

PostgreSQL hacking introduction

... how to write the first patch

Tomáš Vondra / EDB

tomas.vondra@postgresql.org



Why I do this workshop?

Agenda

- basics of the development process
 - structure of community, commitfests, tools
 - development process(es)
 - which patch to write / review?
- workflows:
 - basic: git clone, configure, make, tests, ...
 - advanced: CF app, github CI, buildfarm, cfbot, ...

emergent structure
emergent processes

How the community works?

- organic origins, still evolving
- not dominated by a single company
 - EDB, Crunchy Data, Microsoft, Amazon, Cybertec, data egret, dbi, Dalibo, Fujitsu, Percona, ...
 - typical: company using Postgres internally / providing services
 - rules against "domination" of a single company
- core team (7 people, mostly "steering committee / BoD")
- committers (~30 people, right to commit into git repository)
- CoC, (major) contributors, ...

Where's the community?

- e-mail lists
 - postgresql-general, postgresql-performance, ...
 - postgresql-hackers (<https://www.postgresql.org/list/postgresql-hackers/>)
- IRC
 - <https://www.postgresql.org/community/irc/>
- slack, discord, ...
- conferences
 - <https://www.postgresql.org/about/events/>

How is the development process structured?

- yearly "major" releases (~September)
 - feature freeze in April
- quarterly "minor" releases
 - bugfixes, may be more often if severe issues
- commitfests
 - cycle: 1 month work, 1 month pause
 - "soft" rule (no one prohibits work / commit anytime you wish)
 - <https://commitfest.postgresql.org/>

life cycle of a patch



life cycle of a contributor



idea





idea




writing patch



idea





writing patch




sending patch to
pgsql-hackers





idea



writing patch





sending patch to
pgsql-hackers




reviews / testing





idea



writing patch




sending patch to
pgsql-hackers



reviews / testing





committed








idea



writing patch



sending patch to
pgsql-hackers



reviews / testing



committed





buildfarm gets red








idea



writing patch



sending patch to
pgsql-hackers



reviews / testing



committed



buildfarm gets red



tools

pgsql-hackers [<https://www.postgresql.org/list/pgsql-hackers/>]

pgsql-hackers

The PostgreSQL developers team lives here. Discussion of current development issues, problems and bugs, and proposed new features. If your question cannot likely that only a developer will know the answer, you may re-post your question in this list. You must try elsewhere first!

Search the Archives



(enter a message-id to go directly to that message)

[Advanced Search](#)

View Archives

2023		
November 2023	View Archives	Download mbox
October 2023	View Archives	Download mbox
September 2023	View Archives	Download mbox
August 2023	View Archives	Download mbox
July 2023	View Archives	Download mbox
June 2023	View Archives	Download mbox
May 2023	View Archives	Download mbox
April 2023	View Archives	Download mbox
March 2023	View Archives	Download mbox
February 2023	View Archives	Download mbox
January 2023	View Archives	Download mbox
2022		
December 2022	View Archives	Download mbox
November 2022	View Archives	Download mbox
October 2022	View Archives	Download mbox
September 2022	View Archives	Download mbox
August 2022	View Archives	Download mbox
July 2022	View Archives	Download mbox
June 2022	View Archives	Download mbox
May 2022	View Archives	Download mbox
April 2022	View Archives	Download mbox
March 2022	View Archives	Download mbox



CF app [<https://commitfest.postgresql.org/>]

