



# ISP Operations–Troubleshooting BGP en RouterOS

PRESENTED BY:

**ANDRES OCAMPO,**  
SR NETWORK ENGINEER

## Background:

- 5+ años en Networking / ISP / WISP
- Diseñado/Creado redes en todo América
- Certificado MikroTik y Cisco



## Trayectoria:

**Puntonet** (Operador de NOC / Activaciones / Soporte VIP )

**Telefónica Ec** (Ing. Sr. Soporte Servicio de Datos)

**Fibramax** (Supervisor General de Telecomunicaciones)

## Certificaciones:

**Cisco** (CCNA R&S, CCNA Security, CCNP R&S)

**Mikrotik** (MTCNA, MTCTCE)

**TEC Monterrey** (Maestría en Administración de TI)

IP

# Perfil: Sobre IP ArchiTechs



## Expertos en Networking

Whitebox | ISP | Data Center | Enterprise

- ✓ Consultoría Global
- ✓ Administración de Redes
- ✓ Monitoreo
- ✓ Test de Carga
- ✓ Desarrollo

Locations in: US | Canada | South America

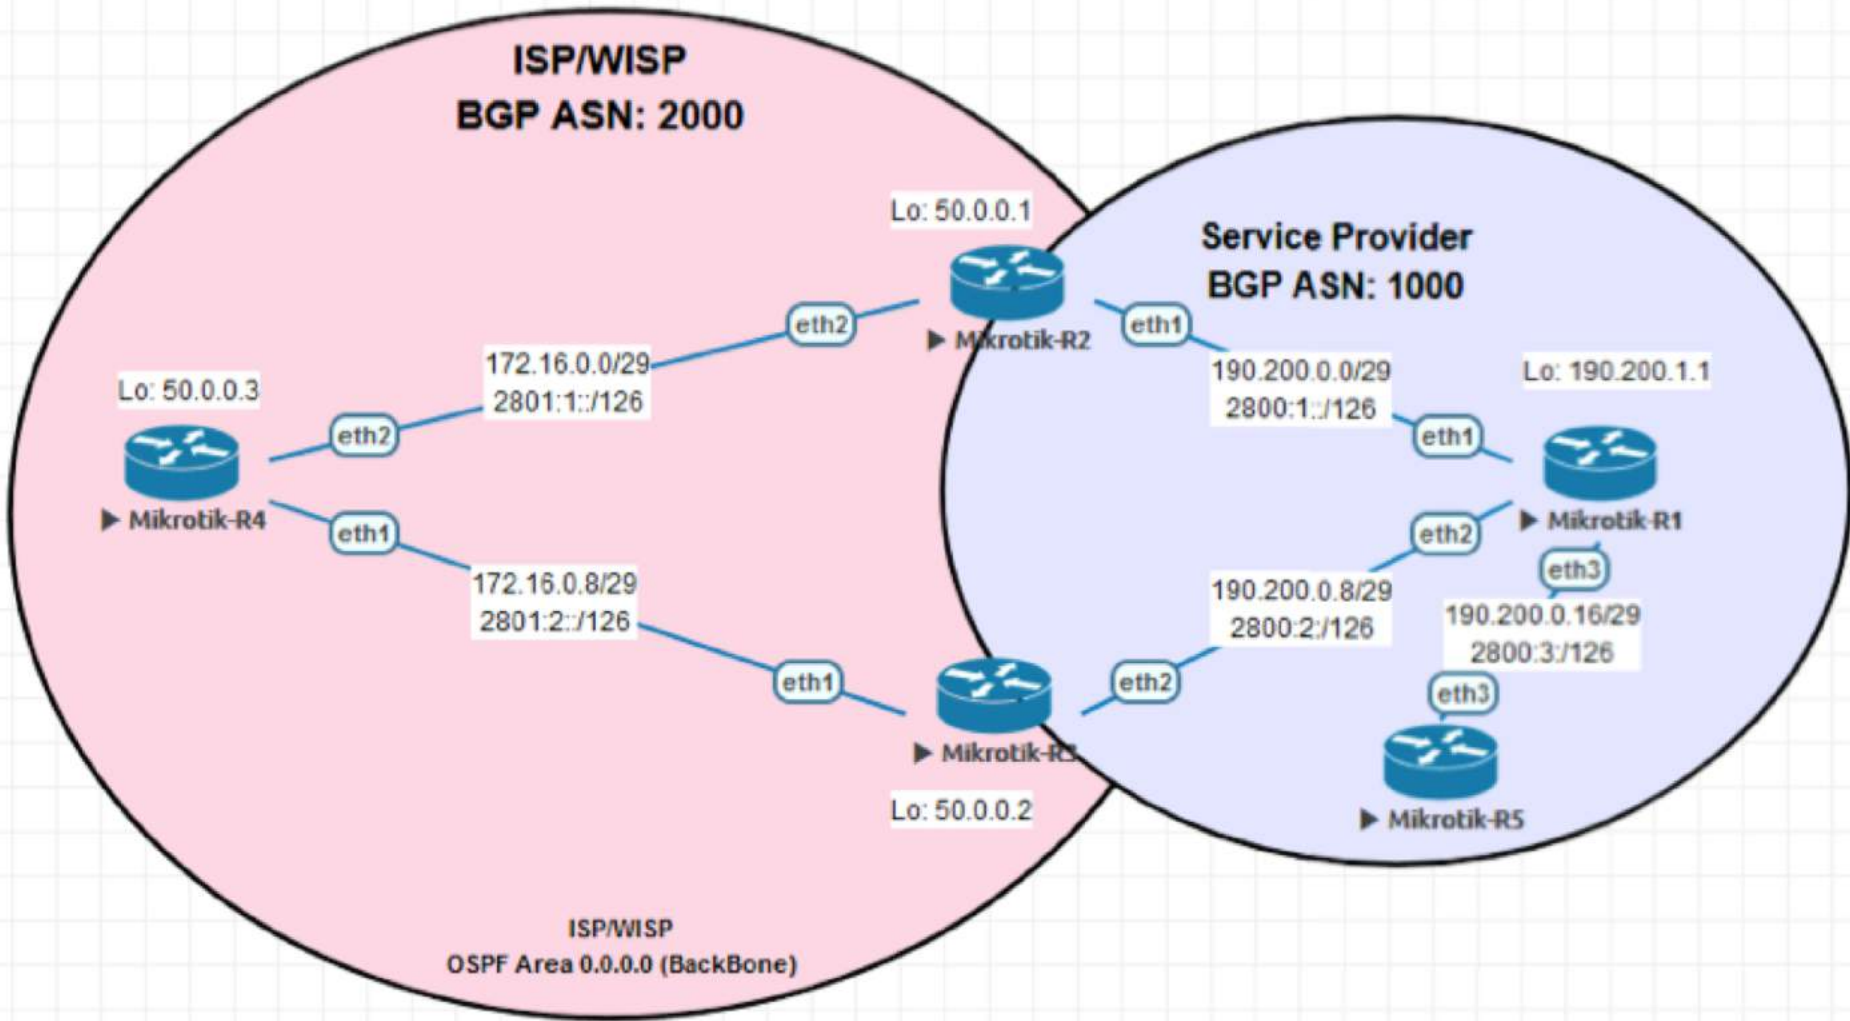
Call us at: +1 855-645-7684

E-mail: [consulting@iparchitechs.com](mailto:consulting@iparchitechs.com)

Web: [www.iparchitechs.com](http://www.iparchitechs.com)

