

PostgreSQL “New Stuff” 2016

Christopher Browne
For GTALUG

Agenda

- PostgreSQL 9.5 was released in January 2016 with plenty of Neat New Stuff
<https://www.postgresql.org/docs/9.6/static/release-9-5.html>
- PostgreSQL 9.6 is being prepped for release “somewhat soon”, with even more
<https://www.postgresql.org/docs/devel/static/release-9-6.html>
- Famous Classifications of New Features
 - **Performance optimizations** - life gets better upon upgrading, “nearly for free”
 - **Administrative improvements** - administrators can do better things
 - **Development features** - you’ll need to modify code to benefit from these
 - Enhancements to built-in **replication**

New Things Blogged

Depesz blog “Waiting for 9.5”, “Waiting for 9.6”, etc has described features in progress; lots of good explanations of interesting upcoming features...

- <https://www.depesz.com/tag/pg96/>
- <https://www.depesz.com/tag/pg95/>
- <https://www.depesz.com/tag/pg94/>
- <https://www.depesz.com/tag/pg93/>
- <https://www.depesz.com/tag/pg92/>
- <https://www.depesz.com/tag/pg91/>
- <https://www.depesz.com/tag/pg90/>

PostgreSQL 9.5 Optimizations

- BRIN indexes, storing compact summary about sets of pages
If a table is near “insert only,” BRIN on increasing columns (dates, IDs) is compact and cheap. Data warehouse systems will like it!
- Abbreviated keys to improve the performance of sorts
- Lots of improvements to buffer pinning, lock scalability - magicky, invisible!!!
- Numerous changes difficult to enumerate; “some queries will now run faster”
- CHECK CONSTRAINTS on foreign tables can improve optimization
- Allow Foreign Data Wrappers to do join pushdown - *Remote DB Federation*
- Use 128 bit integer accumulators for some aggregate functions
- Improved array performance in `plpgsql`

PG 9.5 Administrative Improvements

- Allow setting a **cluster name** - useful if running many PG instances on one host
- Control Linux OOM Killer via environment variables
- Mark parameter changes as “**pending_restart**”
- Row level security over who can see/modify data within a table

PG 9.5 Developer Improvements

- New capability to capture transaction commit timestamps
This would have been a “game changer” for Slony, if available 10 years ago!
- “UPSERT” - allow INSERTs that generate constraint conflicts to be transformed into UPDATEs. It was much tougher to do than one would imagine... because...
CONCURRENCY!!!
- GROUP BY now supports GROUPING SETS, ROLLUP, CUBE
- SELECT can SKIP LOCKED so worker processes that want to squabble over a queue table can grab mutually exclusive sets of tuples and play well together
- REINDEX can now work on a whole schema, not just on a table/index

PostgreSQL 9.5 developer enhancements

- Foreign Schema Import - automatically set up local schema to point to ALL tables in a remote database
- Event Triggers
 - `ddl_command_end` captures DDL change
 - On table rewrites triggered by `ALTER TABLE`
 - On `COMMENT`, `SECURITY LABEL`, `GRANT`, `REVOKE`
- Bunch of JSONB enhancements
 - Myles will be talking about JSONB...
- Add `ASSERT` support in `pl/pgsql`

PG 9.5 Development Improvements

- `pg_export_snapshot()` allows multiple transactions to share transaction context, useful for increasing parallelism, e.g. - dump 10 tables concurrently with consistent transaction context

Parallel Query - in 9.6

- `pg_export_snapshot()` was important to this...
- Parallel query starts with a Large Table needing a Sequential Scan in a read-only transaction
 - Opens several worker processes, each sharing the same snapshot
 - Hash joins and nested loops can be performed in parallel
 - Aggregates are performed in parallel
- Lots more work to be done...
 - Lots of preparatory work took place in 9.4 and 9.5
 - Efforts ongoing to do parallel index generation
 - Later versions should harness this to parallelize more activities

PG 9.5 Replication Improvements

- Add mechanisms to track replication progress
- Option to compress full page writes (reduces I/O at cost of more CPU load)
- New tool: `pg_rewind`

(I haven't used this stuff...)

PostgreSQL 9.6 Optimizations

- Avoid re-vacuuming pages that only contain “frozen” data
- Trigger kernel data writeback *early* to reduce in-memory accumulation of dirty pages that will need to be written later
- Perform checkpoint writes in sorted order (helps lots on *spinning rust*)
- Use foreign key relationships to infer selectivity on join predicates
- Improve sorting performance via batched use of qsort
- Various improvements to GIN, GiST indexes
- Extend relations by several blocks at a time when under contention
- Improve sorting of text where text values are repeated
- Bloom indexes - lossy+compact, good when many columns are query criteria

PostgreSQL 9.6 Administrative

- `pg_stat_activity` now offers clues as to what blocked queries are waiting for
- `pg_config` view allows querying environment variable data used at `./configure` time
- Bsd authentication allows using OpenBSD authentication service
- Put better detail on authentication failures into postmaster log
- Support RADIUS passwords up to 128 bytes long
- `--with-systemd` improves usability of PostgreSQL with SystemD

PostgreSQL 9.6 Developer Visible Improvements

- COPY can receive (and dump, let's say, to disk) the results of INSERT/UPDATE/DELETE ... RETURNING query
- Improvements to Full Text Search (shhhh... Myles' part...)
- Generic Command Progress reporting
Allows writing admin tools that get a clue as to what the server is busy on

PostgreSQL 9.6 Replication Enhancements

- Support having multiple synchronous replicas; can configure to require n of them report success before the master can acknowledge COMMIT
- Support requiring that transactions are applied on replica(s) before the master can acknowledge COMMIT
- Getting closer to having “Logical Replication”
 - Worker process reads changes from WAL
 - This replaces the way Slony, Londiste use triggers to capture changes
 - Some of the infrastructure is in place, but it's not yet “production ready”...