Make it possible to add custom options to EXPLAIN	Needs review		Konstantin Knizhnik (knizhnik)		1	2023-11-02 13:07	2023-10-21 12:16
Performance							
More scalable multixacts buffers and locking	Needs review		Kyotaro Horiguchi (horiguti), Andrey Borodin (x4m), Ivan Lazarev (avaness)		18	2023-10-02 17:34	2023-01-11 17:13
Parallelize correlated subqueries that execute within each worker	Needs review		James Coleman (jcole1989)		11	2023-10-02 14:22	2023-07-04 01:26
In-place persistence change of a relation (fast ALTER TABLE ... SET LOGGED with wal_level=minimal)	Needs review		Kyotaro Horiguchi (horiguti)	Nathan Bossart (bossartn), Justin Pryzby (justinpryzby), Jakub Wartak (jakub.wartak.tomtom)	10	2023-10-02 11:04	2023-09-04 08:37
AcquireExecutorLocks() and run-time pruning	Needs review		Amit Langote (amitlan)	Tom Lane (tgl)	10	2023-10-02 11:03	2023-09-28 08:26
Speed up releasing of locks	Needs review		Andres Freund (andresfreund), David Rowley (davidrowley)	Heikki Linnakangas (heikki)	10	2023-10-02 11:04	2023-11-09 16:18
nbtree performance improvements through specialization on key shape	Needs review		Matthias van de Meent (mmeent)	David Christensen (davidchristensen)	8	2023-10-02 11:05	2023-10-30 16:19
Add sortsupport for range types and btree_gist	Needs review		Christoph Heiss (christoph.heiss)		8	2023-10-02 09:45	2022-11-30 17:25
Improve dead tuple storage for lazy vacuum	Waiting on Author		Masahiko Sawada (masahikosawada), John Naylor (john.naylor)		8	2023-10-02 10:56	2023-10-28 08:56
Reducing planning time when tables have many partitions	Waiting on Author		Yuya Watari (watari)	Thom Brown (thom), Ashutosh Bapat (ashutoshbapat), Andrei Lepikhov (lepikhov)	8	2023-10-02 10:56	2023-09-27 07:28
ALTER TABLE and CLUSTER fail to use a BulkInsertState for toast tables	Waiting on Author	17	Justin Pryzby (justinpryzby)		7	2023-11-07 16:17	2023-11-07 16:17
Reduce timing overhead of EXPLAIN ANALYZE using rdts	Needs review	16	Andres Freund (andresfreund), Lukas Fittl (lfittl), David Geier (geidav)	Lukas Fittl (lfittl), Hannu Krosing (hannukrosing)	7	2023-10-02 11:05	2023-02-20 10:36
asynchronous execution support for Custom Scan	Needs review		KaiGai Kohei (kaigai), kazutaka onishi (onishi)		7	2023-10-02 09:46	2022-12-01 23:35
Add semi-join pushdown to postgres_fdw	Needs review		Alexander Pyhalov (a.pyhalov)	Yuki Fujii (fujiiyuki)	7	2023-10-02 11:05	2023-10-31 11:07

<https://git.postgresql.org> [[link](#)]

