

BALANCEO DE CARGA

Escenarios: tradicional y empleando BGP

Alejandro Teixeira G.

ateixeira@mkx.cl

MikroTik Certified Trainer

MikroTik Trainer ID #TR0163

Presentación personal

2

- Alejandro Teixeira
 - Ingeniero de Telecomunicaciones
 - Magister en Ingeniería de Negocios
 - Magister en Finanzas
- Gerente Comercial
 - AUSTRO INTERNET S.A.
- Gerente General
 - TF CONSULTING LTDA



Presentación personal

3

- Consultor y Entrenador MikroTik RouterOS
- Experiencia con MikroTik desde 2007



at@mkx.cl



AcademyXpertsChile



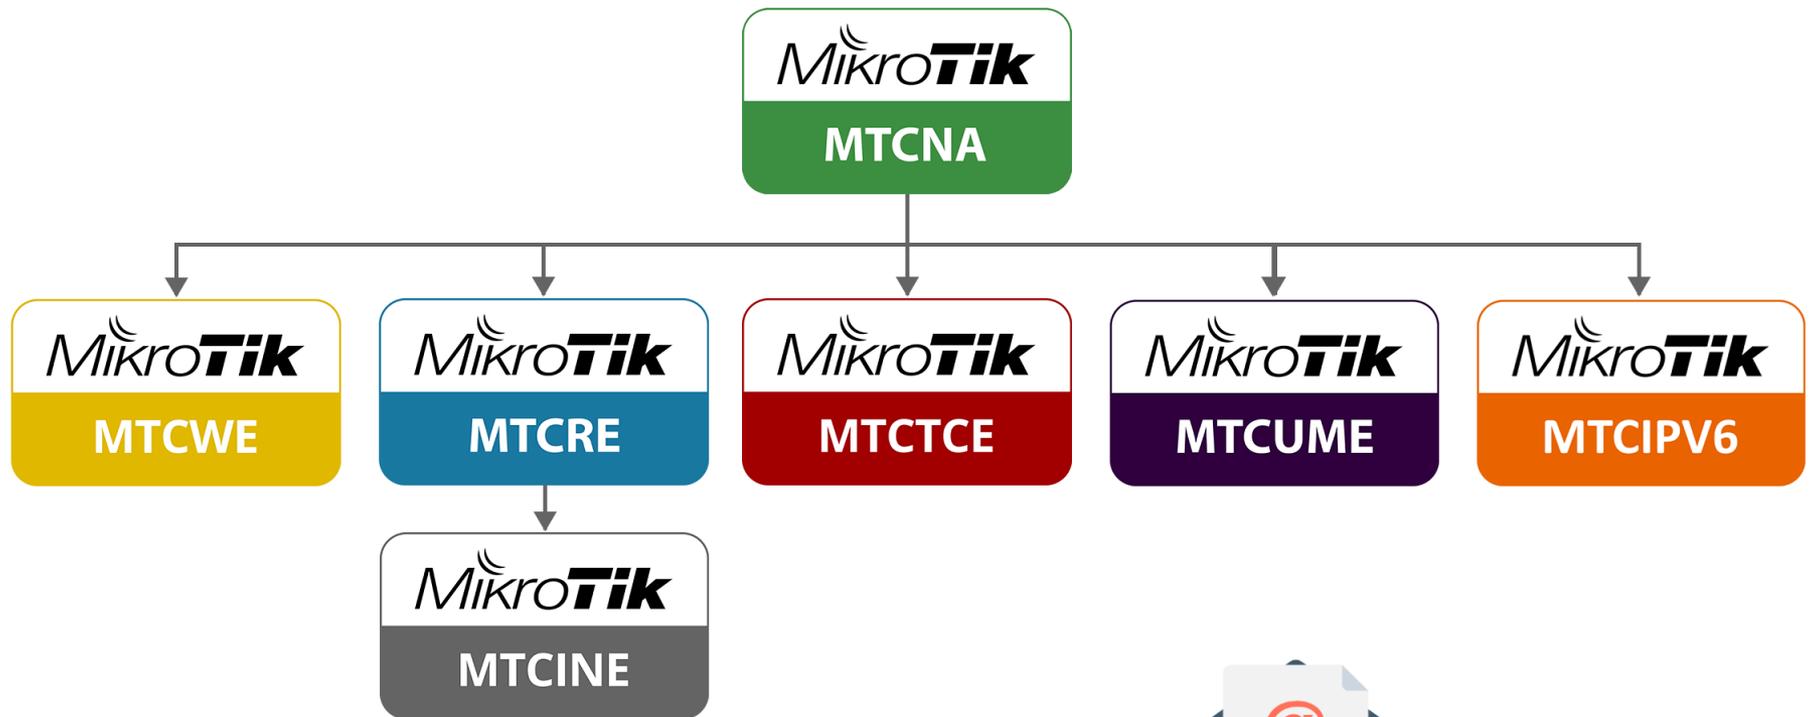
+569-7858-3064



Cursos oficiales MikroTik RouterOS

4

□ Cursos de certificación oficial MikroTik

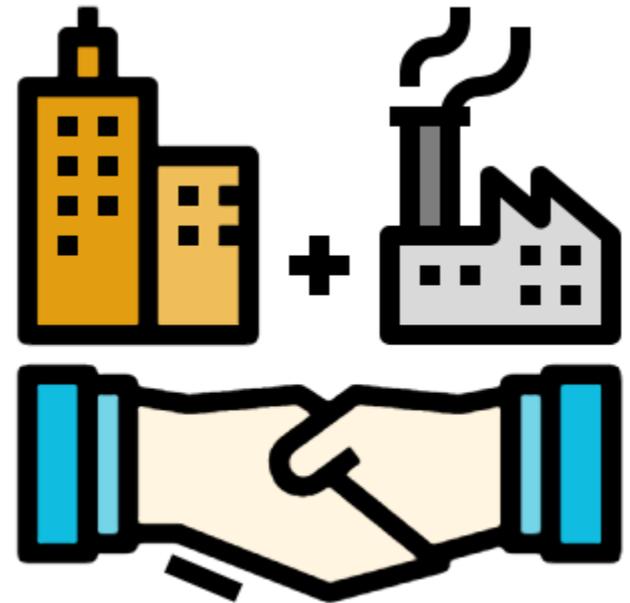


at@mkx.cl

Consultoría y servicios de TI

5

- Asesoría en redes y telecomunicaciones
- Diseño y desarrollo de soluciones infomáticas
- Soporte mensual
 - ▣ Asesoría especializada
 - ▣ Mantenimiento preventivo
 - ▣ Monitoreo
 - ▣ Soporte prioritario
 - ▣ NOC 24/7

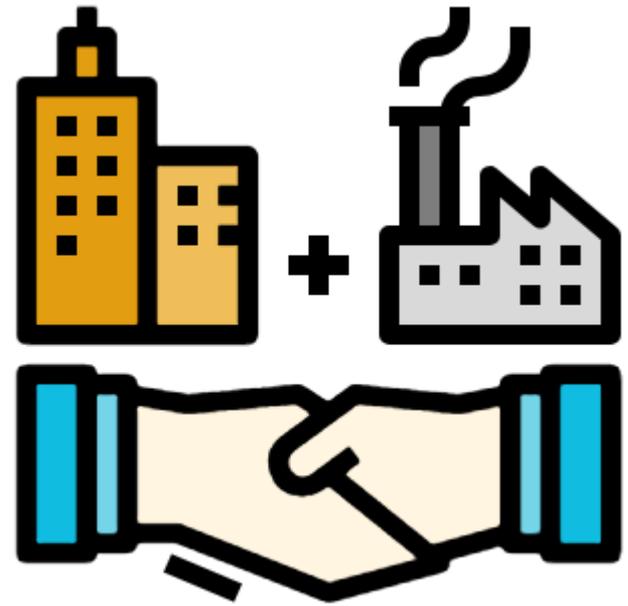


at@eruditum.cl

Consultoría y servicios de TI

6

- Soluciones BGP
- Firewall Avanzado
 - ▣ Anti-DDoS
- Balanceo de Carga
- Calidad de Servicio (QoS)
- VPN
- Soluciones inalámbricas
- Hotspot
- Trámite ante LACNIC
 - ▣ Bloque IP + ASN



at@eruditum.cl

Enlaces de internet y datos

7

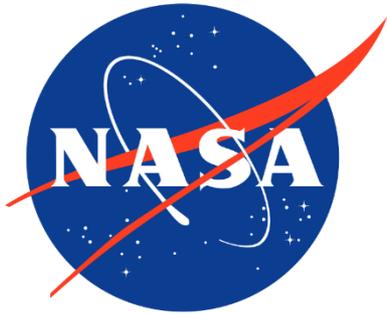
- Enlaces de acceso a internet
- Enlaces de datos
 - ▣ MPLS / VPLS
- Tránsito internacional
- Conexión con CDN de mayor consumo
 - ▣ Google, YouTube, Drive, Gmail
 - ▣ Facebook, Instagram
 - ▣ Microsoft
 - ▣ Netflix



at@austrointernet.cl

Algunos de nuestros clientes

8



Agenda

9

- ¿Qué es balanceo de carga?
 - ▣ Escenario tradicional
 - ▣ En escenarios con BGP
 - ▣ Consideraciones
- Mecanismos de balanceo
 - ▣ ECMP / PCC / NTH
 - ▣ Escenarios con BGP
- Ponderación de enlaces
- Posibles problemas y soluciones
- Failover



