

SERVER-SIDE

WEB SERVICES

HTTP PROTOCOL

MORE HTTP: OTHER FEATURES

HTTP SERVICES

VERSION

v1.3	07 October 2021
v1.2	05 October 2021
v1.1	30 September 2021
v1.0	28 September 2021

ACKNOWLEDGMENTS

THANKS TO:

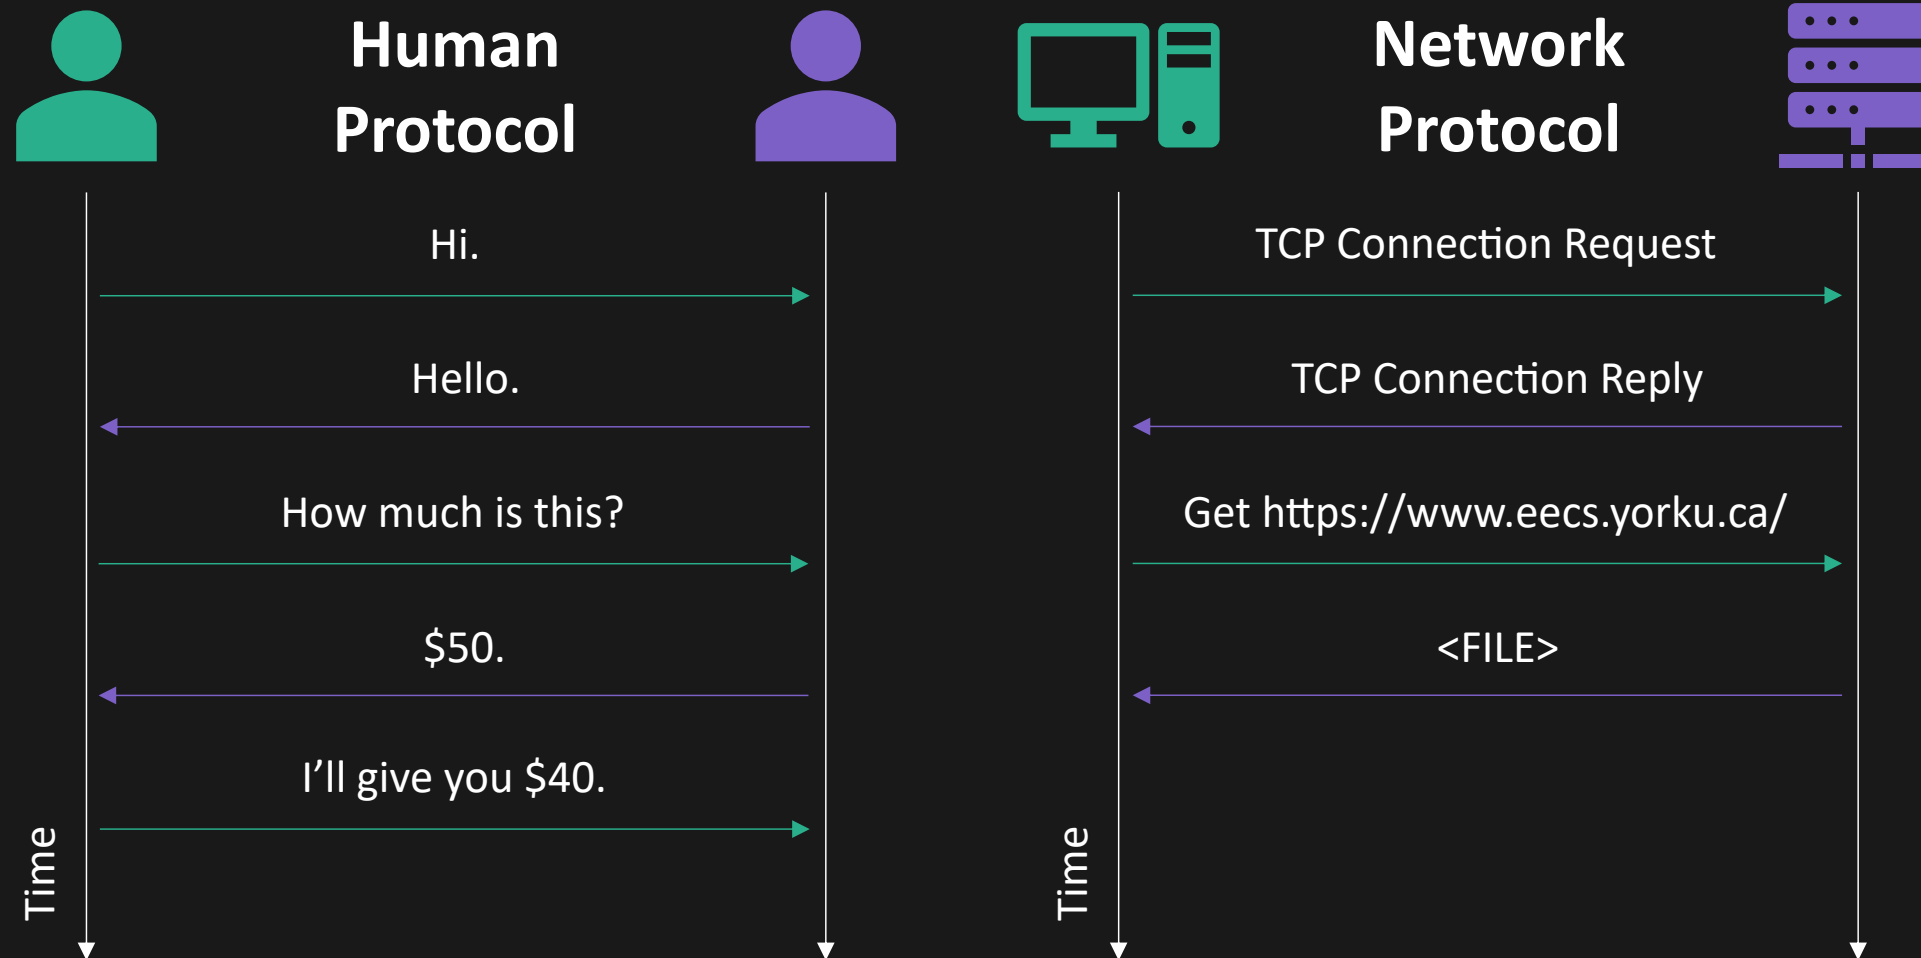
- Hamzeh Roumani, who has shaped EECS-4413 into a leading hands-on CS course at EECS and who generously shared **all** of his course materials and, more importantly, his teaching philosophy with me;
- Parke Godfrey, my long-suffering Master's supervisor and mentor; and
- Suprakash Datta for giving me this opportunity to teach this course.

PRINTABLE VERSION OF THE TALK

[Download PDF](#)

HTTP

THE PROTOCOL

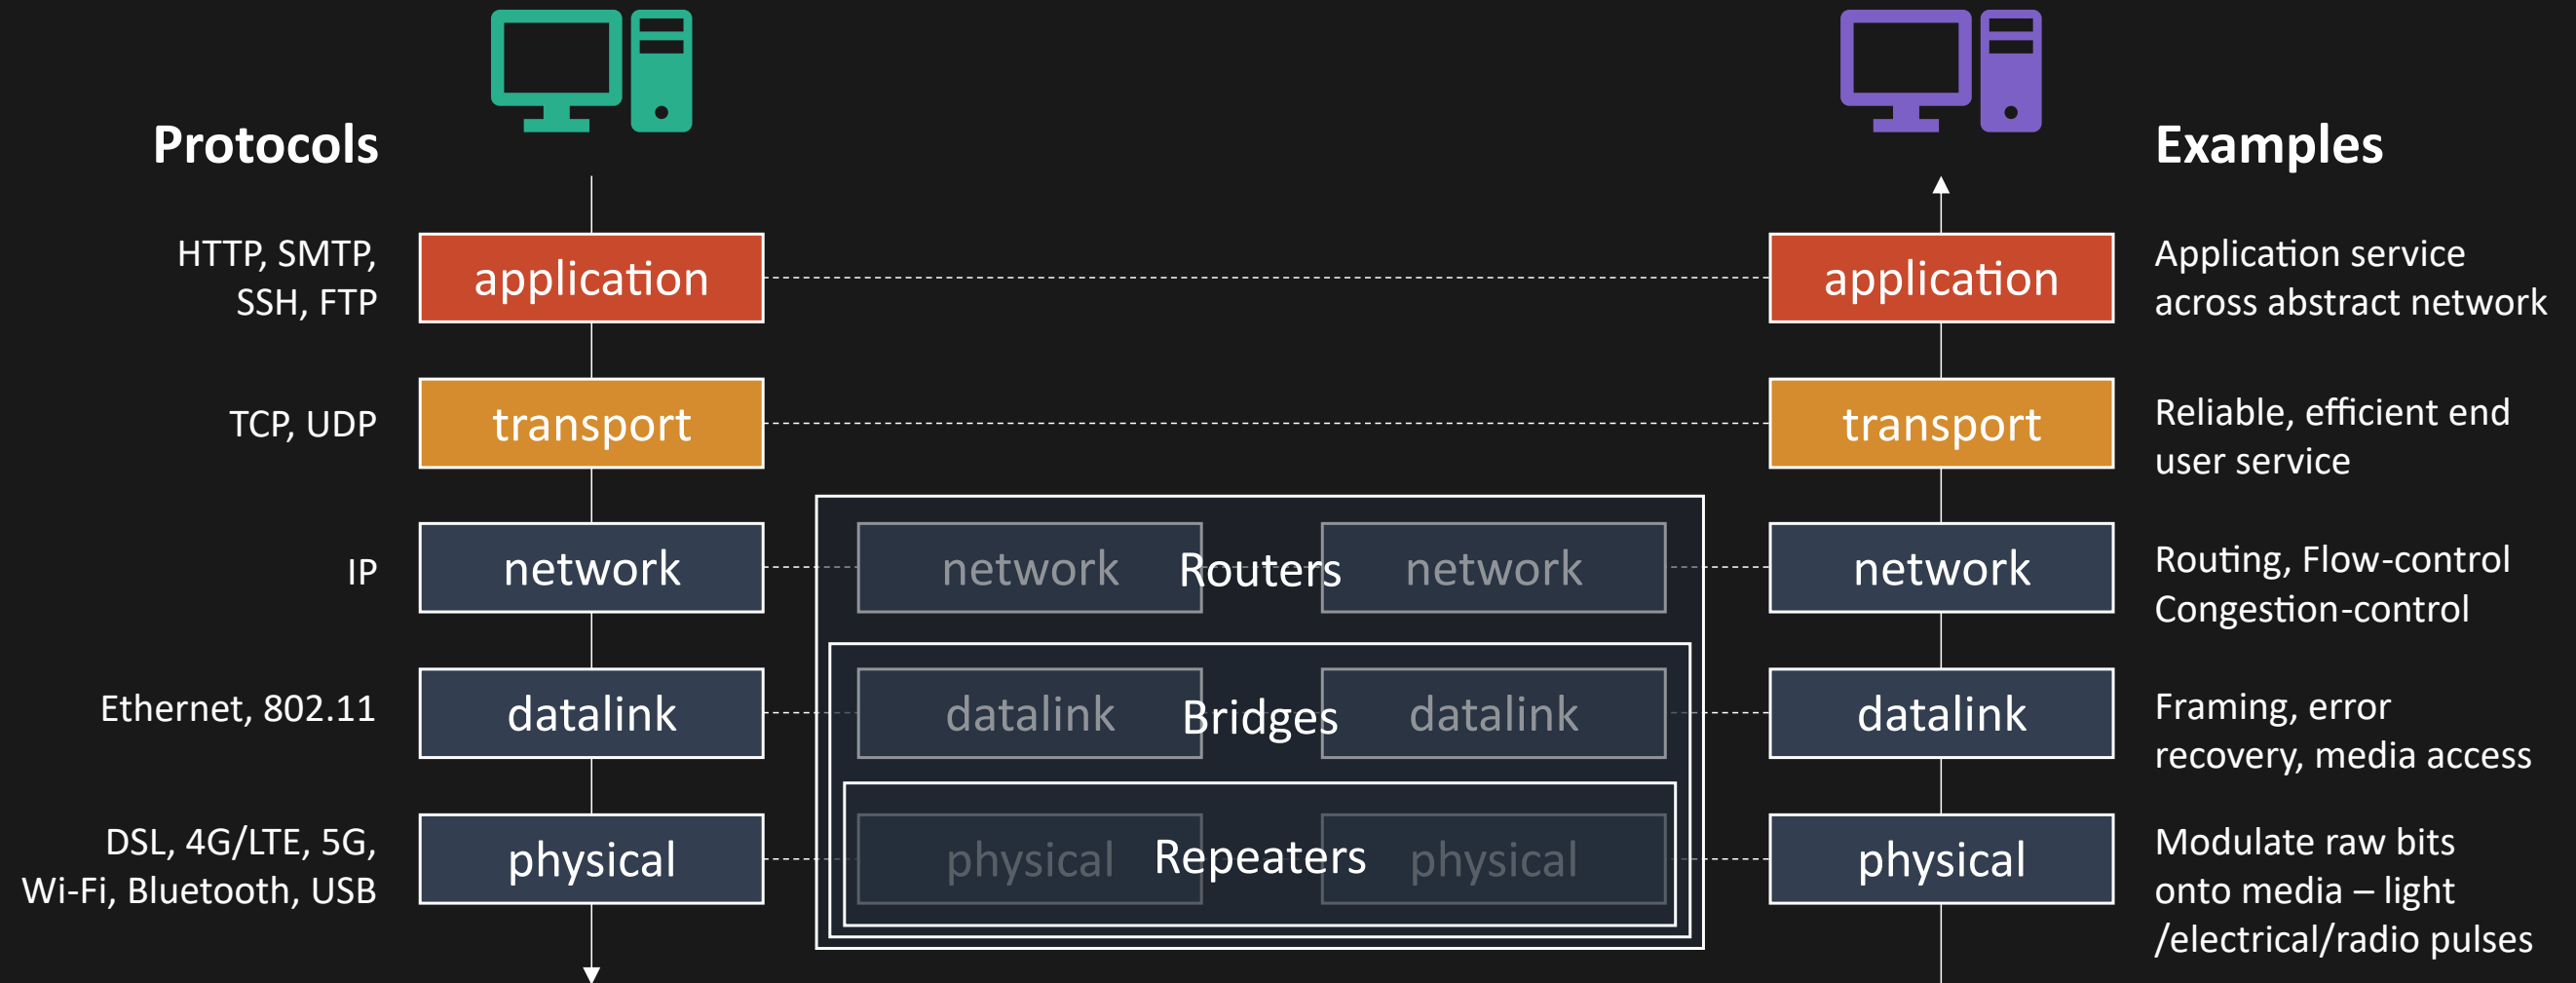


WHAT IS A PROTOCOL?

