UNIVERSITY OF
Southampton

UNIVERSITY OF Southampton

Web Caching

COMP3220 Web Infrastructure

Dr Heather Packer – hp3@ecs.soton.ac.uk

Caching

- Caching stores the result of an operation so that future operations return faster
  - Computation is slow
  - Computation will run multiple times
  - When the output is the same for a particular input

Web Caching

- The temporary storage (caching) of frequently accessed data for rapid access
- Typically caches store “static assets”
  - HTML pages
  - Images
  - Stylesheets, Javascript
- Caches can be located at various points in a network
  - Reduces access time/latency for clients
  - Reduces bandwidth usage across slower links
  - Reduces load on a server

Browser Cache

- Browsers maintain a small cache
- Stored locally
- Cache for a single user or application
- Browser sets a caching policy, deciding what data to cache
  - User specific content
  - Expensive content

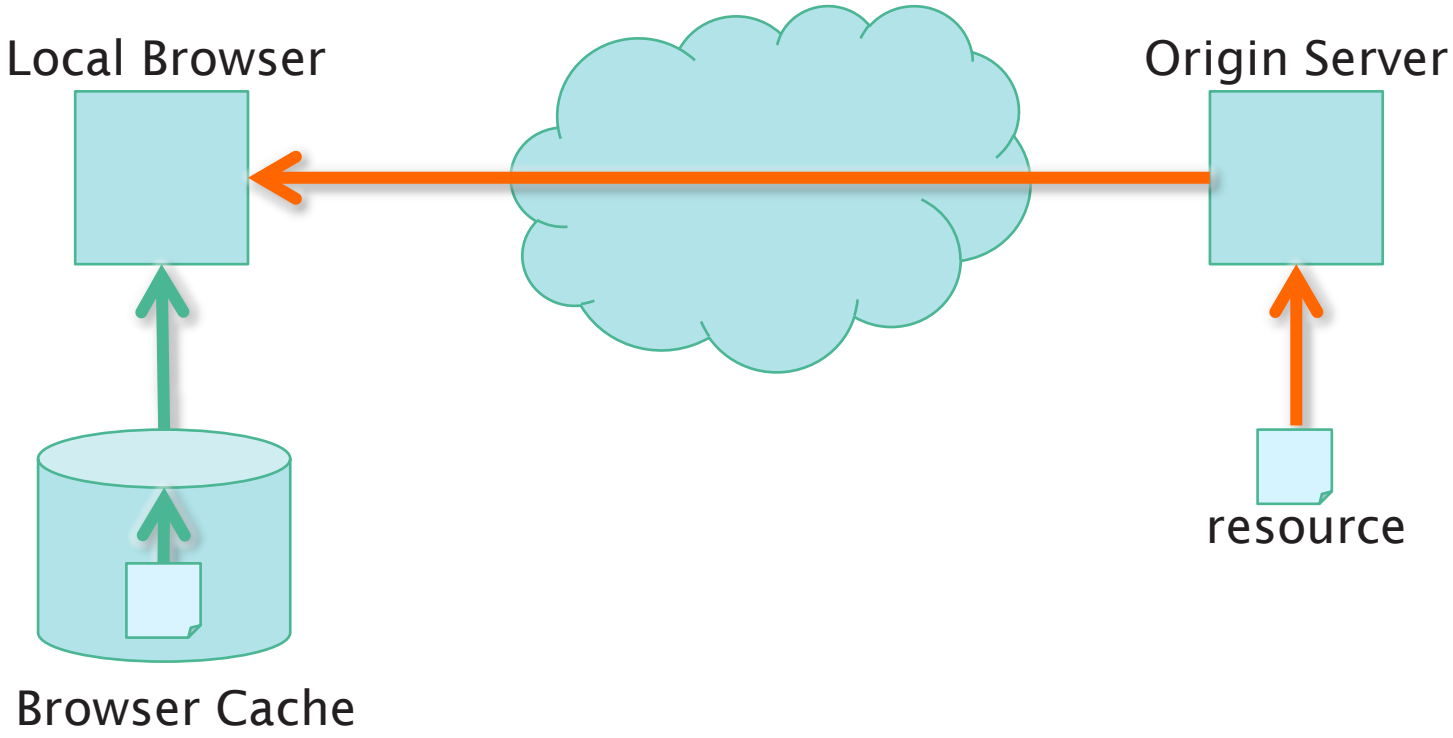

What can be Cached?