**Meta de esta presentación:** Cuando esta presentación finalice, abremos aclarado ciertos conceptos clave:

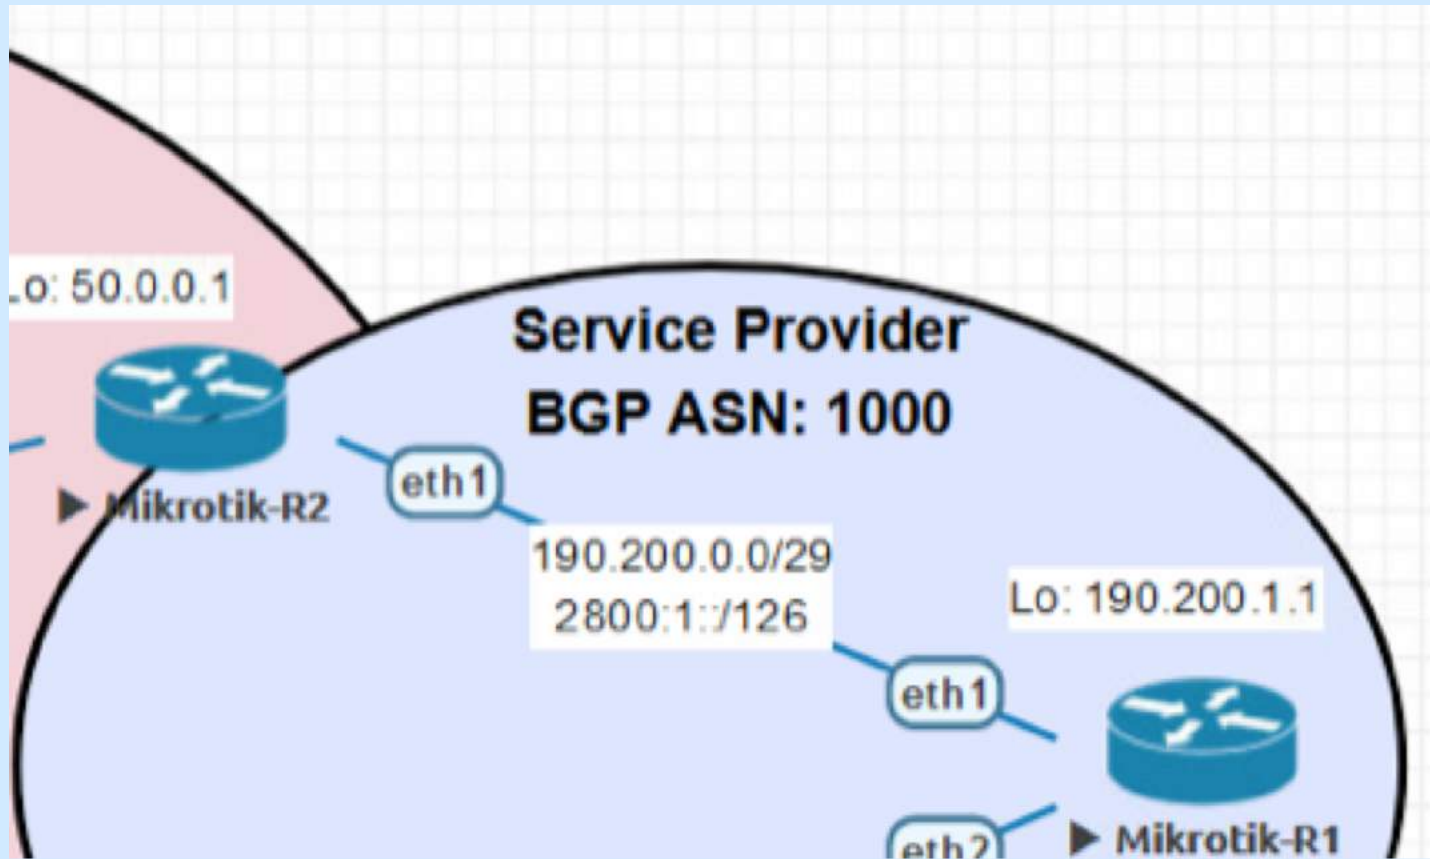
- Cómo resolver problemas de BGP en RouterOS utilizando verificación y logging.
- Que causas afectan a que una sesión BGP se establezca adecuadamente.



