

# Computer Networks

## X\_400487

Lecture 9

Chapter 7: The Application Layer

Q: Latency and packet loss on the Internet today?

Packet loss < 0.4% in most cases, latency below 50ms in most cases\*



Lecturer: Jesse Donkervliet  
with slides from Lin Wang



\*source: <https://data.eric.edu/overload/measuring-bandwidth-america-2011-2021-Fred-Measuring-Bandwidth-America-Report.pdf>

# Online Games

A challenge for computer networks: real-time interactive systems

## UDP or TCP? What is best for your app?

Discussed in class through a series of comparison videos available as part of a blog post at [https://gafferongames.com/post/deterministic\\_lockstep/](https://gafferongames.com/post/deterministic_lockstep/). The entire blog is highly recommended reading.


## Milestone reached!

Creating large-scale distributed systems is difficult!

We can now **start** building applications and systems that communicate over a network!

Advanced courses unlocked:


1. Advanced Systems Programming
2. Advanced Computer Networks
3. Distributed systems (also requires Computer Organization and Operating Systems)



Now we can finally build applications and no longer worry about networking!

Or so we thought!

SSH IMAP FTP  
RPC MQTT SMTP  
XMPP HTTP QUIC




Source: <https://www.youtube.com/watch?v=Qm0jKjG0v00> Does QUIC Make the Web Faster? @USECONV2016

## Where The Application Layer Sits

Application layer protocols are often part of an "app"


But they don't need a GUI, e.g., DNS



### Application Layer Messages


Application layer messages are often split over multiple packets

Or may be aggregated in a packet




### Application Communication Needs

Vary widely with app; must build on Transport services



### Evolution of Internet Applications

Always changing, and growing...



### Every minute:

- 241M emails sent
- 6.3M Zoom meeting minutes
- 649k hours video viewed on YouTube
- 10.4M viewing minutes on Instagram
- 34k Slack Messages
- ...



Created by: @Discovery Today & LTMG

### Application Layer Topics

1. Domain Name System (DNS)
2. Email
3. Web (HTTP, Web caching/proxy)
4. Multimedia applications

### Domain Name System

Port 53

## Domain Name System

An application used by the network itself!

Machines on the internet are identified by their **IP address**

These addresses are difficult for humans to remember!

Q: Can you think of another disadvantage?


<http://4.31.198.44/rfc1035.txt>  
[j.i.r.donkervliet@131.180.77.82](mailto:j.i.r.donkervliet@131.180.77.82)

DNS translates **human readable names** to IP addresses

19

## DNS name space

**Hierarchical structure.**




Top level domains controlled by Internet Corporation for Assigned Names and Numbers (ICANN).

20

## DNS name space

**Hierarchical structure.**




Organizations can request second-level domains from **registrars**.

21

## DNS name space

**Hierarchical structure.**




Countries have their own second-level domains.

22

## DNS name space

**Hierarchical structure.**



'edu' and 'gov' are generic domains, but mostly used by organizations in the United States.

23

## DNS name space

If you control a domain, you can specify arbitrary subdomains.

United Kingdom uses **ac.uk**, for academic use and **co.uk**, for commercial use.

The Netherlands puts everything directly under **.nl**.


24

## Name servers

Q: How does Alice's machine know where to find the name server?

```
DNS server assignment: Automatic (DHCP)
Link speed (Receive/Transmit): 100/100 (Mbps)
IPv4 address: 2a02:444:10891:5b2:ae64:d791:7645
Link-local IPv4 address: fe80:444c:40000:76c1:3535
IPv4 DNS server: 2a02:444:10891:5b2:ae64:d791:7645
IPv4 DNS server: 2a02:444:10891:5b2:ae64:d791:7645
IPv4 address: fe80:444c:40000:76c1:3535
IPv4 DNS server: fe80:444c:40000:76c1:3535
```

To translate a domain name to an IP address, you ask a **name server**.



Q: How does the name server know the location of ietf.org?

## Location of name servers

Hosts learn about the location of name servers via **DHCP**


The **operating system** keeps track of name servers and dynamically selects which one to use

```
Linux
cat /etc/resolv.conf
Windows
ipconfig /all
```

## Dynamic Host Configuration Protocol (DHCP)

MAC addresses are built into NICs. But network addresses are not.


Used to configure other settings such as: **DNS name servers**, addresses of default gateway, time servers, etc.



Q: How to send DHCP offer back to machine without an address?

