

Using Metrics to Track Code Review Performance

An Expreience Report

Daniel Izquierdo-Cortazar, Nelson Sekitoleko,
Jesus M. Gonzalez-Barahona, Lars Kurth

jgb@bitergia.com @jgbarah
Bitergia / LibreSoft (URJC)

Evaluation and Assesment in Software Engineering Conference
(EASE 2017)
Karlskrona (Sweden), June 15-16 2017



Code review increases time-to-deploy

Code review improves code quality...

...but it delays availability of new code

Can we characterize the problem?

Pre-commit peer code review

Research questions

- Characterization of code review activity
- Impact in delaying new code
- Does complexity of changes matter?

Context: two Linux-related projects

Code review in Xen and Netdev

Project	Patch series	Reviewers	Comments
Xen	8,138	235	54,139
Netdev	10,680	240	29,674

(2012-2015)

Patch series: one or more patches
that are reviewed together



Reviewing a patchset of 1 patch

- [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Juergen Gross*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *David Vrabel*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Juergen Gross*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Ian Campbell*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *David Vrabel*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Ian Jackson*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Juergen Gross*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Wei Liu*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Juergen Gross*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Andrew Cooper*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Juergen Gross*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Andrew Cooper*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Juergen Gross*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Andrew Cooper*
 - Re: [Xen-devel] [PATCH] libxc: try to find last used pfn when migrating, *Juergen Gross*

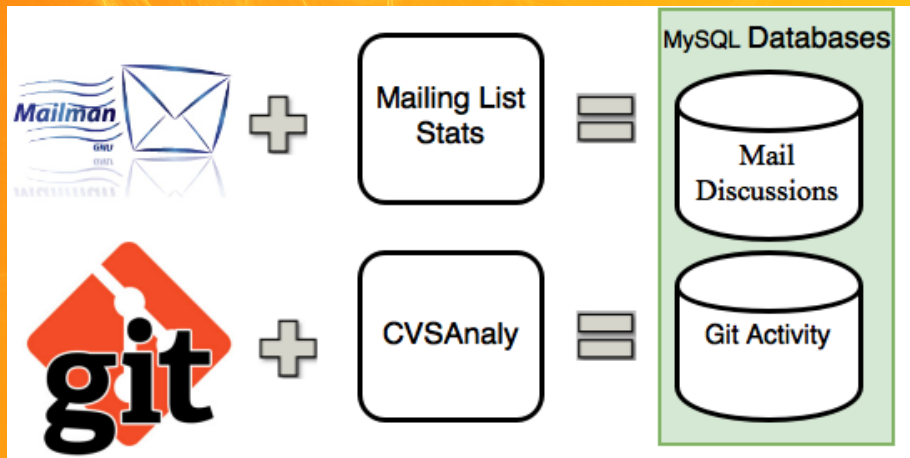
Reviewing a patchset of 9 patches

- [Xen-devel] [PATCH v9 0/9] Introduce HVM without dm and new boot ABI, *Roger Pau Monne*
 - [Xen-devel] [PATCH v9 1/9] xen/x86: set the vPMU interface based on the presence of a lapic, *Roger Pau Monne*
 - Re: [Xen-devel] [PATCH v9 1/9] xen/x86: set the vPMU interface based on the presence of a lapic, *Boris Ostrovsky*
 - [Xen-devel] [PATCH v9 2/9] xen/x86: make sure the HVM callback vector is correctly set, *Roger Pau Monne*
 - [Xen-devel] [PATCH v9 3/9] xen/x86: allow disabling all emulated devices inside of Xen, *Roger Pau Monne*
 - [Xen-devel] [PATCH v9 4/9] libxc: allow creating domains without emulated devices., *Roger Pau Monne*
 - [Xen-devel] [PATCH v9 5/9] xen/x86: allow HVM guests to use hypercalls to bring up vCPUs, *Roger Pau Monne*
 - Re: [Xen-devel] [PATCH v9 5/9] xen/x86: allow HVM guests to use hypercalls to bring up vCPUs, *Julien Grall*
 - [Xen-devel] [PATCH v9 6/9] libxc/xen: introduce a start info structure for HVMlite guests, *Roger Pau Monne*
 - Re: [Xen-devel] [PATCH v9 6/9] libxc/xen: introduce a start info structure for HVMlite guests, *Jan Beulich*
 - [Xen-devel] [PATCH v9 7/9] libxc: switch xc_dom_elfloader to be used with HVMlite domains, *Roger Pau Monne*
 - [Xen-devel] [PATCH v9 8/9] libxl: allow the creation of HVM domains without a device model., *Roger Pau Monne*
 - [Xen-devel] [PATCH v9 9/9] libxl: add support for migrating HVM guests without a device model, *Roger Pau Monne*
 - Re: [Xen-devel] [PATCH v9 9/9] libxl: add support for migrating HVM guests without a device model, *Andrew Cooper*
 - Re: [Xen-devel] [PATCH v9 9/9] libxl: add support for migrating HVM guests without a device model, *Roger Pau Monne*
 - Re: [Xen-devel] [PATCH v9 9/9] libxl: add support for migrating HVM guests without a device model, *Wei Liu*