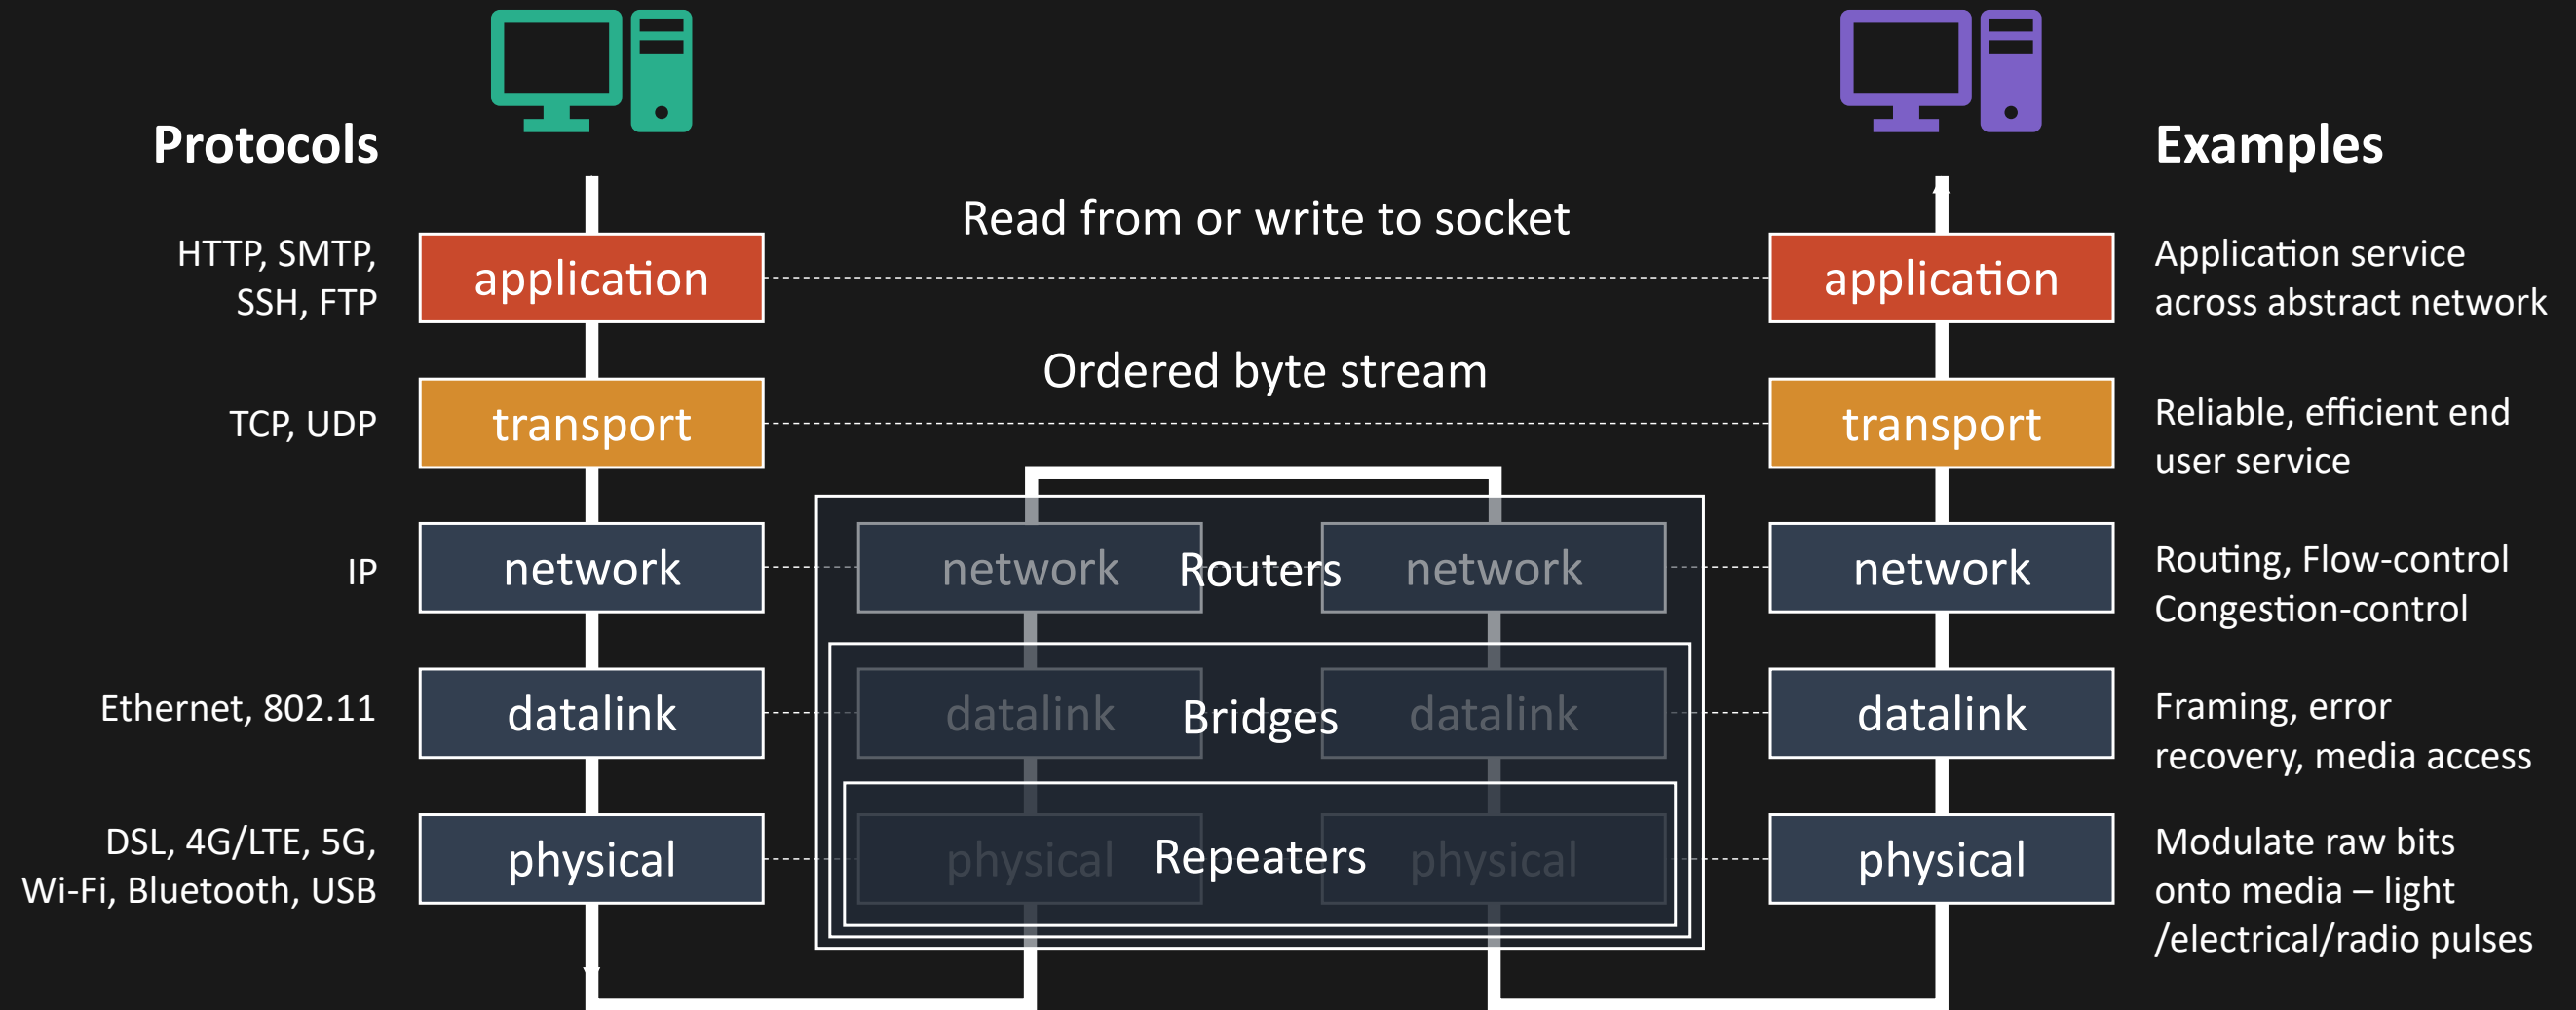
A protocol defines format & order of messages sent and received among network entities, and actions taken on message transmission, receipt.

Source: Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).

THE PROTOCOL STACK



THE PROTOCOL STACK



HYPertext TRAnSFER PROTOCOL (HTTP)

The **Hypertext Transfer Protocol (HTTP)** is an application layer protocol for distributed, collaborative, hypermedia information systems, used primarily with the WWW (World Wide Web) in the client-server model where a web browser is a client communicating with the webserver which is hosting the website. Since 1990, this has become the foundation for data communication. HTTP is a standard and stateless protocol that is used for different purposes as well using extensions for request methods, error codes, as well as headers.

Sources: Hypertext Transfer Protocol – HTTP/1.1 (W3C); Introduction to HTTP (w3schools.in); and Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).

- Created by Tim Berners-Lee at CERN (1991)
- Standardized and much expanded by the IETF.
- HTTP/1.0 (RFC 1945), HTTP/1.1 (RFC 2068), HTTP/2 (RFC 7540)
- Other related protocols: HTTPS and WebSocket.
- Rides on top of the TCP protocol, standard on port 80
- TCP provides: reliable, bi-directional, in-order byte stream
- Goal: transfer objects between systems
- Do not confuse with other WWW concepts:
 - HTTP is not page layout language (that is HTML)
 - HTTP is not object naming scheme (that is URLs)

HTTP IN OPERATION



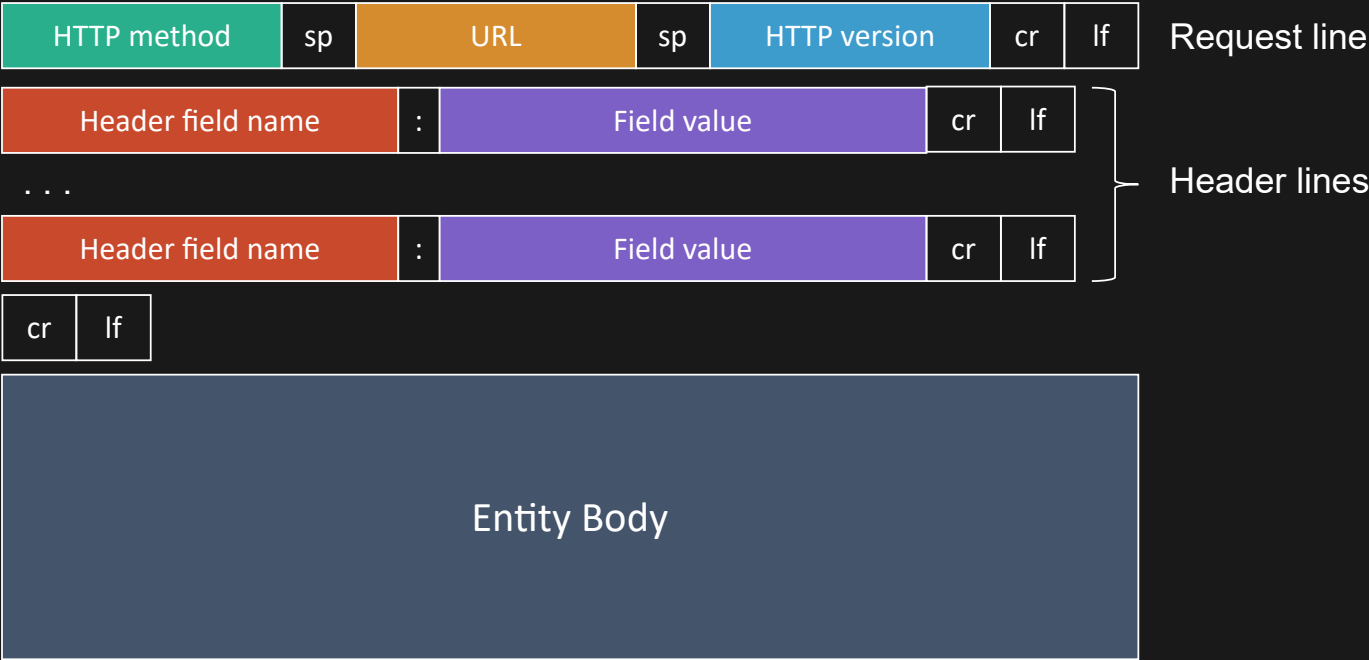
http://www.eecs.yorku.ca/course_archive/2021-22/F/4413/index.html

Time

-
- 1a. HTTP client initiates TCP connection to HTTP server (process) at www.eecs.yorku.ca. Port 80 is default for HTTP server.
 - 1b. HTTP server at host www.eecs.yorku.ca waiting for TCP connection at port 80. “accepts” connection, notifying client.
 2. HTTP client sends HTTP request message (containing URL) into TCP connection socket.
 3. HTTP server receives request message, forms response message containing requested object (/course_archive/2021-22/F/4413/index.html), sends message into socket.
 4. HTTP server closes TCP connection.
 5. HTTP client receives response message containing HTML file, displays HTML. Parsing HTML file, finds references CSS, JS, and Image files.
 6. Steps 1-5 repeated for each separately referenced object.

Sources: Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).

HTTP REQUEST



```
GET /course_archive/2021-22/F/4413/ HTTP/1.1
Host: www.eecs.yorku.ca
Connection: keep-alive
Pragma: no-cache
Cache-Control: no-cache
Upgrade-Insecure-Requests: 1
User-Agent: Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36.
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,image/webp.
Accept-Encoding: gzip, deflate, br
Accept-Language: en-CA,en-US;q=0.9,en;q=0.8
Cookie: mayaauth="..."
```

Sources: Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).

