

# BGP Advanced Concepts


Announcing IPv4 prefixes via IPv6-only networks

Wolfgang Tremmel  
[wolfgang.tremmel@de-cix.net](mailto:wolfgang.tremmel@de-cix.net)



# ***BGP Topics Overview***

- 01 - Prefixes and AS numbers
- 02 - BGP Introduction
- 03a - Setting up iBGP
- 03b - Setting up eBGP
- 04 - Becoming multi-homed
- 05 - BGP Best Path Selection
- 06 - BGP Communities
- 07 - BGP Traffic Engineering
- 08 - BGP Security



## **Advanced Concepts**

Announcing IPv4 prefixes via IPv6-only networks

# *Motivation*





# Announcing Prefixes

Standard BGP announcements

2001:db8:10::/64

198.51.100.0/24


What do we need the next-hop address for?

## BGP Announcement:

Prefix: 192.0.2.0/24

AS-PATH: 65551

Next-Hop: 198.51.100.1



AS65551

2001:db8:10::5

198.51.100.5

# Announcing Prefixes

## Standard BGP announcements

### BGP Announcement:

Prefix: 192.0.2.0/24

AS-PATH: 65001

Next-Hop: 198.51.100.1

AS65001

2001:db8:10::1  
198.51.100.1

2001:db8:10::/64  
198.51.100.0/24

2001:db8:10::3  
198.51.100.3

AS64499

2001:db8:10::5  
198.51.100.5

I want to send traffic  
to 192.0.2.1

Where do I send it to?

Ah! The next-hop  
is 198.51.100.1

So I do an ARP request  
for it



# Announcing Prefixes

## Standard BGP announcements

### BGP Announcement:

Prefix: 192.0.2.0/24  
AS-PATH: 65001  
Next-Hop: 198.51.100.1

AS65001

2001:db8:10::1  
198.51.100.1

2001:db8:10::/64  
198.51.100.0/24

2001:db8:10::3  
198.51.100.3

AS64499


2001:db8:10::5  
198.51.100.5

I want to send traffic to 192.0.2.1  
Where do I send it to?

