

Migrating GNOME to Git

(a human & technical perspective)

Frédéric Péters
<fpeters@gnome.org>

Rencontres Mondiales du Logiciel Libre – Nantes - 10 juillet 2009

Disclaimer

*“ As an aside, I think people that like git do so out of a kind of **software Stockholm syndrome**: you have to learn so much about esoterics like refs, the object database, the index, etc. that you end up feeling empathy for git's idiosyncracies. ”*

Andy Wingo

The Past

original gvfs development in git, in 2006

git as an option discussed when we were doing the subversion migration (Dec 2006)

git mirror (June 2008)

happenings in various teams:

- GTK+ team considered moving to git

- NetworkManager moved to freedesktop

- Empathy developers used their git server and pushed to gnome.org with git-svn

Why Git is better than X?

Cheap Local Branching

Everything is Local

Git is Fast

Git is Small

The Staging Area

Distributed

Any Workflow

GitHub

Easy to Learn

DVCS Survey (1/3)

(December 11th)

“ I do think the survey result will be interpreted as a decision "we must switch to \$dvcs", "we will keep svn". Even if "The survey results will be informational" is mentioned. At least I kind of fear that. ”

(moi)

10h50 : Méthodes de vote : comment consulter les développeurs d'un projet sans fausser le résultat avant de poser la question

DVCS Survey (2/3)

< marnanel> it does look like **git is winning**, doesn't it?

< jclinton> the think the margin by which **git is winning** cannot possibly be affected by a later vote

* mclasen has a really hard time following the mental gymnastics that lead from '**git is a clear winner**' to '**bzr on the server**'

DVCS Survey (3/3)

< zith> so git is the conclusion of the survey and the following edrama?

git transition exploratory team

(January 6th)

From: Owen Taylor

There's a lot of not-very-productive-discussion and angst going around now about what we are doing for VCS.

A group of us who believe that a straight git transition probably makes sense (myself, Federico, Behdad, Elijah, Kristan Hoegsberg) have started looking into the concrete steps necessary to get there.

Preview Repository (1/3)

(January 14th)

From: Owen Taylor

With heroic efforts on the conversion scripts, Kristian has finished an initial run at converting all `svn.gnome.org` repositories and putting them on `git.gnome.org`.

Preview Repository (2/3)

People are invited to test it, to review things

Many will not

Those who do will just have a quick look

Nobody will try building the whole stack from
Git

Preview Repository (3/3)

Some problems were anticipated

- Subversion Externals

- Commit Hooks

Some were not

- Empty directories

But how useful is it to identify potential problems when concerned persons are not informed?

Project Decision (1/2)

(March 19th)

From: Lucas Rocha

The GNOME Release Team would like to announce that git will be the new Version Control System (VCS) for GNOME. In our opinion, the decision reflects the opinion of the majority of our active contributors.

[...]

Project Decision (2/2)

[...]

The new `git.gnome.org` server is now in good shape and contains a functional preview of all GNOME git repositories. The official migration of all our Subversion repositories to git will take place just after the 2.26.1 release, on April 16.

Migration Day-1

From: Xavier Claessens

Date: Wed, 15 Apr 2009 09:52:39 +0200

Subject: Git migration Thursday?

Is it still true that the SVN->GIT migration will take place tomorrow? I guess the SVN will be set read-only during the import, at what time will it happen? Do we have an estimation of how long it takes?

Migration Day (1/12)

16:02 -!- krh has changed the topic to
SVN read-only, git migration in
progress today - 2.26.1 released, hard
code freeze ends, other freezes still
apply for 2.26.x

Migration Day (2/12)

18:09 < owen> We could certainly have alternate mechanisms to figure out a description for doap-less modules, I'm just trying to get *a* mechanism in place so I can answer people who ask about "Unnamed repository; edit this file to name it for gitweb."

Migration Day (3/12)

18:50 < krh> ok, halfway through the d* repos

19:16 < krh> it's probably going to take a little longer

19:16 < krh> we added a git filter-branch post-processing step that takes a little longer

Migration Day (4/12)

20:12 < krh> sri: doing eog right now

21:09 < krh> ok, we're now importing
the remaining repos in most-recently-
committed to order

21:10 < krh> jdahlin: it turns out
that jhbuild is at the top of that
list :)

Migration Day (5/12)

21:19 < claude> krh: is it normal some modules get only some branches imported, like dasher?