HTTP REQUEST METHODS

GET Used to request a particular resource data from the Web server by specifying the parameters as a query string (name and value pairs) in the URL part of the request.

- Are bookmark-able as they appear in the URL.
- Cache-able.
- Saved in the browser history if it is executed using a web browser.
- Character length restrictions (2048 chars maximum).
- Cannot be used to send binary data.
- Data can only be retrieved and have no other effect.
- When communicating sensitive data such as login credentials, should not be used.
- A safe and ideal method to request data only.

HEAD Almost identical to the GET method, but the only difference is that it will not return any response body. The HEAD request becomes useful for testing whether the GET request will actually respond before making the actual GET request.

DELETE Used to delete any specific resource.

TRACE Used for performing a message loop-back, which tests the path for the target resource. It is useful for debugging purposes.

CONNECT Used for establishing a tunnel to the server recognized by a given URI.

POST Used to send data to the Web server in the request body of HTTP.

- Not be bookmark-able as they do not appear in the URL.
- Not cached.
- Not saved as history by the web browsers.
- No restriction on the amount of data to be sent.
- Can be used to send ASCII as well as binary data.
- Use when communicating sensitive data, such as when submitting an HTML form.
- Security depends on the HTTP protocol.
- By using secure HTTP (HTTPS), information is protected.

PUT Requests that the target resource creates or updates its state with the state defined by the representation enclosed in the request. Used to update existing resources with uploaded content or to create a new resource if the target resource is not found. The difference between POST and PUT is that PUT requests are static, which means calling the same PUT method multiple times will not yield a different result.

PATCH Requests that the target resource modifies its state according to the partial update defined in the representation enclosed in the request.

OPTIONS Used for describing the communication preferences for any target resource.

Sources: Introduction to HTTP (w3schools.in).

UNIVERSAL RESOURCE IDENTIFIERS (URI)

URIs are also known as WWW addresses, Uniform Resource Name (URN), and the Uniform Resource Locator (URL). These are formatted, case-insensitive strings that identify a web service, a resource, a website, etc.

```
URI = "http : " "/" host [ ":" port ] [ abs_path ["?" query]]
```

- A standard way to send many name/value pairs in a single string (QUERY_STRING or Form data)
- Specified in RFC 2396 ‘Uniform Resource Identifiers (URI): Generic Syntax’

RULES OF URL-ENCODING

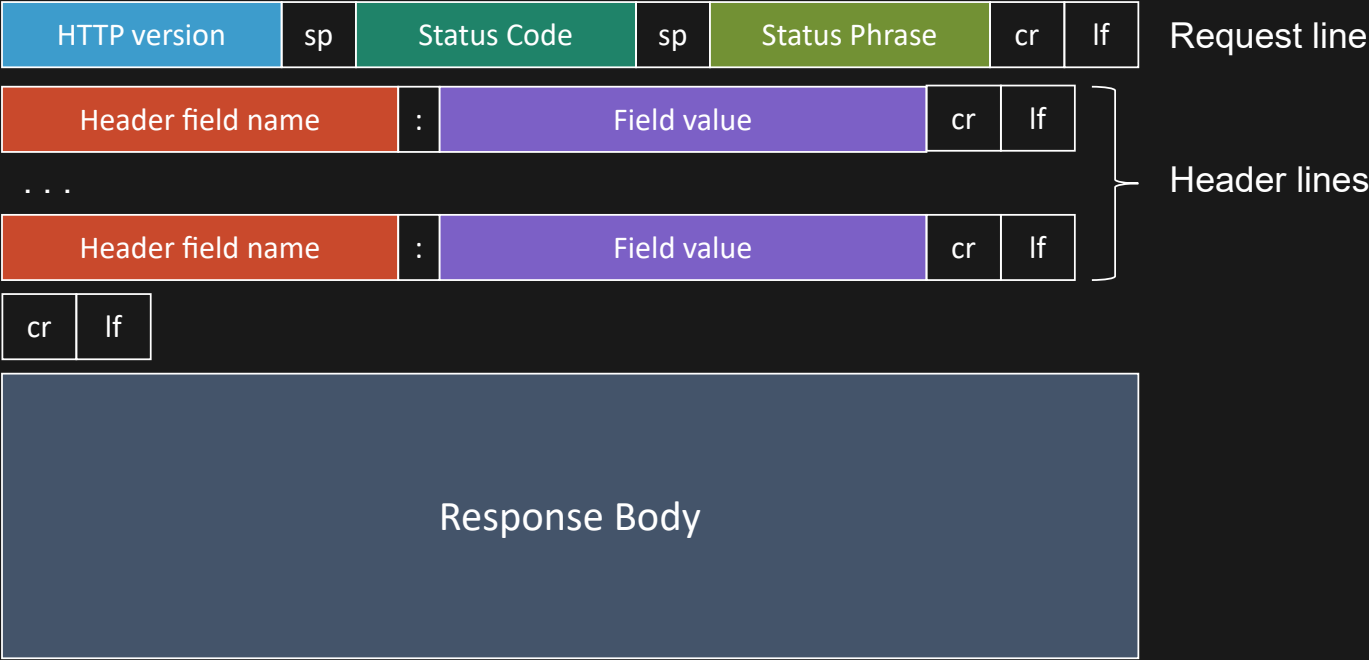
1. All submitted URLs, query strings or form data should be concatenated into single strings of ampersand (&) separated name=value pairs, one pair for each form tag or query parameter. Like this:

```
form_tag_name_1=value_1&form_tag_name_2=value_2&...
```

2. Spaces in a name or value are replaced by a plus (+) sign or “%20” (a percent sign followed by 20). This is because url’s cannot have spaces in them and under METHOD=GET, the form data is supplied in the query string in the url.
3. Other characters (ie, =, &, +) are replaced by a percent sign (%) followed by the two-digit hexadecimal equivalent of the punctuation character in the Ascii character set.
4. Otherwise, it would be hard to distinguish these characters inside a query or form variable from those between the variables in the first rule above.

Sources: Introduction to HTTP (w3schools.in) and CGI - CSC309F Programming on the Web, David A. Penny (2001)..

HTTP RESPONSE



```
HTTP/1.1 200 OK

Date: Tue, 28 Sep 2021 15:48:11 GMT

Server: Apache/2.4.48 (Unix) PHP/7.4.20 OpenSSL/1.0.2k mod_wsgi/4.7.1 Python/3.8

Strict-Transport-Security: max-age=31536000; includeSubDomains

Last-Modified: Tue, 28 Sep 2021 00:21:45 GMT

Accept-Ranges: bytes

Content-Length: 13305

Keep-Alive: timeout=5, max=100

Connection: Keep-Alive

Content-Type: text/html

<!DOCTYPE html>

<html lang="en-US">

  <head>

    ...
```

Sources: Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).

HTTP RESPONSE STATUS CODES

1XX INFORMATIONAL

100 HTTP CONTINUE
101 SWITCHING PROTOCOLS

2XX SUCCESS

200 OK
201 CREATED
202 ACCEPTED
203 NON AUTHORITY INFORMATION
204 NO CONTENT
205 RESET CONTENT
206 PARTIAL CONTENT

4XX CLIENT ERROR

400 BAD REQUEST
401 UNAUTHORIZED
402 PAYMENT REQUIRED
403 FORBIDDEN
404 NOT FOUND
405 METHOD NOT ALLOWED
406 NOT ACCEPTABLE
407 PROXY AUTHENTICATION REQUIRED
408 REQUEST TIME OUT
409 CONFLICT
410 GONE
411 LENGTH REQUIRED
412 PRECONDITION FAILED
413 REQUEST ENTITY TOO LARGE
414 REQUEST URI TOO LARGE
415 UNSUPPORTED MEDIA TYPE

3XX REDIRECTION

300 MULTIPLE CHOICES
301 MOVED PERMANENTLY
302 MOVED TEMPORARILY
303 SEE OTHER
304 NOT MODIFIED
305 USE PROXY

5XX SERVER ERROR

500 INTERNAL SERVER ERROR
501 NOT IMPLEMENTED
502 BAD GATEWAY
503 SERVICE UNAVAILABLE
504 GATEWAY TIME OUT
505 HTTP VERSION NOT SUPPORTED

TRY OUT HTTP (CLIENT SIDE) FOR YOURSELF

1. Telnet to your favorite Web server, e.g.:

```
$ telnet www.eecs.yorku.ca 80
```

Open TCP connection to port 80 (default http server port) at `www.eecs.yorku.ca`. Anything typed is sent to port 80 at `www.eecs.yorku.ca`.

2. Type in a GET HTTP request:

```
$ GET /course_archive/2021-22/F/4413/index.html HTTP/1.1
```

Type this at the prompt (followed by two carriage returns) to send this minimal GET request to the HTTP server.

3. Look at response message sent by the HTTP server!

MORE HTTP

OTHER FEATURES

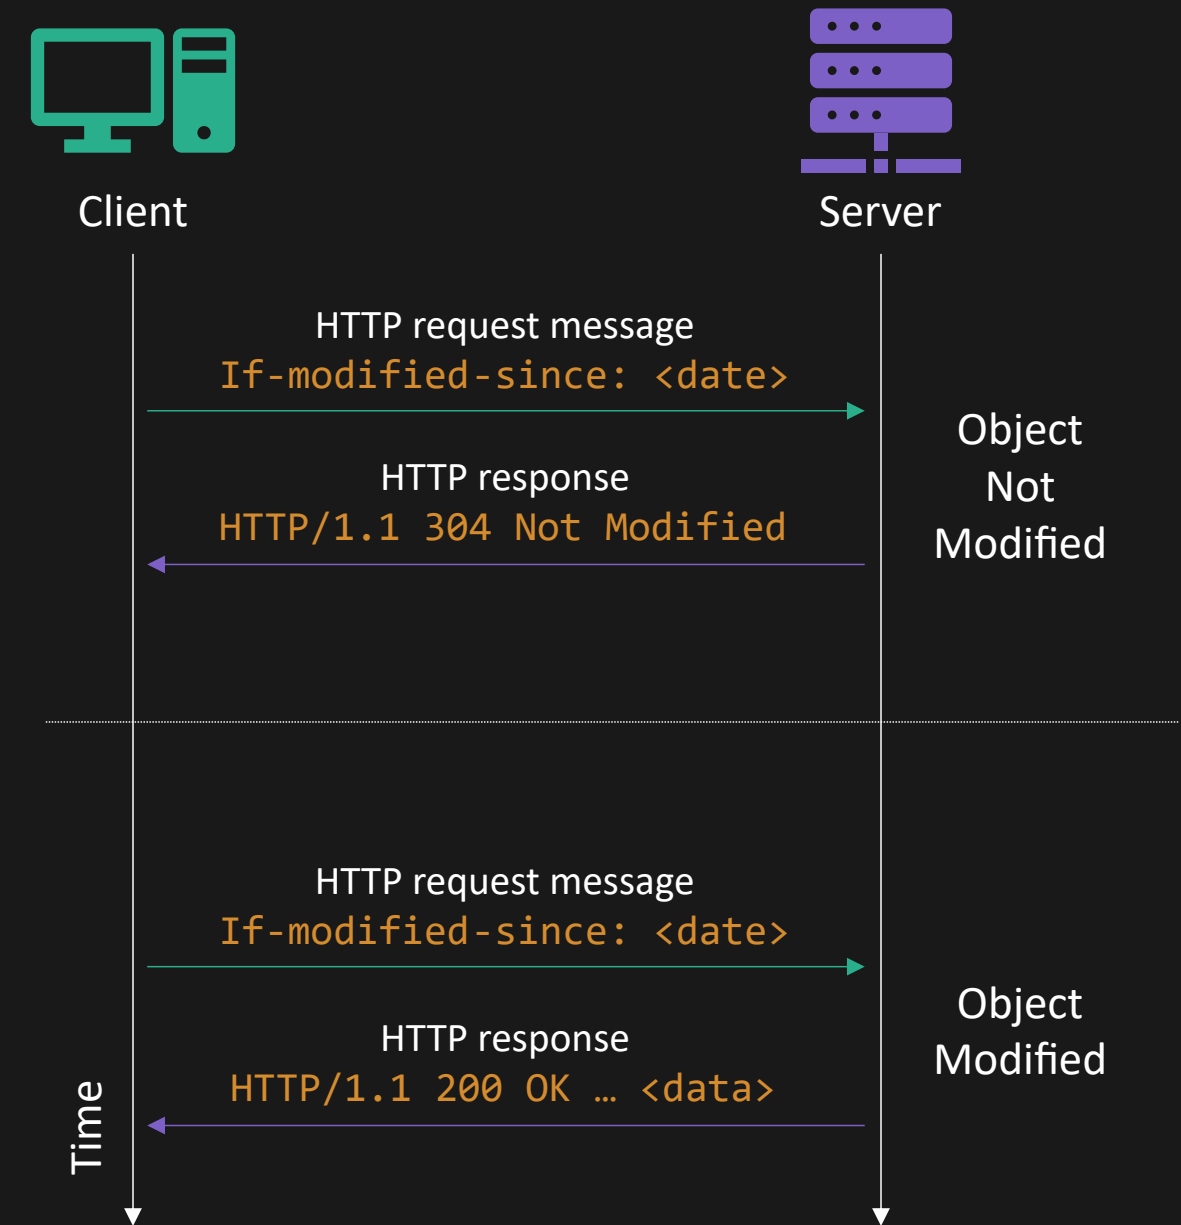
- Conditional GET
- Redirection
- Basic Authentication
- Open Authentication (OAuth)
- Persistence & Cookies
- Keep-Alive (with HTTP/1.1)
- Web Caching

CONDITIONAL GET

IF-MODIFIED-SINCE REQUEST HEADER

The client tells the server it has data and asks the server whether it has a fresher version or the client is up to date.

- Goal: Don't send object if the client has an up-to-date copy (cached).
- Client: Specify the date of cached copy in the HTTP request:
`If-modified-since: <date>`.
- Server: Response contains no object if the cached copy is up to date:
`HTTP/1.1 304 Not Modified`.



