Networking Basics

02b - Ethernet + VLANs

Wolfgang Tremmel
academy@de-cix.net

DE CIX

Where networks meet

Networking Basics

02b - Ethernet + VLANs + QinQ

Wolfgang Tremmel
academy@de-cix.net

DE-CIX

Where networks meet

Networking Basics

DE-CIX Academy

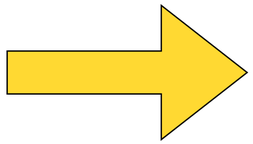

01 - Networks, Packets, and Protocols

02 - Ethernet

02a - Ethernet/VLANs

+ QinQ

03 - IP: the Internet Protocol

03a - IP addresses, prefixes, and routing

03b - Global IP routing

04a - UDP

04b - TCP

05 - Uni-, Broad-, Multi-, and Anycast

DE-CIX

VLANs at DE-CIX

VLANs at DE-CIX

How we use them

- Customers connect to DE-CIX via Ethernet
- Standard connection is a untagged access port

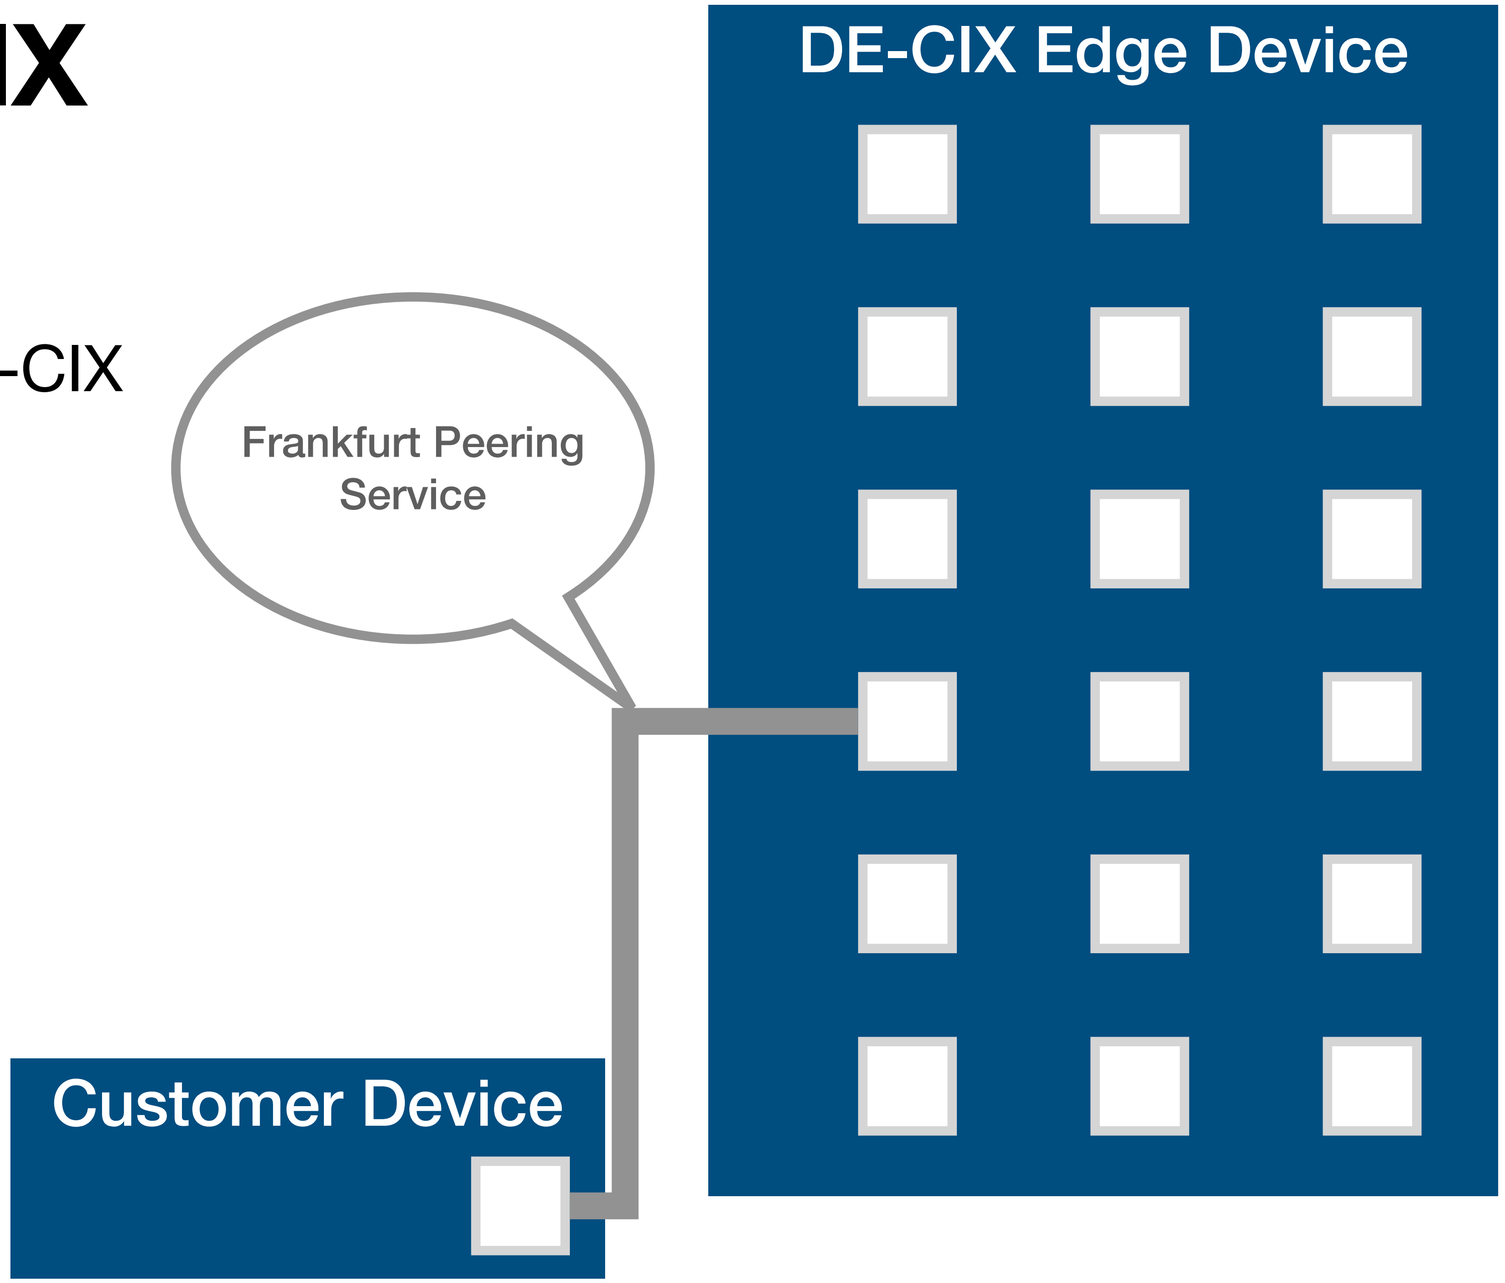

VLANs at DE-CIX

How we use them

- Customers connect to DE-CIX via Ethernet
- Standard connection is a untagged access port
- But we can also deliver via a tagged trunk-like port

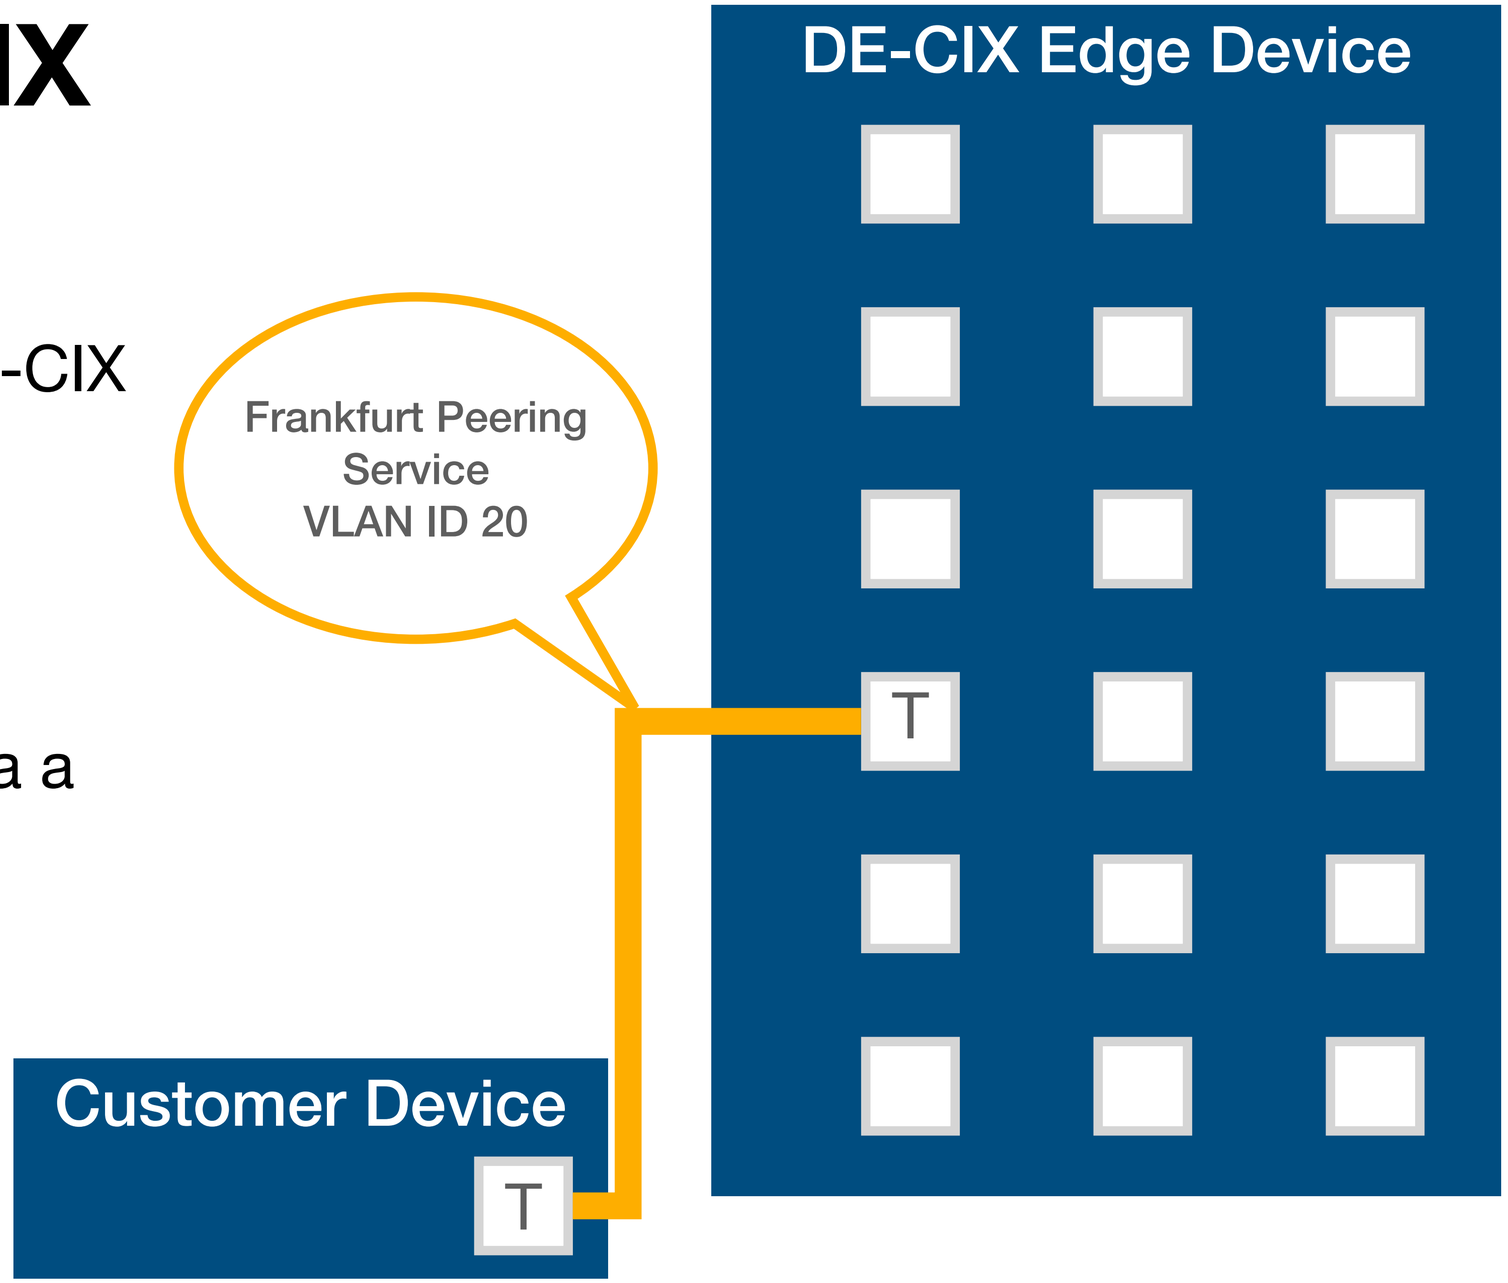

VLANs at DE-CIX

How we use them

- Customers connect to DE-CIX via Ethernet
- Standard connection is a untagged access port
- But we can also deliver via a tagged trunk-like port
- And on a trunk-like port we can deliver multiple services

