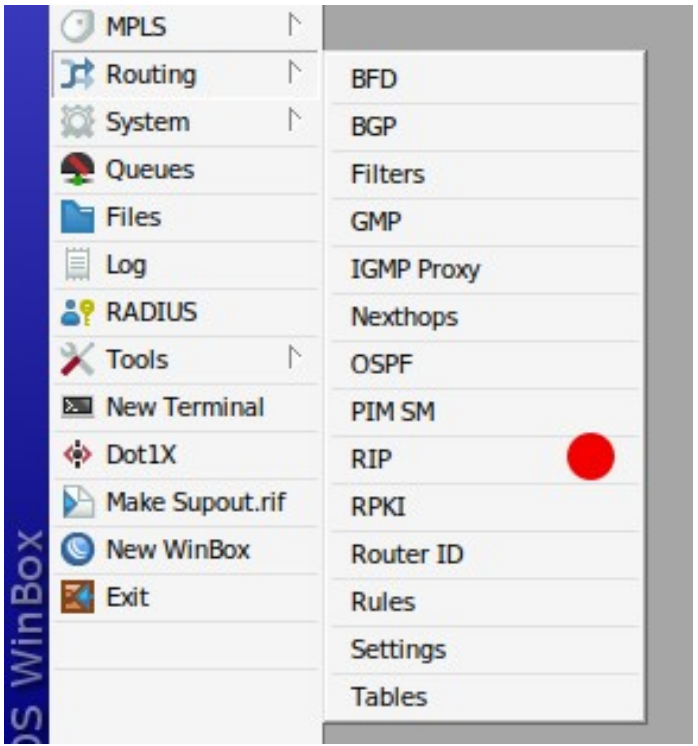

RIP

RPKI

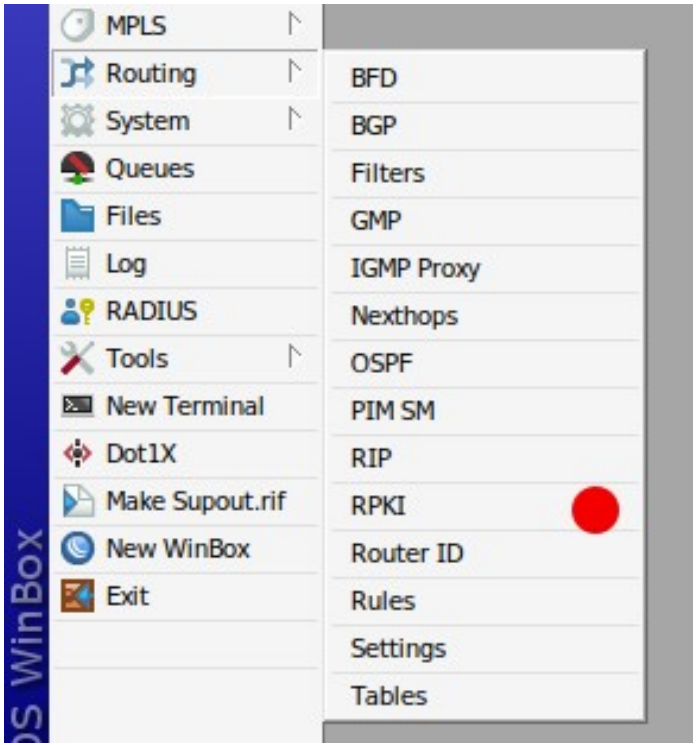

RPKI – Resource Public Key Infastructure

A simple but worthy implementation to validate the authenticity of incoming prefixes.