- BGP require los siguientes requisitos para establecer una sesión entre dos Peers:
  - Local AS / Remote AS
  - Autenticación
  - Multihop (eBGP únicamente)
  - MTU
- Veamos que puede pasar cuando no se cumplen estos requisitos:

# Escenario 1: Remote AS Mismatch

<b>BGP</b>	<b>Remote AS Mismatch</b>	<b>R1 to R2</b>
------------	---------------------------	-----------------





BGP	Remote AS Mismatch	R1 to R2
<b>R1</b> <pre>/routing bgp instance set default as=1000  /routing bgp peer add name=To-MikrotikR2 remote-address=190.200.0.2 \ remote-as=200</pre>		
<b>R2</b> <pre>/routing bgp instance set default as=2000 router-id=50.0.0.1  /routing bgp peer add name=To-MikrotikR1 remote-address=190.200.0.1 \ remote-as=1000</pre>		

**BGP****Remote AS Mismatch****R1 to R2****R1**

```
[admin@Mikrotik-R1] > routing bgp peer print status
Flags: X - disabled, E - established
 0          name="Mikrotik-R2"      instance=default      remote-
address=190.200.0.2  remote-as=200      tcp-md5-key=""      nexthop-
choice=default      multihop=no      route-reflect=no      hold-time=3m
ttl=255 in-filter="" out-filter="" address-families=ip update-
source=ether1      default-originate=always      remove-private-as=no
as-override=no passive=no use-bfd=no state=opensent
```

**R2**

```
[admin@Mikrotik-R2] > routing bgp peer print status
Flags: X - disabled, E - established
 0          name="To-MikrotikR1"    instance=default      remote-
address=190.200.0.1  remote-as=1000    tcp-md5-key=""      nexthop-
choice=default      multihop=no      route-reflect=no      hold-time=3m
ttl=255 in-filter="" out-filter="" address-families=ip update-
source=ether1      default-originate=never      remove-private-as=no
as-override=no passive=no use-bfd=no state=idle
```

