The Bits Must Flow

(net)Working through the abstractions

What happens when you visit a website?

Start with a classic Audience Participation

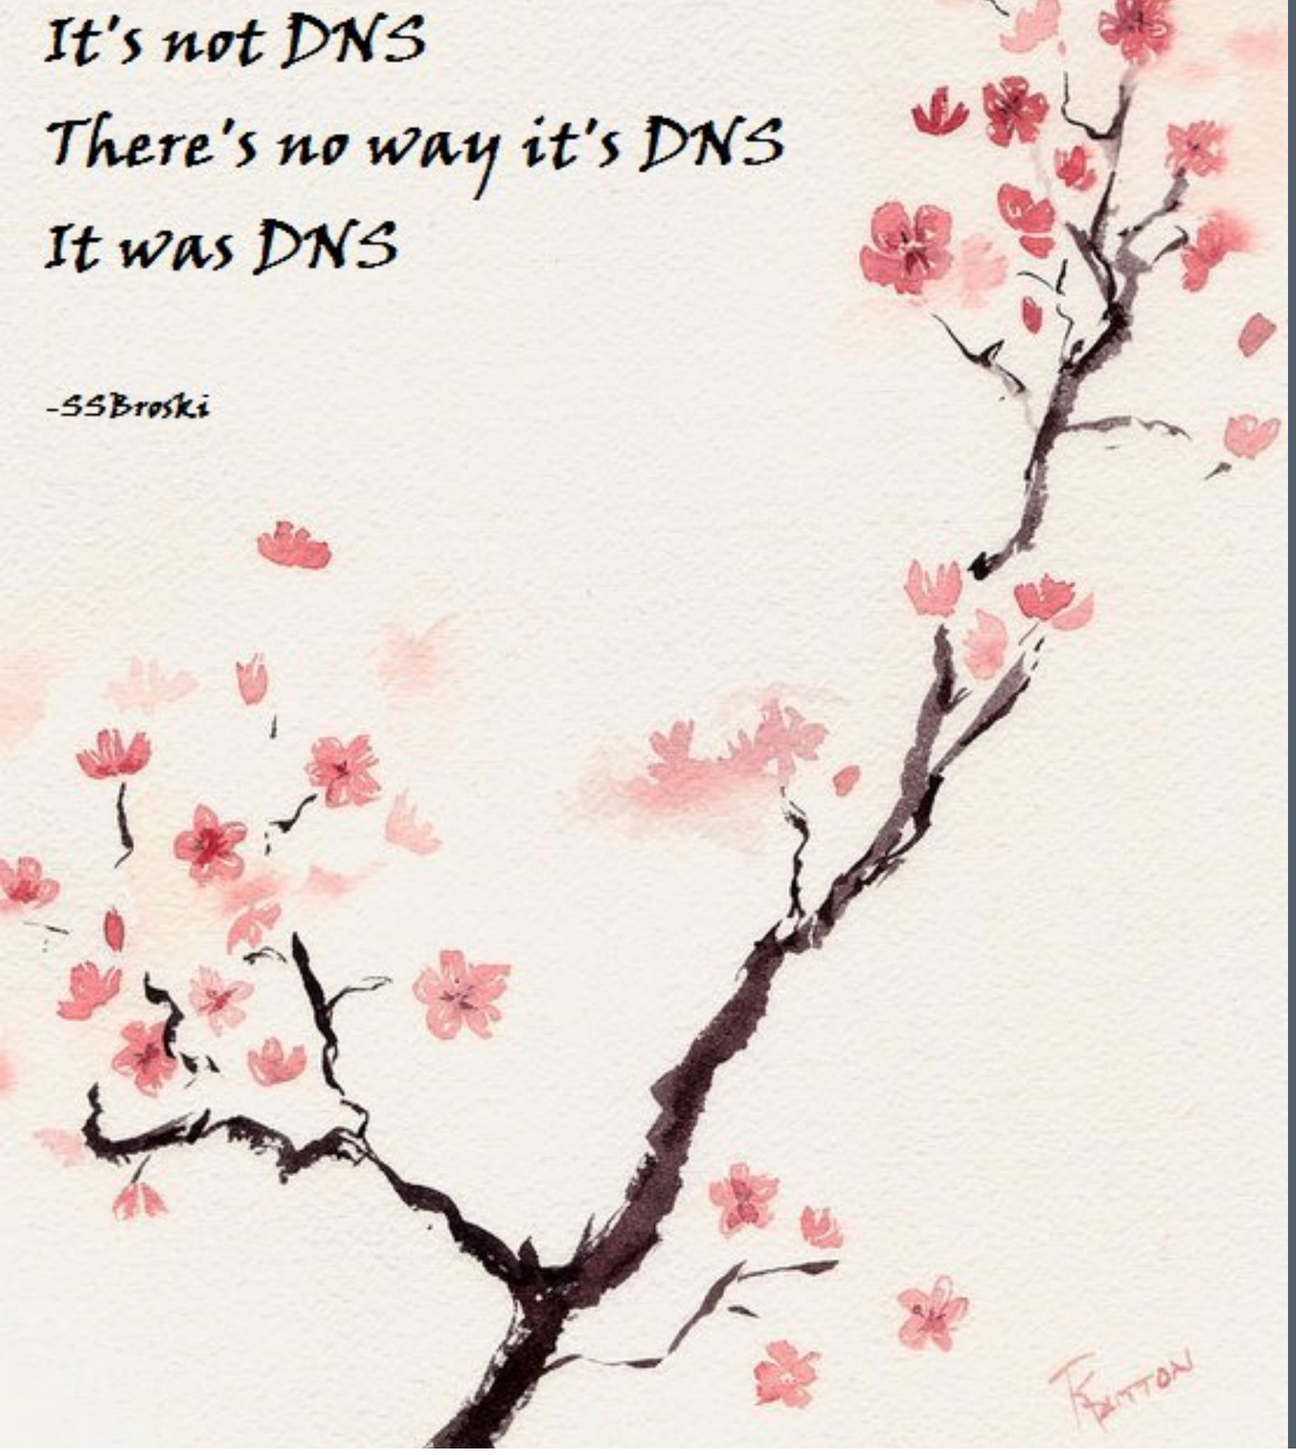

Why DNS?

Internet Traffic Per Second

167,069GB[1]

[1] https://interenetlivestats.com/one-second/#traffic-band

1,336,544,000,000,000

bits per second

OSI 7-Layer model

Application
Presentation
Session
Transport
Network
Data Link
Physical

7 — @CrayZeigh

OSI	TCP/IP
Application	Application
Presentation	
Session	
Transport	Transport
Network	Internet
Data Link	Network Access
Physical	

Why all these layers anyway?

