

Cognitive Interaction Toolkit - Feature #412

multiple cloning of gits - an idea for load reduction for servers + slow connections

2015-06-10 09:44 - Simon Schulz

Status:	Resolved	Start date:	2015-06-10
Priority:	High	Due date:	
Assignee:	Jan Moringen	% Done:	50%
Category:	build-generator	Estimated time:	0.00 hour
Target version:	Toolkit 0.2	Spent time:	0.00 hour
Description			
<p>The build generator and jenkins does multiple checkouts from git repositories. This is annoying and produces high load on the repository servers. It also makes it impossible to be used on slow DSL connections at home (waiting hours for a build...).</p> <p>Another thing is that we have problems with building the distribution at the lsp-csra due to the big size and the amounts of checkouts (high load on repo servers).</p> <p>What about implenting the following caching behaviour ?</p> <p>1) the buildgenerator:</p> <pre>if (does not exist \$CACHEDIR/\$projname){ clone repo to \$CACHEDIR/\$projname }else{ cd \$CACHEDIR/\$projname && git pull }</pre> <p>2) the buildgen enters \$CACHEDIR/\$projname as git url in the jenkins project configuration (that works, this is how i use jenkins at home)</p> <p>It might be necessary to add a git pull to the project files, but i think jenkins does that automatically.</p> <p>By implementin this we would reduce 3x checkout to 1x or even no full checkout at all. This would vastly reduce the load on our repository/network infrastructure and would also reduce the total build times!</p> <p>I would really prefer using more diskpace due to the caching instead of network bandwidth usage.</p> <p>That should not be to hard to implement. I do this manually when i build large distributions at home.</p>			
Related issues:			
related to Cognitive Interaction Toolkit - Feature # 338: using the toolkit i...		Closed	2014-11-22

History

#1 - 2015-06-10 11:53 - Robert Haschke

Nice idea. I do it the same way on my laptop from home. However, this doesn't work anymore if the build-generator and jenkins are on different machines / cannot access the cached repository dirs the same. However, build-generator could do the caching anyway.

Additionally, projects repos should be designed modular, i.e. not required to checkout a monolithic repo if their parts can be (and are) build separately...

#2 - 2015-07-23 15:44 - Patrick Holthaus

That could actually save a lot of traffic (also for the redmine itself!)

#3 - 2015-08-24 21:12 - Florian Lier

- Target version set to Toolkit 0.2

bump

#4 - 2015-10-12 16:57 - Jan Moringen

- Status changed from *New* to *In Progress*

- % Done changed from 0 to 50

build-generator side of this strategy has been prototypically implemented and is currently undergoing testing.

#5 - 2015-10-12 17:02 - Jan Moringen

- related to Feature #338: using the toolkit installer on "slow" (≤ 10 MBit/s) connections - tar snapshots instead of git? added

#6 - 2016-04-08 15:11 - Jan Moringen

- Status changed from *In Progress* to *Resolved*

Generator-side caching with incremental updates is implemented.

Not wiping out workspaces can be used to make the Jenkins-side incremental.