ARP and high Broadcast rates

Why it is a problem for some routers and how to solve it

Wolfgang Tremmel
academy@de-cix.net

DE-CIX

Where networks meet

Networking Basics

DE-CIX Academy

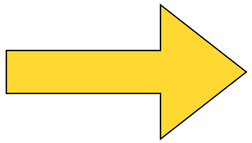

01 - Networks, Packets, and Protocols

02 - Ethernet + 02a - VLANs

ARP - Address Resolution Protocol

03 - Internet Protocol (IP), 03a Routing, 03b Global IP

04a - User Datagram Protocol (UDP)

04b - TCP

04c - ICMP

05 - Unicast, Broadast, Multicast, and Anycast

DE-CIX

Internet Model

Layers working together

- Interaction between Internet and Link layers
- What happens when an IP packet is sent
  - The first time to a new neighbor
  - What you know is its IPv4 address
  - But not its Ethernet MAC

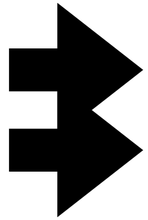

Layer	Name
5	Application
4	Transport
3	Internet
2	Link
1	Physical

DE CIX

Requesting an unknown MAC address

IP Layer and Ethernet layer

Layer	Name
5	Application
4	Transport
3	Internet
2	Link
1	Physical

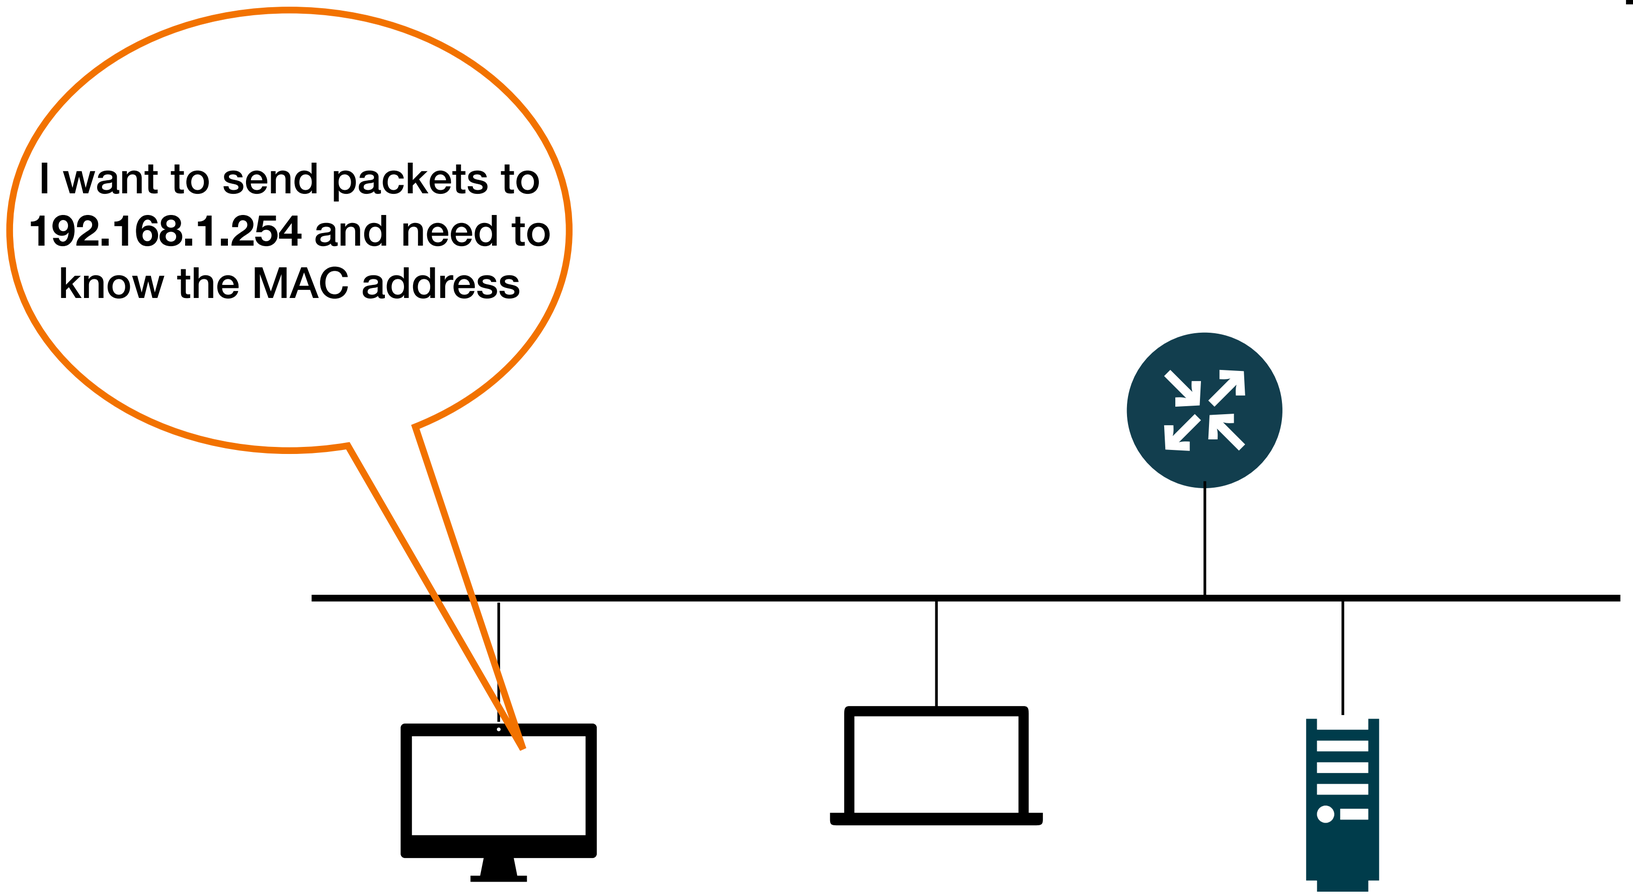

Requesting an unknown MAC address

IP Layer and Ethernet layer

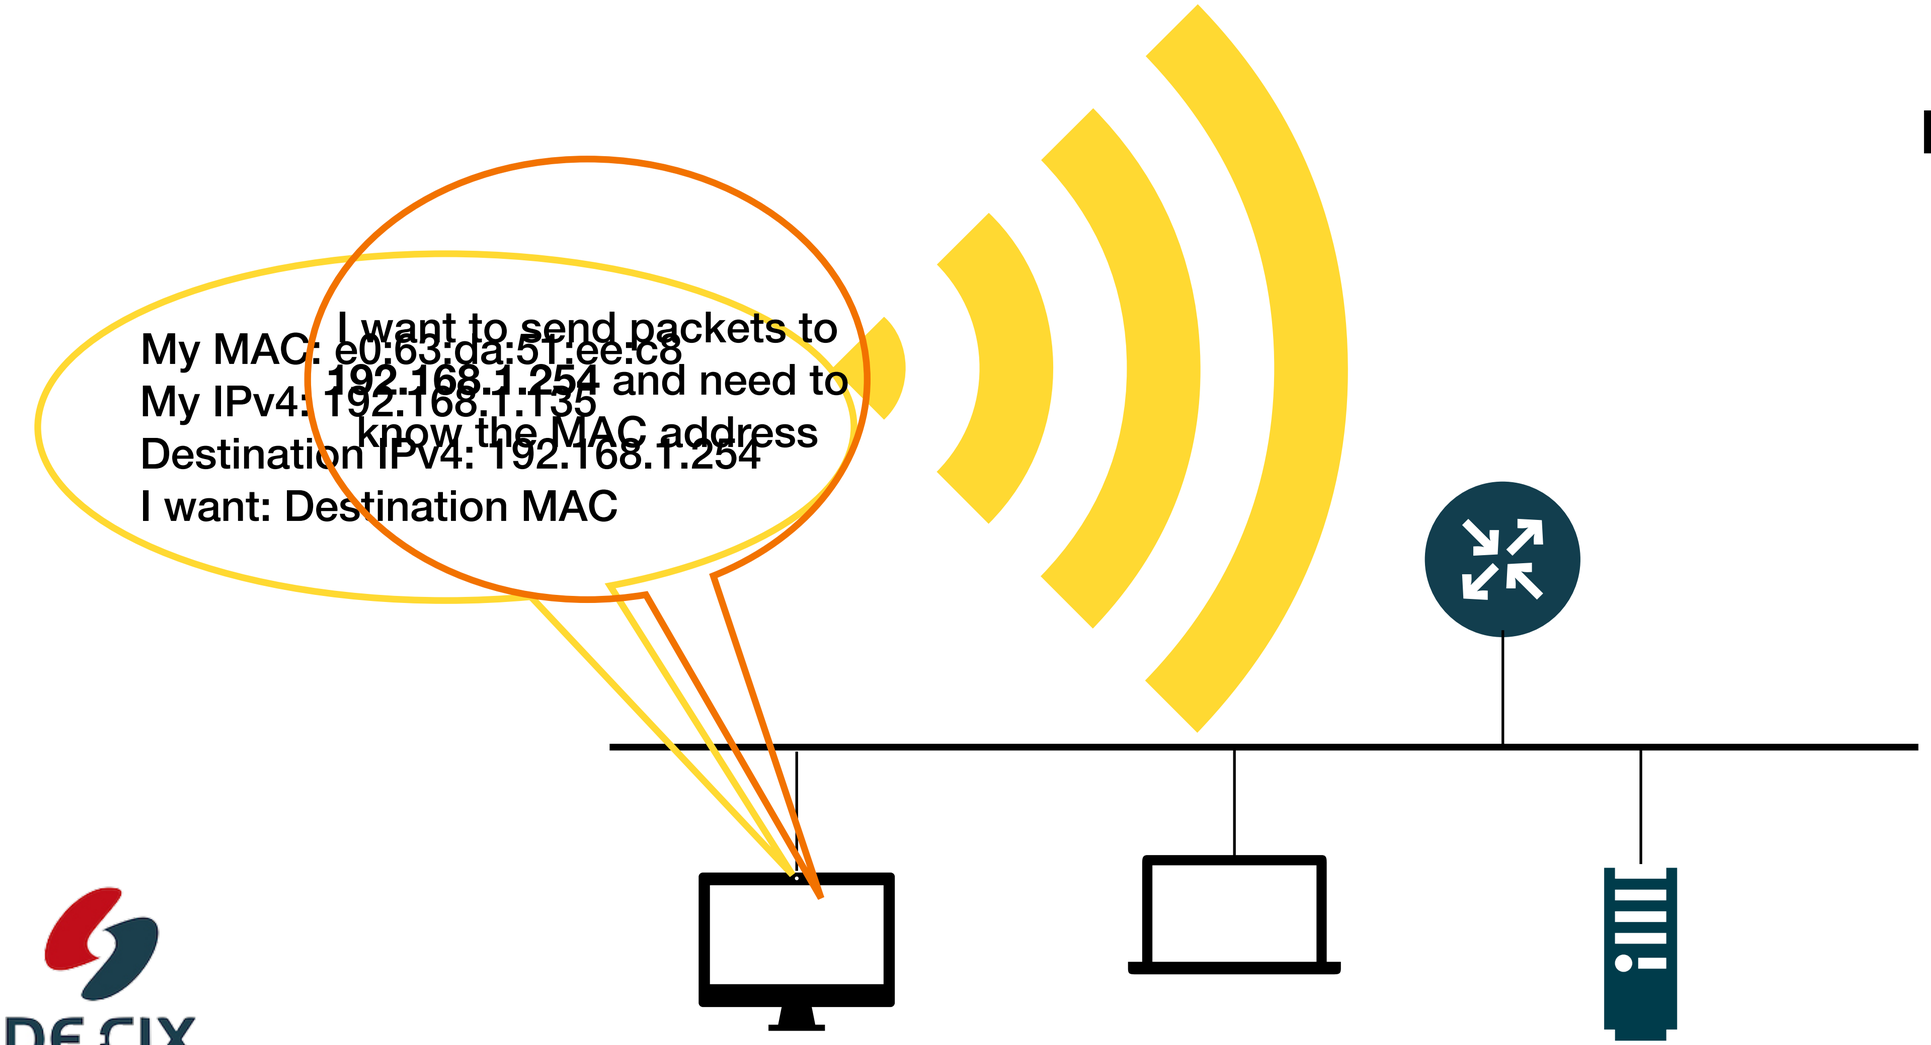

Layer	Name
5	Application
4	Transport
3	Internet
2	Link
1	Physical

Requesting an unknown MAC address

IP Layer and Ethernet layer

My MAC: 40:6c:8f:11:66:b6
My IPv4: 192.168.1.254
Destination IPv4: 192.168.1.254
Destination MAC: e0:63:da:51:ee:c8