Project	Description	Owner	Last Change
2ndquadrant_bdr.git	Obsolete repo for Bi-direction...	Abhijit Menon-Sen	7 years ago summary shortlog log tree
advocacy.git	This repository contains pictu...	Jonathan Katz	11 years ago summary shortlog log tree
automirror.git	The postgresql.org automirror...	Dave Page	11 years ago summary shortlog log tree
bucardo.git	Bucardo		7 years ago summary shortlog log tree
cfbot.git	cfbot repo	Stephen Frost	No commits summary shortlog log tree
check_postgres.git	check_postgres		6 years ago summary shortlog log tree
docbot.git	pg_docbot source code, website...	Andreas 'ads' Scherbaum	5 years ago summary shortlog log tree
fosdem2021-content.git	Public content for the Postgre...	Andreas 'ads' Scherbaum	21 months ago summary shortlog log tree
fosdem2021-static.git	Static content for the Postgre...	Andreas 'ads' Scherbaum	2 years ago summary shortlog log tree
hamn.git	This is the code behind http...	Magnus Hagander	2 months ago summary shortlog log tree
json-datatype.git	Repository for Google Summer...		12 years ago summary shortlog log tree
libusual.git	Low-level utility library...	Marko Kreen	8 years ago summary shortlog log tree
pg_adviser.git	Moving the pgAdviser project...	Gurjeet Singh	13 years ago summary shortlog log tree
pg_bsd_indent.git	Postgres' version of the BSD...	Tom Lane	6 months ago summary shortlog log tree
pg_filedump.git	pg_filedump provides facilitie...	Christoph Berg	8 weeks ago summary shortlog log tree
pgapt.git	Code and repository config...	Christoph Berg	26 hours ago summary shortlog log tree
pgarchives.git	PostgreSQL mailinglist archive...	Magnus Hagander	2 months ago summary shortlog log tree
pgbench-tools.git	Tools for automating multiple...	Greg Smith	12 years ago summary shortlog log tree
pgbouncer.git	PgBouncer master repo	Marko Kreen	8 years ago summary shortlog log tree
pgcac-website.git	PostgreSQL Community Associati...	Dave Page	3 months ago summary shortlog log tree
pgcommitfest.git	CommitFest management code.	Robert Haas	9 years ago summary shortlog log tree
pgcommitfest2.git	New generation pg commitfest...	Magnus Hagander	5 months ago summary shortlog log tree
pgdaychicago-test.git	PGDay Chicago Test Site	Mark Wong	2 days ago summary shortlog log tree
pgdayuk.git	Websites PGDay UK events.	Dave Page	8 weeks ago summary shortlog log tree
pgdu2018-test.git	PgDU 2018 Test	Gary Evans	5 years ago summary shortlog log tree



git clone

```
git clone https://github.com/postgres/postgres.git
```

78 min ago	master	shortlog log tree
25 hours ago	REL_16_STABLE	shortlog log tree
34 hours ago	REL_15_STABLE	shortlog log tree
34 hours ago	REL_12_STABLE	shortlog log tree
34 hours ago	REL_13_STABLE	shortlog log tree
34 hours ago	REL_14_STABLE	shortlog log tree
3 days ago	REL_11_STABLE	shortlog log tree
9 months ago	REL9_2_STABLE	shortlog log tree



configure

```
configure --help
```

```
configure --prefix=$HOME/pgbuild \  
  --enable-depend \  
  --enable-tap-tests --enable-cassert \  
  --with-lz4
```

make targets

`make`

`make -s -j4 install`

`make (check | installcheck)`

`make check-world`

`make html`

`make coverage`

Types of tests

regression tests

- basic SQL script + expected output

isolation tests

- interaction of multiple sessions (ACID, ...)









TAP tests

- more complex stuff (multiple clusters, ...)
- replication, pg_dump, pg_upgrade

Valgrind

cirrus-ci.org [[src/tools/ci/README](https://cirrus-ci.org/src/tools/ci/README)]

tvondra > postgres






Builds					
	free/alloc bufiles in hashjoin code -O 346ace9	Repository postgres	Owner tvondra	Branch buffile-release	11:05
	WIP -O 6be8ea2	Repository postgres	Owner tvondra	Branch index-selectivity	10:28
	fixup: do not prefetch for empty indexes -O 916576e	Repository postgres	Owner tvondra	Branch index-insert-prefetch	01:23
	wip -O be74cc4	Repository postgres	Owner tvondra	Branch index-insert-prefetch	08:56
	prefetch -O fd74085	Repository postgres	Owner tvondra	Branch index-insert-prefetch	14:36
	PoC prefetch for IOS -O 766a41d	Repository postgres	Owner tvondra	Branch index-prefetch-v6	22:28
	index prefetch v6 -O 8659787	Repository postgres	Owner tvondra	Branch index-prefetch-v6	15:14
	test_decoding: Remove unnecessary table in twophase test -O 98e8974	Repository postgres	Owner tvondra	Branch master	14:13



































































cfbot [<http://cfbot.cputube.org/>]

PostgreSQL Patch Tester

[Current commitfest](#) | [Next commitfest](#) | [FAQ](#) | [Statistics](#) | [Highlights](#)

Here lives an experimental bot that converts email threads that are registered in the [Commitfest system](#) into [branches on Github](#), and collates test results from [Cirrus CI](#). Key:  or  = new/recently changed,  or  = stable,  = working.