RPKI

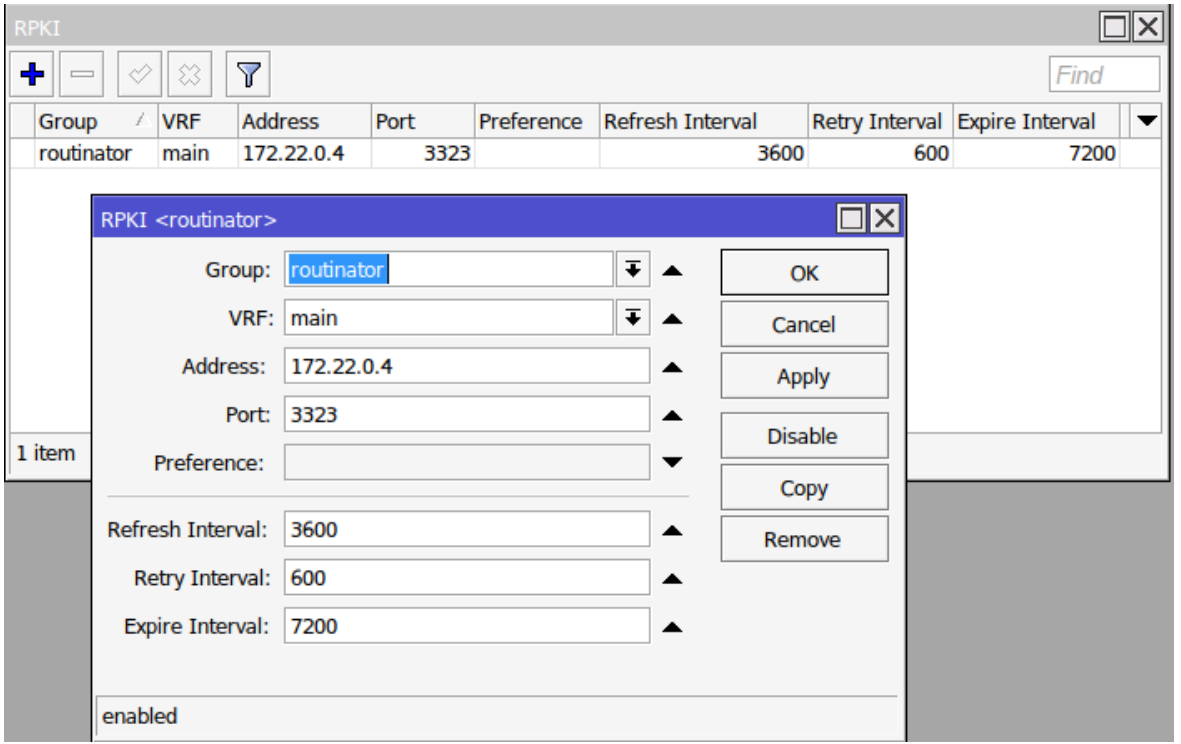

RPKI

You can run the RPKI validator ‘routinator’ within a container if you so wish.