My MAC: e0:63:da:51:ee:c8
My IPv4: 192.168.1.135
Destination IPv4: 192.168.1.254
I want: Destination MAC

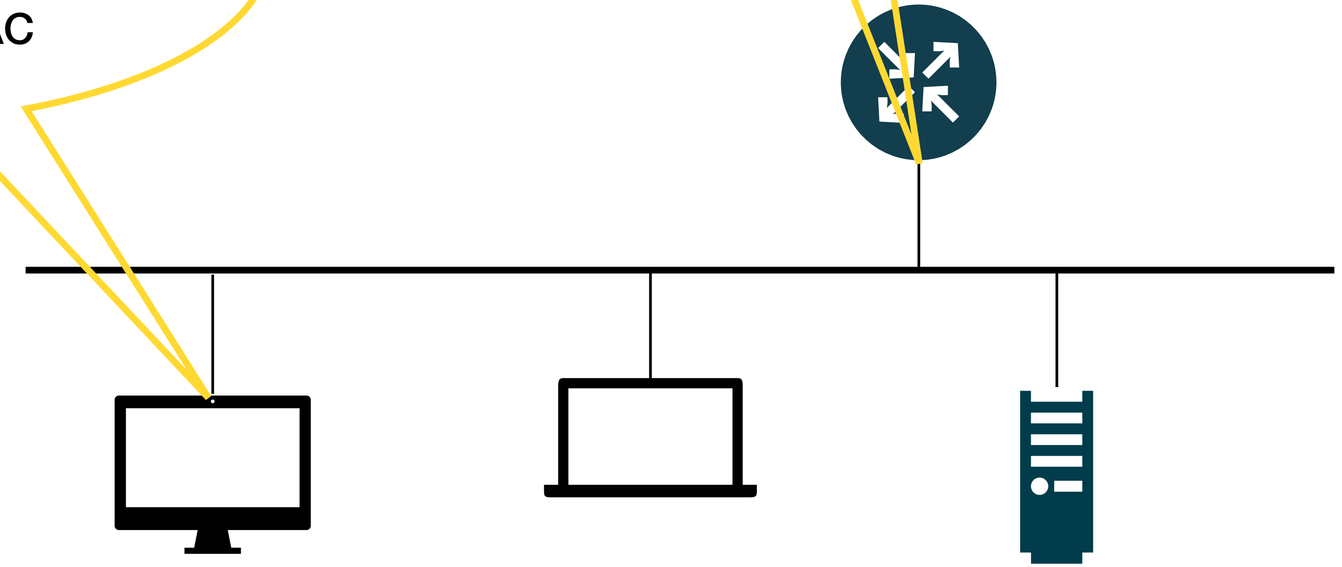

Layer	Name
5	Application
4	Transport
3	Internet
2	Link
1	Physical

DE CIX

Types of communication

Types of communication

Modes of communication

- Unicast
  - 1:1 communication
  - Standard in the Internet

- Broadcast
  - 1:all
  - Discovery

- Anycast
  - 1:1 but don't care which one
  - improve speed or reduncancy

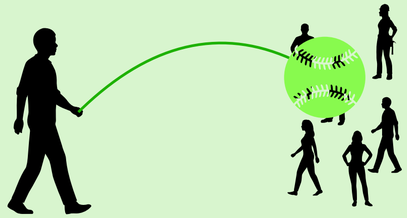

- Multicast
  - 1:group
  - address a group with something in common

DE CIX

Broadcast

Addressing a crowd

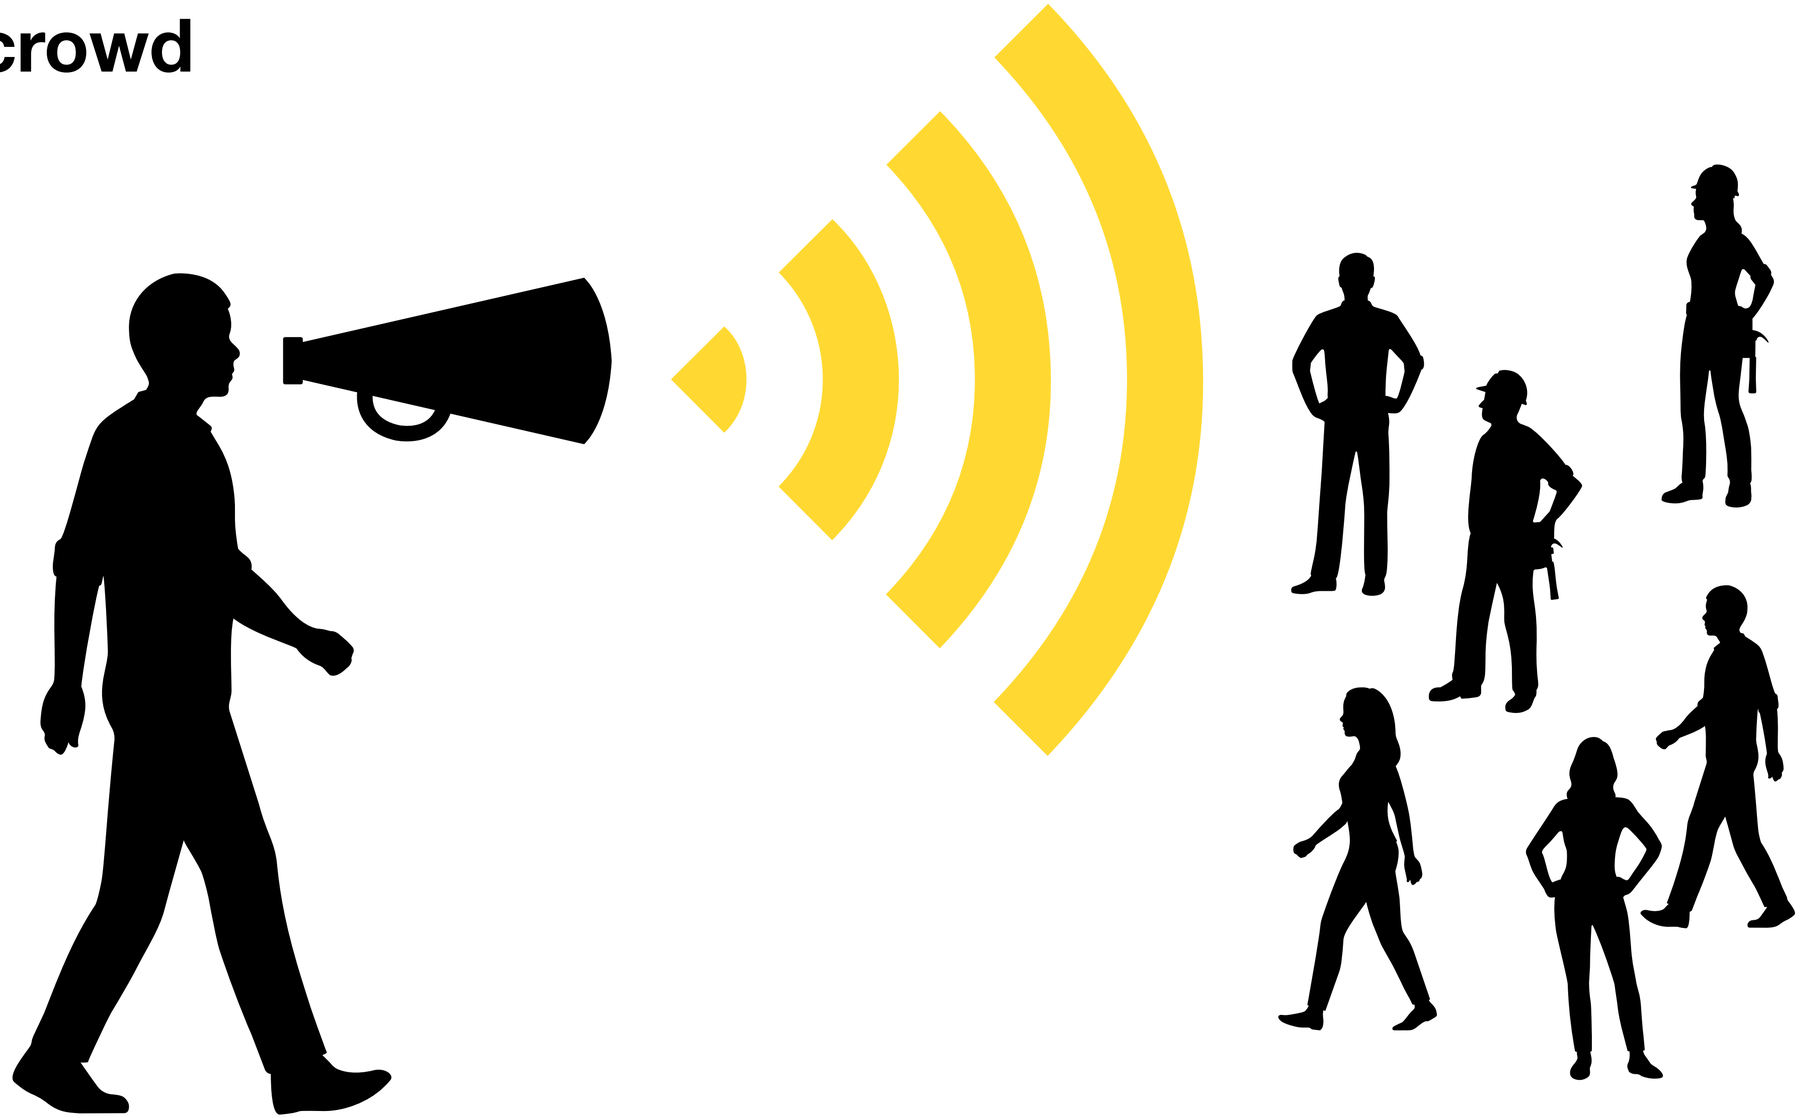

ARP - Address Resolution Protocol

ARP

Address Resolution Protocol

Layer	Name
5	Application

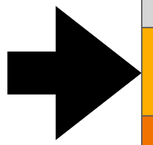

Layer	Name
3	ARP
2	Link
1	Physical

- Nearly as old as the Internet
- Introduced 1982 in RFC826
  - If you read the RFC, it is different than todays RFCs
- What is shown here is ARP how it is used and implemented today
- ARP is a network layer protocol - just like IPv4.
  - It is **not** encapsulated in IPv4 packets!

DE CIX

ARP

Packet format

- We have ARP requests and replies
  - They use the same packet format
- Lets do an ARP request
  - We want to know the MAC of the host with IPv4 192.168.1.254
  - Let's fill out an ARP request packet!