**BGP****Remote AS Mismatch****R1 to R2****R1**

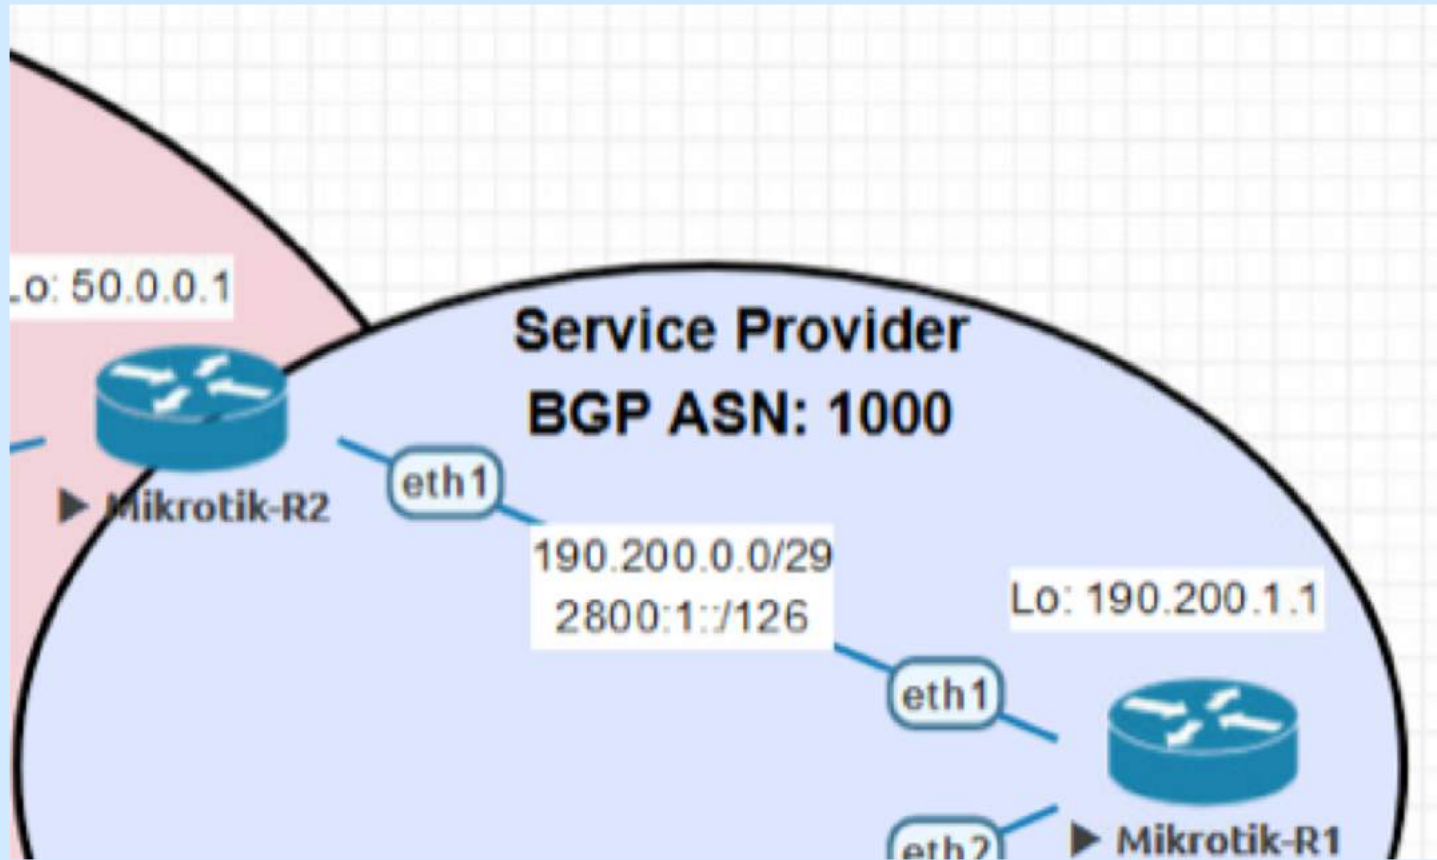
```
02:24:21 route,bgp,info Connection opened by remote
host
02:24:21 route,bgp,info RemoteAddress=190.200.0.2
02:24:21 route,bgp,error Invalid remote AS, expected
200
02:24:21 route,bgp,error RemoteAddress=190.200.0.2
02:24:21 route,bgp,error RemoteAs=2000
```

**R2**

```
02:24:21 route,bgp,info TCP connection established
02:24:21 route,bgp,info RemoteAddress=190.200.0.1
02:24:21 route,bgp,error Received notification
02:24:21 route,bgp,error OPEN error: bad remote-as
```

# Escenario 2: Dirección IP Incorrecta

BGP	Dirección IP Incorrecta	R1 to R2
-----	-------------------------	----------



BGP	Dirección IP Incorrecta	R1 to R2
<b>R1</b> <pre data-bbox="181 411 656 521">/routing bgp instance set default as=1000  /routing bgp peer add name=To-MikrotikR2 remote-address=190.200.0.20 \ remote-as=200</pre>		
<b>R2</b> <pre data-bbox="181 901 1023 1011">/routing bgp instance set default as=2000 router-id=50.0.0.1  /routing bgp peer add name=To-MikrotikR1 remote-address=190.200.0.1 \ remote-as=1000</pre>		

**BGP****Dirección IP Incorrecta****R1 to R2****R1**

```
[admin@Mikrotik-R1] > routing bgp peer print status
Flags: X - disabled, E - established
 0          name="Mikrotik-R2"    instance=default  remote-
address=190.200.0.20  remote-as=2000  tcp-md5-key=""  nexthop-
choice=default  multihop=no  route-reflect=no  hold-time=3m
ttl=255  in-filter=""  out-filter=""  address-families=ip
update-source=ether1  default-originate=always  remove-
private-as=no  as-override=no  passive=no  use-bfd=no
state=idle
```