Ah! The next-hop is 198.51.100.1

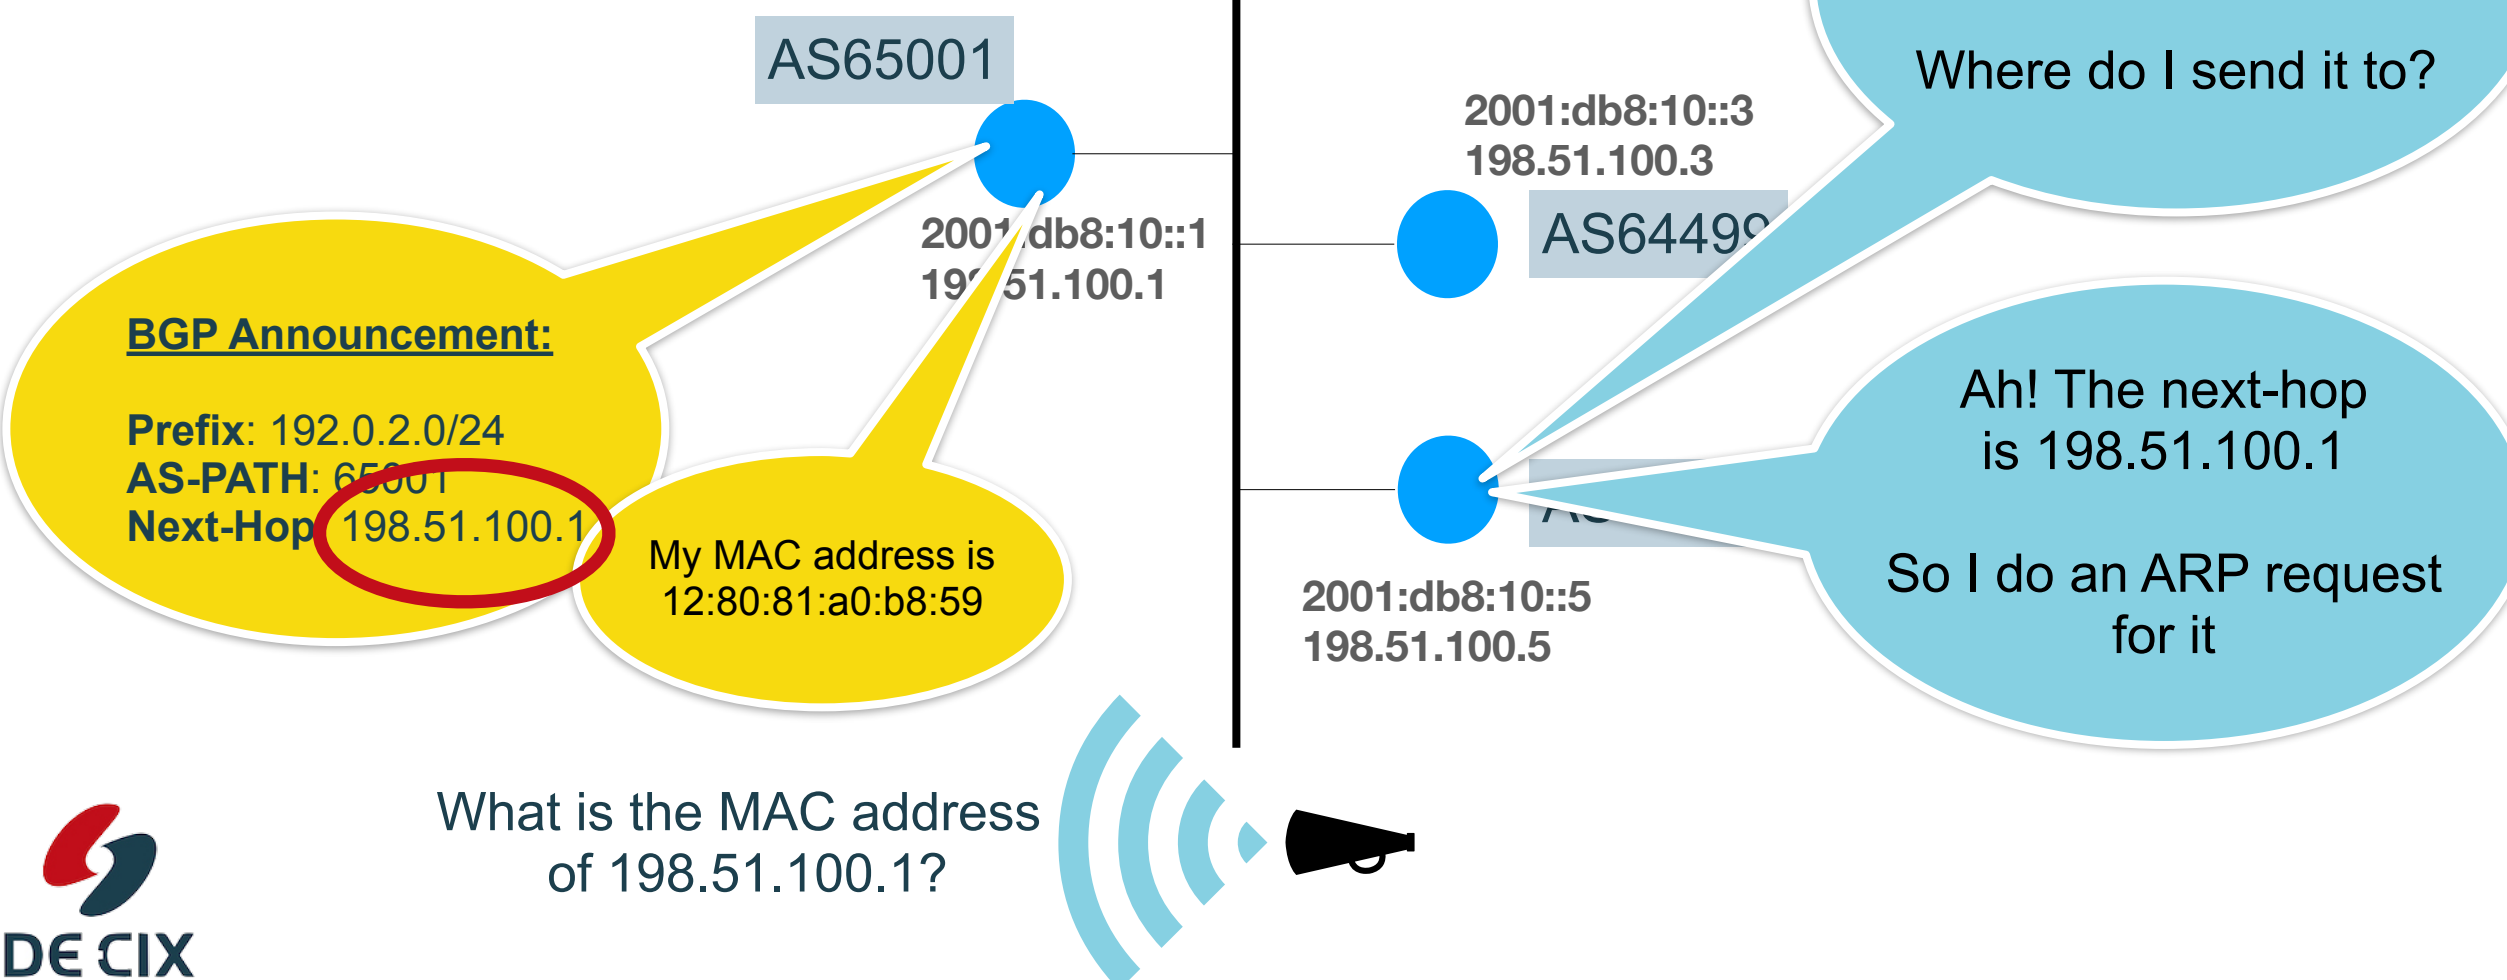
So I do an ARP request for it

What is the MAC address of 198.51.100.1?



# Announcing Prefixes

## Standard BGP announcements



*Do we really need  
IPv4 for that?*

**NO**



# Announcing Prefixes

## Standard BGP announcements

### BGP Announcement:

Prefix: 192.0.2.0/24

AS-PATH: 65001

Next-Hop: 2001:db8:10::1

AS65001

2001:db8:10::1  
198.51.100.1

2001:db8:10::/64  
198.51.100.0/24

2001:db8:10::3  
198.51.100.3

AS64499

2001:db8:10::5  
198.51.100.5

I want to send traffic  
to 192.0.2.1

Where do I send it to?

Ah! The next-hop  
is 2001:db8:10::1

So I do a **Neighbor  
Discovery**



# Announcing Prefixes

## Standard BGP announcements

### BGP Announcement:

Prefix: 192.0.2.0/24

AS-PATH: 65001

Next-Hop: 2001:db8:10::1

AS65001

2001:db8:10::1

2001:db8:10::/64

2001:db8:10::3

AS64499

2001:db8:10::5

I want to send traffic to 192.0.2.1  
Where do I send it to?

Ah! The next-hop is 2001:db8:10::1

So I do a **Neighbor Discovery**



# Announcing Prefixes

## Standard BGP announcements

### BGP Announcement:

Prefix: 192.0.2.0/24

AS-PATH: 65001

Next-Hop: 2001:db8:10::1

AS65001

2001:db8:10::1

2001:db8:10::/64

2001:db8:10::3

AS64499

2001:db8:10::5

I want to send traffic to 192.0.2.1  
Where do I send it to?