RPKI

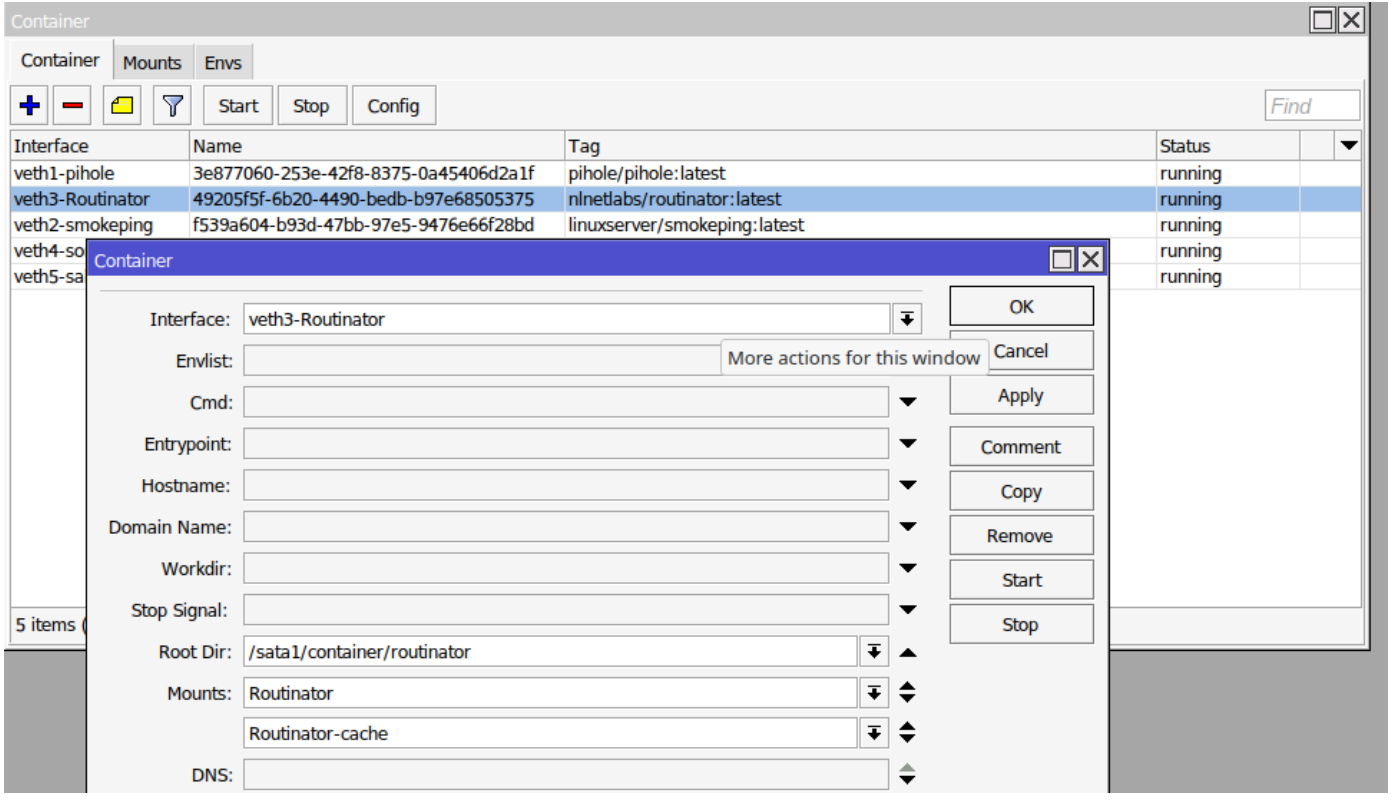

RPKI – Using it via a filter

If (rpki invalid) {reject} else {accept}

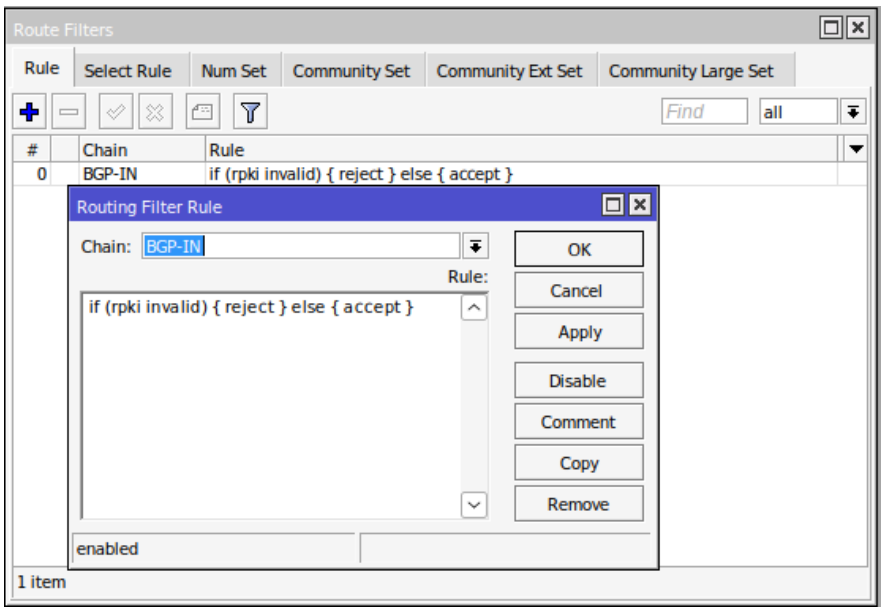

RPKI

More about RPKI later today by

14:30 - Michael Takeuchi

Router ID

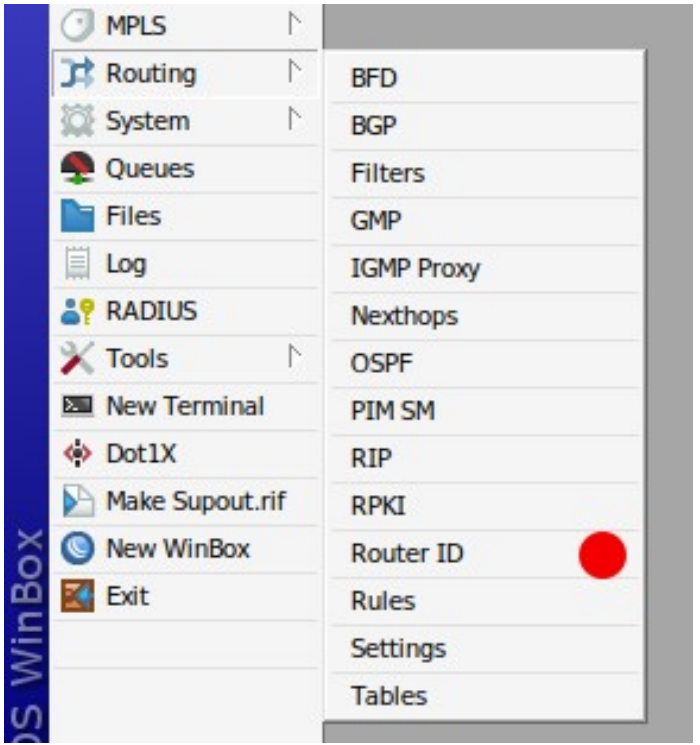

Router ID

A new addition for v7 (as in not found in v6)

Router ID

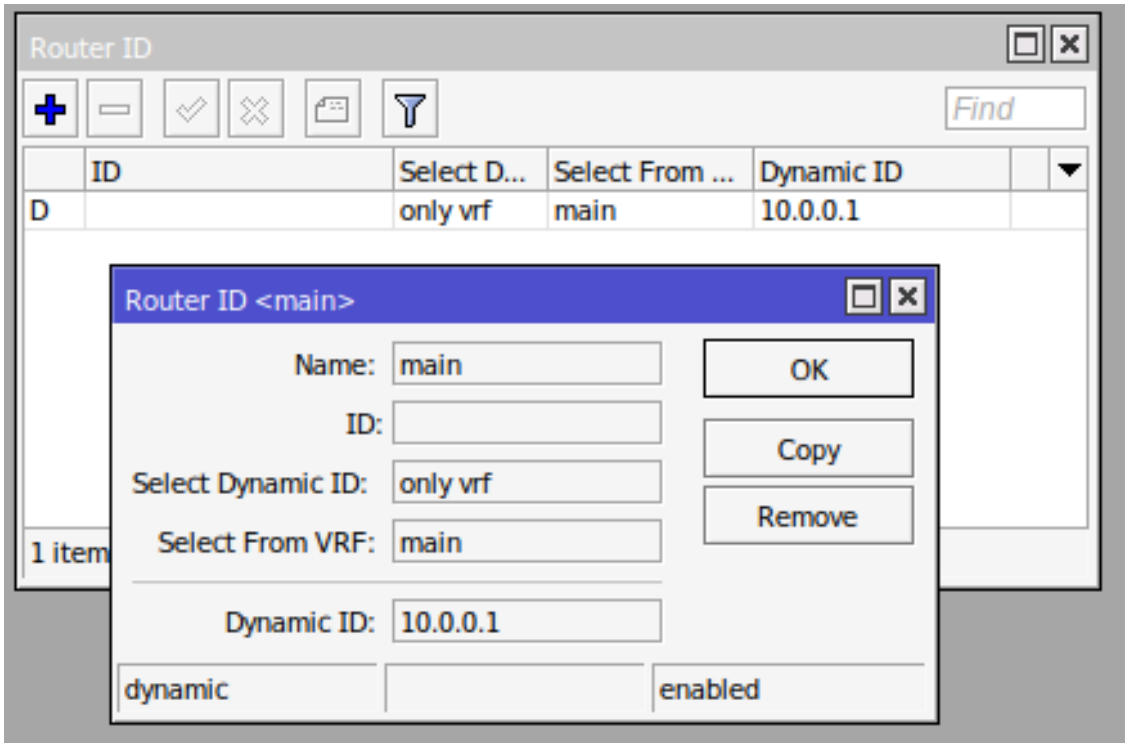

Router ID - v6

A simple config window to aid creation of a router ID typically used in e.g. OSPF.

In OSPF v6 Router ID if not set would automatically be set from the lowest active IP

Router ID - v7

However, in v7 you can either use the dynamically generated ID using the highest active IP with the smallest prefix length or of course fix it like you should!

Credit must be given to Michal, Steve, Thomas and Troels for allowing me to waste time during their MTCRE this week on discovering the specific formula that creates the dynamic router ID

A big thanks to you guys!