SESSION MANAGEMENT

COOKIES

PROBLEM: HTTP IS STATELESS

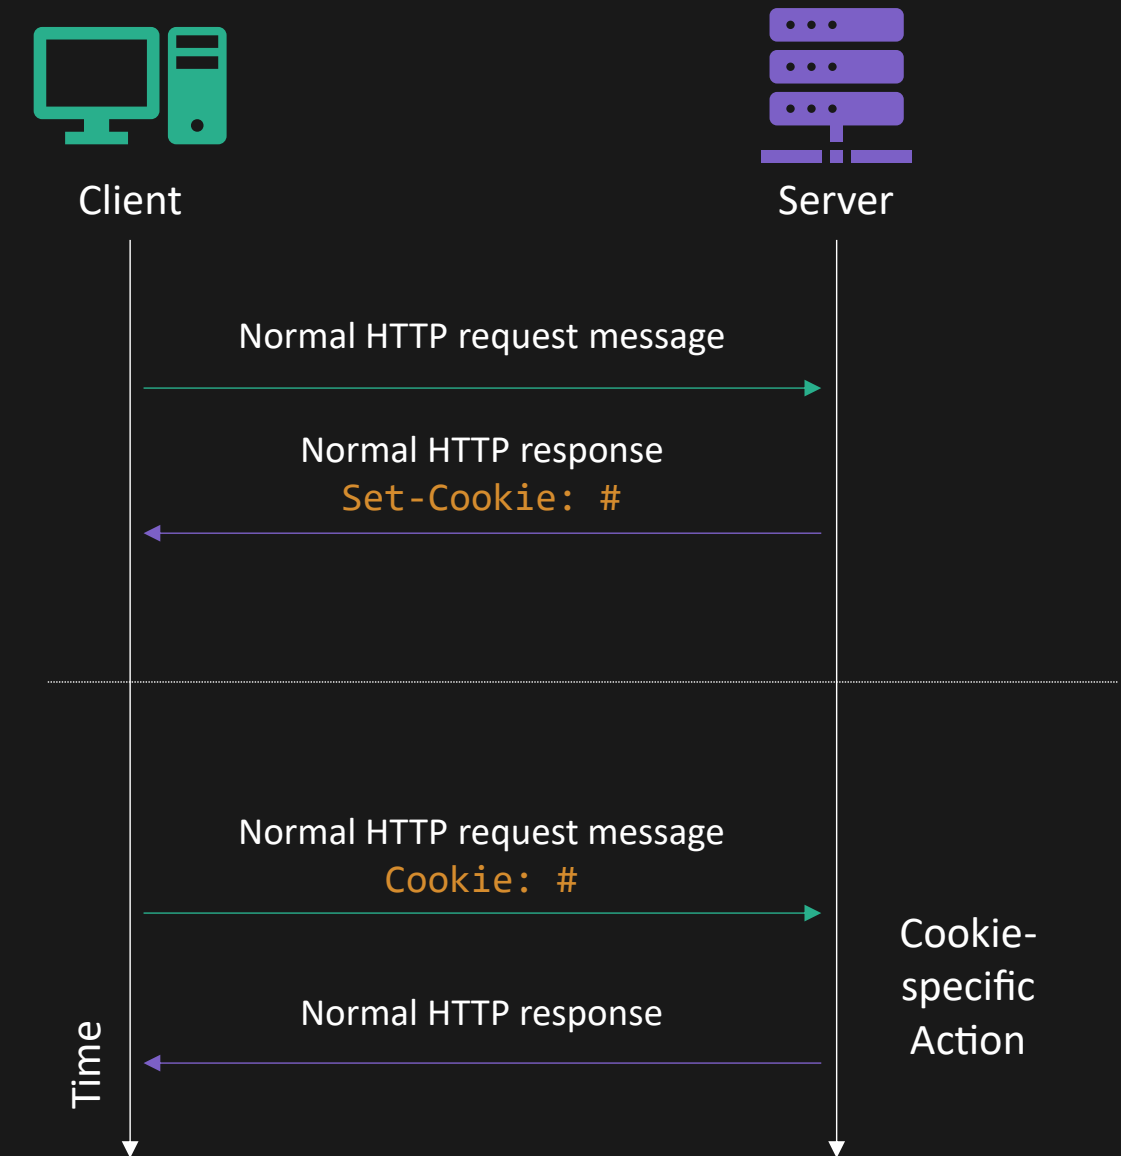
- Server does not maintain status information across client requests
- No way to correlate multiple request from same user

SOLUTION: STORE A COOKIE ON THE CLIENT-SIDE

- Small amount of information (typically server-generated user id)
- Sent by client with each request
- Updated by server with response

CLIENT-SERVER INTERACTION: COOKIES

- Server sends Cookie to client in response message:
`Set-Cookie: 1678453`.
- Client presents Cookie in later requests: `Cookie: 1678453`.
- Server matches presented-cookie with server-stored info:
 - Authentication
 - Remembering user preferences, previous choices



Sources: Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).

SESSION MANAGEMENT

SESSIONS

HTTP cannot maintain state (RESTful) but we can:

CLIENT-SIDE

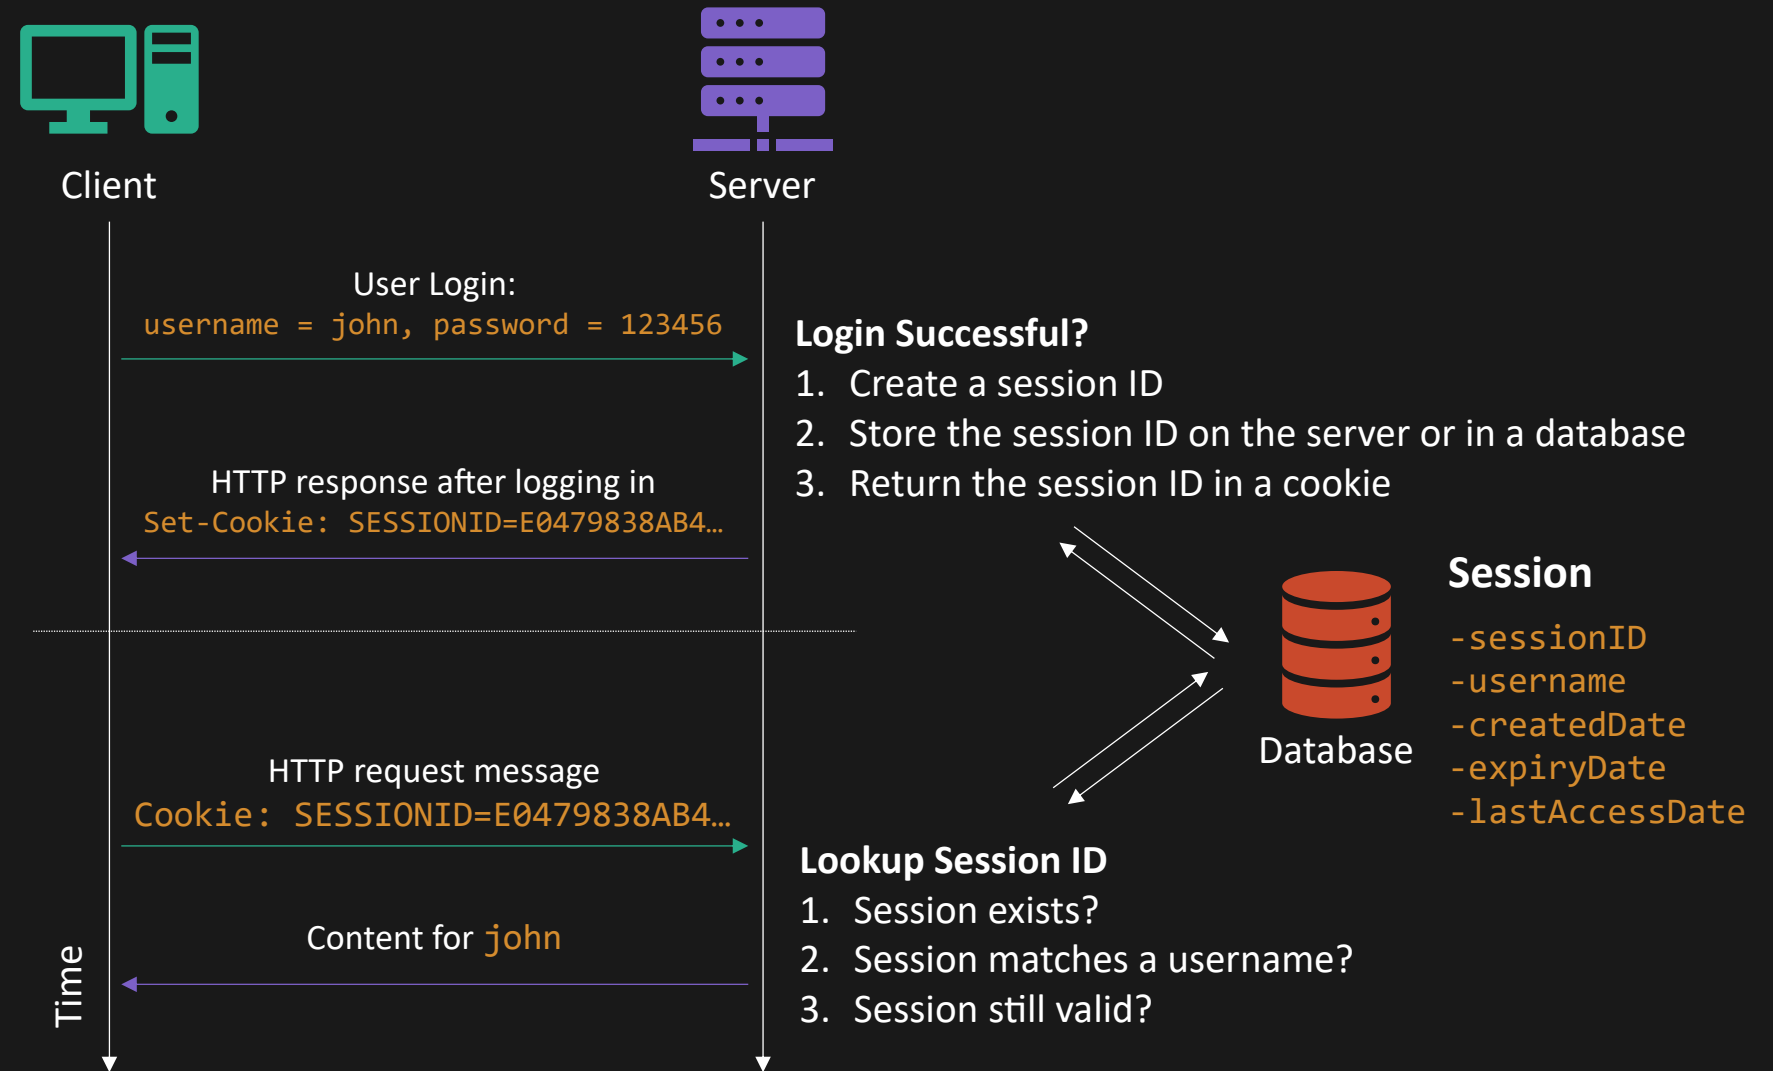
- State maintained by client and sent as needed to server.

NETWORK SIDE

- State shuffled back and forth with every request/response
- Typically through hidden fields, URL Rewriting, or Cookies

SERVER SIDE

- Server keeps it in memory or a database with a key derived from the client's credentials (known through authentication or assigned).
- The key (cookie) is stored in an HTTP header (network side).



SESSION MANAGEMENT

COOKIES VS SESSIONS

COOKIES

A cookie is a small file with the maximum size of 4KB that the web server stores on the client computer. Once a cookie has been set, all page requests that follow return the cookie name and value. A cookie can only be read from the domain that it has been issued from.

SESSIONS

A session is a global variable stored on the server. Each session is assigned a unique id which is used to retrieve stored values. Whenever a session is created, a cookie containing the unique session id is stored on the user's computer and returned with every request to the server. If the client browser does not support cookies, the unique session id is displayed in the URL. Sessions have the capacity to store relatively large data compared to cookies.

KEY DIFFERENCES

COOKIES

Cookies are client-side files that contain user information

Cookies are not secure, as data is stored in a text file on the client. If an unauthorized user gets access to a client's browser, they can tamper with the data.

Cookie ends depending on the lifetime you set for it

A cookie is not dependent on Session.

You don't need to start cookie as it is stored on the client's machine

There is no way to delete or unset a cookie as it resides on the client.

The official maximum cookie size is 4KB

SESSIONS

Sessions are server-side files which contain user information. More secure and tamper-proof.

Sessions are more secured compared to cookies, as they save data in encrypted form and are stored on the webserver.

A session ends when a user closes their browser.

A session is dependent on Cookie.

You need to start the session.