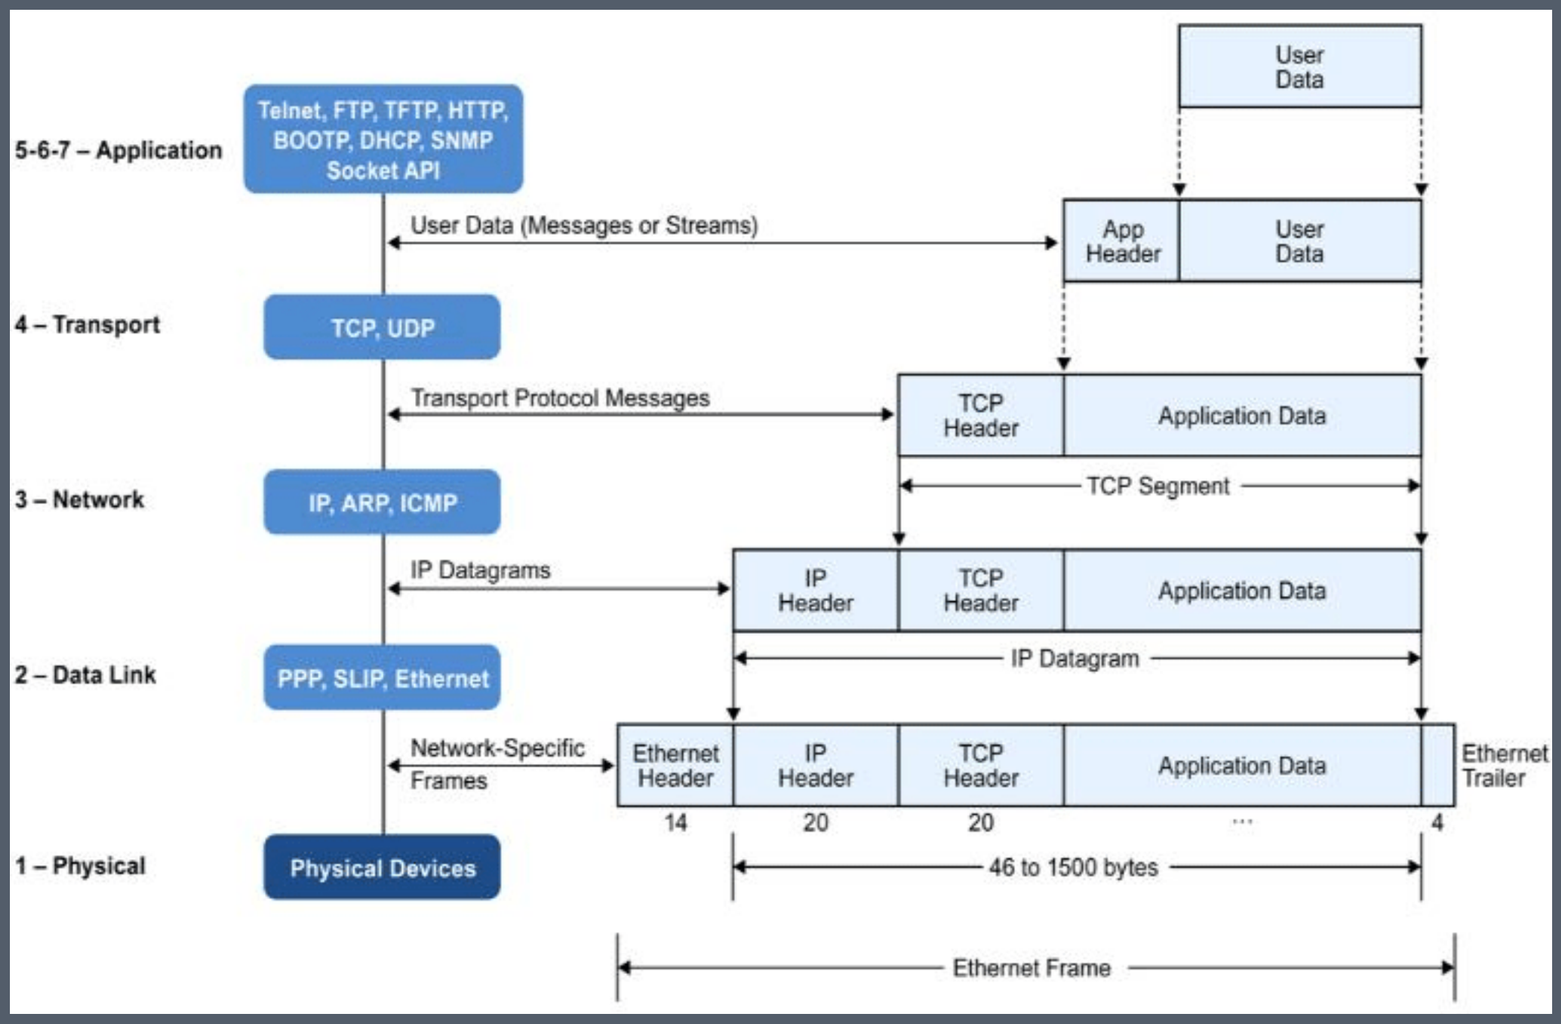

Network Access:

Data Frames help translate digital to physical

Preamble	SFD	Destination MAC	Source MAC	Type	Data and Pad	FCS
7 Bytes	1 Byte	6 Bytes	6 Bytes	2 Bytes	46-1500 Bytes	4 Bytes

MAC Addresses

01:23:45:67:89:ab

but how do you figure out what your destination MAC address is?

MAC Addresses

01:23:45:67:89:ab

→ Identifies the (network) device

but how do you figure out what your destination MAC address is?

MAC Addresses

01:23:45:67:89:ab

→ Identifies the (network) device

→ For same-network devices

but how do you figure out what your destination MAC address is?

