

FiberCli

IPv6

MUM Valencia 2018

José Manuel Román

Carlos Cárdenas

*Mikro***Tik**

Jose Manuel Román

Experiencia

- C.E.O @ Fibercli.com & Cloud Networking Spain
 - 6+ Years Security Consultant and Analyst
 - 7+ Years Teaching Computer Networking
 - 17+ Years on IT Industry
- Master ITIL
 - CISA & CISSP Cert.
 - Mikrotik MTC-ALL-E & Trainer (TR)



Carlos Cárdenas Cebrián

Educación

- Técnico superior en Electrónica Industrial.
- Grado en Ingeniería Informática de Sistemas en la Universidad de Cádiz (Sin completar).
- Entusiasta del networking. Autodidacta
- Certificaciones técnicas en varios fabricantes.
- Certificaciones Mikrotik:

Experiencia

- 5 años como instalador de equipos wireless, sistemas de seguridad y reparando equipos de RF.
- 8 años como técnico de banco de pruebas en el campo de la Aeronáutica, especialidad en radiocomunicaciones y sistemas de ayuda a la navegación.
- 4 años como administrador de red en WISPs.
- Actualmente trabajando en AUNNA IT como administrador de red.



INNOVACIÓN TECNOLÓGICA EN
COMUNICACIONES Y REDES


IPv6

¿Por qué necesitamos IPV6?



IPv4 está prácticamente agotado

- El espacio total de direcciones IPv4 es de 2^{32} = 4,3 billones
- Alrededor de 3,739 billones de usuarios de Internet actualmente
- Dispositivos móviles constantemente conectados a Internet
- Internet de las cosas: ¿Cómo será Internet dentro de unos años?



IPv4 es un recurso cada vez más caro

- El coste del direccionamiento necesario aumenta constantemente.
- RIPE pronto agotará las pocas IPv4 que le quedan y habrá que acudir al mercado
- Problemas de IPs de “segunda mano”:
 - Blacklists
 - GeoIP
 - CDNs
 - Precio €€€


RIPE /24 BLOCK	RIPE /21 BLOCK
	
AUCTION	AUCTION
/24 Block Registered in RIPE	/21 Block Registered in RIPE
Current bid: \$4,608	Current bid: \$35,840
Avg. Cost Per Unit: \$18.00	Avg. Cost Per Unit: \$17.50
Bids: 1	Bids: 0
Closes in 2d 2h	Closes in 2d 3h

Fuente: <https://www.ipv4auctions.com/ripe-blocks>

¿Cómo podemos paliar el problema?


- NAT 444 (Medio Plazo)
 - Rompe modelo de conexión end to end.
 - No permite comunicación con IPv6.
 - Necesidad de identificar al usuario final ante requerimientos judiciales (logs):
 - Se recomienda usar CGNAT determinístico
 - <https://www.abogacia.es/2016/03/14/nat-autoria-en-ciberdelitos/>
- DUAL STACK: NAT444 + IPv6 (Largo Plazo)
 - El mejor mecanismo de transición.
 - Permite acceso a IPv4 e IPv6.
 - Coexistencia entre protocolos.
 - Transparente al usuario final.
 - Solventa algunos problemas del NAT444
 - Identifica a los usuarios ante la ley
 - Google, Youtube, Facebook, Netflix, etc. ya están accesibles por IPv6

CGNAT: Carrier Grade NAT



Recurso: [Maia, Wardner: "From IPv4 scarcity to IPv6 Abundance" MUM EU18](#)

DUAL STACK



- Transparente para el usuario final.
- Autoconfiguración de direcciones IPv6.
- El navegador usará IPv6 automáticamente si está disponible, en caso contrario, IPv4.



IPv6

¿Cómo conseguimos direcciones IPv6?



Conseguir un rango de direcciones IPv6

- Para poder solicitar un rango, la organización deberá:
 - Estar dado de alta en RIPE como LIR.
 - Tener un plan de asignación en un plazo de dos años.
- El rango de direcciones mínimo asignado es un /32:
 - Asignan hasta un /29 sin necesidad de justificación.
 - En función del número de clientes y extensión de la red, es posible solicitar más.

Conseguir un rango de direcciones IPv6

**RIPE NCC**
RIPE NETWORK COORDINATION CENTRE

RIPE Database (Whois) Website

Search IP Address or ASN

Your IP address is: 2a0b:5081:17:0:

Manage IPs and ASNs > Analyse > Participate > Get Support > Publications > About Us >



IPv4, IPv6 and ASNs
LIR Portal
RIPE Database
Resource Transfers and Mergers



RIPE NCC Statistics
RIPEstat
RIPE Atlas
Raw Datasets



RIPE NCC Membership
RIPE Community
Policy Development
Internet Governance



Contact Us
Training and Education
Documentation
Spam, Hacking and Phishing



RIPE Documents
RIPE Labs
IPv6 Info Centre
RIPE NCC Activity Plan



What We Do
Careers
Executive Board
Corporate Governance

News and Announcements

[RIPE NCC IPv4 Address Space Chart - Now Updated](#)
Publication date: 27 Sep 2018

[NRO NC Election Candidates Announced](#)
Publication date: 24 Sep 2018

>

Service Status Archived

✓ All of our services are operating normally.

>

Meetings Recent

RIPE 77
15 Oct 2018 - 19 Oct 2018 Amsterdam, The Netherlands

RIPE NCC General Meeting, Amsterdam, 17-19 October 2018
17 Oct 2018 - 19 Oct 2018 Amsterdam, the Netherlands

SEE 8 / RIPE NCC Regional Meeting
16 Apr 2019 - 17 Apr 2019 Sarajevo, Bosnia and Herzegovina

RIPE 78
20 May 2019 - 24 May 2019 Reykjavik, Iceland

>

Member Information

Become a Member

[List of Our Services](#)

[Billing Procedure and Fee Schedule](#)



Have you got a spam/hacking issue?

We want to help you find the responsible party. Please read our [abuse page](#) to find out who's responsible.

New on RIPE Labs

[Measuring Anycast DNS Services Using RIPE Atlas](#)
Publication date: 01 Oct 2018

GM - Register Your Vote Today!

RIPE NCC General Meeting - ...