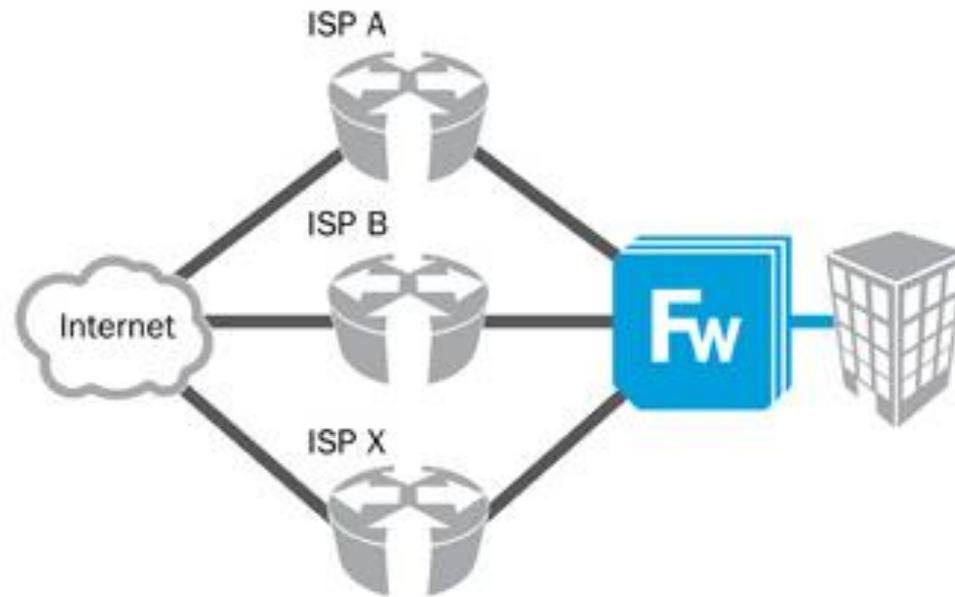
10

¿Qué es balanceo de carga?

¿Qué es balanceo de carga?

11

- Proceso a través del cual **el tráfico saliente es distribuido por múltiples enlaces**



Ejemplo

12

Escenario tradicional					
Canal	VLAN	VRRP	Bonding	LTE	
TU	Tx		Rx	Tx Packet (p/s)	R
	455.4 kbps		11.8 kbps	41	
	41.0 Mbps		761.2 kbps	3 390	
	172.5 kbps		10.2 Mbps	387	
	125.2 kbps		10.3 Mbps	281	
	199.5 kbps		10.2 Mbps	439	
	204.7 kbps		10.2 Mbps	444	

Ejemplo

13

Interface List												
Interface	Interface List	Ethernet	EoIP Tunnel	IP Tunnel	GRE Tunnel	VLAN	VRRP	Bonding	LTE			
✓	✗	🗨️	🔍	Power Cycle								
	Name	Type	MTU	Actual MTU	L2 MTU	Tx	Rx	Tx Pa				
RS	ether1_multicast	Ethernet	1500	1500	1600	5.2 kbps	54.6 Mbps					
R	sfp1 - IMTEL	Ethernet	1500	1500	1580	753.5 Mbps	33.1 Mbps					
R	sfp2 - PIT	Ethernet	1500	1500	1580	466.7 Mbps	3.2 Gbps					
R	sfp3 - Principal	Ethernet	1500	1500	1580	342.1 Mbps	586.0 Mbps					
R	sfp4 - Respaldo	Ethernet	1500	1500	1580	4.1 kbps	873.5 Mbps					
R	sfp5 -Extension-1072	Ethernet	1500	1500	1580	2.4 Gbps	1757.6 Mbps					
R	sfp6 - Clientes	Ethernet	1500	1500	1580	2.5 Gbps	123.3 Mbps					
R	sfp7 - NUCLEO	Ethernet	1500	1500	1580	20.2 Mbps	1158.8 kbps					
R	sfp8 - Peñalolen	Ethernet	1500	1500	1580	200.3 Mbps	7.5 Mbps					

Tráfico total: 4,85 Gbps
3 enlaces de internet
Escenario BGP

Escenarios

14

□ Tradicional

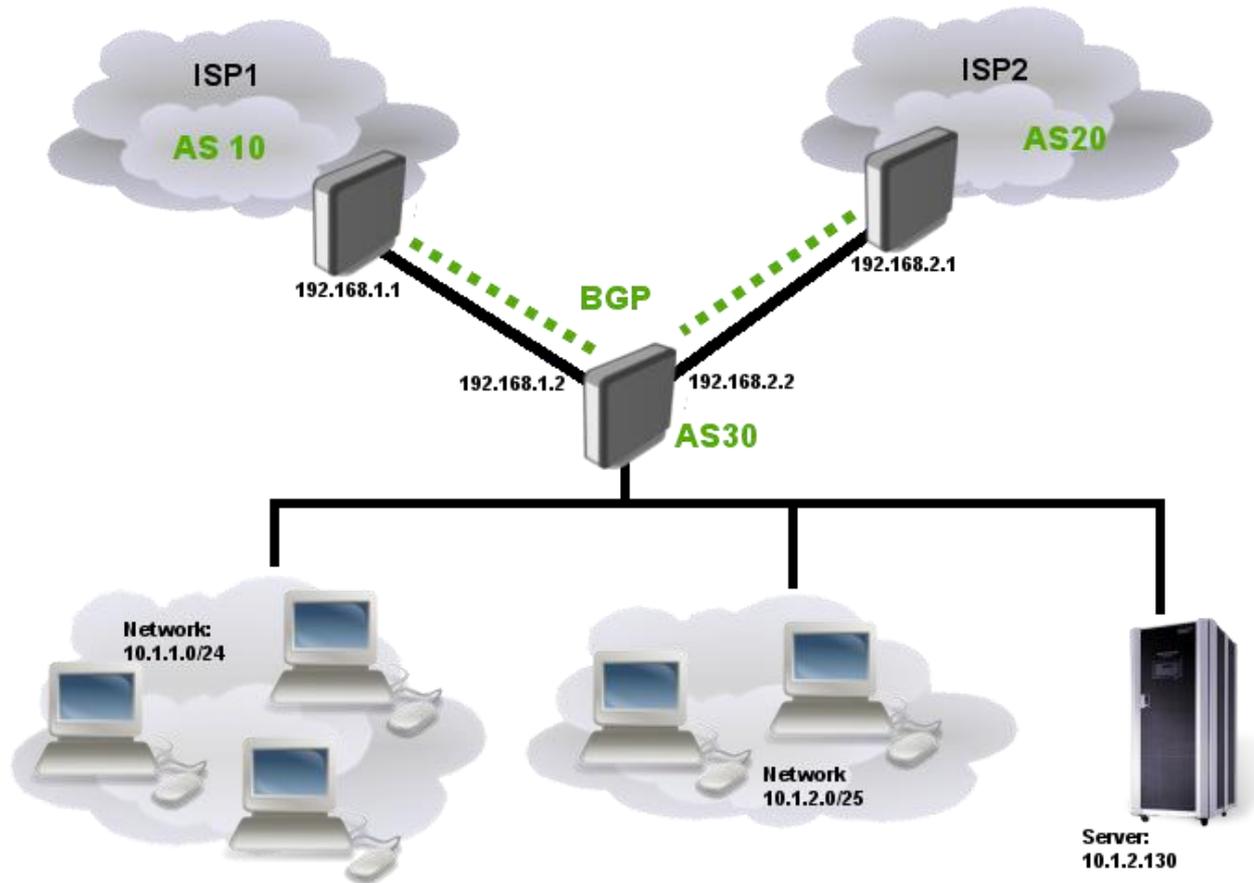
- Cada WAN tiene un pool de IP asignado
- Los pool de IP pertenecen a cada ISP
 - ENTEL
 - MOVISTAR
 - CLARO

□ BGP (multihomed)

- Requiere recursos propios (otorgados por LACNIC)
 - IP
 - ASN
- Pool de IP compartido en todas las WAN

BGP: multihomed

15



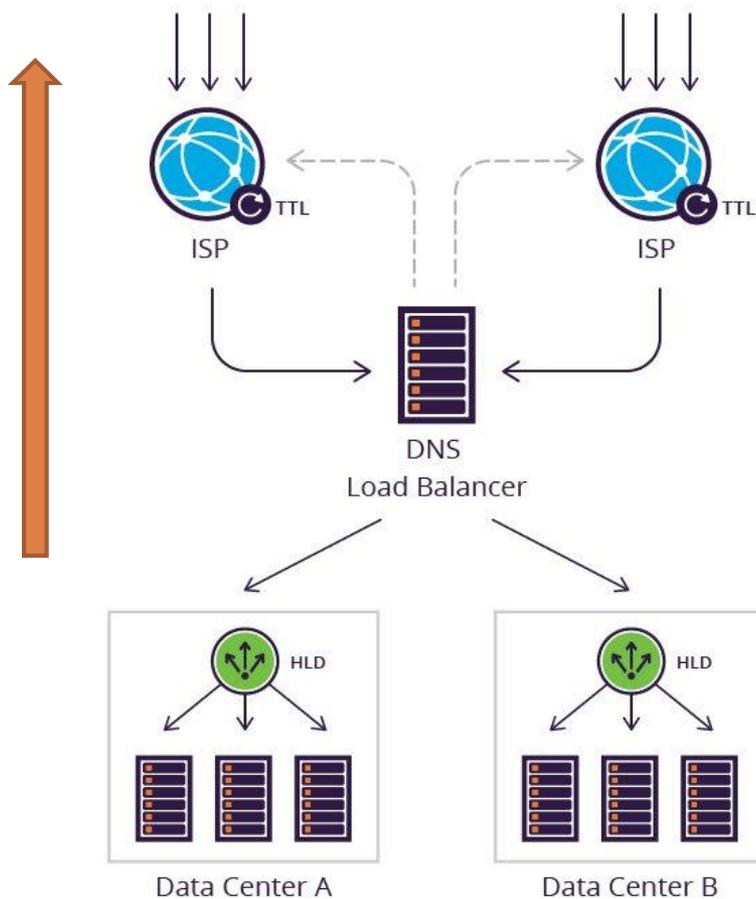
Consideraciones

16

- Cuando se emplea algún mecanismo de balanceo, **lo que realmente se balancea son el número total de conexiones** entre la cantidad de enlaces disponibles.
- **No se balancea el consumo de tráfico**
- **En BGP se anuncian prefijos IP** (segmentos de red) a cada ISP y se aplica prepend para definir las prioridades de subida y bajada, entre otras condiciones
 - ▣ Capacidad para recibir por cualquier WAN el mismo segmento de red