## Recursive and iterative queries

Other name servers are in **iterative mode**



```
(base) jesse@jesses-XPS-2022:~$ dig canvas.vu.nl
;<<> DIG 9.16.1-ubuntu <<> canvas.vu.nl
;; global options: +cmd
;; Got answer:
;;->HEADER<- opcode: QUERY, status: NOERROR, id: 41657
;; flags: qr rd ad; QUERY: 1, ANSWER: 5, AUTHORITY: 0, ADDITIONAL: 0
;; WARNING: recursion requested but not available

;; QUESTION SECTION:
;canvas.vu.nl. IN A

;; ANSWER SECTION:
canvas.vu.nl. 0 IN CNAME vu-vanity.instructives.com.
vu-vanity.instructives.com. 0 IN CNAME canvas-dub-prod-c84-1303699784.eu-west-1.elb.amazonaws.com.
canvas-dub-prod-c84-1303699784.eu-west-1.elb.amazonaws.com. 0 IN A 52.17.144.218
canvas-dub-prod-c84-1303699784.eu-west-1.elb.amazonaws.com. 0 IN A 54.210.29.136
canvas-dub-prod-c84-1303699784.eu-west-1.elb.amazonaws.com. 0 IN A 54.77.55.222

;; Query time: 9 msec
;; SERVER: 172.31.224.1#53(172.31.224.1)
;; WHEN: Mon May 15 14:47:46 CEST 2023
;; MSG SIZE rcvd: 284
```

## DNS Resource Record (RR) Types


Name servers reply with **domain resource records**. A record can contain:

1. **IPv4 address** (record type A)
2. **IPv6 address** (record type AAAA)
3. Domain that accepts **email** (record type MX)
4. **Name server** for this domain (record type NS)
5. **Alias** to Canonical Name (record type CNAME)
6. ...

# Content Delivery Networks


## Content delivery networks

Q: How to make sure users do not all contact the same node?  
 A type of **caching** to increase system scalability.




## Content delivery networks

**Front end forwards requests and distributes load**  
 A type of **caching** to increase system scalability.



## Content delivery networks

Powered by DNS!   
 DNS can be used for load balancing!  
 A type of **caching** to increase system scalability.



## Content delivery networks

```
$ dig @192.5.6.30 ibm.com
;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags: udp: 4096
;; QUESTION SECTION:
;ibm.com.                IN      A
;; AUTHORITY SECTION:
ibm.com.                  172800 IN      NS      usw2.akam.net.
ibm.com.                  172800 IN      NS      usc2.akam.net.
ibm.com.                  172800 IN      NS      eur2.akam.net.
ibm.com.                  172800 IN      NS      ns1-99.akam.net.
ibm.com.                  172800 IN      NS      ns1-206.akam.net.
ibm.com.                  172800 IN      NS      asia3.akam.net.
ibm.com.                  172800 IN      NS      usc3.akam.net.
ibm.com.                  172800 IN      NS      eur5.akam.net.
```

## Content delivery networks


```
$ dig @192.5.6.30 ibm.com
;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:
;; QUESTION SECTION:
;ibm.com.
;; AUTHORITY SECTION:
ibm.com.                  172800 IN      NS      usw2.akam.net.
ibm.com.                  172800 IN      NS      usc2.akam.net.
ibm.com.                  172800 IN      NS      eur2.akam.net.
ibm.com.                  172800 IN      NS      ns1-99.akam.net.
ibm.com.                  172800 IN      NS      ns1-206.akam.net.
ibm.com.                  172800 IN      NS      asia3.akam.net.
ibm.com.                  172800 IN      NS      usc3.akam.net.
ibm.com.                  172800 IN      NS      eur5.akam.net.
```

'akam' means 'Akamai',  
 a CDN company.



## Email How does it work?

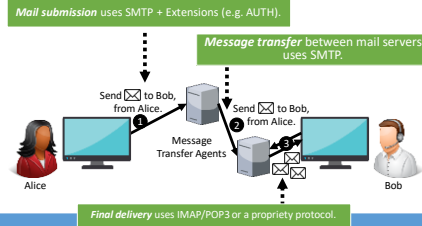
Powered by DNS!



43

## Email How does it work?

Q: Example of a proprietary protocol used for *final delivery*?



44

## Internet Message Access Protocol (IMAP)

RFC 9000

Ports: 143, 993

Q: Does gmail.com use POP3 or IMAP?