Register Now! >

NRO News

[Preview: New NRO Website](#)

>

Conseguir un rango de direcciones IPv6

The screenshot shows the RIPE NCC website interface. At the top left is the RIPE NCC logo. To the right, a user profile for 'Carlos Cardenas Cebrian' is visible. Below the logo, a search bar contains the text 'Your IP address is: 2a0b:5081:17:'. A navigation bar below the search bar contains links: 'Manage IPs and ASNs', 'Analyse', 'Participate', 'Get Support', 'Publications', and 'About Us'. Under 'Manage IPs and ASNs', the 'LIR Portal' is highlighted with a red circle. Below the navigation bar, there are six columns of links. The first column includes 'IPv4, IPv6 and ASNs', 'LIR Portal', 'RIPE Database', and 'Resource Transfers and Mergers'. The second column includes 'RIPE NCC Statistics', 'RIPEstat', 'RIPE Atlas', and 'Raw Datasets'. The third column includes 'RIPE NCC Membership', 'RIPE Community', 'Policy Development', and 'Internet Governance'. The fourth column includes 'Contact Us', 'Training and Education', 'Documentation', and 'Spam, Hacking and Phishing'. The fifth column includes 'RIPE Documents', 'RIPE Labs', 'IPv6 Info Centre', and 'RIPE NCC Activity Plan'. The sixth column includes 'What We Do', 'Careers', 'Executive Board', and 'Corporate Governance'. Below the navigation bar, there are three main sections: 'News and Announcements' with links to 'RIPE NCC IPv4 Address Space Chart - Now Updated' and 'NRO NC Election Candidates Announced'; 'Meetings' with a 'Recent' tab showing 'RIPE 77' and 'RIPE NCC General Meeting, Amsterdam, 17-19 October 2018'; and 'Member Information' with links to 'Become a Member', 'List of Our Services', and 'Billing Procedure and Fee Schedule'. On the right side, there is a 'GM - Register Your Vote Today!' section with a video player and a 'Register Now!' button, and an 'NRO News' section with a link to 'Preview: New NRO Website'.

RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Carlos Cardenas Cebrian

RIPE Database (Whois) Website

Search IP Address or ASN

Your IP address is: 2a0b:5081:17:

Manage IPs and ASNs > Analyse > Participate > Get Support > Publications > About Us >

IPv4, IPv6 and ASNs
LIR Portal
RIPE Database
Resource Transfers and Mergers

RIPE NCC Statistics
RIPEstat
RIPE Atlas
Raw Datasets

RIPE NCC Membership
RIPE Community
Policy Development
Internet Governance

Contact Us
Training and Education
Documentation
Spam, Hacking and Phishing

RIPE Documents
RIPE Labs
IPv6 Info Centre
RIPE NCC Activity Plan

What We Do
Careers
Executive Board
Corporate Governance

News and Announcements

RIPE NCC IPv4 Address Space Chart - Now Updated
Publication date: 27 Sep 2018

NRO NC Election Candidates Announced
Publication date: 24 Sep 2018

Service Status Archived

✓ All of our services are operating normally.

Meetings Recent

RIPE 77
15 Oct 2018 - 19 Oct 2018 Amsterdam, The Netherlands

RIPE NCC General Meeting, Amsterdam, 17-19 October 2018
17 Oct 2018 - 19 Oct 2018 Amsterdam, the Netherlands

SEE 8 / RIPE NCC Regional Meeting
16 Apr 2019 - 17 Apr 2019 Sarajevo, Bosnia and Herzegovina

RIPE 78
20 May 2019 - 24 May 2019 Reykjavik, Iceland

Member Information

Become a Member

List of Our Services

Billing Procedure and Fee Schedule

Have you got a spam/hacking issue?

We want to help you find the responsible party. Please read our [abuse page](#) to find out who's responsible.

New on RIPE Labs

Measuring Anycast DNS Services Using RIPE Atlas

GM - Register Your Vote Today!


RIPE NCC General Meeting - ...

Register Now!

NRO News

Preview: New NRO Website

Conseguir un rango de direcciones IPv6



The screenshot shows the RIPE NCC LIR Portal interface. The top navigation bar includes 'Manage IPs and ASNs', 'Analyse', and 'Participate'. The breadcrumb trail indicates the current location: 'Home > Manage IPs and ASNs > LIR Portal'. On the left sidebar, under 'My LIR', the 'Resources' menu is expanded, showing options like 'My Resources', 'Request Resources', 'Request Transfer', 'IPv4 Transfer Listing Service', 'RPKI Dashboard', and 'RIPE Database'. The 'Request Resources' link is highlighted with a red circle. A red arrow points from this link to the 'IPv6 Addresses' option in the 'What would you like to request?' section, which is also circled in red. Below this section is a 'Continue' button. At the bottom, there is a link to 'supporting notes'.

RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Manage IPs and ASNs > Analyse > Participate >

You are here: Home > Manage IPs and ASNs > LIR Portal

My LIR >

Resources ▾


- My Resources
- Request Resources**
- Request Transfer
- IPv4 Transfer Listing Service
- RPKI Dashboard
- RIPE Database >

What would you like to request?

- IPv6 Addresses**
- IPv4 Addresses
- ASN (Autonomous System Number)
- IXP (Internet Exchange Point) Assignment

Continue

Not sure? Read the [supporting notes](#).



This screenshot shows a dialog box titled 'What kind of IPv6 address space do you want to request?'. It contains a list of radio button options. The first option, 'IPv6 Allocation', is selected and circled in red. A red arrow points from this option down to the 'Continue' button at the bottom of the dialog. The other options are 'IPv6 Provider Independent (PI) Assignment', 'Temporary IPv6 Provider Independent (PI) Assignment', 'IPv6 End User Site Assignment', and 'IPv6 Anycast Assignment'. The 'Continue' button is highlighted with a dark blue background.

What kind of IPv6 address space do you want to request?

- IPv6 Allocation**
- IPv6 Provider Independent (PI) Assignment
- Temporary IPv6 Provider Independent (PI) Assignment
- IPv6 End User Site Assignment
- IPv6 Anycast Assignment

< Continue

Conseguir un rango de direcciones IPv6

What IPv6 subnet size would you like?

/29

Specify the country in which the allocation will be announced

SPAIN

Person with administrative responsibility (admin-c)

-RIPE

Person with technical responsibility (tech-c)

CCC57-RIPE

Carlos Cardenas Cebrian

Do you already have an IPv6 range?