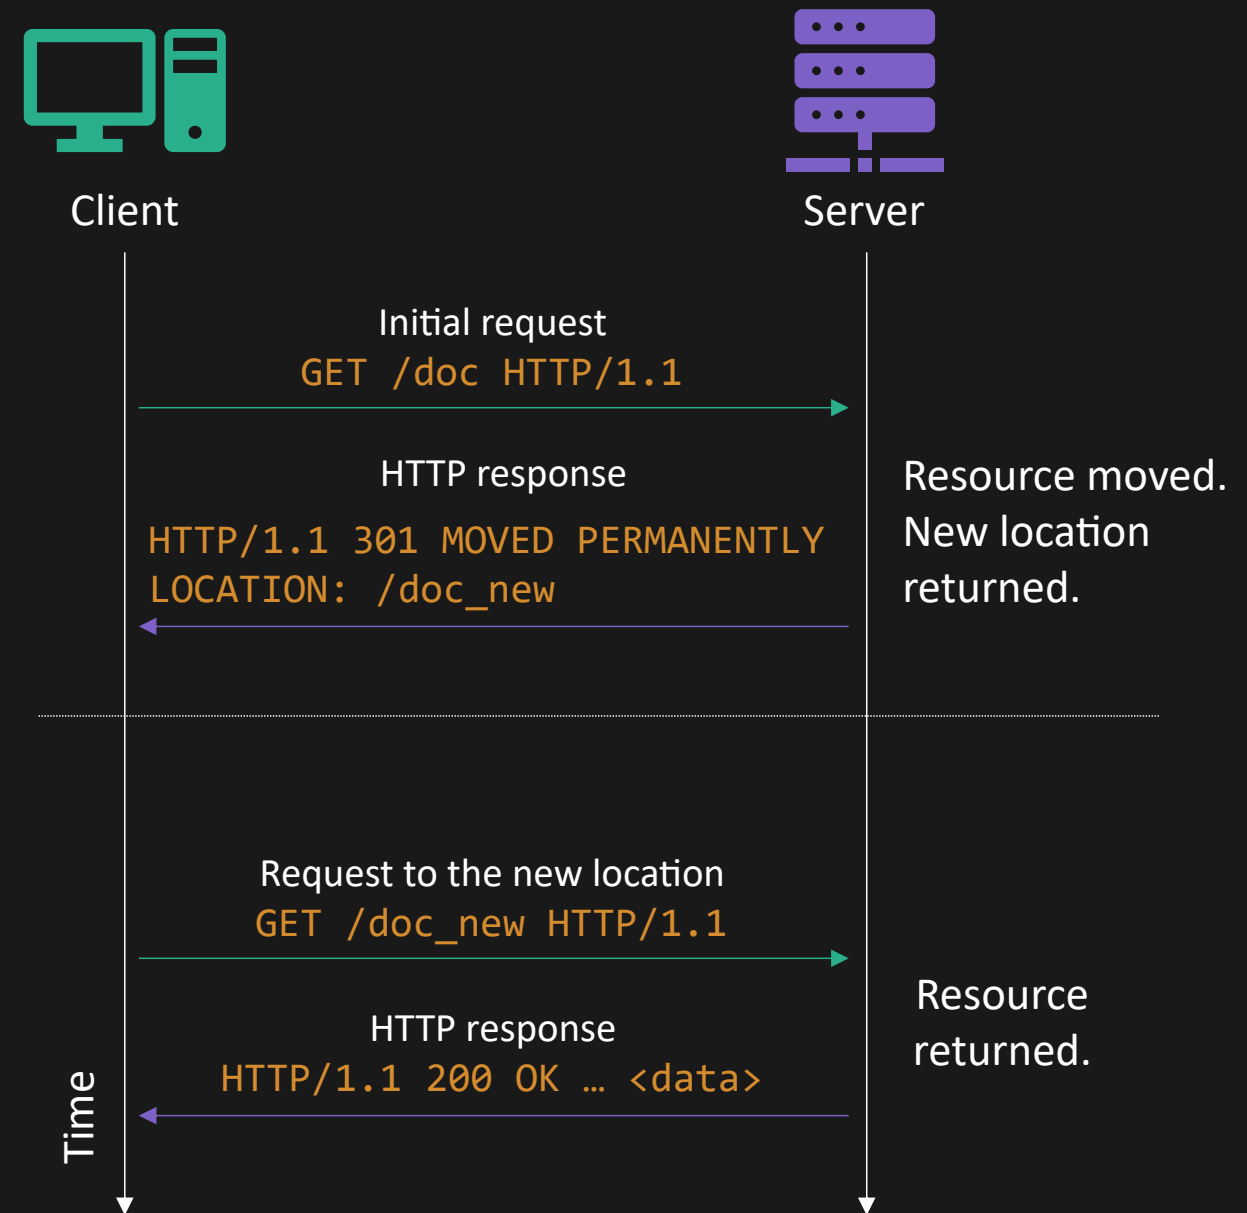
You can destroy (invalidate) a session from the server.

Within a session you can store as much data as you like. The only limits you can reach is the maximum memory consumable at one time.

REDIRECTION

In HTTP, redirection is triggered by a server sending a special redirect response to a request. Redirect responses have status codes that start with `3`, and a `Location` header holding the URL to redirect to.

When browsers receive a redirect, they immediately load the new URL provided in the `Location` header. Besides the small performance hit of an additional round-trip, users rarely notice the redirection.

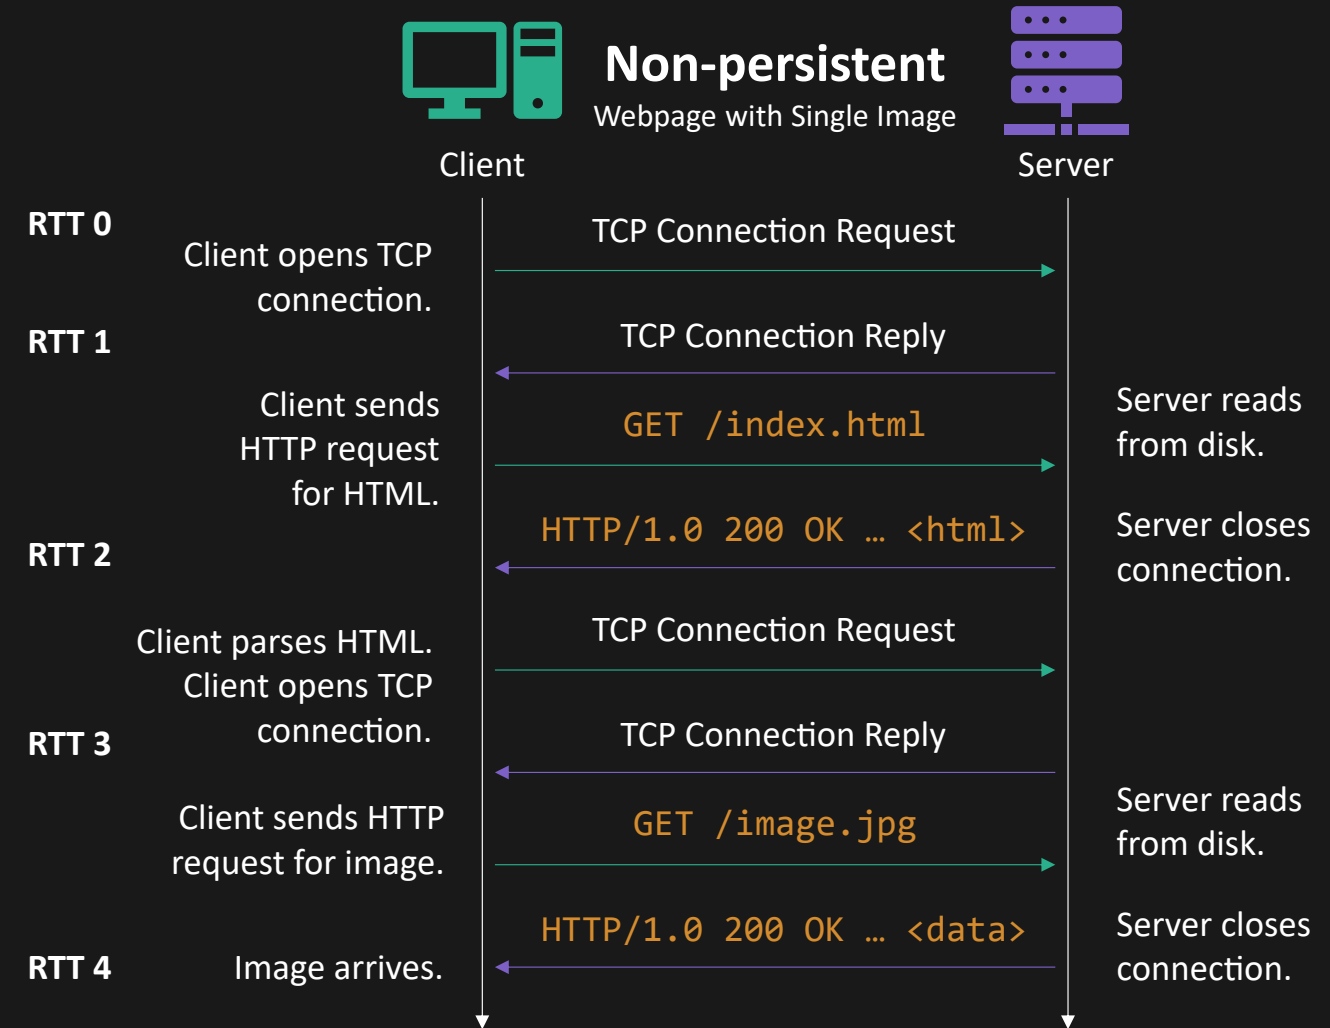


KEEP-ALIVE - PERSISTENT CONNECTIONS

HTTP 1.0 PROBLEM:

- Each request opens new connection
 - Starting up is slow
 - Takes several packets
- Short transfers are hard on TCP
 - Stuck in “slow start” phase of TCP connection
 - Loss recovery is poor when windows are small
- Lots of extra connections
 - Increases server state/processing

Sources: Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).



KEEP-ALIVE - PERSISTENT CONNECTIONS

HTTP 1.1 SOLUTION:

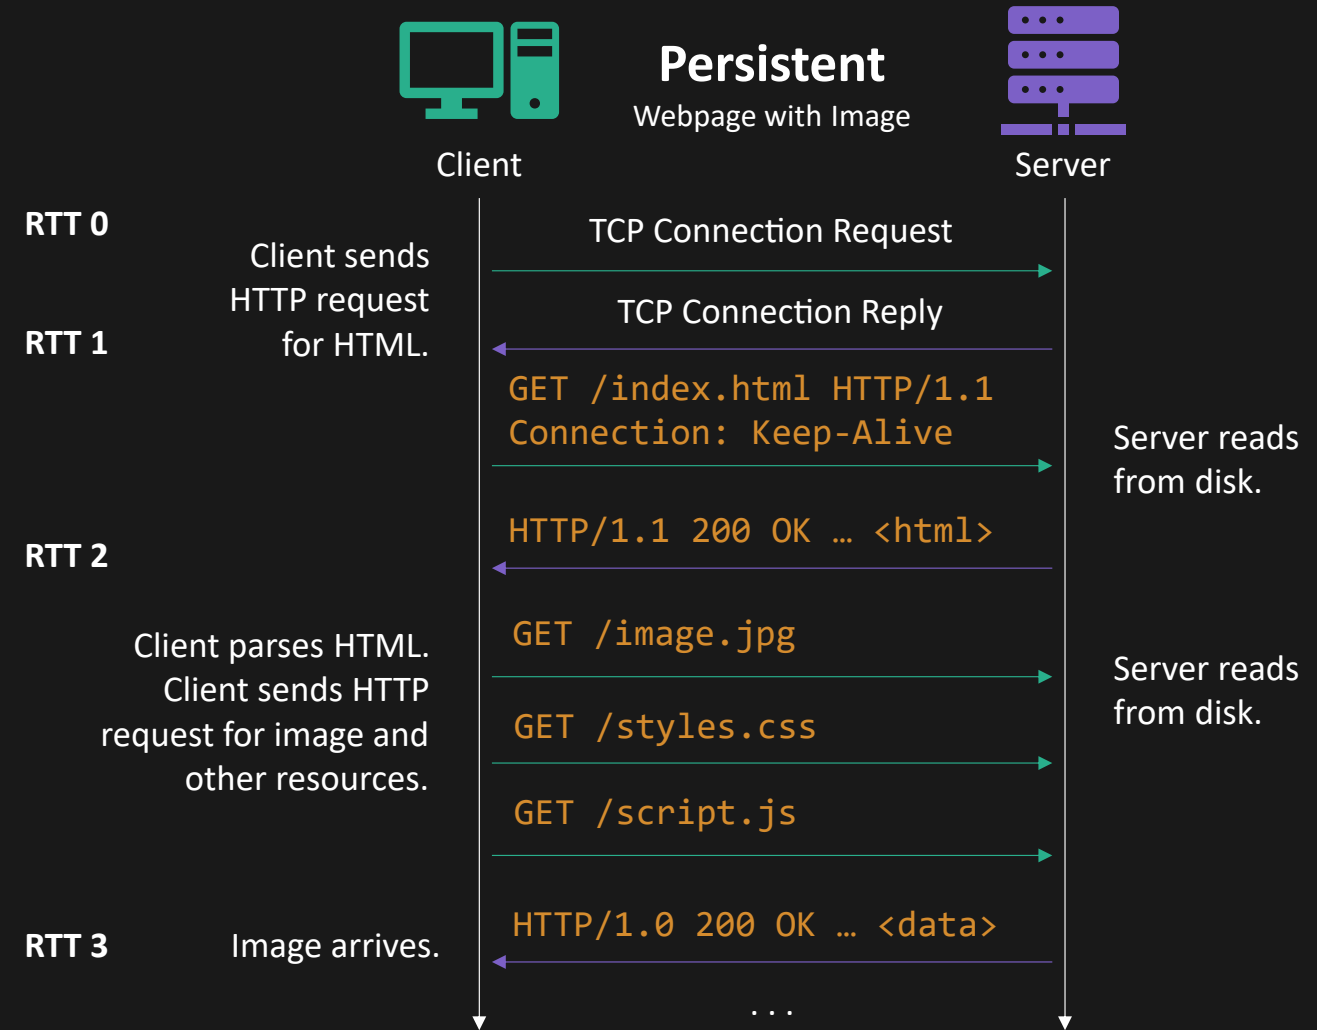
- Keeps connection open for a time after server response so that multiple requests can ride on single connection
- Reduced connection setup overhead.

```
GET /index.html HTTP/1.1
Connection: Keep-Alive
```

```
... Multiple HTTP requests ...
```

```
GET /script.js HTTP/1.1
Connection: Close
```

Sources: Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).



NON-PERSISTENT

- HTTP/1.0
- Server parses request, responds, and closes TCP connection
- 2 RTTs to fetch each object
- Each object transfer suffers from slow start
- Most 1.0 browsers used parallel TCP connections.