Consideraciones

17



- Balanceo correcto
 - 2 conexiones
 - No es equitativo

5 Mbps



1 Mbps



18

Mecanismos de balanceo

Herramientas involucradas

19

- Escenario tradicional
 - `/ip firewall mangle`
 - `/ip route`
 - `/system scripts`
 - `/system scheduler`
 - `/ip firewall filter`
- Escenario BGP
 - `Filtros BGP`
 - `Tabla de rutas`
 - `/ip firewall mangle (opcional)`

ECMP: Equal Cost Multi-Path

20

- Simple de implementar
- Balanceo persistente por conexión (IP SRC/DST)
- Se agregan todas las puertas de enlace
 - ▣ En la misma ruta
- No se puede controlar la forma de balanceo
- No funciona con puertas de enlace iguales
 - ▣ Mas de una por proveedor
- Failover automatico
 - ▣ No es práctico

ECMP: Implementación

21

```
/ip route add  
gateway=1.1.1.1,2.2.2.1  
check-gateway=ping
```

New Route

General Attributes

Dst. Address: 0.0.0.0/0

Gateway: 200.200.200.1

150.150.150.1

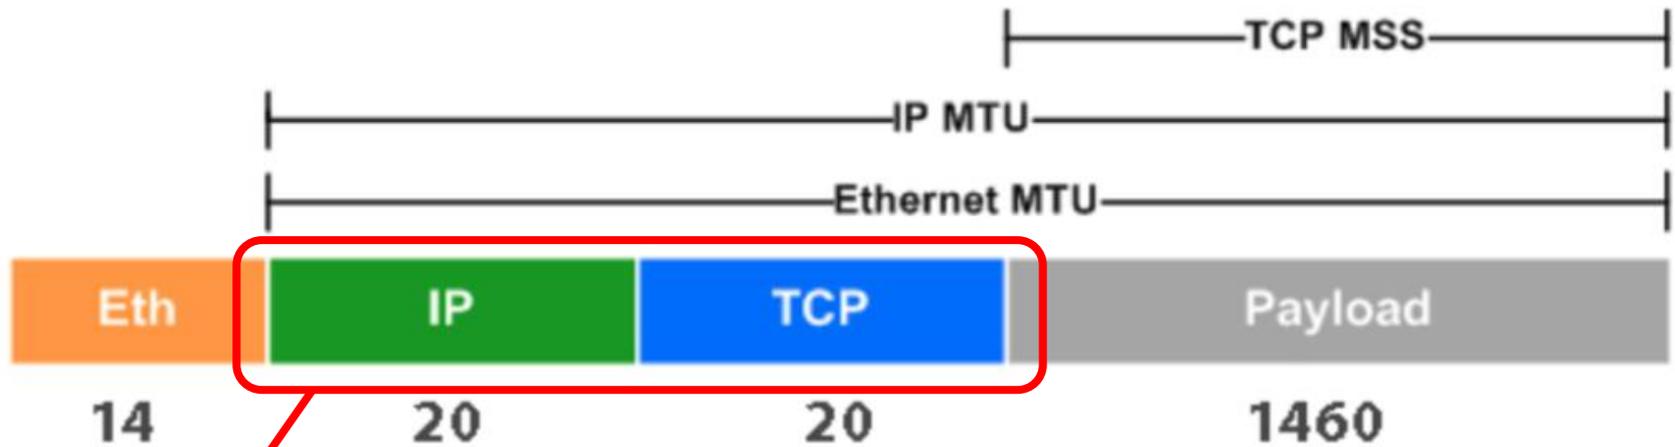
Check Gateway:

Type: unicast

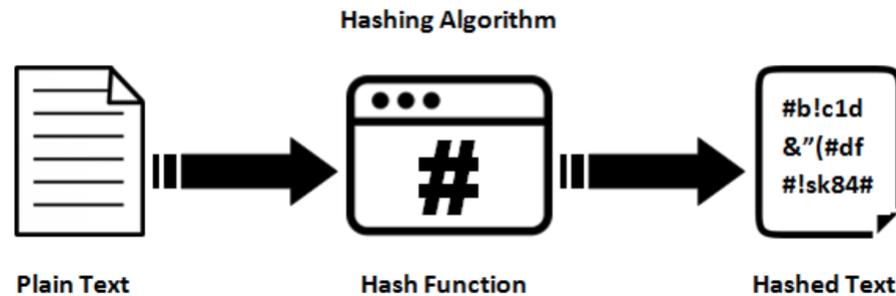
PCC: Peer Connection Classifier

22

- Sofisticado
- No distribuye de forma equitativa la cantidad de conexiones (aleatorio)
- PCC toma algunos campos de la cabecera IP
 - ▣ Alguna combinación
 - IP (origen/destino) y puerto (origen/destino)
- Clasificadores
 - ▣ src-address
 - ▣ dst-address
 - ▣ src-port
 - ▣ dst-port



- **IP SRC**
- **IP DST**
- **PORT DST**
- **PORT SRC**



Escenario: both addrees and ports

24

- **IP SRC: 10.25.78.54**
 - ▣ 00001010.00011001.01001110.00110110
 - ▣ 169.430.582
- **IP DST: 8.8.8.8**
 - ▣ 00001000.00001000.00001000.00001000
 - ▣ 134.744.072
- **PORT SRC: 58952**
- **PORT DST: 53**

IP SRC + IP DST + PORT SCR + PORT DST

304.233.659 MOD COUNT(WAN)

MOD = RESIDUO DE LA DIVISIÓN

25

- $10 \text{ MOD } 2 = 0$
- $10 \text{ MOD } 3 = 1$
- $10 \text{ MOD } 4 = 2$
- $10 \text{ MOD } 5 = 0$
- $10 \text{ MOD } 6 = 4$
- $10 \text{ MOD } 7 = 3$
- $10 \text{ MOD } 8 = 2$
- $10 \text{ MOD } 9 = 1$
- $10 \text{ MOD } 10 = 0$

SQL MOD() function

Syntax : $\text{MOD}(N, M)$

Example : $\text{MOD}(7, 2)$

Output : 1

© w3resource.com

PCC: MANGLE

26

Mangle Rule <>

General Advanced Extra Action Statistics

Src. Address List: private

Dst. Address List: private

Layer7 Protocol:

Content:

Connection Bytes:

Connection Rate:

Per Connection Classifier: both addresses : 3 / 0

Src. MAC Address:

Out. Bridge Port:

In. Bridge Port:

Ingress Priority:

CLASIFICADOR

RESIDUO

COUNT(WAN)

PCC: ¿Cómo funciona?

27

- Número de enlaces wan: 3
- Posibles resultados: 0, 1 o 2

- Hash 1: $14.350 / 3 = 1$
- Hash 2: $3480 / 3 = 0$
- Hash 3: $13.468 / 3 = 1$
- Hash 4: $5.390 / 3 = 2$
- Hash 5: $7.894 / 3 = 1$

- Link 0: 1
- Link 1: 3
- Link 2: 1



to-scl-2	642.7 GiB	2895 992 045	BALANCER-SCL
to-scl-2	660.6 GiB	2930 834 269	BALANCER-SCL
to-scl-4	644.1 GiB	2895 769 277	BALANCER-SCL
to-scl-4	647.6 GiB	2886 430 152	BALANCER-SCL

TABLA DE RUTAS

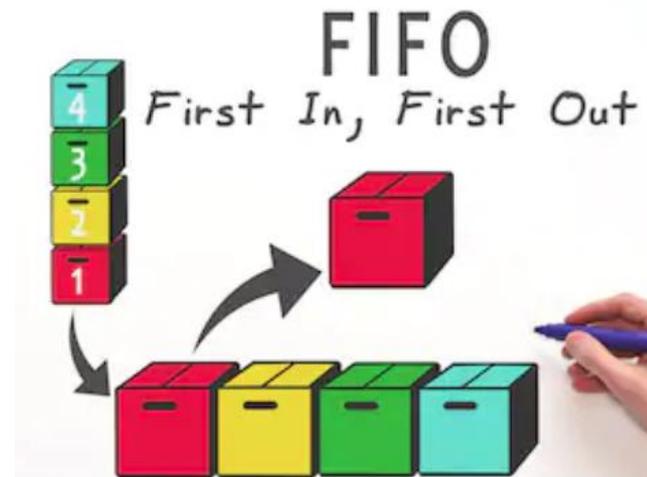
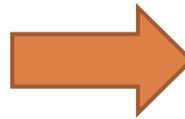
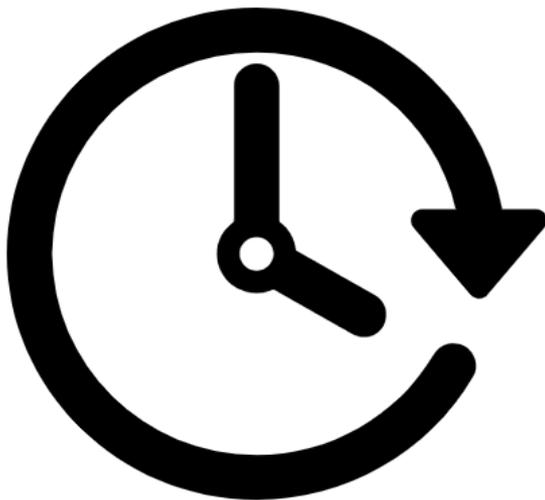
28