Enter IPv6 range

Maintainers, persons and roles are RIPE Database objects,
[click here to create them.](#)

☒ I confirm that I am going to make assignments and/or sub-allocations from the allocation, in accordance with [the IPv6 Allocation and Assignment policy](#)

☒ I confirm that the network that will be using these Internet number resources has an active element in the RIPE NCC service region.

< Continue

IPv6

El protocolo IPv6



IPv6 Básico

- Direcciones IPv6: 2^{128} (128 bits)
 - IPv4: 2^{32} (32 bits)
- Cada subred debe de ser como mínimo un /64
- Asignación a clientes finales:
 - /64 (1 subred)
 - /48 (65536 subredes)
- Asignación mínima al LIR /32
 - 65536 /48s
 - 16777216 /56
- No hace uso de broadcast: Multicast e ICMPv6

Dirección IPv6

Una dirección IPv6 (en hexadecimal)

2001:0DB8:AC10:FE01:0000:0000:0000:0000

↓ ↓ ↓ ↓ 

2001:0DB8:AC10:FE01:: Se pueden omitir los ceros




1000000000000001:0000110110111000:101011000010000:1111111000000001:
0000000000000000:0000000000000000:0000000000000000:0000000000000000

- 8 campos de 16 bit cada uno
- Usa notación hexadecimal (4 dígitos por campo)
- Cada grupo de 4 dígitos se separa por “:”

Mas info: [RFC5952](https://www.rfc-editor.org/rfc/rfc5952)


Subredes en IPv6 (/48 tomadas desde el /32 asignado por RIPE)



- 8 campos de 16 bit cada uno
- Usa notación hexadecimal (4 dígitos por campo)
- Cada grupo de 4 dígitos se separa por “:”

Mas info: [RFC5952](https://www.rfc-editor.org/rfc/rfc5952)


Subredes en IPv6



Mas info: [RFC5952](https://tools.ietf.org/html/rfc5952)

Cabecera IPv6


IPv4 Header



Legend

- Field's Name Kept from IPv4 to IPv6
- Fields Not Kept in IPv6
- Name and Position Changed in IPv6
- New Field in IPv6

IPv6 Header



Tipos de direcciones IPv6



Tipos de direcciones

Type	Range
Link local	fe80::/10
Global Unicast	2000::/3
Multicast	ff00::/8
Unique local	fc00::/7

Direcciones especiales

Type	Range
Loopback	::1/128
Documentation	2001:db8::/32
6to4	2002::/16
Unspecified address	::/128
Teredo	2001::/32
Anycast	Any unicast (enable "no DAD")

IPv6 es compatible



```
root@ns:~# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
   inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
   link/ether 08:00:27:00:00:00 brd ff:ff:ff:ff:ff:ff
   inet 10.200.1.166/24 scope global eth0
       valid_lft forever preferred_lft forever
   inet6 2a0b:5081:3000:0000:0000:0000:0000:0000/48 scope global
       valid_lft forever preferred_lft forever
   inet6 fe80::20c:29ff:fe63:c346/64 scope link
       valid_lft forever preferred_lft forever
```


IPv6

Plan de
direccionamiento



Plan de direccionamiento IPv6

- No es necesario preocuparse por agotar el espacio de direcciones IPv6.
- Agregación de rutas es siempre recomendable (Tabla de rutas más óptima).
- Importante: ¡Las Link Local address no aparecen en un traceroute!



BCOP RIPE-690

- Recomendaciones sobre tamaños de asignación y tiempo de persistencia:
 - Reservar para cada enlace PtP un /64 pero direccionar como /127.
 - Asignar al menos un /48 o un /56 por cliente.
 - Tiempo de concesión largo (mejor estático): Trazabilidad del cliente.
 - Numerar las interfaces WAN del CPE con GUAs (Global Unicast Address).


Mas Info: [RIPE BCOP-690](#)

Dimensionamiento de las subredes

- Asignar prefijo con tamaño necesario para el número de clientes a servir.
- Procurar que se pueda hacer agregación de rutas. Tabla rutas óptima.
- Reservar uno o dos prefijos contiguos para posibles necesidades.
- Usar un /64 para direccionar los Loopbacks.
- Reservar otro /64 para servicios (DNS, WWW, etc.).

Mas Info: [RIPE BCOP-690](#)


Subredes en IPv6



Mas info: [RFC5952](https://tools.ietf.org/html/rfc5952)

Escenario

- Anunciamos el prefijo asignado por RIPE 2001::/16 desde el router Core
- Numeramos todos los interfaces con GUAs para que funcione traceroute y ping
- Asignamos subred a cada Antena, según número de clientes y tamaño del prefijo a delegar a cada uno.
- En nuestro escenario asignamos un /48 a cada antena y delegamos un /60 a cada cliente final (16 subnets /64)




Direccionamiento

Prefijo	Uso	Localización
2000::100/127	Infraestructura	IX BGP Carrier
2000::103/127	Infraestructura	IX BGP Core
2000:a::/112	Infraestructura	OSPFv3 CPD-Antenas
2001:1::/48	Clientes	Antena 1
2001:2::/48	Clientes	Antena 2
2001:3::/48	Clientes	Antena 3
2001:100::/48	Infraestructura	Direccionamiento Interno

- Reservamos a partir del 2001:0::/48 para ampliaciones clientes

Escenario


- Peer BGP IPv6 entre Carrier y Core
- Tunnel 6to4 entre el core y los CPD's
- OSPFv3 con las antenas
- PPPoE para las conexiones de los clientes
- Y asignación mediante SLAAC en la LAN



Escenario

ATENCION!!


El tunel 6to4 estaría en
una red IPv4 pero para
simplificar el laboratorio
se ha conectado
directamente



BGP IPv6

Aplicaremos una
configuración BGP
entre dos cores con
IPv6





BGP Carrier



En este lab configuraremos unos BGP carrier para ver cómo se propagan los prefijos IPv6 y poder llegar a los equipos finales de cliente desde éste, a través de nuestra infraestructura.

BGP 1

The screenshot shows the 'IPv6 Address List' window with 5 items. The first two items are selected: '2000::100/127' on interface 'ether3' and '2000::103/127' on interface 'ether1'. Below the list, two configuration windows are open for these addresses. The first window is for '2000::100/127' with interface 'ether3'. The second window is for '2000::103/127' with interface 'ether1'. Both windows have checkboxes for 'EUI64', 'Advertise', and 'No DAD', all of which are unchecked. The status at the bottom is 'enabled' and 'Global'.