PERSISTENT

- Default for HTTP/1.1
- On same TCP:
 - connection: server,
 - parses request, responds, parses new request, ...
- Client sends requests for all referenced objects as soon as it receives base HTML.
- Fewer RTTs and less slow start.
- Prefetching

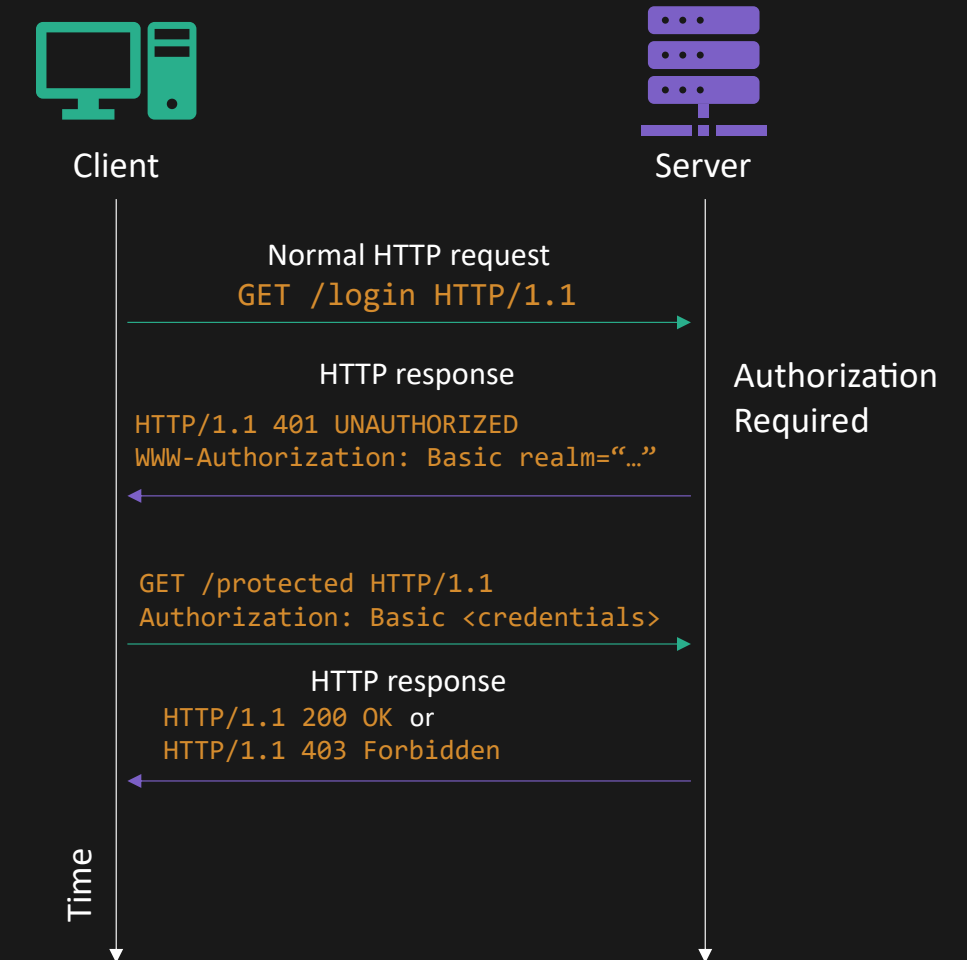
Sources: Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).

BASIC AUTHENTICATION

- When challenged, the client sends their user ID and password in the clear to the server.
- Not secure enough (snooping is easy) but useful for simple things

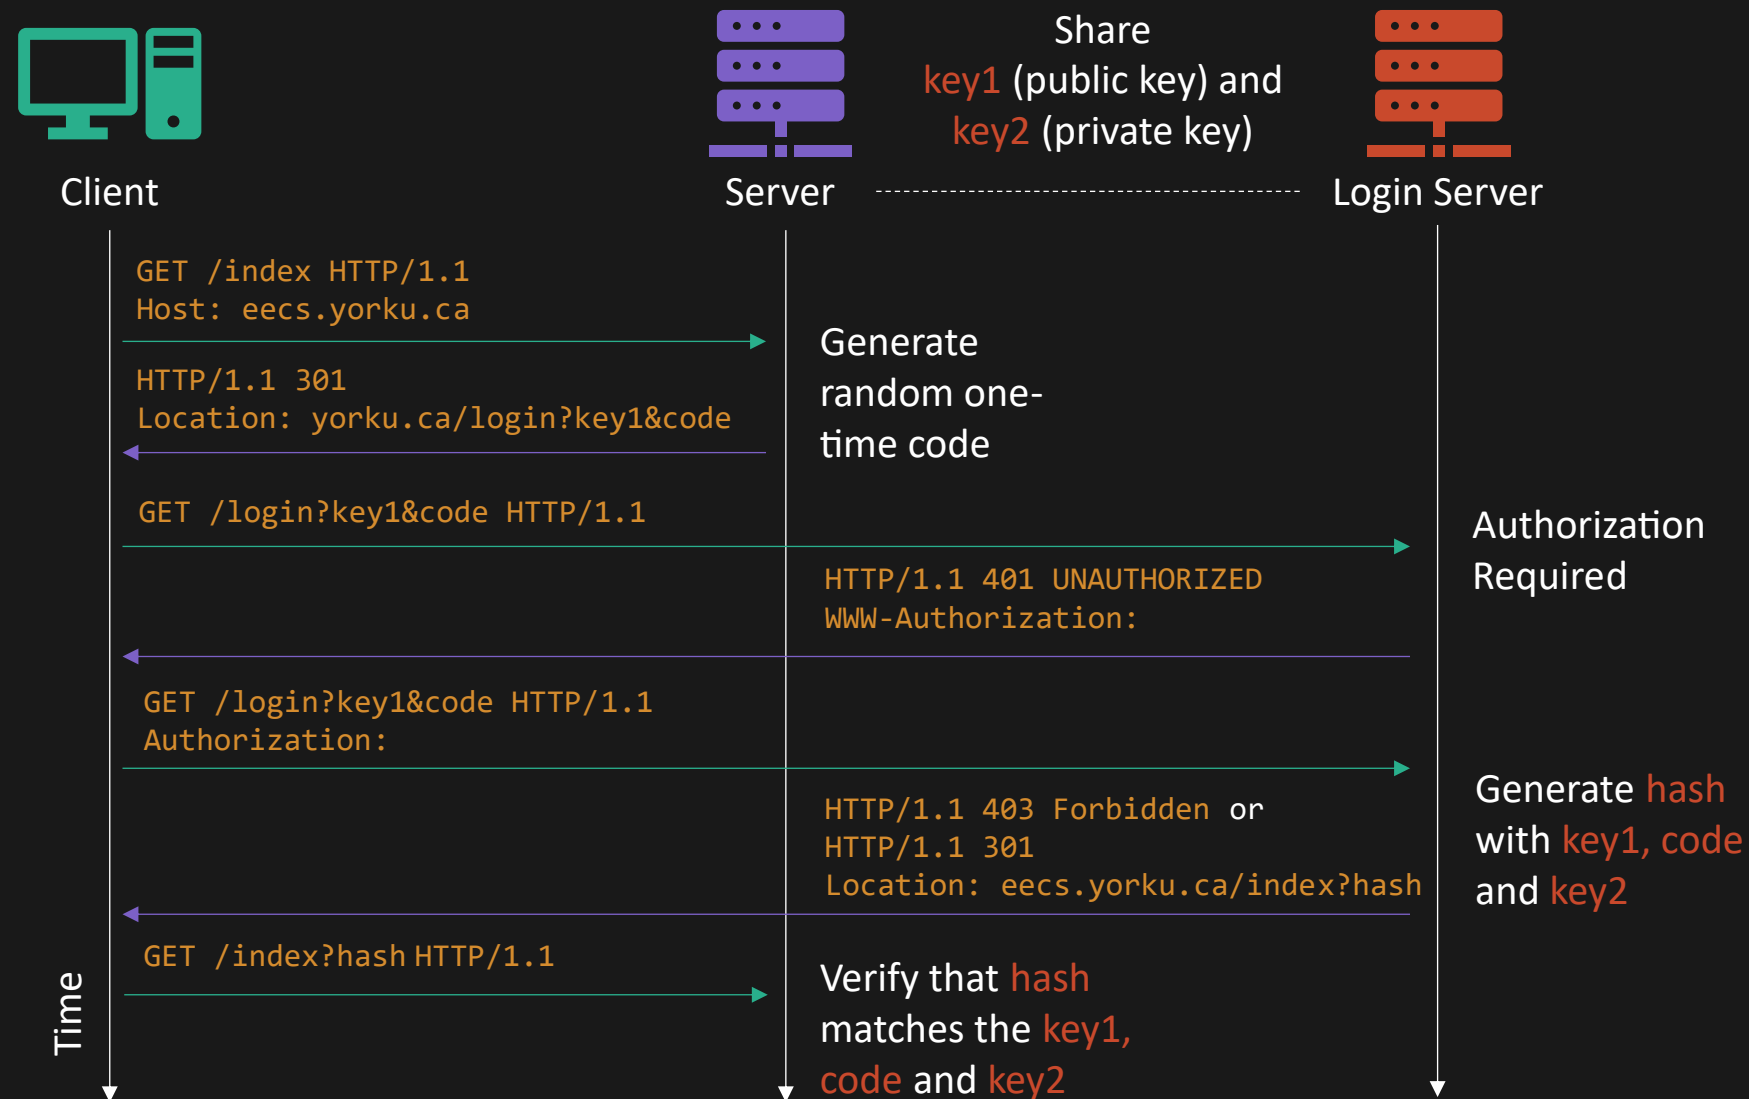
CLIENT-SERVER INTERACTION: AUTHENTICATION

- Authentication goal: Control access to server documents
- Stateless: client must present authorization in each request
- Authorization: typically name, password
 - Sends `Authorization` header line in request
 - If no authorization presented, server refuses access, sends `WWW-Authenticate` header line in response.
- Browser caches name and password so that user does not have to repeatedly enter it.



Sources: Protocols and HTTP, Web Browsers and Servers, Caching, DNS, David A. Penny (2001).

OPEN AUTHENTICATION (OAUTH)



WEB CACHES (PROXY SERVERS)


WEB CACHING

- Improve performance
 - Scalability
 - Response time
 - Load balancing
 - Availability
 - Saves network and server resources
- Proxy cache
 - Done at the client-side

GOAL

Fill client request without going to origin server.

- User sets browser: Web accesses via web cache
- Client sends all HTTP requests to web cache
 - If object at web cache, web cache immediately returns object in HTTP response
 - Else requests object from origin server, then returns HTTP response to client



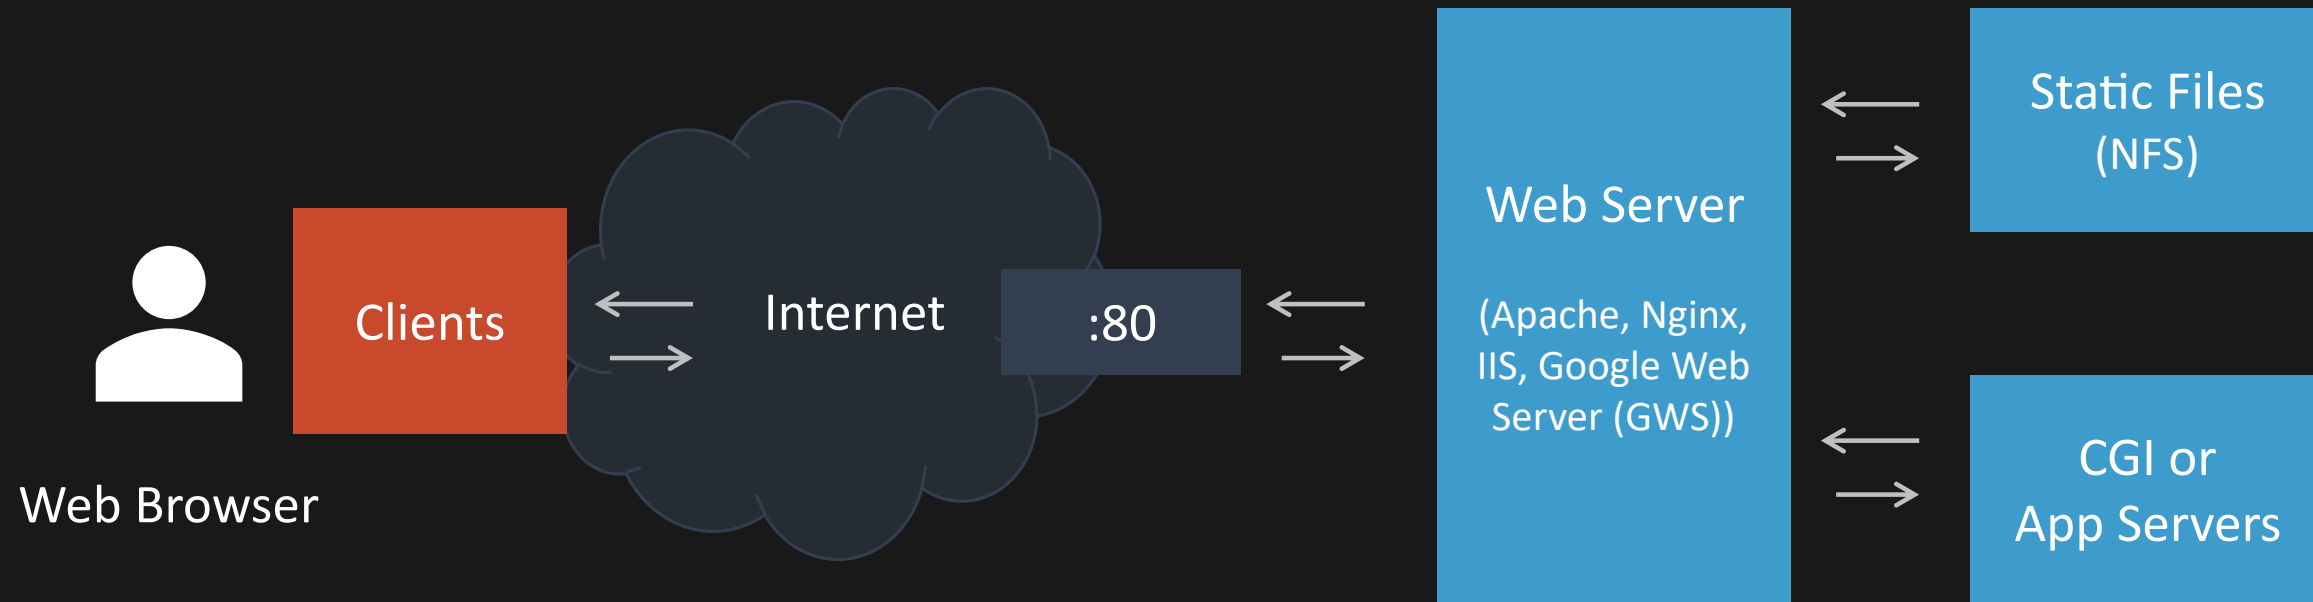
HTTP

SERVICES

HTTP CLIENT

Web browser = TCP client + HTTP + HTML/CSS/JS + DOM

```
1 public class HTTPClient {
2     public static void main(String[] args) throws Exception {
3         PrintStream out = System.out;
4         try (Scanner in = new Scanner(new URL(args[0])).openStream())
5             while (in.hasNextLine()) {
6                 out.println(in.nextLine());
7             }
8     }
9 }
10 }
```



HTTP SERVER

Web server = TCP Server + Port 80 + HTTP.