ARP²

[2] Address Resolution Protocol

arp -a

need a way to separate traffic

ARP[2]

→ Mapping IPs and MAC addresses

[2] Address Resolution Protocol

arp -a

need a way to separate traffic

ARP[2]

→ Mapping IPs and MAC addresses

→ Broadcasts to find neighbors

[2] Address Resolution Protocol

arp -a

need a way to separate traffic

Broadcast & Switchs v Hubs

VLANs

→ Limiting Broadcast Domains

Broadcast & Switchs v Hubs

VLANs

→ Limiting Broadcast Domains

→ IEEE 802.1q

Broadcast & Switchs v Hubs

VLANs

→ Limiting Broadcast Domains

→ IEEE 802.1q

→ up to 4096 VLANs[3]

[3] VXLAN addresses this but that's A Whole Other Thing

Broadcast & Switchs v Hubs

VLANs

→ Limiting Broadcast Domains

→ IEEE 802.1q

→ up to 4096 VLANs[3]

→ Native or Tagged

[3] VXLAN addresses this but that's A Whole Other Thing

Broadcast & Switchs v Hubs

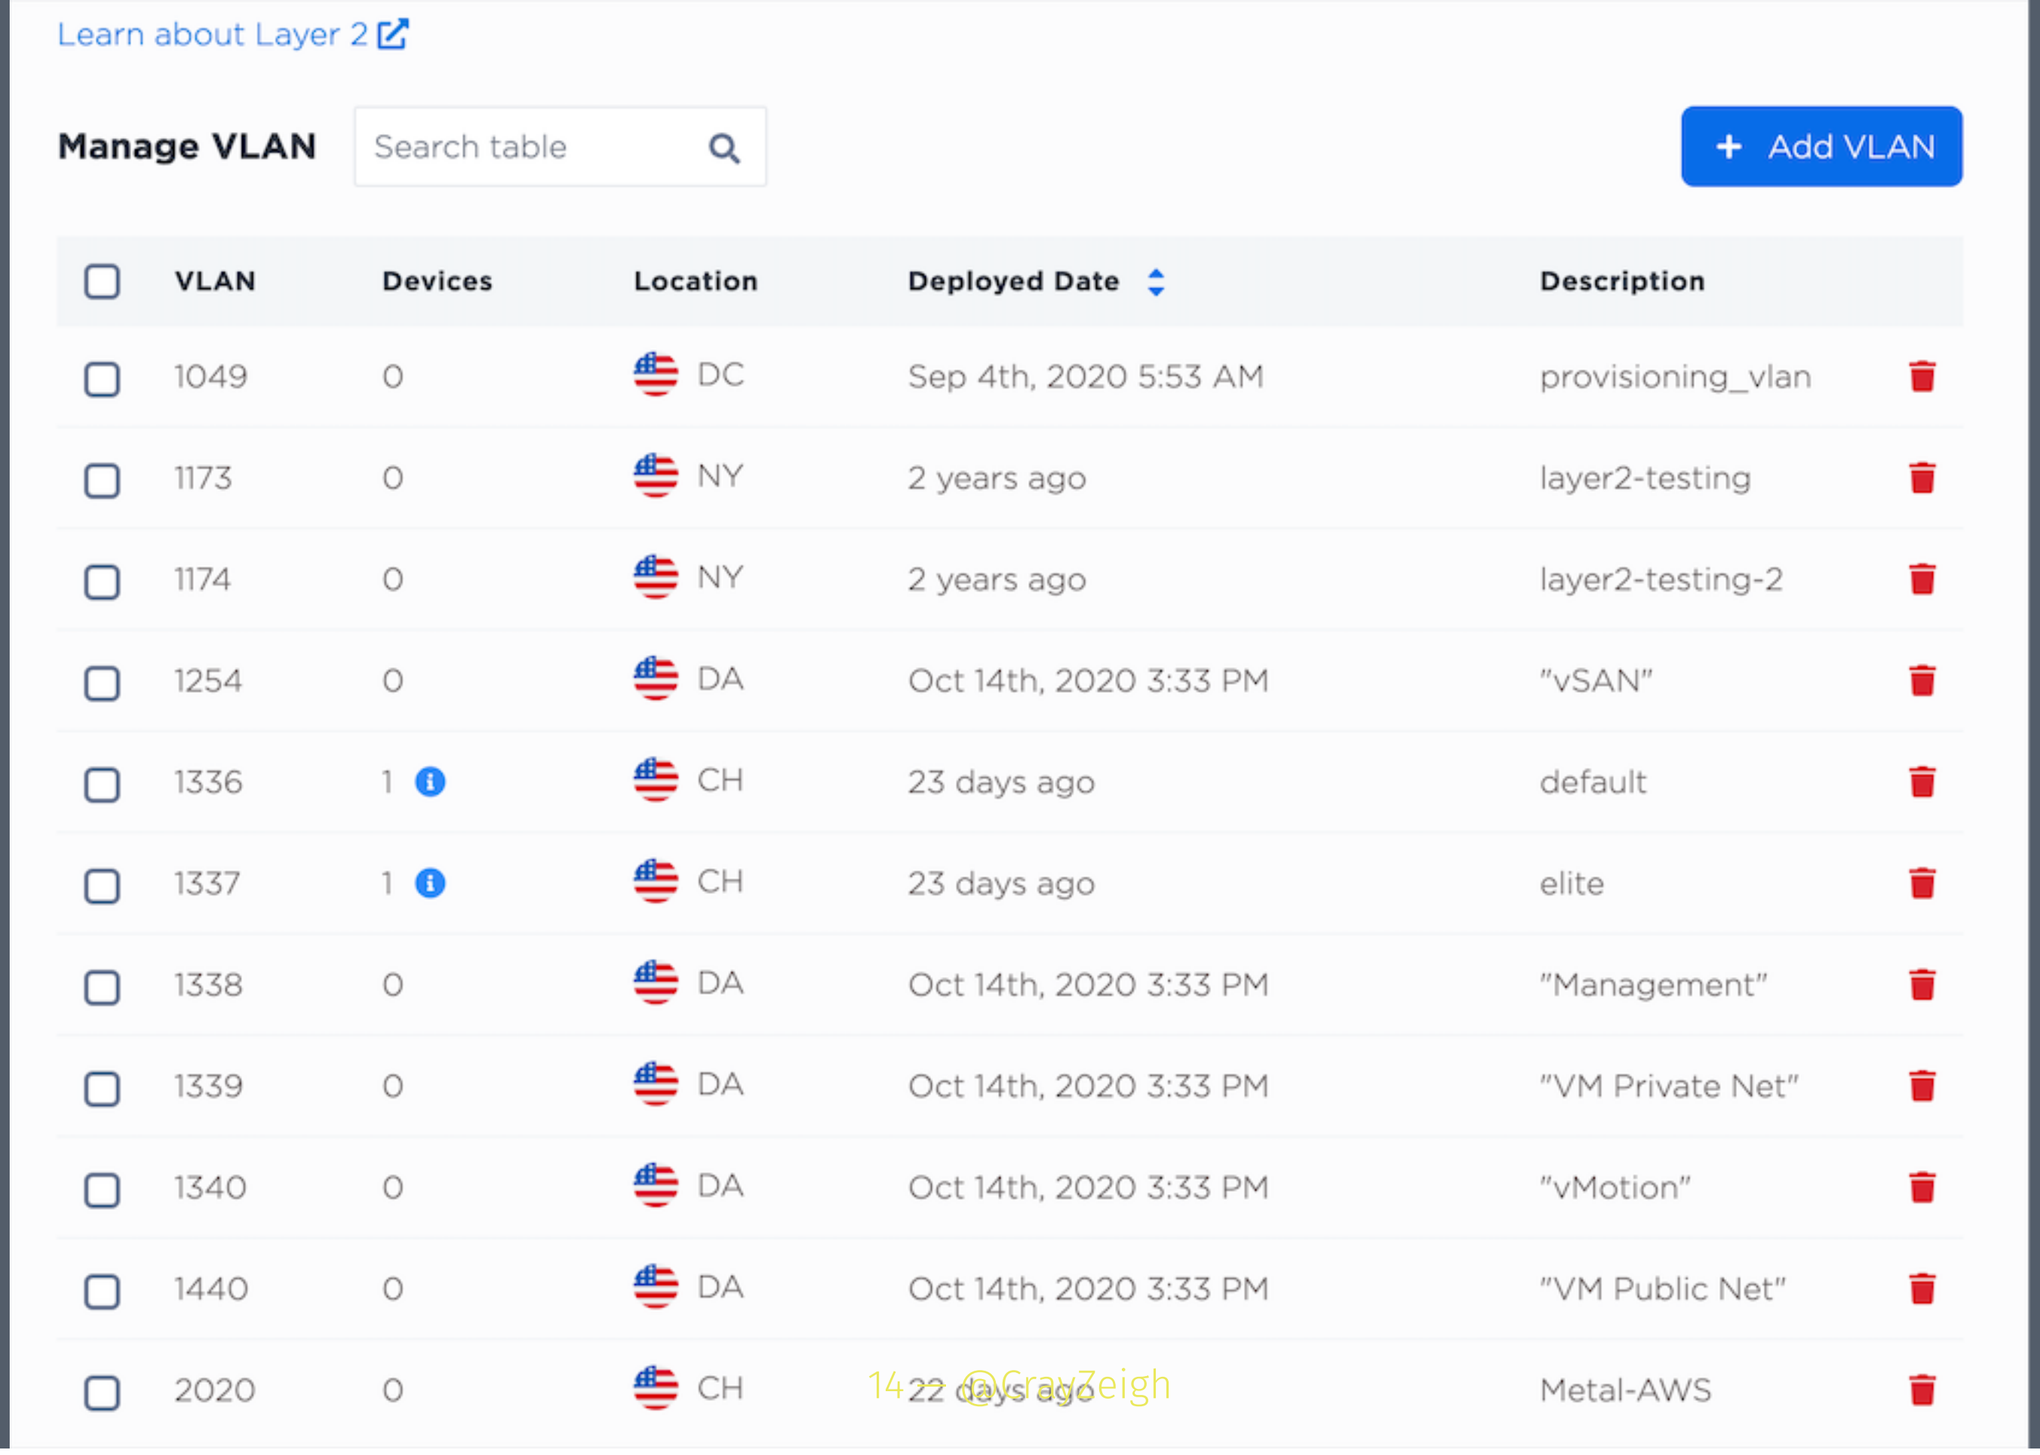

AWS Connections as Layer 2 Segregating networks building mulitple kinds of VM traffic