Ready for Committer

45/4633	ALTER DEFAULT PRIVILEGES only affects objects created by the target role	Laurenz Albe	 H	     
45/3659	Add SPLIT PARTITION/MERGE PARTITIONS commands	Dmitry Koval	  H	     
45/3511	Add non-blocking version of PQcancel	Jelte Fennema-Nio	 H	     
45/4632	Clarify that you can create unique constraints on ONLY partitioned table	Laurenz Albe	 H	     
45/4240	Cross-database SERIALIZABLE safe snapshots	Thomas Munro	 H	     
45/4209	Evaluate arguments of correlated SubPlans in the referencing ExprState	Andres Freund	 H	     
45/4585	Fix a wrong comment in setrefs.c	Richard Guo	 H	     
45/4593	Fix output of zero privileges in psql	Erik Wienhold	 H	     
45/4127	Infinite interval values	Ashutosh Bapat, Joseph Koshakow, jian he	 H	     



buildfarm [<https://buildfarm.postgresql.org>]



PostgreSQL BuildFarm Status

Shown here is the latest status of each farm member for each branch it has reported on in the last 30 days.

Use the farm member link for history of that member on the relevant branch.

Legend	= cassert	= debug	= gssapi	= krb5	= llvm	= meson	= nls	= openssl
	= pam	= perl	= python	= tap-tests	= tcl	= thread-safety	= vpath	= xml

Branch: HEAD			
Alias	System	Status	Flags
hachi	Debian sid clang 11.0.1 aarch64	00:20 ago OK [b7f315c] Config	
komodoensis	Debian Sid gcc recent x86_64	00:29 ago OK [b7f315c] Config	
snakefly	AmazonLinux@Graviton2 AL2 gcc 7.3.1 aarch64	00:37 ago OK [b7f315c] Config	
dragonet	Debian Sid clang clang 3.9 x86_64	00:29 ago OK [b7f315c] Config	
crake	Fedora 38 gcc 13.2.1 x86_64	00:32 ago OK [b7f315c] Config	
idiacanthus	Debian Sid clang 5 x86_64	00:35 ago OK [b7f315c] Config	
desmoxytes	Debian Sid gcc recent x86_64	00:34 ago OK [b7f315c] Config	
taipan	Debian Sid gcc 13 x86-64	00:36 ago OK [b7f315c] Config	
bushmaster	Debian Sid clang 17 x86-64	00:36 ago OK [b7f315c] Config	
myna	Photon 3.0 Gcc 6.3.0 x86_64	00:24 ago OK [b7f315c] Config	
butterflyfish	Photon 2.0 Gcc 6.3.0 x86_64	00:24 ago OK [b7f315c] Config	
hippopotamus	openSUSE openSUSE 15.3 gcc gcc 7.5 x86_64	00:36 ago OK [b7f315c] Config	
indri	macOS+MacPorts 14.0 clang 15.0.0 arm64	00:26 ago OK [b7f315c] Config	
longfin	OS X 13.6 clang 14.0.3 x86_64	00:37 ago OK [b7f315c] Config	
olingo	Debian sid clang recent x86-64	00:36 ago OK [b7f315c] Config	
grassquit	Debian sid gcc recent x86-64	00:36 ago OK [b7f315c] Config	
flaviventris	Debian Sid gcc snapshot x86_64	00:36 ago OK [b7f315c] Config	



coverage [<https://coverage.postgresql.org/>]