- Built-in static file serving
- Built-in scalability
- Built-in security (HTTPS + auth) via `.htaccess`
- Built-in telemetry (logs and error logs)
- Extensibility: PHP (can violate view migration); CGI (good & language agnostic); App Servers (best): Tomcat JSP, WebSphere, WebLogic, NodeJS, ASP.NET, ...).

COMMON GATEWAY INTERFACE (CGI)

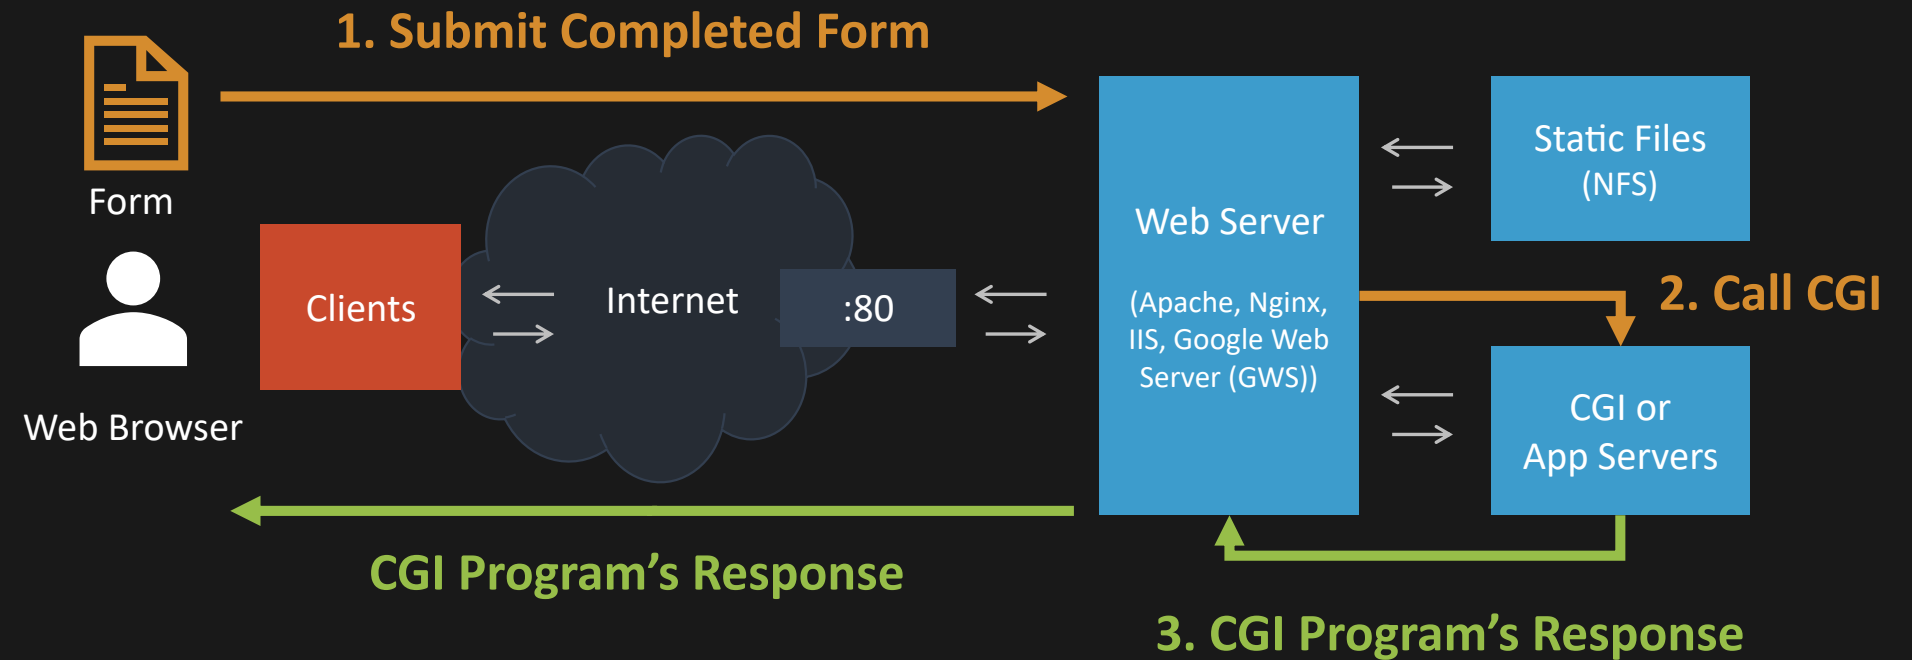
The Common Gateway Interface (CGI) is a simple interface for running external programs, software or gateways under an information server in a platform-independent manner.

1. The client makes a request by specifying a URL and additional info.
2. The webserver:
 - a. Receives the request. (in the URL)
 - b. Identifies the request as a CGI request.
 - c. Locates the program corresponding to the request.
 - d. Starts up the handling program (heavy weight process creation!!)
 - e. Feeds request parameters to the handler (through stdin or environment variables).
3. The Handler:
 - a. Executes with the given request parameters
 - b. The Output of the handler is sent via stdout back to the webserver for rerouting back to the requesting web browser.
 - c. Output is typically a web page. Could be plaintext, JSON, or XML.
 - d. Terminates.

Sources: The Common Gateway Interface (CGI) Version 1.1 and CGI - CSC309F Programming on the Web, David A. Penny (2001).

COMMON GATEWAY INTERFACE (CGI)

The Common Gateway Interface (CGI) is a simple interface for running external programs, software or gateways under an information server in a platform-independent manner.



Sources: The Common Gateway Interface (CGI) Version 1.1 and CGI - CSC309F Programming on the Web, David A. Penny (2001).

MORE ON CGI

CGI defines the abstract parameters, known as metavariables, which describe the client’s request. Together with a concrete programmer interface this specifies a platform-independent interface between the script and the HTTP server.

METAVARIABLES

AUTH_TYPE	REMOTE_IDENT	QUERY_STRING
CONTENT_LENGTH	REMOTE_USER	REMOTE_ADDR
CONTENT_TYPE	REQUEST_METHOD	REMOTE_HOST
GATEWAY_INTERFACE	SCRIPT_NAME	SERVER_PROTOCOL
PATH_INFO	SERVER_NAME	SERVER_SOFTWARE
PATH_TRANSLATED	SERVER_PORT	

EXAMPLE

```
GET /~vwchu/test.cgi?key=value&a=1234 HTTP/1.1
```

Item	Value
GATEWAY_INTERFACE	CGI/1.1
SERVER_PROTOCOL	HTTP/1.1
SERVER_SOFTWARE	Apache/2.4.48 (Unix) PHP/7.4.20 OpenSSL/1.0.2k mod_wsgi/4.7.1 Python/3.8
SERVER_NAME	www.eecs.yorku.ca
SERVER_PORT	443
REQUEST_METHOD	GET
SCRIPT_NAME	/~vwchu/test.cgi
QUERY_STRING	key=value&a=1234
REMOTE_ADDR	130.63.96.170
REMOTE_PORT	39256
REMOTE_USER	vwchu

Sources: CGI - CSC309F Programming on the Web, David A. Penny (2001)..

ASIDE

REFACTORING OUR CODE GOALS

- Code maintainability
- Modularity
- Encapsulation & abstraction
- Separation of concerns
- Handling complexity
- We shouldn't have to rebuild the machinery that runs the webserver every time we want to create a web service!

```

1 public class HTTPServer extends Thread {
2     ...
3     public void run() {
4         log.printf("Connected to %s:%d\n", client.getInetAddress(),
5                     client.getPort());
6
7         try {
8             Socket client = this.client;
9             Scanner req = new Scanner(client.getInputStream());
10            PrintStream res = new PrintStream(client.getOutputStream(), true);
11        } {
12            String request = req.nextLine();
13            String method, uri, version;
14
15            StringWriter response = new StringWriter();
16            PrintWriter writer = new PrintWriter(response);
17
18            try (Scanner parse = new Scanner(request)) {
19                method = parse.next().toUpperCase();
20                uri = parse.next();
21                version = parse.next().toUpperCase();
22            }
23
24            int status = 200;
25
26            Map<String, String> reqHeaders = getHeaders(req);
27            Map<String, Object> resHeaders = new HashMap<>();
28
29            resHeaders.put("Server", "Java HTTP Server : 1.0");
30            resHeaders.put("Date", new Date());
31            resHeaders.put("Content-Type", "text/plain");
32
33            try {
34                if (!method.equals("GET") && !method.equals("HEAD") &&
35                    !method.equals("POST")) {
36                    status = 501;
37                } else if (!version.equals("HTTP/1.1")) {
38                    status = 505;
39                } else {
40                    String[] components = getURISplit(uri);
41                    uri = components[0];
42                    Map<String, String> qs = getQueryStrings(components[1]);
43                    String body = null;
44
45                    if (method.equals("POST")) {
46                        body = getRequestBody(req);
47                    }
48                    ...

```

```

1         switch (uri) {
2             case "/":
3                 response.println("Hello! Welcome to the server.");
4                 break;
5
6             case "/calc":
7                 ...
8
9             case "/add": case "/subtract": case "/multiply":
10            case "/divide": case "/exponent":
11                ...
12
13            case "/students":
14                ...
15
16            default:
17                status = 404;
18        }
19    }
20    } catch (Exception err) {
21        log.println(err.getMessage());
22        err.printStackTrace(log);
23        status = 500;
24    }
25
26    if (status != 200 && response.toString().length() == 0) {
27        writer.println(httpResponseCodes.get(status));
28    }
29
30    String responseText = response.toString();
31    resHeaders.put("Content-Length", responseText.getBytes().length);
32    sendHeaders(res, status, resHeaders);
33
34    if (!method.equals("HEAD")) {
35        res.println(responseText);
36    }
37
38    res.flush(); // flush character output stream buffer
39    } catch (Exception e) {
40        log.println(e);
41    } finally {
42        log.printf("Disconnected from %s\n", clientAddress);
43    }
44    }
45
46    public static void main(String[] args) throws Exception {
47        ...
48    }
49    }

```

ASIDE

REFACTORING OUR CODE GOALS

- Code maintainability
- Modularity
- Encapsulation & abstraction
- Separation of concerns
- Handling complexity
- We shouldn't have to rebuild the machinery that runs the webserver every time we want to create a web service!

```

1 public class HTTPServer extends Thread {
2     ...
3     public void run() {
4         log.printf("Connected to %s:%d\n", client.getInetAddress(),
5                     client.getPort());
6
7         try {
8             Socket client = this.client;
9             Scanner req = new Scanner(client.getInputStream());
10            PrintStream res = new PrintStream(client.getOutputStream(), true);
11        } {
12            String request = req.nextLine();
13            String method, uri, version;
14
15            StringWriter response = new StringWriter();
16            PrintWriter writer = new PrintWriter(response);
17
18            try (Scanner parse = new Scanner(request)) {
19                method = parse.next().toUpperCase();
20                uri = parse.next();
21                version = parse.next().toUpperCase();
22            }
23
24            int status = 200;
25
26            Map<String, String> reqHeaders = reqHeaders(req);
27            Map<String, String> resHeaders = resHeaders(res);
28
29            resHeaders.put("Server", "Java HTTP Server : 1.0");
30            resHeaders.put("Date", new Date());
31            resHeaders.put("Content-Type", "text/plain");
32
33            try {
34                if (!method.equals("GET") && !method.equals("HEAD") &&
35                    !method.equals("POST")) {
36                    status = 501;
37                } else if (!version.equals("HTTP/1.1")) {
38                    status = 505;
39                } else {
40                    String[] components = getURISComponents(uri);
41                    uri = components[0];
42                    Map<String, String> qs = getQueryStrings(components[1]);
43                    String body = null;
44
45                    if (method.equals("POST")) {
46                        body = getRequestBody(req);
47                    }
48                    ...

```

The server code and the application code are all mixed together in the `run` method.

```

1         switch (uri) {
2             case "/":
3                 response.println("Hello! Welcome to the server.");
4                 break;
5
6             case "/calc":
7                 ...
8
9             case "/add": case "/subtract": case "/multiply":
10             case "/divide": case "/exponent":
11                 ...
12
13             case "/students":
14                 ...
15
16             default:
17                 status = 404;
18         }
19     }
20 } catch (Exception err) {
21     log.println(err.getMessage());
22     err.printStackTrace(log);
23     status = 500;
24 }
25
26 if (status != 200 && response.toString().length() == 0) {
27     log.println("Error: " + status);
28 }
29
30 String responseText = response.toString();
31 resHeaders.put("Content-Length", responseText.getBytes().length);
32 sendHeaders(res, status, resHeaders);
33
34 if (!method.equals("HEAD")) {
35     res.println(responseText);
36 }
37
38 res.flush(); // flush character output stream buffer
39 } catch (Exception e) {
40     log.println(e);
41 } finally {
42     log.printf("Disconnected from %s\n", clientAddress);
43 }
44 }
45
46 public static void main(String[] args) throws Exception {
47     ...
48 }
49 }

```


ASIDE

REFACTORING OUR CODE GOALS

- Code maintainability
- Modularity
- Encapsulation & abstraction
- Separation of concerns
- Handling complexity
- We shouldn't have to rebuild the machinery that runs the webserver every time we want to create a web service!

```
1 public class HTTPServer extends Thread {
2     ...
3     public void run() {
4         ...
5         status = doRequest(method, uri, reqHeaders, qs, body, status, resHeaders, writer);
6         ...
7     }
8     private int doRequest(String method, String uri, Map<String, String> reqHeaders,
9                           Map<String, String> qs, String body, int status, Map<String, Object> resHeaders,
10                          PrintWriter response) throws Exception {
11         switch (uri) {
12             case "/":
13                 response.println("Hello! Welcome to the server.");
14                 break;
15             case "/calc":
16                 ...
17             case "/add":      case "/subtract": case "/multiply":
18             case "/divide": case "/exponent":
19                 ...
20             case "/students":
21                 ...
22             default:
23                 return 404;
24         }
25         return status;
26     }
27     ...
28 }
```

First, move the application code out of the `run` method into a separate `doRequest` method. Note that the `doRequest` method requires a lot of arguments, too many.

