

# Web Server Design

## Lecture 5 – Range and Partial Content

Old Dominion University

Department of Computer Science

CS 431/531 Fall 2022

**Sawood Alam** <salam@cs.odu.edu>

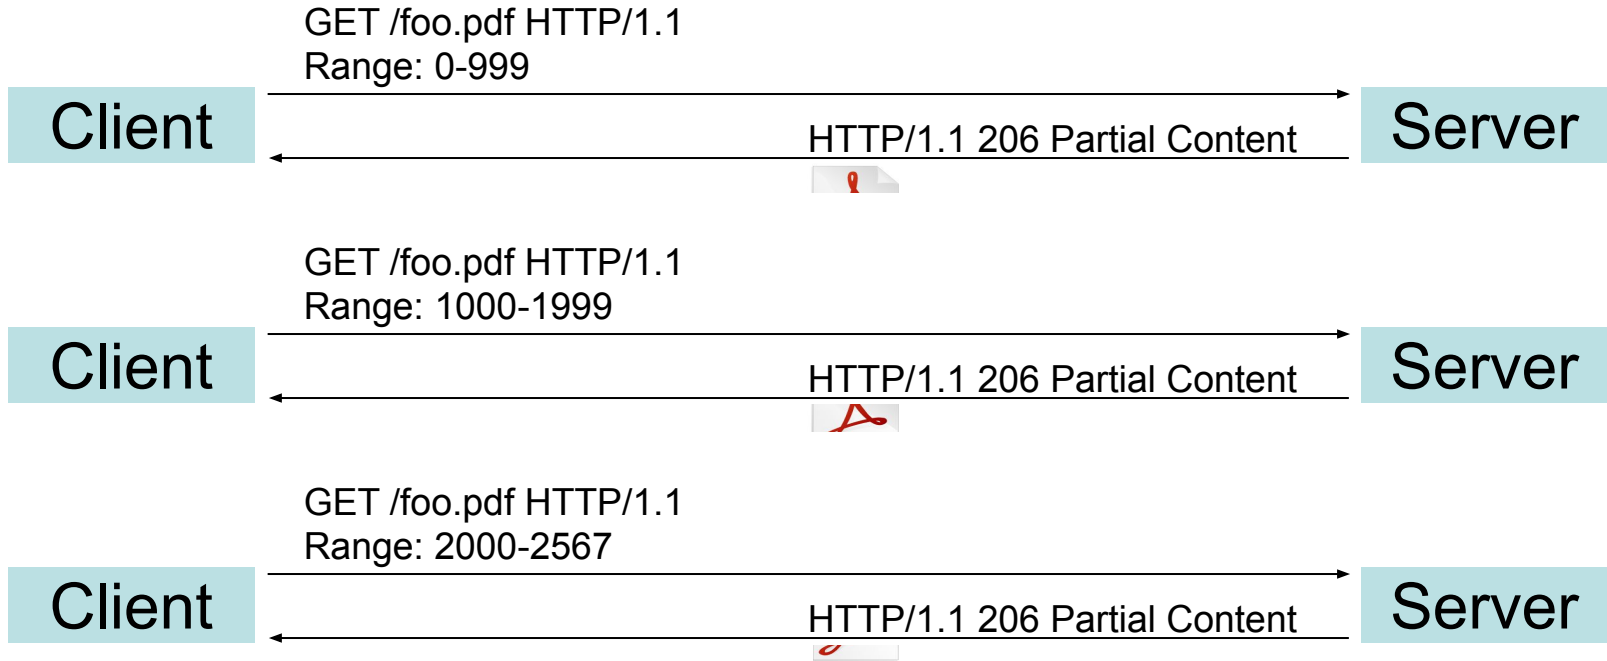
2022-09-28

Original slides by Michael L. Nelson

# GET



# GET



“The partial GET method is intended to reduce unnecessary network usage by allowing partially-retrieved entities to be completed without transferring data already held by the client.”