	Address	From Pool	Interface	Advertise
G	2000::100/127		ether3	no
G	2000::103/127		ether1	no
DL	fe80::5054:ff:fe45:83a...		ether3	no
DL	fe80::5054:ff:fe5a:bfa...		ether1	no
DL	fe80::5054:ff:fe8a:9cc...		ether2	no

5 items (1 selected)

IPv6 Address <2000::100/127>

Address: 2000::100/127

From Pool:

Interface: ether3

☐ EUI64

☐ Advertise

☐ No DAD

enabled Global

IPv6 Address <2000::103/127>

Address: 2000::103/127

From Pool:

Interface: ether1

☐ EUI64

☐ Advertise

☐ No DAD

enabled Global

```

/ipv6 address
add address=2000::100/127
advertise=no interface=ether3
/ipv6 address
add address=2000::103/127
advertise=no interface=ether1
  
```

BGP 2

The screenshot shows the 'IPv6 Address List' window with 3 items. The first item is selected: '2000::101/127' on interface 'ether1'. Below the list, a configuration window is open for this address with interface 'ether1'. The window has checkboxes for 'EUI64', 'Advertise', and 'No DAD', all of which are unchecked. The status at the bottom is 'enabled' and 'Global'.

	Address	From Pool	Interface	Advertise
G	2000::101/127		ether1	no
DL	fe80::5054:ff:feac:ecd...		ether2	no
DL	fe80::5054:ff:fee8:2b4...		ether1	no

3 items (1 selected)

IPv6 Address <2000::101/127>

Address: 2000::101/127

From Pool:

Interface: ether1

☐ EUI64

☐ Advertise

☐ No DAD

enabled Global

```

/ipv6 address
add address=2000::101/127
advertise=no interface=ether1
  
```

BGP 1

BGP

Instances VRFs Peers Networks Aggregates VPN4 Routes Adv

Name	AS	Router ID	Out Filter	Confede...
default	65001	10.0.0.133		

BGP Instance <default>

Name: default OK

AS: 65001 Cancel

Router ID: 10.0.0.133 Apply

☐ Redistribute Connected Disable

☐ Redistribute Static Comment

☐ Redistribute RIP Copy

☐ Redistribute OSPF Remove

☒ Redistribute Other BGP

Out Filter: []

Confederation: []

Confederation Peers: []

Cluster ID: []

Routing Table: []

☒ Client To Client Reflection

☐ Ignore AS Path Length

enabled

```
/routing bgp instance
set default as=65001
redistribute-other-bgp=yes
router-id=10.0.0.133
```

BGP 2

BGP

Instances VRFs Peers Networks Aggregates VPN4 Routes Adv

Name	AS	Router ID	Out Filter	Confede...
default	65002	10.0.0.134		

BGP Instance <default>

Name: default OK

AS: 65002 Cancel

Router ID: 10.0.0.134 Apply

☐ Redistribute Connected Disable

☐ Redistribute Static Comment

☐ Redistribute RIP Copy

☐ Redistribute OSPF Remove

☒ Redistribute Other BGP

Out Filter: []

Confederation: []

Confederation Peers: []

Cluster ID: []

Routing Table: []


☒ Client To Client Reflection

☐ Ignore AS Path Length

enabled

```
/routing bgp instance
set default as=65002
router-id=10.0.0.134
```

BGP 1




Name	Instance	Remote Address	Remote AS	M...	R...	TTL	Remote ID
peer1	default	2000::101	65002	no	no	d...	10.0.0.134
peer2	default	2000::102	65000	no	no	d...	10.0.0.132

2 items

```
/routing bgp peer
add address-families=ipv6
name=peer1 remote-address=2000::101
remote-as=65002
```

```
/routing bgp peer
add address-families=ipv6
name=peer2 remote-address=2000::102
remote-as=65000
```

BGP 2




Name	Instance	Remote Address	Remote AS	M...	R...	TTL	Remote ID
peer1	default	2000::100	65001	no	no	d...	10.0.0.133

1 item

```
/routing bgp peer
add address-families=ipv6
name=peer1 remote-address=2000::100
remote-as=65001
```

Core

Configuramos las direcciones IPv6 que se publicarán en nuestro core y como las verán los carriers, para el Core lo configuramos de la siguiente forma, primero las IPv6 para la sesión BGP



```
/ipv6 address  
add address=2000::102/127  
advertise=no  
interface=ether1
```

Core

Definimos nuestro AS y el ID del router

BGP

Instances VRFs Peers Networks Aggregates VPN4 Routes Ad...

Name AS Router ID Out Filter Confede...

default	65000	10.0.0.132		
---------	-------	------------	--	--

BGP Instance <default>

Name: default

AS: 65000

Router ID: 10.0.0.132

☐ Redistribute Connected

☐ Redistribute Static

☐ Redistribute RIP

☐ Redistribute OSPF

☐ Redistribute Other BGP

Out Filter:

Confederation:

Confederation Peers:

Cluster ID:

Routing Table:

☒ Client To Client Reflection

☐ Ignore AS Path Length

enabled

```
/routing bgp instance  
set default as=65000  
router-id=10.0.0.132
```


Core

The screenshot shows a BGP configuration window with a table of peers and a detailed configuration for 'peer1'.

Name	Instance	Remote Address	Remote AS	M...	R...	TTL	Remote ID	Uptime
peer1	default	2000::103	65001	no	no	d...	10.0.0.133	20:16:10

BGP Peer <peer1>

General | Advanced | Status

Name: peer1
Instance: default
Remote Address: 2000::103
Remote Port:
Remote AS: 65001
TCP MD5 Key:
Nexthop Choice: default
☐ Multihop
☐ Route Reflect
Hold Time: 180 s
Keepalive Time:
TTL: default
Max Prefix Limit:
Max Prefix Restart Time:
In Filter:
Out Filter:
AllowAS In:
☐ Remove Private AS
☐ AS Override
Default Originate: never
☐ Passive
☐ Use BFD


enabled established

Configuramos el peer del carrier

```
/routing bgp peer  
add address-families=ipv6  
name=peer1  
remote-address=2000::103  
remote-as=65001
```

Core

Añadimos el rango de red que queremos publicar, y comprobamos que esté marcado ipv6



```
/routing bgp network
```

```
add network=2001::/16 synchronize=no
```