Methodology

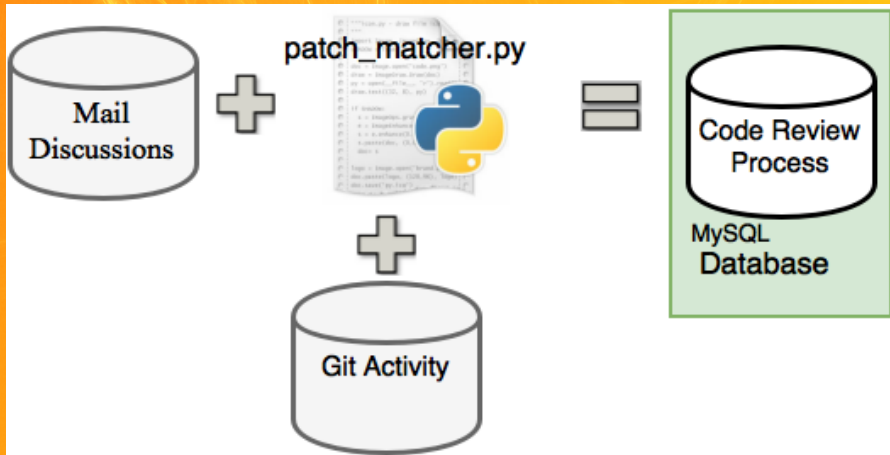


Methodology: data retrieval



<http://metricsgrimoire.github.io>

Methodology: matching & merging



Methodology: the devil is in the details

Low-hanging fruits are not enough

- Time to merge: from first message to commit
- Review threads difficult to detect: specific algorithm
- Matching of subjects to commits difficult: several patterns
- Heuristics for abandoned reviews
- Testing hypothesis... with the community

Methodology: matching & merging

Year	Total patches	Matched to commits
Xen		
2012	1907	1108
2013	2503	1549
2014	2092	1558
2015	2773	2114
Netdev		
2012	4031	2267
2013	3660	2165
2014	3224	2100
2015	3123	2057


Methodology: analysis




- Time to merge
- Time to commit
- Time to review
- No. of comments
- No. of patch series
- Main Devs
- Main Orgs
- etc...

Results: code review activity

Xen



Netdev




Patch series by number of patches

Results: code review activity


Xen

Evolution of the mean (blue) and median (green) comments per patch



Netdev


Evolution of the mean (blue) and median (green) comments per patch




Comments per patch

Results: delay due to code review

Xen



Netdev




Time to merge

Results: impact of complexity


Xen

Netdev

Time to merge for patch series per number of patches (without outliers)

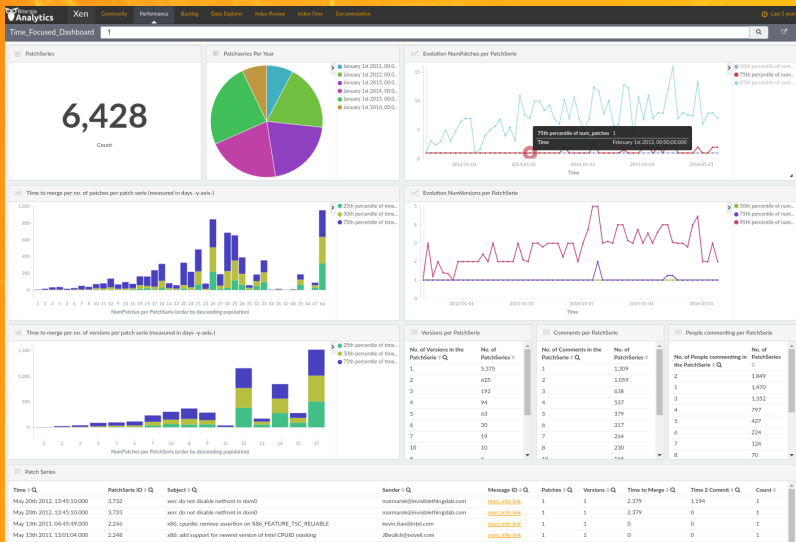


Time to merge for patch series per number of patches (without outliers)



Time to merge per number of patches in patch series

Further work: continuous monitoring



<http://xen.biterg.io>

Summary

Characterization, checked with practitioners:

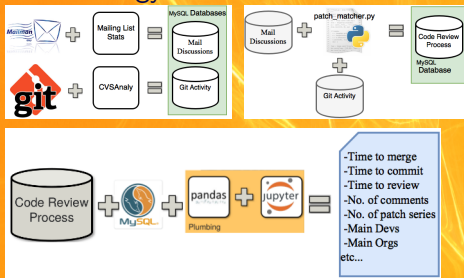
- Activity: patchsets, patches per patchset, comments
- Delay: time to merge
- Impact from complexity: time to merge by number of patches in patchset

All the analysis can be automated


Using Metrics to Track Code Review Performance

*Daniel Izquierdo-Cortazar, Nelson Sekitoleko,
Jesus M. Gonzalez-Barahona, Lars Kurth*

Methodology



Results:



Main metrics:

- Activity: patch series per # patches, # comments
- Delay: time to merge

Takeaways:

- Simple, automatable metrics
- Complexity matters
- Data sets available

License

©2016 Bitergia

Some rights reserved.

This presentation is distributed under the
“Attribution-ShareAlike 3.0” license, by Creative Commons,
available at

<http://creativecommons.org/licenses/by-sa/3.0/>