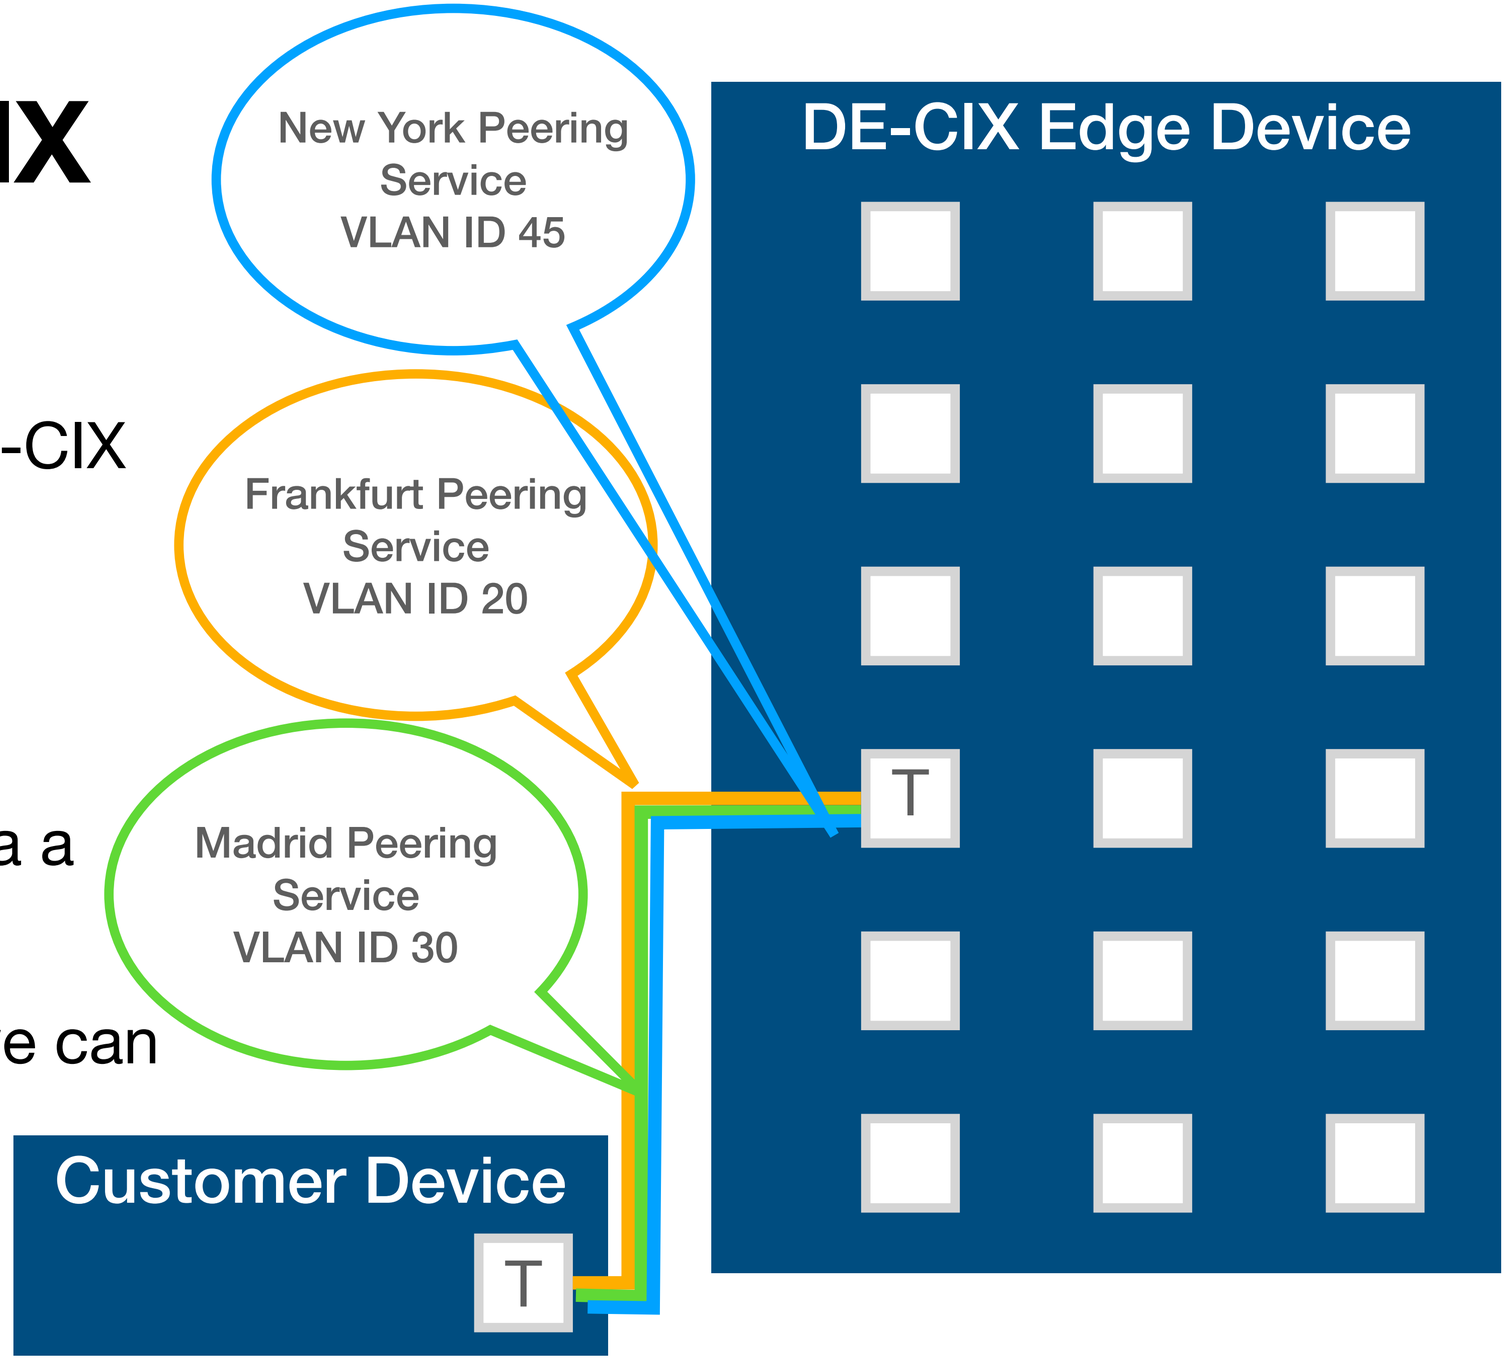

VLANs at DE-CIX

Connect to the Cloud

- The same way we connect customers to (multiple) Cloud service providers
- At DE-CIX the VLAN ID on each end does not have to be the same!

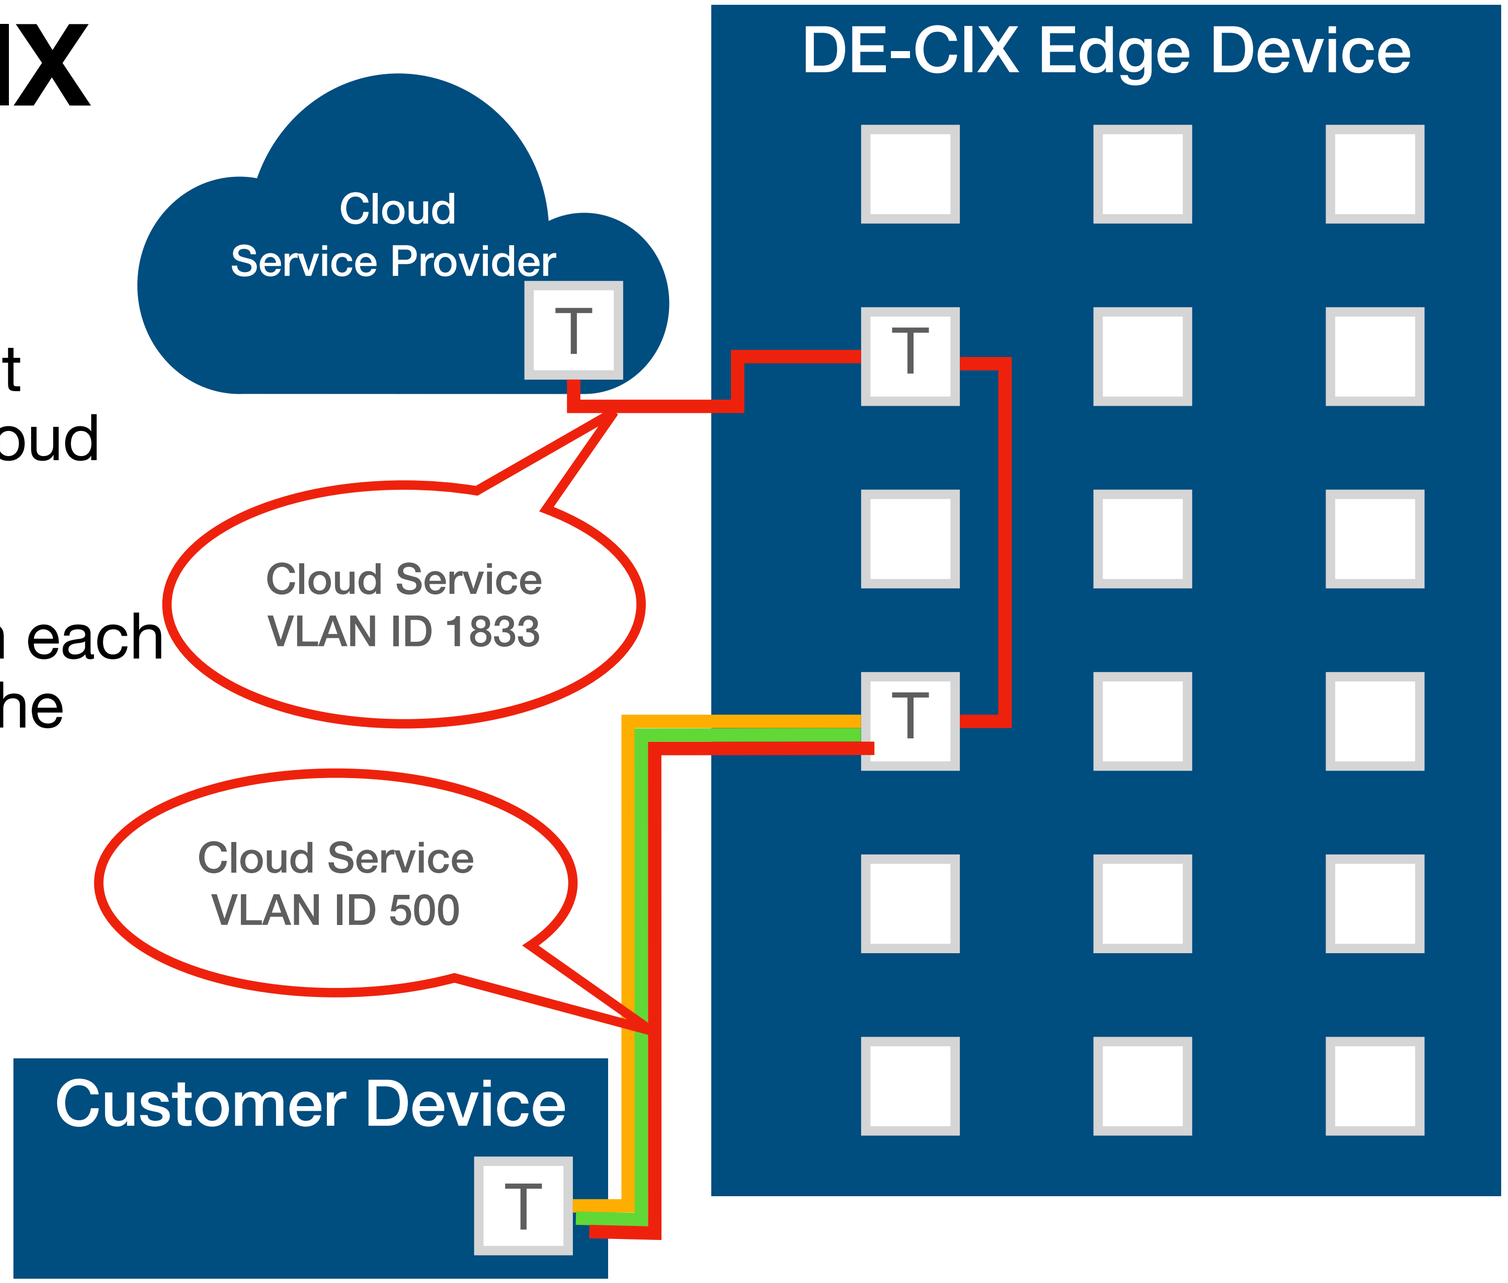

VLANs at DE-CIX

Connect Microsoft Azure

- The same way we connect customers to (multiple) Cloud service providers
- At DE-CIX the VLAN ID on each end does not have to be the same!
- Microsoft Azure Express Route needs some special configuration