21:20 < krh> claude: no, that shouldn't happen

21:22 < krh> claude: ok, the imported broke down on the dasher cvs repo

21:26 < krh> oh, it has ()'s in a tag name

Migration Day (6/12)

21:34 < bkor> krh: eh, not to be too critical, but you did compare branches (etc) right?

21:35 < krh> bkor: yes, but not all repos

21:35 < krh> the verification is even more heavy weight than the import process

21:49 < claude> krh: you might also check billreminder import

Migration Day (7/12)

23:26 < rubenv> krh: did the
conversion stop?

23:27 < krh> rubenv: no it's still
running

23:29 < krh> we're doing gnome-vfs
right now

23:30 < krh> and evolution

Migration Day (8/12)

03:40 < owen> krh: is it intentional that you are running the import from nfs?

03:41 < krh> owen: no

03:55 < owen> krh: So, why don't you stop things, then I'll reboot the machine and give it, eh, say 20G of memory

04:03 < owen> OK, we're back up with 20G, leaving aside /tmp, you probably can just mkdir /dev/shm/import and run the import from there

Migration Day (9/12)

04:13 < krh> it doesn't look like it's much faster

04:13 < krh> 2 minutes

04:14 < owen> do you know how long it was before?

04:14 < krh> I'll try that now

04:20 < krh> 2m45 without the ramdisk

Migration Day (10/12)

04:41 < krh> so the repacking hits
100% cpu

04:41 < krh> and parsecvs hits a 100%

04:42 < krh> oh

04:42 < krh> parsecvs has two stages

04:42 < krh> first one is 100% cpu

04:42 < krh> second write out the repo
and only gets around 10% cpu

Migration Day (11/12)

04:46 < owen> I'm still more disturbed by git-filter-branch only getting 5% cpu in general, though I guess all the time is in subprocesses and forking them

04:46 < krh> yeah

04:47 < krh> I guess we could have written a specialized tool for that

Migration Day (12/12)

05:30 < jclinton> krh: the order changed?

05:30 < krh> jclinton: yeah, we're doing it in parallel now

05:31 < krh> and somewhat random

05:31 < krh> to get the big ones out of the way

05:31 < krh> but will return to the last-recently-updated order

05:32 < krh> it turns out that git filter-branch is very slow for what it does

Migration Day+1 (1/7)

13:33 < fredp> claudio: I think the remaining issues are things to be fixed by hand, and owen & krh are probably asleep now.

13:34 < claudio> fredp: gyrus migration seems to have been truncated somewhere in 2006, newer commits don't show up (at least in the cgit web interface)

13:34 < fredp> claudio: you should probably mail to gnome-infrastructure@; I already pointed there a few modules that are missing files.

Migration Day+1 (2/7)

14:54 < fredp> owen: what about seahorse/seahorse-plugins, will they be handled manually after all others?

14:55 < owen> fredp: not sure the plan, I'd have to ask krh what the status is... I know it was discussed earlier

Migration Day+1 (3/7)

16:01 < sandy|out> krh: hey man,
Tomboy's broken in the same way :-)

16:01 < sandy|out>
<http://git.gnome.org/cgiit/tomboy/tag/?id=HEAD>

16:01 < sandy|out> maybe we can just
not import that old CVS tag which was
unfortunately named "HEAD"

Migration Day+1 (4/7)

16:39 < sandy|out> krh: just to be sure, you guys intentionally didn't migrate svn:ignore to .gitignore, right?

16:40 < krh> sandy|out: it's not a 1-to-1 conversion so we punted on that

16:48 < behdad> would have been easier if you copied them in .gitignore and people would just refine them

Migration Day+1 (5/7)

16:52 < krh> behdad: how do svn ignores work again?

16:55 < krh> I did this migration so I wouldn't have to learn svn

16:55 < krh> it's totally backfiring

Migration Day+1 (6/7)

19:02 < abock_> krh: any update on Banshee's migration?