Rules

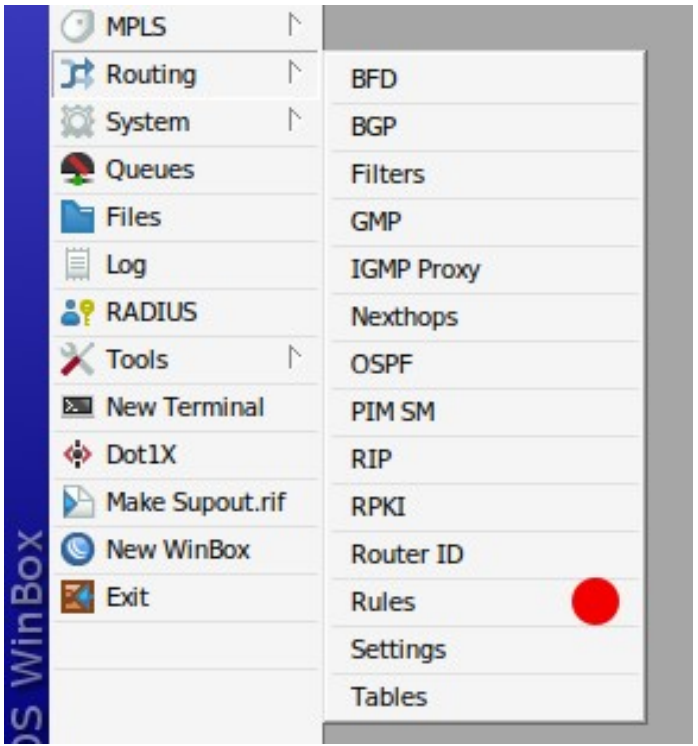

Rules - v6

IP → Route

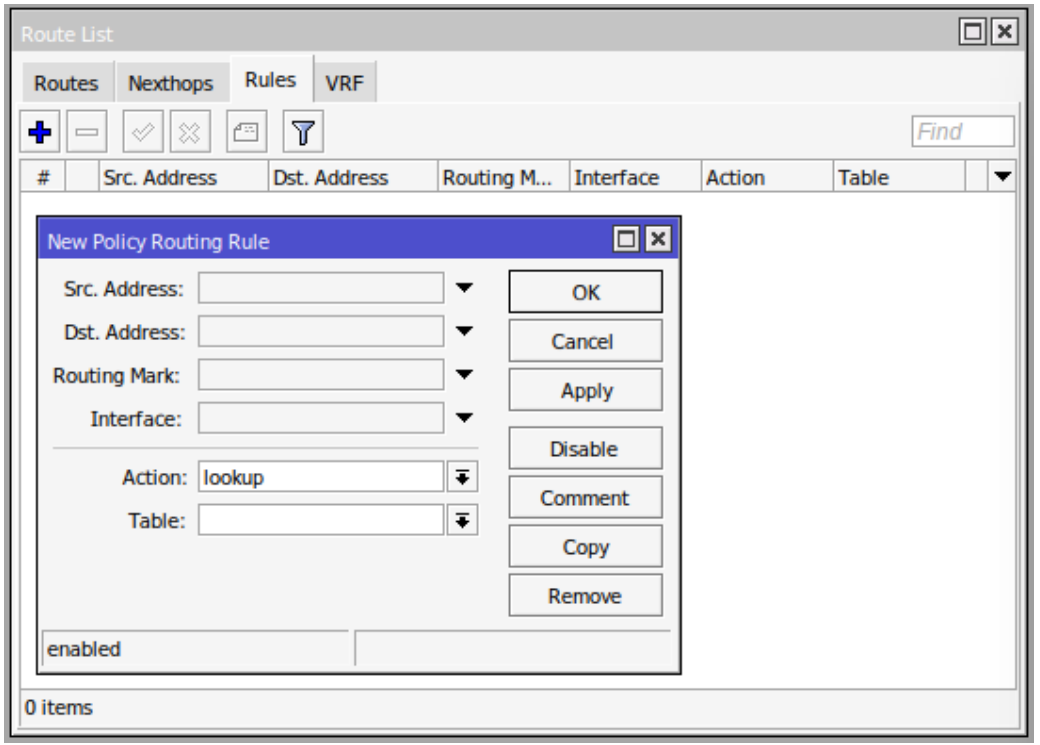

Rules - v7

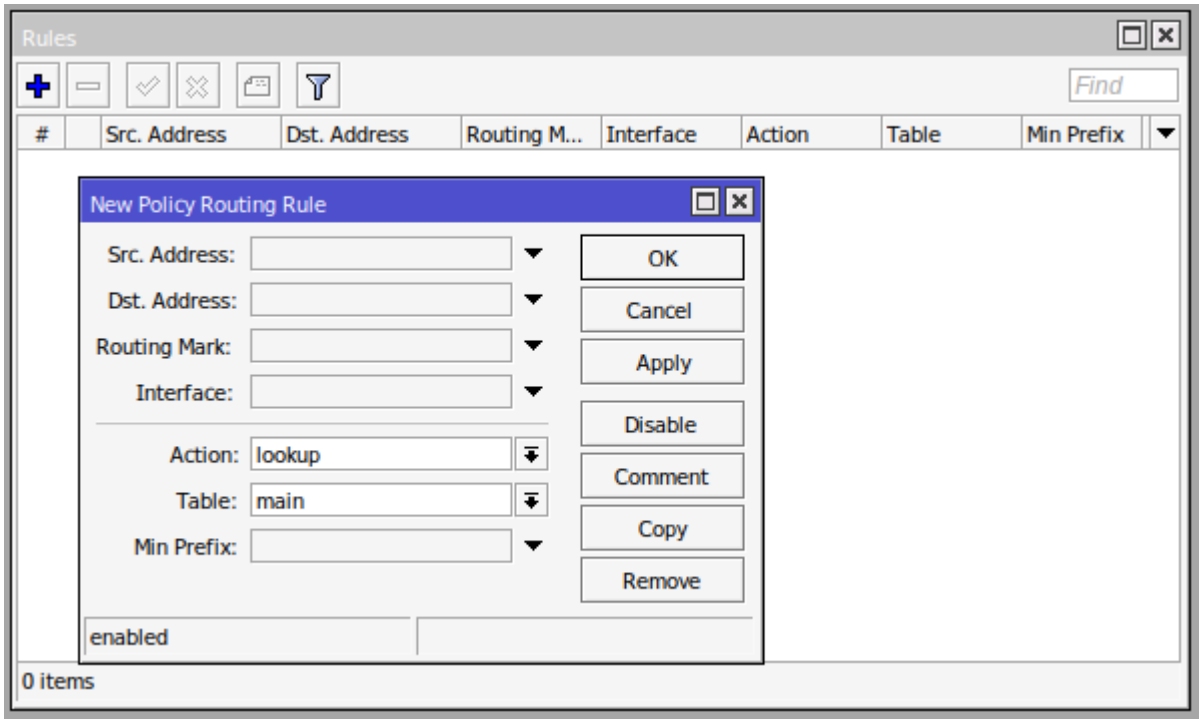