# Server Signaling to the Client

- The server tells the client it does/does not support partial GETs *on a particular URI* through the “Accept-Ranges” header
  - “bytes” or “none”
  - RFC 7233
- A client can request partial GETs without a signal from the server, but the client must be prepared for the server to say no:
  - Accept-Ranges: none

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
HEAD /~mln/teaching/cs595-s06/a1-test/1/ HTTP/1.1
Host: www.cs.odu.edu

HTTP/1.1 200 OK
Date: Sun, 12 Feb 2006 20:59:01 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 1936
Content-Type: text/html
X-Pad: avoid browser bug

Connection closed by foreign host.
```

# Client Requesting a Range

- “Range:” header
  - Section 3.1, RFC 7233
  - Server MAY ignore the header
    - 200 vs. 206 tells the story
- Byte range examples (Section 2.1, RFC 7233):

Examples of byte-ranges-specifier values (assuming an entity-body of length 10000):

- The first 500 bytes (byte offsets 0-499, inclusive):

`bytes=0-499`

- The second 500 bytes (byte offsets 500-999, inclusive):

`bytes=500-999`

- The final 500 bytes (byte offsets 9500-9999, inclusive):

`bytes=-500`

- Or `bytes=9500-`

- The first and last bytes only (bytes 0 and 9999):

`bytes=0-0,-1`

- Several legal but not canonical specifications of the second 500 bytes (byte offsets 500-999, inclusive):

`bytes=500-600,601-999`

`bytes=500-700,601-999`

# Server Response:

## 206 Partial Content Response Code & Headers Content-Range, Content-Length

- Server responds with:
  - HTTP/1.1 206 Partial Content
- The “Content-Length” response header is now the number of bytes *in this response*, not the number of bytes in the entire entity
- “Content-Range” header is present
  - Content-Range: bytes *start-end/total*

# Long-Lived Session

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=0-9
```

```
HTTP/1.1 206 Partial Content
Date: Sun, 19 Feb 2006 21:59:31 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 10
Content-Range: bytes 0-9/1936
Content-Type: text/html
```

```
<html>
<he
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=10-19
```

```
HTTP/1.1 206 Partial Content
Date: Sun, 19 Feb 2006 21:59:42 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 10
Content-Range: bytes 10-19/1936
Content-Type: text/html
```

```
ad>
<title
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=20-29
```

```
HTTP/1.1 206 Partial Content
Date: Sun, 19 Feb 2006 21:59:52 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 10
Content-Range: bytes 20-29/1936
Content-Type: text/html
```

```
>
Michael
Connection closed by foreign host.
```

(Line breaks added after entity data for clarity)

# Same Session, Pipelined

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=0-9

GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=10-19

GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=20-29
Connection: close
```

(No line breaks added)

```
HTTP/1.1 206 Partial Content
Date: Mon, 20 Feb 2006 01:11:58 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 10
Content-Range: bytes 0-9/1936
Content-Type: text/html
```

```
<html>
<heHTTP/1.1 206 Partial Content
Date: Mon, 20 Feb 2006 01:11:58 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 10
Content-Range: bytes 10-19/1936
Content-Type: text/html
```

```
ad>
<titleHTTP/1.1 206 Partial Content
Date: Mon, 20 Feb 2006 01:11:58 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 10
Content-Range: bytes 20-29/1936
Connection: close
Content-Type: text/html
```

```
>
Michael Connection closed by foreign host.
```



# Requesting a Range that Exceeds the Entity Size

```
$ telnet www.cs.odu.edu 80 | tee 7-3.out
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=0-2000
Connection: close
```

```
HTTP/1.1 206 Partial Content
Date: Mon, 20 Feb 2006 01:19:05 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 1936
Content-Range: bytes 0-1935/1936
Connection: close
Content-Type: text/html
```

