



Web portal caching

1 Contents

2	Web portal caching	2
3	Introduction	2
4	How HTTP caching works	2
5	Caching in WebKit	3
6	Client/Server implementation strategies	4
7	Application cache	4
8	Custom HTTP application caching server running locally	5
9	Separatedly-maintained locally accessible copy of the portal con-	
10	tents	5

11 Web portal caching

12 Introduction

13 The purpose of this document is to evaluate the available strategies to implement
14 a custom, single-purpose browser restricted to a single portal website that hosts
15 several HTML/JS applications.

16 The portal and the visited applications should be available even if no Internet
17 connection is available.

18 If a connection to the Internet is available, the locally-stored contents should be
19 refreshed.

20 Locally-stored copies should be used to speed up loading even when the connec-
21 tion to the Internet is available.

22 The portal and the applications store all their runtime data using the `local-`
23 `storage`¹ or `IndexedDB`² mechanisms and how that is synchronized is out of the
24 scope of this document, which instead focuses on how to manage static assets.

25 How HTTP caching works

26 Caching is a very important and complex feature in modern web engines to
27 improve page load time and reduce bandwidth consumption. [RFC7234](#)³ defines
28 the mechanisms that control caching in the HTTP protocol regardless of its
29 transport or serialization, which means that the same mechanisms apply to
30 HTTPS and HTTP2 in the same way.

31 HTTP has provisions for several use cases:

- 32 • preventing highly dynamic resources from being cached
- 33 • letting clients know for how long is acceptable to use cached data

¹<https://html.spec.whatwg.org/multipage/webstorage.html>

²<https://www.w3.org/TR/IndexedDB/>

³<https://tools.ietf.org/html/rfc7234>

- 34 • optimizing validation of cached entries to skip the download of the body
35 if the copy on the client still matches the one on the server
- 36 • informing clients about resources that can be safely used even if stale when
37 no connection is available and which ones must return an error

38 Caching is generally available only for the `GET` method and is controlled by the
39 server for every single HTTP resource by adding the `Cache-control` header to its
40 responses: this instruct the client (the web engine) on the ways it can store the
41 retrieved contents and re-use them to skip the download on subsequent requests.

42 One of the most important uses of the `Cache-control` header is to disable any kind
43 of caching on highly dynamic generated resources, by specifying the `no-store`
44 value.

45 The `public` and `private` directives instruct clients that the resource can be stored
46 in the local cache (`public` also allows for caching in intermediate proxy servers,
47 a feature which is progressively getting obsolete as it conflicts with the confidential-
48 ity requirements of HTTPS/TLS).

49 The `Expire` header and the `max-age` directive let the server instruct the client for
50 how long it can consider the cached resource valid. The client can completely
51 skip any network access as long as the cached resource is “fresh”, otherwise it
52 has to validate it against the server, but this does not mean that a complete
53 re-download is always needed: using conditional requests, that is using the `If-`
54 `Modified-Since` or `If-None-Match` headers to pass the values of the `Last-Modified`
55 or `ETag` headers from the previous request, the download of the body is skipped
56 if the values match and only headers will be transferred with a `304 Not Modified`
57 response.

58 The HTML5 specification recently introduced the concept of [application cache](#)⁴
59 which caters for an additional, higher-level use case: pro-actively downloading
60 all the resources needed by an HTML application for offline usage.

61 This works by adding a `manifest` attribute to the `<html>` element of the main
62 application page, and from there indicate the URL of a specially formatted
63 resource that lists all the URLs the client needs to pro-actively retrieve in order
64 to be able to run the application correctly when offline. The caching model
65 used by this specification is somewhat less refined than the one used by the
66 HTTP specification and for this reason it needs some special attention on how
67 to ensure that the application is properly refreshed when changes are made on
68 the server.

69 The more complex and powerful [Service Workers](#)⁵ specification is meant to re-
70 place this, but it is not supported yet by all modern browsers (works in Firefox
71 and Chrome, WebKit and Edge don't support it yet). The specification has
72 been stable for more than a year, despite not being finalized yet. The WebKit

⁴<https://html.spec.whatwg.org/multipage/browsers.html#offline>

⁵<https://www.w3.org/TR/service-workers/>

73 team has not yet shown a clear interest in implementing it, which may be the
74 reason why the specification is still in the current status.