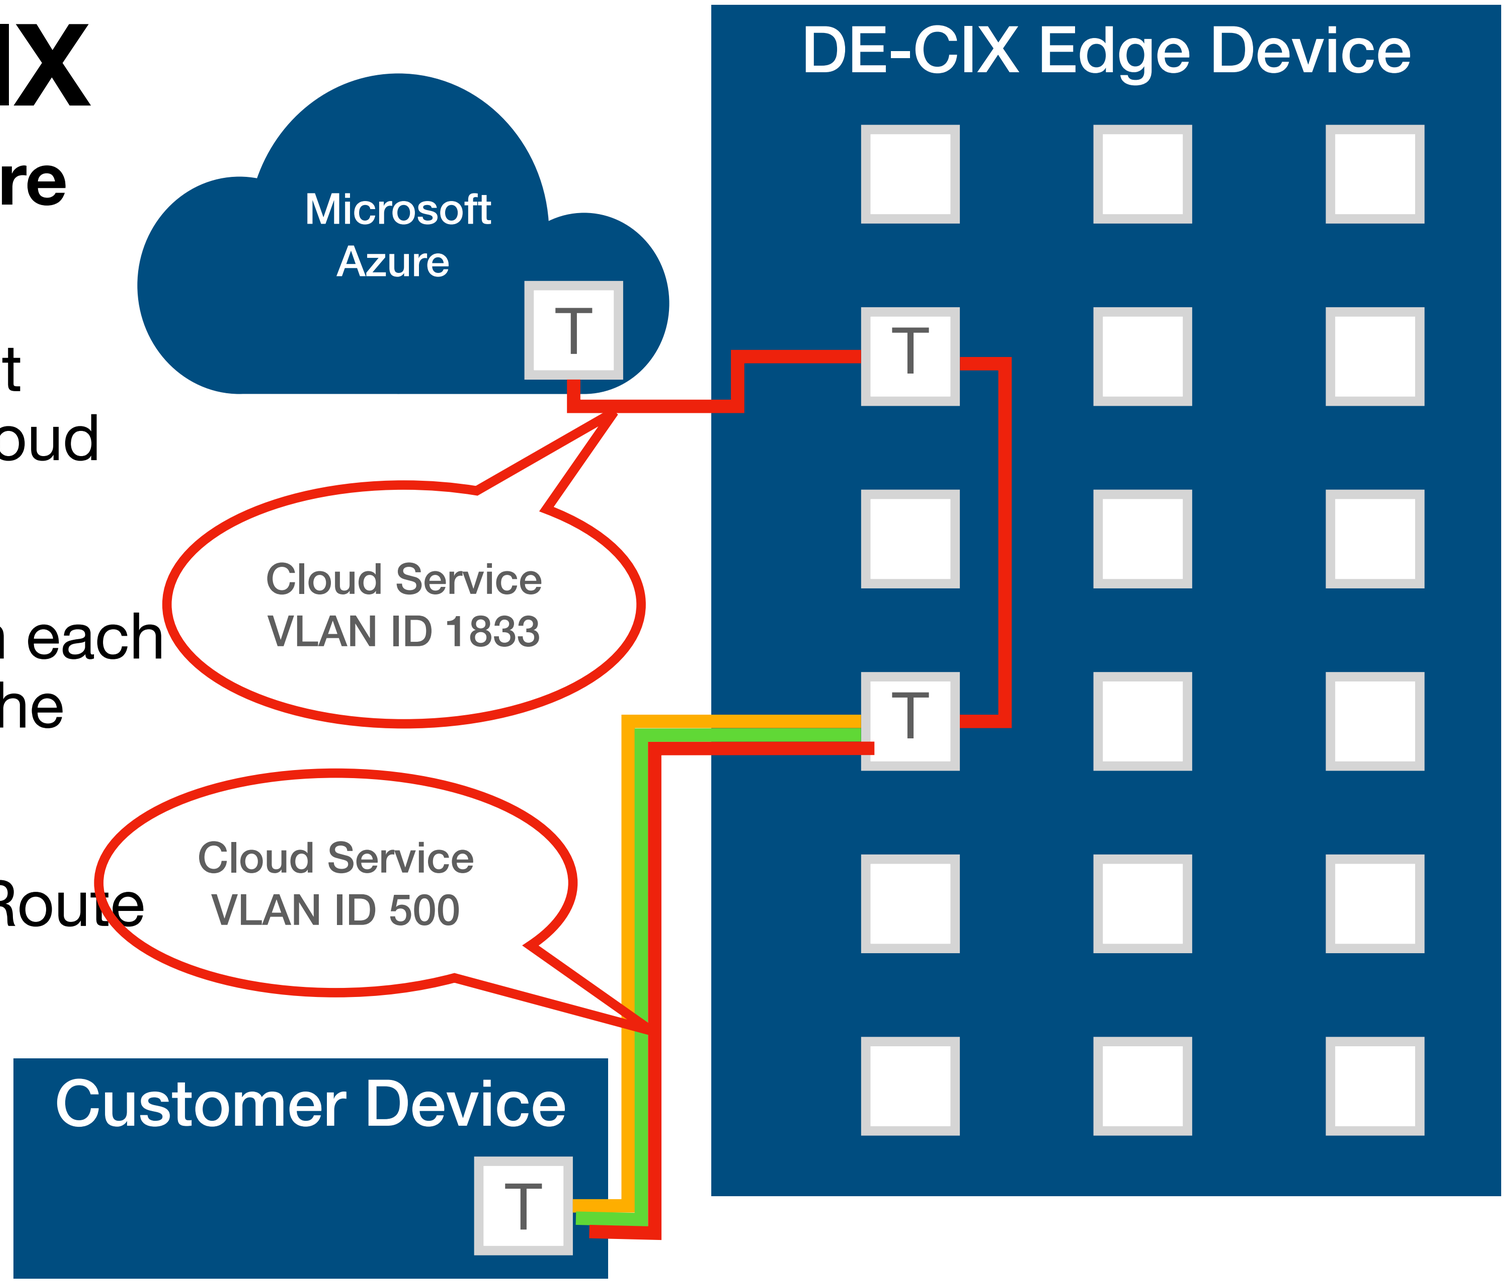

VLANs at DE-CIX

Microsoft Azure Express Route

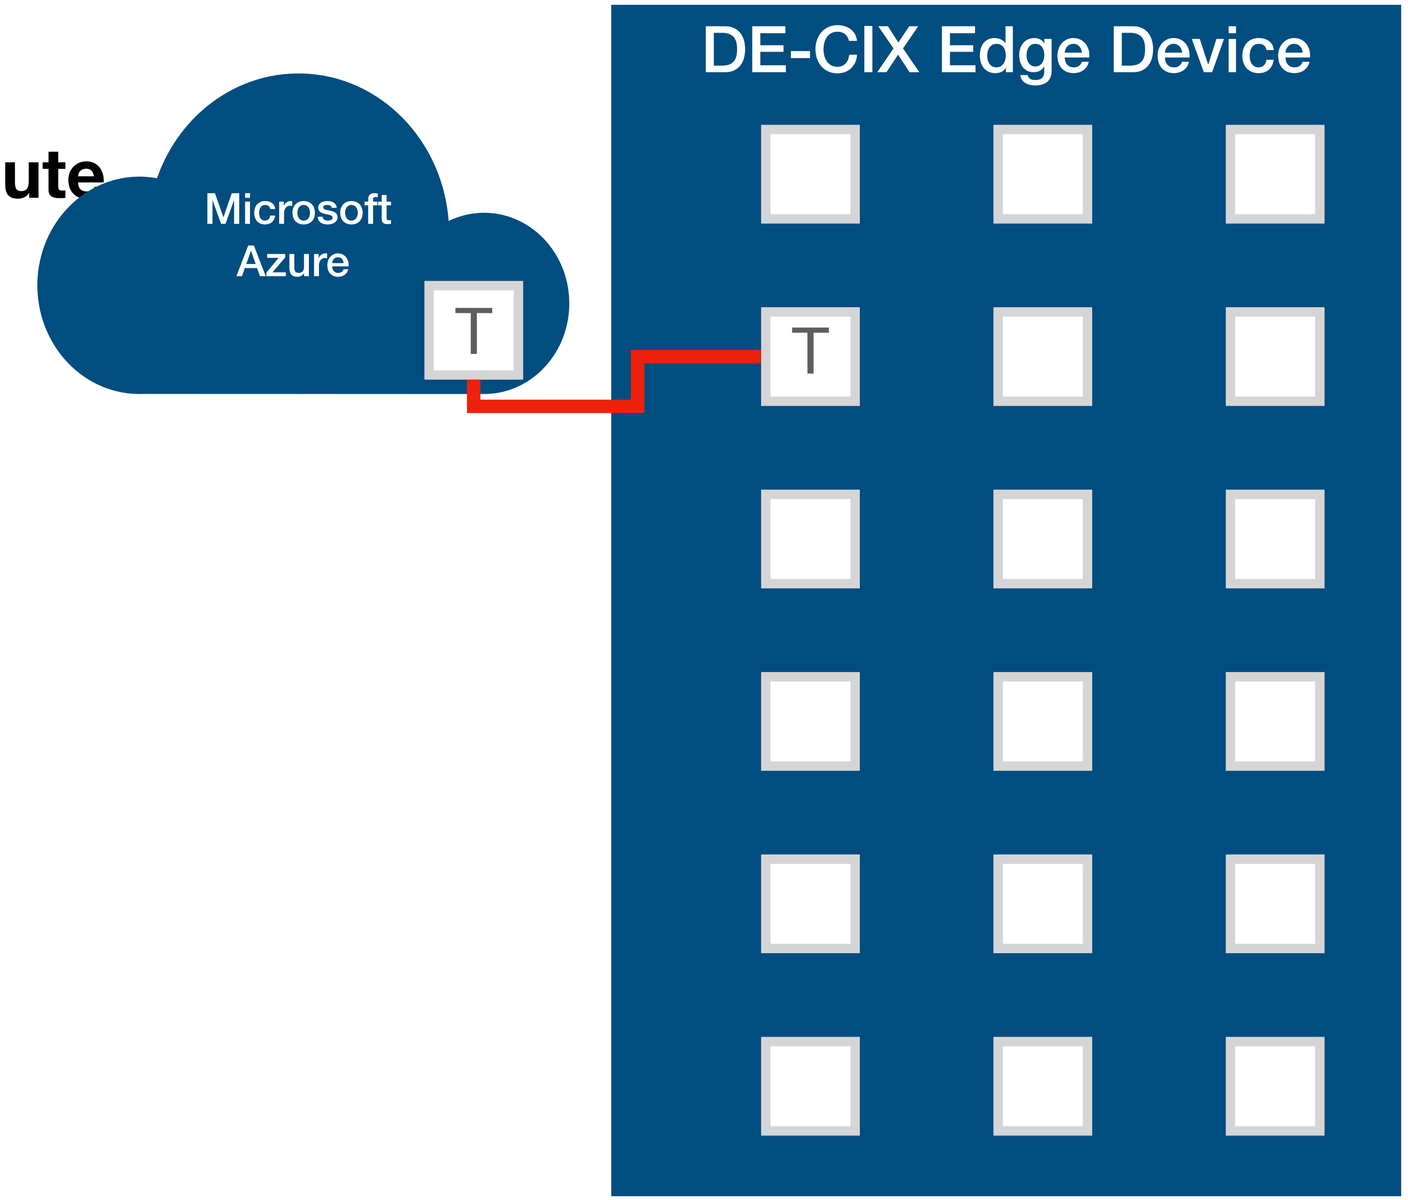

Microsoft Azure Express Route

Two redundant connections

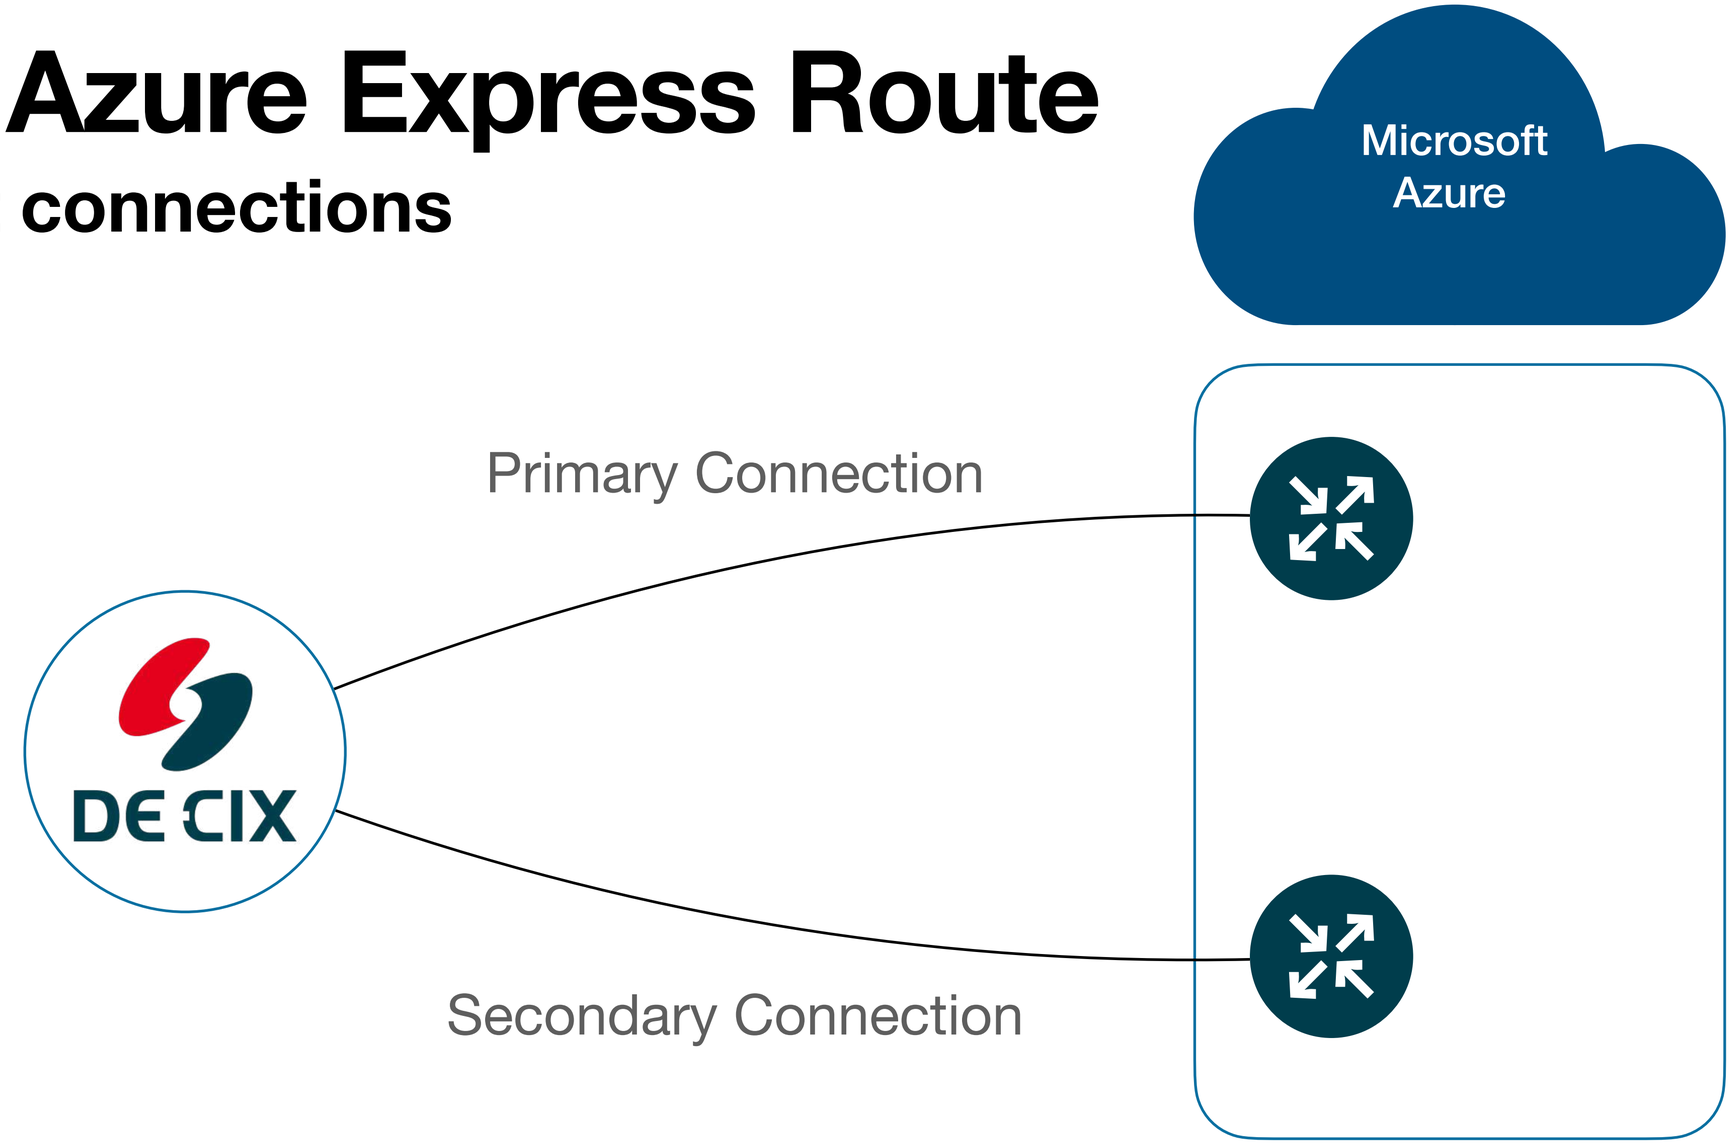

DE CIX

Microsoft Azure Express Route

On the customer side

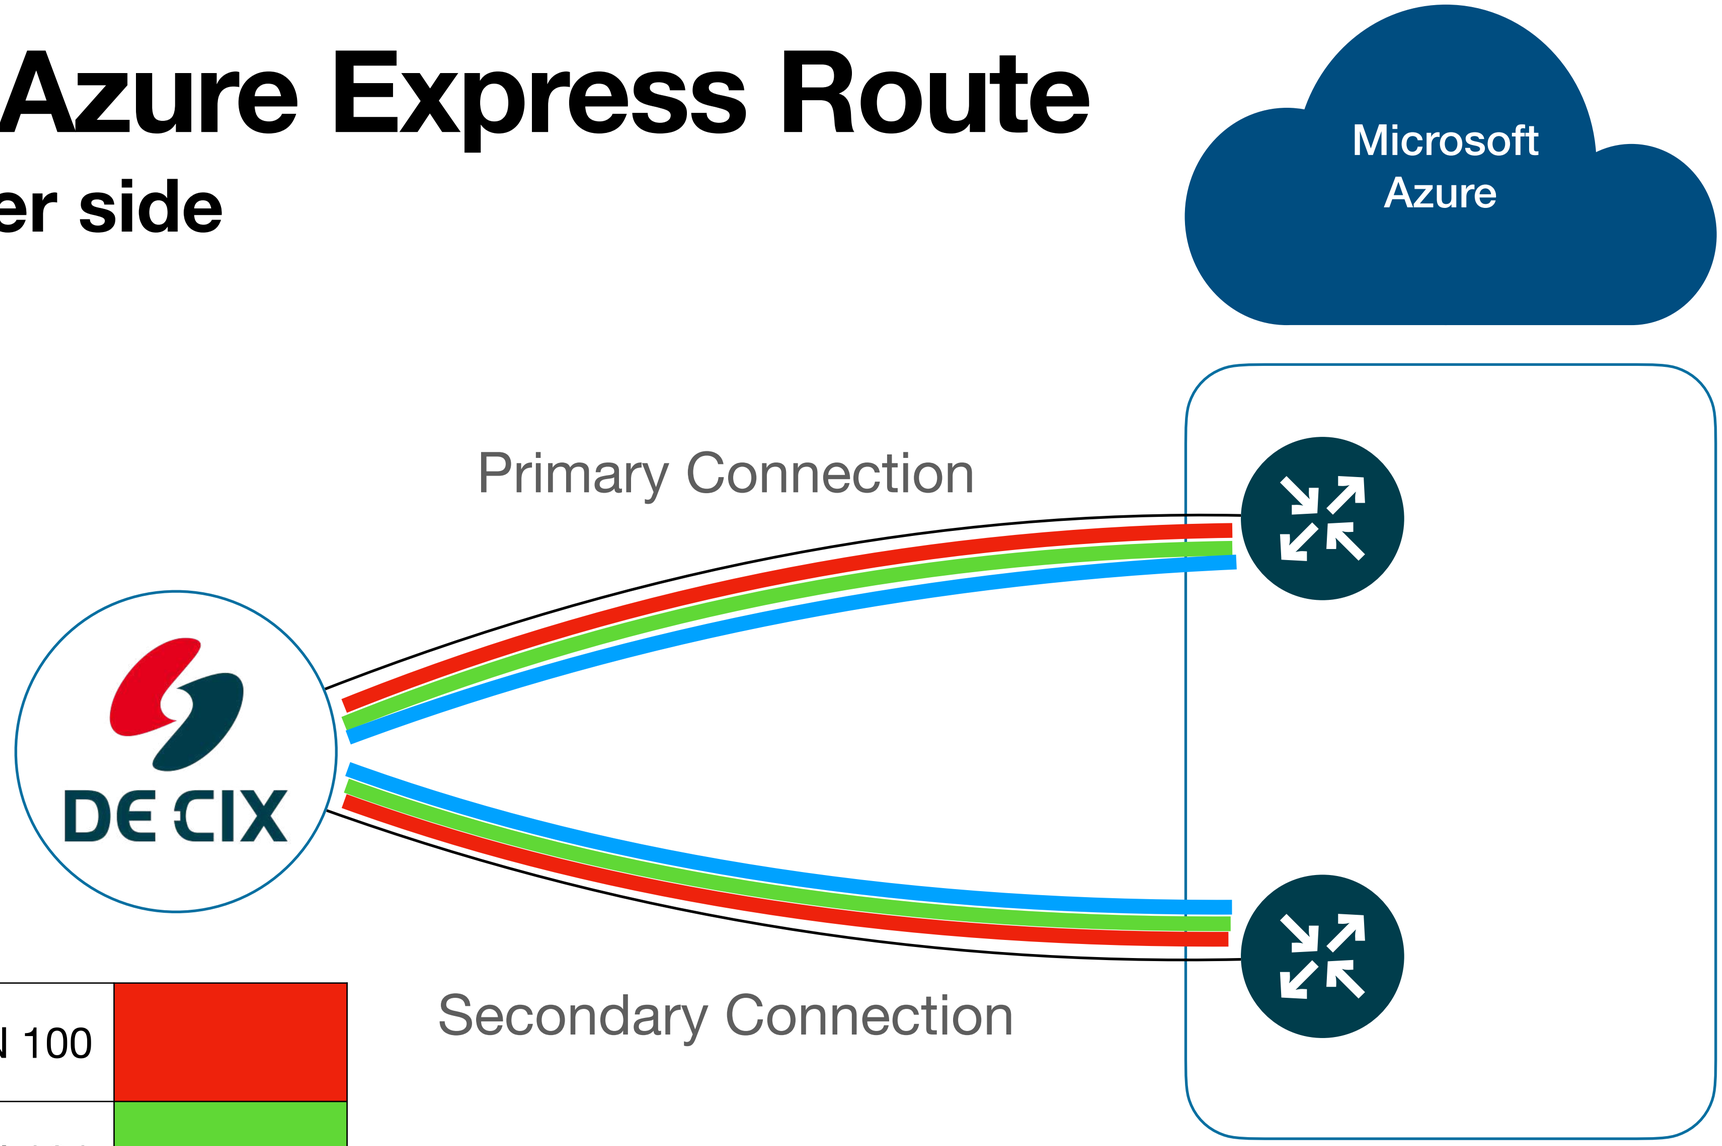

Customer A	VLAN 100	
Customer B	VLAN 101	
Customer C	VLAN 102	

DE CIX

Microsoft Azure Express Route

On the customer side

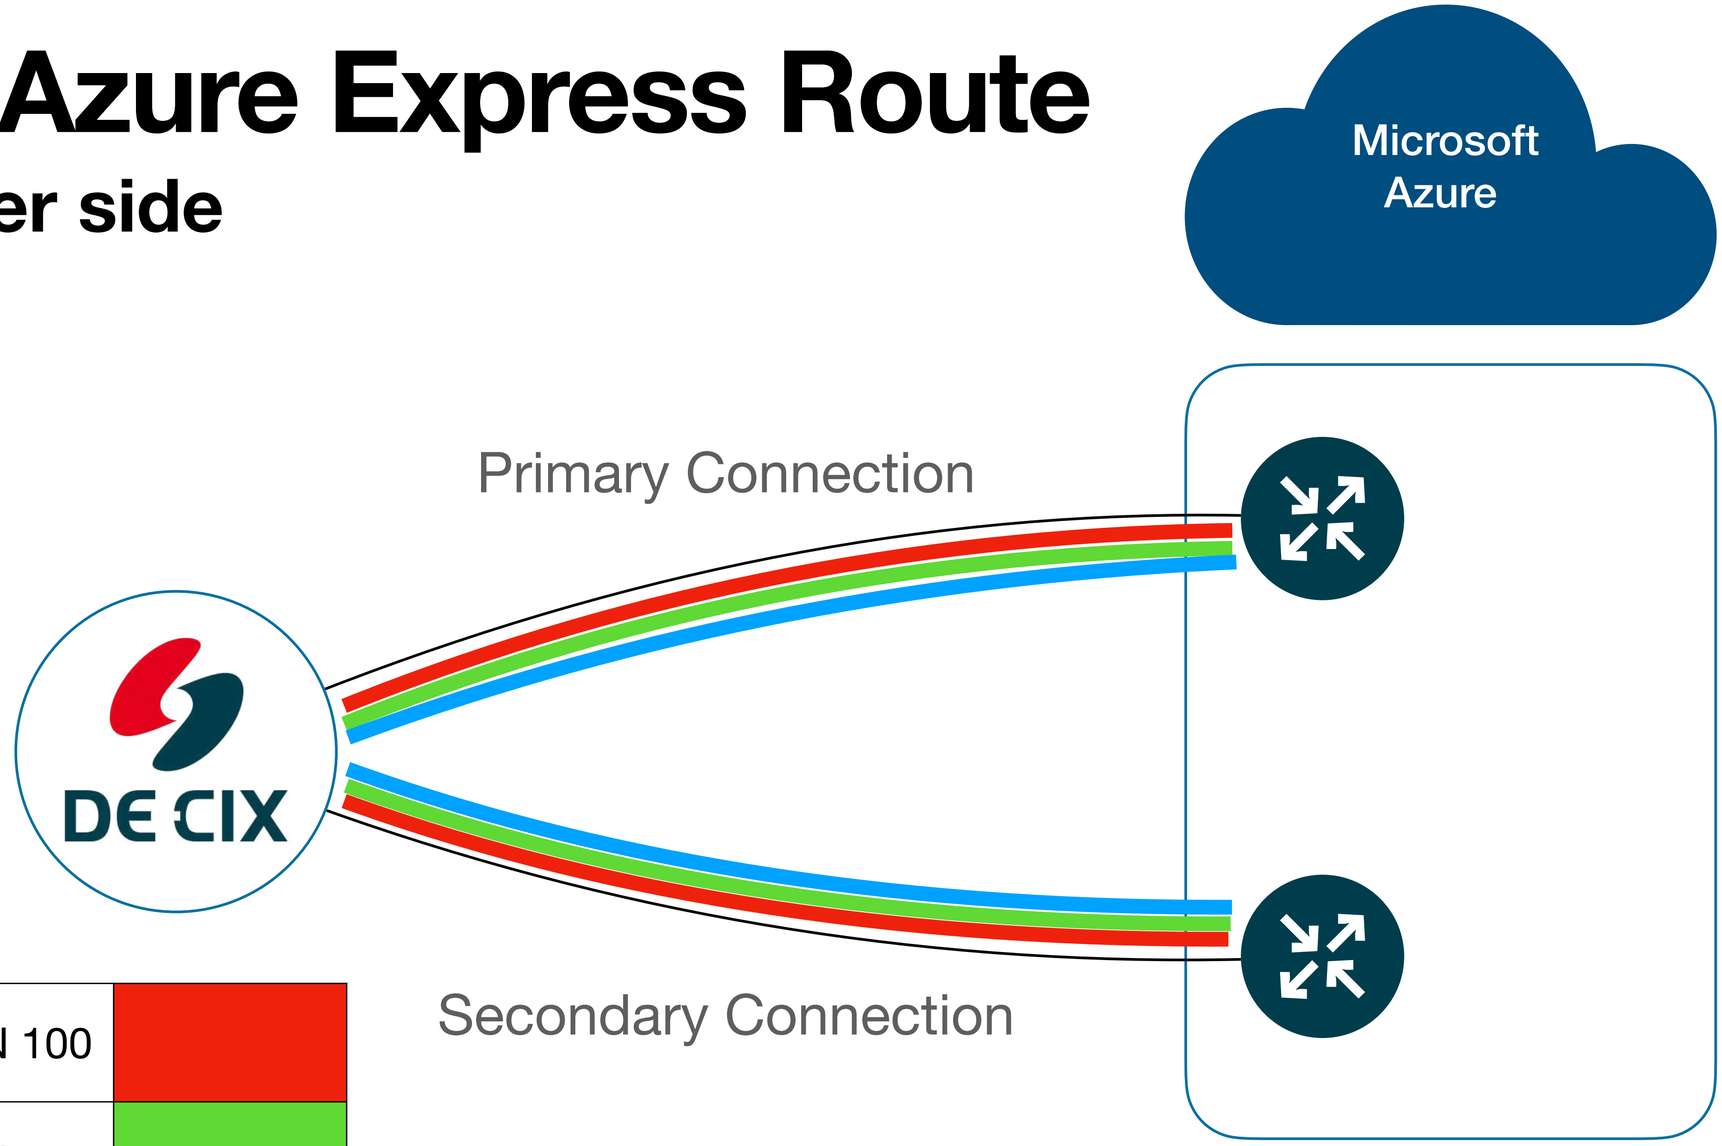

Customer A	VLAN 100	
Customer B	VLAN 101	
Customer C	VLAN 102	

DE CIX

Microsoft Azure Express Route

On the customer side

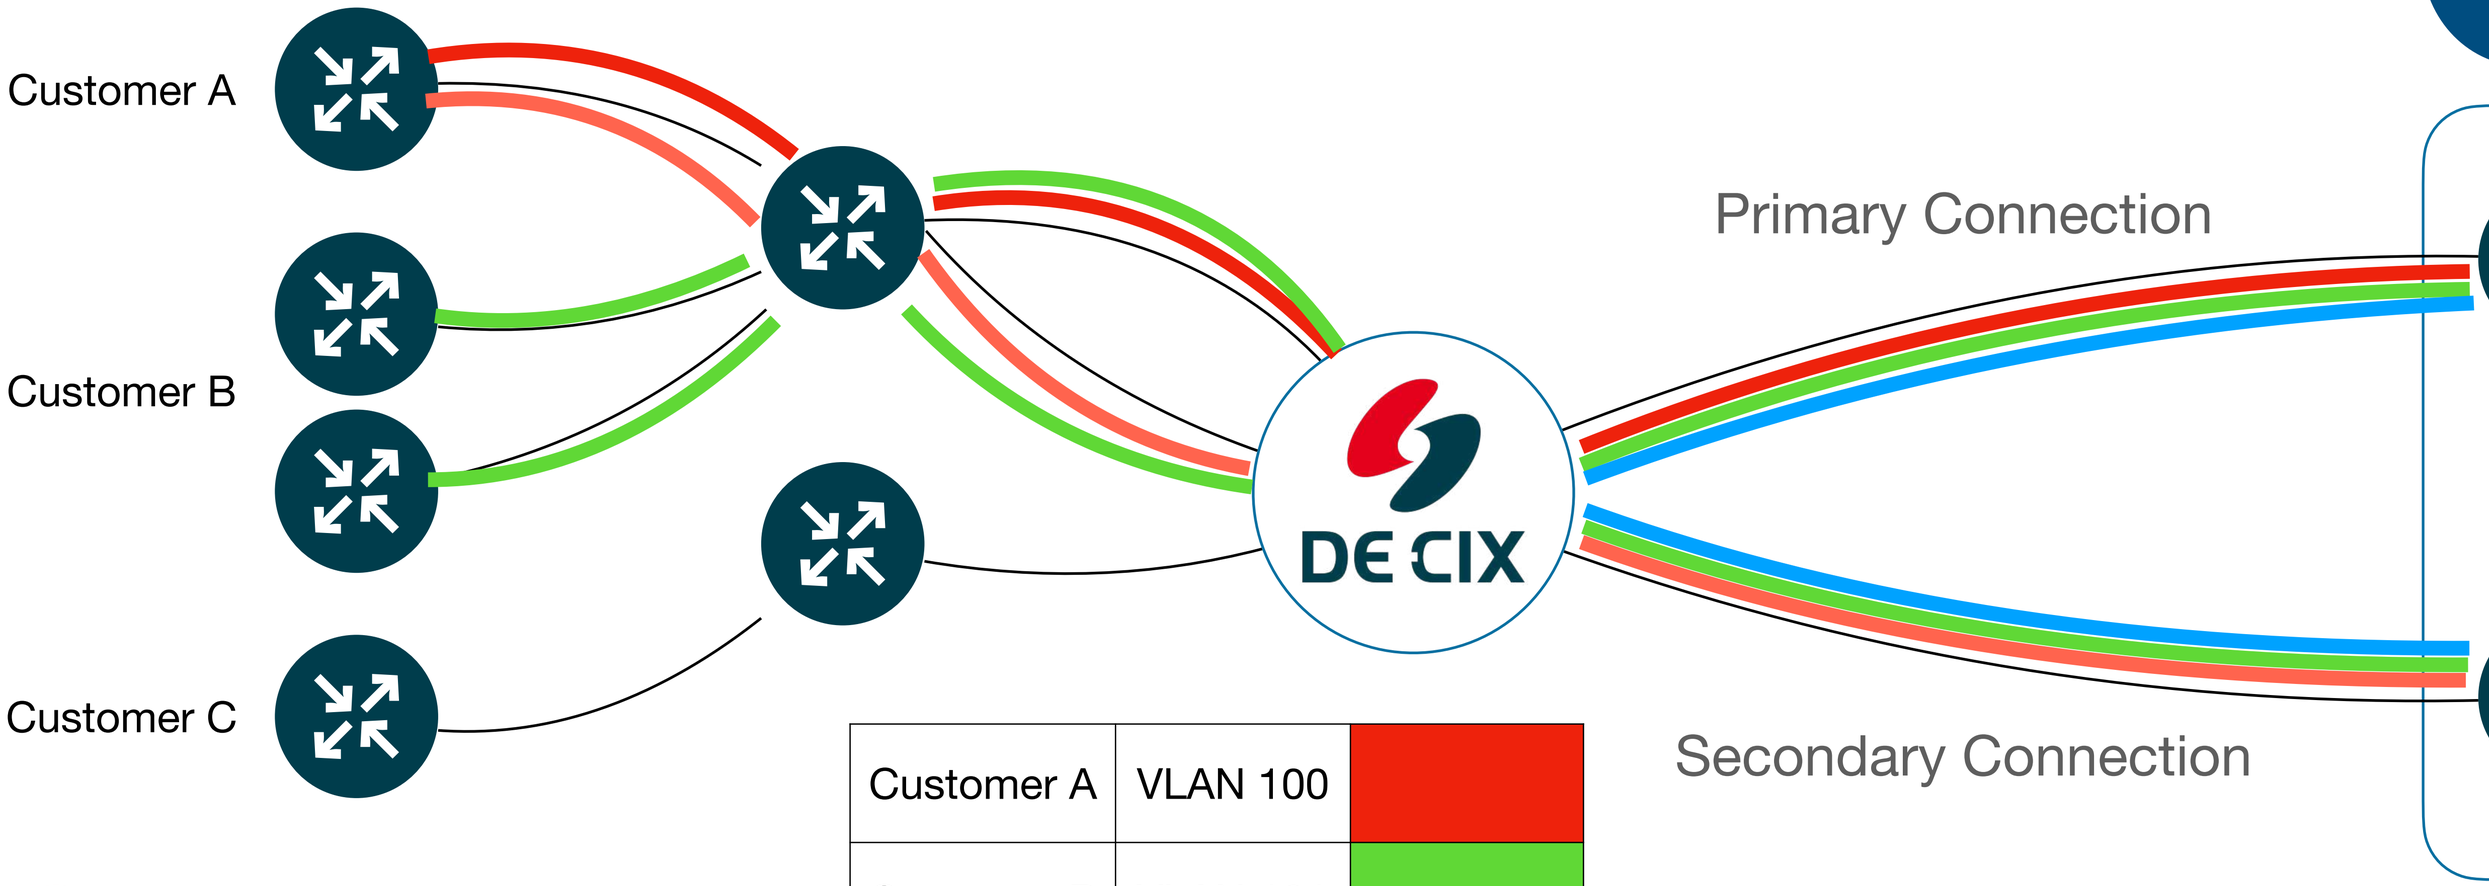

Customer A	VLAN 100	
Customer B	VLAN 101	
Customer C	VLAN 102	

DE CIX

Microsoft Azure Express Route

On the customer side

