

# Web Requests, Down To The Atom


**Chris Zacharias**


Web Unleashed 2019

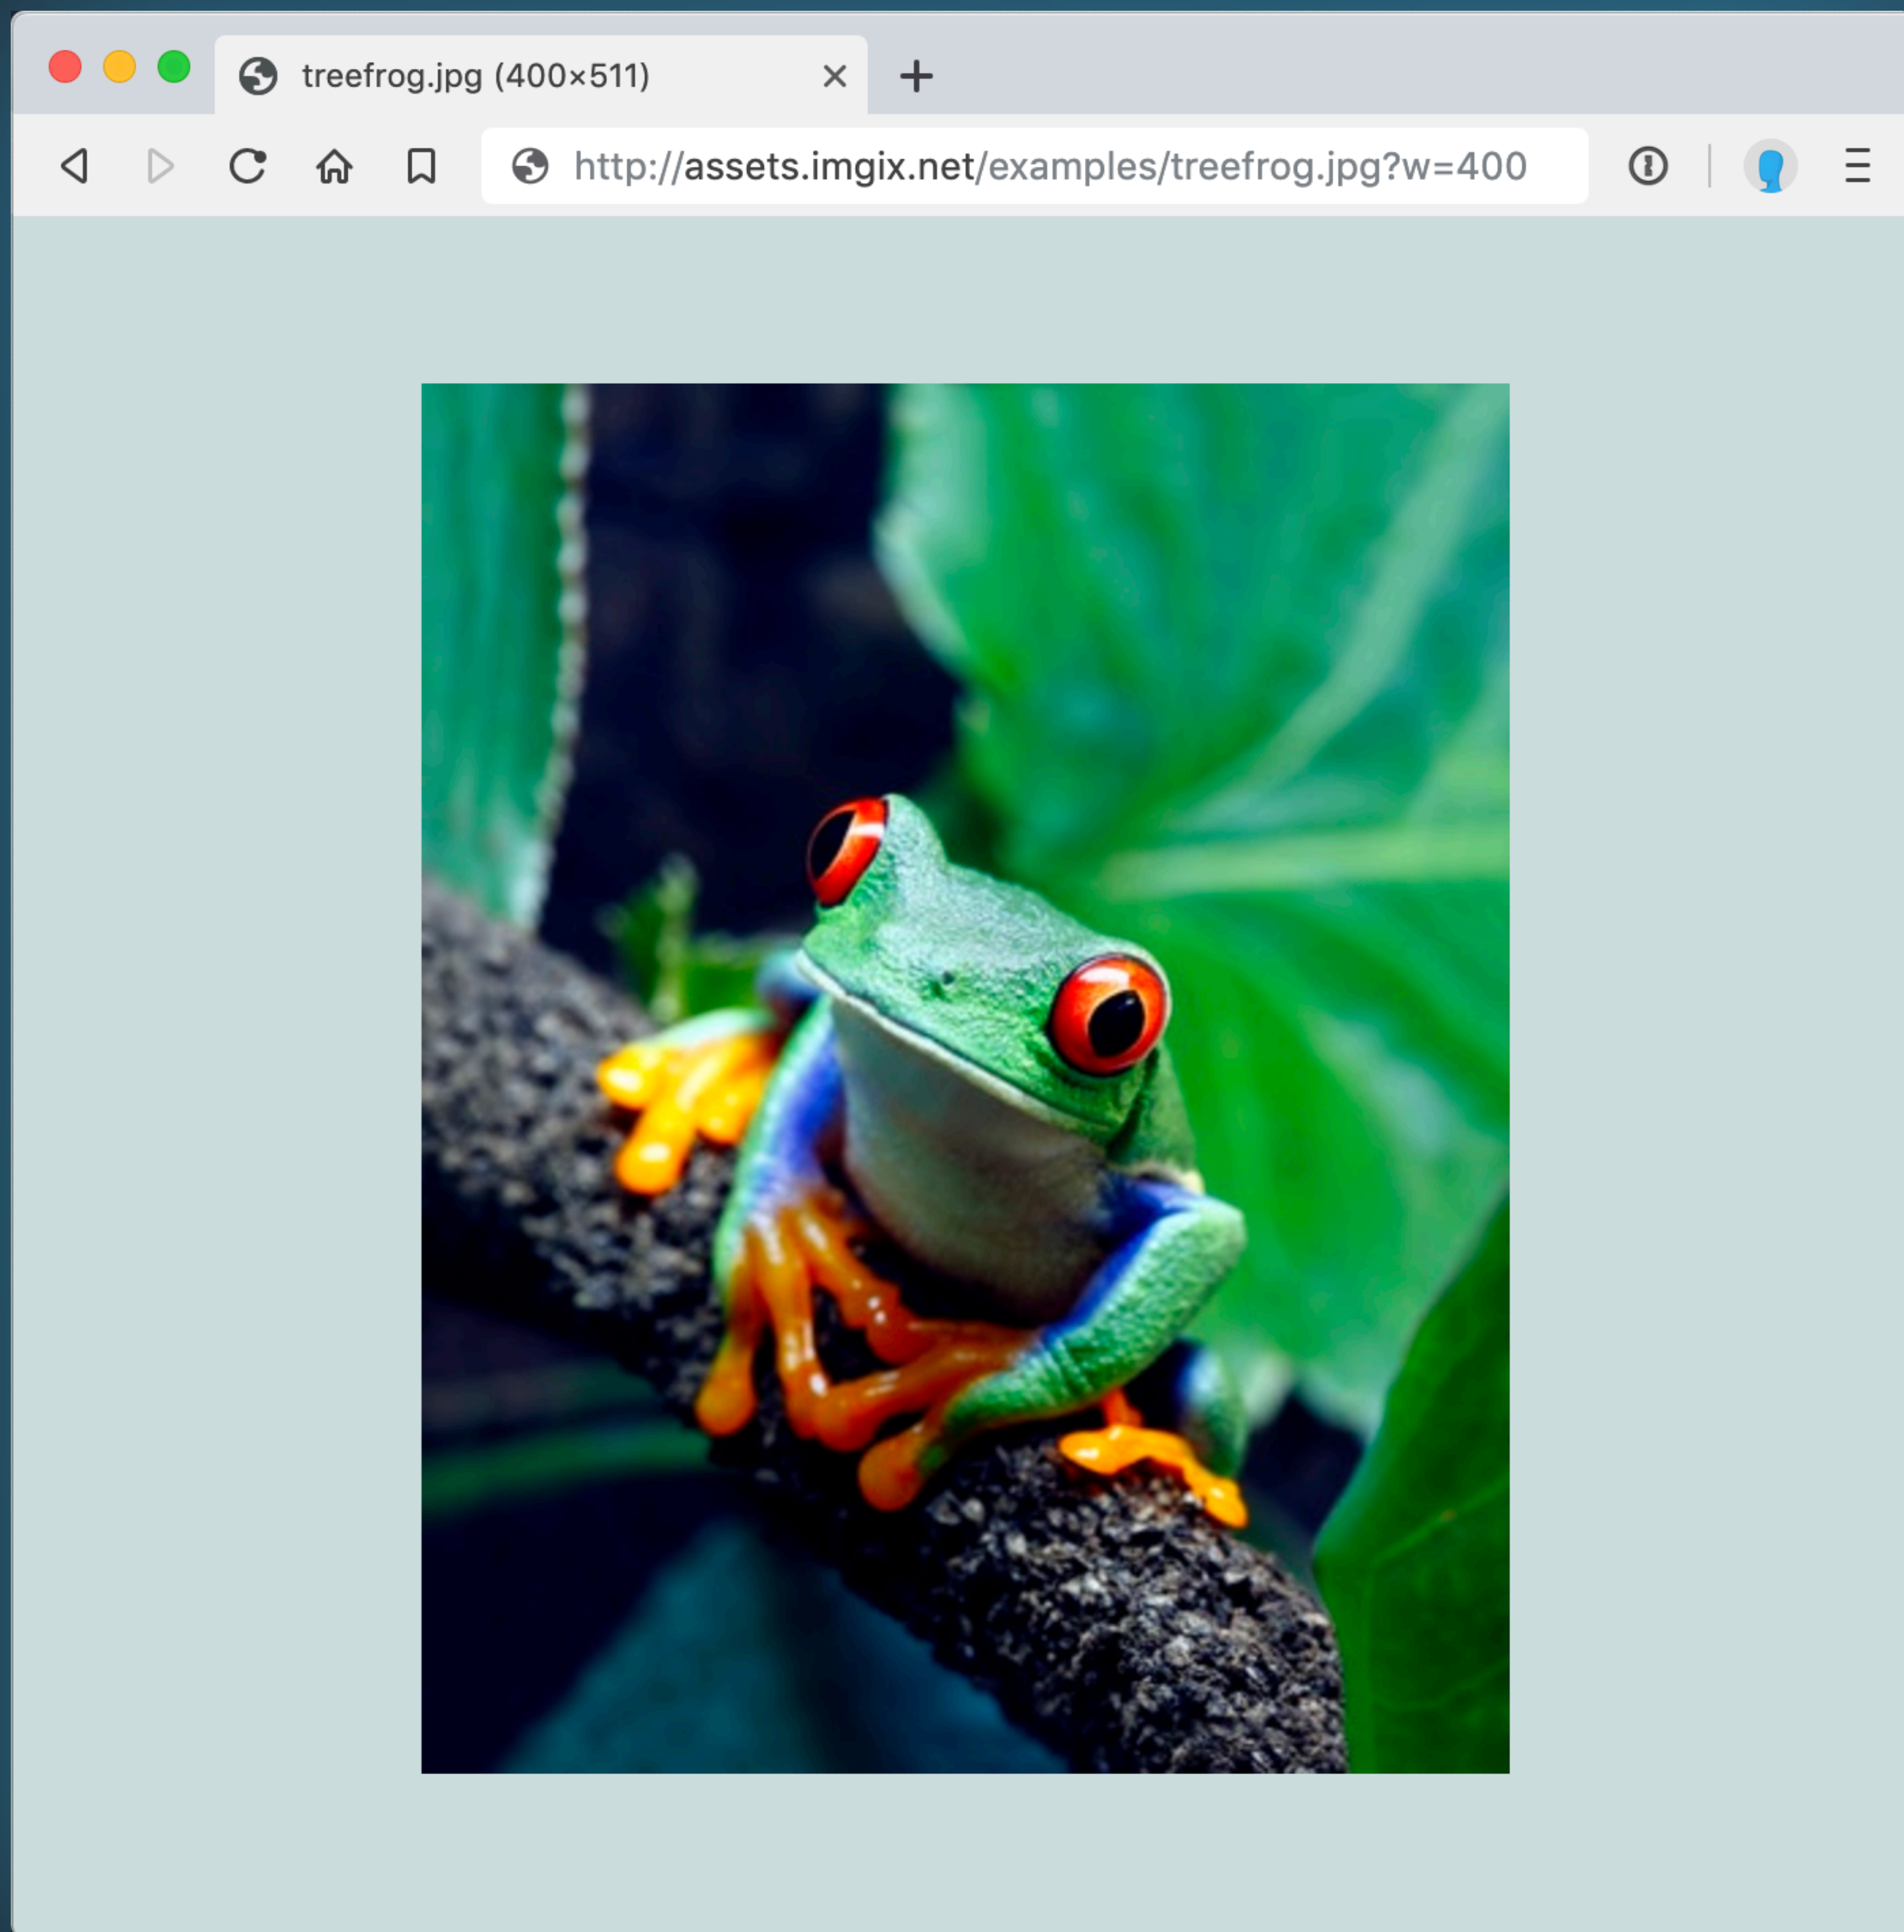
Toronto

## Understanding how the Internet fundamentally works is important.

- Identify performance opportunities at every layer.
- Make smarter decisions around compression and optimization.
- Become better at triaging and debugging network-related issues.
- Recognize potential attack vectors and vulnerabilities.
- Understand when and where costs occur.
- Communicate better with other engineers.









A web request begins with a URL.

 <http://assets.imgix.net/examples/treefrog.jpg?w=400>

The URL answers three simple questions.

 `http://assets.imgix.net/examples/treefrog.jpg?w=400`

What kind of  
request is it?

Where do I send the  
request?

What is it that I  
want to request?

The URL answers three simple questions.



`http://assets.imgix.net/examples/treefrog.jpg?w=400`

Protocol

Network Location

Path and Query

What kind of  
request is it?

Where do I send the  
request?

What is it that I  
want to request?