IP

Packets wrap your digital data and know where to send it

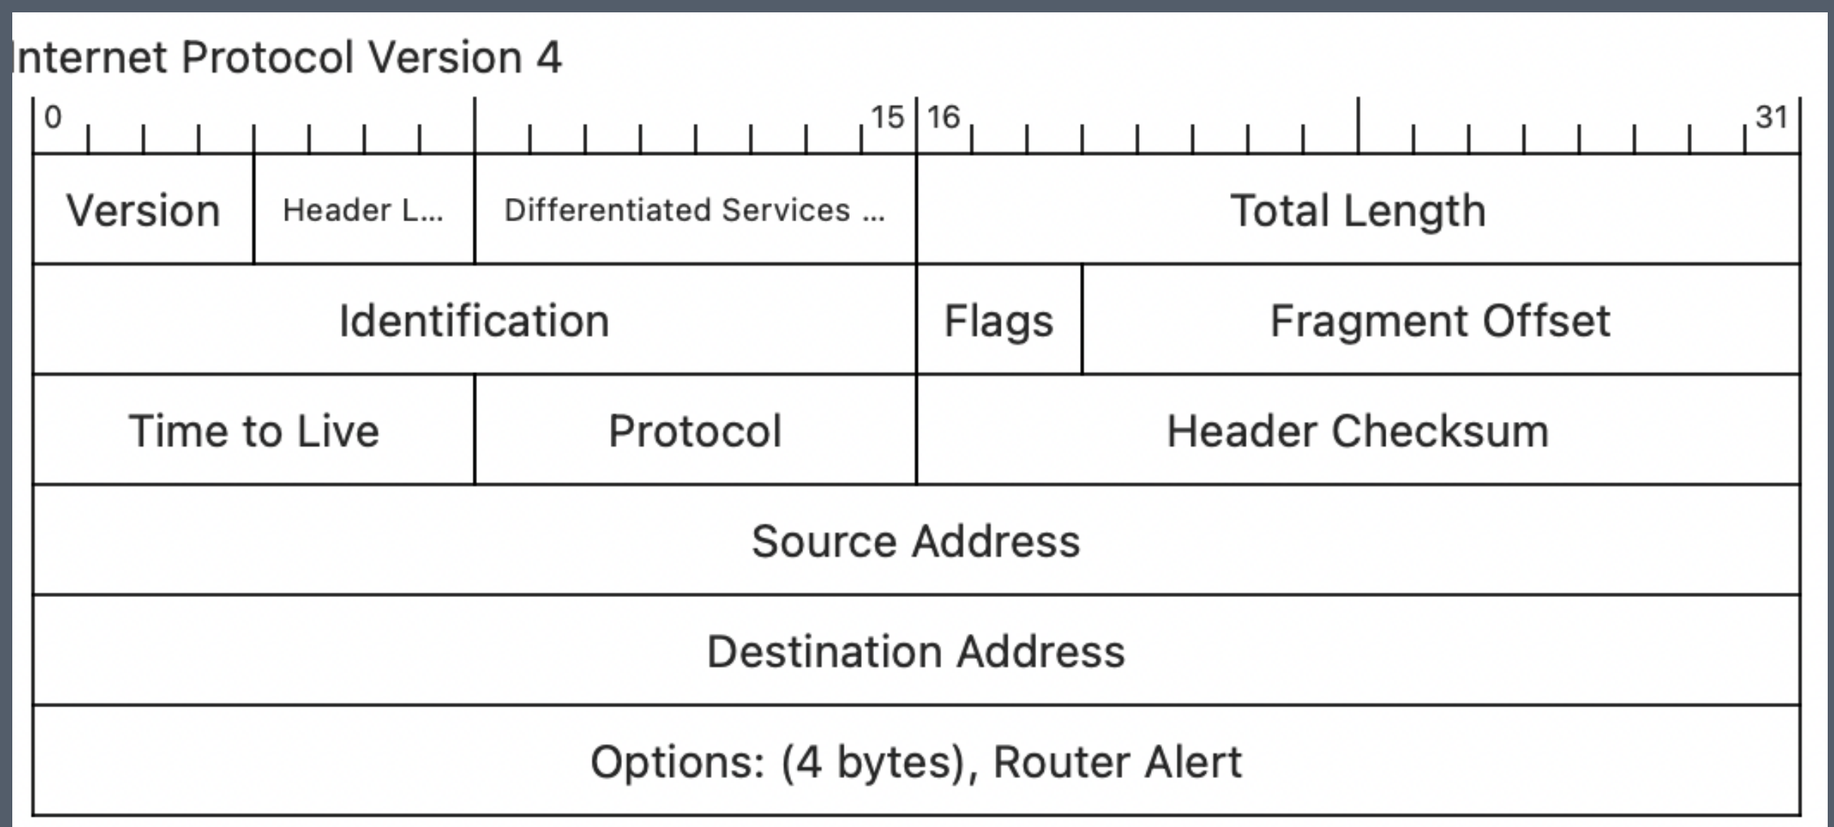

Classes & CIDR

Probably never really dealt with Classed IP addresses though there is some holdover in the reserved private IP space, 10.x, 172.16.x - 172.31.x, 192.168.x

^ Previously given classed space you could have C (256), next up B (65k) and A (16.8m)

Classes & CIDR

→ Classess Inter Domain Routing

Probably never really dealt with Classed IP addresses though there is some holdover in the reserved private IP space, 10.x, 172.16.x - 172.31.x, 192.168.x
^ Previously given classed space you could have C (256), next up B (65k) and A (16.8m)

Classes & CIDR

→ Classess Inter Domain Routing

→ Replaced previous "class a/b/c" IP addressing to help address IP address availability

Probably never really dealt with Classed IP addresses though there is some holdover in the reserved private IP space, 10.x, 172.16.x - 172.31.x, 192.168.x
^ Previously given classed space you could have C (256), next up B (65k) and A (16.8m)

Classes & CIDR

→ Classess Inter Domain Routing

→ Replaced previous "class a/b/c" IP addressing to help address IP address availability

→ Helps determine destination locality i.e. routing

Probably never really dealt with Classed IP addresses though there is some holdover in the reserved private IP space, 10.x, 172.16.x - 172.31.x, 192.168.x
^ Previously given classed space you could have C (256), next up B (65k) and A (16.8m)

CIDR Notation

10.10.10.10/24

IP Address/Network Bits

IP: 10.10.10.10
SM: 255.255.255.0

Converts to Binary:

```
IP: 00001010.00001010.00001010.00001010
SM: 11111111.11111111.11111111.00000000
```

in Subnet Mask:

1s = Network Space

0s = Host Space

Special IPs

Not aobut private, multicast or research IPs that's a different thing

Think of "network" as "any for routing purposes this cannot be used in any other way

Special IPs

→ Broadcast (10.10.10.255)

Not aobut private, multicast or research IPs that's a different thing
Think of "network" as "any for routing purposes this cannot be used in any other way

Special IPs

→ Broadcast (10.10.10.255)

→ Host bits all 1

Not aobut private, multicast or research IPs that's a different thing

Think of "network" as "any for routing purposes this cannot be used in any other way

Special IPs

→ Broadcast (10.10.10.255)

→ Host bits all 1

→ Network (10.10.10.0)

Not aobut private, multicast or research IPs that's a different thing

Think of "network" as "any for routing purposes this cannot be used in any other way

Special IPs

→ Broadcast (10.10.10.255)

→ Host bits all 1

→ Network (10.10.10.0)

→ Host bits all 0

Not aobut private, multicast or research IPs that's a different thing

Think of "network" as "any for routing purposes this cannot be used in any other way

All together

CIDR Notated IP Address of a Host	10.10.10.10/24
Network	10.10.10.0/24
Broadcast IP	10.10.10.255
Available Host IPs	10.10.10.1 - 254

Bigger Networks

CIDR	192.168.1.100/22
Network	192.168.0.0/22
Broadcast	192.168.3.255
Available Hosts	192.168.0.1 - 3.254

More Weird Ones

Weird Ones Explained

Weird Ones Explained

→ /30

Weird Ones Explained

→ /30

→ Costs 4 IPs, but only grants 2 hosts

Weird Ones Explained

→ /30

→ Costs 4 IPs, but only grants 2 hosts

→ Broadcast & Network still apply

Weird Ones Explained

→ /30

→ Costs 4 IPs, but only grants 2 hosts

→ Broadcast & Network still apply

→ Might use today for compatibility reason or because you like IP addresses

Weird Ones Explained

Weird Ones Explained

→ /31

Weird Ones Explained

→ /31

→ Creates 2 adjacet IPs, only "costs" 2 IPs

Weird Ones Explained

→ /31

→ Creates 2 adjacet IPs, only "costs" 2 IPs

→ Proposed in RFC3021 in 2000(!) to combat dwindline IP availability

Anycast App Idea

Anycast App Idea

Weird Ones Explained

→ /32

→ Single IP address

Anycast App Idea

Weird Ones Explained

→ /32

→ Single IP address

→ still very useful mainly for additional or public IPs

Anycast App Idea

Routing

Source: 10.10.10.10/24
Destination: 10.10.10.100

Routing

Source: 10.10.10.10/24
Destination: 10.10.10.100

1. Checks network space to see address is local

Routing

Source: 10.10.10.10/24
Destination: 10.10.10.100

1. Checks network space to see address is local
2. Sends local ARP broadcast to find MAC of destination

Routing

Source: 10.10.10.10/24
Destination: 10.10.10.100

1. Checks network space to see address is local
2. Sends local ARP broadcast to find MAC of destination
3. wraps packet in frame with newly discovered MAC

