

UNIVERSITY OF
Southampton

Content Negotiation

COMP3220 Web Infrastructure

Dr Nicholas Gibbins – nmg@ecs.soton.ac.uk

HTTP content negotiation

HTTP can deliver different representations of a resource based on client preferences

Two styles of content negotiation:

- Server-driven (where the server makes the final choice of representation)
- Client-driven (where the client makes the final choice of representation)

Three areas for negotiation

- Media type (Accept: and Content-Type:)
- Language (Accept-Language: and Content-Language:)
- Encoding (Accept-Encoding: and Content-Encoding:)

Server-driven content negotiation

1. Client tells the server what it is able to accept in a request header
2. Server chooses an appropriate representation to return to the client
3. Server tells the client what its choice was in a response header

Client request headers all have the same structure:

headername: choice₁; q=quality₁, choice₂; q=quality₂, choice₃; q=quality₃ ...

The header name and the values for *choice_n* all depend on what is being negotiated

The quality values *quality_n* are numeric values between 0.0 and 1.0

Media type negotiation

Client request header: Accept:

Server response header: Content-Type:

Choice values are Internet Media Types:


- text/plain
- text/html
- application/pdf
- image/jpeg
- image/png
- image/svg+xml
- audio/mpeg
- ...

Media type negotiation



```
GET / HTTP/1.1  
Host: example.org  
Accept: text/html; q=1.0, text/plain; q=0.5
```

```
HTTP/1.1 200 OK  
Content-Type: text/html  
  
<html>  
  <head>  
    <title>Example, Inc. Homepage</title>  
  </head>  
  <body><h1>welcome to Example!</h1>...</body>  
</html>
```



Language negotiation

Client request header: Accept-Language:

Server response header: Content-Language:

Choice values are typically ISO 639 language codes:

- en for English
- de for German
- zh for Chinese

Convention for indicating regional dialects by adding ISO 3166-1 country codes:

- en-GB for British English
- de-CH for Swiss German


Language negotiation



```
GET / HTTP/1.1
Host: example.org
Accept-Language: de; q=1.0, en-gb; q=0.5
```

```
HTTP/1.1 200 OK
Content-Type: text/html
Content-Language: de

<html>
  <head>
    <title>Example, Inc. Homepage</title>
  </head>
  <body><h1>willkommen zu Example!</h1></body>
</html>
```



Encoding negotiation

Client request header: Accept-Encoding:

Server response header: Content-Encoding:

Typically used to specify whether the body of a HTTP message is compressed

Possible choice values:

- br (Brotli compressed data format)
- deflate (zlib compression)
- gzip (LZ77 compression – the most common choice)
- compress (LZW compression)
- identity (no encoding)

The Vary: header


Origin servers can indicate what parts of the request message (aside from the method, Host: header, and request target) they'll use to select an appropriate representation



```
GET / HTTP/1.1  
Host: example.org
```

```
HTTP/1.1 200 OK  
Content-Type: text/html  
Vary: accept, accept-language
```

```
...
```




The Content-Location: header

Used to indicate a direct URI to use to access a content-negotiated resource



```
GET /logo HTTP/1.1  
Host: example.org  
Accept: image/png; q=1.0, image/gif; q=1.0, image/jpeg; q=0.8
```

```
HTTP/1.1 200 OK  
Content-Type: image/gif  
Content-Location: /logo.gif  
...
```



Client-driven content negotiation

1. Client requests a resource representation
2. Server returns 300 Multiple Choices with a list of URIs for alternative representations in the body
3. Client requests a representation of one of those URIs

No standard for how the server lists the alternative representations!


Client-driven content negotiation



```
GET / HTTP/1.1  
Host: example.org
```

```
HTTP/1.1 300 Multiple Choices  
Content-Type: application/json
```

```
[{"type": "image/jpeg",  
  "uri": "http://example.org/me.jpg"},  
 {"type": "image/png",  
  "uri": "http://example.org/me.png"}]
```




```
GET / HTTP/1.1  
Host: example.org
```

How not to negotiate content



```
GET / HTTP/1.1  
Host: example.org  
User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS x 10.15;  
rv:80.0) Gecko/20100101 Firefox/80.0
```

You're a Firefox browser, so I'd better send you the Firefox version of that resource



How not to negotiate content (Browser Wars, Part II)

Best viewed with



Client Hints

Content negotiation using the User-Agent: header is not considered good practice

Accept: is a very crude basis for content negotiation

Client Hints is a proposed HTTP extension that allows browsers to state their capabilities and preferences

- Device memory
- Downlink speed
- Viewport width
- Device pixel ratio
- Reduced data usage



A close-up, low-angle shot of a computer screen. The screen is filled with a blue, grid-like pattern. In the foreground, a white rectangular box contains the text 'http://www' in a bold, black, sans-serif font. The text is slightly out of focus, and the background is a blurred grid of blue squares.

http://www

Exercise: Content Negotiation

Exercise: Content Negotiation

```
curl -H "[header]" [uri]
```

- Generates a HTTP request including the specified header

Use curl to study the following requests:

```
curl -v https://www.debian.org/
```

```
curl -v -H "Accept-Language: de" https://www.debian.org/
```

```
curl -v -H "Accept-Language: fi" https://www.debian.org/
```

Next Lecture: Conditional Requests