- Cache friendly
  - Logos and brand images
  - Style sheets
  - Javascript files, site and library
  - Fonts
  - Downloadable content
  - Media files
- Be careful caching:
  - Data
  - HTML pages
  - Frequently modified Javascript and CSS
  - Content requested with authentication cookies
- Never cache
  - Sensitive data
  - User-specific data that frequently changes

Caching using Conditional requests

- Last-Modified: Tue 17 Nov 2020 08:00:20 GMT
  - GET Requests: If-Modified-Since:
- Etag: “0123456789ABCDEF…”
  - GET Request: If-None-Match:
- HTTP Status code: 304 Not Modified
- PUT/PATCH Methods
  - If-Unmodified-Since:
  - If-Match:
  - 412 Precondition Failed

Controlling Caches with HTTP: Last-Modified Header

GET

GET / HTTP/1.1

Host: comp3220.ecs.soton.ac.uk

Accept: */*

HTTP/1.1 200 OK

Date: Wed 18 Nov 2020 17:43:20 GMT

Connection: keep-alive

Content-Type: text/html; charset=UTF-8

Content-Length: 4003

Last-Modified: Tue 17 Nov 2020 08:00:20 GMT

Conditional GET

GET / HTTP/1.1

Host: comp3220.ecs.soton.ac.uk

Accept: */*

If-Modified-Since: Tue 17 Nov 2020 08:00:20 GMT

HTTP/1.1 304 Not Modified

Date: Wed 15 Nov 2017 07:55:10 GMT

Connection: keep-alive

Last-Modified: Tue 17 Nov 2020 08:00:20 GMT

Caching Headers

- HTTP Response header Cache-Control:
- Flags
  - no-store
  - no-cache
  - max-age=
  - must-revalidate
  - public
  - Private
- Use of Cache-Control headers to be determined by web site architect/designer

HTTP with Cache-Control Header

GET

GET / HTTP/1.1

Host: comp3220.ecs.soton.ac.uk

Accept: */*

HTTP/1.1 200 OK

Date: Wed 18 Nov 2017 17:43:20 GMT

Connection: keep-alive

Content-Type: text/html; charset=UTF-8

Content-Length: 4003

Last-Modified: Tue 17 Nov 2020 08:00:20 GMT

Cache-Control: max-age=86400

GET

* No request sent *

Different Web Caching Solutions

- Caches can be located at various points in a network
  - Browser Cache
    - Embedded in the browser
  - Proxy Cache
  - Reverse Proxy Cache

Proxy Cache

- Cache located close to the clients (hosted by University or Internet Service Provider)
  - Decrease bandwidth usage
  - Decreases network latency
- Scale provides the main advantage: many users within the ISP may all be asking for the same web pages
- ISPs use this approach to decrease bandwidth across their networks