Routing

Source: 10.10.10.10/24
Destination: 10.10.10.100

1. Checks network space to see address is local
2. Sends local ARP broadcast to find MAC of destination
3. wraps packet in frame with newly discovered MAC
4. sends data frame through switch to destination "directly"

Routing

Source: 10.10.10.10/24
Destination: 1.1.1.1

Routing

Source: 10.10.10.10/24
Destination: 1.1.1.1

1. Checks network space and see's address is remote

Routing

Source: 10.10.10.10/24
Destination: 1.1.1.1

1. Checks network space and see's address is remote
2. Forwards packet to the local router (usually default gateway) through switch

Routing

Source: 10.10.10.10/24
Destination: 1.1.1.1

1. Checks network space and see's address is remote
2. Forwards packet to the local router (usually default gateway) through switch
3. Wraps packet in frame with router's mac address and desired destination's IP

Routing

Source: 10.10.10.10/24
Destination: 1.1.1.1

1. Checks network space and see's address is remote
2. Forwards packet to the local router (usually default gateway) through switch
3. Wraps packet in frame with router's mac address and desired destination's IP
4. Switch forwards frame to the router, router re-wraps the paket with a frame pointing to the next router in line

Routing

Source: 10.10.10.10/24
Destination: 1.1.1.1

1. Checks network space and see's address is remote
2. Forwards packet to the local router (usually default gateway) through switch
3. Wraps packet in frame with router's mac address and desired destination's IP
4. Switch forwards frame to the router, router re-wraps the paket with a frame pointing to the next router in line
5. Eventually, router for 1.1.1.1 will recieve the packet, and wrap in a frame with the appropriate destination's MAC

How does the router know where the next stop is?

Routing Tables

How does the router know where the next stop is?

Routing Tables

→ (also used locally on your hosts)

How does the router know where the next stop is?

Routing Tables

→ (also used locally on your hosts)

→ 3 general types of routes

How does the router know where the next stop is?

Routing Tables

→ (also used locally on your hosts)

→ 3 general types of routes

→ Connected (networks assigned to router interfaces)

How does the router know where the next stop is?

Routing Tables

→ (also used locally on your hosts)

→ 3 general types of routes

→ Connected (networks assigned to router interfaces)

→ Static (manually set, default gateway usually)

How does the router know where the next stop is?

Routing Tables

→ (also used locally on your hosts)

→ 3 general types of routes

→ Connected (networks assigned to router interfaces)

→ Static (manually set, default gateway usually)

→ Learned (Shared with peers, BGP)

BGP

Border Gateway Protocol

BGP

Border Gateway Protocol

→ Advertises routes between (TCP) peered Autonomous Systems

BGP

Border Gateway Protocol

→ Advertises routes between (TCP) peered Autonomous Systems

→ Routes can be aggregate "supernets" to save table space

BGP

Border Gateway Protocol

→ Advertises routes between (TCP) peered Autonomous Systems

→ Routes can be aggregate "supernets" to save table space

→ Helps determine "best" route to destination since multiple routes may contain the same prefixes

BGP

Border Gateway Protocol

→ Advertises routes between (TCP) peered Autonomous Systems

→ Routes can be aggregate "supernets" to save table space

→ Helps determine "best" route to destination since multiple routes may contain the same prefixes