	Dst. Address	Gateway	Di...	Routing Mark
AS	▶ 0.0.0.0/0	209.244.0.3 recursive via 190.82.8...	1	vlan-501
AS	▶ 0.0.0.0/0	208.67.220.222 recursive via 186....	1	vlan-401
AS	▶ 0.0.0.0/0	208.67.222.220 recursive via 186....	1	vlan-301
AS	▶ 0.0.0.0/0	208.67.222.123 recursive via 186....	1	vlan-201
AS	▶ 0.0.0.0/0	208.67.220.123 recursive via 186....	1	vlan-101
AS	▶ 0.0.0.0/0	172.23.105.1 reachable ipip_vc5	1	to-scl-5
AS	▶ 0.0.0.0/0	172.23.104.1 reachable ipip_vc4	1	to-scl-4
AS	▶ 0.0.0.0/0	172.23.103.1 reachable ipip_vc3	1	to-scl-3
AS	▶ 0.0.0.0/0	172.23.102.1 reachable ipip_vc2	1	to-scl-2
AS	▶ 0.0.0.0/0	172.23.101.1 reachable ipip_vc1	1	to-scl-1
AS	▶ 0.0.0.0/0	208.67.220.123 recursive via 186....	1	full-balancer

NTH

30

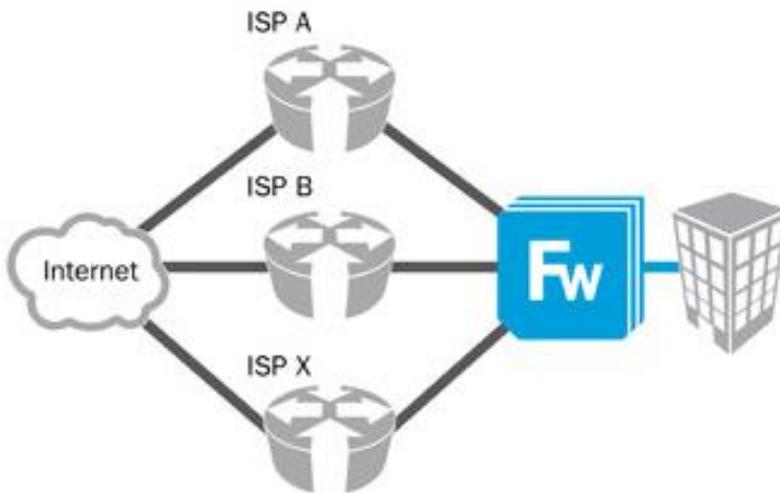
- Permite distribuir de forma **equitativa**
- Cada regla del NTH tiene su propio **contador**
- Cuando una regla **recibe un paquete** se **incrementa el contador**
- Cuando el contador llega al máximo se reinicia



Resumen: balanceo tradicional

31

- ECMP / PCC / NTH
 - ▣ Conexiones persistentes
 - ▣ Subida y bajada por la misma WAN
 - ▣ Conexión depende de la IP de la interfaz WAN



	Dst. Address	Gateway	Di...	Routing Mark
AS	▶ 0.0.0.0/0	209.244.0.3 recursive via 190.82.8...	1	vlan-501
AS	▶ 0.0.0.0/0	208.67.220.222 recursive via 186....	1	vlan-401
AS	▶ 0.0.0.0/0	208.67.222.220 recursive via 186....	1	vlan-301
AS	▶ 0.0.0.0/0	208.67.222.123 recursive via 186....	1	vlan-201
AS	▶ 0.0.0.0/0	208.67.220.123 recursive via 186....	1	vlan-101
AS	▶ 0.0.0.0/0	172.23.105.1 reachable ipip_vc5	1	to-scl-5
AS	▶ 0.0.0.0/0	172.23.104.1 reachable ipip_vc4	1	to-scl-4
AS	▶ 0.0.0.0/0	172.23.103.1 reachable ipip_vc3	1	to-scl-3
AS	▶ 0.0.0.0/0	172.23.102.1 reachable ipip_vc2	1	to-scl-2
AS	▶ 0.0.0.0/0	172.23.101.1 reachable ipip_vc1	1	to-scl-1
AS	▶ 0.0.0.0/0	208.67.220.123 recursive via 186....	1	full-balancer

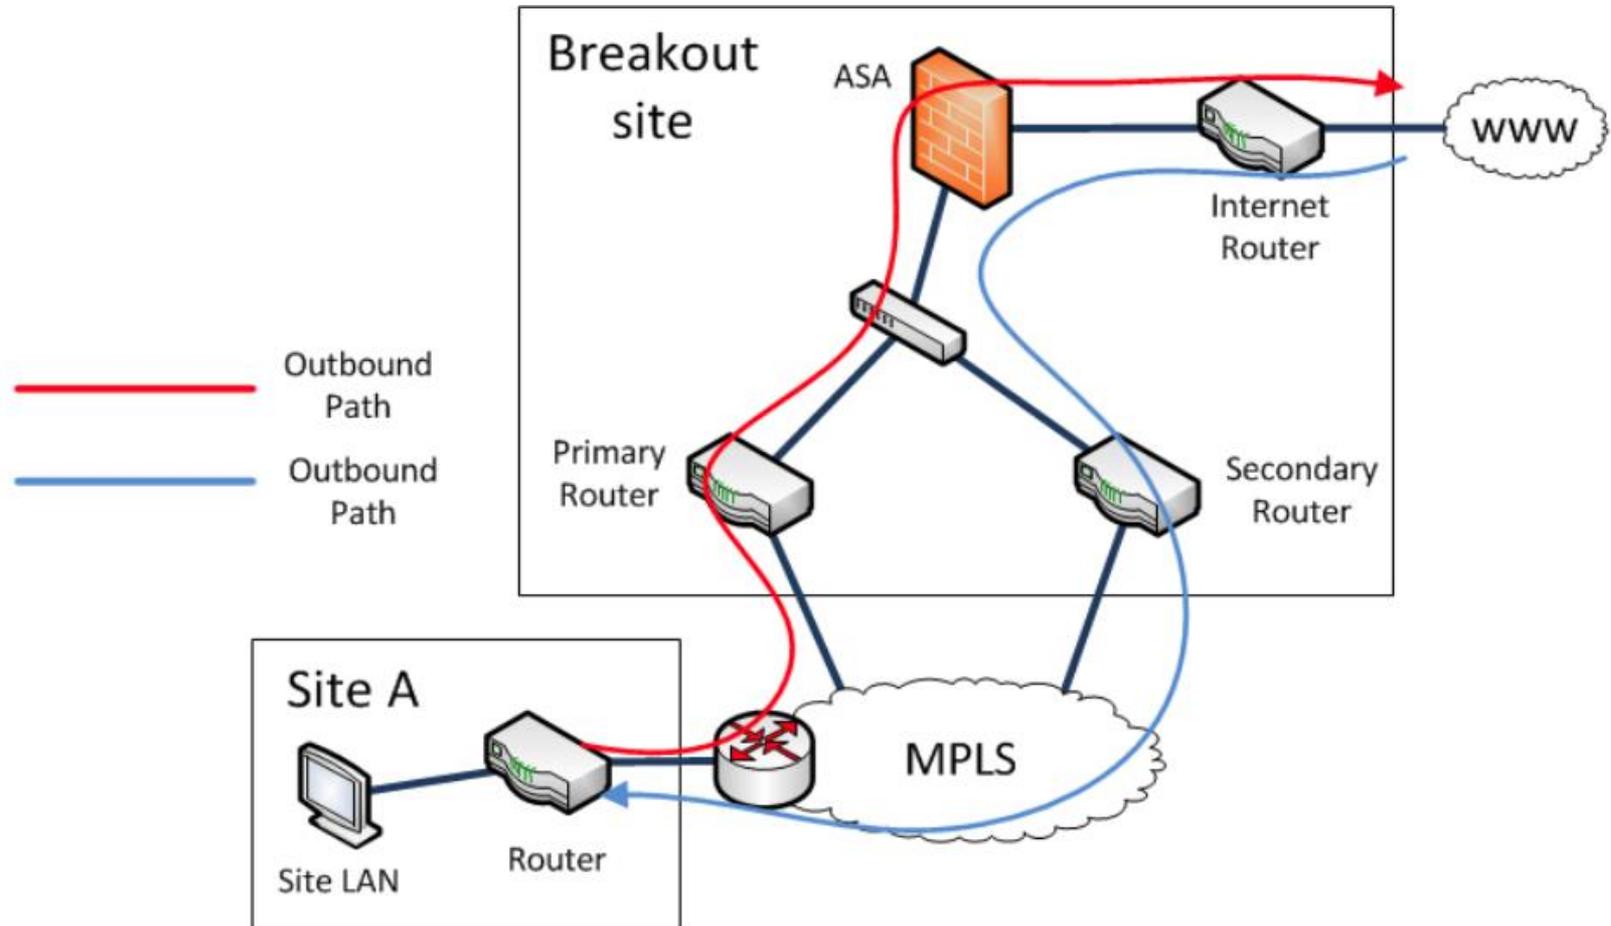
Balanceo BGP: multihomed

32

- Las conexiones NO son persistentes
 - ▣ 10.88.77.22 → 8.8.8.8
 - ▣ Sale por WAN 1 / Regresa por WAN 2
- Se puede aplicar PCC / ECMP y NTH
 - ▣ No se suele hacer
- Se debe establecer una sesión BGP con el upstream
 - ▣ Se anuncian los prefijos de red propios de la empresa o de tránsito (terceros)
 - ▣ Se aplican filtros
 - Entrantes
 - Salientes