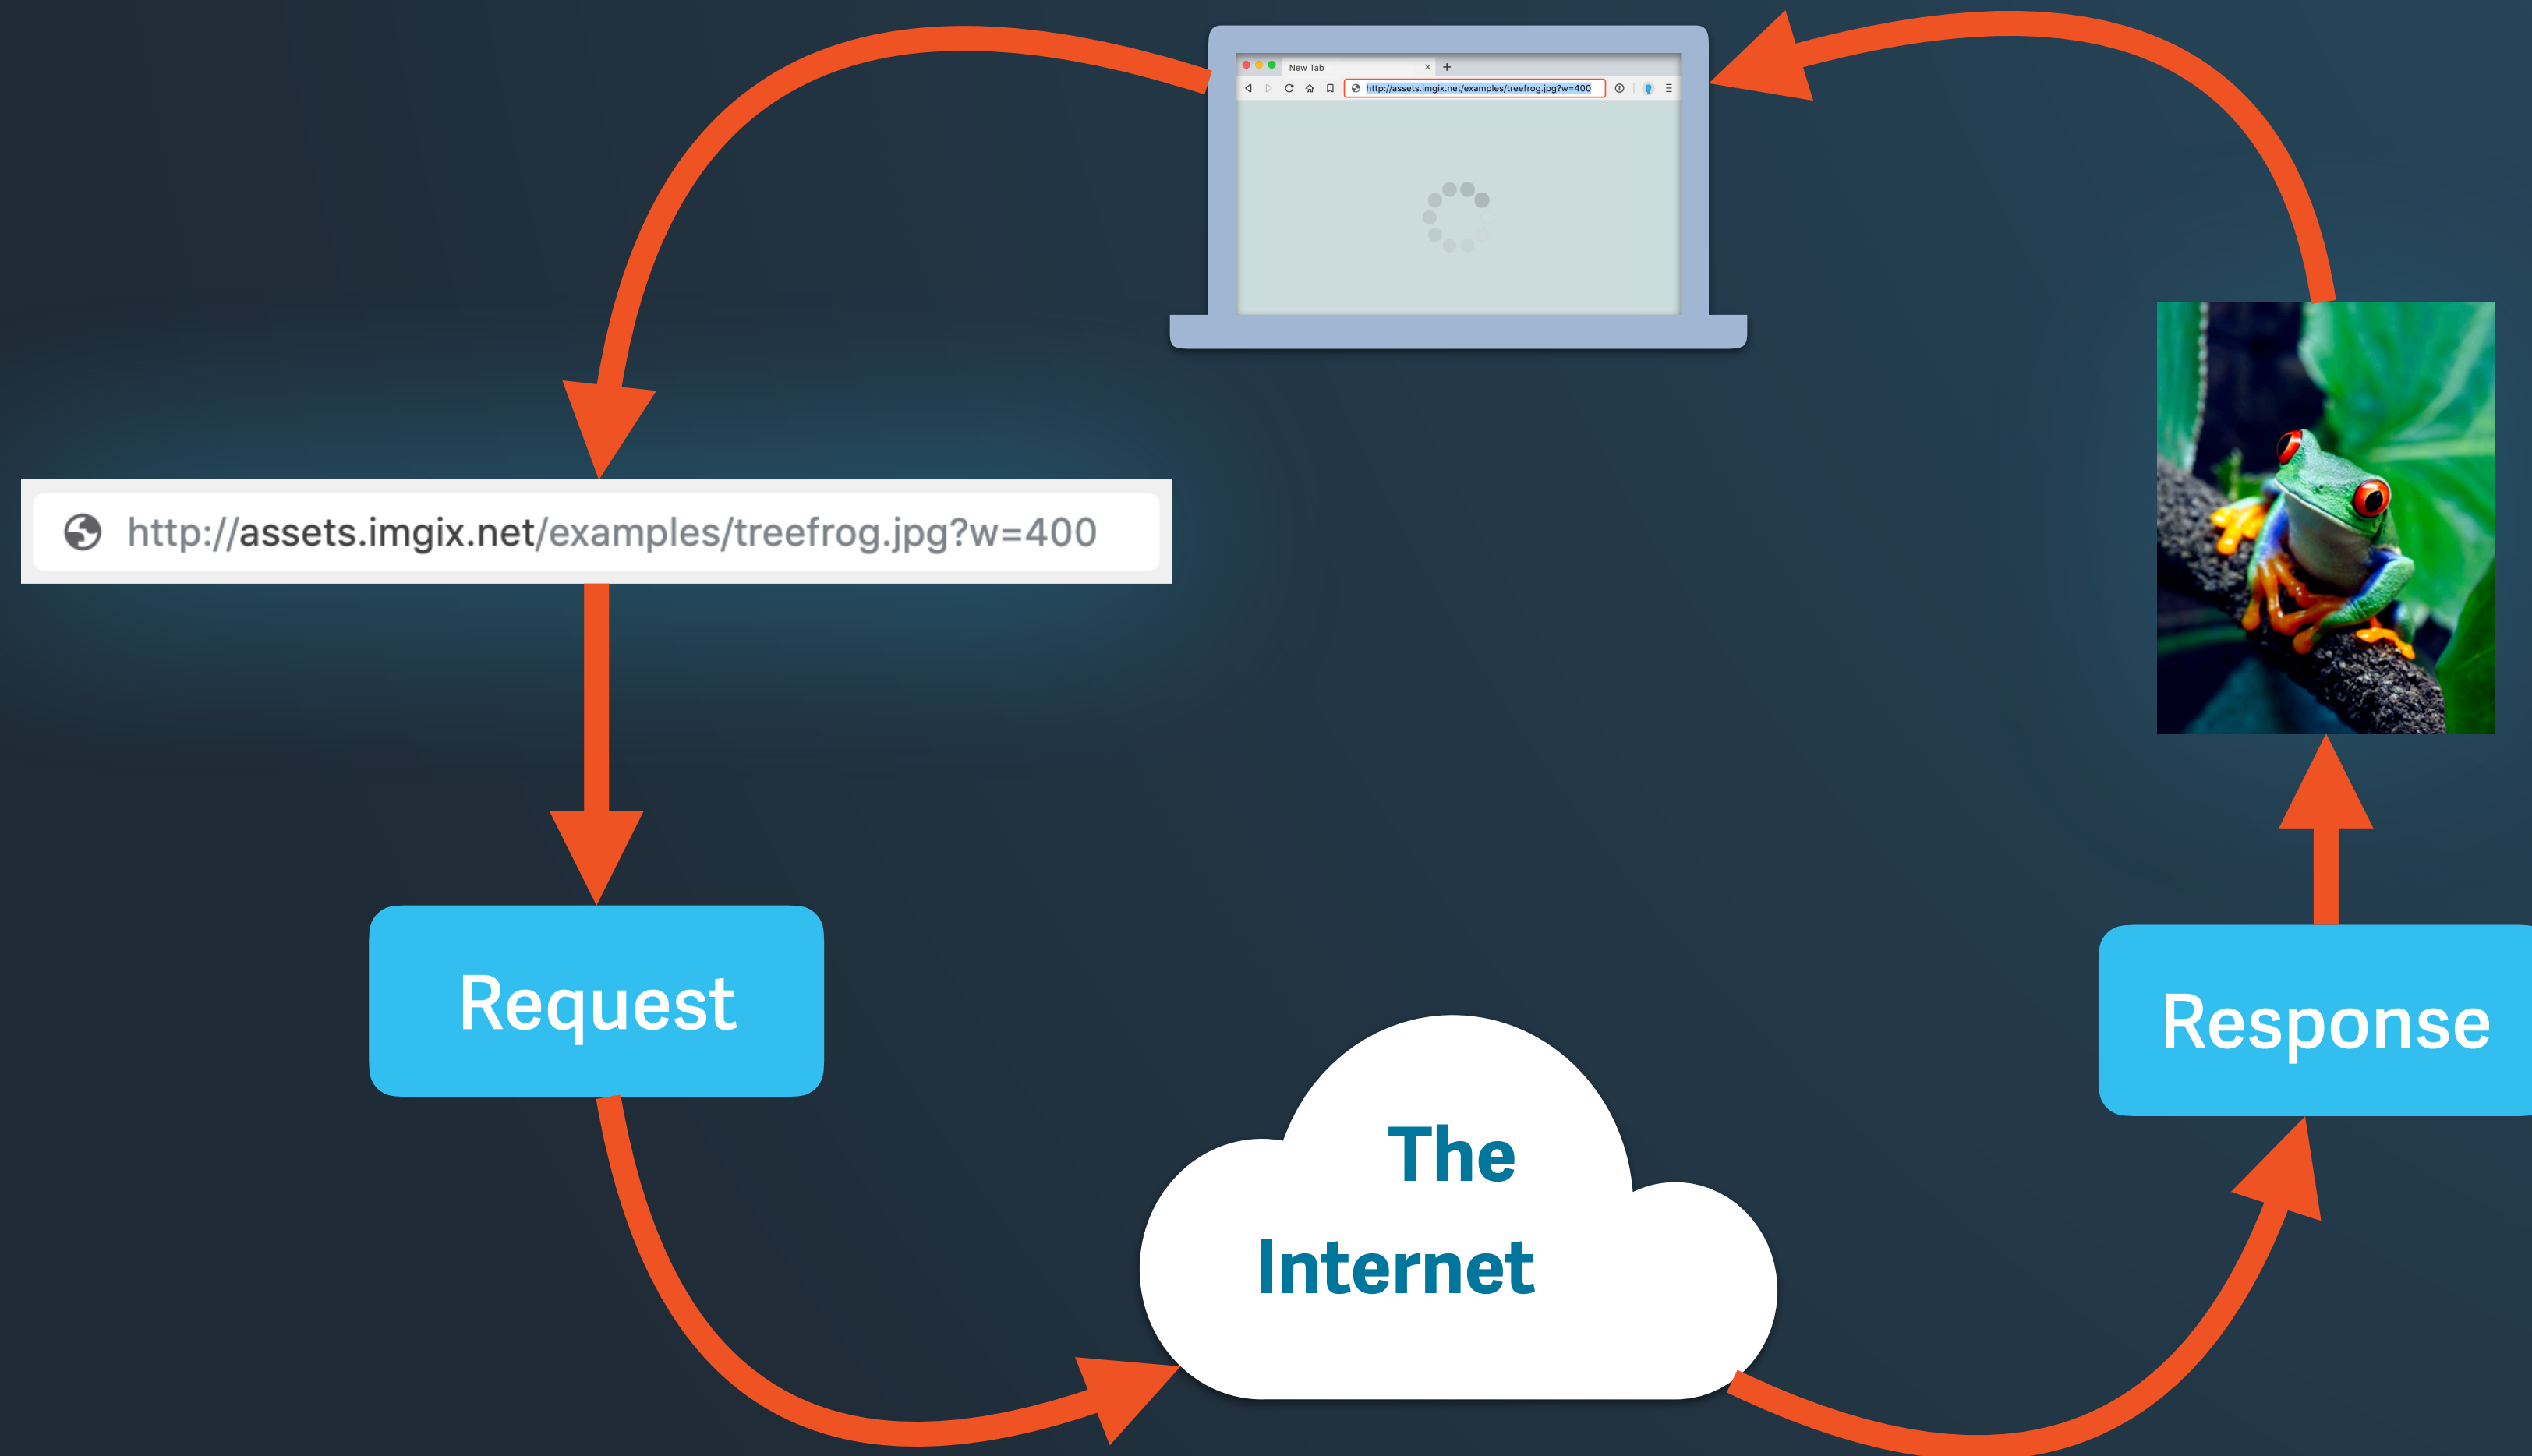


The browser builds an HTTP request out of the URL we provide.




An HTTP GET request.


The request is sent out onto the Internet and (hopefully) we get a response.



The Internet is just a collection of networks that connect clients to servers.

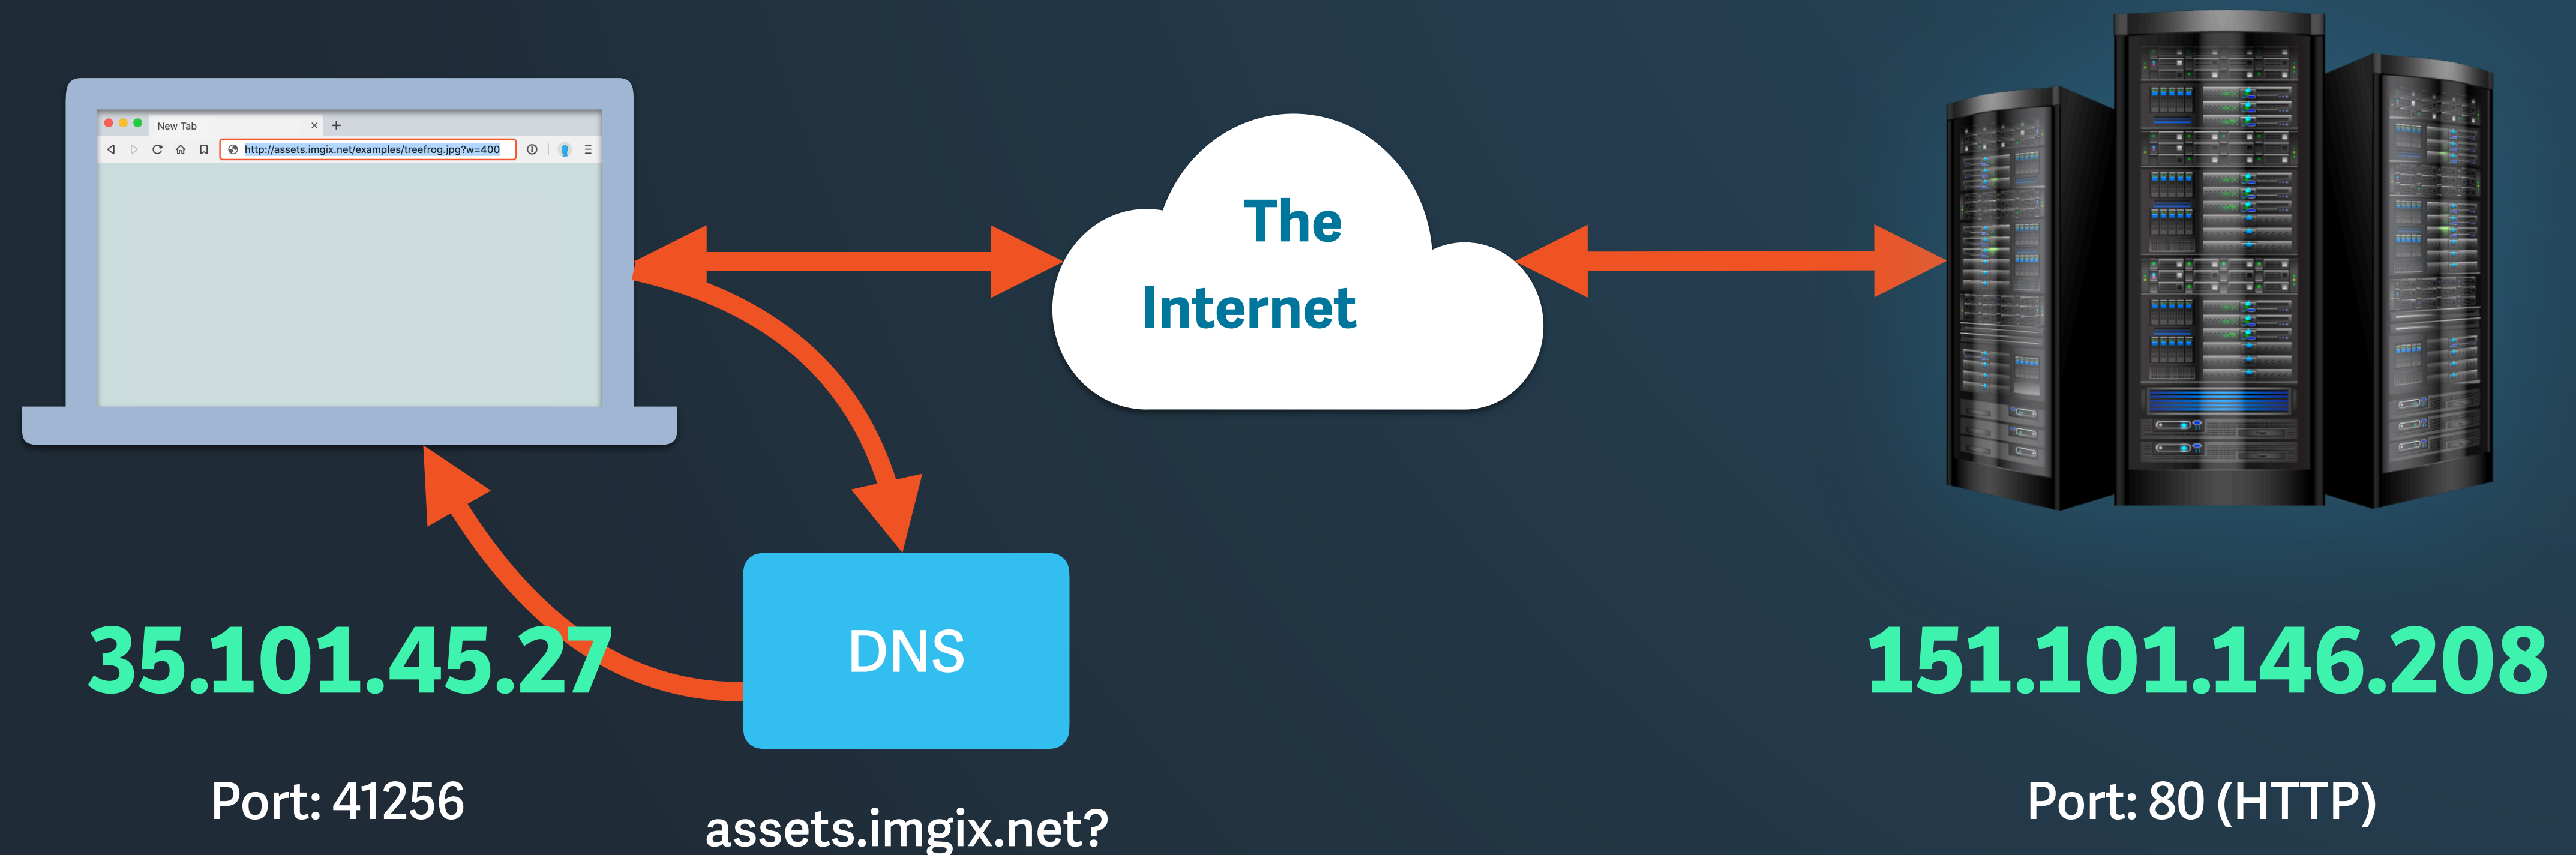


By "hopping" between networks, we can find our way to the right server.