Core

Añadimos la ruta para nuestro rango de ip's esté apuntando a una salida.

```
/ipv6 route  
add distance=1  
dst-address=2001::/16  
gateway=fe80::1%ether1
```

IPv6 Route List

	Dst. Address	Gateway
AS	2001::/16	fe80::1%ether1 reachable
DAC	2000::102/127	ether1 reachable

IPv6 Route <2001::/16>

General Attributes

Dst. Address: 2001::/16

Gateway: fe80::1%ether1  reachable

Check Gateway:

Type: unicast


Distance: 1

Scope: 30

Target Scope: 10

Core

Comprobamos que éste el peering establecido y comprobaremos sí se propagan hacia el carrier




The screenshot shows a window titled "BGP" with several tabs: Instances, VRFs, Peers, Networks, Aggregates, VPN4 Routes, and Advertisements. The "Peers" tab is selected. Below the tabs is a toolbar with icons for adding, removing, checking, and deleting peers, as well as buttons for Refresh, Refresh All, Resend, and Resend All. A search bar labeled "Find" is on the right. The main area contains a table with the following data:

Name	Instance	Remote Address	Remote AS	M...	R...	TTL	Remote ID	Uptime	Prefix Co...	State
peer1	default	2000::103	65001	no	no	d...	10.0.0.133	20:49:32		established

At the bottom left, it says "1 item".

BGP 1


BGP 2



The screenshot shows a BGP configuration window with the 'Peers' tab selected. It displays two configured peers, peer1 and peer2, both in an established state.

Name	Instance	Remote Address	Remote AS	M...	R...	TTL	Remote ID	Uptime	Prefix Co...	State
peer1	default	2000::101	65002	no	no	d...	10.0.0.134	18:54:29		established
peer2	default	2000::102	65000	no	no	d...	10.0.0.132	20:48:31	1	established

2 items




The screenshot shows the same BGP configuration window, but now only peer1 is listed, indicating peer2 has been removed.

Name	Instance	Remote Address	Remote AS	M...	R...	TTL	Remote ID	Uptime	Prefix Co...	State
peer1	default	2000::100	65001	no	no	d...	10.0.0.133	18:54:15	1	established

1 item

Tunnel 6to4

¿Por qué no usar
una red existente
IPv4 para enlazar
los CPD's con
nuestro Core?



Tunnel 6to4



Conectaremos los CPD's con el Core mediante una red WAN IPv4 para hacer la distribución de los CPD's más sencilla al usar las redes actuales.

En este laboratorio no se representará la red IPv4 y se conectarán directamente los interfaces simulando las ip's como direcciones públicas

Core

Address List

Address	Network	Interface
172.16.100.1/30	172.16.100.0	ether2

2 items

Address <172.16.100.1/30>

Address: 172.16.100.1/30

Network: 172.16.100.0

Interface: ether2

OK Cancel Apply Disable Comment Copy Remove

enabled

```
/ip address
add address=172.16.100.1/30
    interface=ether2
    network=172.16.100.0
```

CPD

Address List

Address	Network	Interface
172.16.100.2/30	172.16.100.0	ether1

2 items

Address <172.16.100.2/30>

Address: 172.16.100.2/30

Network: 172.16.100.0

Interface: ether1

OK Cancel Apply Disable Comment Copy Remove

enabled

```
/ip address
add address=172.16.100.2/30
    interface=ether1
    network=172.16.100.0
```

Core

The screenshot shows the configuration window for the '6to4' interface on the 'Core' device. The interface is a 6to4 Tunnel with MTU 1280 and L2 MTU 65535. The local address is 172.16.100.1 and the remote address is 172.16.100.2. The DSCP is set to 10, and the 'Clamp TCP MSS' checkbox is checked.

Name	Type	Actual MTU	L2 MTU	Tx
R 6to4	6to4 Tunnel	1280	65535	
;;; Sw4-BGP				
R ether1	Ethernet	1500		
;;; Sw5.CPD				
R ether2	Ethernet	1500		

Interface <6to4>

General Status Traffic

Name: 6to4

Type: 6to4 Tunnel

MTU: 1280

Actual MTU: 1280

L2 MTU: 65535

Local Address: 172.16.100.1

Remote Address: 172.16.100.2

IPsec Secret:

Keepalive:

DSCP: 10

Dont Fragment: no

☒ Clamp TCP MSS

```

/interface 6to4
 add local-address=172.16.100.1
 mtu=1280 name=6to4 remote-address=
 172.16.100.2
  
```

CPD

The screenshot shows the configuration window for the '6to4' interface on the 'CPD' device. The interface is a 6to4 Tunnel with MTU 1280 and L2 MTU 65535. The local address is 172.16.100.2 and the remote address is 172.16.100.1. The DSCP is set to 10, and the 'Clamp TCP MSS' checkbox is checked.

Name	Type	Actual MTU	L2 MTU	Tx
R 6to4	6to4 Tunnel	1280	65535	
;;; Sw5-Core				
R ether1	Ethernet	1500		
;;; Sw7-Ant				
R ether2	Ethernet	1500		

Interface <6to4>

General Status Traffic

Name: 6to4

Type: 6to4 Tunnel

MTU: 1280

Actual MTU: 1280

L2 MTU: 65535

Local Address: 172.16.100.2

Remote Address: 172.16.100.1

IPsec Secret:

Keepalive:

DSCP: 10

Dont Fragment: no

☒ Clamp TCP MSS

```

/interface 6to4
 add local-address=172.16.100.2
 mtu=1280 name=6to4 remote-address=
 172.16.100.1
  
```

IPv6

Para obtener la dirección IPv6 correspondiente al túnel 6to4 hay que hacer una conversión entre la ipv4 y el pool a usar. En la siguiente dirección se puede obtener:

<http://silmor.de/ipaddrcalc.html#ip46>

The screenshot shows a web browser window with the address bar displaying "No es seguro | silmor.de/ipaddrcalc.html#ip46". The page content includes:

- IPv4 to IPv6 Transitional**
 - [hide]
 - IPv4 in IPv6 Addresses**

In some configurations IPv4 addresses can be written or used in IPv6 notation or they become part of an IPv6 address. This form allows you to convert from IPv4 to IPv6 and back. Depending on your application you may have to shift the IPv6 segments.

IPv4 Address: Please use dotted decimal notation.