Conexiones NO son persistentes

33



Balanceo BGP: multihomed

34

□ Saliente

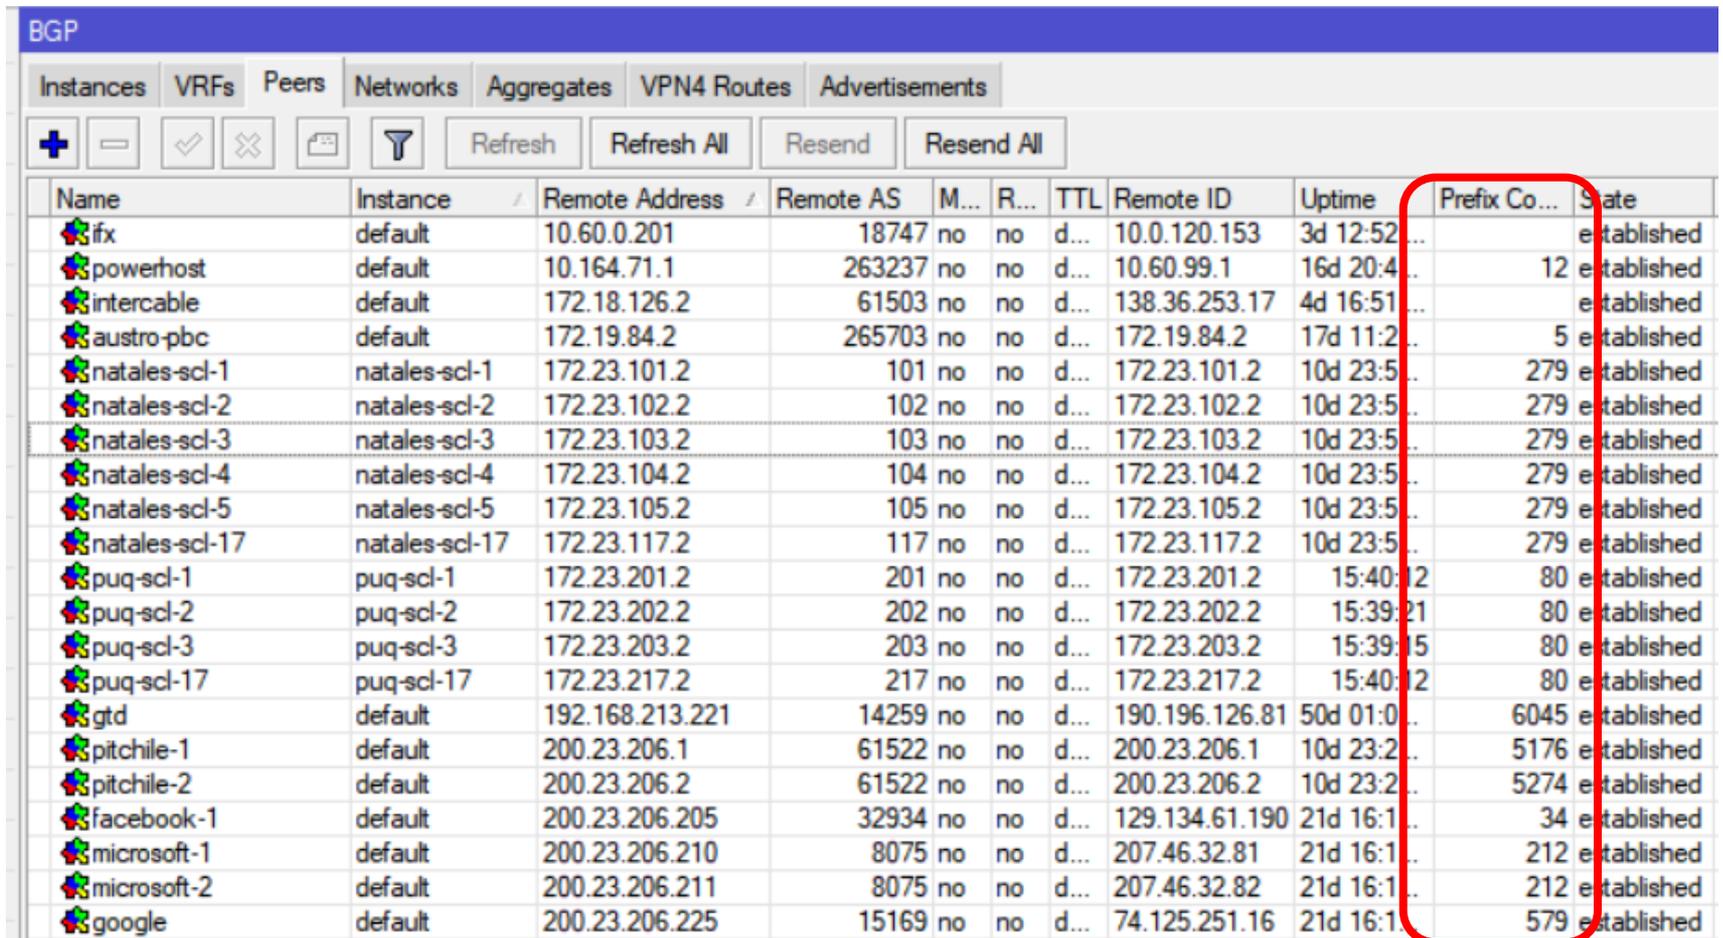
- Se basa en la tabla de rutas
 - Filtros entrantes

□ Entrante

- Se basa en la **información entregada a los ISP** de tránsito (upstream)
- Prefijo + máscara de subred
- Prepend del AS = AS-PATH
- **Decisión propia de los ISP (no se puede alterar)**

Sesión BGP

35



The screenshot shows a BGP configuration tool interface with a table of BGP peers. The table has columns for Name, Instance, Remote Address, Remote AS, M..., R..., TTL, Remote ID, Uptime, Prefix Co..., and State. A red box highlights the 'Prefix Co...' column, which represents the number of prefixes received from each peer.

Name	Instance	Remote Address	Remote AS	M...	R...	TTL	Remote ID	Uptime	Prefix Co...	State
ifx	default	10.60.0.201	18747	no	no	d...	10.0.120.153	3d 12:52...		established
powerhost	default	10.164.71.1	263237	no	no	d...	10.60.99.1	16d 20:4...	12	established
intercable	default	172.18.126.2	61503	no	no	d...	138.36.253.17	4d 16:51...		established
austro-pbc	default	172.19.84.2	265703	no	no	d...	172.19.84.2	17d 11:2...	5	established
natales-scl-1	natales-scl-1	172.23.101.2	101	no	no	d...	172.23.101.2	10d 23:5...	279	established
natales-scl-2	natales-scl-2	172.23.102.2	102	no	no	d...	172.23.102.2	10d 23:5...	279	established
natales-scl-3	natales-scl-3	172.23.103.2	103	no	no	d...	172.23.103.2	10d 23:5...	279	established
natales-scl-4	natales-scl-4	172.23.104.2	104	no	no	d...	172.23.104.2	10d 23:5...	279	established
natales-scl-5	natales-scl-5	172.23.105.2	105	no	no	d...	172.23.105.2	10d 23:5...	279	established
natales-scl-17	natales-scl-17	172.23.117.2	117	no	no	d...	172.23.117.2	10d 23:5...	279	established
puq-scl-1	puq-scl-1	172.23.201.2	201	no	no	d...	172.23.201.2	15:40:12	80	established
puq-scl-2	puq-scl-2	172.23.202.2	202	no	no	d...	172.23.202.2	15:39:21	80	established
puq-scl-3	puq-scl-3	172.23.203.2	203	no	no	d...	172.23.203.2	15:39:15	80	established
puq-scl-17	puq-scl-17	172.23.217.2	217	no	no	d...	172.23.217.2	15:40:12	80	established
gtd	default	192.168.213.221	14259	no	no	d...	190.196.126.81	50d 01:0...	6045	established
pitchile-1	default	200.23.206.1	61522	no	no	d...	200.23.206.1	10d 23:2...	5176	established
pitchile-2	default	200.23.206.2	61522	no	no	d...	200.23.206.2	10d 23:2...	5274	established
facebook-1	default	200.23.206.205	32934	no	no	d...	129.134.61.190	21d 16:1...	34	established
microsoft-1	default	200.23.206.210	8075	no	no	d...	207.46.32.81	21d 16:1...	212	established
microsoft-2	default	200.23.206.211	8075	no	no	d...	207.46.32.82	21d 16:1...	212	established
google	default	200.23.206.225	15169	no	no	d...	74.125.251.16	21d 16:1...	579	established