Layer	Name
3	ARP
2	Link
1	Physical

Byte	0	1
0	Hardware type (HTYPE)	
2	Protocol type (PTYPE)	
4	Hardware addess length	Protocol address length
6	Operation	
8	Sender hardware address	
10		
12		
14	Sender IPv4 address	
16		
18	Destination hardware address	
20		
22		
24	Destination IPv4 address	
26		

DE CIX

ARP Request

Packet format

- Hardware type is "1" for Ethernet
- Protocol type is "0x800" same as Ethertype field
- Operation is "1" for "Request"
- Sender IPv4 and MAC
- Destination IPv4 is what we know
- Destination MAC is unknown
  - We put in the broadcast address

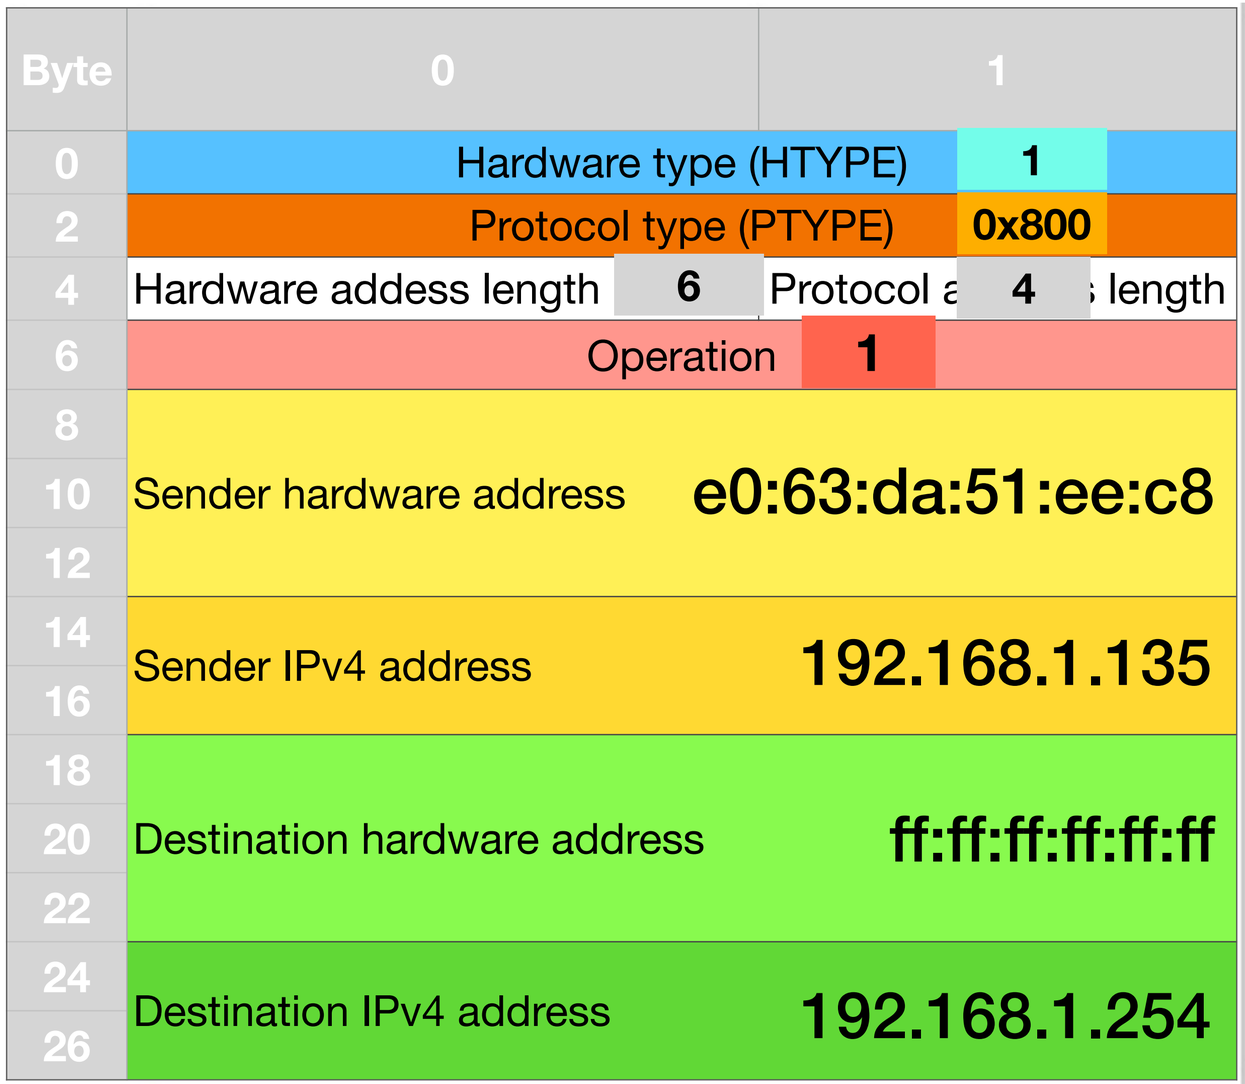

ARP Request

Sending it

- The ARP packet goes into the payload of an Ethernet frame
- Ethertype is 0x806
- Source MAC is the senders MAC address
- Destination MAC is the Ethernet Broadcast address

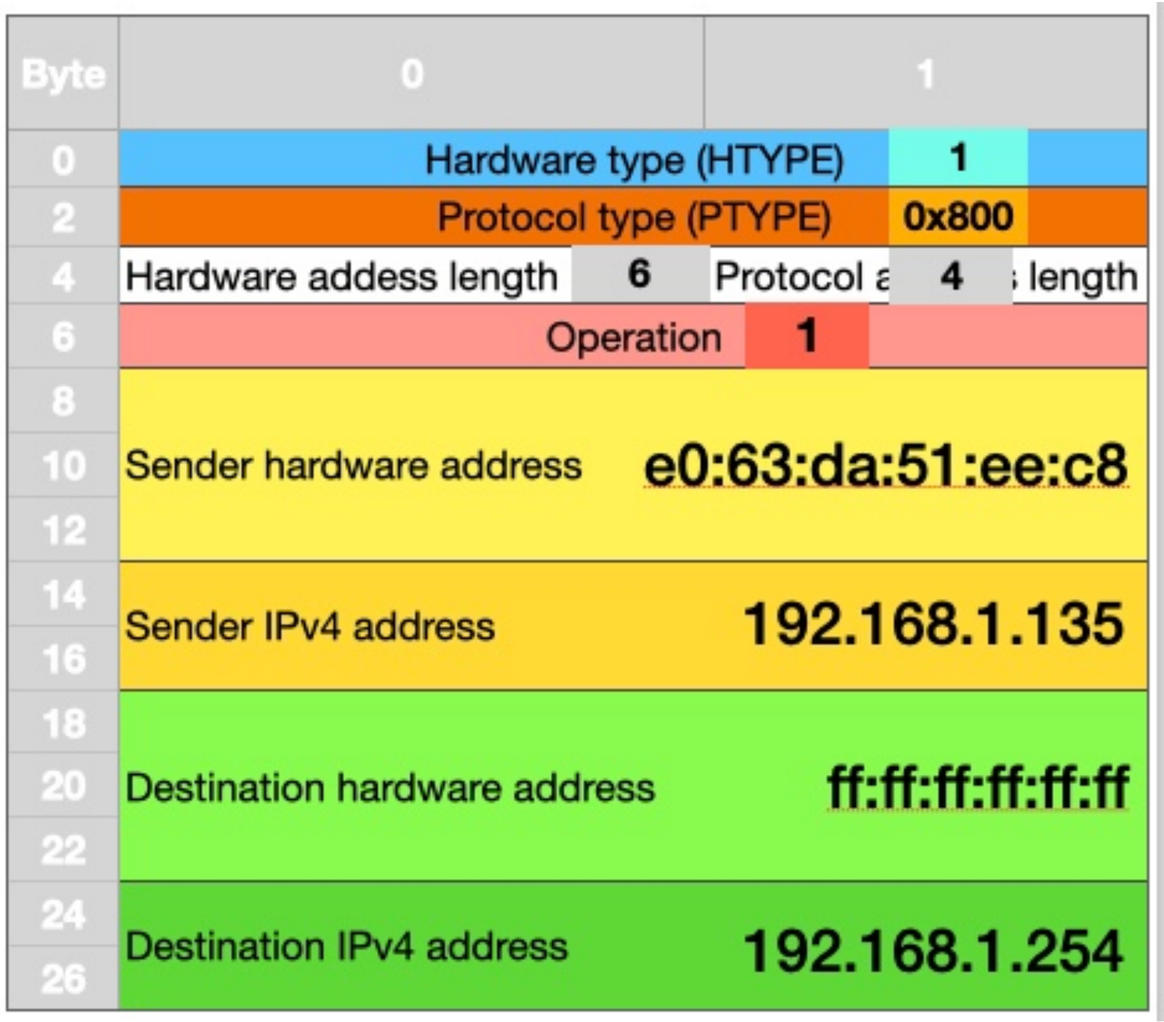

Preamble	SFD	Destination MAC Address	Source MAC Address	Ethertype	Payload	Checksum
10101010 10101010 10101010 10101010	10101011	48 Bits 6 Octets	48 Bits 6 Octets	16 Bits 2 Octets	46-1500 Octets	32 Bits 4 Octets
10101010 10101010 10101010 10101010	10101011	ff:ff:ff:ff:ff:ff	e0:63:da:51:ee:c8	0x806		

DE CIX

ARP Reply

Packet format

- Hardware type is "1" for Ethernet
- Protocol type is "0x800" same as Ethertype field
- Operation is "2" for "Reply"
- Destination IPv4 is the requestors
- Destination MAC also
- Sender IPv4 is the one asked for
- Sender MAC is the requested one

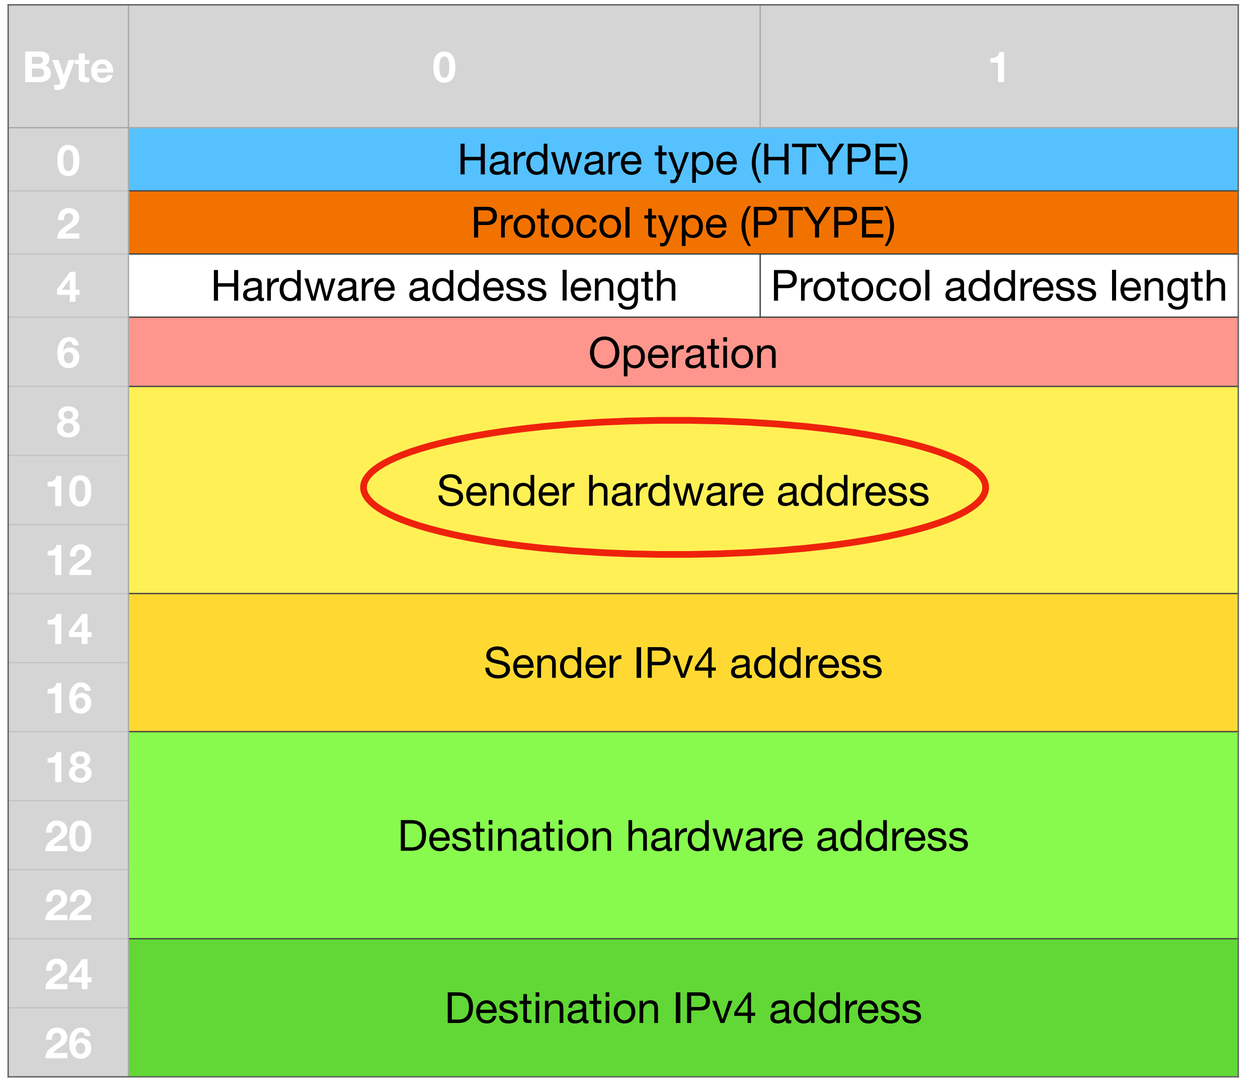

DE CIX

What's next?