**R2**

```
[admin@Mikrotik-R2] > routing bgp peer print status
Flags: X - disabled, E - established
 0          name="To-MikrotikR1"  instance=default  remote-
address=190.200.0.1  remote-as=1000  tcp-md5-key=""  nexthop-
choice=default  multihop=no  route-reflect=no  hold-time=3m
ttl=255  in-filter=""  out-filter=""  address-families=ip
update-source=ether1  default-originate=never  remove-private-
as=no  as-override=no  passive=no  use-bfd=no  state=opensent
```



## Operations: **BGP troubleshooting – log comparison**

**BGP**

**Dirección IP Incorrecta**

**R1 to R2**

**R1**

```
02:47:59 route,bgp,info Failed to open TCP connection: Network  
is unreachable
```

```
02:47:59 route,bgp,info RemoteAddress=190.200.0.20
```

**R2**

```
02:37:41 route,bgp,info TCP connection established
```

```
02:37:41 route,bgp,info RemoteAddress=190.200.0.1
```

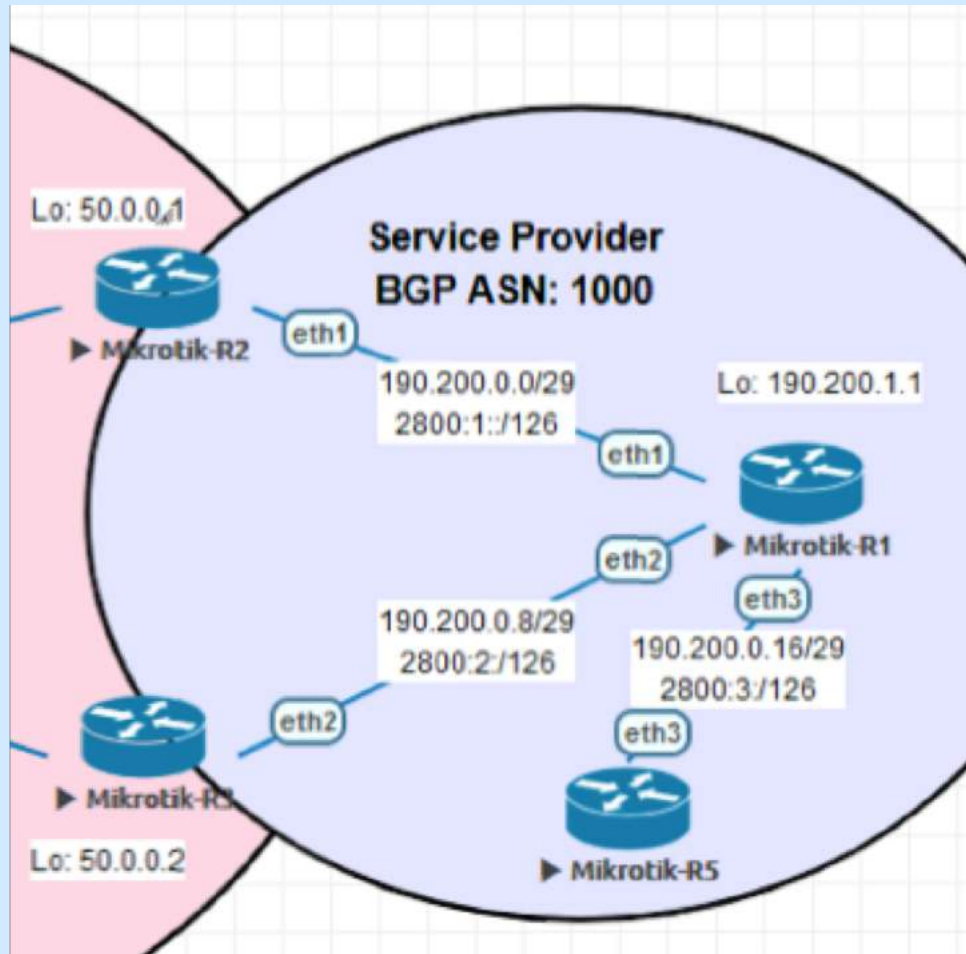
```
02:37:41 route,bgp,info Connection terminated
```

```
02:37:41 route,bgp,info RemoteAddress=190.200.0.1
```