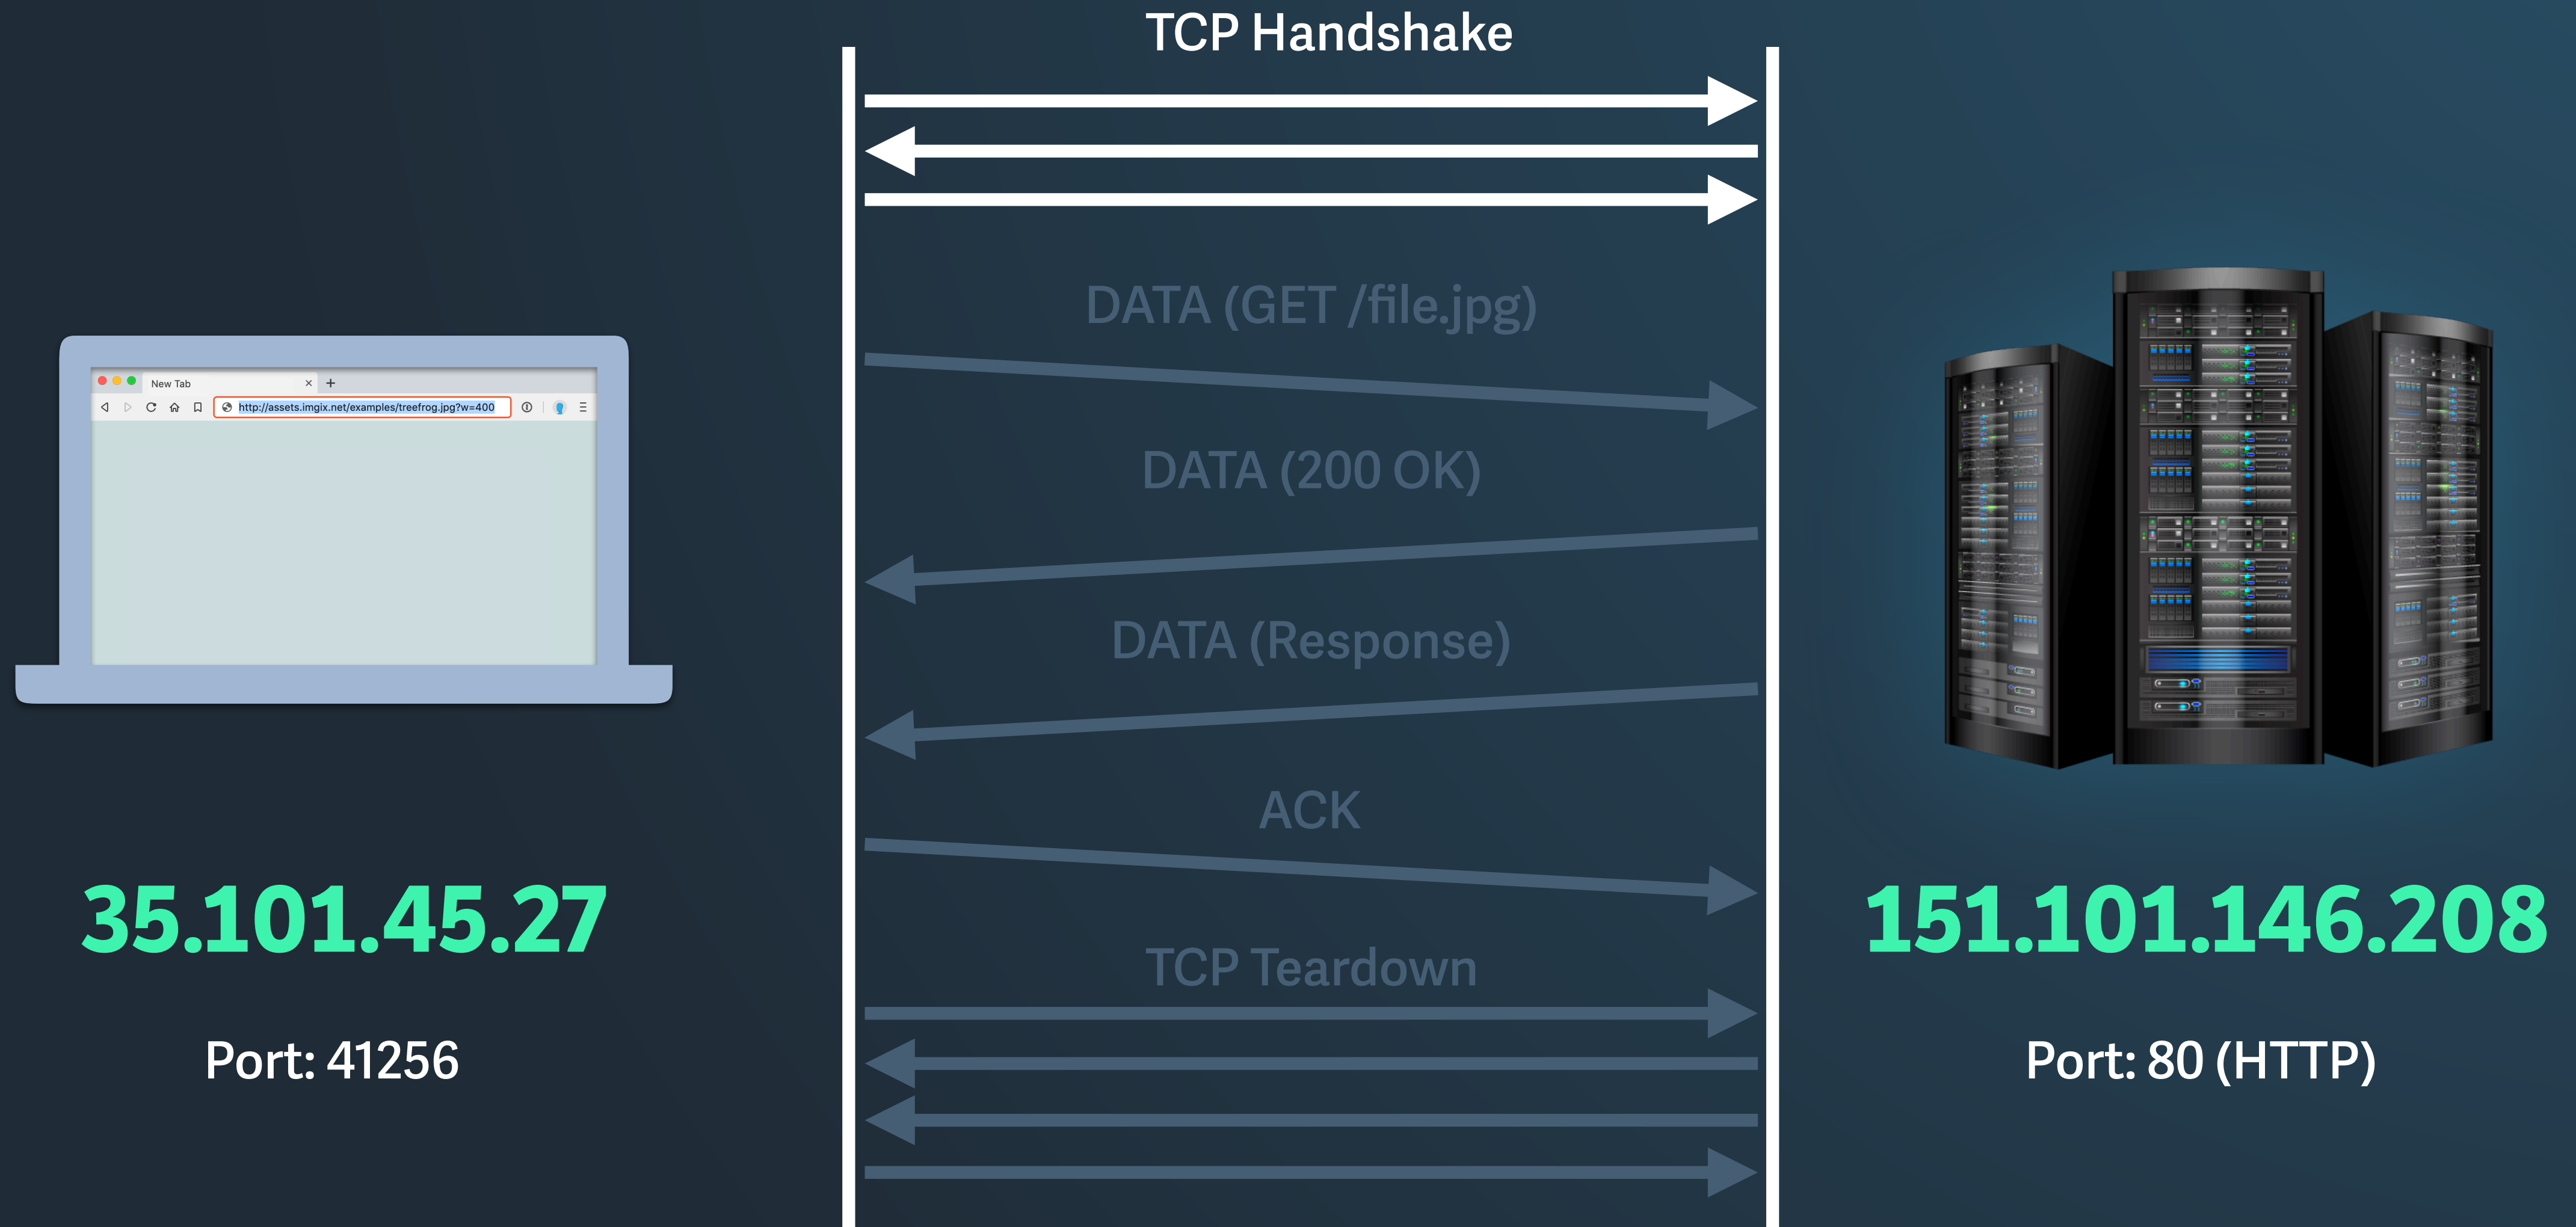


To perform a request, we need a network connection, an IP address and a port.

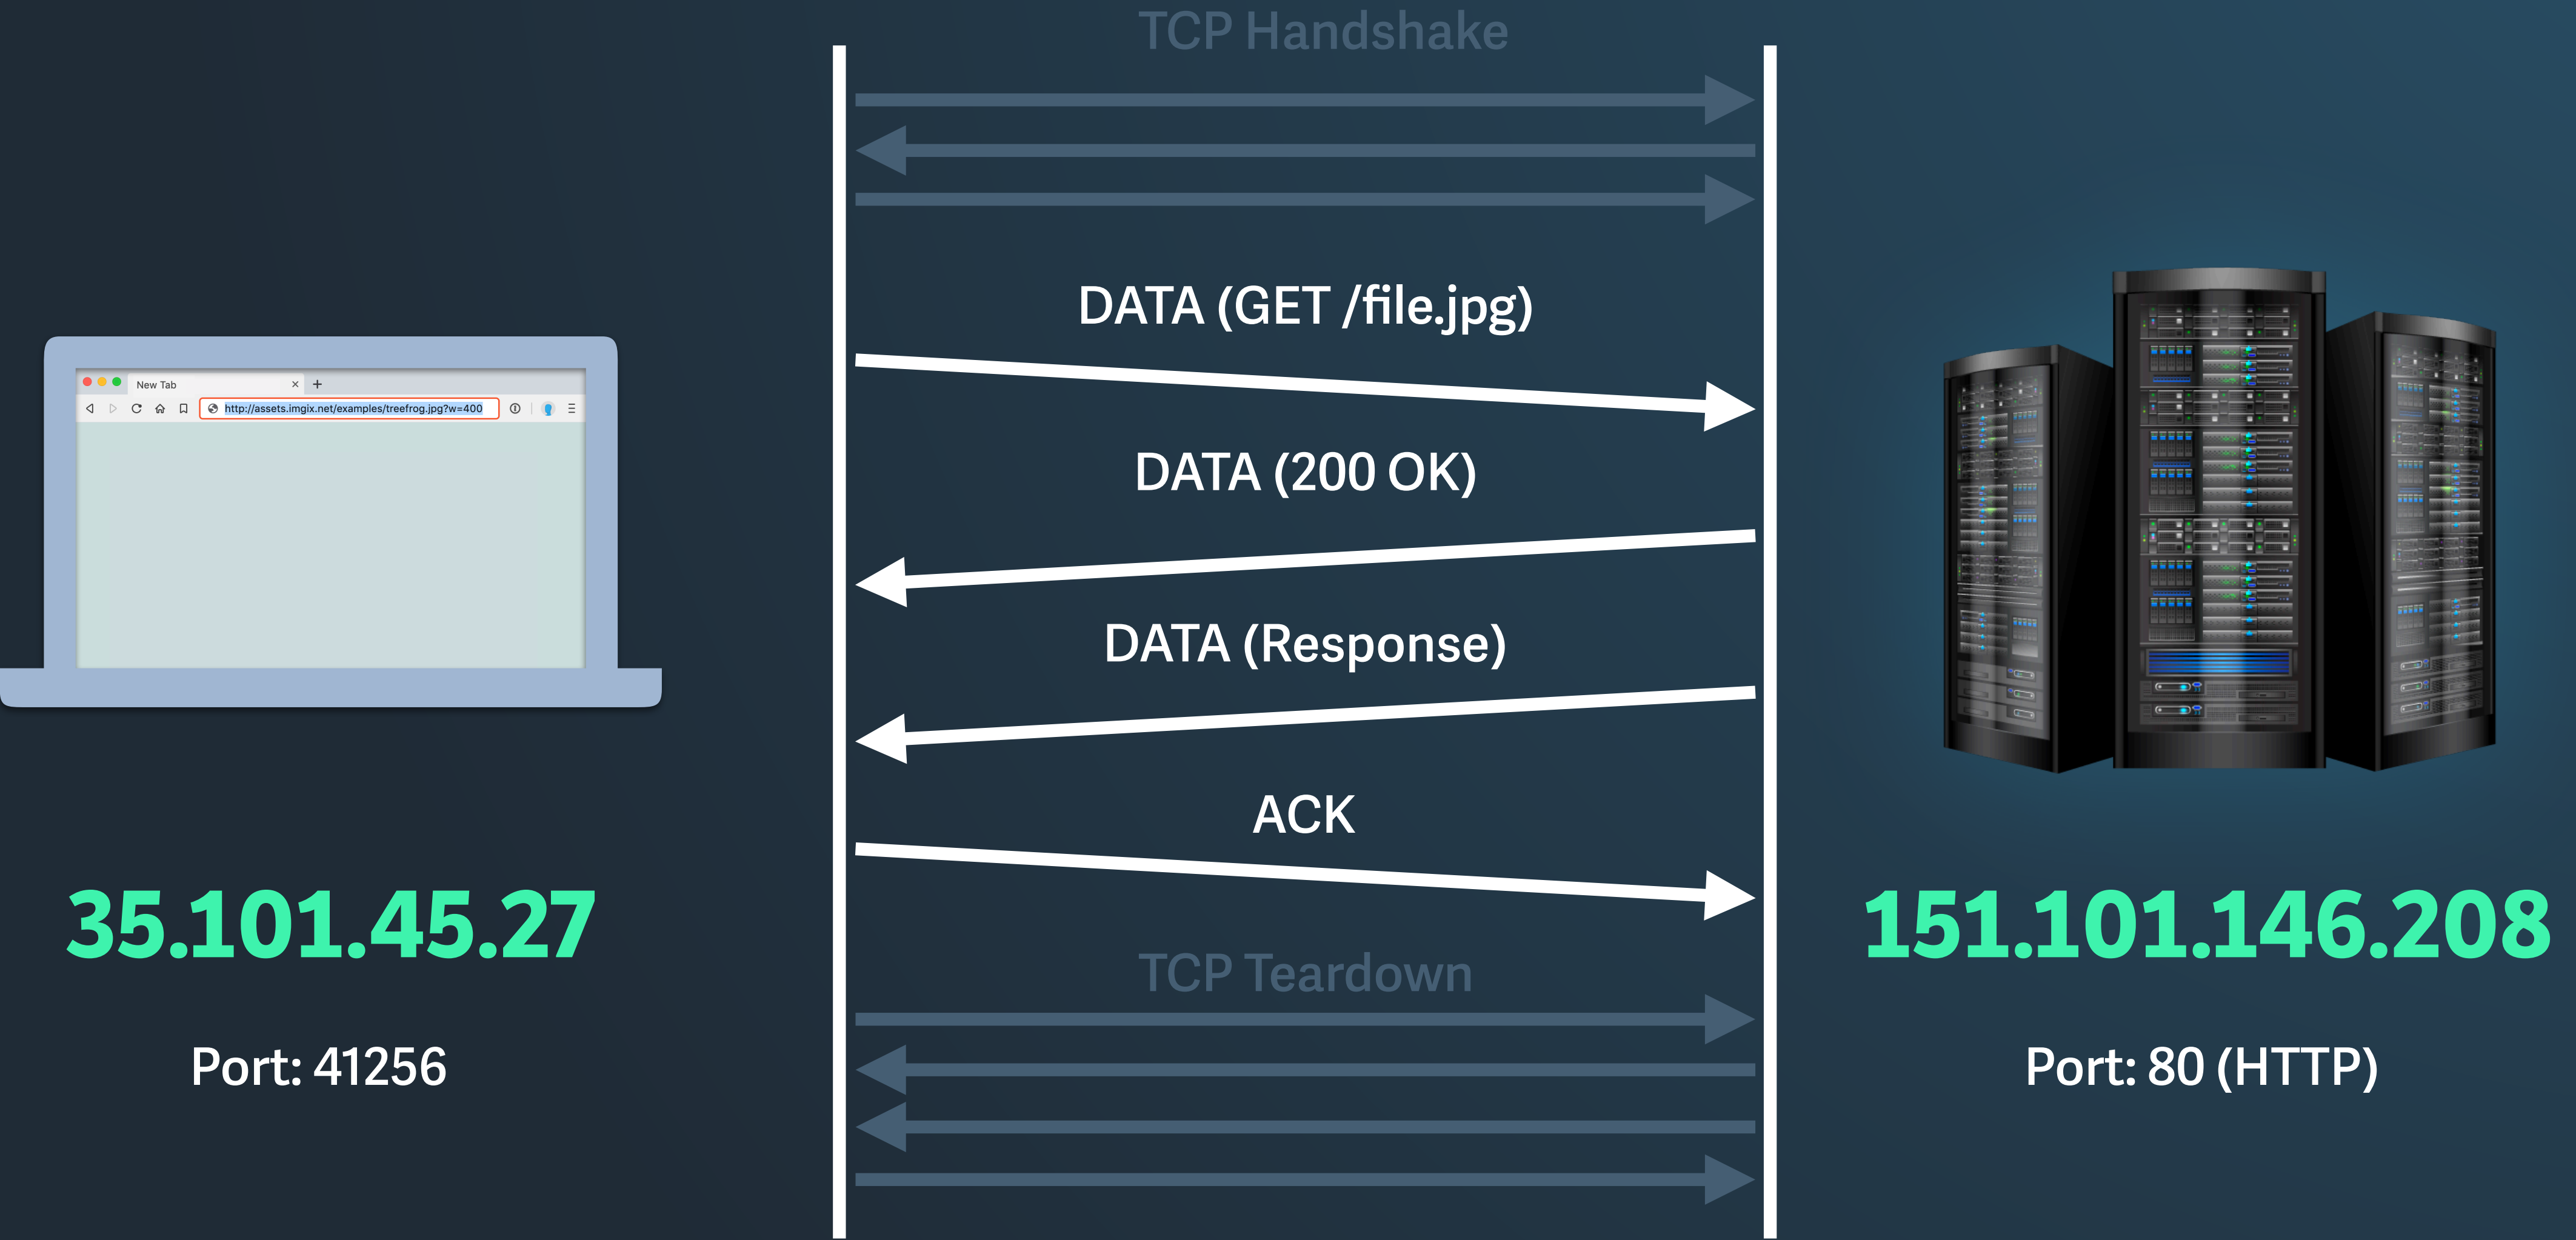





Once we have those things, we can establish a TCP connection.



Packets are sent between client and server to transfer data.