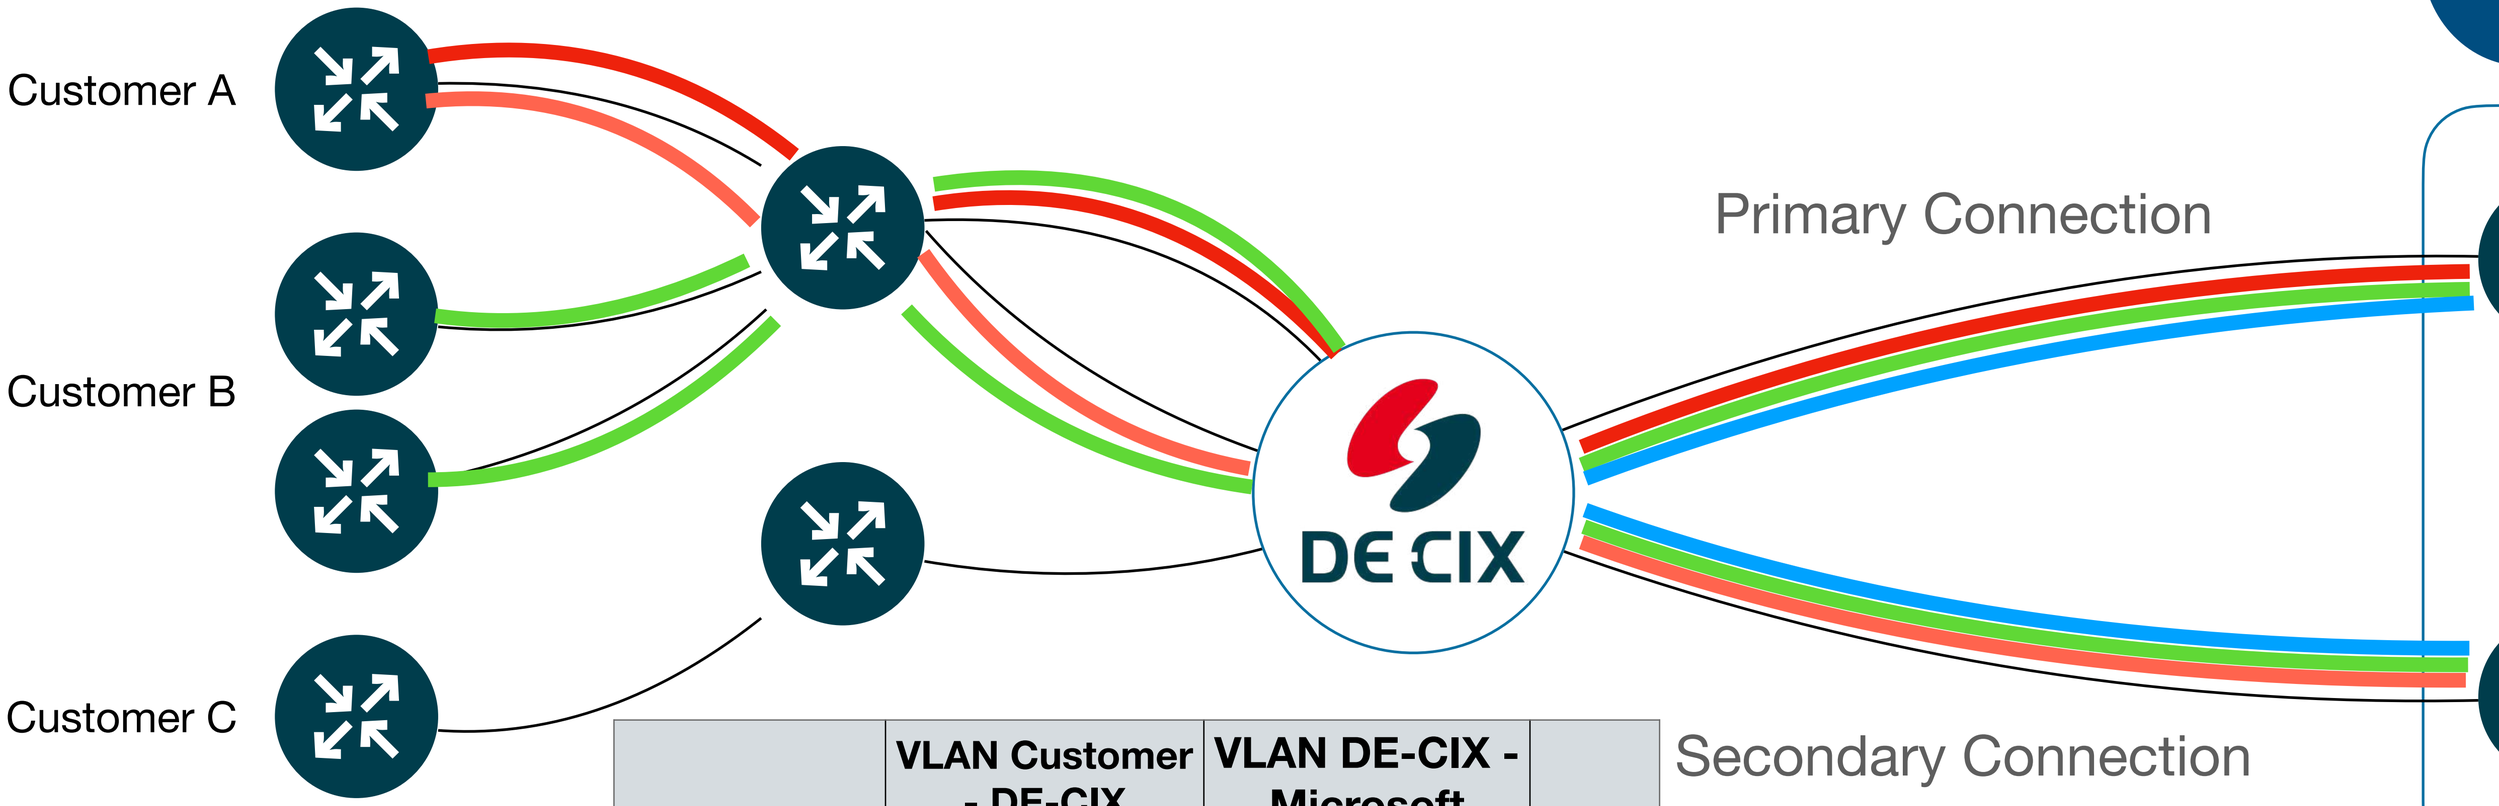

	VLAN Customer - DE-CIX	VLAN DE-CIX - Microsoft	
Customer A	VLAN 2334	VLAN 100	
	VLAN 2235		
Customer B	VLAN 17	VLAN 101	
Customer C	VLAN 551	VLAN 102	

DE CIX

Now, that all looks easy.

Doesn't it?

DE CIX

Multiple Services

Microsoft offers multiple services via Express Route

- Azure private
- Azure public
- both can be reached via one (redundant) connection (called Express Route)
- how to separate the two
  - (you might have guessed this)
  - VLANs!
  - To be more precise: a 2nd (inner) VLAN tag is used

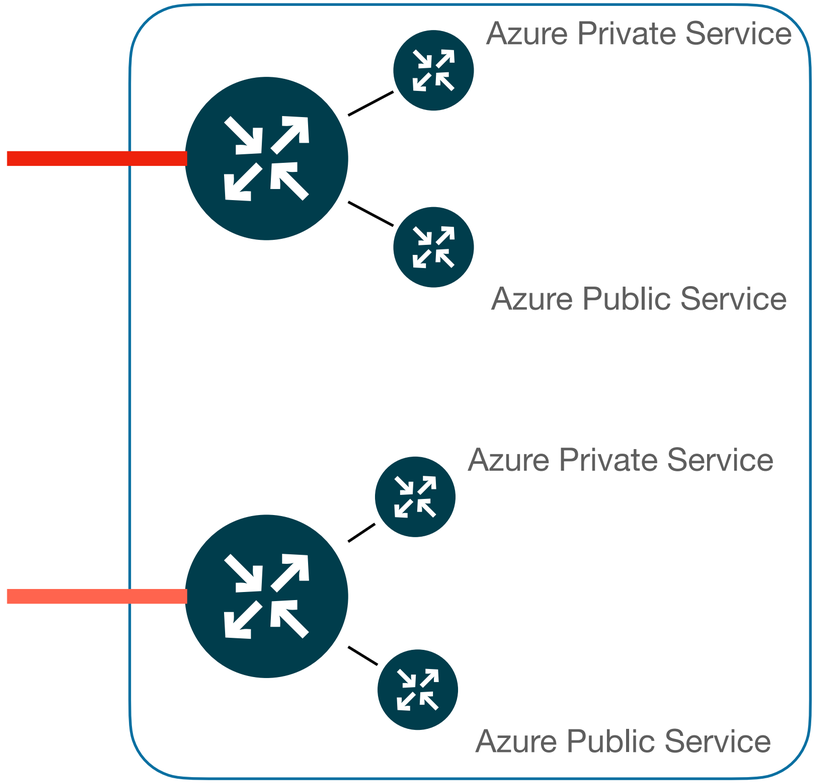

DE CIX

Multiple VLANs - one inside the other

...usually its only two

- The payload of Ethernet is IP
- You can add a VLAN tag to deliver multiple Ethernets on one physical port