→ Leveragable for anycast/edge performance increases

BGP

Troubleshooting BGP problems

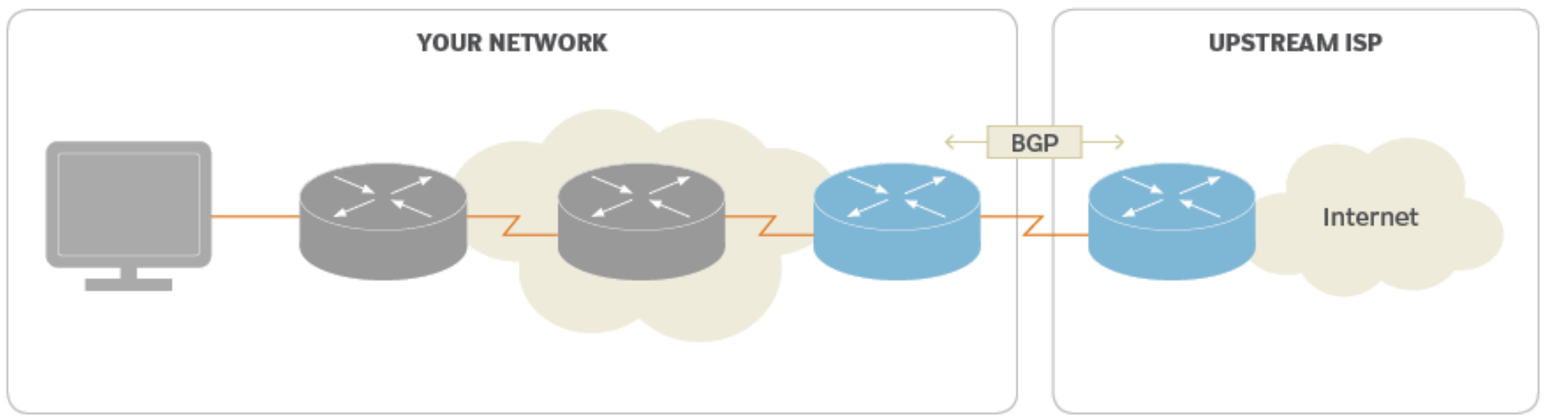

BGP

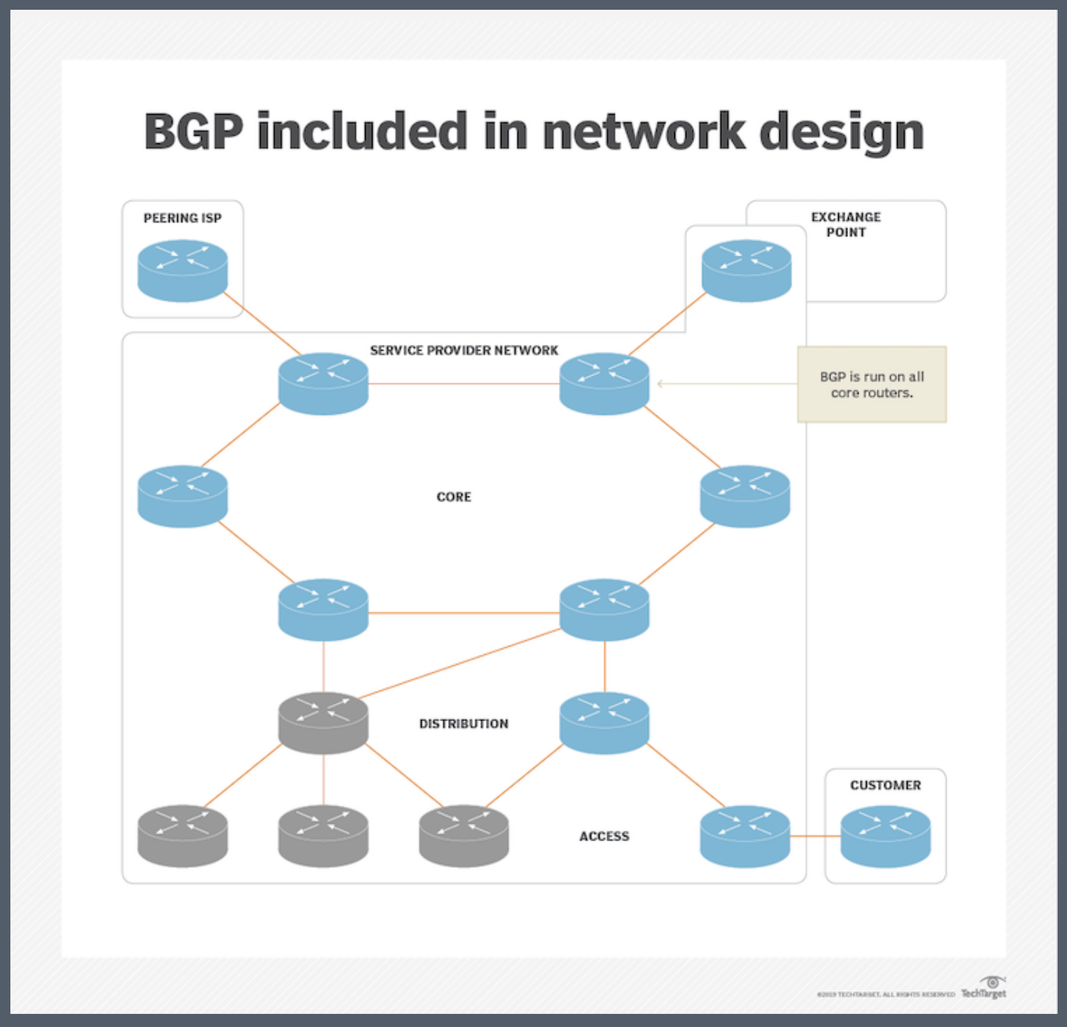
BGP included in network design
PEERING ISP
EXCHANGE POINT
