

Redmine - Patch #37600

Cache attachments and thumbnails in web browsers

2022-08-24 05:15 - Go MAEDA

Status:	New	Start date:	
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	Performance	Estimated time:	0.00 hour
Target version:	Candidate for next major release		
Description			
<p>The attached patch adds "max-age" value to "Cache-Control" field in response headers of AttachmentController#download and AttachmentController#thumbnail in order to reduce the number of requests to a Redmine server and speed up page loading especially when an issue or a Wiki page has dozens of inline images and thumbnails.</p>			
Before:			
<pre>Cache-Control: private</pre>			
After:			
<pre>Cache-Control: max-age=86400, private</pre>			
<p>In my observation, some users attach about 100 or more images to an issue. Even if they access the issue again shortly after it was accessed, their browser sends many GET requests to AttachmentController#download and AttachmentController#thumbnail to revalidate their local cache. Each of those requests that receive "304 Not Modified" consumes a few tens of milliseconds of server time.</p>			
<p>By adding "max-age: 86400" to "Cache-Control" header field, web browsers will cache images and thumbnails and will not make revalidate requests for 24 hours, which improves page loading speed and reduces server load.</p>			
<p>In Redmine, the content of attachments and thumbnails with the same URL is immutable. So there is no need to frequently revalidate the local cache and it is safe to keep the cache for some time.</p>			
<p>It is safe to keep caches of attachments and thumbnails for some time because the content of attachments and thumbnails with the same URL is immutable in Redmine. And there is no need to frequently revalidate the local cache.</p>			

History

#1 - 2022-08-25 08:29 - Go MAEDA

- Target version set to Candidate for next major release

#2 - 2022-08-25 10:44 - Holger Just

lirc we generally avoid to allow explicit caching for anything because different user may have different access permissions and those could change quickly.

Thus, when allowing caching, we should make sure that the contents of the cache is not leaked to different users which may potentially happen at different places:

- public / centralized caches (such as enterprise proxy servers or CDNs) - if we allow caching here, the cache key must be tied to the user id in the session. As far as I'm aware, this is not possible right now with "public" information just from the single request. As such, we can not allow those caches at all.
- local browser-based caches - Here, we *may* potentially allow some kind of caching. If I understand the Cache-Control rules correctly, this is what you aim to enable here.

We should still investigate if this change allows for leaking the cache contents to different users e.g. after a logout or when multiple users are using a single browser. I'm not entirely sure, if this scenario is something we **have to** protect against, but I'd like to discuss the potential ramifications of this for an informed decision.

Another different (and possibly more common) scenario might be when users first access a Redmine as Anonymous and later login in, which would cause them to see different attachments.

Generally, if we decide that it is possible to allow private caches for some data, I think it would indeed improve performance for attachments as they are in fact immutable after creation.

Files

cache-attachments-and-thumbnails.patch	3.05 KB	2022-08-24	Go MAEDA
--	---------	------------	----------