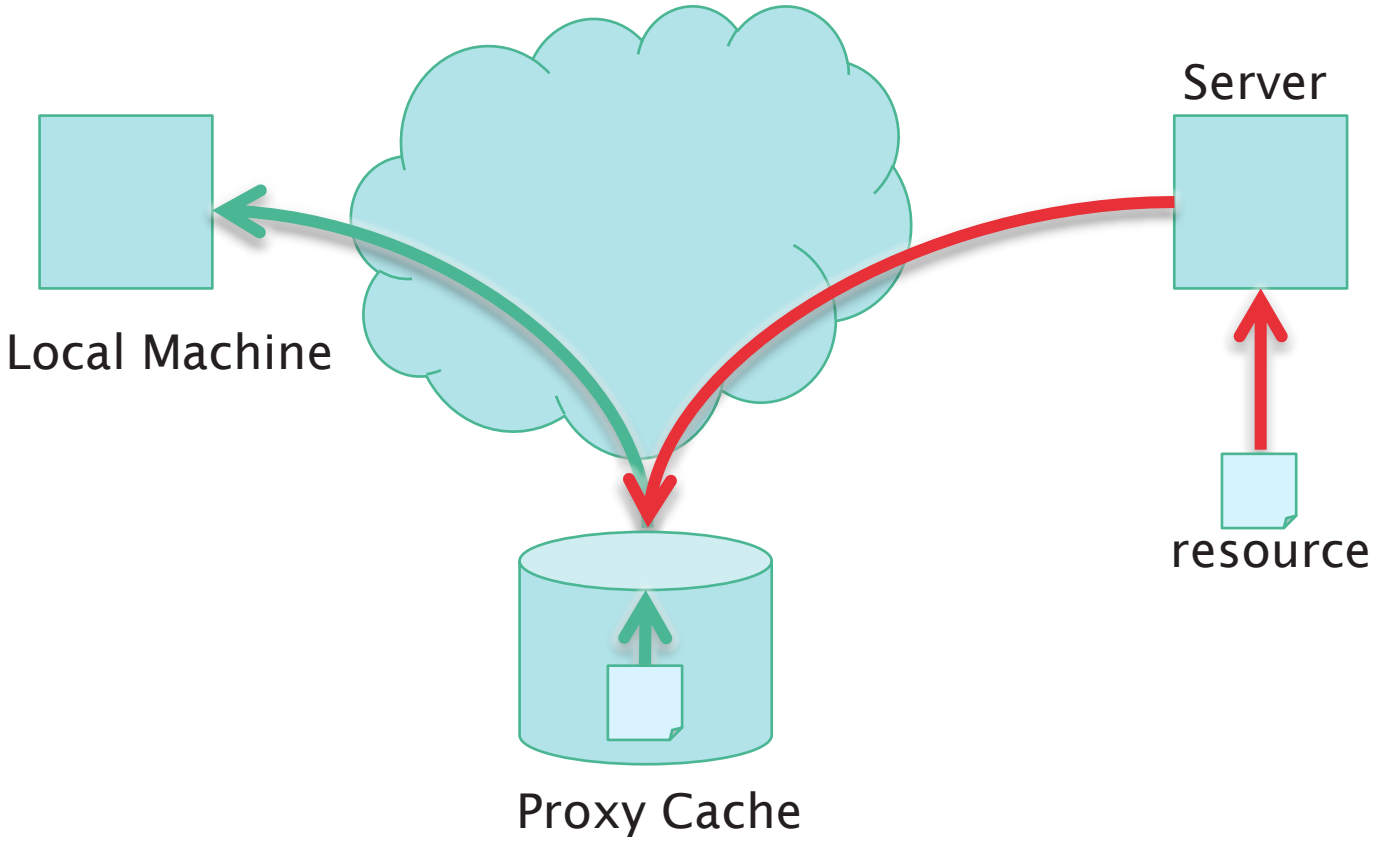

Reverse Proxy Cache

- Cache proxy located closer to the origin web server
- Usually deployed by a Web host
- Decreases load on the Web service (e.g. database)
- Several reverse proxy caches implemented together can for a Content Delivery Network

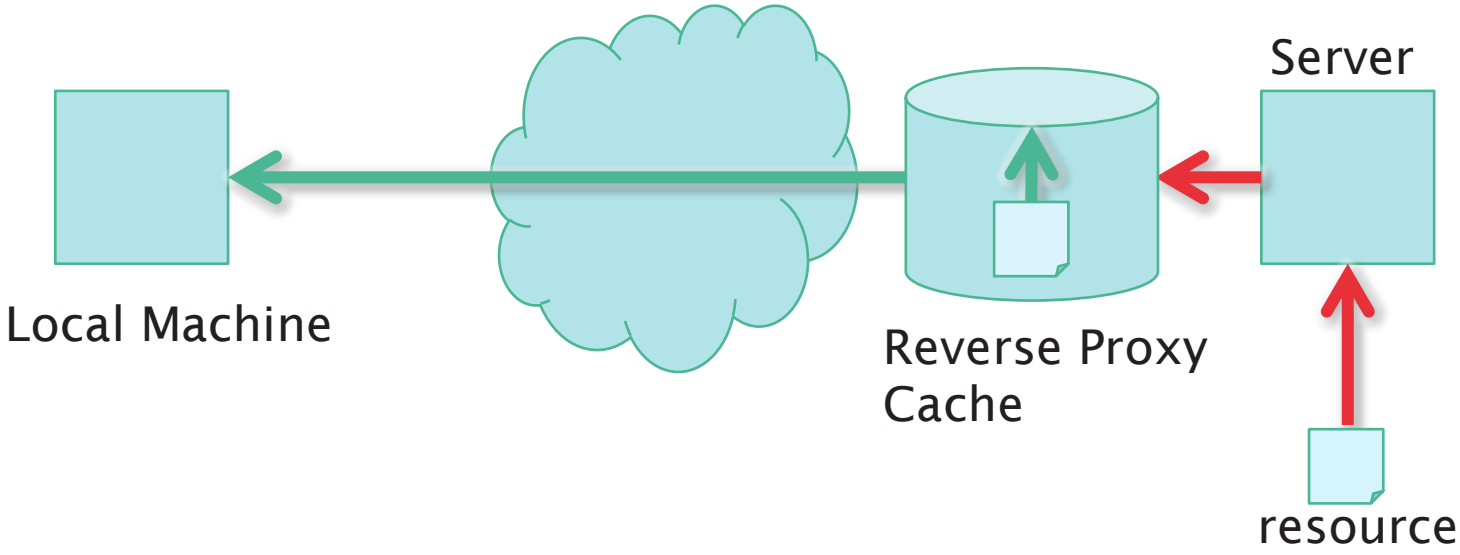

Learning Outcomes

- Web caching reduces latency and bandwidth
- Identify HTTP headers and cache-control headers used for caching