- You can also add a 2nd VLAN tag which is only "seen" once the outer one has been processed

Attribution: Fanghong. derivative work: Greyhood
https://commons.wikimedia.org/wiki/File:Matryoshka_transparent.png

DE CIX

Ethernet Frame

VLAN tagged

- Ethertype values:

0x0800	IPv4
0x86dd	IPv6
0x0806	ARP
0x8100	VLAN Tagged

Preamble	SFD	Destination MAC Address	Source MAC Address	Ethertype	Payload	Checksum
10101010101010101010101010101011		48 Bits 6 Octets	48 Bits 6 Octets	16 Bits 2 Octets	46-1500 Octets	32 Bits 4 Octets

Preamble	SFD	Destination MAC Address	Source MAC Address	VLAN Header (801.1Q)		Ethertype	Payload	Checksum
10101010101010101010101010101011		48 Bits 6 Octets	48 Bits 6 Octets	0x8100	VLAN ID	16 Bits 2 Octets	42 - 1500 Octets	32 Bits 4 Octets

Ethernet Frame

Double tagged

0x0800	IPv4
0x86dd	IPv6
0x0806	ARP
0x8100	VLAN Tagged
0x88a8	Multiple VLAN Tagged

Preamble	SFD	Destination MAC Address	Source MAC Address	Ethertype	Payload	Checksum
10101010101010101010101010101011		48 Bits 6 Octets	48 Bits 6 Octets	16 Bits 2 Octets	46-1500 Octets	32 Bits 4 Octets

Preamble	SFD	Destination MAC Address	Source MAC Address	VLAN Header (801.1Q)		Ethertype	Payload	Checksum
10101010101010101010101010101011		48 Bits 6 Octets	48 Bits 6 Octets	0x8100	VLAN ID	16 Bits 2 Octets	42 - 1500 Octets	32 Bits 4 Octets

Preamble	SFD	Destination MAC Address	Source MAC Address	VLAN Header (outer)		VLAN Header (inner)		Ethertype	Payload	
10101010101010101010101010101011		48 Bits 6 Octets	48 Bits 6 Octets	0x88a8	VLAN ID	0x8100	VLAN ID	16 Bits 2 Octets	42 - 1500 Octets	

How to double-tag?