Prefijos recibidos = caminos posibles

Tabla de rutas basada en las sesiones BGP

36

Route List

Routes Nexthops Rules VRF

BGP is yes

	Dst. Address	Gateway	Pref. Source	BGP AS Path
DAb	▶ 216.252.222.0/24	200.23.206.225 reachable vlan-25_pit	160.20.188.1	15169
Db	▶ 216.252.222.0/24	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,15169
DAb	▶ 216.252.220.0/24	200.23.206.225 reachable vlan-25_pit	160.20.188.1	15169
Db	▶ 216.252.220.0/24	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,15169
DAb	▶ 216.252.220.0/22	200.23.206.225 reachable vlan-25_pit	160.20.188.1	15169
Db	▶ 216.252.220.0/22	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,15169
Db	▶ 216.241.29.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
DAb	▶ 216.241.29.0/24	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,19228,18747
Db	▶ 216.241.29.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
Db	▶ 216.241.26.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
Db	▶ 216.241.26.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
DAb	▶ 216.241.26.0/24	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,19228,18747
DAb	▶ 216.241.25.0/24	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,19228,18747
Db	▶ 216.241.25.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
Db	▶ 216.241.25.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
DAb	▶ 216.241.17.0/24	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,19228,18747
Db	▶ 216.241.17.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
Db	▶ 216.241.17.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
Db	▶ 216.241.14.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
Db	▶ 216.241.14.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
DAb	▶ 216.241.14.0/24	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,19228,18747
Db	▶ 216.241.13.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
DAb	▶ 216.241.13.0/24	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,19228,18747
Db	▶ 216.241.13.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
Db	▶ 216.241.11.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
Db	▶ 216.241.11.0/24	200.23.206.20 reachable vlan-25_pit	160.20.188.1	265703,265703,265703,265703,265703,61522,18747
DAb	▶ 216.241.11.0/24	192.168.213.221 reachable ether1_gtd-internet	160.20.188.1	14259,19228,18747

Tabla de rutas: MAIN

38

- **Múltiples caminos** para alcanzar un destino
- Solo **1 ruta activa** para un destino específico
- Se tomará siempre la **ruta mas específica**
- Se tomará la ruta que tenga el **AS-PATH mas corto**
 - ▣ Se puede modificar con los filtros entrantes
 - ▣ Prepend AS-PATH
- Distancia no es útil para la decisión
 - ▣ **BGP Best Path Selection Algorithm**
 - ▣ https://wiki.mikrotik.com/wiki/Manual:BGP_Best_Path_Selection_Algorithm

Best Path Selection Algorithm

BEST PATH ALGORITHM

1. Router is ignoring received path if the route is not **valid**. Route is valid if:
 - NEXT_HOP of the route is valid and reachable
 - AS_PATH received from external peers does not contain the local AS
 - route is not rejected by routing filters

For more information read [nexthop selection and validation](#).

2. The first path received is automatically considered 'best path'. Any further received paths are compared to first received to determine if the new path is better.
3. Prefer the path with the highest **WEIGHT**.

WEIGHT parameter is local to the router on which it is configured. A route without assigned WEIGHT have a default value of 0.

4. Prefer the path with the highest **LOCAL_PREF**. It is used only within an AS.

A path without **LOCAL_PREF** attribute have a value of 100 by default.

5. Prefer the path with the shortest **AS_PATH**. (skipped if `ignore-as-path-len` set to **yes**)

Each **AS_SET** counts as 1, regardless of the set size. The **AS_CONFED_SEQUENCE** and **AS_CONFED_SET** are not included in the **AS_PATH** length.

6. Prefer the path that was locally originated via [aggregate](#) or [BGP network](#)
7. Prefer the path with the lowest **ORIGIN** type.

Interior Gateway Protocol (IGP) is lower than Exterior Gateway Protocol (EGP), and EGP is lower than INCOMPLETE

in other words **IGP < EGP < INCOMPLETE**

8. Prefer the path with the lowest **multi-exit discriminator** (MED).

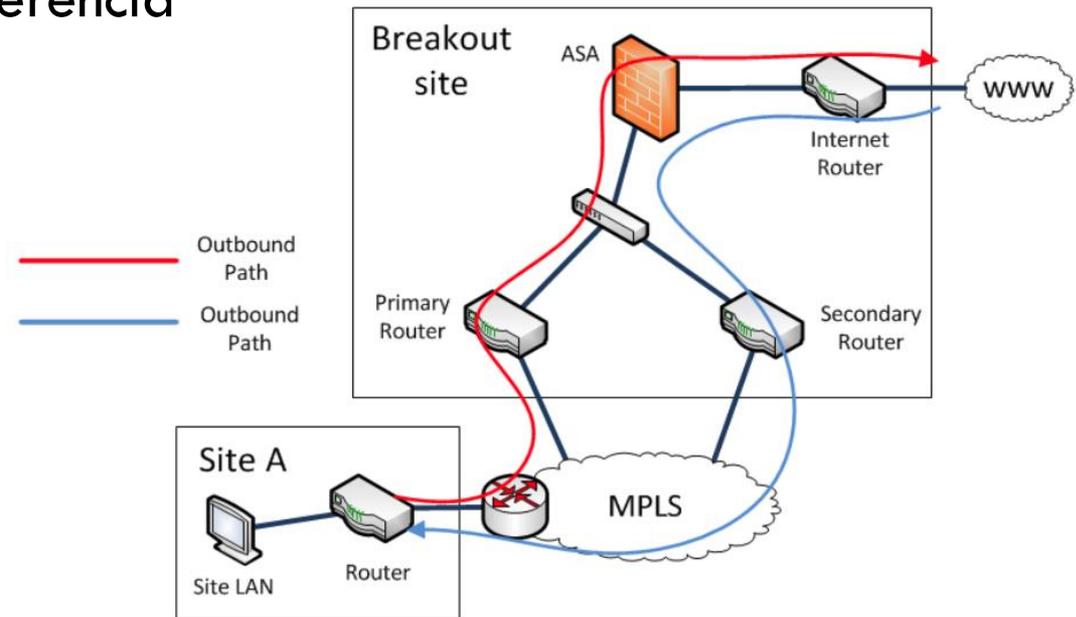
The router compare MED attribute only for paths that have the same neighboring (leftmost) AS. Paths without explicit MED value are treated as with MED of 0

9. Prefer **eBGP** over **iBGP** paths
10. Prefer the route that comes from the BGP router with the lowest **router ID**. If a route carries the **ORIGINATOR_ID** attribute, then the **ORIGINATOR_ID** is used instead of router ID.
11. Prefer the route with the shortest **route reflection cluster list**. Routes without a cluster list are considered to have a cluster list of length 0.
12. Prefer the path that comes from the lowest neighbor address

Best Path Selection Algorithm

40

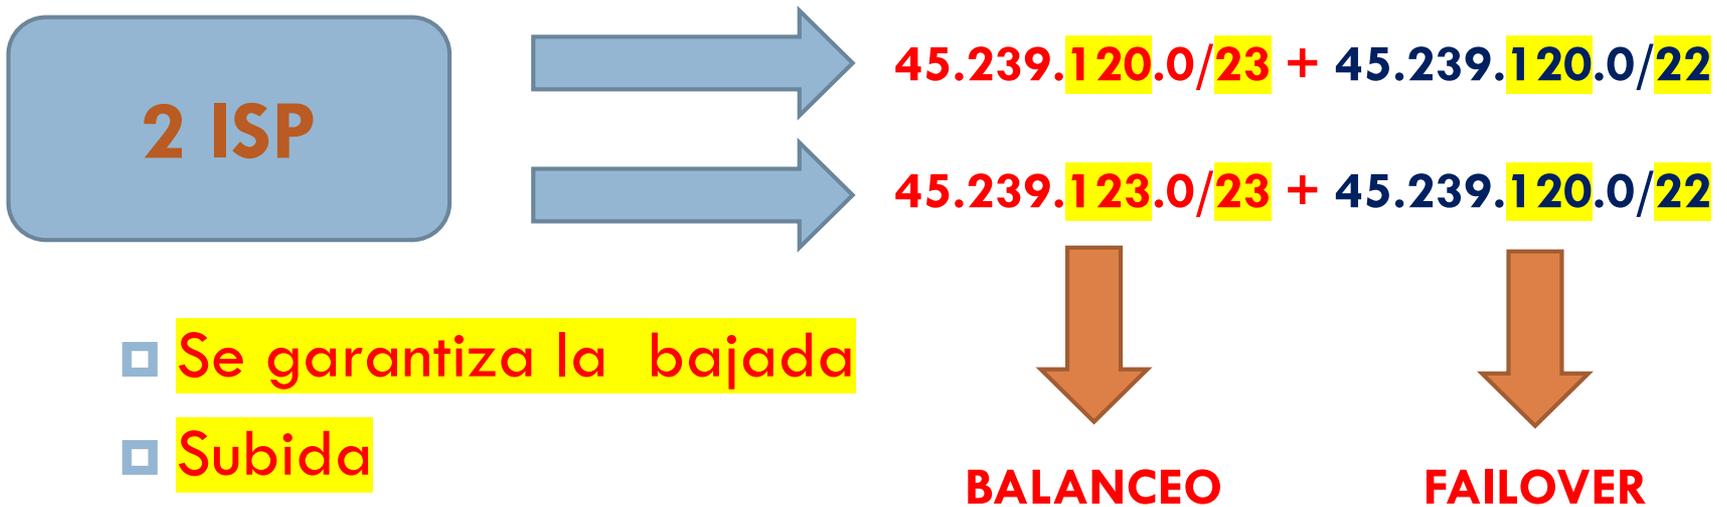
- Ejecutado por **cada router de la RED** del UPSTREAM
- **No podemos controlar** la interpretación del upstream
- Podemos intentar **“persuadir”** a los upstream
 - Tomar alguna preferencia
 - **Balanceo entrante**



BGP: balanceo entrante / Escenario

41

- Basado en que LACNIC otorgó 1 prefijo /22
- Usando **filtros de salida**
 - ▣ Modificar los anuncios (publicaciones)