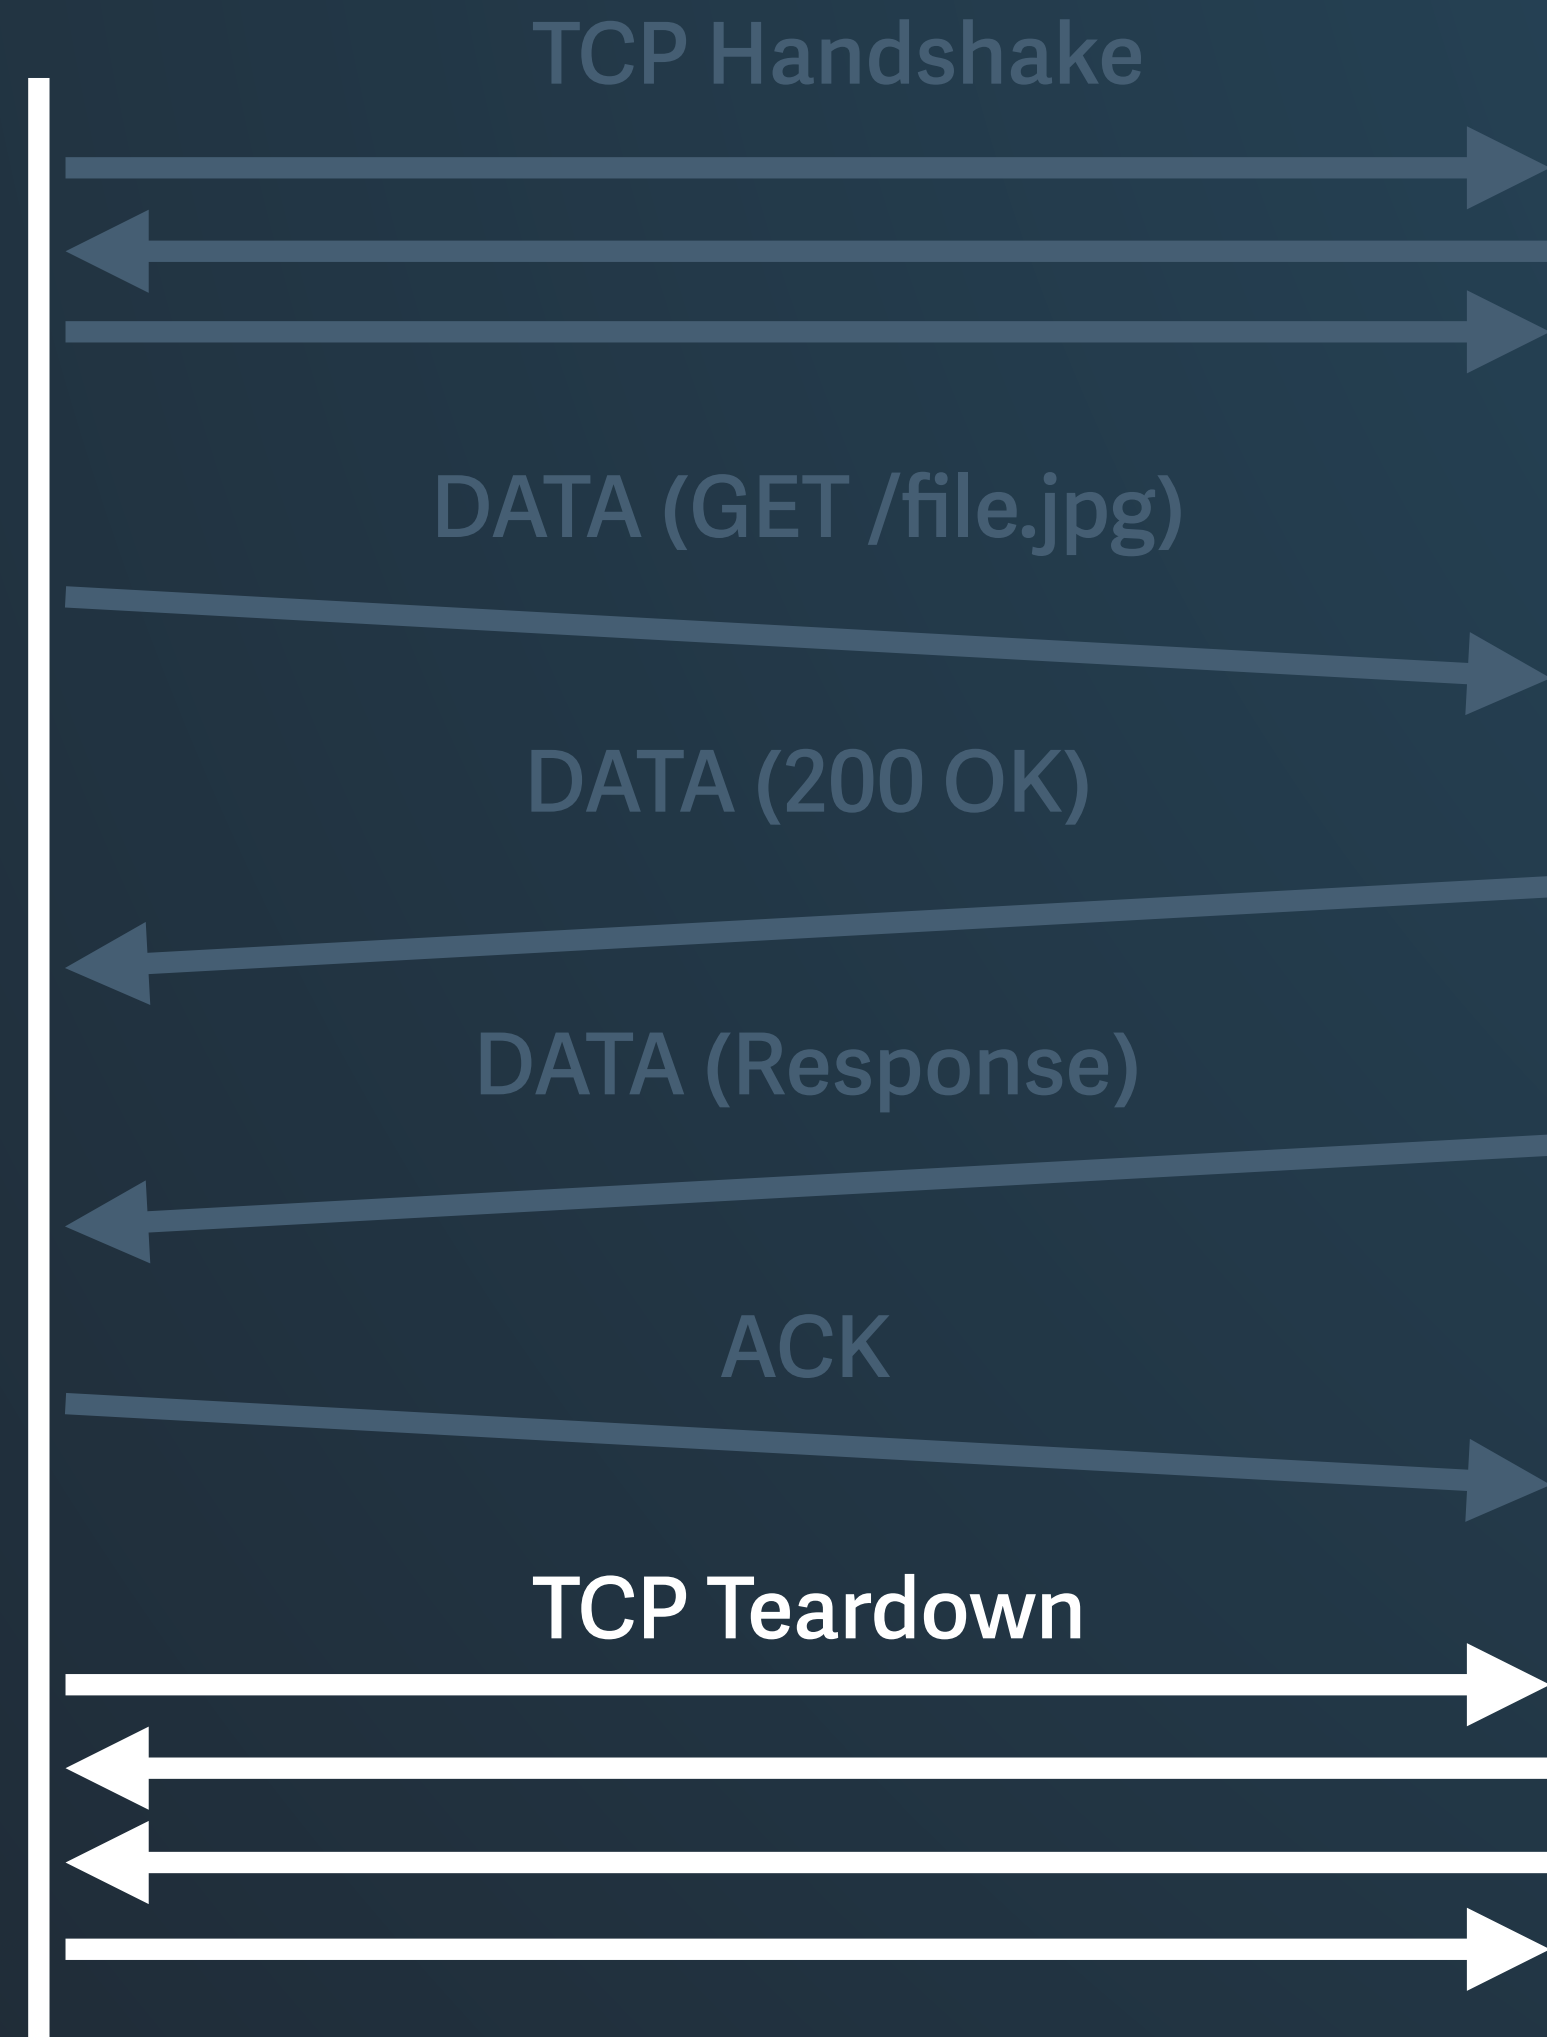


The connection is shut down when it is no longer needed.



**35.101.45.27**

Port: 41256



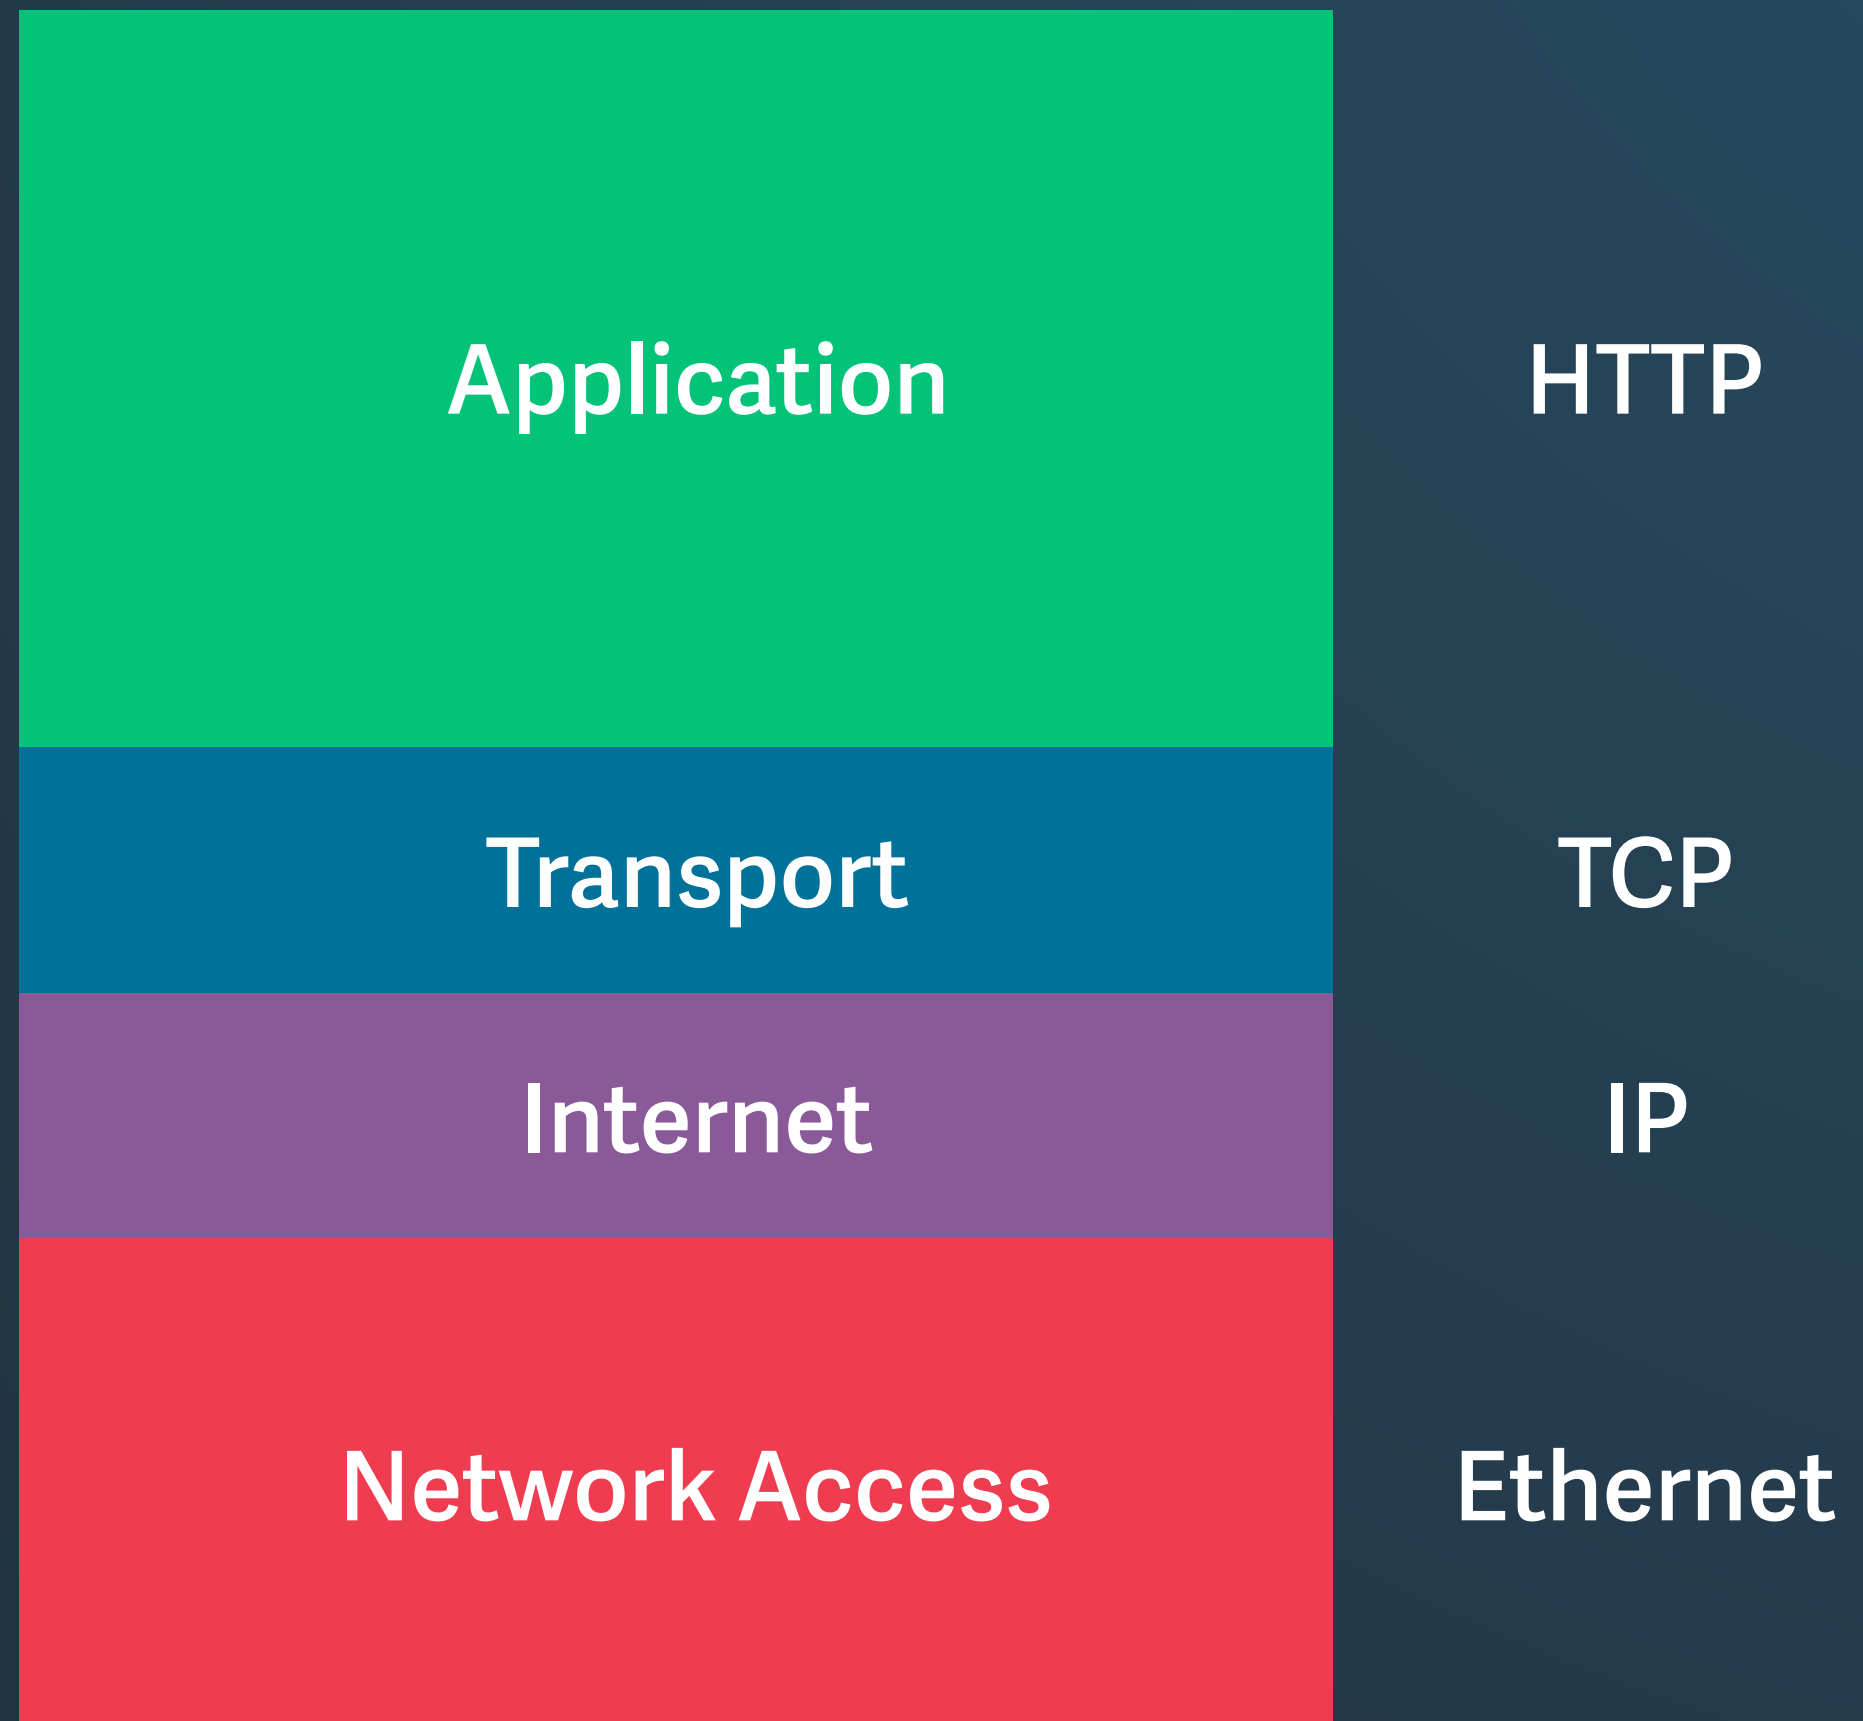
**151.101.146.208**

Port: 80 (HTTP)

**To describe the networking stack,  
we make use of abstract models.**

This is the most common understanding of the network stack.

