netWorking through the abstractions





l'm fen (they/them)

social: @crayzeigh@hachyderm.io slides: speaking.crayzeigh.com



Deve oper Advocate





Internet Traffic Per Second

@crayzeigh@hachyderm.io



bits per second 133655440

Y



What happens when you visit a Webster

It's not DNS There's no way it's DNS It was DNS

-SSBroski





OS 7-Laver mode

Application Presentation Session Transport Network Data Link Physical



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



Why all these layers anyway?



Retwork A CCESS Data Frames link digital to physical



Layer 2: Switching Sending data to local devices

@crayzeigh@hachyderm.io

12

Frame Header

Preamble	SFD	Dest. MAC	Src. MAC	Туре	Data
7 Bytes	1 Byte	6 Bytes	6 Bytes	2 Bytes	46–15 Bytes

A & Pad FCS

500 4 Bytes

S

MAC¹ Accress

01F23F45F67F89Fab

1. Media Access Control

– Identifies the (network) device – For devices on the *local* network

@crayzeigh@hachyderm.io

16



Mapping IPs and MAC addresses



 Mapping IPs and MAC addresses Necessary for your device to talk to your router



- Mapping IPs and MAC addresses Necessary for your device to talk to your router
- Finds neighbors through broadcast



- Mapping IPs and MAC addresses
- Necessary for your device to talk to your router
- Finds neighbors through broadcast - arp -a for your local table





Used for limiting broadcast domains



Used for limiting broadcast domains up to 4096 VLANs¹

1. VXLAN addresses this limitation but that's A Whole Other Thing



- Used for limiting broadcast domains
- up to 4096 VLANs^{1}
- Can be native (default traffic on a switch port) or tagged (logically divided in packet header)

VXLAN addresses this limitation but that's A Whole Other Thing 1.

Learn	about	Layer 2	\mathbf{C}
-------	-------	---------	--------------

Mana	ige VLAN	Search table	Q		+ Add VLAN
	VLAN	Devices	Location	Deployed Date 🌲	Description
	1049	0	DC	Sep 4th, 2020 5:53 AM	provisioning_vlan
	1173	0	Service NY	2 years ago	layer2-testing
	1174	0	Service NY	2 years ago	layer2-testing-2
	1254	0	豊 DA	Oct 14th, 2020 3:33 PM	"vSAN"
	1336	1 🚯	👙 СН	23 days ago	default 🛑
	1337	1 🚯	👙 СН	23 days ago	elite
	1338	0	👙 DA	Oct 14th, 2020 3:33 PM	"Management"
	1339	0	👙 DA	Oct 14th, 2020 3:33 PM	"VM Private Net"
	1340	0	豊 DA	Oct 14th, 2020 3:33 PM	"vMotion"
	1440	0	👙 DA	Oct 14th, 2020 3:33 PM	"VM Public Net"
	2020	0	👙 СН	22 days ago	Metal-AWS

Internet Protoco Packets wrap your digital data and route it remotely

Layer 3: Routing data to remot

Directing data to remote destinations

Pv4 Packet Header

Internet Protocol Version 4

0		15	16		
Version	Header L	Differentiated Services	Total Leng		
	Identification			Fragment	
Time t	o Live	Protocol	Header Chec		
Source Address					
Destination Address					
Options: (4 bytes), Router Alert					







- Classless Inter-Domain Routing



 Classless Inter-Domain Routing Helps determine destination locality, i.e. routing



PCBSSES 30 DR

- Classless Inter-Domain Routing
- Helps determine destination locality, i.e. routing
- CIDR replaced "class a/b/c" IP addressing to help address IP address availability





10101010/24



HOST NETWORK



10.10.10/24

IP Address: 10.10.10.10 Subnet Mask: 255.255.255.0
Converts to Binary

IP: 00001010.00001010.00001010.00001010
SM: 1111111.111111.111111.00000000

In the Subnet Mask:

1 = Network

0 = Host

Special Ps

- Broadcast (ex. 10.10.10.255)
 - host bits are all 1's
 - For sending data to all hosts in a network
- Network (ex. 10.10.10.0)
 - host bits are all 0's
 - only used for forwarding data between routers

Putting it all together

CIDR

10.10.10/24

Network

10.10.10.0/24

Broadcast IP

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 IP

192.168.3.255

Available Host IPs

192.168.0.1 - 192.168.3.254

Weird Ones

Select

Select

/32 (1 IP) - \$0.15/hr

/31 (2 IPs) - \$0.3/hr

/30 (4 IPs) - \$0.6/hr

Description (optional)



Weird Ones Explained

- "Costs" 4 IPs, but only gives 2 host addresses
- Broadcast & Network IPs still required
- Used for legacy compatibility or you just really like holding IP addresses

addresses ired u just

Weird Ones Explained /31

- Creates 2 adjacent host IPs
- Only "costs" 2 IPs
- Proposed in RFC3021 (in 2000) to combat dwindling IP availability

Weird Ones Explained



Single IP Address useful generally for isolating public internet traffic





IPv6 Packet Header



– 128 bits long (vs 32 bits for v4)



128 bits long (vs 32 bits for v4) $2^{32} \approx 4.3 \mathrm{x10}^9$



128 bits long (vs 32 bits for v4) $2^{32} \approx 4.3 \times 10^9 \approx 4.3$ billion



