

ChiliProject - Task # 97: Setting up CI Server

Status:	Closed	Priority:	Normal
Author:	Wieland Lindenthal	Category:	ChiliProject - Organization
Created:	2011-01-21	Assignee:	
Updated:	2011-09-04	Due date:	
Remote issue URL:			
Description: Let's bundle all requirement for a CI server.			
Here are some first basic thoughts:			
* Coverage of all branches			
** Nice to have: automatic creation of new build tasks for newly pushed branches.			
* Run tests for new commits			
* Feeds to subscribe to			
* Integration of test results (feeds?) into ChiliProject website.			
I think it would be nice to have this setup till the first release, as we have to be stable after it latest.			
@Eric: You had some ideas about reverting every commit that made the test fail, right?			

Associated revisions

2011-08-21 12:51 pm - Felix SchÄ¼fer

Merge pull request #97 from finnlabs/pulls/577/invalid-watcher-user-fix

Invalid watcher user error when adding an invalid user as watcher #577

History

2011-01-21 07:56 pm - Eric Davis

> * Coverage of all branches

Lets wait on this until we have our git workflow figured out. We still aren't sure about how to handle long lived branches under active development yet. At least the main development and stable branches should be tested.

> * Run tests for new commits

Yes and make sure to purge and load all of the testing SCMs so they are tested also.

> * Integration of test results (feeds?) into ChiliProject website.

Maybe later, we should think about how to handle this. I know there is a Hudson plugin but I want to do a code review of everything that is installed on the server first.

Other ideas (probably for later):

* Generate rcov and metric_fu metrics on the code and post to a public site. Will help to find poor quality code and refactorings

* Generate rdoc/yardoc files and publish for developer/API documentation. I'm still thinking about how to do this. (Sidebar: I'd love to have all of the REST API documented in the code and use Yard to extract it automatically into an API documentation set. Would help keeps code/docs synced.)

* Log build success/failures to IRC

I don't think automatically reverting every commit is a good idea. I'd rather have a person review it and make sure it really is an error before reverting.

2011-01-31 02:36 pm - Gregor Schmidt

I would like to help out on that.

I started setting up a dedicated machine to run Hudson/Jenkins. My plan is to support basically the same set of configurations, that are currently running on ci.finn.de, i.e.

- * Ruby 1.8.7, 1.9.2 and JRuby

- * PostgreSQL, MySQL, SQLite

Just like Eric proposed, I would target ChiliProject's master at first.

I am planning to improve the setup (compared to ci.finn.de), such that

- * test repositories used for the scm adapter tests are handled automatically

- * at least two builds may run at the same time

I hope to have this finished and running on a public site by the end of next week. I hope, that this installation would become ChiliProject's official CI.

Nevertheless, please continue collecting ideas and requirements.

2011-02-01 12:37 pm - Gregor Schmidt

Is there a need/interest to run the LDAP tests in ci?

2011-02-01 12:46 pm - Holger Just

It would be nice to have that, but not an immediate requirement for now. </Holger's opinion>

I could setup a simple java-based OpenDS or an OpenLDAP server and feed it the test-ldifs. In the long term, the setup of an LDAP instance should be part of the running test suite, but that would require changes of the tests and some scripting.

2011-02-01 10:49 pm - Felix Schäfer

- Status set to Open

2011-02-08 05:05 pm - Gregor Schmidt

I'm making good progress on setting everything up. I'm planning to have the following jobs:

1. chiliproject_master

Listening to changes on the master branch, building with 1.8.6 and 1.8.7 with mysql, postgresql and sqlite3

2. chiliproject_master_experimental

Listening to changes on the master branch, building with 1.9.2 and jruby 1.5.6 with mysql, postgresql and sqlite3, as well as mysql2 gem for ruby 1.8.6, 1.8.7, 1.9.2.

Separating them into two jobs clearly communicates, which configurations are considered 'supported'.

Later, I guess, there should be similar jobs for the unstable and stable branches.

I would welcome any feedback on that setup.

Additionally I ran into the following questions:

- * Which IRC-Bot should be used for notifications?
- * What should be reported?
- ** Every ci run and its status for the non-experimental jobs?
- ** Every broken run and once, when it is fixed?
- ** Only changes to the build status (once when getting red and once when coming back to green)?
- * Is it true, that the mysql2 gem is considered experimental in the context of ChiliProject or should it be supported? - It currently fails for one test.
- * Did I miss something important?

2011-02-08 05:26 pm - Holger Just

Gregor Schmidt wrote:

> I'm making good progress on setting everything up.

Awesome!

> # chiliproject_master
> # chiliproject_master_experimental

I think this distinction makes sense. We later need to synchronize the documentation of actually supported configurations with the builds.

> Later, I guess, there should be similar jobs for the unstable and stable branches.

@master@ is stable (well kind of). Don't know if makes sense to have jobs for the @release@ branch. I guess they don't hurt as these shouldn't run often. And when they run, they provide an additional safety net for releases.