Sends commands to *mail server* to manipulate mailboxes

Common commands:

1. LOGIN. Log into server
2. FETCH. Fetch messages from a folder
3. CREATE/DELETE. Create or delete a folder
4. EXPUNGE. Remove messages marked for deletion

Uses mostly plain text!

Replaced POP3 protocol

Security through TLS (not covered in the course)

46

## Simple Mail Transfer Protocol (SMTP)

RFC 5321

Ports: 25, 587

SMTP uses ASCII

You can use TELNET to talk to a mail server!

```
S: 220 ee.uwa.edu.au SMTP service ready
C: HELO abcd.com
S: 250 cs.washington.edu says hello to ee.uwa.edu.au
C: MAIL FROM: <alice@cs.washington.edu>
S: 250 sender ok
C: RCPT TO: <bob@ee.uwa.edu.au>
S: 250 recipient ok
C: DATA
S: 354 Send mail; end with "." on a line by itself
C: -
```

Basic SMTP does not support binary data!

Basic SMTP does not include authentication!

Many extensions exist to address these issues.

47

## Multipurpose Internet Mail Extensions (MIME)

Developed for email, now used more broadly

Adds headers to email:

- MIME-Version
- Content-Description
- Content-Id
- Content-Transfer-Encoding
- Content-Type**

If MIME-Version in header  
check Content-Type  
Else  
plain text

48

## MIME Content-Type

1. Text: text/plain, text/html
2. Images: image/jpeg, image/gif
3. Video: video/mp4, video/mpeg
4. **Multipart**: multipart/mixed, multipart/alternative

Used to create messages with multiple data types (e.g., an email with attachment).

Basic SMTP does not support binary data!

49

## Multipurpose Internet Mail Extensions (MIME)

Developed for email, now used more broadly

Adds headers to email:

- MIME-Version
- Content-Description
- Content-Id
- Content-Transfer-Encoding**
- Content-Type


If MIME-Version in header  
check Content-Type  
Else  
plain text

Modern SMTP protocol supports binary data

## Sending binary data via ASCII-only SMTP

When MIME was introduced, servers were not expecting non-ASCII data.

Q: How to send binary via a server that can only handle ASCII?




Base64 encoding converts binary data into ASCII

## Base64 encoding


Q: How large is the overhead of base64 encoding?

Used to convert binary data to and from ASCII.  
Alphabet: [A-Za-z0-9+/  
6 bits are translated into 1 character.



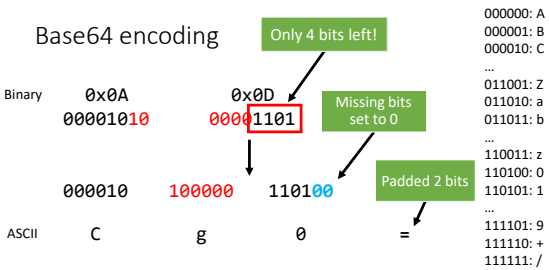
## Base64 encoding

binary to ASCII example



## Base64 encoding


Only 4 bits left!



## Base64 encoding

Missing bits set to 0

Padded 4 bits





Page last updated at 10:52 GMT, Wednesday, 14 October 2009 11:52 UK

E-mail this to a friend

Printable version

## Berners-Lee 'sorry' for slashes

The forward slashes at the beginning of internet addresses have long annoyed net users and now the man behind them has apologised for using them.

Sir Tim Berners-Lee, the creator of the World Wide Web, has confessed that the // in a web address were actually "unnecessary".




Tim Berners-Lee started the web to help scientists communicate

<http://news.bbc.co.uk/2/hi/technology/8306611.stm>

62

## HTTP Request/Response


HTML documents hosted by servers.  
Clients sends request for document from server.



63

## Evolution of HTTP

Optimizations are gradually incorporated to improve performance/security



64

## HTTP Protocol

Similar to chat application from the lab!

Originally a simple text-based protocol  
Many options added over time


Try it yourself:

```
$ telnet en.wikipedia.org 80
GET wiki/HTML HTTP/1.0
```

<https://tools.ietf.org/html/rfc2616>

65

## HTTP Request via TELNET



66

## HTTP Request Methods

Methods: GET, POST, PUT, HEAD, ...

```
$ curl -v -L --http1.1 https://vu.nl -o /dev/null
...
> GET / HTTP/1.1
> Host: vu.nl
> User-Agent: curl/7.64.1
> Accept: */*
...
https://www.w3.org/TR/2010/WD-html5-20100624/
Specifies the protocol, the domain name, and a path.
```

67






HTTP Persistent connection


Persistent connections allow browsers to issue multiple requests over the same TCP connection

HTTP 1.1



HTTP Performance Problem Head of Line Blocking (HOL)


Each request needs to wait for the previous one to complete!