Ah! The next-hop is 2001:db8:10::1

So I do a **Neighbor Discovery**

Hello Multicast-Address FF02::1:ff00:1

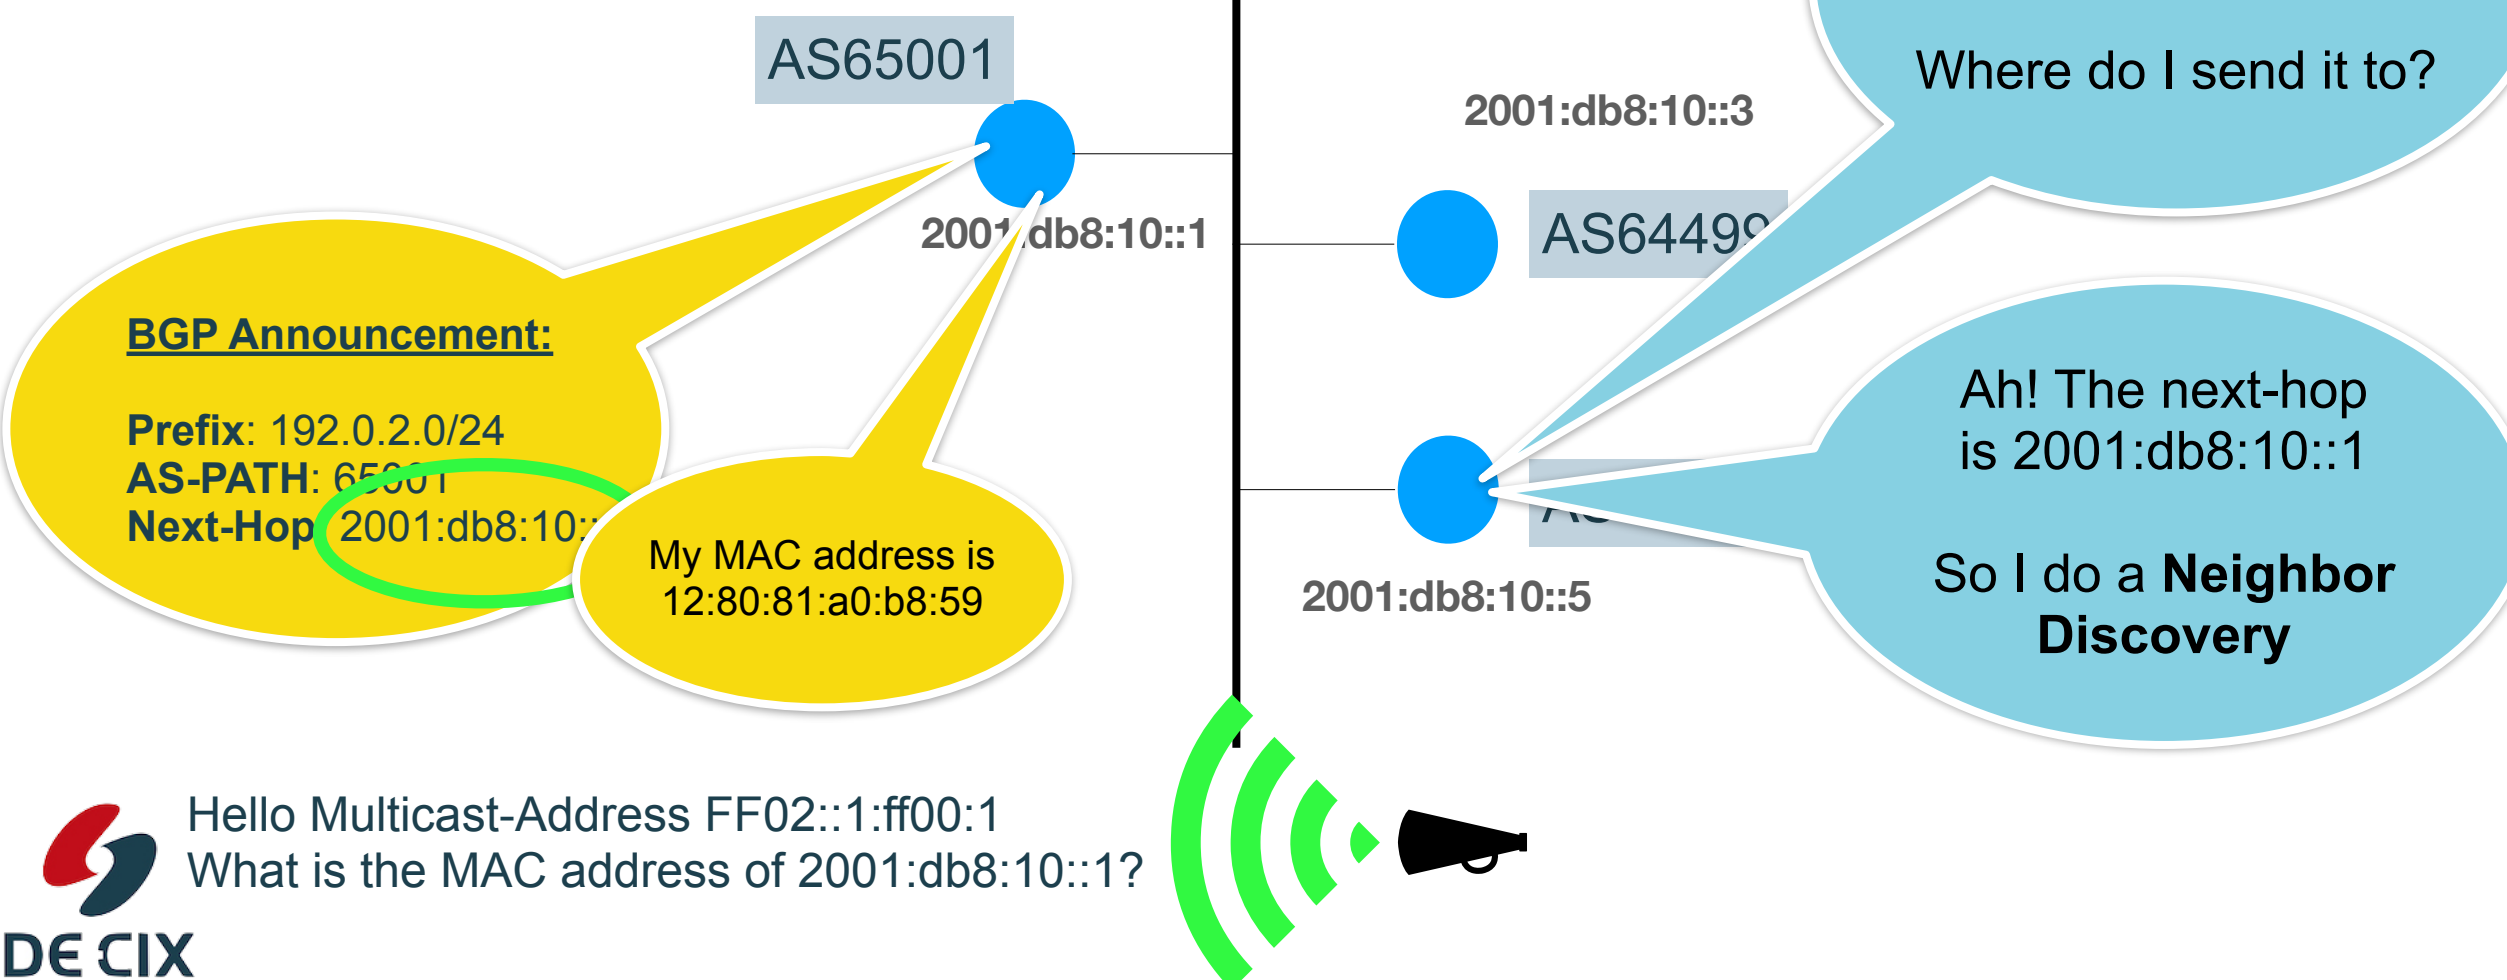
What is the MAC address of 2001:db8:10::1?