> * Which IRC-Bot should be used for notifications?

On [\[\[IRC|chat.chiliproject.org\]\]](http://irc.chat.chiliproject.org) runs an rbot named @pepperbot@. We could use that one or have an additional one on the ci server. If in doubt, I would try to use pepperbot to reduce the number of bots on the channel.

> * What should be reported?

During the old finnbot times on #redmine, we reported every build status which was considered spammy behavior by many people. So I would vote for your second option of reporting every broken build and once when it's fixed. That way we have a sufficiently high visibility of errors without spamming the channel too much.

The unstable branch (and possibly feature branches) are expected to break a bit more often, so I guess here it is sufficient to report only changes of the build status.

> * Is it true, that the mysql2 gem is considered experimental in the context of ChiliProject or should it be supported? - It currently fails for one test.

I personally would consider mysql2 even more supported than mysql (see the [\[\[installation\]\]](#) guide.) But in the end, it should work with both variants.

> * Did I miss something important?

We might have some longer running feature branches in the future. E.g. for thing like rails 3 migration, permission overhaul, ... These could live inside the main repository as there are probably many of the core devs working on them. These branches should be automatically build (the same as @unstable@.) I know Michael has build something to catch up on new branches automatically.

2011-02-08 05:35 pm - Gregor Schmidt

Holger Just wrote:

> > * Is it true, that the mysql2 gem is considered experimental in the context of ChiliProject or should it be supported? - It currently fails for one test.

>

> I personally would consider mysql2 even more supported than mysql (see the [[installation]] guide.) But in the end, it should work with both variants.

In that case, I guess, the mysql2 runs should become part of the cp_master job and not cp_master_experimental.

2011-02-08 09:58 pm - Eric Davis

Gregor Schmidt wrote:

> * What should be reported?

> ** Every ci run and its status for the non-experimental jobs?

> ** Every broken run and once, when it is fixed?

> ** Only changes to the build status (once when getting red and once when coming back to green)?

Failures are a must. I would actually say each failure should be reported, not just on changes. I'd be fine without successes since we have commit bots already.

Using this "commit stream" as an example

* Success

* Success

* Fail - notify (first failure)

* Fail - notify (failure)

* Success - notify (first success after a failure)

* Success

* Failure - notify (first failure)

* Success - notify (first success after a failure)

> * Is it true, that the mysql2 gem is considered experimental in the context of ChiliProject or should it be supported? - It currently fails for one test.

I haven't tried the mysql2 gem myself so whatever everyone else suggests.

> * Did I miss something important?

It might be good to add rbx later on if someone wants to help support it.

2011-02-17 12:59 am - Eric Davis

Gregor Schmidt:

How is the CI server coming? We all green yet?

2011-02-17 06:52 am - Gregor Schmidt

It's going good. Since merging the mysql2 adapter fix we are green for 1.8.6 and 1.8.7 for the pg, mysql, mysql2 and sqlite3-ruby gem. The merges from Redmine 1.1 also fixed all tests for 1.9.2. So currently only JRuby is not working. It's one general JRuby problem failing with all adapters, nothing else for MySQL but larger and smaller failures and errors for Postgres and SQLite. You may inspect the 1.8.x test results on the CI preview at <http://ci.nach-vorne.eu/>. Unfortunately that machine is not very powerful, so I had to deactivate parallel builds.

The hardware, that should run the new ci is nearly ready and it will hopefully be ready to go at the end of this week, i.e. Friday. That one should be able to run up to 4 job configurations in parallel.

2011-02-18 09:16 pm - Gregor Schmidt

Apologies, again. Unfortunately our hosting provider was kind of unresponsive lately. We're still waiting to get the system running. I or Holger will post an update, once we've got something new to tell.

Sorry,

Gregor

2011-02-24 11:39 pm - Eric Davis

- *Target version deleted (1.1.0 "Bell")*

Bumping this out of 1.1.0. It's important but not really part of the code to be released.

2011-02-28 08:27 am - Gregor Schmidt

- *Status changed from Open to In Progress*

As already mentioned on IRC a basic version is set up on <http://ci.chiliproject.org>. But there is still some major work to do.

It is still missing the combination of 1.8.6 with the mysql gem. I will also set up ci jobs for the newly created unstable and stable branches and the IRC integration.

Also I plan to write a wiki page covering -all- most of the things happening during a ci run and how it basically works. This should help to get the reasons things may behave differently on the ci.cp.o than on `_your machine_`.

Also, I will document all the versions of external components. -- Looks like this will become a rather large document.

Whenever something goes wrong on the ci or things behave strangely. Ping me on IRC or via mail.

Thank you for your patience.

2011-09-01 05:01 am - Felix Schäfer

What's the status of the docs here?

2011-09-04 09:34 pm - Holger Just

- *Status changed from In Progress to Closed*

The CI humming along smoothly. Thanks for your awesome work Gregor!

I moved the outstanding documentation part to a new issue in #605.