IPv6 Address: Please use hexadecimal notation with the relevant 32 bits to the far right.
 - 6to4 and 6RD Network Prefix**

6to4 and 6rd are transitional mechanisms that will be used until native IPv6 is universally available. With both mechanisms you can assign an IPv6 prefix to an entire network based on the IPv4 address of the gateway. Both use the 6in4 encapsulation to transport IPv6 packets inside IPv4 packets between the border gateway of the local network and the gateway servers outside.

6to4 is a public service, everybody can configure a gateway to use it - no subscription is necessary, since gateways will always know where to route responses based on the prefix. All 6to4 prefixes are in the 2002::/16 network and are /48 bits long (16bits for 2002::/16 and 32bits from the IPv4 address of the gateway). Unfortunately this service has become quite unreliable since public gateway servers seem to be unable to scale with the demand for prefixes.

6rd is the provider internal equivalent of 6to4. The provider establishes a gateway (or cluster of gateways) in its internal network and customer gateways are configured to use this gateway. The provider side prefix can be considerable longer than with 6to4 (/32 is normal), but it is also quite common to use only some bits of the IPv4 address - normally IPv4 addresses for customers are either assigned from a limited pool of public addresses (a /16 being the norm) or from one of the "private" pools (e.g. 10.0.0.0/8), so the leftmost bits of every customer IP will be identical and can be ignored. For this mechanism to work you have to be a subscriber of an ISP that provides this mechanism to its customers. The values that go into this calculation may or may not have some resemblance to what you can find out using the whois service, but the provider is free to use sub-nets, so you will need information directly from the provider.
- Form fields:
 - Provider prefix IPv6: /16 Use 2002::/16 for 6to4 and whatever you ISP gave you for 6rd
 - Customer IPv4: IP: using bits; use your public IPv4 address (PPP: your own, not the peer address) and the bits value from the ISP or 32 for 6to4
 - IPv6 Customer prefix: /46
- Buttons:
 - ISP -> Customer: convert the ISP prefix and public IPv4 to a IPv6 customer prefix
 - Customer -> ISP: convert the IPv6 customer prefix to ISP prefix and public IPv4 (the provider prefix length will be used as is)

Core

IPv6 Address List

	Address	From Pool	Interface	Advertise
G	2002:b041:9004::/3		6to4	no
DL	fe80::3:ac10:6401/64		6to4	no
DL	fe80::5054:ff:fe4e:c03...		ether1	no
G	2000::102/127		ether1	no
DL	fe80::5054:ff:fe2a:af8...		ether2	no
DL	fe80::5054:ff:fe6c:df6c...		ether3	no

IPv6 Address <2002:b041:9004::/3>

Address: 2002:b041:9004::/3

From Pool:

Interface: 6to4

☐ EUI64

☐ Advertise

☐ No DAD

OK Cancel Apply Disable Comment Copy Remove

enabled Global

```
/ipv6 address
add address=2002:b041:9004::/3
advertise=no interface=6to4
```

CPD

IPv6 Address List

	Address	From Pool	Interface	Advertise
G	2002:b041:9008::/3		6to4	no
DL	fe80::3:ac10:6402/64		6to4	no
DL	fe80::5054:ff:fe7c:55d...		ether2	no
DL	fe80::5054:ff:febd:fdcb...		ether1	no

IPv6 Address <2002:b041:9008::/3>

Address: 2002:b041:9008::/3

From Pool:

Interface: 6to4

☐ EUI64

☐ Advertise

☐ No DAD

OK Cancel Apply Disable Comment Copy Remove

enabled Global

```
/ipv6 address
add address=2002:b041:9008::/3
advertise=no interface=6to4
```

Core

CPD

Ping

General Advanced

Ping To: 2002:b041:9008::

Interface: ▼

☐ ARP Ping

Packet Count: ▼

Timeout: 1000 ms

Start Stop Close New Window

Seq # ^	Host	Time	Reply Size	TTL	Status	▼
0	2002:b041:9008::	0ms	50	64	echo reply	
1	2002:b041:9008::	0ms	50	64	echo reply	
2	2002:b041:9008::	0ms	50	64	echo reply	
3	2002:b041:9008::	0ms	50	64	echo reply	
4	2002:b041:9008::	0ms	50	64	echo reply	
5	2002:b041:9008::	1ms	50	64	echo reply	

Ping

General Advanced

Ping To: 2002:b041:9004::

Interface: ▼

☐ ARP Ping

Packet Count: ▼

Timeout: 1000 ms

Start Stop Close New Window

Seq # ^	Host	Time	Reply Size	TTL	Status	▼
0	2002:b041:9004::	1ms	50	64	echo reply	
1	2002:b041:9004::	1ms	50	64	echo reply	
2	2002:b041:9004::	0ms	50	64	echo reply	
3	2002:b041:9004::	1ms	50	64	echo reply	
4	2002:b041:9004::	0ms	50	64	echo reply	
5	2002:b041:9004::	1ms	50	64	echo reply	
6	2002:b041:9004::	0ms	50	64	echo reply	

Core

Ahora que tenemos por donde van a ir las direcciones ip, las apuntamos en esa ruta.

```
/ipv6 route  
add distance=1  
dst-address=2001::/16  
gateway=6to4
```

The screenshot shows a network configuration window titled "IPv6 Route List". It contains a table of routes and a detailed configuration panel for a selected route.

	Dst. Address	Gateway
DAC	2000::102/127	ether1 reachable
AS	2001::/16	6to4 reachable
DAC	2000::/3	6to4 reachable


Below the table, the selected route "IPv6 Route <2001::/16>" is shown in a configuration panel with the following fields:

- General tab: Dst. Address: 2001::/16
- Attributes tab: Gateway: 6to4, reachable
- Check Gateway: (empty)
- Type: unicast
- Distance: 1
- Scope: 30
- Target Scope: 10

OSPFv3

Usar OSPFv3 para
enrutar las direcciones
IPv6 hasta los
concentradores PPPoE





OSPFv3



La versión 3 es mucho más sencillo de configurar, en pocos pasos está funcionando.

CPD

The screenshot shows the 'IPv6 Address List' window in CPD. The table lists several addresses, with '2000::a:1/112' selected. Below the table, the configuration dialog for this address is shown, with 'Interface' set to 'ether2'.