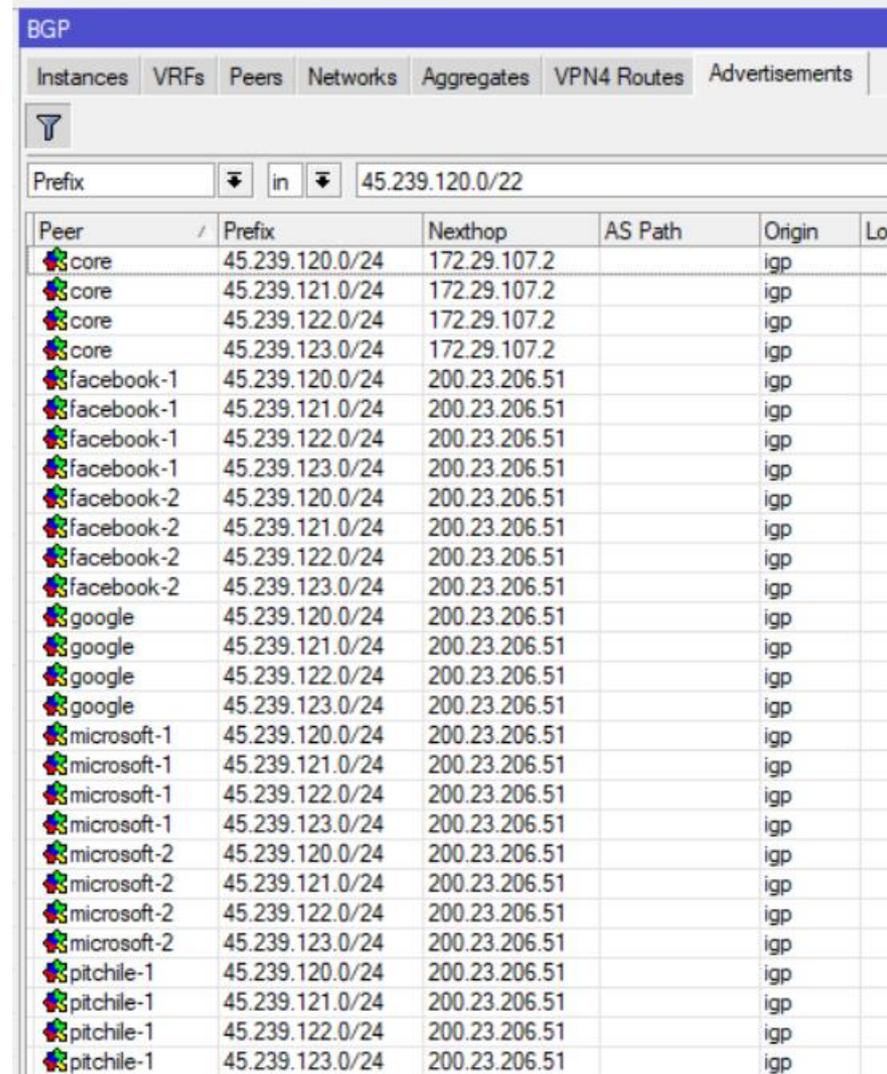


- ▣ Se garantiza la bajada
- ▣ Subida
 - Depende de la MAIN
 - Se puede formar salida con mangle

Publicación por defecto: persistencia

42

- En este escenario **NO** se puede **garantizar la persistencia**
- Bajo este escenario **se puede saturar** alguno de los enlaces dado
- Es un buen escenario para **failover automático**
- **Balanceo** basado en la **decisión de los upstream**



The screenshot shows a BGP management interface with a table of advertisements. The table has columns for Peer, Prefix, Nexthop, AS Path, Origin, and Location. The data is as follows:

Peer	Prefix	Nexthop	AS Path	Origin	Lo
core	45.239.120.0/24	172.29.107.2		igp	
core	45.239.121.0/24	172.29.107.2		igp	
core	45.239.122.0/24	172.29.107.2		igp	
core	45.239.123.0/24	172.29.107.2		igp	
facebook-1	45.239.120.0/24	200.23.206.51		igp	
facebook-1	45.239.121.0/24	200.23.206.51		igp	
facebook-1	45.239.122.0/24	200.23.206.51		igp	
facebook-1	45.239.123.0/24	200.23.206.51		igp	
facebook-2	45.239.120.0/24	200.23.206.51		igp	
facebook-2	45.239.121.0/24	200.23.206.51		igp	
facebook-2	45.239.122.0/24	200.23.206.51		igp	
facebook-2	45.239.123.0/24	200.23.206.51		igp	
google	45.239.120.0/24	200.23.206.51		igp	
google	45.239.121.0/24	200.23.206.51		igp	
google	45.239.122.0/24	200.23.206.51		igp	
google	45.239.123.0/24	200.23.206.51		igp	
microsoft-1	45.239.120.0/24	200.23.206.51		igp	
microsoft-1	45.239.121.0/24	200.23.206.51		igp	
microsoft-1	45.239.122.0/24	200.23.206.51		igp	
microsoft-1	45.239.123.0/24	200.23.206.51		igp	
microsoft-2	45.239.120.0/24	200.23.206.51		igp	
microsoft-2	45.239.121.0/24	200.23.206.51		igp	
microsoft-2	45.239.122.0/24	200.23.206.51		igp	
microsoft-2	45.239.123.0/24	200.23.206.51		igp	
pitchile-1	45.239.120.0/24	200.23.206.51		igp	
pitchile-1	45.239.121.0/24	200.23.206.51		igp	
pitchile-1	45.239.122.0/24	200.23.206.51		igp	
pitchile-1	45.239.123.0/24	200.23.206.51		igp	

44

Podderación de enlaces

¿Dónde puedo aplicar ponderación?

45

- ECMP
 - ▣ Simple de implementar
- PCC
 - ▣ Complejo de implementar
 - ▣ Robusto y escalable
- NTH
 - ▣ No permite implementar

Ponderación: ECMP

46

Route <0.0.0.0/0>

General Attributes

Dst. Address: 0.0.0.0/0

Gateway:	160.160.160.1	▼	◆
	160.160.160.1	▼	◆
	160.160.160.1	▼	◆
	170.170.170.1	▼	◆

Check Gateway: ▼

Type: unicast ▼

Distance: 1 ▲

Scope: 30

Target Scope: 10

Routing Mark: ▼

Pref. Source: ▼

disabled active static

OK
Cancel
Apply
Enable
Comment
Copy
Remove

Ponderación: PCC

47

- add action=mark-connection chain=prerouting comment="SIN PONDERACION - 2 WAN / PCC" **connection-mark=no-mark** in-interface=ether2_laptop new-connection-mark=mc_to-wan1 **per-connection-classifier=both-addresses-and-ports:8/x**

both addresses and ports:8/0	mc_to-wan1	0 B	0	CON PONDERAC...
both addresses and ports:8/1	mc_to-wan1	0 B	0	CON PONDERAC...
both addresses and ports:8/2	mc_to-wan1	250 B	2	CON PONDERAC...
both addresses and ports:8/3	mc_to-wan2	0 B	0	CON PONDERAC...
both addresses and ports:8/4	mc_to-wan2	0 B	0	CON PONDERAC...
both addresses and ports:8/5	mc_to-wan3	0 B	0	CON PONDERAC...
both addresses and ports:8/6	mc_to-wan3	0 B	0	CON PONDERAC...
both addresses and ports:8/7	mc_to-wan4	0 B	0	CON PONDERAC...

48

Posibles problemas (tradicional)

Posibles problemas balanceando

49

- ❑ Problemas con páginas HTTPS
- ❑ Problemas con conexiones al mismo router
- ❑ Acceso a servidores internos por una WAN específica (DST-NAT)
- ❑ Forzar un host a salir por una wan específica
- ❑ Varios links con el mismo gateway

Problema con páginas cifradas (https)

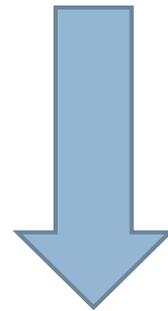
50

- Sitios que usan cifrado (bancos, google, youtube, facebook, etc) no permiten establecer múltiples conexiones provenientes de diferentes enlaces simultáneamente.
 - ▣ Phishing
- Solución
 - ▣ Mantener la **persistencia**
 - ▣ Usar ECMP o PCC con **BOTH ADDRESS**
 - Excluir los puertos origen y destino

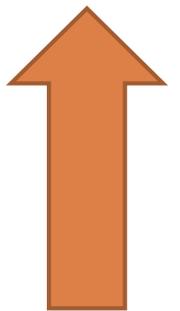
Conexiones dirigidas al router

51

- Paquetes salientes utilizan la misma decisión de ruteo que los paquetes que atraviesan el router.
- Las respuestas de un paquete que fue recibida por wan1 podría ser enviada de vuelta y enmascarada por wan2
- Solución
 - Marcar conexiones entrantes
 - Forzar a salir por la misma wan