HTTP1.1 Pipelined requests

Reduces Head of Line Blocking!

HTTP 1.1



HTTP/2


- Binary instead of plaintext.
  - Easier for machines to parse
  - More difficult for humans to read
  - Q: Why would it be easier for machines?
- Multiplexed streams over a single TCP connection.
  - Supports out-of-order responses!
- Server push allows the server to send resources before the client asks for it explicitly.



<https://books.ietf.org/html/rfc7540>


Head-of-Line Blocking in HTTP/2

Despite *pipelining* (HTTP1.1) and *out-of-order responses* (HTTP/2), HTTP/2 performance still suffers from a type of Head of Line blocking



Head-of-Line Blocking in HTTP/2

Despite *pipelining* (HTTP1.1) and *out-of-order responses* (HTTP/2), HTTP/2 performance still suffers from a type of Head of Line blocking




QUIC orders data per stream [RFC 9051](#)

### HTTP/3 (HTTP + QUIC)

Each HTTP request can use a separate stream; within a stream, data is delivered in order; across streams no such guarantee is made

HTTP/3 uses the **QUIC** protocol

QUIC performs multiplexing, uses UDP



[RFC 6455](#)

### WebSockets

Application layer protocol

Q: Can the application layer contain protocols?

A socket-like interface on the application layer.

Full-duplex connection between server and client.

Q: Can you think of a use-case?

Increasingly complex 'apps' on the Web that need to send data continuously.

Examples:

- irc-ws.chat.twitch.tv
- ws.todoist.com

irc-ws.chat.twitch.tv	other	1.10 MB
-----------------------	-------	---------

<https://tools.ietf.org/html/rfc6455>

[RFC 6455](#)

### WebSockets

Application layer protocol

A socket-like interface on the application layer.  
Full-duplex connection between server and client.

Q: Can you think of a use-case?

Increasingly complex 'apps' on the Web that need to send data continuously.

Examples:


- irc-**ws**.chat.twitch.tv
- ws**.todoist.com

irc-ws.chat.twitch.tv	other	1.10 MB
-----------------------	-------	---------

'ws' stands for WebSocket

### Stacking


#### Application layer protocols



### Stacking


#### Application layer protocols

Application layer can continue stacking protocols



### Stacking

#### Application layer protocols



### Starting a WebSocket over HTTP

```
GET /chat HTTP/1.1
Host: example.com:80
Upgrade: websocket
Connection: Upgrade
Sec-WebSocket-Key: dGh1IHhvbXBsZS5jb2ZjZQ==
Sec-WebSocket-Version: 13

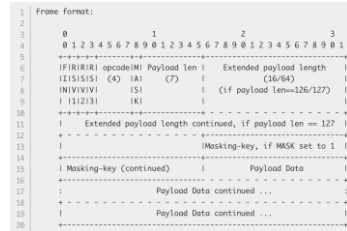
HTTP/1.1 101 Switching Protocols
Upgrade: websocket
Connection: Upgrade
Sec-WebSocket-Accept:
s3pPLMB1Tx09kYGzZhZRbK+x0o=

Reply from server if it accepts
```


Client requests to switch to WebSocket protocol

Reply from server if it accepts

### WebSocket frame format



### HTTP is the new "narrow waist"



E.g., REST APIs

Method	Description
GET	Read a Web page
HEAD	Read a Web page's header
POST	Append to a Web page
PUT	Store a Web page

Q: Advantages over using TCP directly?

Answers include:

- Provides set of methods
- Provides security
- Provides naming

Source: Peterson, Larry L., and Bruce S. Davie. Computer networks: a systems approach. Elsevier, 2007.