DE CIX

Customer side

Directly connected to DE-CIX

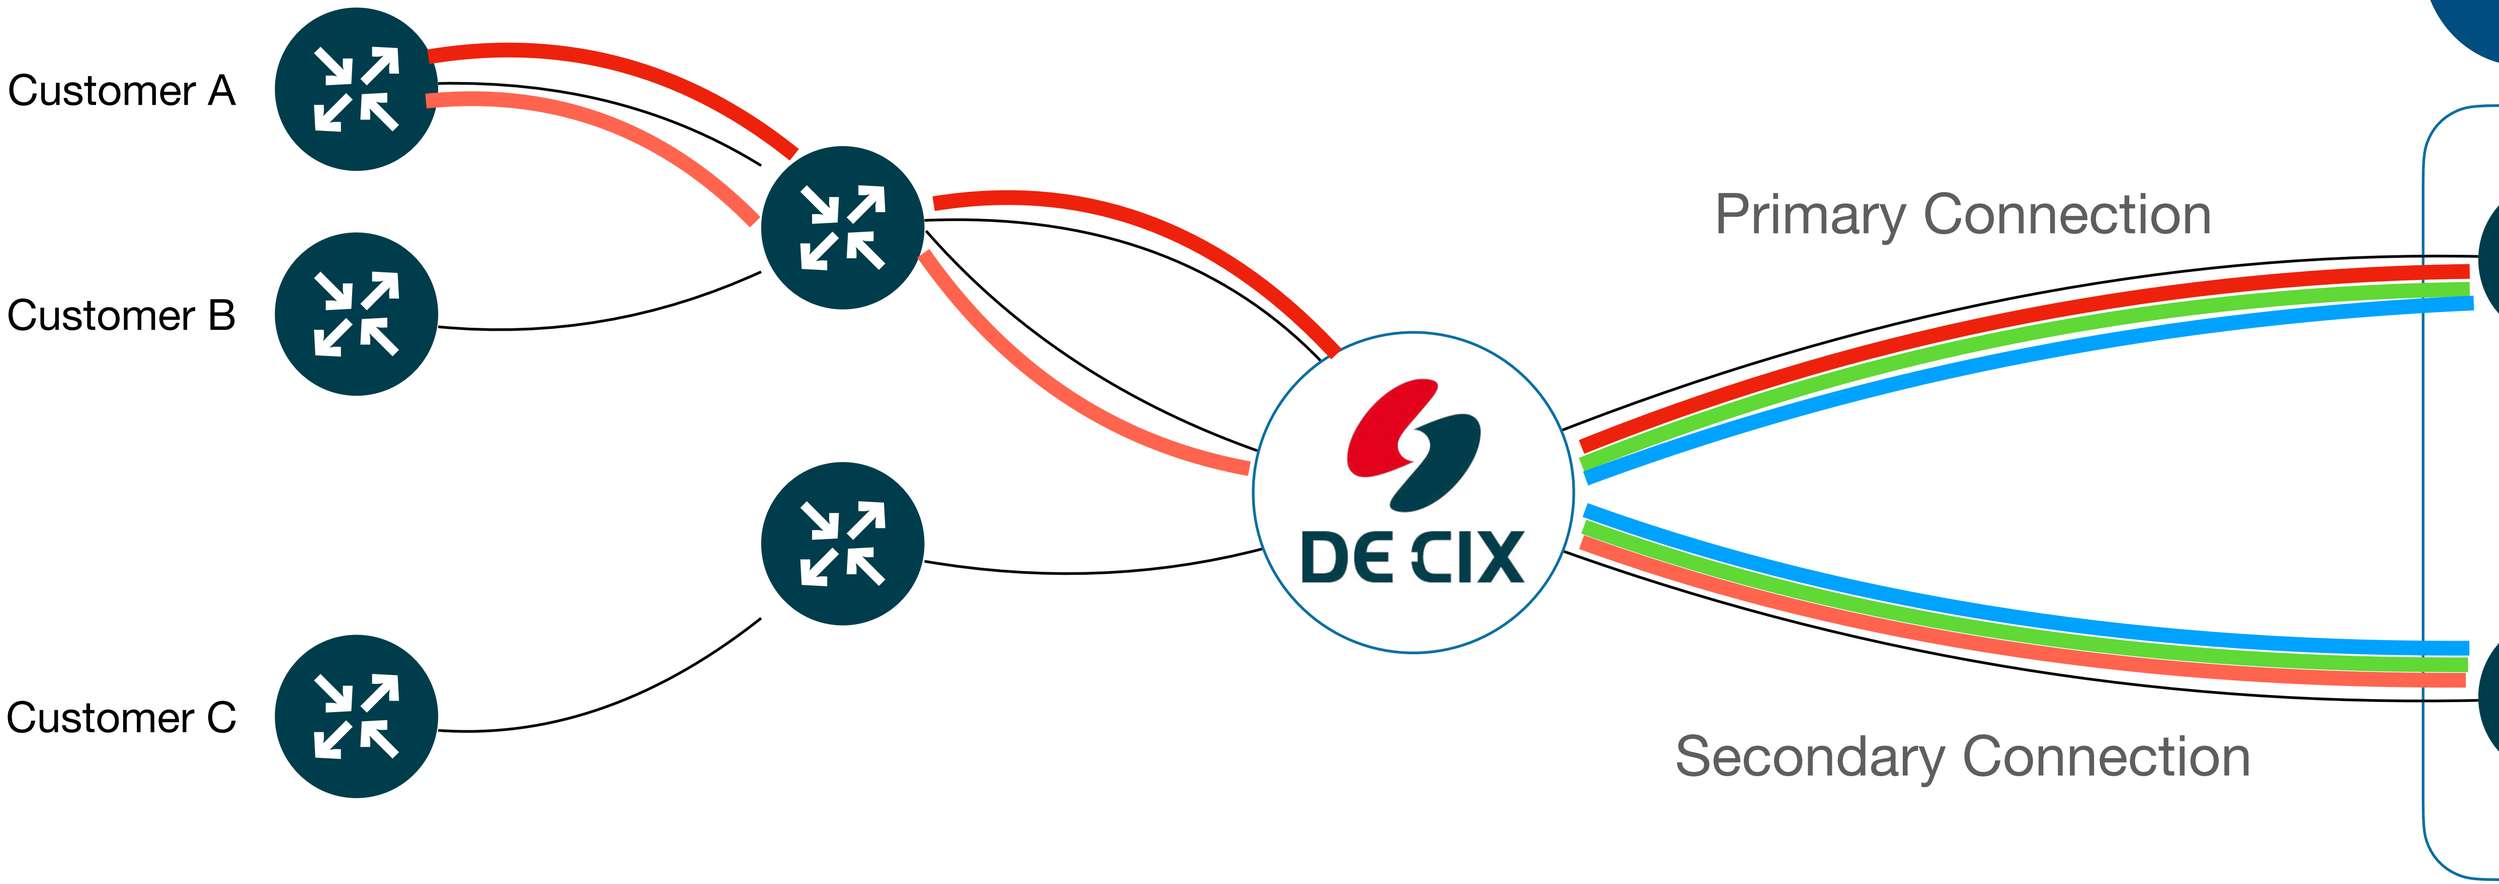

DE CIX

Customer side

Directly connected to DE-CIX

- All modern routers can do double-tagging
- Customer has to set ***both tags*** on the router
- Example for Cisco IOS XR

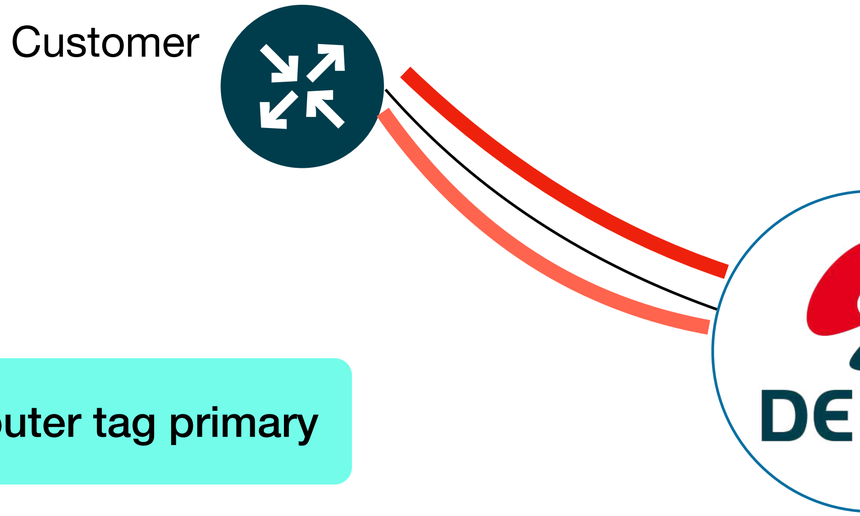

outer tag primary

inner tag (same for both)

MTU must be large enough

outer tag secondary

```
interface TenGigE0/0/2 0
 description to DE-CIX Platform
 mtu 9000
!
interface TenGigE0/0/2/0.36
 description to MS Azure Private Peerin  Primary
 ipv4 address 192.168.100.1 255.255.255.252
 ipv6 address fdb1:6268:1455:100::1/126
 encapsulation dot1q 36 second-dot1q 11
!
interface TenGigE0/0/2/0.37
 description to MS Azure Private Peering Secondary
 ipv4 address 192.168.100.5 255.255.255.252
 ipv6 addr    fdb1:6268:1455:100::5/126
 encapsulation dot1q 37 second-dot1q 11
```

DE CIX

Customer side

Directly connected to DE-CIX

- All modern routers can do double-tagging
- Customer has to set ***both tags*** on the router
- Example for Juniper

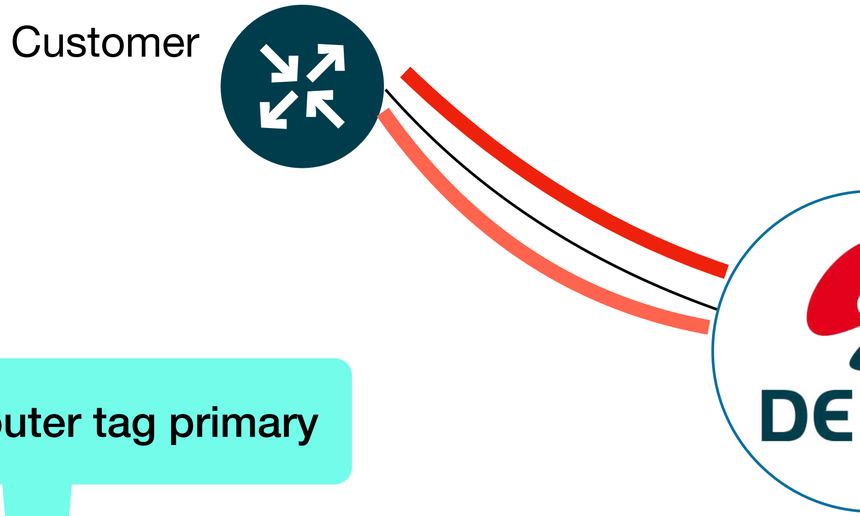

outer tag primary

inner tag (same for both)

MTU must be large enough

```
ge-0/0/1 {
    description "to MSAzurePrivatePeering"
    flexible-vlan-tagging
    mtu 9000;
    unit 0 {
        vlan-tags outer 36 inner 11;
        family inet {
            address 192.168.100.1/30;
        }
        family inet6 {
            address fdb1:6268:1455:100::1/126;
        }
    }
}
```