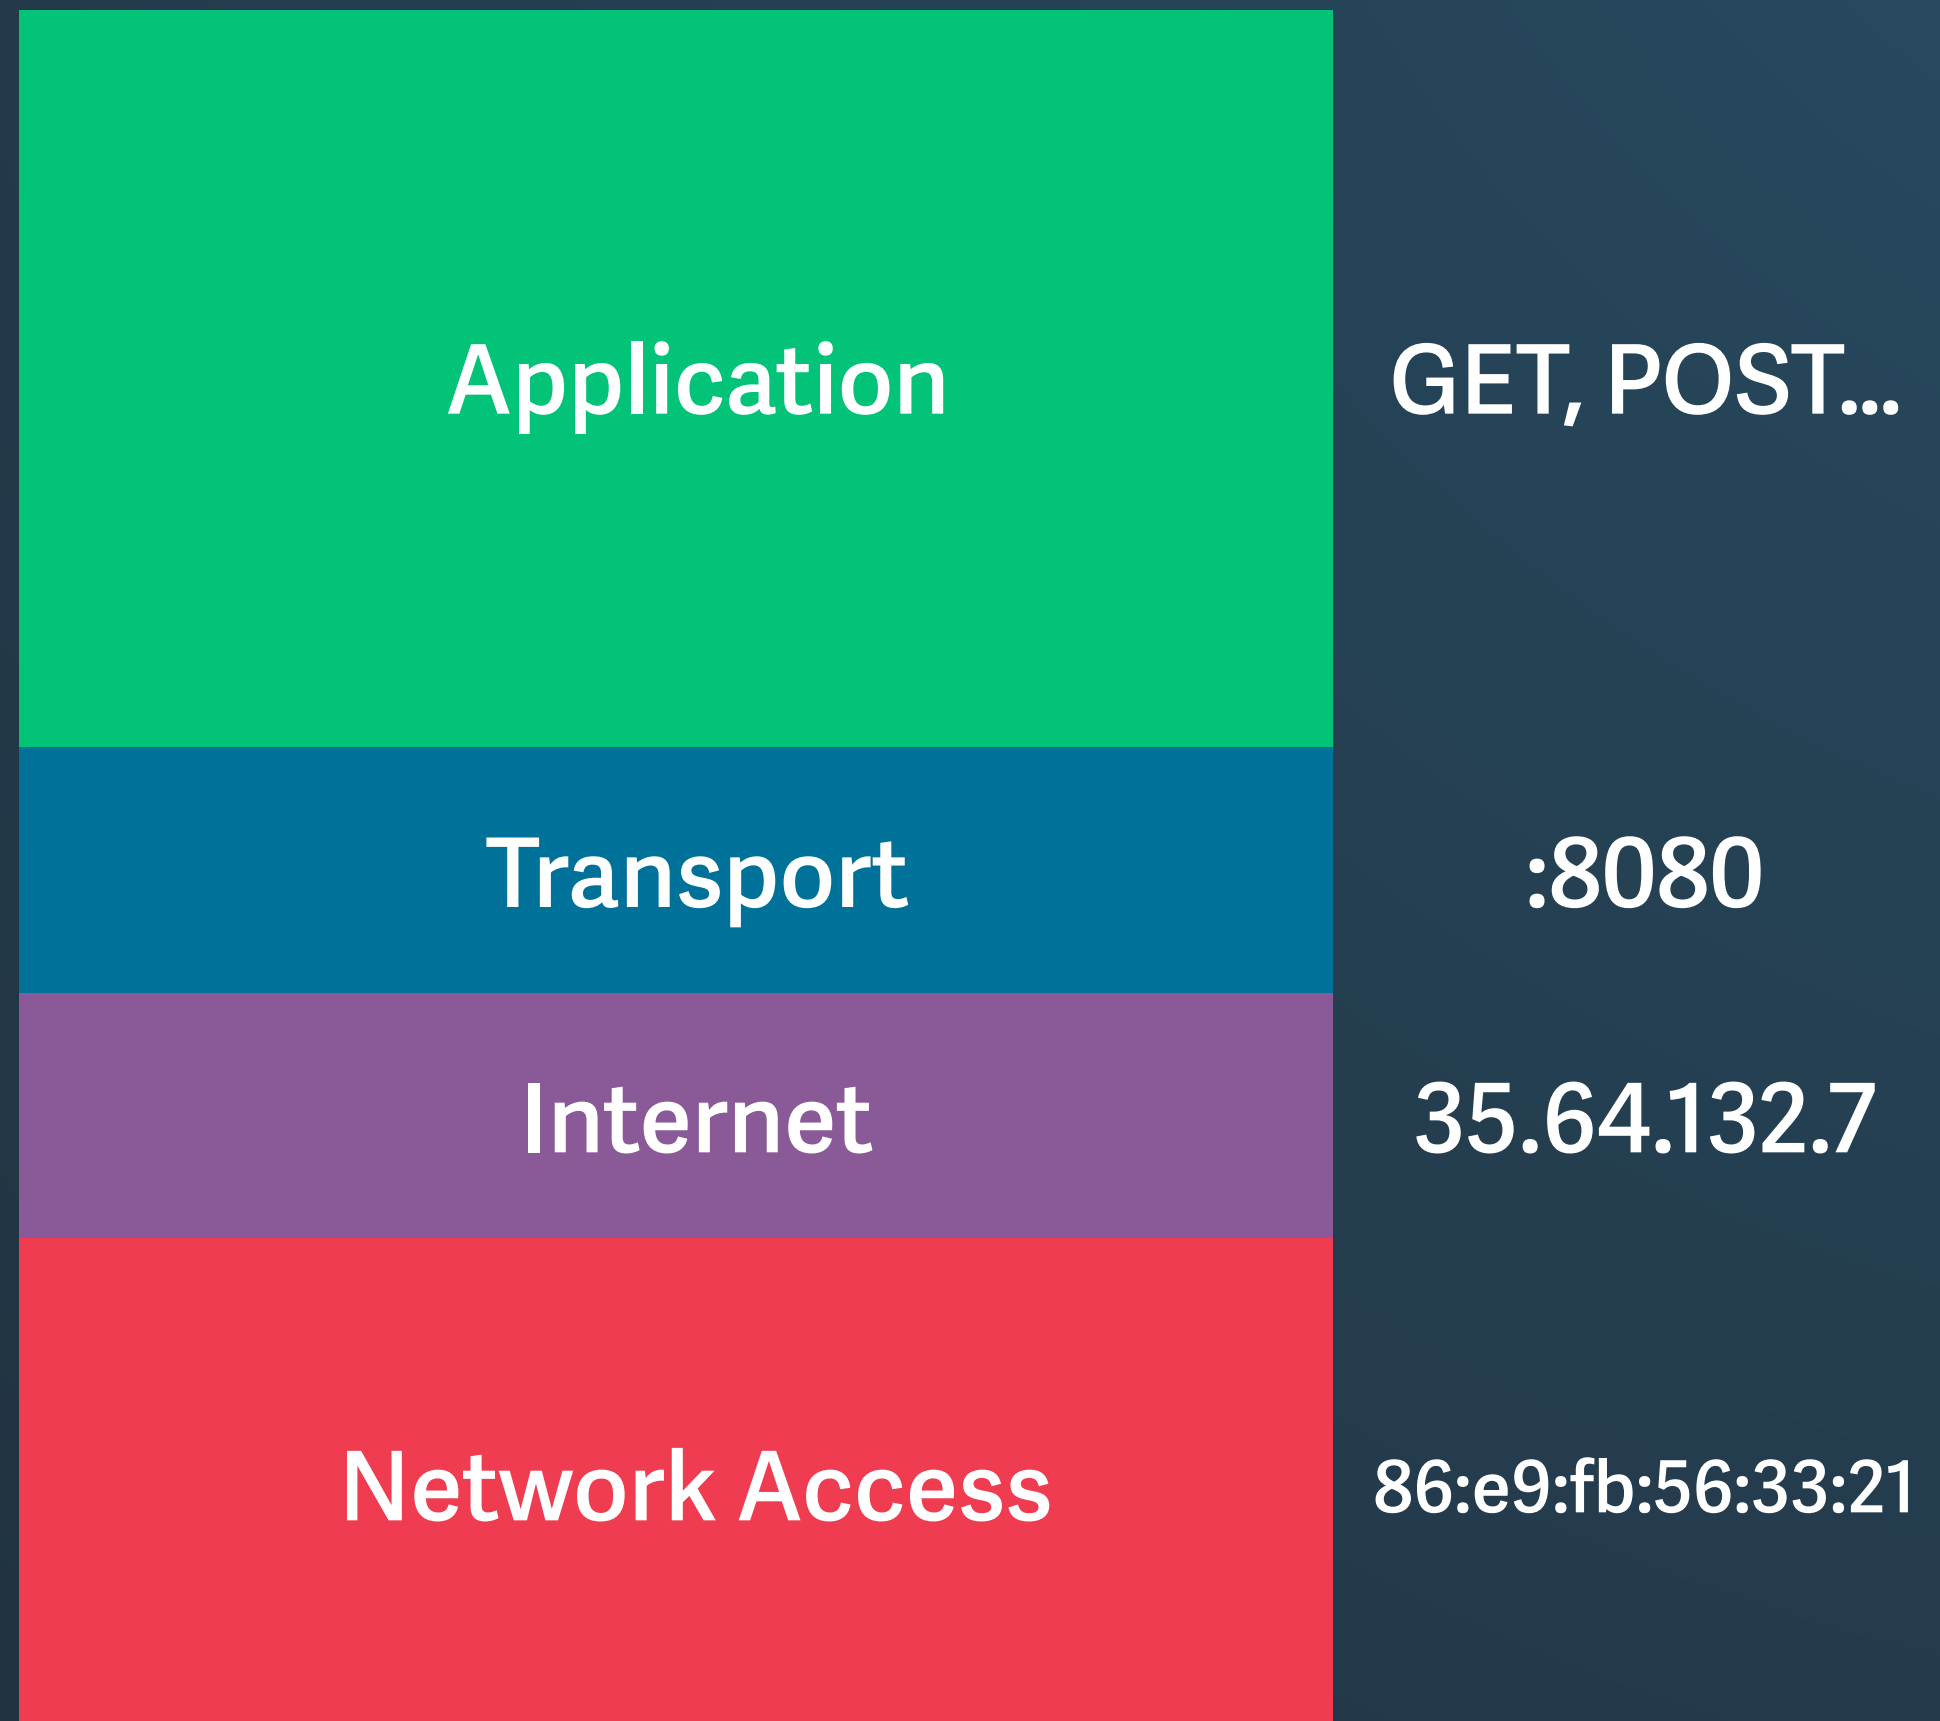
### TCP / IP Model





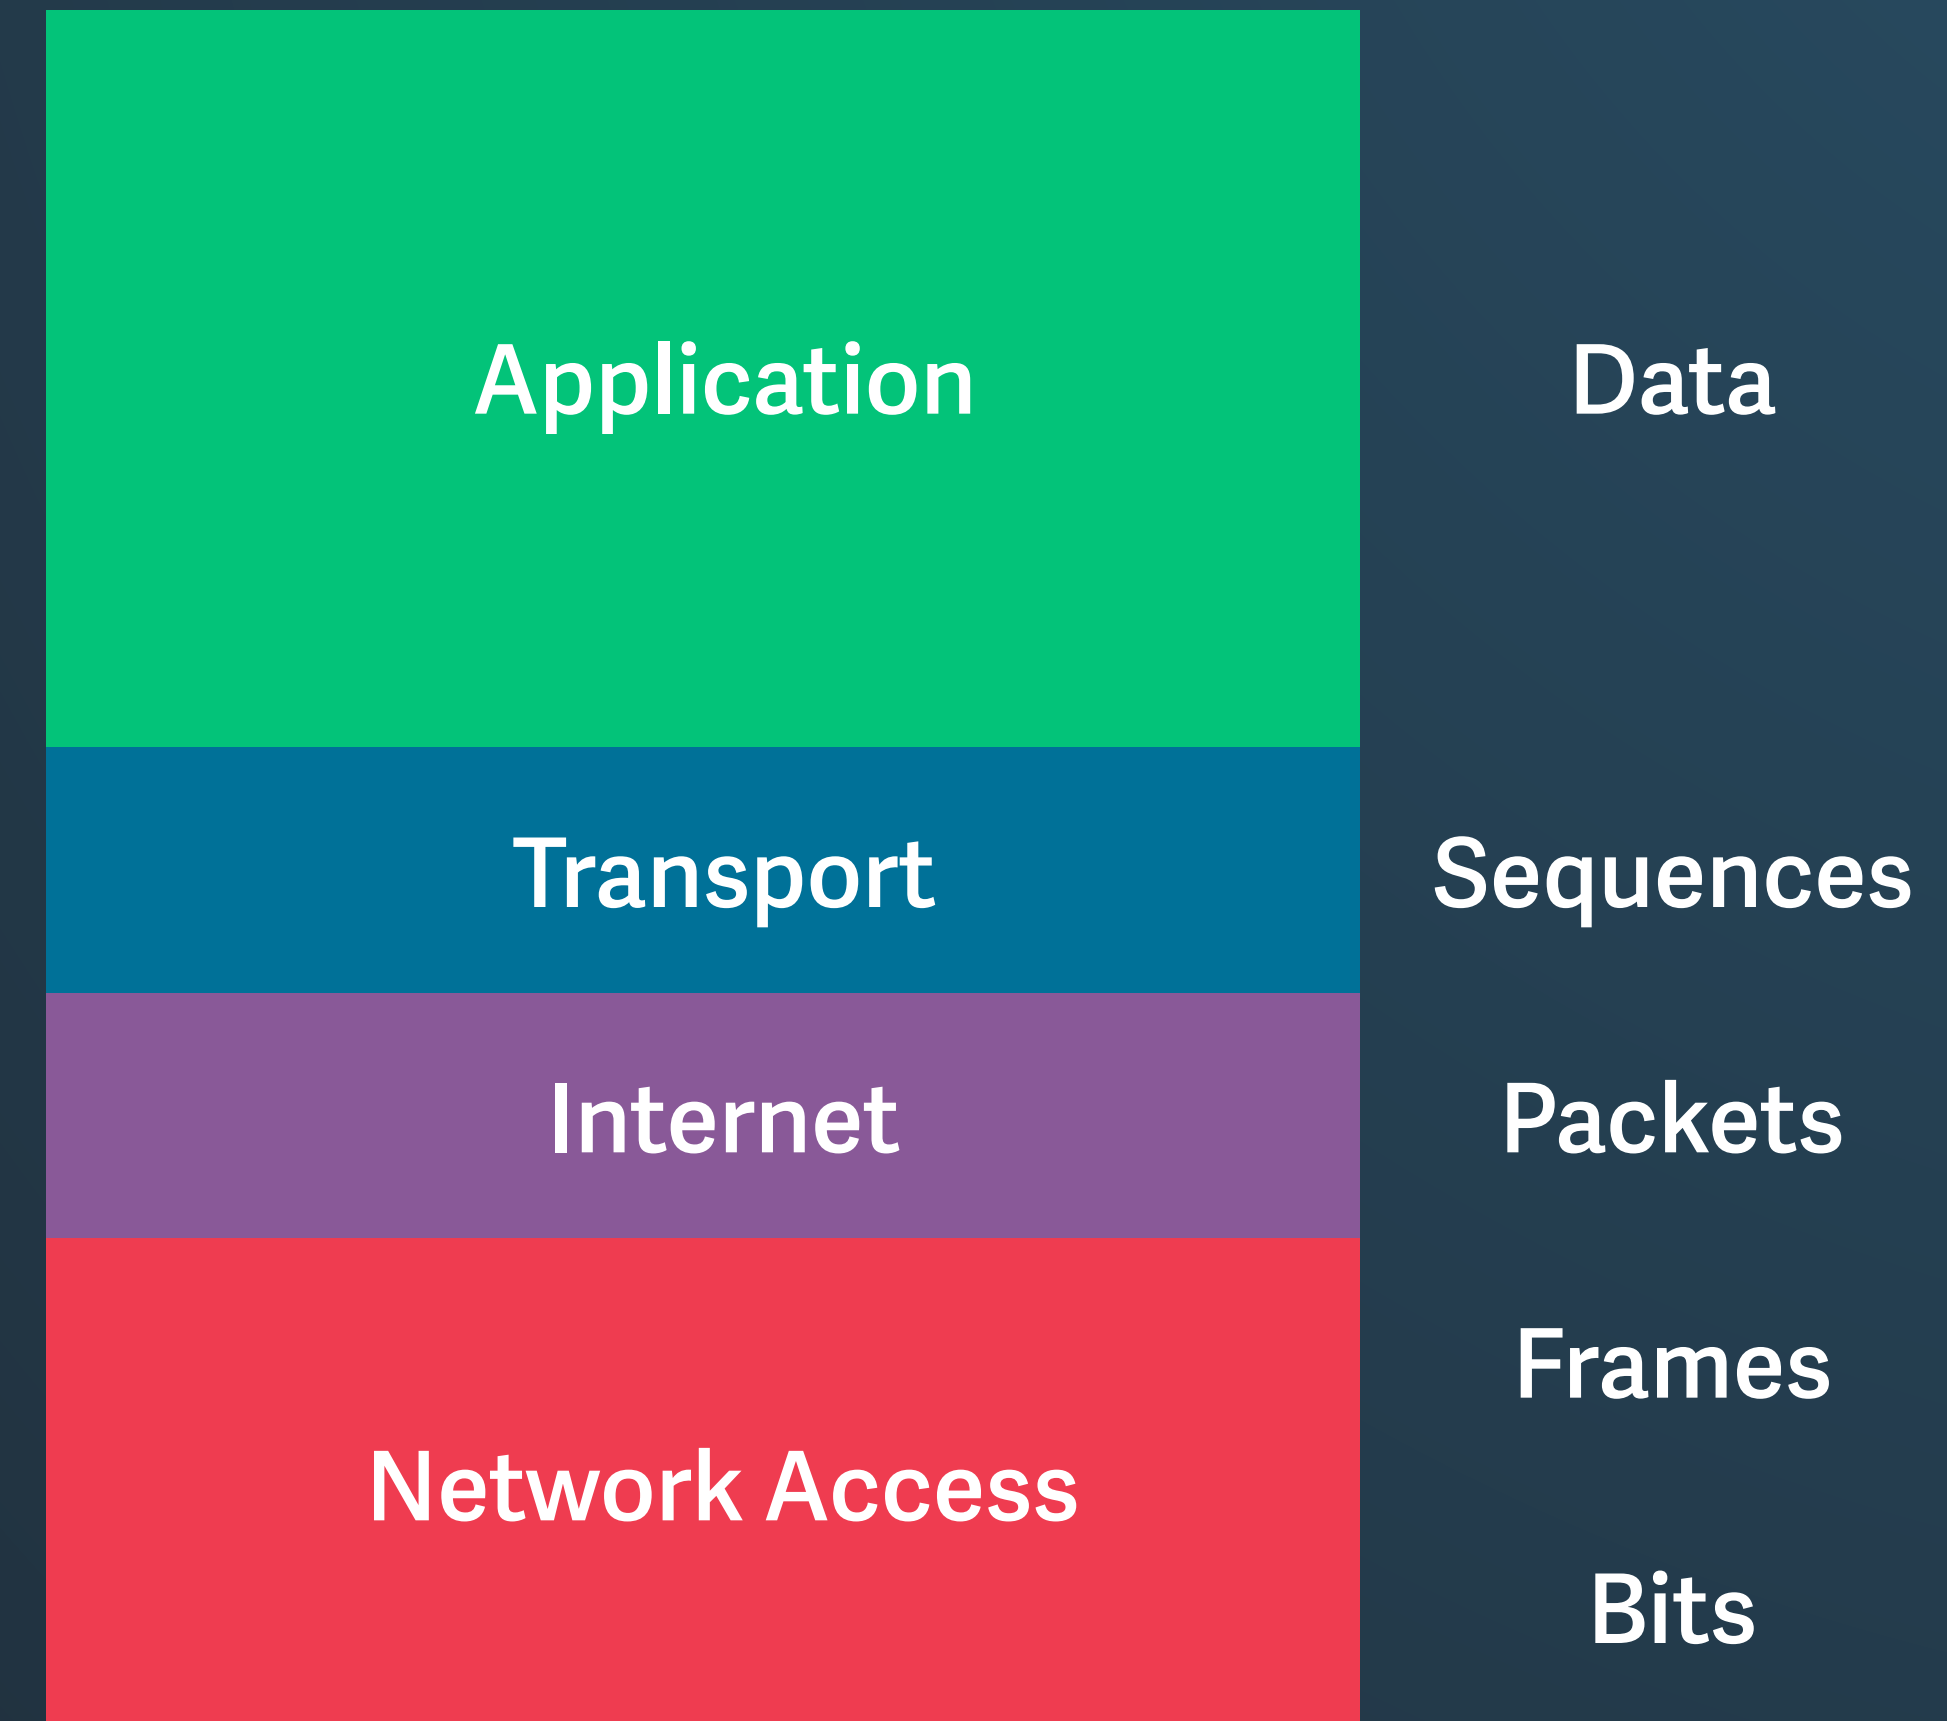
This is the most common understanding of the network stack.