19:03 < krh> abock_: I'm looking into it now

19:03 < krh> and the other casualties of the migration :)

21:35 < owen> pbor: yeah. We one of several modules that have been mentioned that need investigation . Somebody needs compile a list.

21:36 < owen> (billreminder, gnomemm, evince, gnome-common, at least)

Migration Day+1 (7/7)

22:32 < owen> sandy: Sorry, I'm sort of on strike against doing work on the docs, despite having various opinions on what should theoretically be there

22:34 < sandy> I hope they give into your demands (normally when you're on strike you have demands, right?)

22:35 < owen> Well, my main demand is that other people do the work, which seems to be working partially (thanks!), even if other people who said they would help out there when we started the migration didn't

Migration Day+2

21:07 < fredp> owen: would you by change have an idea of an ETA for the few modules with problems? (mostly will it happen before Monday?)

21:08 < owen> fredp: I haven't talked to krh about his plans for when he is going to working on it. He probably has other things to do this weekend, but I don't know for sure.

21:08 < fredp> ok, thanks.

Migration Day+6

Date: Tue, 21 Apr 2009 09:30:18 -0400

From: Kristian Høgsberg
<krh@redhat.com>

Subject: Update on broken repos

I've been working through the list of broken repos and making some progress.

There's still a handful left to fix, but here's the status so far for the repos that have been brought to my attention:

(...)

Migration Day+6

Not yet fixed:

banshee (tags/banshee/<version> got lost)

evince (not sure)

evolution (ETOOMANYTAGS)

gnome-common (branch-name (gnome-core/) and ETOOMANYTAGS)

gnome-perl-introspection (svn subrepos)

gnomemm (svn subrepos)

gtkhtml (not sure)

gyrus (references vendor branch)

seahorse (svn subrepos)

Migration Day+6

```
commit 1e15bd794837...
```

```
Author: Frederic Peters  
<fpeters@od.be>
```

```
Date: Thu Apr 16 22:48:29 2009 +0200
```

```
switched modulesets (> 2.20) to  
use git.gnome.org
```

Why Git is better than X?

Cheap Local Branching

Everything is Local

Git is Fast

Git is Small

The Staging Area

Distributed

Any Workflow

GitHub

Easy to Learn

Why Git is better than X?

Cheap Local Branching

~~Everything is Local~~

Git is Fast

Git is Small

~~The Staging Area~~

Distributed

~~Any Workflow~~

GitHub

~~Easy to Learn~~

Git is Fast & Small (1/4)

From: Ankitkumar Rameshchandra Patel

Why would I need to clone the entire package/application, when I just work on single translation po file?

It's very difficult for localizers running on low speed internet connection to clone the entire packages like gtk, evolution, etc. just to submit a single po file.

Can we have a different repo based on CVS/SVN set for translation files only?

Git is Fast & Small (2/4)

From: Nurali Abdurahmonov

Hi all. In uzbekistan internet traffic very expensive. In git I spend more traffic than svn. I use dial-up connection. With this connection I cannot clone big modules like gtk. Can I clone "po" folder only? Or can I use svn for my translations?

Git is Fast & Small (3/4)

From: Funda Wang

Hello,

I want to know if there is certain easy way of only checking out po dir certain module with git. While we are using CVS and SVN, we as committers can only checkout po dir of certain module, which will reduce network traffic a lot for large modules such as gtk+.

Thanks.

Git is Fast & Small (4/4)

From: Xavi de Blas

[...]

if i do this:

```
git clone ssh://[...]
```

i clone all the projects page, but i only want to clone mine (chronojump) here what i want to clone:

```
http://git.gnome.org/cgi/gnomeweb-wml/tree/projects.gnome.org/chronojump
```

sorry, i'm sure it has to be really simple but i haven't read solution on documentation and i cannot chat

Why Git is better than X?