DE-CIX

Customer side

Connected via an ISP

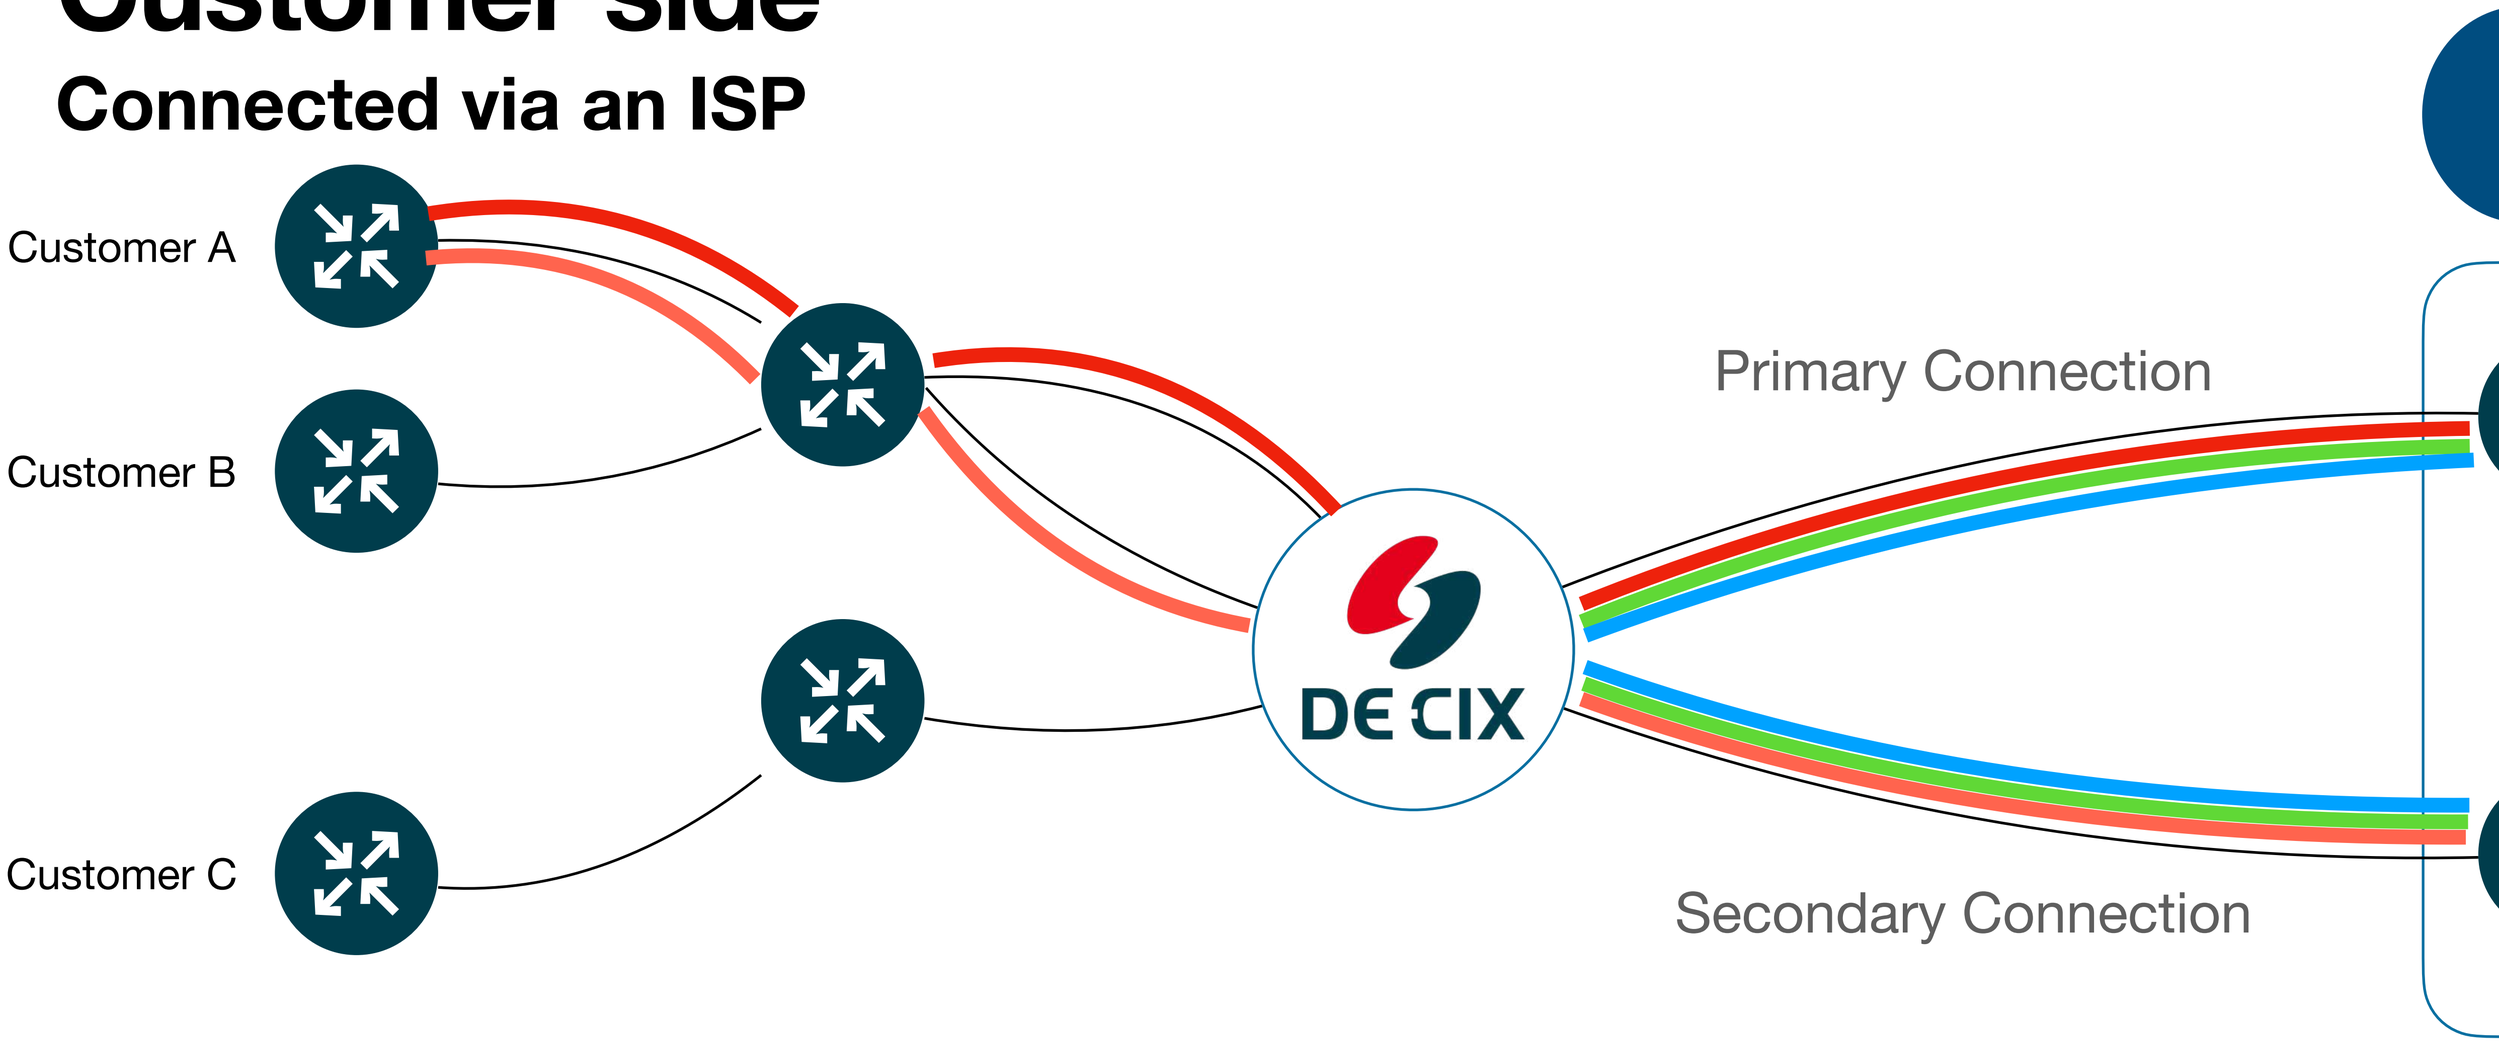

DE CIX

Customer side

Connected via an ISP

- Customer connects via an ISP to DE-CIX
- Here several options exist
- Here: ISP takes care of the outer VLAN(s)
  - Customer single-tags only the inner VLAN
  - For the required redundancy, customer needs two routers

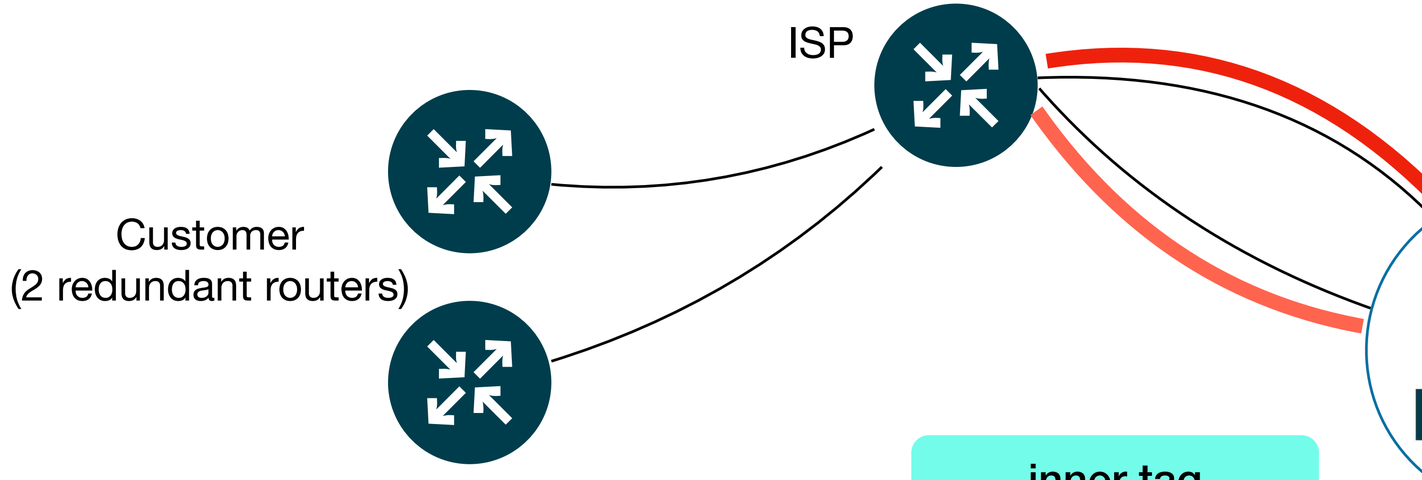

MTU must be large enough

inner tag (same on both routers)

```
interface TenGigE0/0/2/0
 description to ISP
 mtu 9000
!
interface TenGigE0/0/2/0.11
 description to MS Azure Private Peering Primary
 ipv4 address 192.168.100.1 255.255.255.252
 ipv6 address fdb1:6268:1455:100::1/126
 encapsulation dot1q 11
```

DE-CIX

Conclusion

Please remember....

Facts about double tagged VLANs

- Ethernet is a **broadcast** network
- VLANs set up **virtual LANs** on a **common physical infrastructure**
  - VLAN **IDs** run from 1 - 4094
  - DE-CIX uses VLANs for multiple service delivery on one physical port
- You can **double-tag** VLANs to deliver multiple services inside a VLAN
  - like Microsoft does with Express Route

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 used in the presentation

Ethernet today

- Ethernet
  - Wikipedia entry for Ethernet
  - IEEE Standard for Ethernet
- Various types of Ethernet
  - 10Base5
  - 10Base2
  - 10Base-T
- more speed
  - FastEthernet - 100Mbit/s
  - GigabitEthernet - 1000Mbit/s / 1GBit/s
  - 10 Gigabit Ethernet - 10GBit/s
  - 100 Gigabit Ethernet (and 40 Gigabit Ethernet)
- Currently used hardware
  - Twisted pair cables: Cat5, Cat6, RJ45 connector
  - Optical fibres: Single-mode and multi-mode
  - Ethernet switch

VLANs

- Wikipedia entry for
  - VLANs: https://en.wikipedia.org/wiki/Virtual_LAN
  - IEEE 802.1Q (VLAN standard): https://en.wikipedia.org/wiki/IEEE_802.1Q
  - Q-in-Q, IEEE 802.1ad: https://en.wikipedia.org/wiki/IEEE_802.1ad
  - Private VLAN (port isolation): https://en.wikipedia.org/wiki/Private_VLAN
- Some RFCs (Request for comment = Internet standards) about VLANs:
  - RFC3069 VLAN Aggregation for Efficient IP Address Allocation
  - RFC4554 Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks
- IEEE Standards (may not be freely available):
  - IEEE 802.1Q-2014: https://ieeexplore.ieee.org/servlet/opac?punumber=6991460
  - IEEE 802.1ad: http://www.ieee802.org/1/pages/802.1ad.html

Standards

- IEEE standards
  - 802.3-2018 current standard, also here
  - IEEE 802 committee website

- Registered information:
  Ethertype list at IANA, Public register at IEEE

- Some Internet RFCs regarding Ethernet
  - IP over Ethernet: RFC894, RFC895
  - IPv6 over Ethernet: RFC1972, RFC2464