75 **Caching in WebKit**

76 WebKit currently has several caches:

- 77 • a non-persistent, in-memory cache of rendered pages which is set to 2
78 pages if the total RAM is bigger or equal to 512MB
- 79 • a non-persistent, in-memory [decoded/parsed object cache](#)⁶, set to 128MB
80 if the total RAM is bigger or equal to 2GB and progressively lowered as
81 the amount of total RAM decreases
- 82 • a persistent, on-disk [resources cache](#)⁷ of 500MB if there are more than
83 16GB free on the disk, progressively scaling down to 50MB if less than
84 1GB is available.

85 Those sizes are computed automatically but they can be customized to fit any
86 requirements.

87 When a new resource needs to be cached WebKit makes sure that the upper
88 bound is respected and frees older cache entries in a LRU pattern to make
89 enough room to accomodate the resource which is about to be downloaded.

90 Downloaded contents to be stored in the on-disk URL cache are directly saved
91 in the filesystem, using the normal buffering that the kernel does for every
92 application to improve performance and minimize eMMC wear. This is further
93 minimized by the fact that only contents marked for caching by the server using
94 the appropriate HTTP headers will be cached: highly dynamic contents like
95 news tickers won't be marked as cacheable so they won't impact the eMMC at
96 all.

97 The application cache is handled separately and it is unlimited by default, but
98 this is a setting that can be changed. All the resources are stored in a SQLite
99 database as data blobs, except for audio and video resources where the only the
100 metadata is stored in the database and the contents are stored separately.

101 To use the application cache effectively in WebKitGTK+ some implementation
102 work would be required to limit the maximum size as the WebKit core hooks
103 are currently not used by the WebKitGTK+ port, and the WebKit core itself
104 does not currently provide any expiration policy for the cached contents.

105 **Client/Server implementation strategies**

106 Multiple strategies can be used to implement the previously defined system and
107 affect the design of the client and of the contents offered by the portal server.

⁶<https://trac.webkit.org/browser/trunk/Source/WebKit2/Shared/CacheModel.cpp#L83>

⁷<https://trac.webkit.org/browser/trunk/Source/WebKit2/Shared/CacheModel.cpp#L158>

108 **Application cache**

109 The main HTML page of the portal links to an appcache manifest that instruct
110 the browser to pro-actively fetch all the needed resources.

111 All subsequent accesses to the portal will be served from the cached copy, re-
112 gardless of the availability of an Internet connection.

113 If the portal is accessed when an Internet connection is available, the browser will
114 retrieve the appcache manifest from the server in the background and check for
115 modifications: if a new version is detected the portal resources will be refreshed
116 in the background and will be used for subsequent accesses to the portal.

117 Each application will have its own appcache manifest, so it will be locally cached
118 after the first visit.

119 To ensure that the portal is available on first-boot even if no Internet connection
120 is available, during the process of generating the system image the browser will
121 be launched using a special mode that will cause it to connect to the portal, pop-
122 ulate the application cache and exit as soon as the `ApplicationCache::updateready`
123 event is fired. An ad-hoc program using WebKit may be used instead of adding
124 a special mode to the browser.

125 This is the simplest and most portable approach on the client side, as all the
126 caching logic is provided by the portal server using standard W3C mechanisms.

127 **Custom HTTP application caching server running locally**

128 Alternatively, the browser can be instructed to connect to a custom HTTP proxy
129 server running locally instead of directly to the portal server.

130 Since TLS authentication cannot work appropriately through proxy servers, it is
131 taken care by the proxy server itself, with the browser talking to the local proxy
132 over unencrypted HTTP and the proxy converting HTTP requests to HTTPS.

133 This means that unencrypted communications will only happen locally between
134 trusted components, while all the network traffic will be encrypted. Just like for
135 any other HTTP error, the proxy can return error pages to the browser in case
136 of TLS error (for instance, if the server certificate is expired) or return cached
137 contents if available.

138 The custom proxy is then responsible for connecting to the portal server and
139 retrieving updated contents from there, locally caching it with any kind of expiry
140 and refresh policy desired, and processing cached resources when needed, for
141 instance by rewriting links from HTTPS to HTTP.

142 The browser needs to be configured to reduce its own caching to a minimum,
143 since the smart proxy already does it.

144 During the manufacturing process the proxy cache will be preloaded with the
145 resources hosted by the portal server.

146 This is the most flexible approach.

147 **Separately-maintained locally accessible copy of the portal contents**

148 Instead of having a locally running custom HTTP caching proxy, the portal
149 contents are stored as plain files on the system. The browser will contain custom
150 logic to load the local HTML file instead of the portal URL when no Internet
151 connection is available.

152 A separate process will periodically compare the locally-stored HTML file and
153 resources against the portal server and refresh the local copy.

154 This is the least flexible choice, and the locally stored copies cannot be used as
155 cache to speed up rendering when the connection to the Internet is available.