wan1



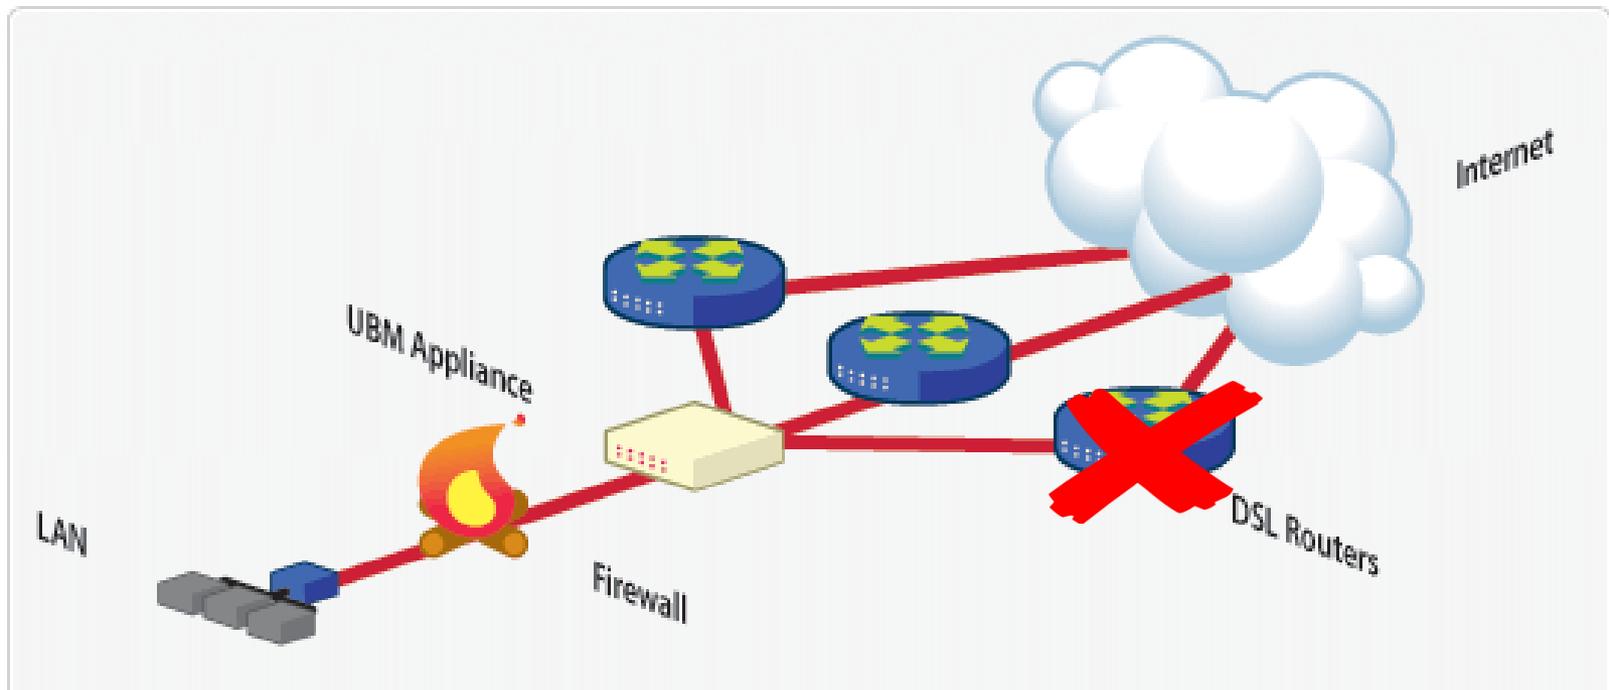
wan2

52

Failover

Failover

53



Cómo manejar Failover

54

□ Formas comunes

□ Scripts

- Haciendo ping a links externos.

□ Netwatch

- Haciendo ping a links externos.

□ Sin scripts ni Netwatch

□ Routes Nexthop Lookup

- Rutas recursivas
- Usando scope=10
- check-gateway=ping

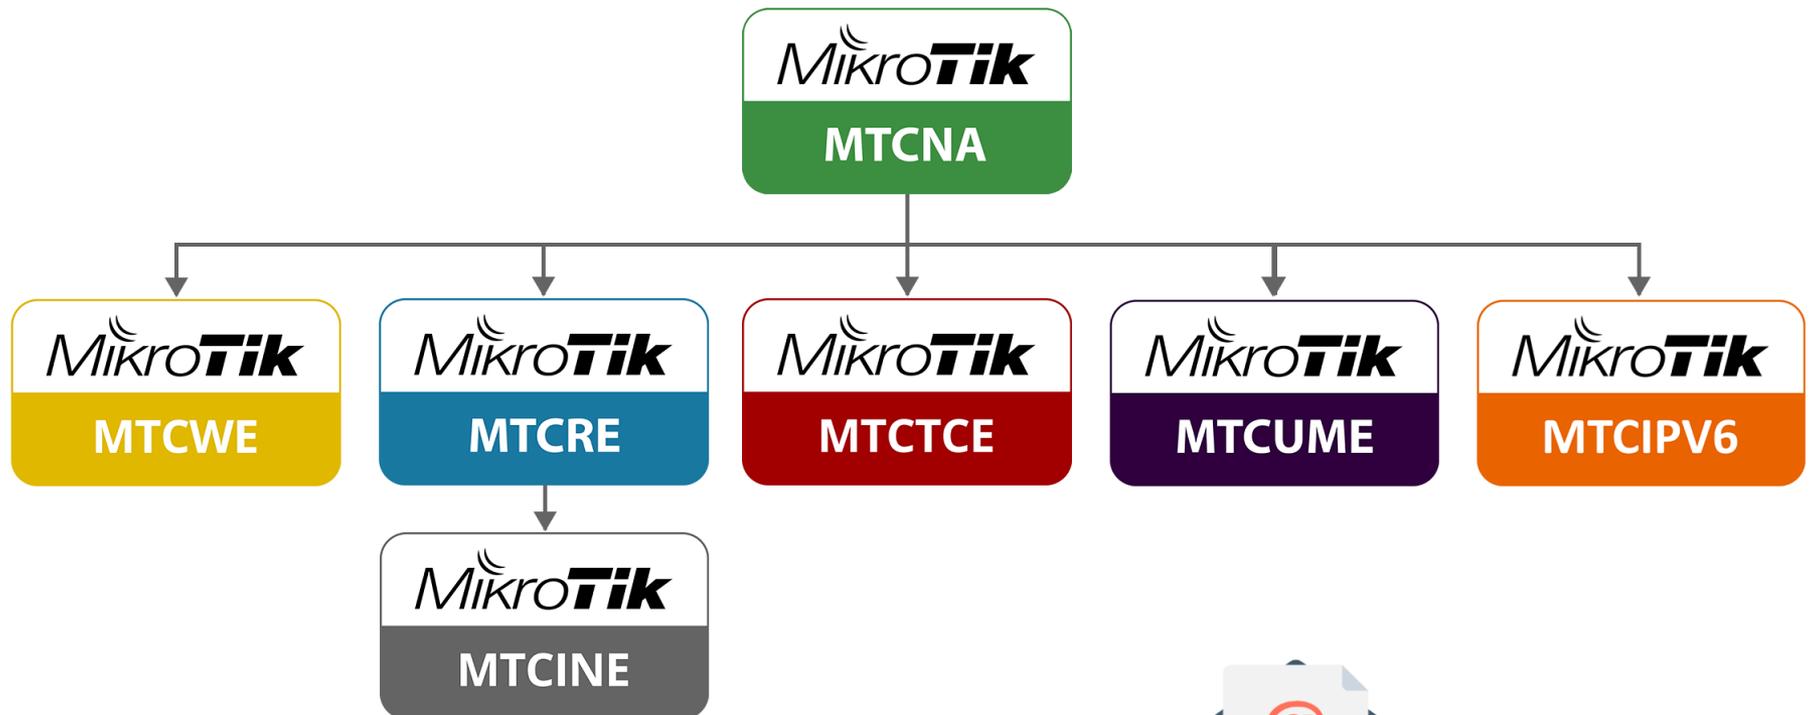


RUTAS RECURSIVAS

Cursos oficiales MikroTik RouterOS

55

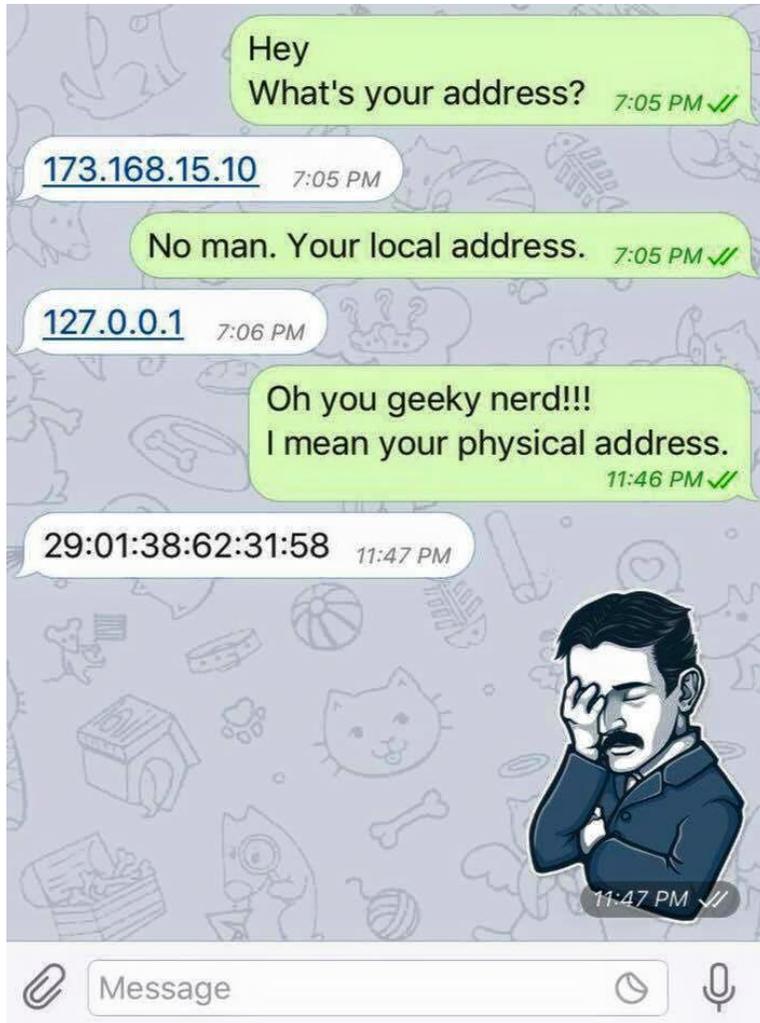
□ Cursos de certificación oficial MikroTik



at@mkx.cl

¿Preguntas?

56



Gracias por su atención