Do I have to ask every time?

- Your host has now received the MAC address to an IPv4 address
- This is stored in the ***arp table***
  - With some timeout value (Cisco routers: 4 hours, Juniper routers: 25 minutes)
  - After that it gets refreshed (exact calculation is implementation specific and might be complicated)
- And another broadcast is made

DE CIX

Broadcast rate

Broadcast rate

High broadcast rate considered harmful

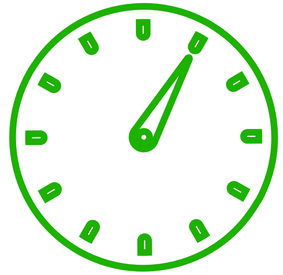

- On a normal LAN, the ARP rate is quite low
  - about 0-5 ARPs per second (measured on my home LAN)
- On a peering LAN with many routers connected it can go extremely high
  - DE-CIX Frankfurt: 700 ARPs per second and more
  - Of which each must be processed by every router connected
  - Some routers rate-limit ARP traffic
  - To protect themselves against attacks

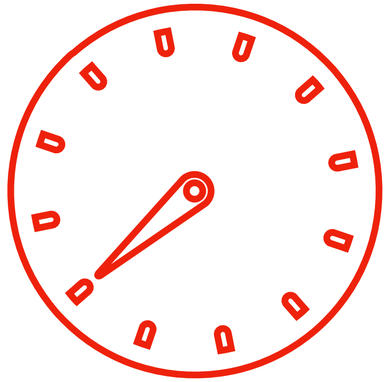

DE CIX

How a router works

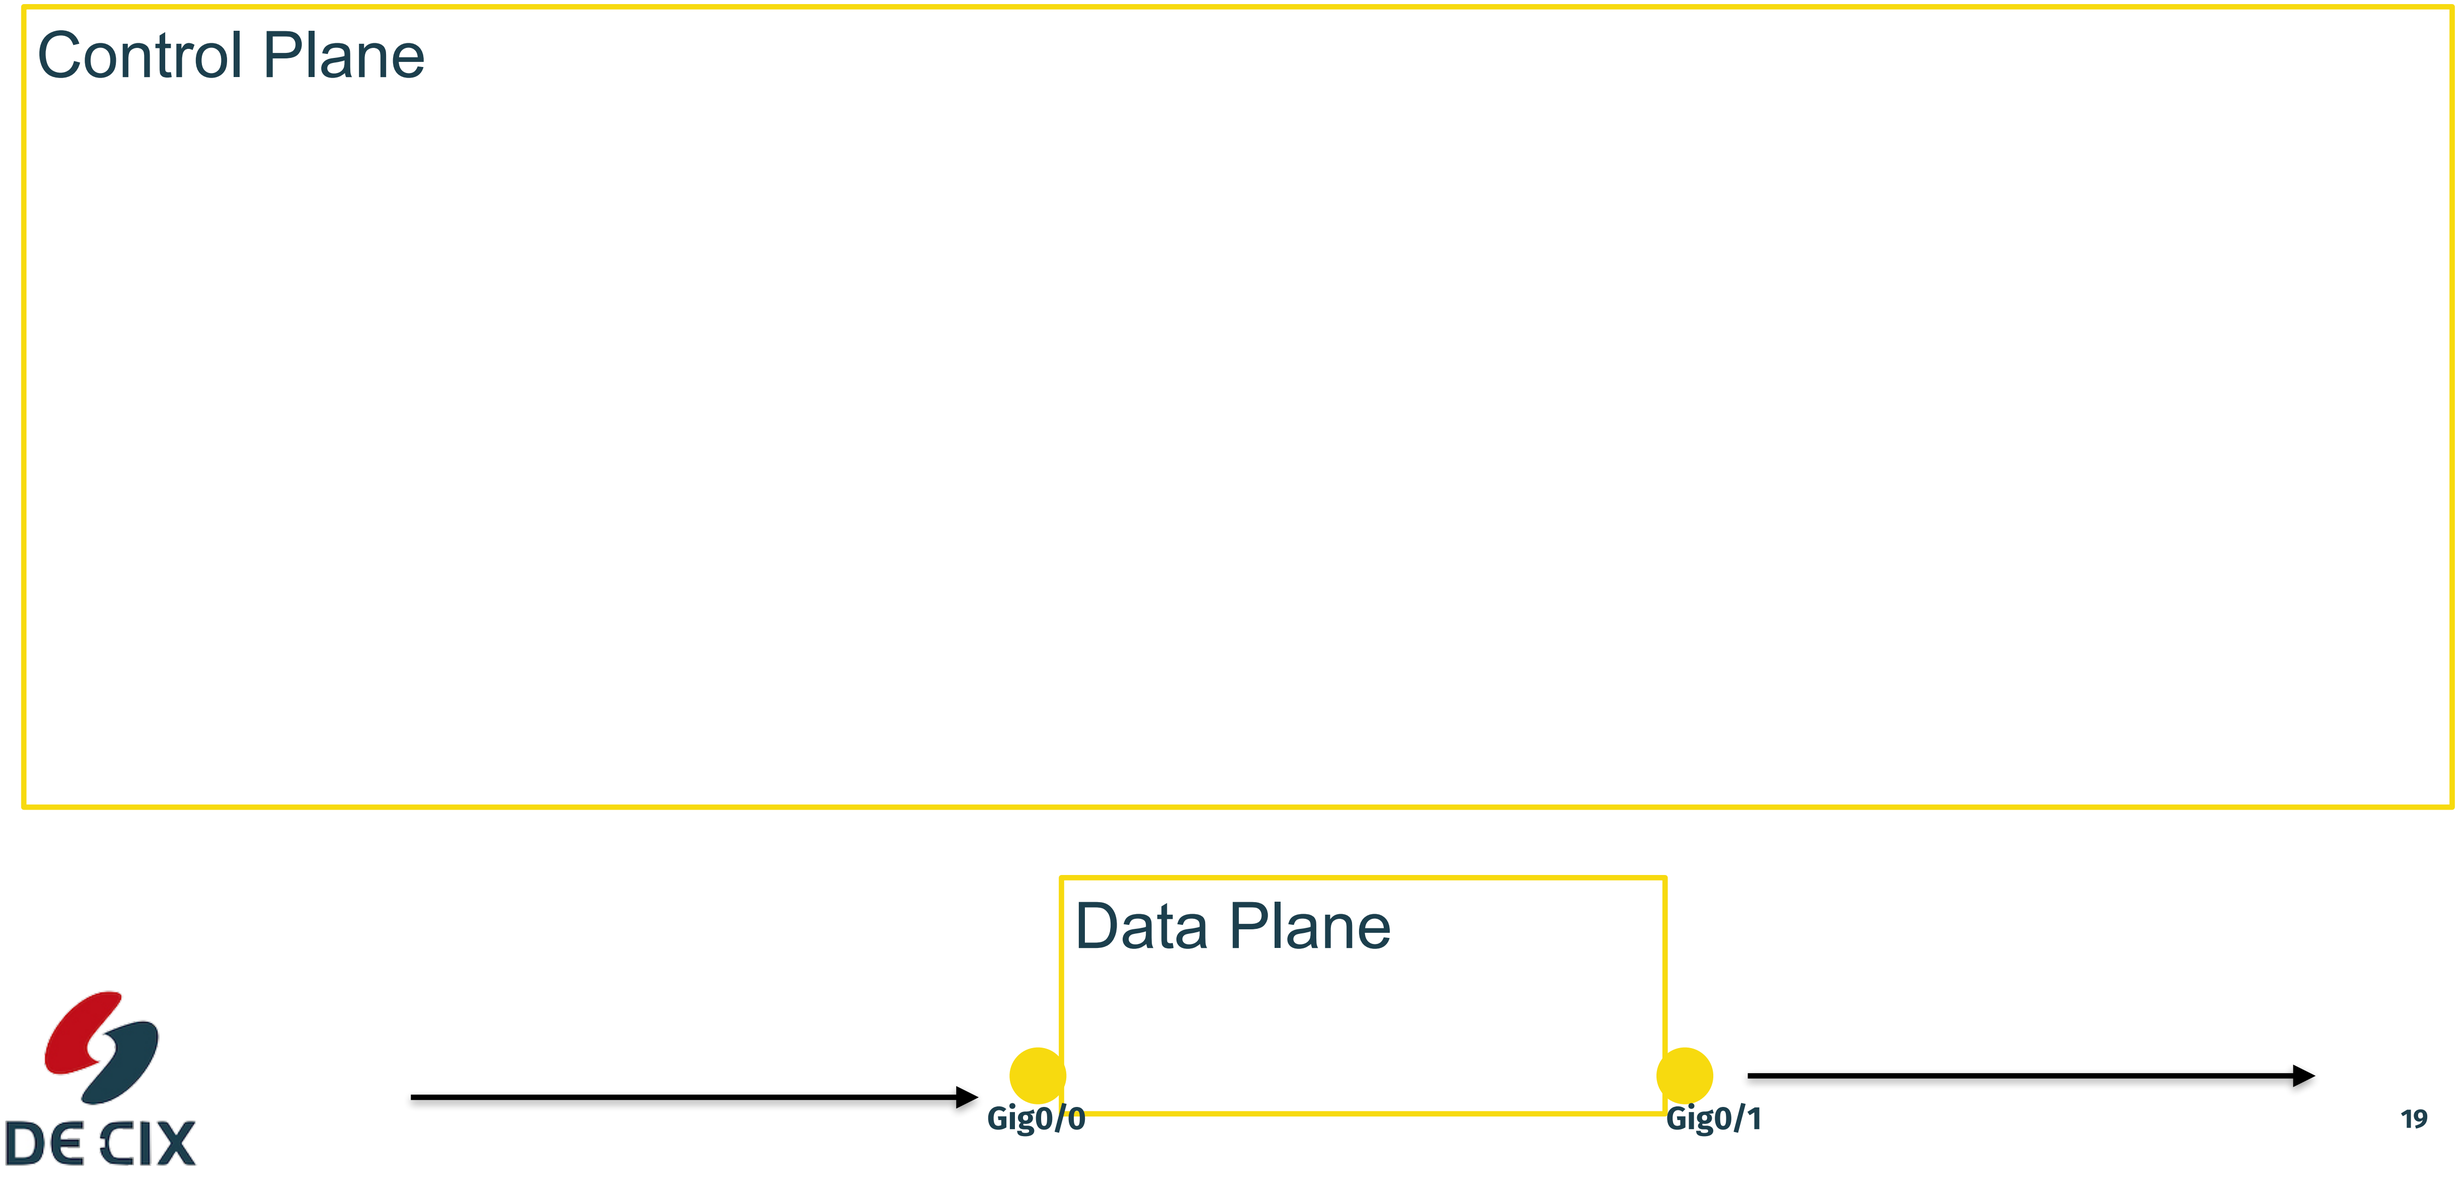

How a router works

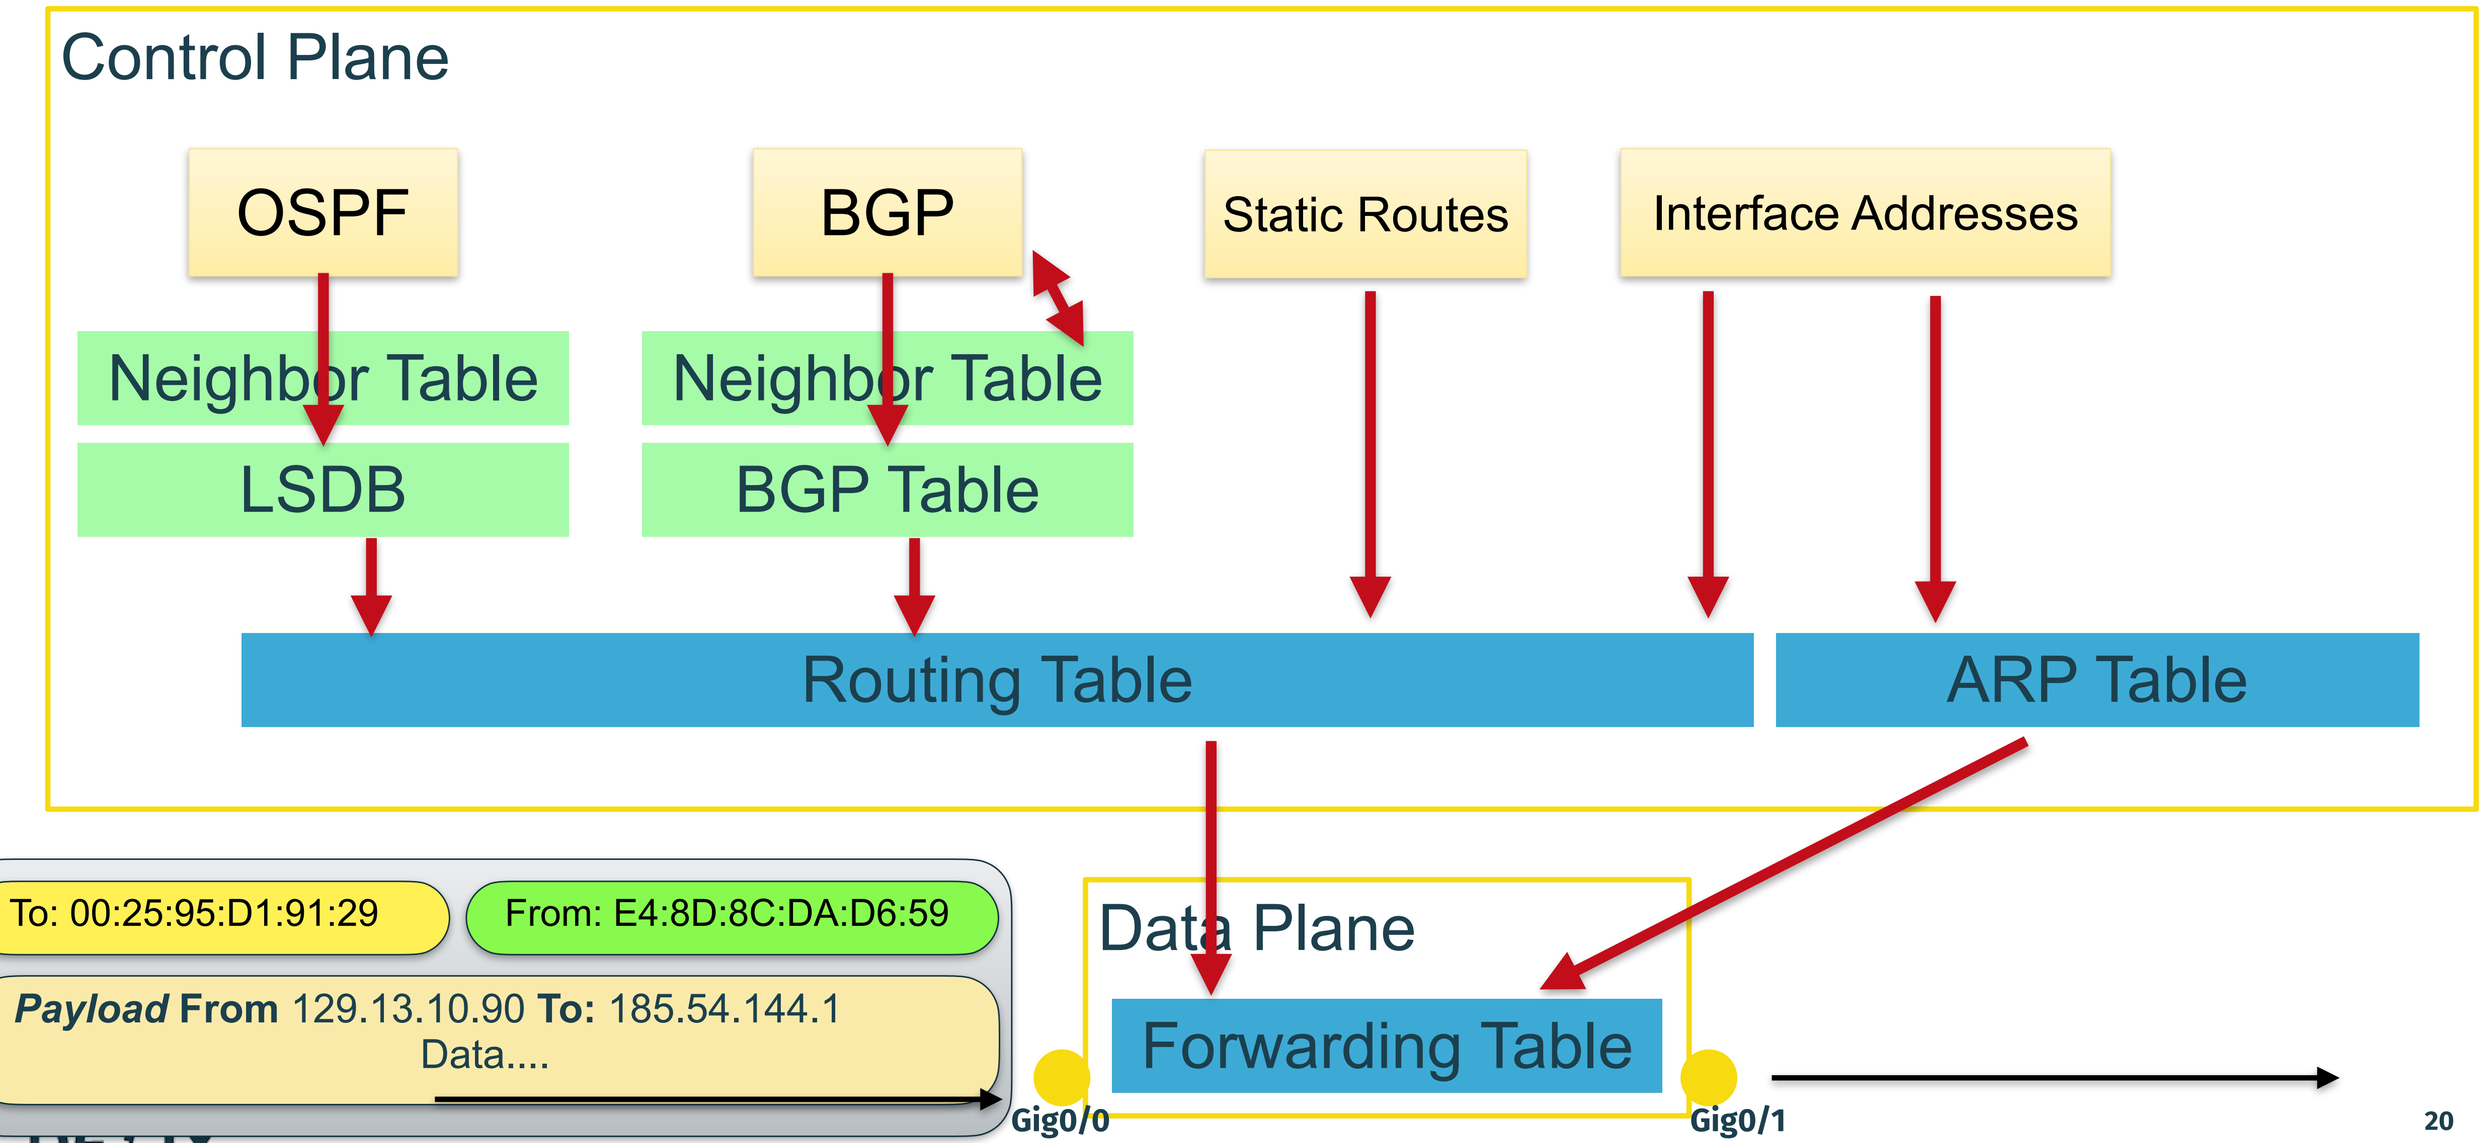

How a router works

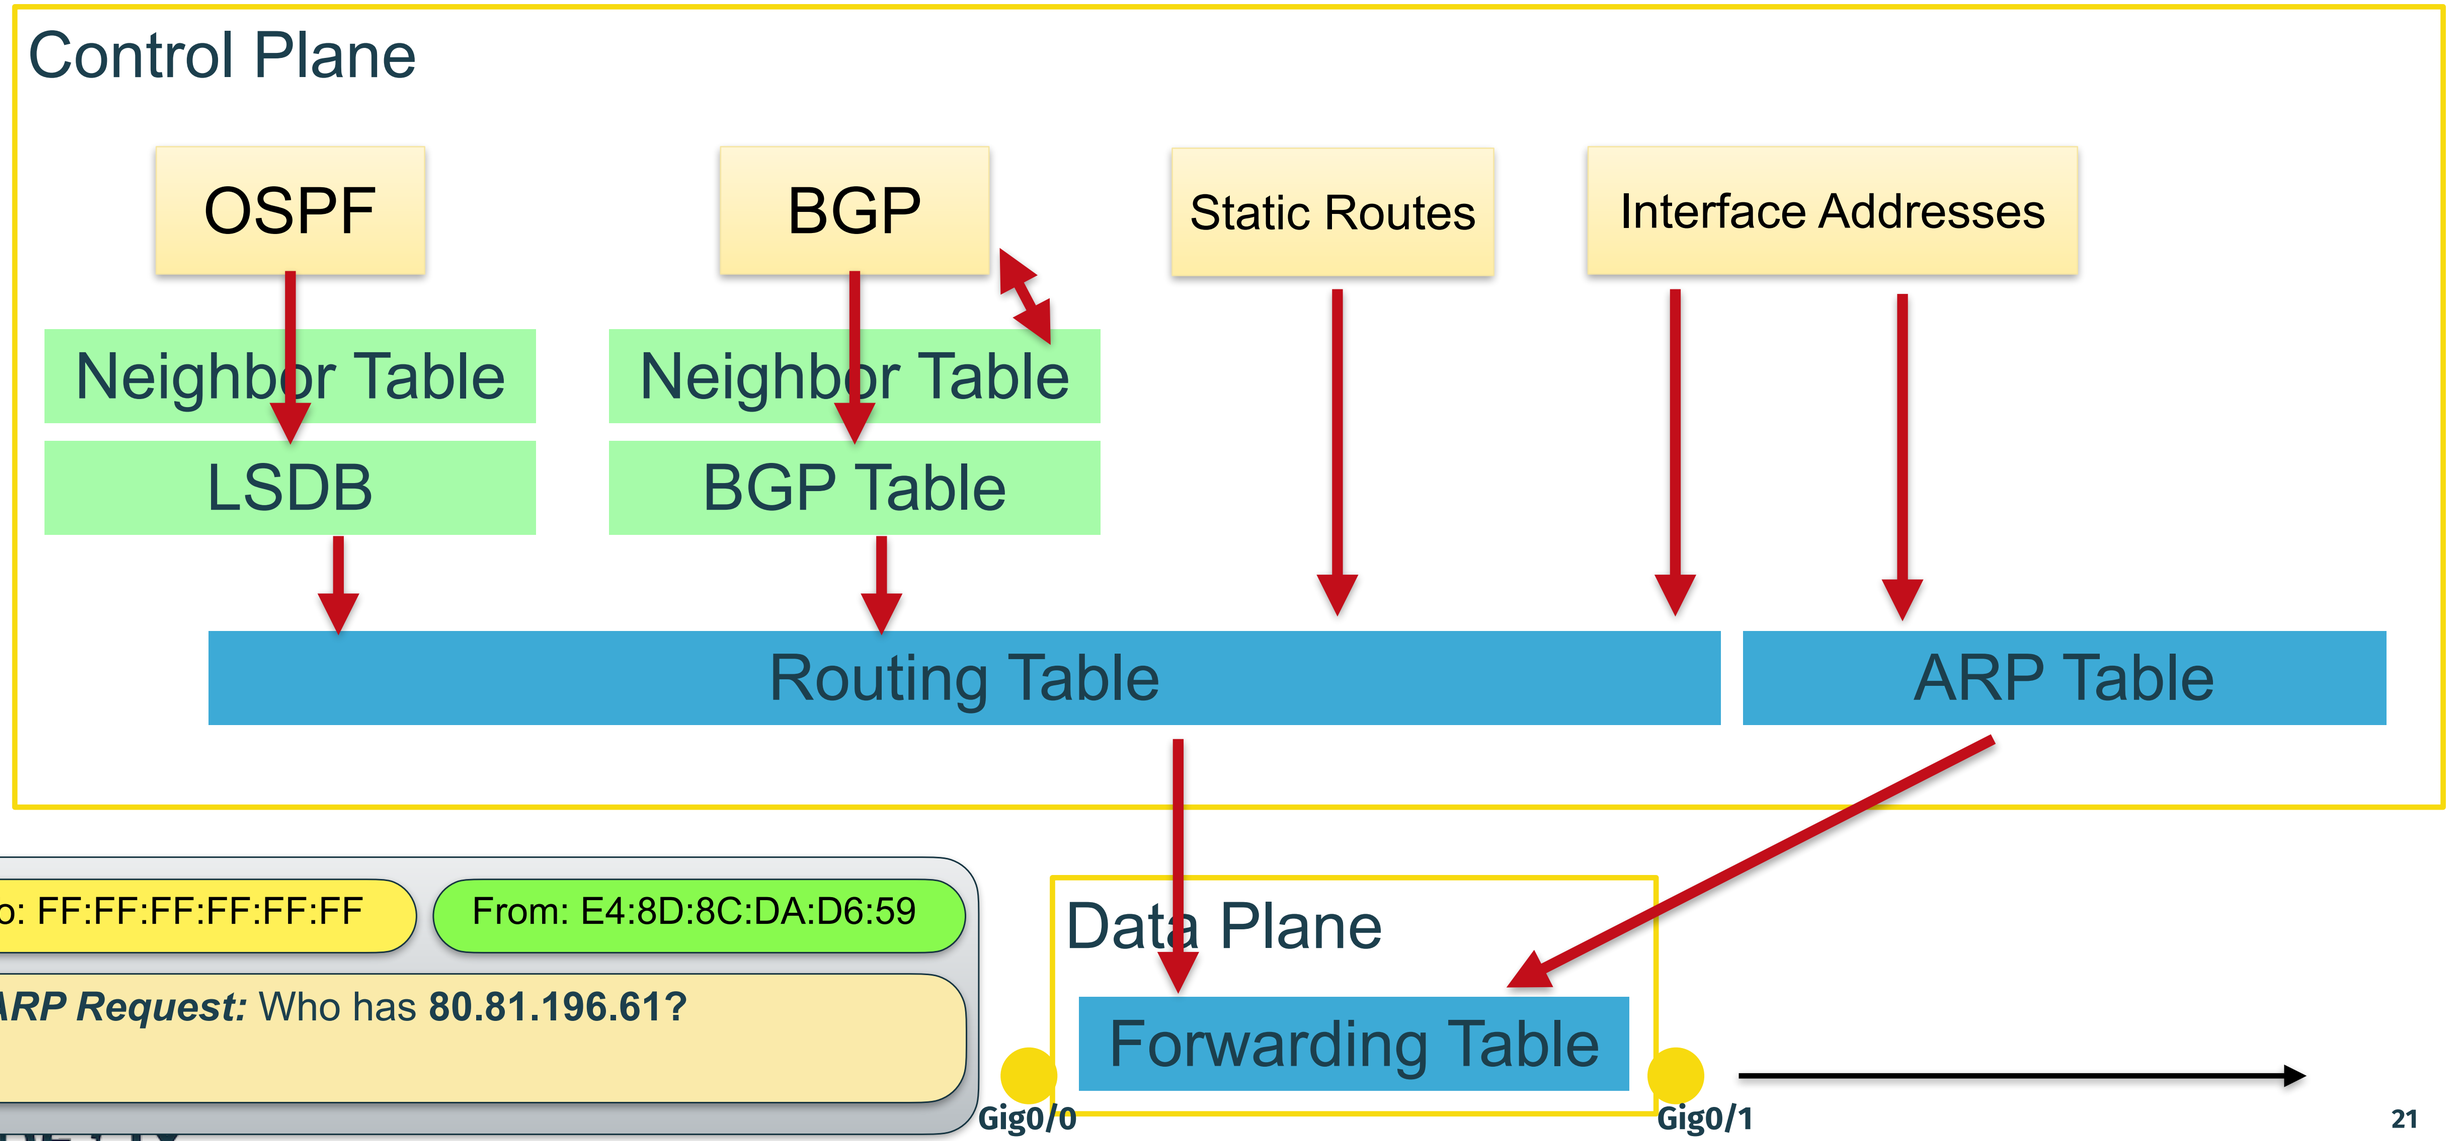

Protecting the control plane

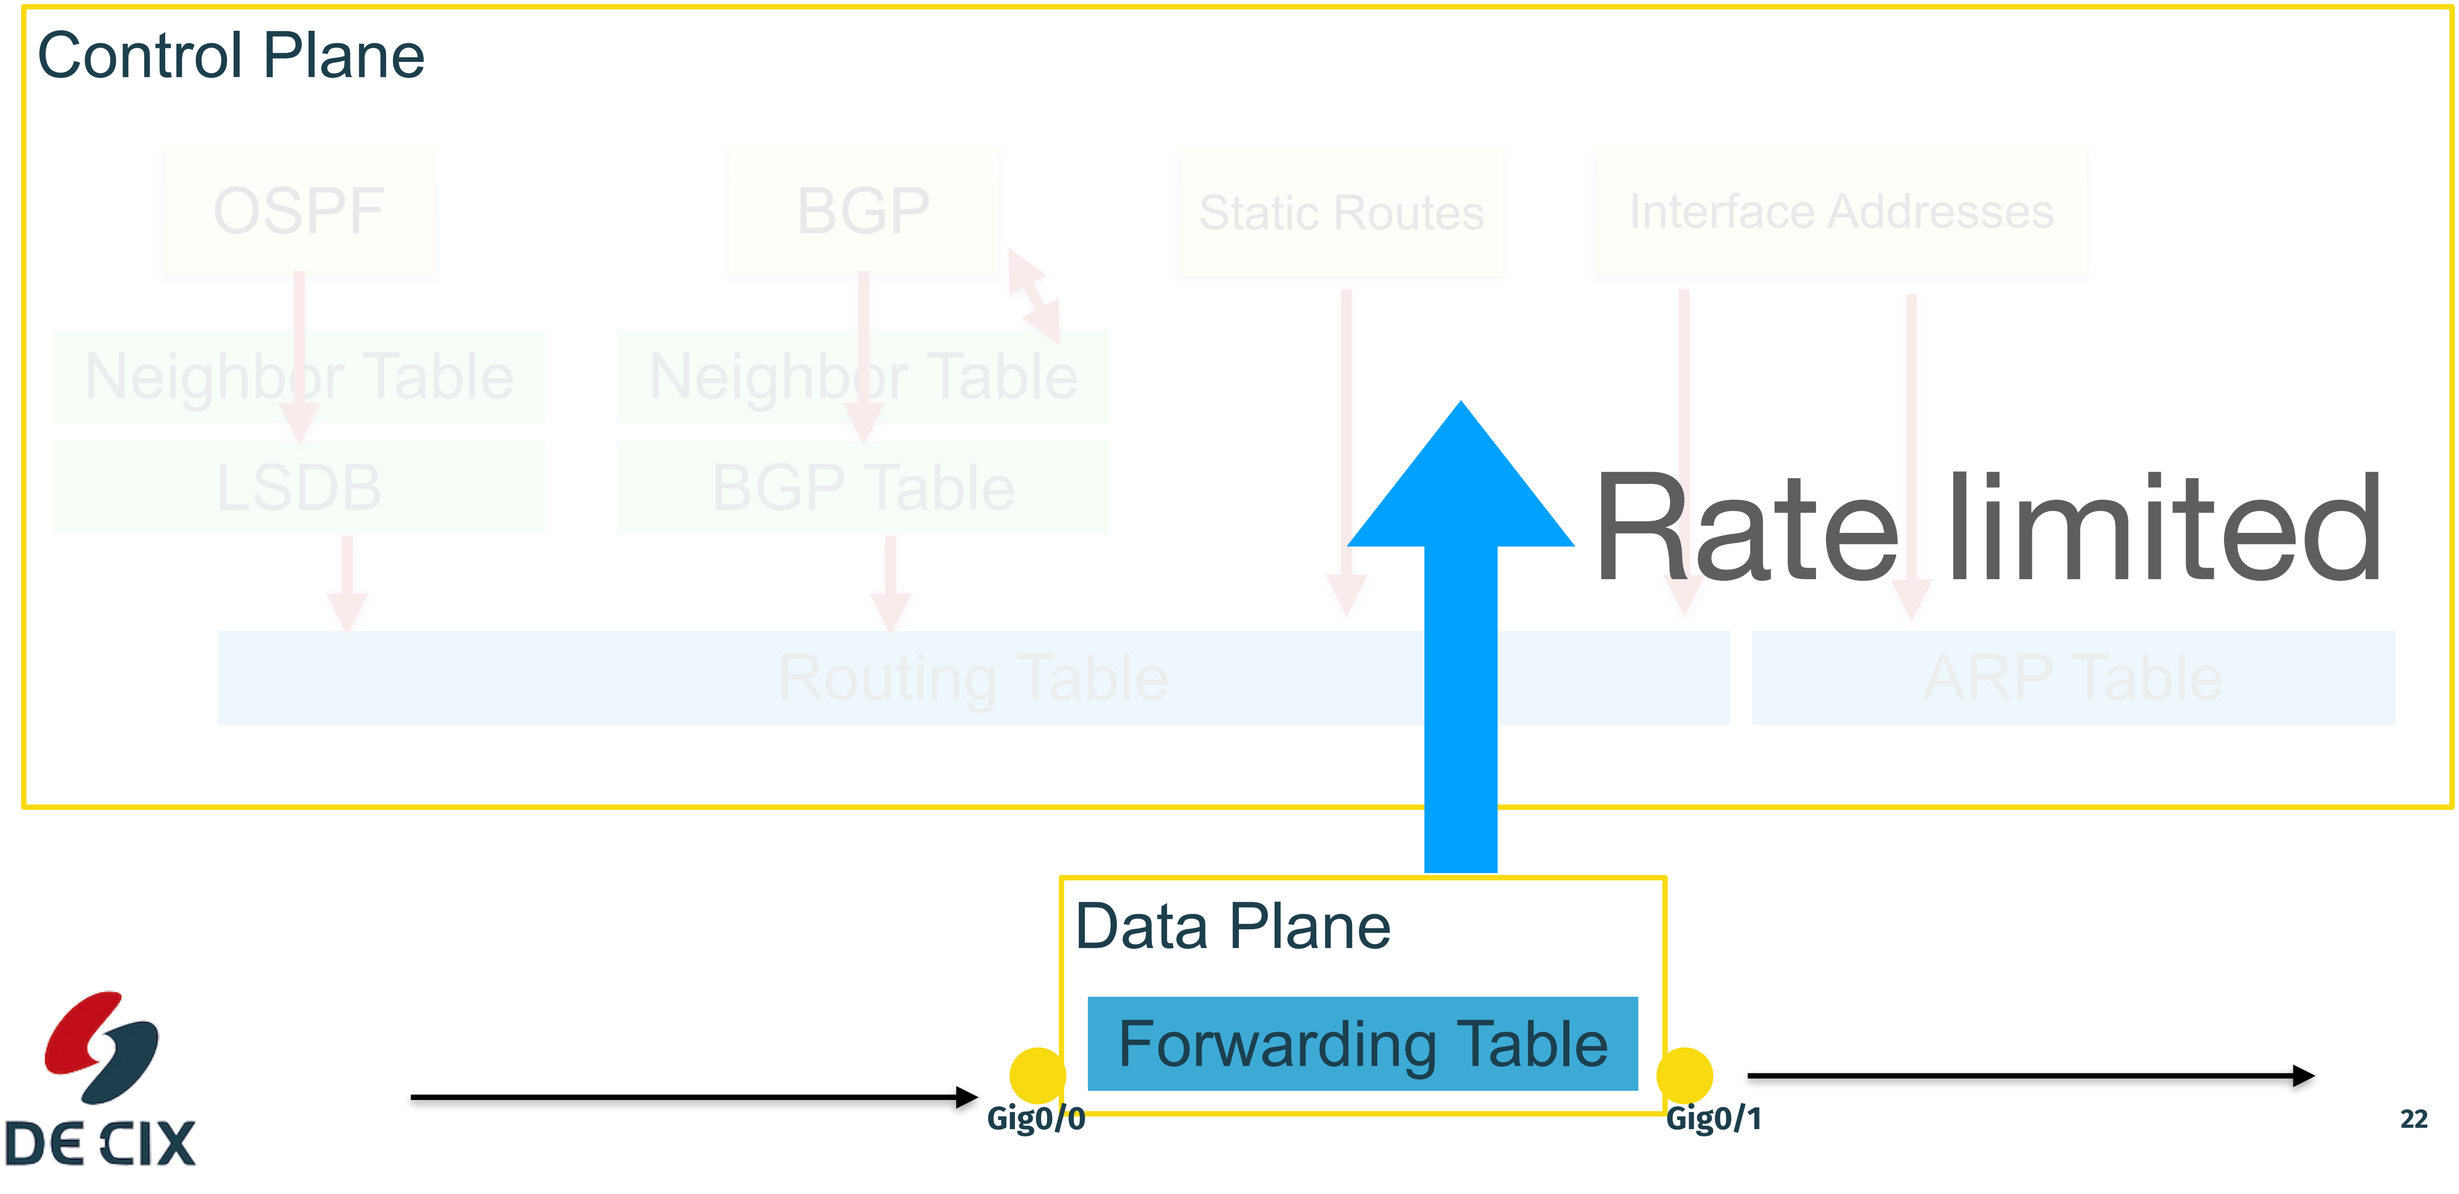

Protecting the control plane

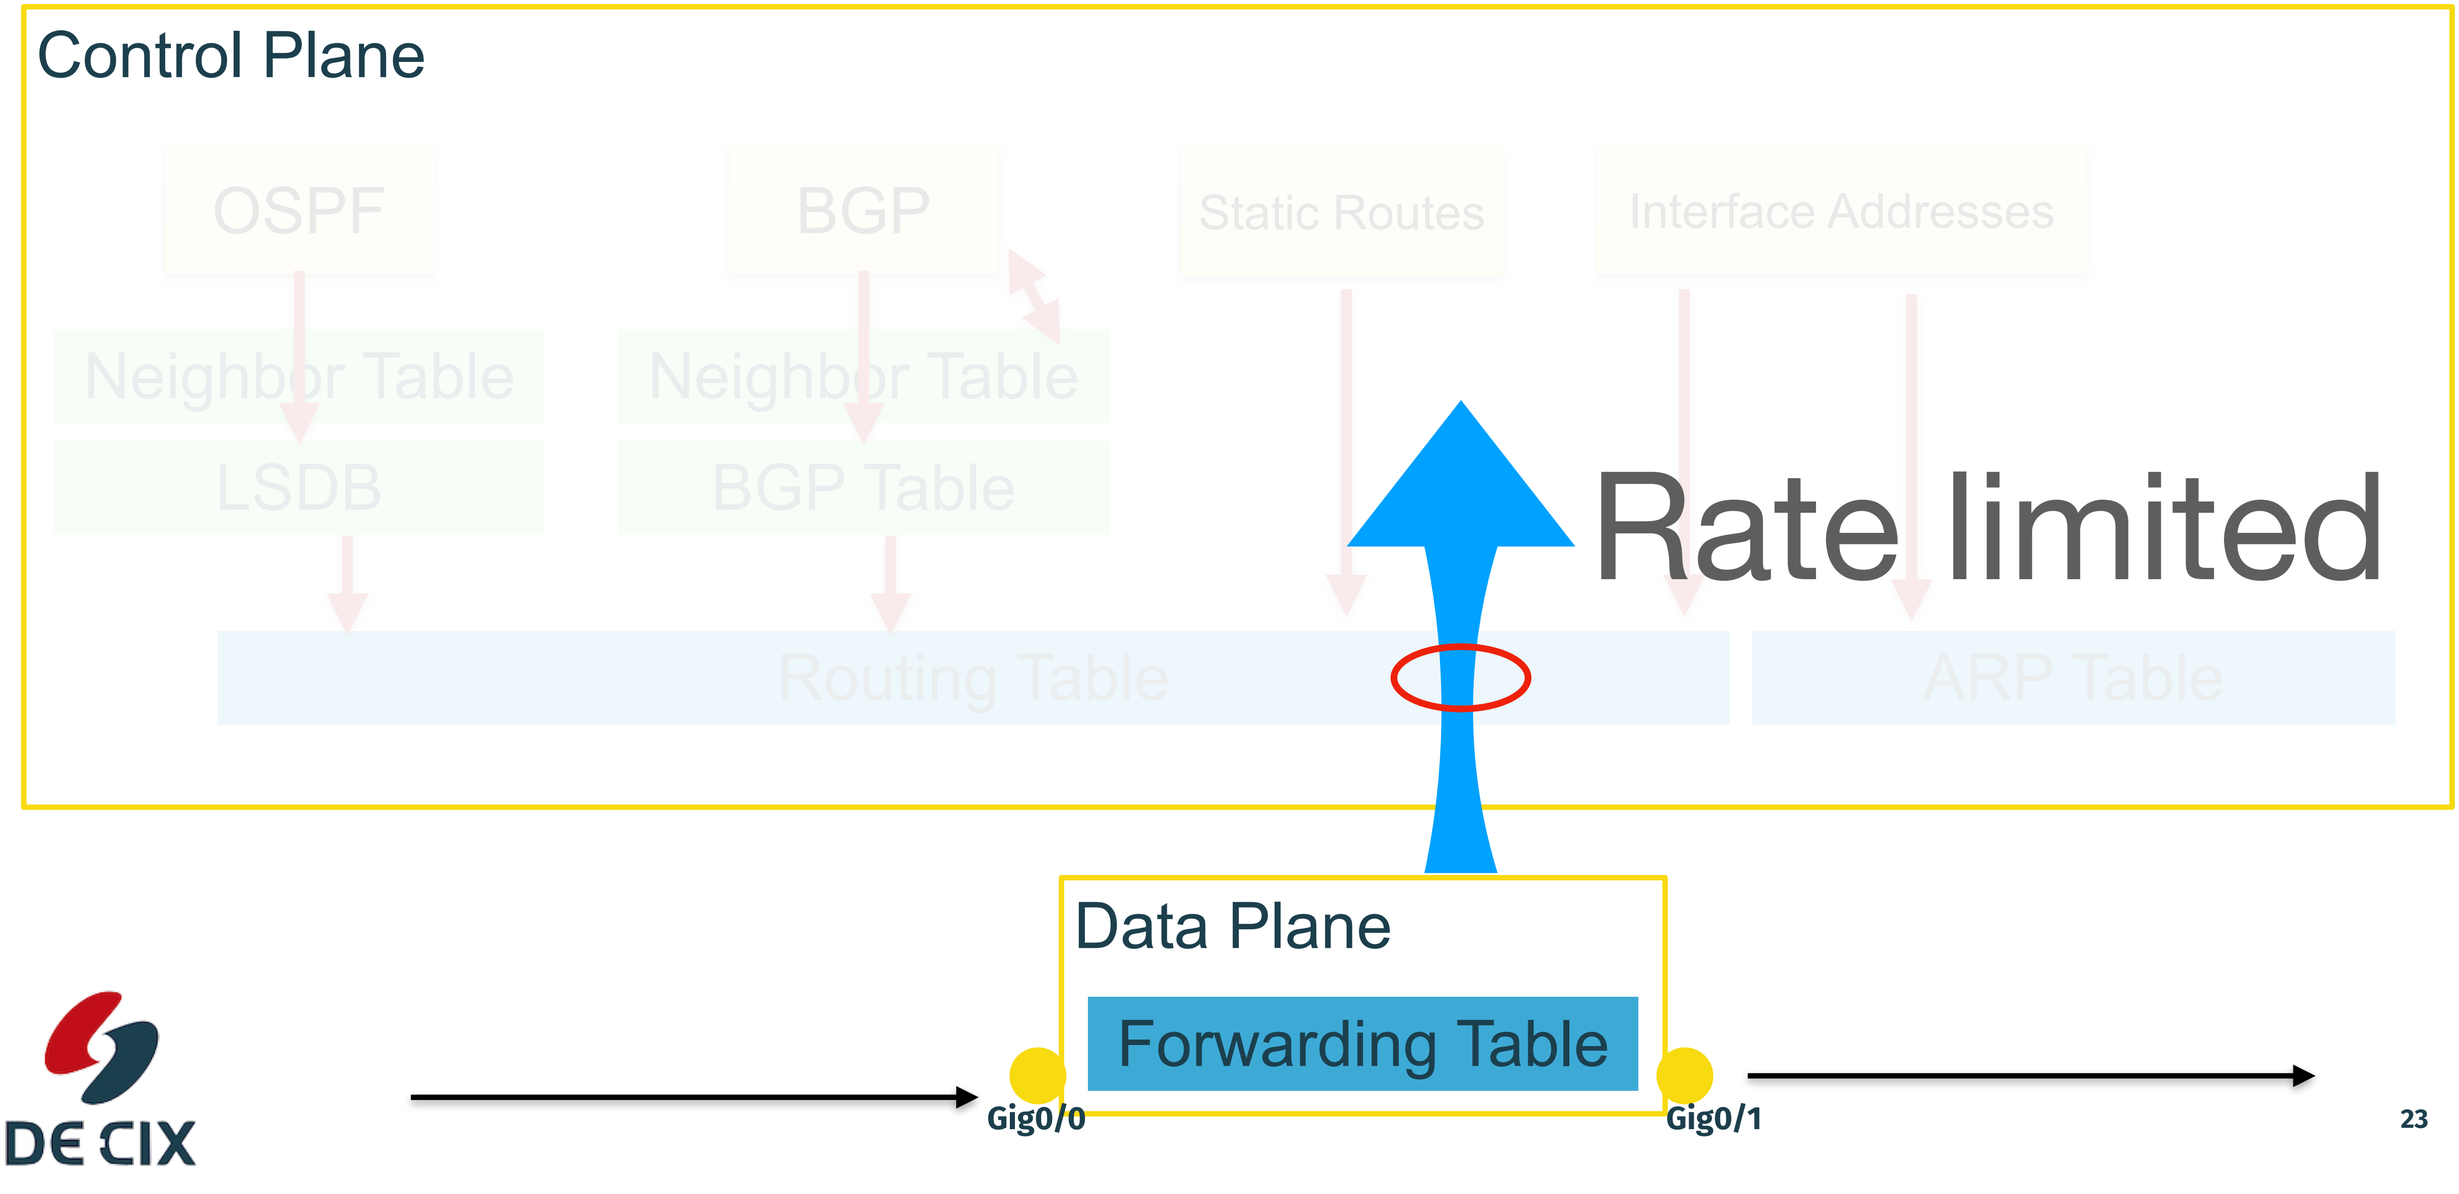

Solution: Adjust rate limit for DE-CIX FRA

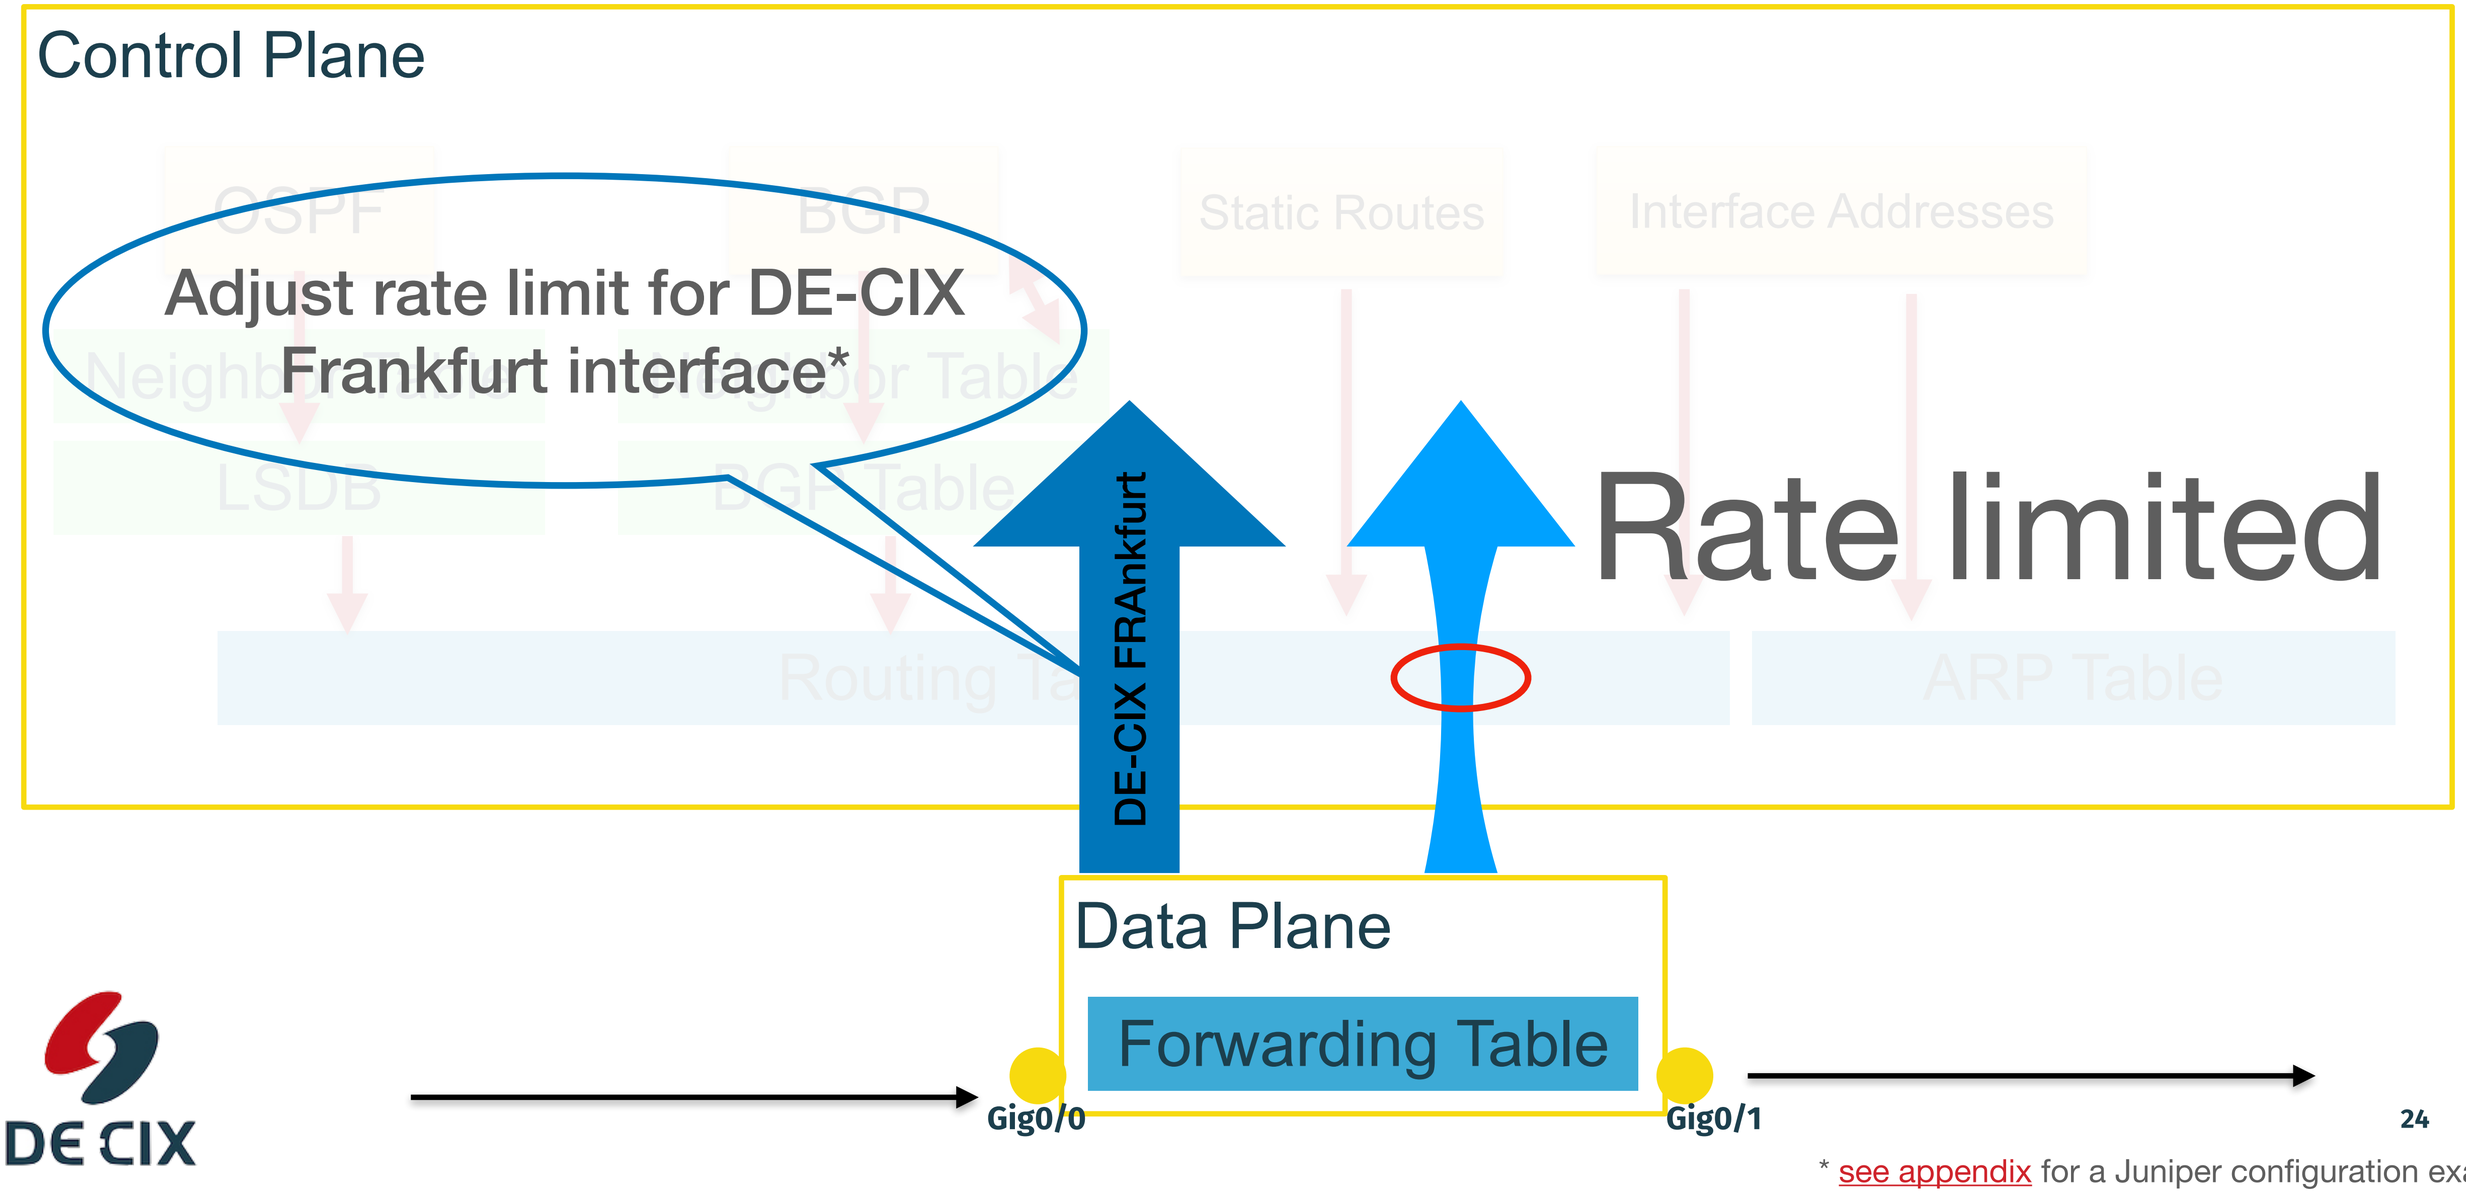

* see appendix for a Juniper configuration example

Conclusion

Conclusion

The ARP protocol

- The ARP protocol gives you the Ethernet address of a neighbor system
- ARP is IPv4 specific
- It works via Broadcast
- High Broadcast rates can overload a router
- Juniper polices the acceptable Broadcast rate
- This policing can be configured and adjusted to the needs of a peering LAN