Settings

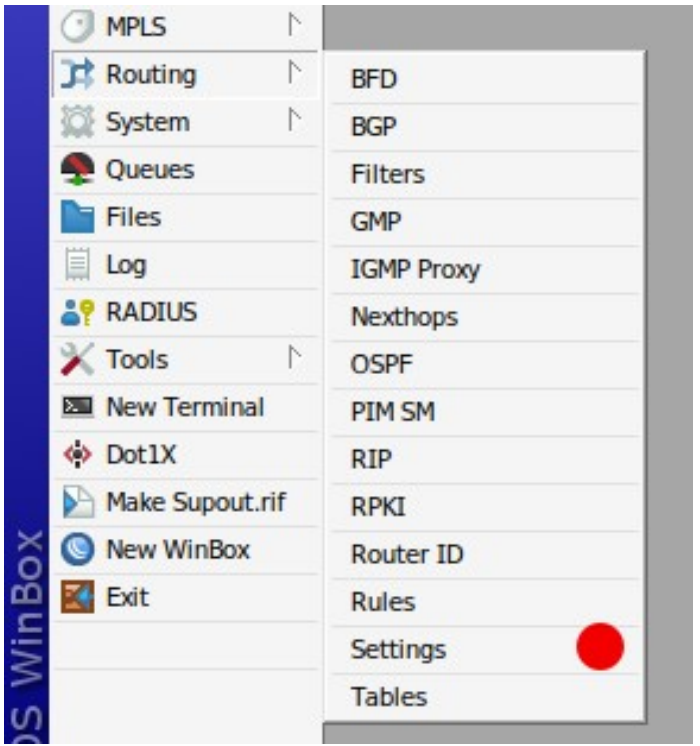

Settings – hmm

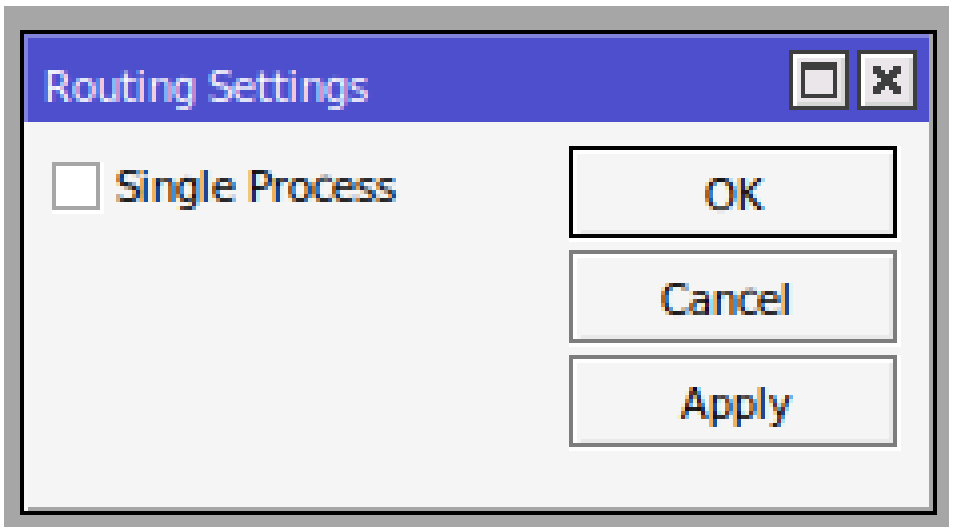

Tables

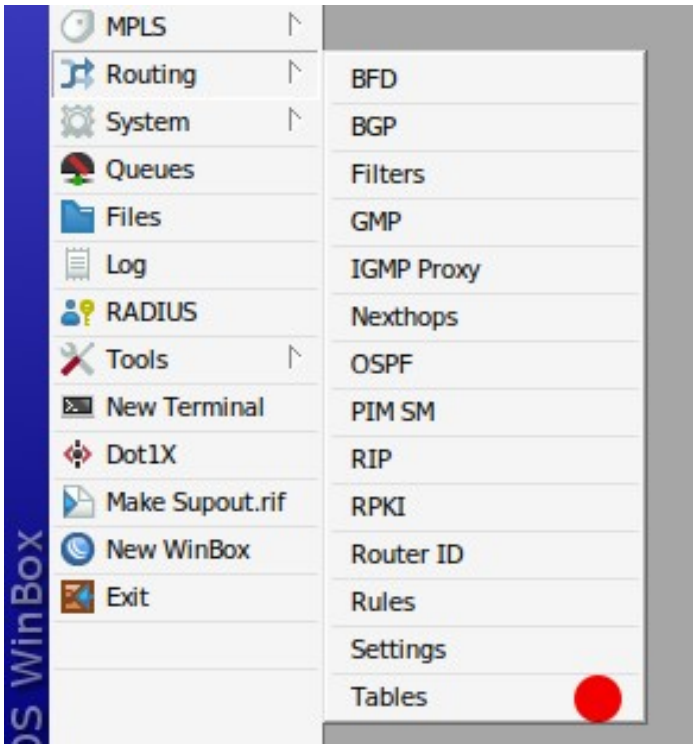

Tables

What used to be known as ‘Route Mark’.

You now have to create a ‘table’ and add to the FIB for it to be accessible/usable in mangle or the routing table