128 bits long (vs 32 bits for v4) $2^{32} \approx 4.3 \text{x} 10^9 \approx 4.3$ billion $2^{128} \approx 3.4 \mathrm{x10}^{28}$



128 bits long (vs 32 bits for v4) $2^{32} \approx 4.3 \times 10^9 \approx 4.3$ billion $2^{128} \approx 3.4 \text{x} 10^{28} \approx \text{way bigger}$



- 128 bits long (vs 32 bits for v4) introduced to handle IP shortage – written as 8 groups of 4 bytes in hex



fddd:f00d:b33f:0000:0000:0000:0000:0001



fddd:f00d:b33f:0:0:0:0:1



fddd:f00d:b33f::1



00ff:0000:0000:0000:0001:0000:0000:0001

Which one is correct?

ff:0:0:0:1:0:0:1

- 1. ff::1:0:0:1
- 2. ff::1::1
- 3. ff:0:0:0:1::1

Which one is correct?

ff:0:0:0:1:0:0:1

- 1. ff::1:0:0:1
- 2. ff::1::1 (ambiguous)
- RFC 5952 1.

Subnetting

	Netv		Но	ost Bi		
	Routing Pre	efix	Subnet ID		Interfa	ce Ide
XXXX	XXXX	XXXX	XXXX	XXXX	XXXX	X

@crayzeigh@hachyderm.io

Bits dentifier XXXX XXXX

Subnetting

ifconfig | grep inet6 inet6 ::1 prefixlen 128 inet6 fe80::1%lo0 prefixlen 64 scopeid 0x1 inet6 fe80::1470:43d6:1243:7a20%en0 prefixlen 64 secured scopeid 0x4 inet6 2601:19e:8380:20dc:4ea:6990:6d8e:68a3 prefixlen 64 autoconf secured inet6 2601:19e:8380:20dc:5ce:b0e:70ad:988 prefixlen 64 deprecated autoconf temporary inet6 2601:19e:8380:20dc::f42a prefixlen 64 dynamic inet6 2601:19e:8380:20dc:141b:7764:190d:9240 prefixlen 64 deprecated autoconf temporary inet6 2601:19e:8380:20dc:d477:fa12:2a44:7983 prefixlen 64 autoconf

temporary

• • •

Subnetting

address	obtaine
··1	loopbac
fe80::1470:43d6:1243:7a20	self-assi
2601:19e:8380:20dc:4ea:6990:6d8e:68a3	self-assi
2601:19e:8380:20dc::f42a	DHCPv6

@crayzeigh@hachyderm.io

ed

ck (special)

signed (local)

signed (global)

r6 (global)

Routing Sending data to remote addresses

Source: 10.10.10.10/24 Destination: 10.10.10.100

Source PH1010101010 etwork 255-255-255-0

Destination: 10.10.10.100

Source: 10.10.10.10/24 Destination: 10.10.10.10.100

- 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"

ocal of destination vered MAC ination

Source 10.10.10.10/24 Destination: 1111



Source IP: 10.10.10.10 Network: 255.255.255.0

Destination: 1111

Source: 10.10.10/24 **Destination**: 1.1.1.1

- Checks network space and see's address is remote 1.
- 2. Wraps packet in frame with router's mac address and desired destination IP, forwards to router
- 3. Forwards frame to the router, router re-wraps packet with a frame pointing to the next router in line
- 4. And so on until the local router for 1.1.1.1 receives the packet and wraps in the final destination frame

How does the router know where the next hop in line is?

Routing Tables

- 3 Types of Routes:
- **Connected**: The networks connected to your interfaces
- **Static**: Manually set, e.g. Default Route/Gateway
- **Learned**: Learned from advertising peers, e.g. via BGP



fen 蜜 💓

@crayzeigh@hachyderm.io

@malanalysis things up takes bgp

to err is dns, but to globally f

Dec 15, 2022 at 11:05 · 🚱 · Toot! · 🔂 10 · 🚖 19





Full-mesh internal BGP



BGP

BGP included in network design



@crayzeigh@hachyderm.io

73

SI



Anycast Benefits

P address or hostnar	me					
147.75.40.38						Test
LOCATION	REQ	MIN	MAX	AVG	STD DEV	LOSS
Frankfurt	3	8.42 ms	8.52 ms	8.47 ms	0.04 ms	0%
Amsterdam	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	3	1.71 ms	2.53 ms	2.06 ms	0.34 ms	0%
Dallas	3	39.06 ms	39.95 ms	39.42 ms	0.38 ms	0%
San Francisco	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%

Anycast Benefits

147.75.40.38

London 147.75.40.38	3	7.9 ms	8.88 ms	8.27 ms	0.43 ms	0%
New York	3	1.71 ms	2.53 ms	2.06 ms	0.34 ms	0%
Dallas	3	39.06 ms	39.95 ms	39.42 ms	0.38 ms	0%
San Francisco	3	2.2 ms	2.65 ms	2.38 ms	0.2 ms	0%

What Happens when you visit a website?





@crayzeigh@hachyderm.io





use discount code ato2023 for \$300 credit



http://eqix.co/metal

Getting Started

check out our youtube channel for tips and ideas



http://eqix.co/start-metal

thanks!



social: @crayzeigh@hachyderm.io slides: speaking.crayzeigh.com



organizer.