### TCP / IP Model

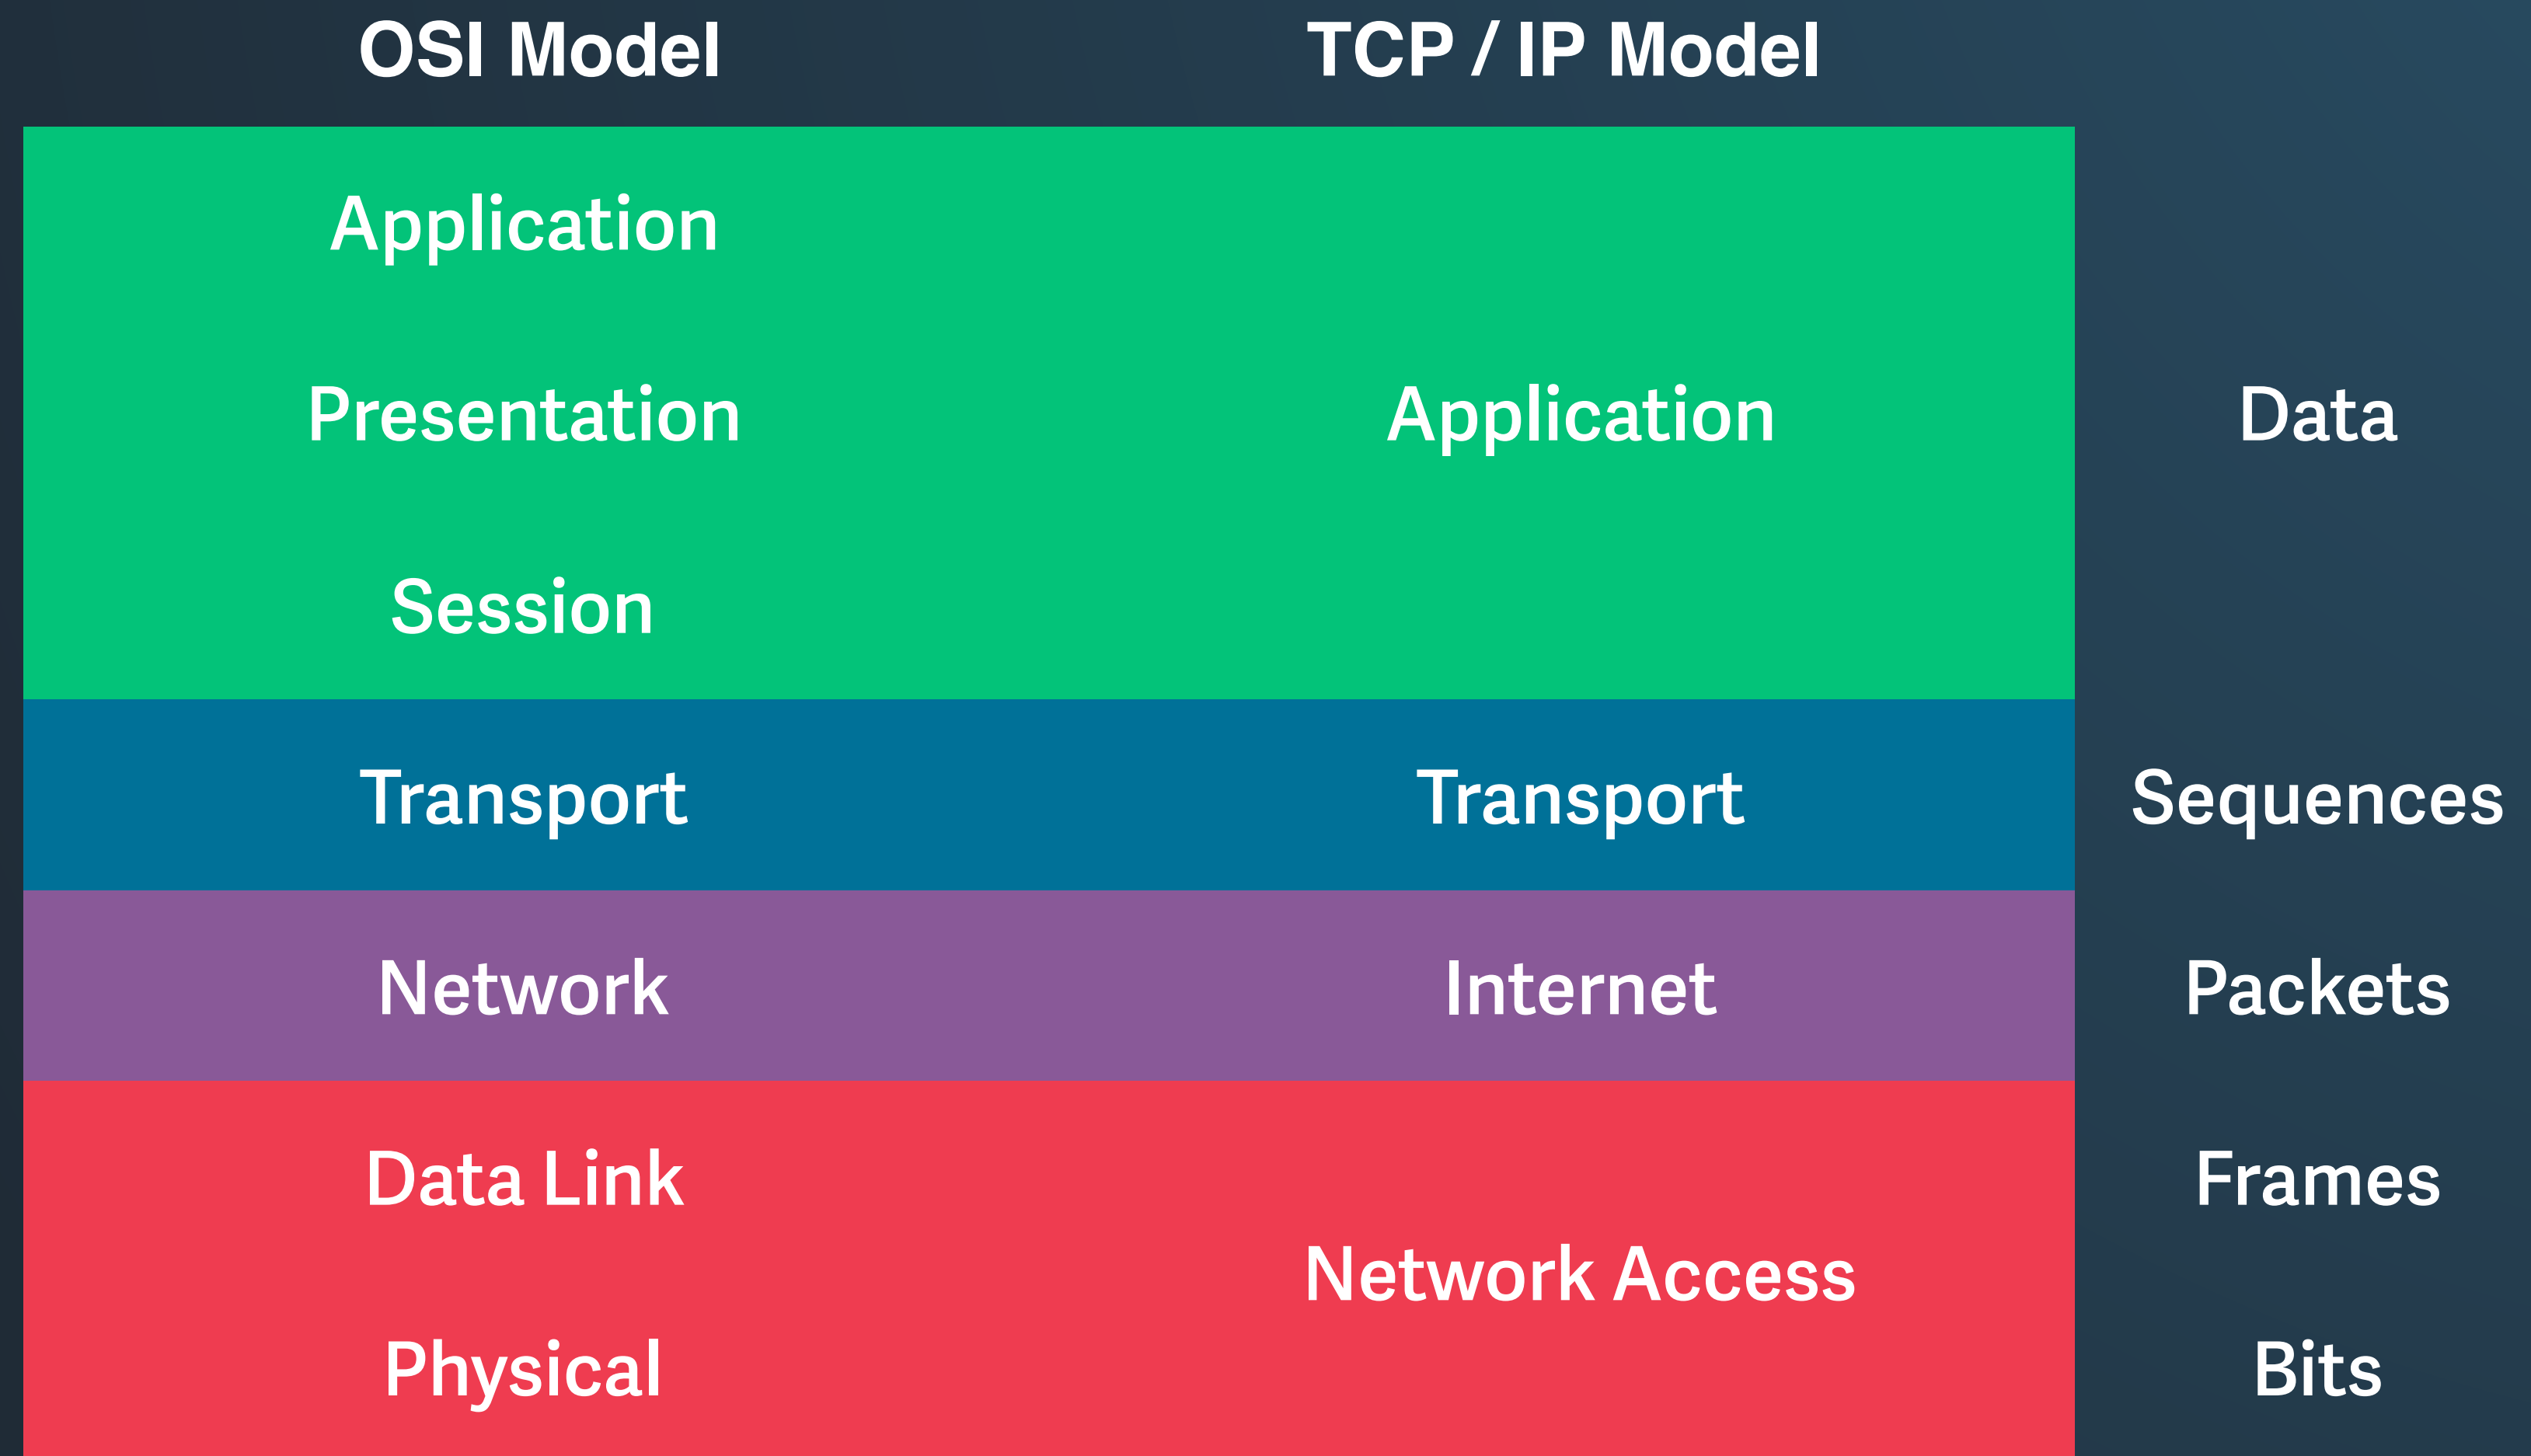


This is the most common understanding of the network stack.

## TCP / IP Model



The OSI model is often used to talk about networking in the abstract.



You may frequently hear about features being “L7”, “L4”, etc.

	OSI Model	TCP / IP Model	
Layer 7	Application	Application	Data
Layer 6	Presentation		
Layer 5	Session		
Layer 4	Transport	Transport	Sequences
Layer 3	Network	Internet	Packets
Layer 2	Data Link	Network Access	Frames
Layer 1	Physical		Bits

**Given that the Internet is stateless, we need to provide a lot of extra information with our request.**



# The data in a web request is encoded as a stream of bytes.

G	E	T		/	e	x	a	m	p	l	e	s	/	t	r	e	e	f	r	o	g	.	j	p	g	?	w	=	4	0	0		H	T	T	P	/	1	.
---	---	---	--	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	--	---	---	---	---	---	---	---

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39

1	\r	\n	H	o	s	t	:		a	s	s	e	t	s	.	i	m	g	i	x	.	n	e	t	\r	\n	U	s	e	r	-	A	g	e	n	t	:		c
---	----	----	---	---	---	---	---	--	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	----	----	---	---	---	---	---	---	---	---	---	---	---	--	---

40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79

u	r	l	/	7	.	5	4	.	0	\r	\n	A	c	c	e	p	t	:		i	m	a	g	e	/	w	e	b	p	,	*	/	*	\r	\n	A	c	c	e
---	---	---	---	---	---	---	---	---	---	----	----	---	---	---	---	---	---	---	--	---	---	---	---	---	---	---	---	---	---	---	---	---	---	----	----	---	---	---	---

80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119

p	t	-	E	n	c	o	d	i	n	g	:		b	r	,	g	z	i	p	,	d	e	f	l	a	t	e	\r	\n	\r	\n
---	---	---	---	---	---	---	---	---	---	---	---	--	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	----	----	----	----

120 121 122 123 124 125 126 127 128 129 130 131 132 133 134 135 136 137 138 139 140 141 142 143 144 145 146 147 148 149 150 151

# The stream is broken apart into chunks.

G	E	T		/	e	x	a	m	p	l	e	s	/	t	r
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

	H	T	T	P	/	1	.	1	\r	\n	H	o	s	t	:
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47

t	\r	\n	U	s	e	r	-	A	g	e	n	t	:		c
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79

p	t	:		i	m	a	g	e	/	w	e	b	p	,	*
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111

i	n	g	:		b	r	,	g	z	i	p	,	d	e	f
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143

e	e	f	r	o	g	.	j	p	g	?	w	=	4	0	0
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

	a	s	s	e	t	s	.	i	m	g	i	x	.	n	e
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63

u	r	l	/	7	.	5	4	.	0	\r	\n	A	c	c	e
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95

/	*	\r	\n	A	c	c	e	p	t	-	E	n	c	o	d
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127

l	a	t	e	\r	\n	\r	\n
144	145	146	147	148	149	150	151

Those chunks are given IDs to keep them in order.