```
<html>
<head>
<title>
[rest of html deleted]
```

# Conditional Request: If-Range

- Combine two requests into one:
  - “Informally, its meaning is as follows: if the representation is unchanged, send me the part(s) that I am requesting in Range; otherwise, send me the entire representation.”
  - Section 3.2
- Combine with “Range:” header
  - Server **MUST** ignore If-Range if the Range header missing
- Pass in ETag or date
  - If-Range: 1f4de2-790-43dd0cc3
  - If-Range: Sun, 29 Jan 2006 18:43:15 GMT

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=0-9
If-Range: Sun, 30 Jan 2006 18:43:15 GMT
Connection: close
```

```
Connection closed by foreign host.
HTTP/1.1 200 OK
Date: Mon, 20 Feb 2006 01:25:03 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 1936
Connection: close
Content-Type: text/html
```

```
<html>
<head>
<title>
[rest of html deleted]
```

# 416 Requested Range Not Satisfiable

```
$ telnet www.cs.odu.edu 80 | tee 7-6.out
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=2000-2999
Connection: close
```

```
HTTP/1.1 416 Requested Range Not Satisfiable
Date: Mon, 20 Feb 2006 01:52:07 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 0
Content-Range: bytes */1936
Connection: close
Content-Type: text/html
```

# Ignore Syntactically Invalid Range Headers (Apache 1.3.X)

- “The recipient of a byte-range-set that includes one or more syntactically invalid byte-range-spec values MUST ignore the header field that includes that byte-range-set.”
  - Section 14.35.1

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=aldksjfalstdjfl
Connection: close
```

```
HTTP/1.1 200 OK
Date: Mon, 20 Feb 2006 01:57:36 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 1936
Connection: close
Content-Type: text/html
```

```
<html>
<head>
<title>
[rest of html deleted]
```

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=aldksjfalsdjfl
Connection: close
```

```
HTTP/1.1 416 Requested Range Not Satisfiable
Date: Mon, 19 Feb 2007 17:07:24 GMT
Server: Apache/2.2.0
Connection: close
Transfer-Encoding: chunked
Content-Type: text/html; charset=iso-8859-1
```

```
180
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>416 Requested Range Not Satisfiable</title>
</head><body>
<h1>Requested Range Not Satisfiable</h1>
<p>None of the range-specifier values in the Range
request-header field overlap the current extent
of the selected resource.</p>
<hr>
<address>Apache/2.2.0 Server at www.cs.odu.edu Port 80</address>
</body></html>
```

```
0
Connection closed by foreign host.
```

# Ignore Syntactically Invalid Range Headers ??? (Apache 2.2.X)

## RFC 7233, section 3.1:

An origin server **MUST** ignore a Range header field that contains a range unit it does not understand. A proxy **MAY** discard a Range header field that contains a range unit it does not understand.

```
$ telnet www.cs.vt.edu 80
Trying 198.82.184.178...
Connected to cs.vt.edu.
Escape character is '^]'.
GET /index.html HTTP/1.1
Host: www.cs.vt.edu
Range: bytes=ladksjlasdjf
Connection: close
```

```
HTTP/1.1 200 OK
Date: Wed, 10 Oct 2018 18:40:49 GMT
Server: Apache
Strict-Transport-Security: max-age=31536000
X-Content-Type-Options: nosniff
Last-Modified: Sun, 07 Oct 2018 05:01:38 GMT
ETag: "a03b-5779c6646f0ea"
Accept-Ranges: bytes
Content-Length: 41019
Vary: Accept-Encoding
Cache-Control: max-age=600
Expires: Wed, 10 Oct 2018 18:50:49 GMT
X-RouteInfo: cmsw-prod-03
Content-Type: text/html; charset=UTF-8
Connection: close
```

```
<!DOCTYPE html>
<html lang="en">
<head>
[deletia]
```