LCOV - code coverage report

Current view: top level

Test: PostgreSQL 17devel

Date: 2023-11-11 19:10:41

Legend: Rating: low: < 75 % medium: >= 75 % high: >= 90 %

	Hit	Total	Coverage
Lines:	393083	498534	78.8 %
Functions:	23558	26430	89.1 %

Directory	Line Coverage ↕	Functions ↕
contrib/adminpack	<div><div></div></div> 45.7 % 79 / 173	70.8 % 17 / 24
contrib/amcheck	<div><div></div></div> 67.0 % 971 / 1449	94.3 % 50 / 53
contrib/auth_delay	<div><div></div></div> 0.0 % 0 / 13	0.0 % 0 / 3
contrib/auto_explain	<div><div></div></div> 89.3 % 92 / 103	100.0 % 6 / 6
contrib/basebackup_to_shell	<div><div></div></div> 86.5 % 90 / 104	100.0 % 14 / 14
contrib/basic_archive	<div><div></div></div> 46.8 % 52 / 111	90.0 % 9 / 10
contrib/bloom	<div><div></div></div> 90.1 % 456 / 506	96.7 % 29 / 30
contrib/bool_plperl	<div><div></div></div> 100.0 % 11 / 11	100.0 % 5 / 5
contrib/btree_gin	<div><div></div></div> 93.3 % 181 / 194	100.0 % 202 / 202
contrib/btree_gist	<div><div></div></div> 94.7 % 2071 / 2186	97.6 % 520 / 533
contrib/citext	<div><div></div></div> 100.0 % 156 / 156	100.0 % 35 / 35
contrib/cube	<div><div></div></div> 85.8 % 939 / 1094	94.2 % 97 / 103
contrib/dblink	<div><div></div></div> 86.7 % 935 / 1078	97.4 % 75 / 77
contrib/dict_int	<div><div></div></div> 97.4 % 38 / 39	100.0 % 5 / 5
contrib/dict_xsyn	<div><div></div></div> 96.2 % 100 / 104	100.0 % 8 / 8
contrib/earthdistance	<div><div></div></div> 90.9 % 20 / 22	100.0 % 5 / 5
contrib/file_fdw	<div><div></div></div> 76.5 % 247 / 323	90.9 % 20 / 22
contrib/fuzzystrmatch	<div><div></div></div> 35.5 % 385 / 1083	87.2 % 41 / 47
contrib/hstore	<div><div></div></div> 83.7 % 1490 / 1780	76.7 % 132 / 172
contrib/hstore_plperl	<div><div></div></div> 100.0 % 58 / 58	100.0 % 6 / 6
contrib/hstore_plpython	<div><div></div></div> 97.0 % 64 / 66	100.0 % 6 / 6
contrib/intarray	<div><div></div></div> 89.3 % 1295 / 1450	92.6 % 138 / 149
contrib/isn	<div><div></div></div> 76.0 % 381 / 501	88.9 % 40 / 45
contrib/jsonb_plperl	<div><div></div></div> 94.3 % 115 / 122	100.0 % 10 / 10



inspiration

how to come up with a patch idea?

- ideally something that you find interesting / useful
 - something you need at work, personal interest, ...
 - e.g. slow query => thinking about how to speed it up
- go through patches in the CF app
 - review an existing patch, help the author, write a related patch, ...
- coverage report
 - code with low coverage - write a couple tests, find/fix bugs
- for example
 - what if the memory context only uses mmap() and not malloc()?

sources

Sources

- Developer FAQ
https://wiki.postgresql.org/wiki/Developer_FAQ
- Server programming
<https://www.postgresql.org/docs/16/server-programming.html>
- Internals
<https://www.postgresql.org/docs/16/server-programming.html>
- PostgreSQL 14 Internals / Egor Rogov [PDF, paperback, ...]
<https://postgrespro.com/community/books/internals>
- The Internals of PostgreSQL / Hironobu SUZUKI
for database administrators and system developers
<https://www.interdb.jp/pg/>

tomas.vondra@postgresql.org