SEQ 0

G	E	T		/	e	x	a	m	p	l	e	s	/	t	r
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

SEQ 16

e	e	f	r	o	g	.	j	p	g	?	w	=	4	0	0
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

SEQ 32

	H	T	T	P	/	1	.	1	\r	\n	H	o	s	t	:
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47

SEQ 48

	a	s	s	e	t	s	.	i	m	g	i	x	.	n	e
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63

SEQ 64

t	\r	\n	U	s	e	r	-	A	g	e	n	t	:		c
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79

SEQ 80

u	r	l	/	7	.	5	4	.	0	\r	\n	A	c	c	e
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95

SEQ 96

p	t	:		i	m	a	g	e	/	w	e	b	p	,	*
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111

SEQ 112

/	*	\r	\n	A	c	c	e	p	t	-	E	n	c	o	d
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127


SEQ 128

i	n	g	:		b	r	,	g	z	i	p	,	d	e	f
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143

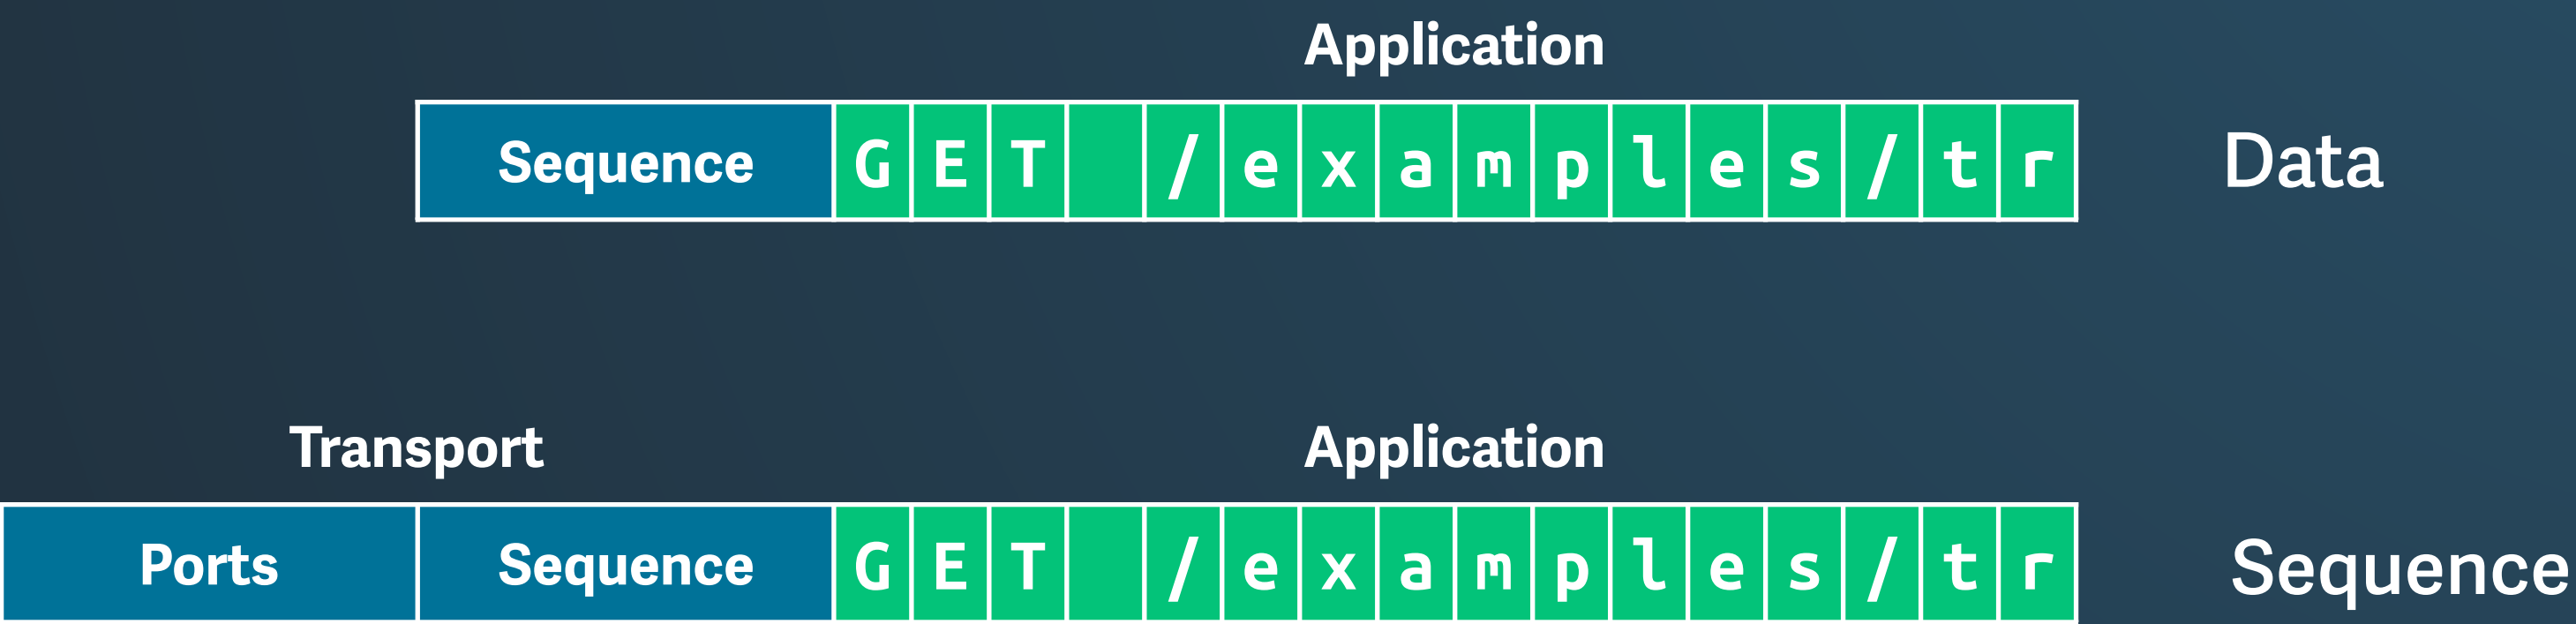
SEQ 144

l	a	t	e	\r	\n	\r	\n
144	145	146	147	148	149	150	151

The chunks need to be packaged for sending across the Internet.

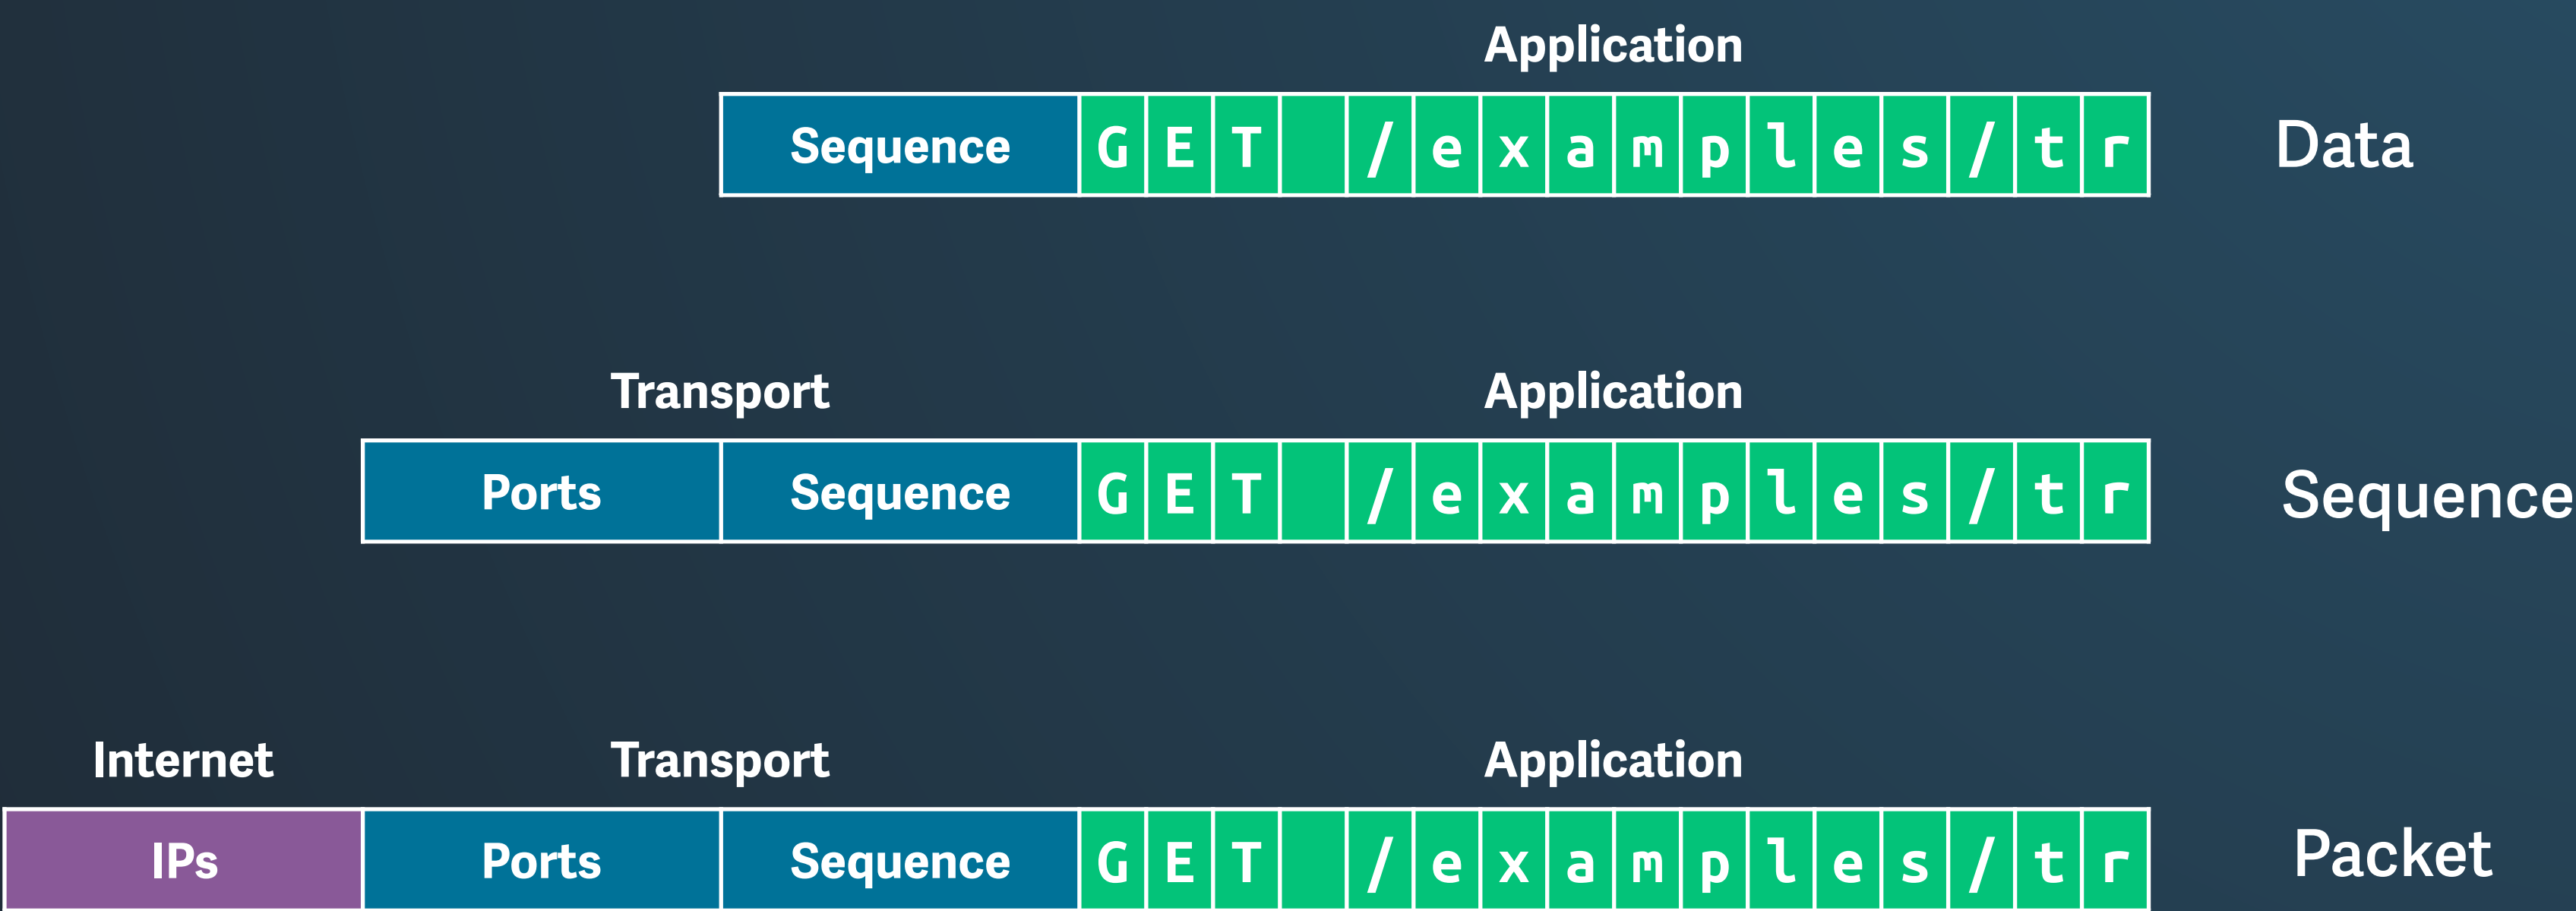


The TCP headers are added to handle the connection between two hosts.