ASIDE

REFACTORING OUR CODE GOALS

- Code maintainability
- Modularity
- Encapsulation & abstraction
- Separation of concerns
- Handling complexity
- We shouldn't have to rebuild the machinery that runs the webserver every time we want to create a web service!

```

1 public class HTTPServer extends Thread {
2     ...
3     private void doRequest(
4         Request req,
5         Response res) throws Exception {
6         ...
7     }
8     ...
9     class Request {
10         final String method;
11         final String uri;
12         final Map<String, String> headers;
13         final Map<String, String> qs;
14         final String body;
15
16         public Request(
17             String method,
18             String uri,
19             Map<String, String> headers,
20             Map<String, String> qs,
21             String body
22         ) {
23             this.method = method;
24             this.uri = uri;
25             this.headers = headers;
26             this.qs = qs;
27             this.body = body;
28         }
29     }

```

```

1     ...
2     class Response {
3         private int status;
4         final Map<String, Object> headers;
5         final StringWriter response = new StringWriter();
6         final PrintWriter out = new PrintWriter(response);
7
8         public Response(int status, Map<String, Object> headers) {
9             this.status = status;
10            this.headers = headers;
11        }
12        public int getStatus() {
13            return status;
14        }
15        public void setStatus(int status) {
16            this.status = status;
17        }
18    }
19 }

```

So, we will group them into `Request` and `Response` classes.

ASIDE

REFACTORING OUR CODE GOALS

- Code maintainability
- Modularity
- Encapsulation & abstraction
- Separation of concerns
- Handling complexity
- We shouldn't have to rebuild the machinery that runs the webserver every time we want to create a web service!

```
1 public abstract class HTTPServer extends Thread {
2     ...
3     protected abstract void doRequest(Request req, Response res) throws Exception;
4     ...
5 }
6
7 public class MainService extends HTTPServer {
8     public MainService(Socket client) {
9         super(client);
10    }
11    protected void doRequest(Request req, Response res) throws Exception {
12        switch (req.uri) {
13            case "/":
14                res.out.println("Hello! Welcome to the server.");
15                break;
16            case "/calc":
17                ...
18            case "/add":      case "/subtract": case "/multiply":
19            case "/divide": case "/exponent":
20                ...
21            case "/students":
22                ...
23            default:
24                res.setStatus(404);
25        }
26    }
27    public static void main(String[] args) throws Exception {
28        ...
29    }
30 }
```

Take the `doRequest` method entirely out of the `HTTPServer` class into a subclass, `MainService`.

ASIDE

REFACTORING OUR CODE GOALS

- Code maintainability
- Modularity
- Encapsulation & abstraction
- Separation of concerns
- Handling complexity
- We shouldn't have to rebuild the machinery that runs the webserver every time we want to create a web service!

```

1 public abstract class HTTPServer extends Thread {
2     ...
3     protected abstract void doRequest(RequestContext req, ResponseContext res) throws Exception;
4     public void run() {
5         ...
6         RequestContext request = new RequestContext(client.getInputStream());
7         ResponseContext response = new ResponseContext(client.getOutputStream());
8         request.processRequest(response);
9         doRequest(request, response);
10        ...
11    }
12 }
13
14 public class RequestContext {
15     ...
16 }
17
18 public class ResponseContext {
19     ...
20 }

```

Put the `Request` class and all of the methods in `HTTPServer` class related to handling a request out of the `HTTPServer` class into a new class called `RequestContext`. Do the same for the `Response` class and all of the methods in `HTTPServer` class related to forming a response, resulting in the `ResponseContext` class.

ASIDE

REFACTORING OUR CODE GOALS

- Code maintainability
- Modularity
- Encapsulation & abstraction
- Separation of concerns
- Handling complexity
- We shouldn't have to rebuild the machinery that runs the webserver every time we want to create a web service!

```
1 public interface Service {
2     public void doRequest(RequestContext request, ResponseContext response) throws Exception;
3 }
4
5 public class CalcOpService implements Service {
6     public void doRequest(RequestContext request, ResponseContext response) { ... }
7 }
8
9 public class CalcService implements Service {
10    public void doRequest(RequestContext request, ResponseContext response) { ... }
11 }
12
13 public class HelloService implements Service {
14    public void doRequest(RequestContext request, ResponseContext response) {
15        response.body.println("Hello! Welcome to the server.");
16    }
17 }
18
19 public class StudentsService implements Service {
20    public void doRequest(RequestContext request, ResponseContext response) { ... }
21 }
```

Split the `MainService` class into each of the different services. One service per class. Each class has its own `doRequest` method.

ASIDE

REFACTORING OUR CODE GOALS

- Code maintainability
- Modularity
- Encapsulation & abstraction
- Separation of concerns
- Handling complexity
- We shouldn't have to rebuild the machinery that runs the webserver every time we want to create a web service!

That leaves our `MainService` class with the barest of `run` methods and a map that registers the available service handlers to their respective resource paths. Later, we will see how Tomcat does it via annotations.

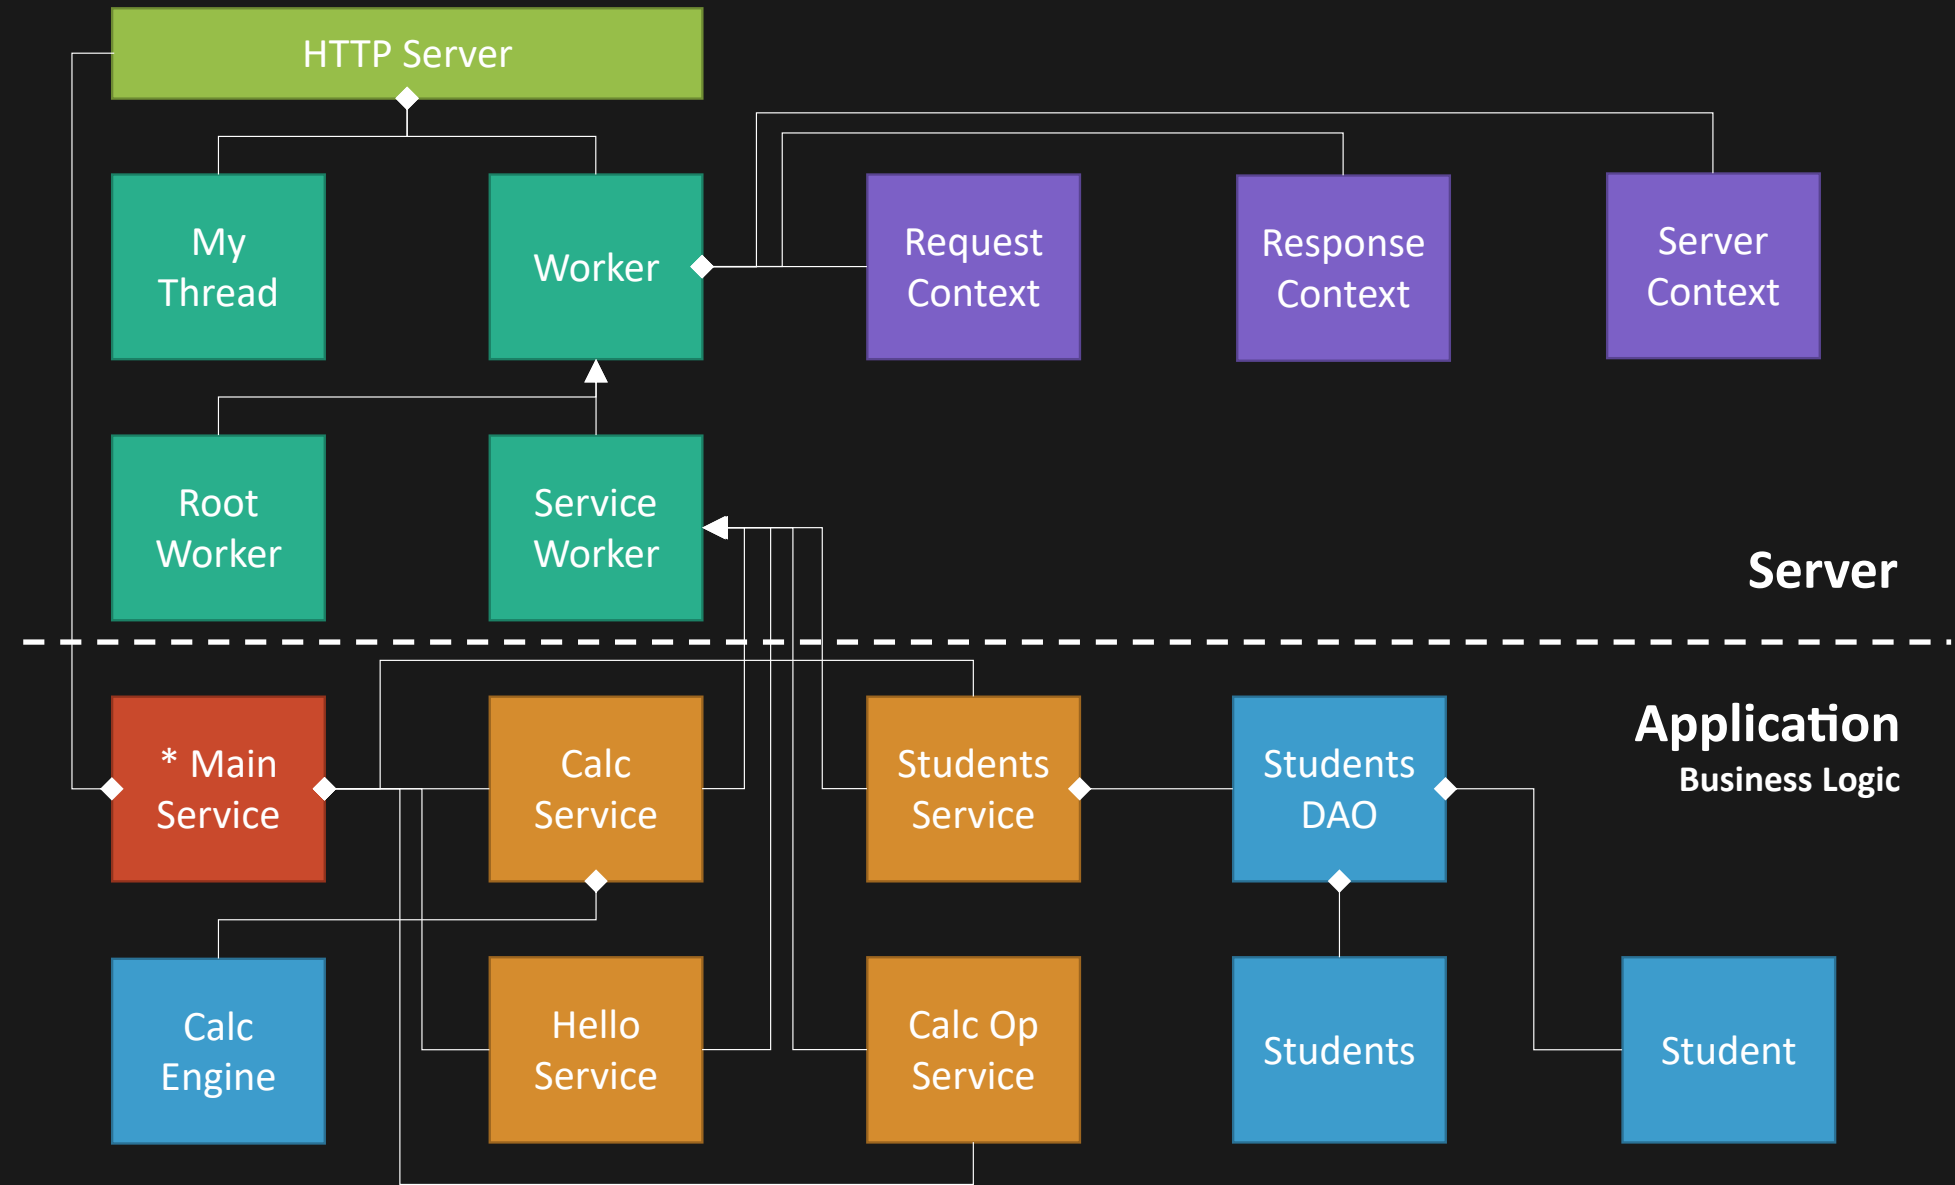
```

1 public class MainService extends HTTPServer {
2     protected static final Map<String, Service> services = new HashMap<>();
3     static {
4         CalcOpService coService = new CalcOpService();
5         services.put("/", new HelloService());
6         services.put("/calc", new CalcService());
7         services.put("/add", coService);
8         services.put("/subtract", coService);
9         services.put("/multiply", coService);
10        services.put("/divide", coService);
11        services.put("/exponent", coService);
12        services.put("/students", new StudentsService());
13    }
14    public MainService(Socket client) {
15        super(client);
16    }
17    protected void doRequest(RequestContext req, ResponseContext res) throws Exception {
18        if (services.containsKey(req.uri)) {
19            services.get(req.uri).doRequest(req, res);
20        } else {
21            res.setStatus(404);
22        }
23    }
24    public static void main(String[] args) throws Exception {
25        int port = 0;
26        InetAddress host = InetAddress.getLocalHost(); // .getLoopbackAddress();
27        try (ServerSocket server = new ServerSocket(port, 0, host)) {
28            log.printf("Server listening on %s:%d\n", server.getInetAddress(), server.getLocalPort());
29            while (true) {
30                (new MainService(server.accept())).start();
31            }
32        }
33    }
34 }

```

If we continue refactoring, we might end up with something like this...

Despite all this added complexity, it's okay. Because to us, web app developers, the web server's code is a black-box that we access via public APIs.



This slide is intentionally left blank.

Return to [Course Page](#) or [Part II](#).