# Unknown Apache version, 200 response

# www.cs.odu.edu no longer supports range requests

```
$ openssl s_client -connect www.cs.odu.edu:443
```

```
CONNECTED(00000003)
```

```
[ssl deletia]
```

```
HEAD /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
```

```
Host: www.cs.odu.edu
```

```
Connection: close
```

```
HTTP/1.1 200 OK
```

```
Server: nginx
```

```
Date: Wed, 10 Oct 2018 18:36:58 GMT
```

```
Content-Type: text/html
```

```
Connection: close
```

```
Vary: Accept-Encoding
```

```
Front-End-Https: on
```

```
closed
```

**Note absence of:** Accept-Range: bytes

# How not to respond

```
$ openssl s_client -connect www.cs.odu.edu:443
CONNECTED(00000003)
[ssl deletia]
GET /~mln/teaching/cs595-s06/a1-test/2/index.html HTTP/1.1
Host: www.cs.odu.edu
Range: bytes=0-9
RENEGOTIATING
2196:error:1409E0E5:SSL routines:SSL3_WRITE_BYTES:ssl handshake
failure:/SourceCache/OpenSSL098/OpenSSL098-52.8.4/src/ssl/s3_pkt.c:566:
```

Can we blame nginx for this too?  
I'm not sure; this isn't even an HTTP event



# If you can't 206, then 200 is more polite

```
$ curl -I -H "Range: bytes=0-9"  
http://2.bp.blogspot.com/-a4fGsbaFbDI/W74ubMwUCyI/AAAAAAAAmHY/gD2OJak16NY7njw  
P3teOMWNI91-SmNrgCLcBGAs/s1600/plots.png  
HTTP/1.1 200 OK  
Access-Control-Expose-Headers: Content-Length  
ETag: "v9877"  
Expires: Thu, 11 Oct 2018 18:44:23 GMT  
Content-Disposition: inline;filename="plots.png"  
Content-Type: image/png  
Vary: Origin  
Access-Control-Allow-Origin: *  
Timing-Allow-Origin: *  
X-Content-Type-Options: nosniff  
Date: Wed, 10 Oct 2018 18:44:23 GMT  
Server: fife  
Content-Length: 39828  
X-XSS-Protection: 1; mode=block  
Cache-Control: public, max-age=86400, no-transform  
Age: 361
```

# Even if with bad syntax

```
$ curl -I -H "Range: bytes=ladsjlasdjf"  
http://2.bp.blogspot.com/-a4fGsbafBdI/W74ubMwUCyI/AAAAAAAAmHY/gD2OJak16NY7nj  
wP3teOMWniT9l-SmNrgCLcBGAs/s1600/plots.png  
HTTP/1.1 200 OK  
Access-Control-Expose-Headers: Content-Length  
ETag: "v9877"  
Expires: Thu, 11 Oct 2018 18:44:23 GMT  
Content-Disposition: inline;filename="plots.png"  
Content-Type: image/png  
Vary: Origin  
Access-Control-Allow-Origin: *  
Timing-Allow-Origin: *  
X-Content-Type-Options: nosniff  
Date: Wed, 10 Oct 2018 18:44:23 GMT  
Server: fife  
Content-Length: 39828  
X-XSS-Protection: 1; mode=block  
Cache-Control: public, max-age=86400, no-transform  
Age: 258
```

# Ranges other than bytes?

## HTTP Range Unit Registry

### Registration Procedure(s)

IETF Review

### Reference

[\[RFC7233\]](#)

### Available Formats



CSV

Name	Description	Reference	Notes
bytes	a range of octets	<a href="#">[RFC7233]</a>	Section 2.1
none	reserved as keyword, indicating no ranges are supported	<a href="#">[RFC7233]</a>	Section 2.3

RFC 7233 introduces a registry: <https://www.iana.org/assignments/http-parameters/http-parameters.xhtml>

Currently only “bytes” and “none”, but applications could define “pages”, “seconds”, “versions”, or other custom offsets.