	Address	From Pool	Interface	Advertise
G	2002:b041:9008::/3		6to4	no
DL	fe80::3:ac10:6402/64		6to4	no
DL	fe80::5054:ff:fe7c:55d...		ether2	no
DL	fe80::5054:ff:febd:fdcb...		ether1	no
G	2000::a:1/112		ether2	no
DL	fe80::5054:ff:fef2:9a5/64		ether3	no

6 items (1 selected)

IPv6 Address <2000::a:1/112>

Address: 2000::a:1/112

From Pool: [dropdown]

Interface: ether2

☐ EUI64
☐ Advertise
☐ No DAD

Buttons: OK, Cancel, Apply, Disable, Comment, Copy, Remove

enabled Global

```
/ipv6 address
add address=2000::a:1/112
advertise=no interface=ether2
```

Antena

The screenshot shows the 'IPv6 Address List' window in Antena. The table lists several addresses, with '2000::a:2/112' selected. Below the table, the configuration dialog for this address is shown, with 'Interface' set to 'ether1'.

	Address	From Pool	Interface	Advertise
G	2000::a:2/112		ether1	no
DL	fe80::5054:ff:fe40:6f0...		ether1	no
DL	fe80::5054:ff:fe6a:55a...		ether3	no
DL	fe80::5054:ff:fee8:b26...		ether2	no

IPv6 Address <2000::a:2/112>

Address: 2000::a:2/112

From Pool: [dropdown]

Interface: ether1

☐ EUI64
☐ Advertise
☐ No DAD

Buttons: OK, Cancel, Apply, Disable, Comment, Copy, Remove

enabled Global

```
/ipv6 address
add address=2000::a:2/112
advertise=no interface=ether1
```

CPD

The screenshot shows the OSPFv3 configuration window for interface ether2. The window has tabs for Interfaces, Instances, Areas, Area Ranges, Virtual Links, Neighbors, and NBMA. The Interfaces tab is active, showing a table with columns Area, Interface, Cost, Priority, and Network Type. The table contains one entry: backbone, ether2, 10, 1, default. Below the table, the configuration details for ether2 are shown. The Area is set to backbone, Interface to ether2, Cost to 10, Priority to 1, and Network Type to default. The Instance ID is 0. There are checkboxes for Passive and Use BFD, both of which are unchecked. The Retransmit Interval is 5 s, Transmit Delay is 1 s, Hello Interval is 10 s, and Router Dead Interval is 40 s. At the bottom, there are buttons for OK, Cancel, Apply, Disable, Copy, and Remove. The status bar at the bottom shows 'enabled', 'passive', 'inactive', and 'State: designated ro...'.

```
/routing ospf-v3 interface
add area=backbone interface=ether2
```

Antena

The screenshot shows the OSPFv3 configuration window for interface ether1. The window has tabs for Interfaces, Instances, Areas, Area Ranges, Virtual Links, Neighbors, and NBMA. The Interfaces tab is active, showing a table with columns Area, Interface, Cost, Priority, and Network Type. The table contains one entry: backbone, ether1, 10, 1, default. Below the table, the configuration details for ether1 are shown. The Area is set to backbone, Interface to ether1, Cost to 10, Priority to 1, and Network Type to default. The Instance ID is 0. There are checkboxes for Passive and Use BFD, both of which are unchecked. The Retransmit Interval is 5 s, Transmit Delay is 1 s, Hello Interval is 10 s, and Router Dead Interval is 40 s. At the bottom, there are buttons for OK, Cancel, Apply, Disable, Copy, and Remove. The status bar at the bottom shows 'enabled', 'passive', 'inactive', and 'State: backup'.

```
/routing ospf-v3 interface
add area=backbone interface=ether1
```


CPD

OSPFv3

Interfaces Instances Areas Area Ranges Virtual Links Neighbors NBMA

+ - ✓ ✗ 📁 🔍

Name	Router ID	Running
* default	10.0.10.1	yes

OSPFv3 Instance <default>

General Metrics Status

Name: default

Router ID: 10.0.10.1

Redistribute Default Route: never

Redistribute Connected Routes: no

Redistribute Static Routes: as type 1

Redistribute RIP Routes: no

Redistribute BGP Routes: no

Redistribute Other OSPF Routes: no

OK Cancel Apply Disable Comment Copy Remove

enabled default

```
/routing ospf-v3 instance
set redistribute-static=as-type-1
router-id=10.0.10.1
```

Antena

OSPFv3

Interfaces Instances Areas Area Ranges Virtual Links Neighbors NBMA

+ - ✓ ✗ 📁 🔍

Name	Router ID	Running
* default	10.0.10.2	yes

OSPFv3 Instance <default>

General Metrics Status

Name: default

Router ID: 10.0.10.2

Redistribute Default Route: never

Redistribute Connected Routes: no

Redistribute Static Routes: as type 1

Redistribute RIP Routes: no

Redistribute BGP Routes: no

Redistribute Other OSPF Routes: no

OK Cancel Apply Disable Comment Copy Remove

enabled default

```
/routing ospf-v3 instance
set redistribute-static=as-type-1
router-id=10.0.10.2
/ipv6 route add dst-address=2001:1::/48
type=unreachable
```


CPD


OSPFv3				
Virtual Links Neighbors NBMA Neighbors LSA Routes AS Bo				
				
Instance ▲	Router ID	Address	Interface	
 default	10.0.10.2	fe80::5054:f...	ether2	
 default	10.0.10.3	fe80::5054:f...	ether2	
2 items				

Antena

OSPFv3				
Virtual Links Neighbors NBMA Neighbors LSA Routes AS B				
				
Instance ▲	Router ID	Address	Interface	
 default	10.0.10.1	fe80::5054:f...	ether1	
 default	10.0.10.3	fe80::5054:f...	ether1	
2 items				

PPPoE

No vamos a
inventar la rueda




PPPoE

Conectar los clientes a las antenas mediante el PPPoE, una forma eficiente y asentada en el mercado

Antenas

Asignamos el pool de prefijos IPv6 que va a delegar esa antena hacia los CPEs

```
/ipv6 pool  
add name=IPv6_PPPE  
prefix=2001:1::/48  
prefix-length=60
```



Antenas

Configuración sencilla de un servidor
PPPoE

```
/ppp profile
```

```
add dhcpv6-pd-pool=IPv6_PPPoE
```

```
name=IPv6 only-one=no
```

```
use-ipv6=required
```

```
/ppp secret
```

```
add name=pppv6 password=ipv6
```

```
profile=IPv6 service=pppoe
```

```
/interface pppoe-server server
```

```
add default-profile=IPv6
```

```
disabled=no interface=ether2
```

```
service-name=PPPoEv6
```

The screenshot displays three configuration windows in Mikrotik WinBox:

- PPP Secret <pppv6>**: Name: pppv6, Password: ipv6, Service: pppoe, Profile: IPv6. Status: enabled.
- PPP Profile <IPv6>**: Name: IPv6, Local Address: (empty), Remote Address: (empty), Remote IPv6 Prefix Pool: (empty), DHCPv6 PD Pool: IPv6_PPPoE. Status: enabled.
- PPPoE Service <PPPoEv6>**: Service Name: PPPoEv6, Interface: ether2, Max MTU: (empty), Max MRU: (empty), MRRU: (empty), Keepalive Timeout: 10, Default Profile: IPv6, One Session Per Host: unchecked, Max Sessions: (empty), PADO Delay: (empty) ms, Authentication: mschap2, mschap1, chap, pap (all checked). Status: enabled.