SERVICE PROVIDER NETWORK
BGP is run on all core routers.
CORE
DISTRIBUTION
CUSTOMER
ACCESS
TechTarget

BGP

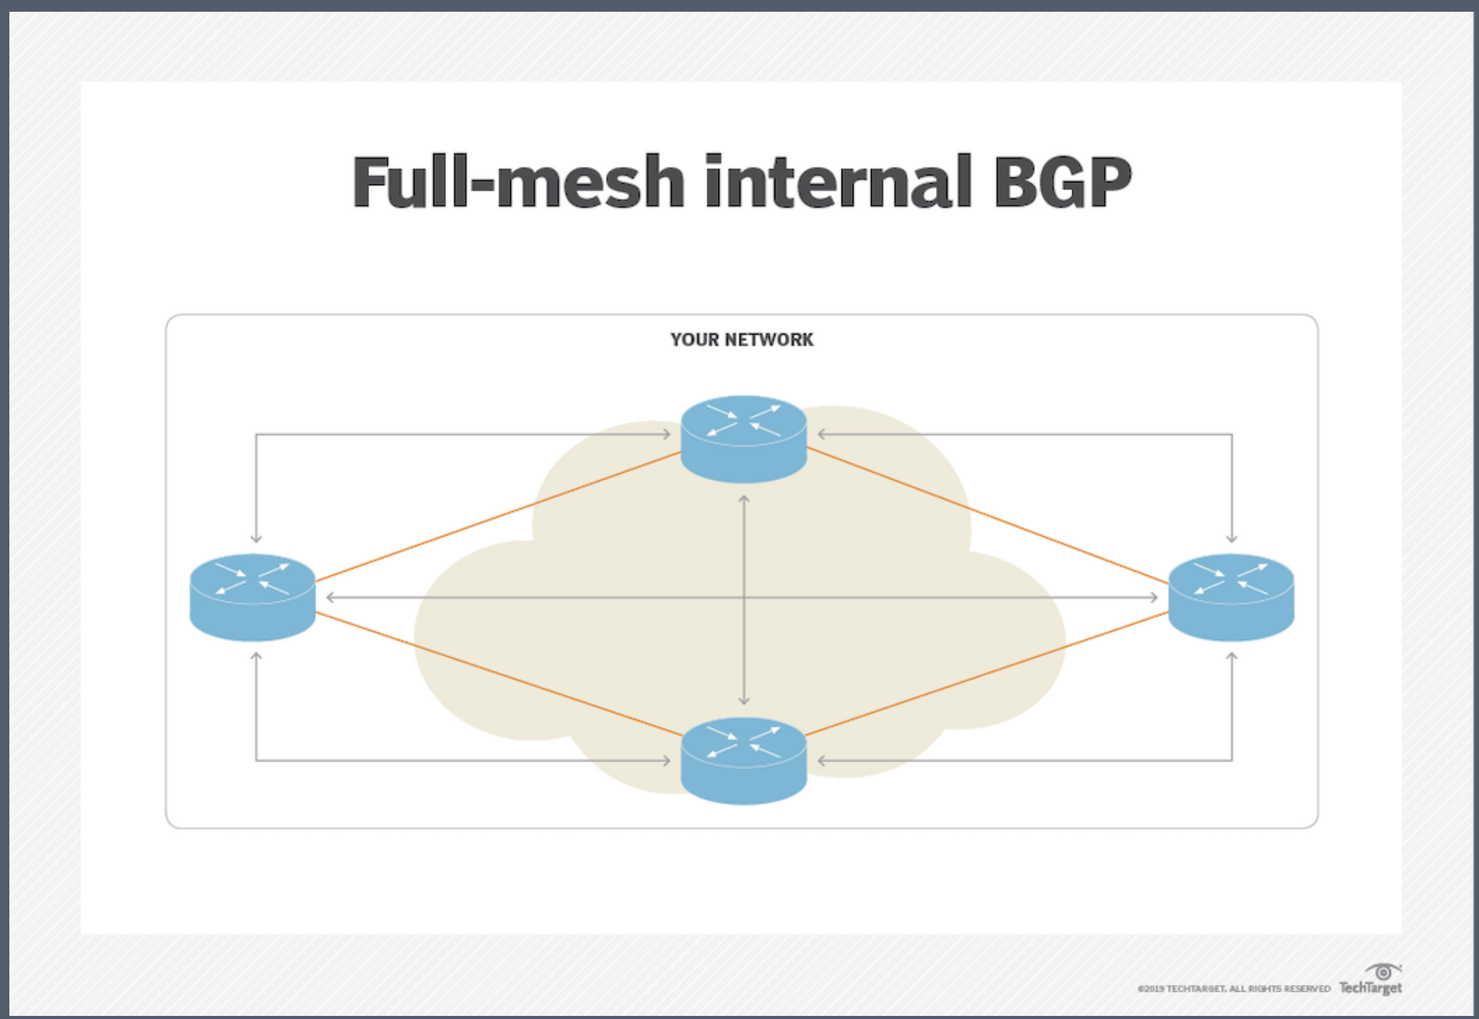
Full-mesh internal BGP
YOUR NETWORK
©2019 TECHTARGET. ALL RIGHTS RESERVED
TechTarget