# Escenario 3: MultiHop

<b>BGP</b>	<b>Multihop</b>	<b>R5 to R2</b>
------------	-----------------	-----------------



**BGP****MultiHop****R5 to R2****R5**

```
/routing bgp instance  
set default as=1000
```

```
/routing bgp peer  
add name=To-Mikrotik-R2 remote-address=190.200.0.2 remote-  
as=2000 update-source=ether3
```

**R2**

```
/routing bgp instance  
set default as=2000 router-id=50.0.0.1
```

```
/routing bgp peer  
add multihop=yes name=To-Mikrotik-R5 remote-  
address=190.200.0.18 remote-as=1000 update-source=ether1
```

**BGP****Multihop****R5 to R2****R5**

```
[admin@Mirktoik-R5] > routing bgp peer print status
Flags: X - disabled, E - established
name="To-Mikrotik-R2" instance=default remote-address=190.200.0.2
remote-as=2000 tcp-md5-key="" nexthop-choice=default multihop=no
route-reflect=no hold-time=3m ttl=255 in-filter="" out-filter=""
address-families=ip update-source=ether3 default-originate=never
remove-private-as=no as-override=no passive=no use-bfd=no
state=active
```

**R2**

```
[admin@Mikrotik-R2] > routing bgp peer print status
Flags: X - disabled, E - established
name="To-Mikrotik-R5" instance=default remote-address=190.200.0.18
remote-as=1000 tcp-md5-key="" nexthop-choice=default multihop=yes
route-reflect=no hold-time=3m ttl=255 in-filter="" out-filter=""
address-families=ip update-source=ether1 default-originate=never
remove-private-as=no as-override=no passive=no use-bfd=no
state=opensent
```

**BGP****Multihop****R5 to R2****R5**

```
03:27:20 route,bgp,info Connection opened by remote host
```

```
03:27:20 route,bgp,info      RemoteAddress=190.200.0.2
```

```
03:27:20 route,bgp,info Reject connection: EBGP peer is not on  
a shared network and multihop is not configured
```

**R2**

```
03:27:00 route,bgp,info      RemoteAddress=190.200.0.18
```

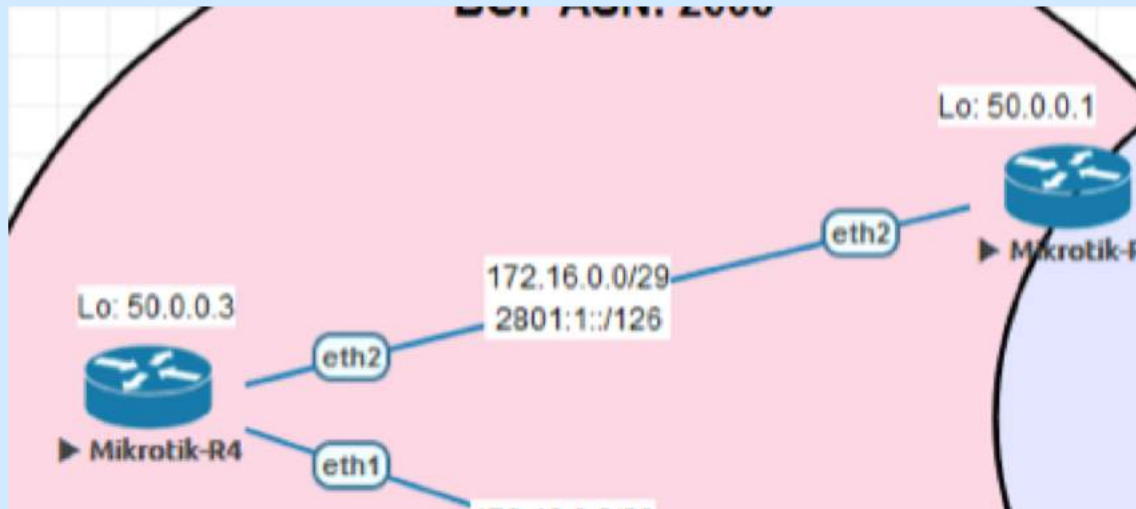
```
03:27:20 route,bgp,info TCP connection established
```

```
03:27:20 route,bgp,info      RemoteAddress=190.200.0.18
```

```
03:27:20 route,bgp,info Connection terminated
```