DE CIX

Thank you!

academy@de-cix.net

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

DE-CIX Management GmbH | Lindleystr. 12 | 60314 Frankfurt | Germany
Phone + 49 69 1730 902 0 | sales@de-cix.net | www.de-cix.net

Links and further reading

Links and further reading

- Internet protocol - https://en.wikipedia.org/wiki/Internet_Protocol
- Protocol stack - https://en.wikipedia.org/wiki/Protocol_stack
  - Transport Layer: https://en.wikipedia.org/wiki/Transport_layer
  - Datagram: https://en.wikipedia.org/wiki/Datagram
- IP Network Model: https://en.wikipedia.org/wiki/Internet_protocol_suite
- IPv4
  - IPv4 - https://en.wikipedia.org/wiki/IPv4
- IPv6
  - IPv6 itself - https://en.wikipedia.org/wiki/IPv6
  - IPv6 header - https://en.wikipedia.org/wiki/IPv6_packet
  - Transmission of IPv6 over Ethernet: https://tools.ietf.org/html/rfc2464
- History of Internet and IP
  - Internet Hall of Fame - https://internethalloffame.org
  - Defense Advanced Research Projects Agency (DARPA) - https://www.darpa.mil
  - ARPANET - https://www.darpa.mil/about-us/timeline/arpanet
  - The "Protocol Wars" - https://en.wikipedia.org/wiki/Protocol_Wars

Links and further reading

ARP and Broadcast

- Wikipedia articles:
  - https://en.wikipedia.org/wiki/Broadcasting_(networking)
  - https://en.wikipedia.org/wiki/Broadcast_address
- Ethernet related:
  - https://en.wikipedia.org/wiki/Broadcast_storm
- IP related:
  - Address Resolution Protocol: RFC826, Wikipedia

Internet RFCs (Standards)

- There are too many RFCs dealing with IPv4 and IPv6 to be listed here
- Just go to https://tools.ietf.org/html/ and use the search field
- How does something become RFC? https://www.rfc-editor.org/pubprocess/
- The IETF - Internet Engineering Task Force

Juniper configuration example

- You need to define a "policer" first (in the *firewall* context):

```
policer DECIX-arp {
       if-exceeding {
           bandwidth-limit 2m;
           burst-size-limit 2m;
       }
       then discard;
}
```

Allowed ARP traffic in bits/s

Burst size in bytes

- Then you apply the policer to your DE-CIX interface:

```
interface ae5 {
       unit 2800 {
           description "DE-CIX Peering Interface";
           vlan-id 2800;
           family inet {
               policer {
                   arp DECIX-arp;
               }
           }
       }
}
```

Interface of your DE-CIX connection

VLAN you are using to connect to DE-CIX Frankfurt

Reference to your policer