Anycast

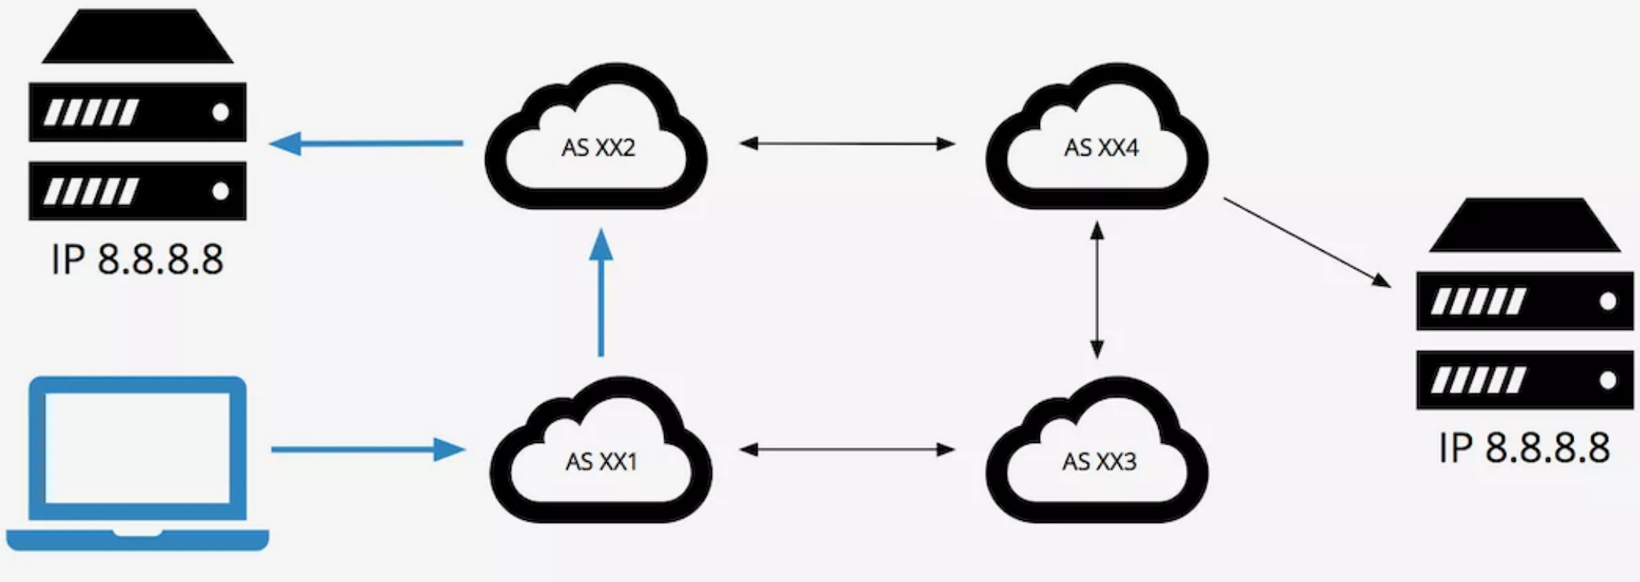

Anycast

Shortest Path: 8.8.8.8 AS path xx1 xx2
8.8.8.8 AS path xx1 xx3 xx4
8.8.8.8 AS path xx1 xx2 xx4
8.8.8.8 AS path xx1 xx3 xx4 xx2

Anycast Benefits

IP address or hostname

147.75.40.38 Test

LOCATION	REQ	MIN	MAX	AVG	STD DEV	LOSS
Frankfurt 147.75.40.38	3	8.42 ms	8.52 ms	8.47 ms	0.04 ms	0%
Amsterdam 147.75.40.38	3	0.73 ms	1.19 ms	0.97 ms	0.19 ms	0%
London 147.75.40.38	3	7.9 ms	8.88 ms	8.27 ms	0.43 ms	0%
New York 147.75.40.38	3	1.71 ms	2.53 ms	2.06 ms	0.34 ms	0%
Dallas 147.75.40.38	3	39.06 ms	39.95 ms	39.42 ms	0.38 ms	0%
San Francisco 147.75.40.38	3	2.2 ms	2.65 ms	2.38 ms	0.2 ms	0%
Singapore 147.75.40.38	3	1.23 ms	1.83 ms	1.46 ms	0.27 ms	0%
Sydney 147.75.40.38	3	92.49 ms	92.57 ms	92.54 ms	0.04 ms	0%
Tokyo 147.75.40.38	3	74.49 ms	74.6 ms	74.57 ms	0.05 ms	0%
Bangalore 147.75.40.38	3	147.75 ms	148.62 ms	148.05 ms	0.41 ms	0%

What happens when you visit a website?

Hi! 👋 I'm Aaron

dev advocate: EQUINIX

organizer:

sometimes host:

Twitter: @CrayZeigh
Slides: speaking.crayzeigh.com