Cheap Local Branching

~~Everything is Local~~

~~Git is Fast~~

~~Git is Small~~

~~The Staging Area~~

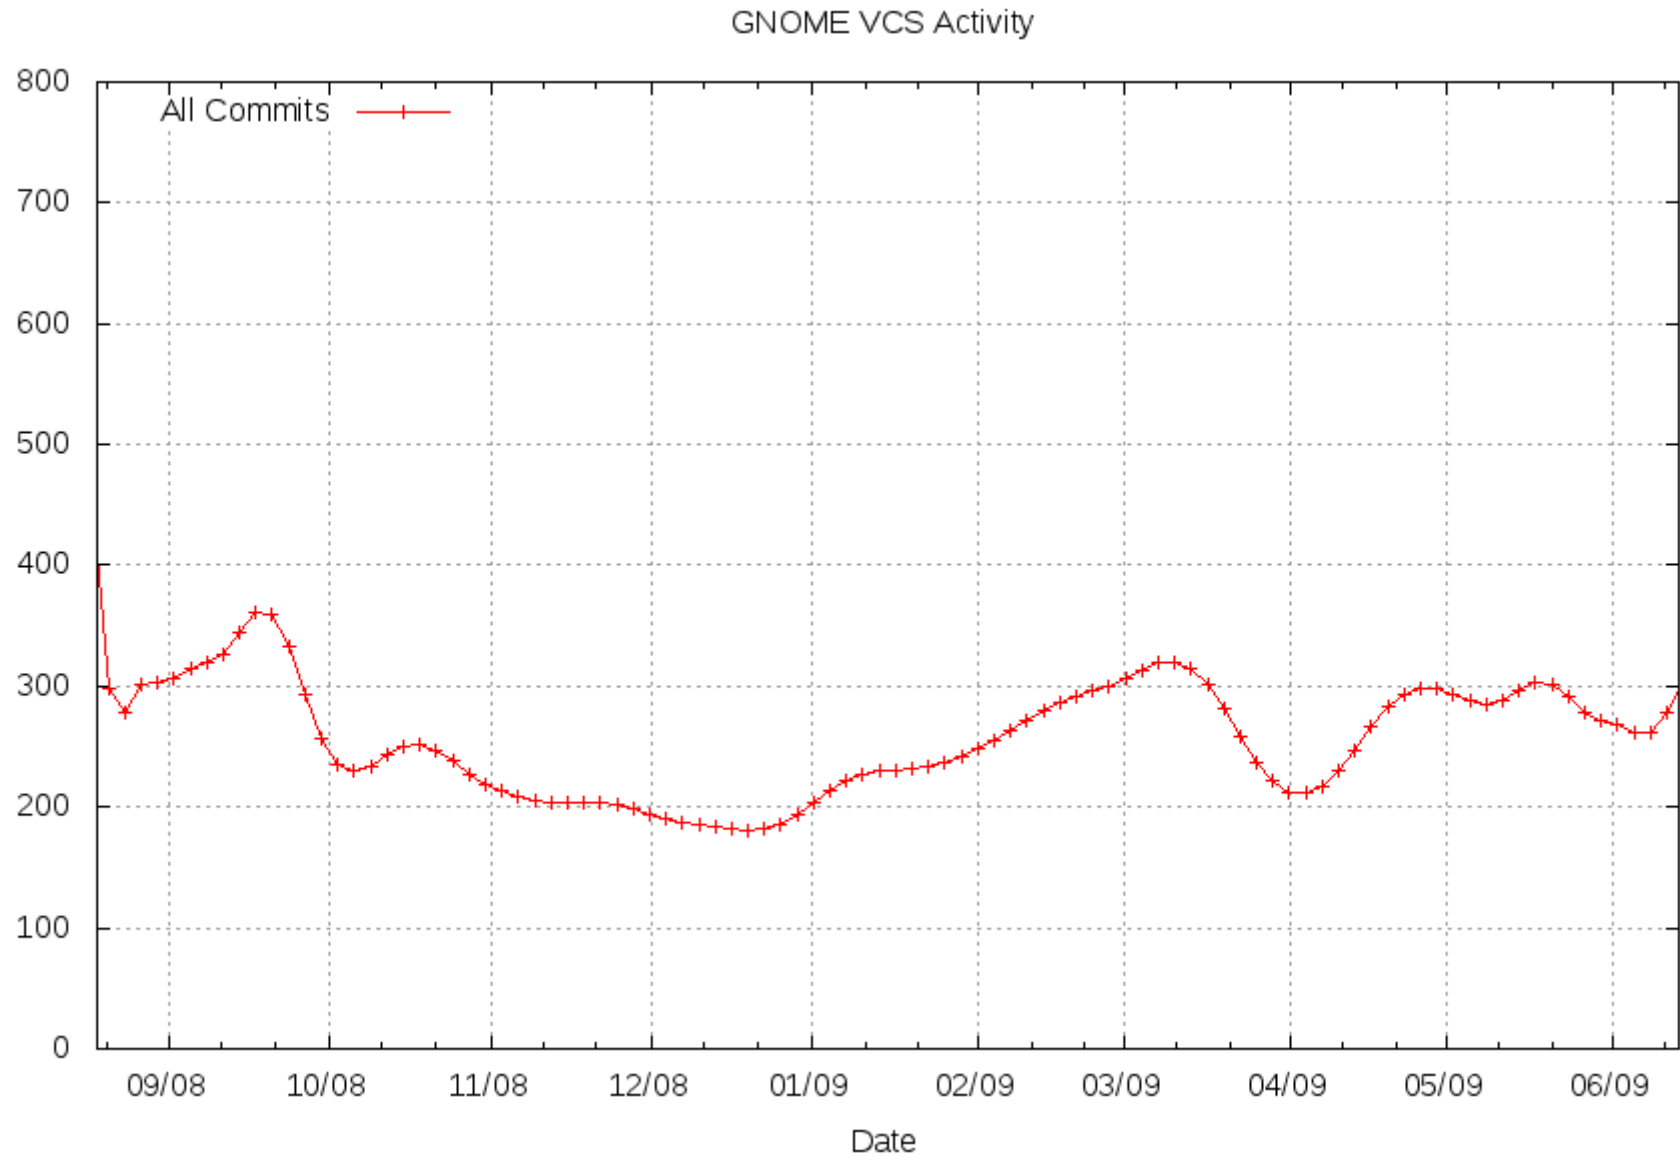
Distributed

~~Any Workflow~~

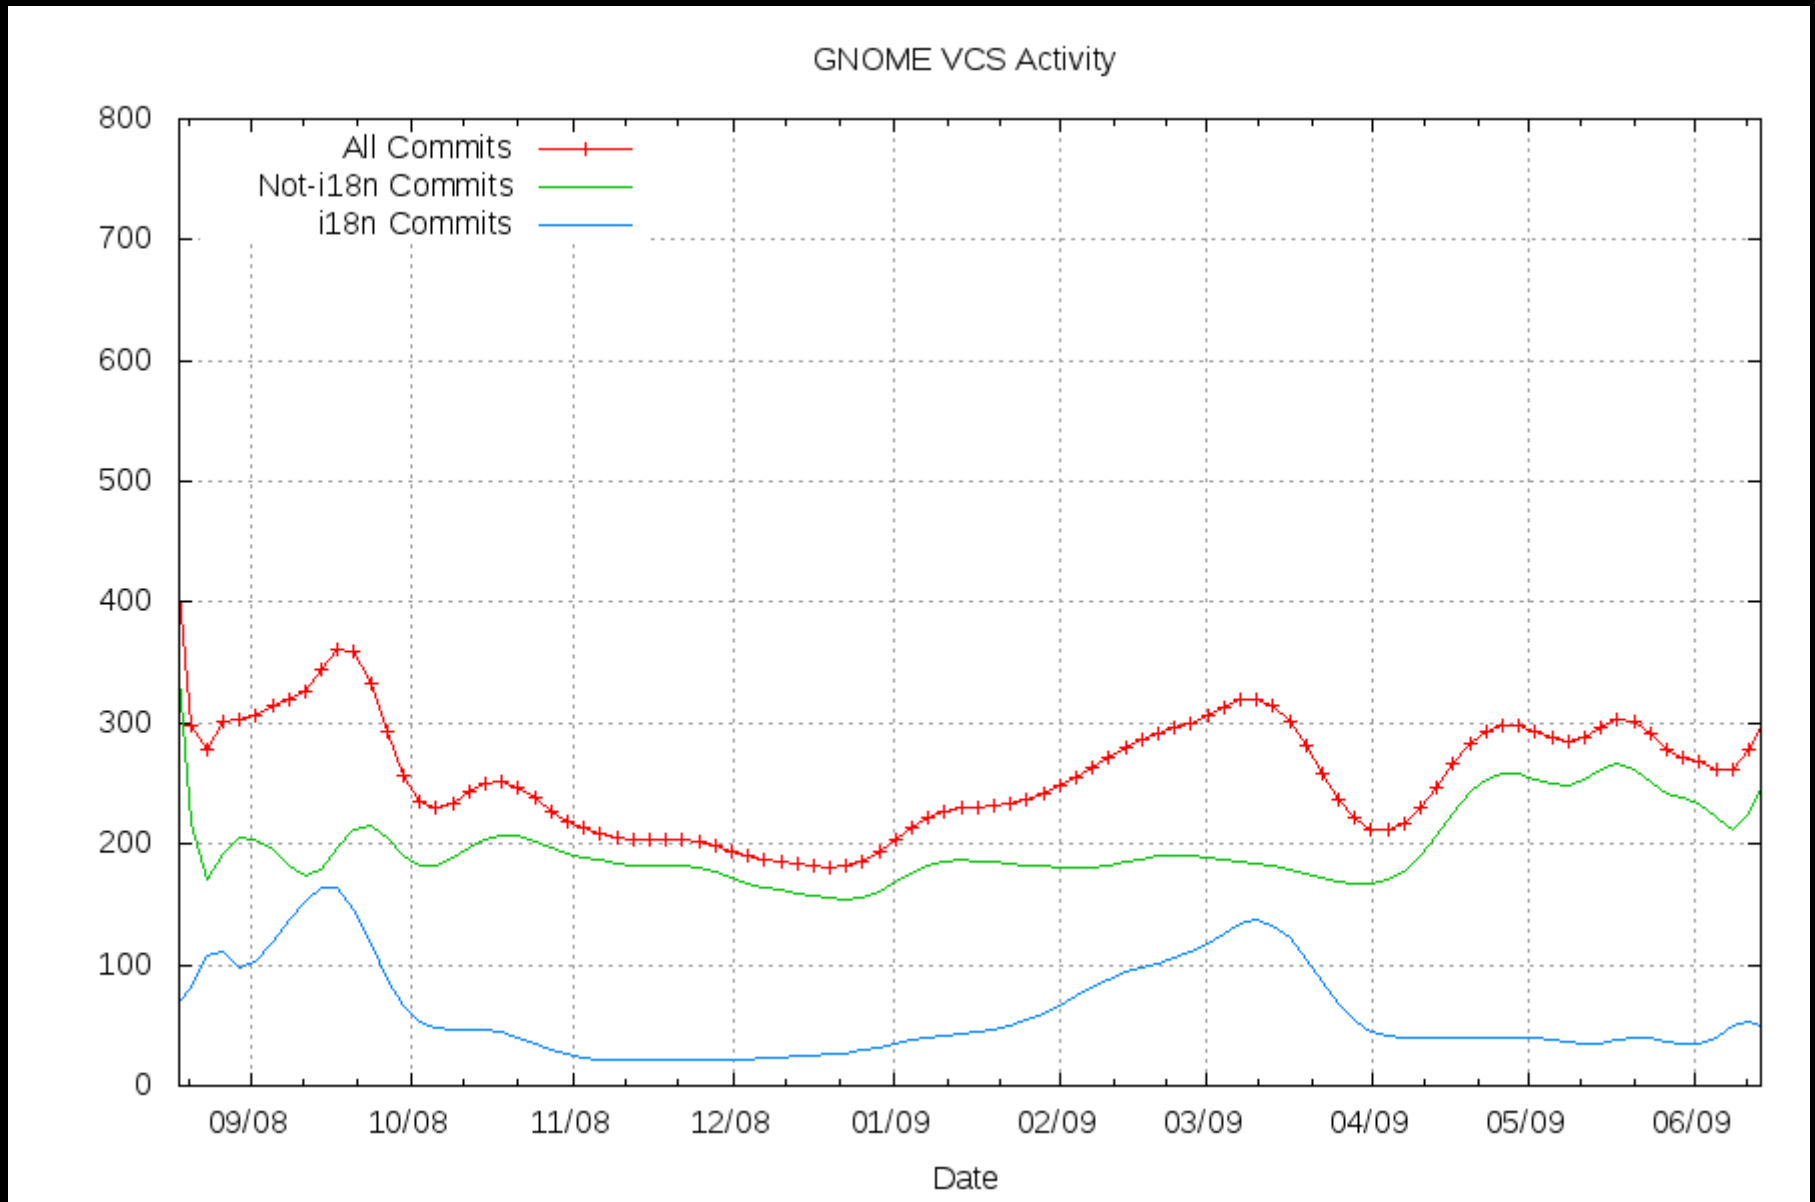
GitHub

~~Easy to Learn~~