Tables

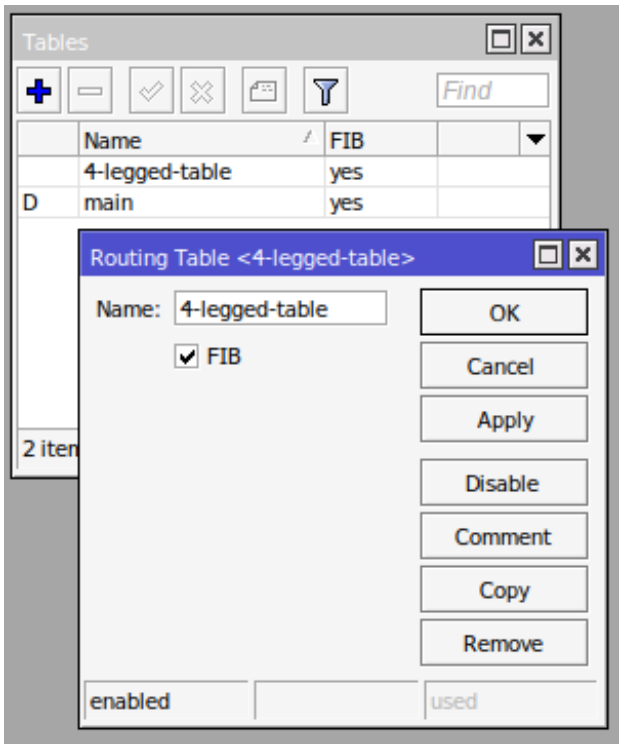

CLI →

Not everything is available via winbox at present.

The CLI (Command Line Interface) is the only option for these extra routing functions and tools

(Based on <= v7.14)

CLI →

The missing functions / tools are as follows :

- Fantasy
- IS-IS
- Route
- Stats
- Discourse
- Reinstall-fib

CLI →

```
[user@demo] /routing> <tab>
bfd             id              pimsm           settings        reinstall-fib
bgp             igmp-proxy      rip             stats
Fantasy         isis            route           table
filter          nexthop         rpki            discourse
gmp             ospf            rule            export
```

Fantasy

Create / generate prefixes to add to your route table for testing purposes

Fantasy

```
[user@demo] /routing> fantasy add count=255 \
disabled=no dst-address=10.0.0.0/8 gateway=ether1 \
name=Prague-test prefix-length=24

[user@demo] /routing> fantasy remove 0
```

IS-IS

At last a new IGP to RouterOS...! I'd love to show you more but I'll leave it to Kevin Myers (Stub area 51) later on this morning to show you this new shiney tool in the box –

11:00- 11:45

Route

An alternative to the IP → Route table

Route

```
[user@demo] /routing> route print
Flags: A - ACTIVE; c - CONNECT, o - OSPF, d - DHCP; H - HW-OFFLOADED
Columns: DST-ADDRESS, GATEWAY, AFI, DISTANCE, SCOPE, TARGET-SCOPE, IMMEDIATE-GW
    DST-ADDRESS       GATEWAY           AFI   DISTANCE  SCOPE  TARGET-SCOPE  IMMEDIATE-GW

Ad  0.0.0.0/0         185.208.89.1      ip4          1     30             10
185.208.89.1%ether2
Ac  10.0.0.1/32       lo                ip4          0     10                lo

Ao  10.0.0.2/32       10.0.1.2%ether1   ip4        110     20             10
10.0.1.2%ether1
Ao  10.0.0.3/32       10.0.1.2%ether1   ip4        110     20             10
10.0.1.2%ether1
```

Stats

A more detailed route debugging tool

Stats

```
[user@demo] /routing> stats process print
Columns: TASKS, PRIVATE-MEM-BLOCKS, SHARED-MEM-BLOCKS, PSS, RSS, VMS, ID, PID, RPID, PROCESS-TIME, KERNEL-TIME, CUR-BUSY
 # TASKS                          PRIVATE-M  SHARED-ME  PSS        RSS        VMS       ID       PID  RPID  PROCE  KERNE  CUR-
 0 routing tables                 1536.0KiB  1792.0KiB  1085.0KiB  3384.0KiB  19.2MiB  main      96     0  140ms  300ms  10ms
   rib
 1 fib                            256.0KiB   0          757.0KiB   3048.0KiB  19.2MiB  fib      116     1  50ms   250ms
 2 routing policy configuration   256.0KiB   256.0KiB   688.0KiB   2856.0KiB   19.2MiB  bgp      125     1  50ms    100ms
```

Discourse

No Idea!

Discourse

```
[user@demo] /routing>
```

Reinstall-FIB

I have assumed this cli command refreshes the FIB (Forwarding information base)

Reinstall-FIB

```
[user@demo] /routing> reinstall-fib <enter> [y/N]:
Refreshing FIB

...and crash!
```

Wait.. what is that I can hear? It’s Lorenzo!

<next>

The End / Fini / Kaput / No more / TTFN

Thank you for listening!