### Application Layer Topics

1. Domain Name System (DNS)
2. Email
3. Web (HTTP, QUIC, WebSocket)
4. Multimedia applications

### Video dominates



Video constitutes around 70 percent of all global mobile network traffic in 2022

- 28,000 people watching Netflix
- 500 hours of content uploaded to YouTube
- 2 million Twitch views
- 3.4 million Snaps created



Source: <https://www.ersoson.com/en/reports-and-essays/mobility-report1/data-for-what-happens-in-a-minute>

### Streaming Video Requires Compression

1024 height x 2048 width = 2M pixels  
 1 pixel = 1 byte  
 30 frames per second → 60 MB/s = 480 Mbps

Without compression, only possible over wired fiber-optic channels

Compression reduced bandwidth requirement by an order of magnitude

### Internet connection speed recommendations

To watch TV shows and movies on Netflix, we recommended having a stable internet connection with a download speed shown below in megabits per second (Mbps).

Video quality	Resolution	Recommended speed
High definition (HD)	720p	3 Mbps or higher
Full high definition (FHD)	1080p	5 Mbps or higher
Ultra high definition (UHD)	4K	15 Mbps or higher

Source: <https://help.netflix.com/en/node/306>

98

Large compression rates  $> \times 10$ .

### Digital audio compression

Audio typically compressed before sending.


**Lossy compression** achieves higher compression rates than **lossless compression**, but **loses data**.

Q: Why is lossy compression acceptable?

Lossy encoders based on how humans perceive sound.


99

### Human hearing frequency range



100

### Human hearing masked signals



101

### Digital video JPEG compression

Changes RGB to  $YC_bC_r$ .

Y is luminance.

$C_b, C_r$  are chrominances.

Q: Why change to this format?

Eyes are **less** sensitive to chrominance than to luminance.

JPEG reduces size of  $C_b$  and  $C_r$ .

Total compression rate  $\times 2$ .

102


Large compression rates  $> \times 50$ .

### Digital video

Q: What is the use of **bidirectional frames**?

MPEG compresses over a sequence of frames, further using motion tracking to remove temporal redundancy


1. I (Intra-coded) frames are self-contained
2. P (Predictive) Looks for comparable **macro blocks** in previous frames. **How long to search is up to the implementation.**
3. B (Bidirectional) frames may base prediction on previous frames and **future** frames.



103

## Run-Length Encoding Part of JPEG Compression


A lossless compression technique.



104

## Huffman Encoding Prefix code: no code word is prefix of other code word

Q: Why is this useful?



Less than half the original size!  $\frac{54}{128} < 0.42$


105

## Huffman Encoding Part of JPEG Compression

11 101 101 100 0111 0110 11 0101 0111 0100 0011 100 11 0010 0001 0000  
(54 bits) "application layer"



Form a tree by select the two smallest nodes, and combining them into a new node, until only the root is left.



106


## Networking Challenges for Multimedia Applications

107

## Challenge 1 YouTube NETFLIX Streaming stored media

How to handle **transmission errors**?


1. Use reliable transport (e.g., TCP).
  - Increases jitter significantly.
2. Use **forward error correction** (error correction in the application layer).
  - Increases jitter, decoding complexity, and overhead.



3. Interleave media
  - Slightly increases jitter and decoding complexity.

108

## Masking errors by interleaving media



Smaller gaps are less noticeable.

Playback on user's machine      Network      Server

109

## Challenge 1 Streaming stored media



Low-water mark prevents **stalls** in playback.

High-water mark gives client time to prevent **running out of buffer space**.

110

## Challenge 2 Streaming live media

Streaming live media is similar to the stored case plus:

1. Can't stream faster than **live rate** to get ahead
  - Usually need larger buffer to absorb jitter
2. Often have many users viewing at the same time
  - UDP with multicast greatly improves efficiency. It is rarely available, so **many TCP connections are used**.

111

## Challenge 3 Streaming interactive media

Real-time conferencing has two or more connected live media streams, e.g., voice over IP, Skype video call  
Requires low jitter **and** low latency.

1. Benefits from network support (Quality of Service).
2. Large bandwidth (no congestion).

Difficult to provide across long distances/multiple networks

112

## Take-Home Message

- Many responsibilities and pseudo layers hidden in Application Layer
  - From OSI: Presentation, Session. Others: WebSocket, RTP, etc.
- Important behind-the-scenes applications exist (e.g., DNS)
- Traditional "killer apps" for the Internet:
  - Email
  - The Web
- HTTP is the new "narrow waist"
  - Improved over time (HTTP/2 [SPDY], HTTP/3 [QUIC])
- Today's Internet is increasingly used for multimedia applications
  - Provide new challenges (high bandwidth, low latency, low jitter)

113