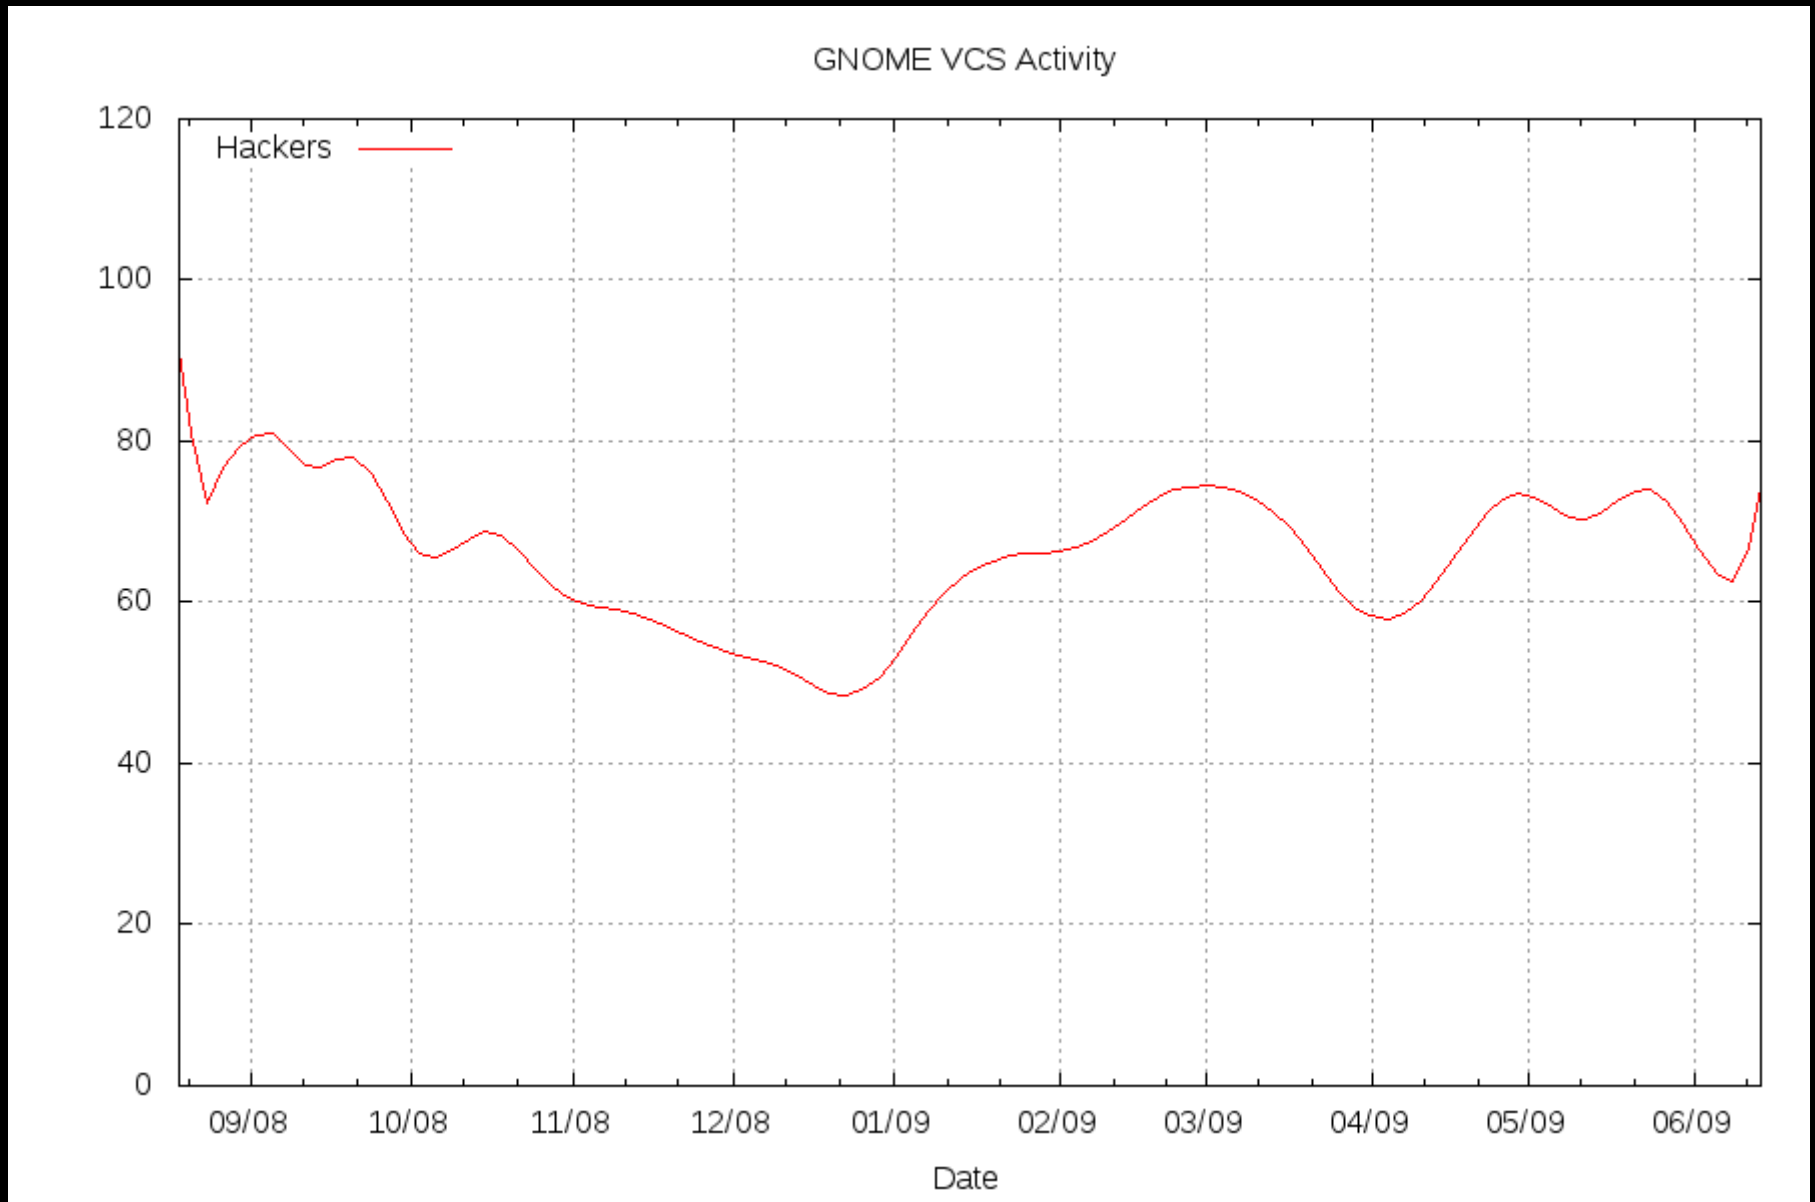
Today (1/3)



Today (2/3)



Today (3/3)



Tomorrow (?)

(July 7th)

From: Behdad Esfahbod

KDE has a git bof tomorrow in the morning. They are considering moving and like to hear about our experience. The Gitorious developers will be there.

Would be good for our sysadmins and git people to show up.

Thanks

Q/A