DE CIX

# Announcing Prefixes

## Standard BGP announcements



### BGP Announcement:

Prefix: 192.0.2.0/24  
AS-PATH: 65001  
Next-Hop: 2001:db8:10::1

My MAC address is  
12:80:81:a0:b8:59

I want to send traffic  
to 192.0.2.1  
Where do I send it to?

Ah! The next-hop  
is 2001:db8:10::1  
So I do a **Neighbor  
Discovery**

Hello Multicast-Address FF02::1:ff00:1  
What is the MAC address of 2001:db8:10::1?



*How to configure this?*






# BGP

→ In BGP, transport (how to connect to my neighbor) and announced prefixes are two different things


→ Command (here: Cisco)  
"neighbor xx:xx::xx" is how to connect to my neighbor



```
router bgp 196610
  neighbor 2a02:c50:db8:7:1:0:aaaa:0
    remote-as 64501
    address-family ipv4 unicast
      next-hop-self
      soft-reconfiguration inbound always
    !
  address-family ipv6 unicast
    soft-reconfiguration inbound always
  !
!
!
```

# BGP

→ Prefixes announced are configured in the "address-family" context



```
router bgp 196610
  neighbor 2a02:c50:db8:7:1:0:aaaa:0
  remote-as 64501
  address-family ipv4 unicast
    next-hop-self
    soft-reconfiguration inbound always
  !
  address-family ipv6 unicast
    soft-reconfiguration inbound always
  !
  !
  !
```

# BGP

→ So, "address-family ipv6 unicast" controls the announcement of IPv6 prefixes

```
router bgp 196610
  neighbor 2a02:c50:db8:7:1:0:aaaa:0
  remote-as 64501
  address-family ipv4 unicast
    next-hop-self
    soft-reconfiguration inbound always
  !
  address-family ipv6 unicast
    soft-reconfiguration inbound always
  !
  !
  !
```



DE CIX

# BGP

- So, "address-family ipv6 unicast" controls the announcement of IPv6 prefixes
- And "address-family ipv4 unicast" the announcement of IPv4 prefixes

```
router bgp 196610
  neighbor 2a02:c50:db8:7:1:0:aaaa:0
  remote-as 64501
  address-family ipv4 unicast
    next-hop-self
    soft-reconfiguration inbound always
  !
  address-family ipv6 unicast
    soft-reconfiguration inbound always
  !
  !
  !
```



DE CIX

# BGP

→ So, "address-family ipv6 unicast" controls the announcement of IPv6 prefixes

→ And "address-family ipv4 unicast" the announcement of IPv4 prefixes

→ here the "next-hop-self" is important so an IPv6 next-hop is set

```
router bgp 196610
  neighbor 2a02:c50:db8:7:1:0:aaaa:0
  remote-as 64501
  address-family ipv4 unicast
  next-hop-self
  soft-reconfiguration inbound always
!
  address-family ipv6 unicast
  soft-reconfiguration inbound always
!
!
!
```



DE CIX



# *Conclusion*

# Conclusion

- The next-hop address in BGP is needed for routers to find out the Layer 2 MAC address to send traffic to
- Usually for IPv4 an IPv4 next-hop address is used, and for IPv6 an IPv6 next-hop address is used
- RFC5549 first defined that the independence of prefix-type (IPv4 or IPv6) and type of next-hop address
- So IPv4 traffic can be transported across a LAN with no IPv4 configured, using an IPv6 next-hop address.

# Thank you!

[academy@de-cix.net](mailto:academy@de-cix.net)



Interested in more webinars? Please subscribe to our mailing list at <https://lists.de-cix.net/www/subscribe/academy>

# Links and further reading

# Links and further reading

- [RFC5549](#): Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop (obsolete)
- [RFC8950](#): Advertising IPv4 Network Layer Reachability Information (NLRI) with an IPv6 Next Hop







# DE CIX

<https://de-cix.net/academy>

Interested in more webinars? Please subscribe to our mailing list at <https://lists.de-cix.net/wws/subscribe/academy>