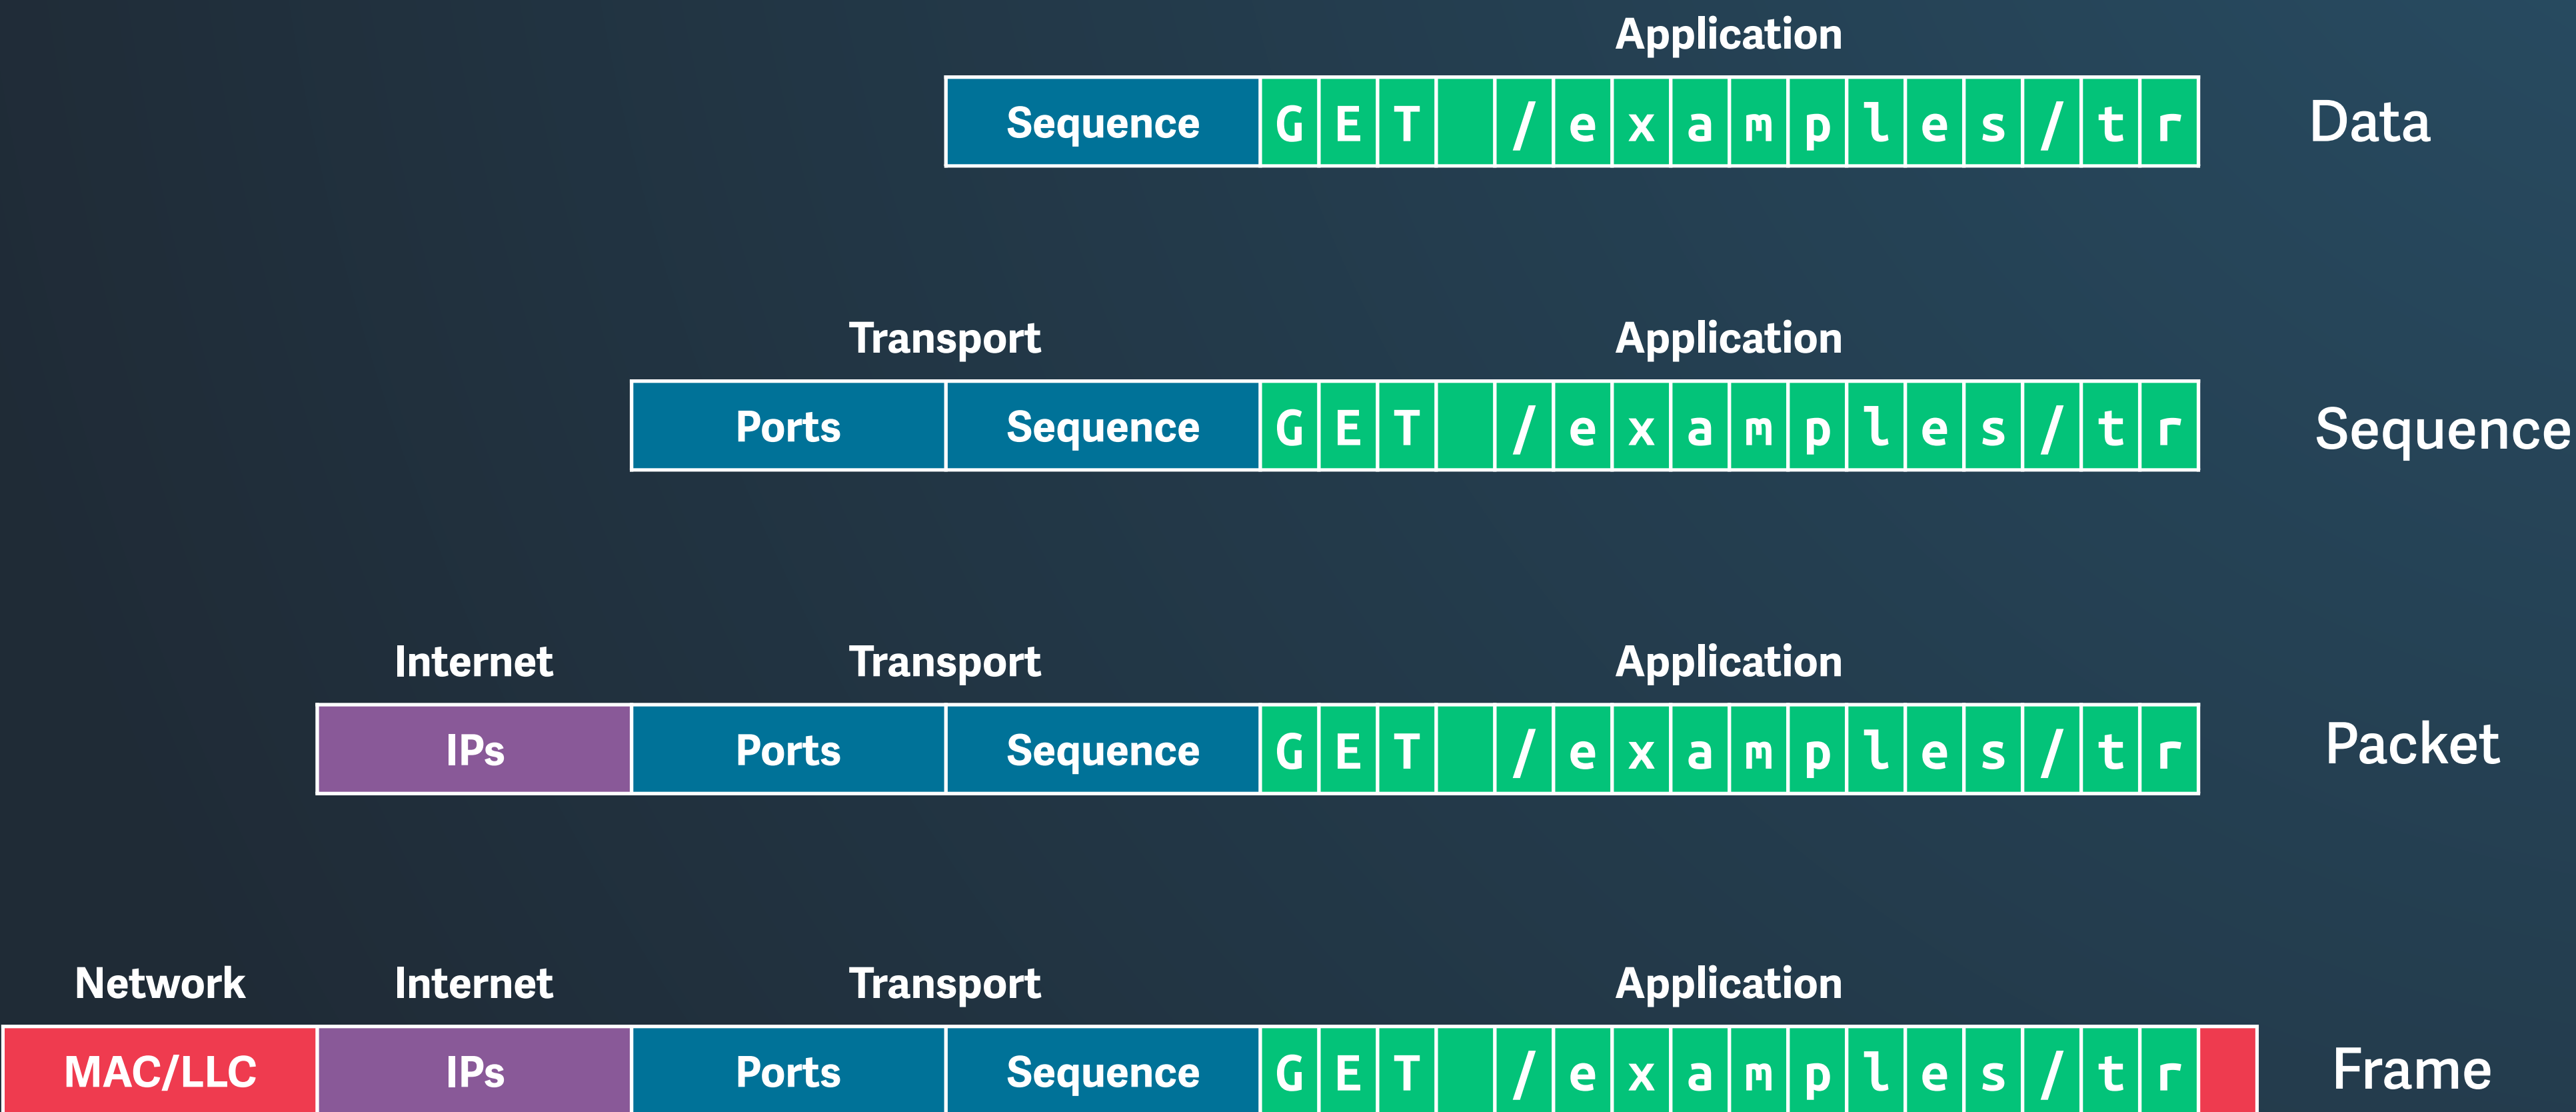





# The IP headers provide the necessary context to navigate the Internet.



The MAC/LLC headers manage how packets navigate the local network.



The MAC/LLC headers manage how packets navigate the local network.



The standard “MTU” or Maximum Transmission Unit is 1526 bytes.

**To understand how data is sent over a physical wire,  
we are going to work backwards from the atom.**

Electrically conductive atoms have a single electron in their outermost shell.




By adding poles with positive and negative charges...







... we cause the atoms nearest to the positive pole to lose an electron.




Atoms missing an electron pull electrons from other atoms.




The negative pole provides electrons to atoms that end up without one.




This creates an electrical current.



In direct current, electrons flow in one direction at a rate known as voltage.




A higher voltage means more electrons are flowing in the current.







The voltage can be made to vary over time.



A measurement of the voltage is taken regularly on a clock cycle.



By thresholding the voltages to discrete values, we can derive bits.

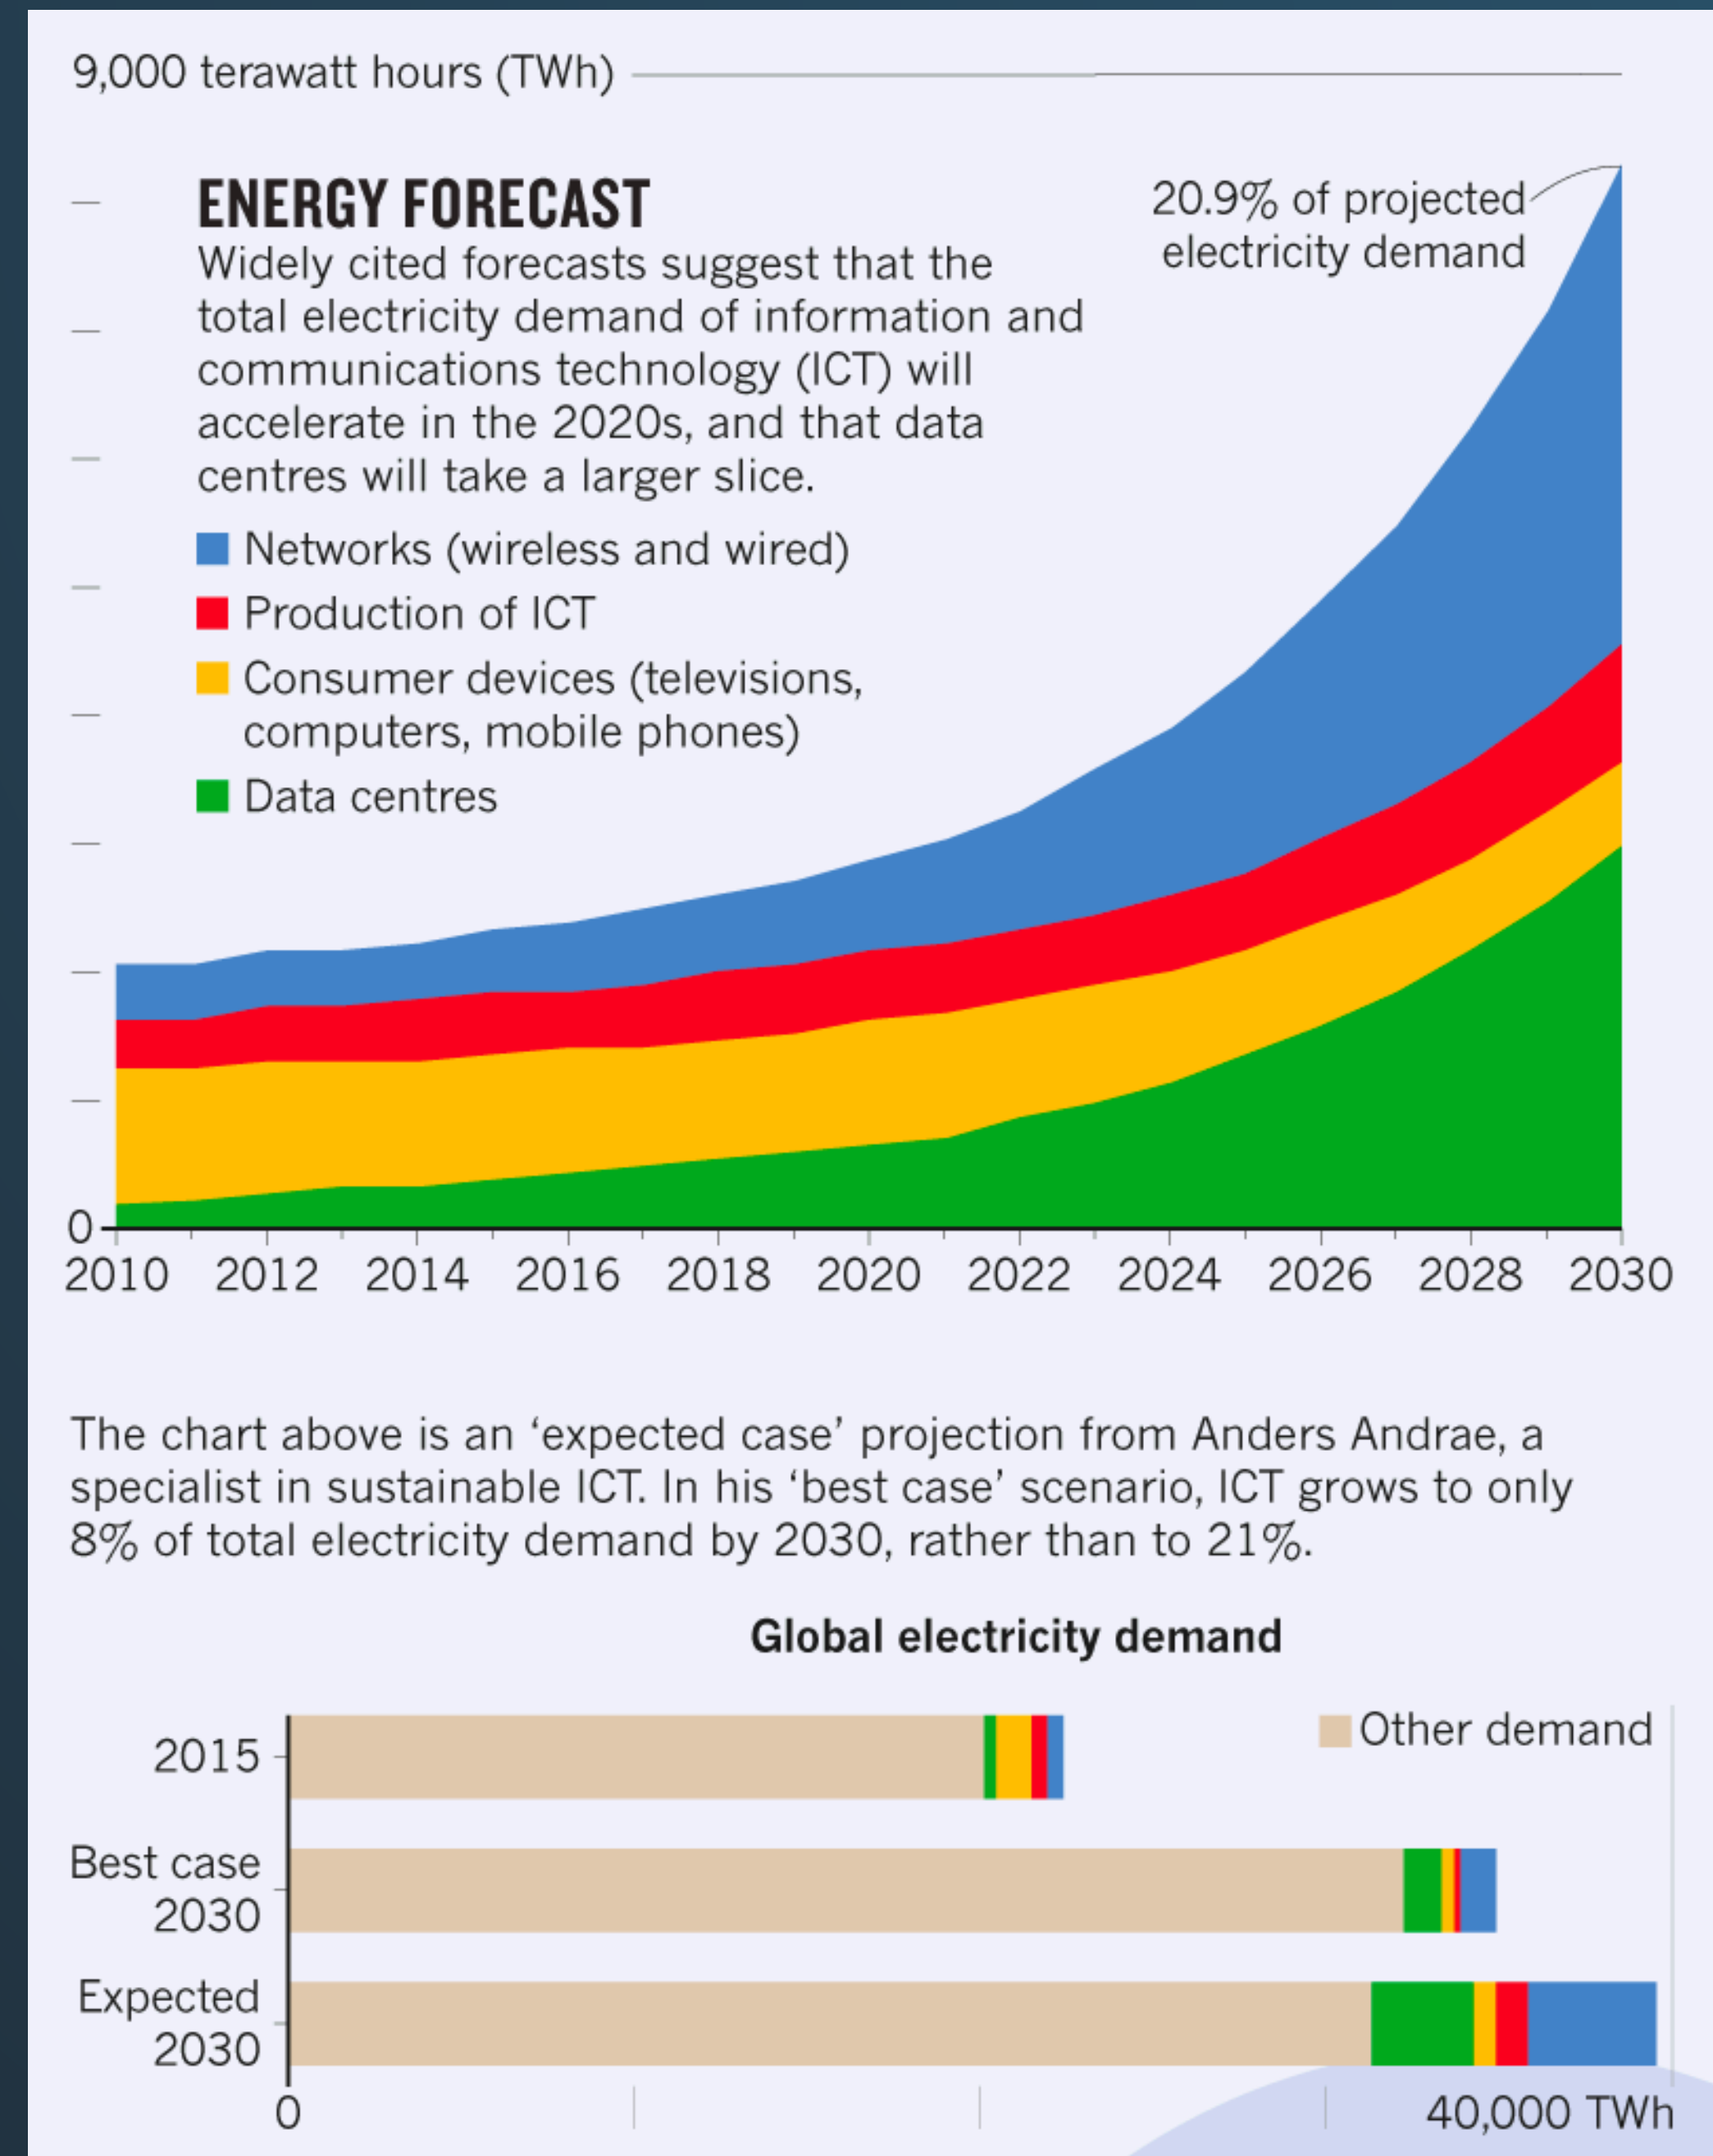


01011010 = "Z"


					0 0 0	0 0 1	0 1 0	0 1 1	1 0 0	1 0 1	1 1 0	1 1 1
					0	1	2	3	4	5	6	7
					0	1	2	3	4	5	6	7
0	0	0	0	0	NUL	DLE	SP	0	@	P	`	p
0	0	0	1	1	SOH	DC1	!	1	A	Q	a	q
0	0	1	0	2	STX	DC2	"	2	B	R	b	r
0	0	1	1	3	ETX	DC3	#	3	C	S	c	s
0	1	0	0	4	EOT	DC4	\$	4	D	T	d	t
0	1	0	1	5	ENQ	NAK	%	5	E	U	e	u
0	1	1	0	6	ACK	SYN	&	6	F	V	f	v
0	1	1	1	7	BEL	ETB	'	7	G	W	g	w
1	0	0	0	8	BS	CAN	(	8	H	X	h	x
1	0	0	1	9	HT	EM	)	9	I	Y	i	y
1	0	1	0	10	LF	SUB	*	:	J	Z	j	z
1	0	1	1	11	VT	ESC	+	;	K	[	k	{
1	1	0	0	12	FF	FS	,	<	L	\	l	
1	1	0	1	13	CR	GS	-	=	M	]	m	}
1	1	1	0	14	SO	RS	.	>	N	^	n	~
1	1	1	1	15	SI	US	/	?	O	_	o	DEL

**Over 2 TRILLION GIGABYTES of data will be  
transmitted over the Internet in 2020.**

It will take a lot of electricity to transit that much data.








**Carbon emissions due to the Internet surpassed those of the airline industry in 2015.**

**Do your part. Use fewer bytes.**

# The Hidden Language of Computer Hardware and Software



Charles Petzold

**Thank you.**

**Chris Zacharias**

[chris@imgix.com](mailto:chris@imgix.com)

[@zacman85](#)