# Escenario 4: Autenticación

<b>BGP</b>	<b>Multihop</b>	<b>R4 to R2</b>
------------	-----------------	-----------------



<b>BGP</b>	<b>Authentication Mismatch</b>	<b>R4 to R2</b>
<p><b>R4</b></p> <pre data-bbox="185 449 1277 544">/routing bgp instance set default as=2000 router-id=50.0.0.3  /routing bgp peer add name=Ro-Mikrotik-R2 remote-address=50.0.0.1 remote-as=2000 update-source=Lo0</pre>		
<p><b>R2</b></p> <pre data-bbox="185 935 1711 1319">/routing bgp instance set default as=2000 router-id=50.0.0.1  /routing bgp peer add default-originate=if-installed name=To-MikrotikR4 remote-address=50.0.0.3 remote-as=2000 update- source=Lo0 tcp-md5-key=demo</pre>		



**BGP****Authentication Mismatch****R4 to R2****R4**

```
[[admin@Mikrotik-R4] > routing bgp peer print statu
Flags: X - disabled, E - established
 1  name="Ro-Mikrotik-R2" instance=default remote-
address=50.0.0.1 remote-as=2000 tcp-md5-key="" nexthop-
choice=default multihop=no route-reflect=no hold-time=3m
ttl=255 in-filter="" out-filter="" address-families=ip
default-originate=never remove-private-as=no as-override=no
passive=no use-bfd=no state=opensent
```

**R2**

```
[admin@Mikrotik-R2] > routing bgp peer print status
Flags: X - disabled, E - established
 1  name="To-MikrotikR4" instance=default remote-
address=50.0.0.3 remote-as=2000 tcp-md5-key="demo" nexthop-
choice=default multihop=no route-reflect=no hold-time=3m
ttl=255 in-filter="" out-filter="" address-families=ip update-
source=Lo0 default-originate=if-installed remove-private-as=no
as-override=no passive=no use-bfd=no state=connect
```



**BGP**

**Authentication Mismatch**

**R<sub>4</sub> to R<sub>2</sub>**

**R<sub>4</sub>**

No existen logs sobre este error

**R<sub>2</sub>**

No existen logs sobre este error

# Escenario 5: Otras causas posibles

- Si luego de revisar todo, la sesión sigue con problemas, entonces hay otras posibles causas de este fallo:
  - MTU mismatch
  - Reglas de Firewall
  - Reglas de NAT

- Causas:
  - L3 MTU diferente entre los peers
  - L2 MTU diferente entre ambos equipos
  - Path MTU discovery determinó un tamaño incorrecto de Ventana para la sesión TCP de BGP
  - BGP Path Maximum Transmission Unit Discovery (PMTUD) puede fallar por bloqueos de paquetes PMTUD ICMP (firewall o ACL)
- Efectos
  - Sesión BGP inestable

- Causas:
  - Reglas de Firewall bloqueando la sincronización de BGP (TCP/179).
  - NAT que enmascara la IP original del neighbor con otra del mismo router.
- Efectos
  - No se puede levantar la sesión BGP

- Causas:
  - Envío de muchos prefijos hacia un router con limitación de cantidad de prefijos
- Efectos
  - Si se supera el Max-Prefix-limit el router terminará y bloqueará la conexión BGP hacia este peer.



Design: **Questions?**

# Preguntas??