CPE

Cliente de PPPoE en el CPE que conecta con la antena

```
/interface pppoe-client
add add-default-route=yes
disabled=no interface=ether1
name=pppoe-out1
password=ipv6 user=pppv6
```

The screenshot shows the Mikrotik WinBox interface for configuring a PPP client. The main window is titled 'PPP' and has tabs for 'Interface', 'PPPoE Servers', 'Secrets', 'Profiles', 'Active Connections', and 'L2TP Secrets'. Below these are buttons for '+', '-', 'check', 'X', 'file', 'filter', 'PPP Scanner', 'PPTP Server', and 'SSTP Server'. A table lists the configured interfaces:

Name	Type	Actual MTU	L2 MTU	Tx
pppoe-out1	PPPoE Client			

The 'Interface <pppoe-out1>' dialog box is open, showing the 'Dial Out' tab. The fields are as follows:


- Service: (empty dropdown)
- AC Name: (empty dropdown)
- User: pppv6
- Password: ipv6
- Profile: default
- Keepalive Timeout: 10
- ☐ Dial On Demand
- ☐ Use Peer DNS
- ☒ Add Default Route
- Default Route Distance: 1
- Allow: ☒ mschap2, ☒ mschap1, ☒ chap, ☒ pap

Buttons on the right include OK, Cancel, Apply, Disable, Comment, Copy, Remove, Torch, and PPPoE Scan.

CPE


Ponemos un DHCPv6 Client para obtener el prefijo para el cliente.

```
/ipv6 dhcp-client
add add-default-route=yes
interface=pppoe-out1
pool-name=ip
pool-prefix-length=60
request=prefix
```



SLAAC

Asignación de
direcciones IPv6 a
los clientes finales



SLAAC




Cómo distribuir el prefijo asignado al CPE entre los equipos de los clientes

CPE

Se configura la interfaz para repartir el prefijo asignado

```
/ipv6 address
add from-pool=ip
interface=ether3
```



CPE

La IPv6 asignada al equipo en la red del CPE

The screenshot displays the Windows Network and Sharing Center. The main window shows the network status for 'Red 2' (Red doméstica), which is connected to the Internet. The status bar indicates 'Sin acceso a Internet' (No Internet access), 'Unido' (Connected), and 'Conexión de área local' (Local area connection).

Below the main window, two smaller windows are open:

- Estado de Conexión de área local** (Local Area Connection Status): This window shows the connection status for the local area connection. It indicates that IPv4 connectivity is 'Sin acceso a la red' (No network access) and IPv6 connectivity is 'Sin acceso a Internet' (No Internet access). The media state is 'Habilitado' (Enabled), and the speed is '1,0 Gbps'. The activity section shows 'Enviados' (Sent) bytes as 81.343 and 'Recibidos' (Received) bytes as 50.210.
- Detalles de la conexión de red** (Network Connection Details): This window provides a detailed view of the network connection. It lists various properties and their values, including the physical address, IP addresses, and DNS settings.

Propiedad	Valor
Sufijo DNS específico p...	
Descripción	Conexión de red Intel(R) PRO/1000 MT
Dirección física	52-54-00-2D-2A-79
Habilitado para DHCP	Sí
Dirección IPv4 de config...	169.254.91.156
Máscara de subred IPv4	255.255.0.0
Puerta de enlace predet...	
Servidor DNS IPv4	
Servidor WINS IPv4	
Habilitado para NetBios ...	Sí
Dirección IPv6	2001:2::4465:6810:258e:5b9c
Dirección IPv6 temporal	2001:2::10f9:7cb1:942d:9e21
Vínculo: dirección IPv6 l...	fe80::4465:6810:258e:5b9c%11
Puerta de enlace predet...	fe80::5054ff:feab:7853%11
Servidores DNS IPv6	fec0:0:0:ffff::1%1 fec0:0:0:ffff::2%1 fec0:0:0:ffff::3%1

CPE

Se comprueba que se llegue
desde los Carriers hasta los
equipos del cliente final

BGP

Instances VRFs Peers Networks Aggregates VPN4 Routes Advertisements

+ - ✓ ✗ [Icon] [Icon] Refresh Refresh All Resend Resend All Find

Name	Instance	Remote Address	Remote AS	M...	R...	TTL	Remote ID	Uptime
peer1	default	2000::101	65002	no	no	d...	10.0.0.134	20:26:
peer2	default	2000::102	65000	no	no	d...	10.0.0.132	22:20:

2 items

Ping (Running)

General Advanced

Ping To: 2001:2::4465:6810:258e:5b9c

Interface: [Dropdown]

☐ ARP Ping

Packet Count: [Dropdown]

Timeout: 1000 ms

Start Stop Close New Window

Seq #	Host	Time	Reply Size	TTL	Status
142	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
143	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
144	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
145	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
146	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
147	2001:2::4465:6810:258e:5...	3ms	50	126	echo reply
148	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
149	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
150	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
151	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
152	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
153	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply
154	2001:2::4465:6810:258e:5...	2ms	50	126	echo reply

155 items 155 of 155 packet... 0% packet loss Min: 2 ms Avg: 2 ms Max: 9 ms

¿Preguntas?

IPv6

MUM Valencia 2018

*Mikro***Tik**

¡Gracias!

IPv6

MUM Valencia 2018

jose.roman@fibercli.com

ccardenas@aunnait.es



AUNNA

INNOVACIÓN TECNOLÓGICA EN
COMUNICACIONES Y REDES