



PostgreSQL v11, v12

Teodor Sigaev,
Oleg Bartunov,
Alexander Korotkov

FOSDEM PostgreSQL Developer Meeting 2018



4.46	JSON data handling in SQL.	174
4.46.1	Introduction.	174
4.46.2	Implied JSON data model.	175
4.46.3	SQL/JSON data model.	176
4.46.4	SQL/JSON functions.	177
4.46.5	Overview of SQL/JSON path language.	178
5	Lexical elements.	181
5.1	<SQL terminal character>.	181
5.2	<token> and <separator>.	185

SQL/JSON

PostgreSQL implementation (1 year of development)

- Uses native data types JSON, JSONB
- JSONPATH data type for SQL/JSON path language
- Nine functions SQL/JSON functions for constructing:
 - JSON_OBJECT, JSON_ARRAY,
JSON_OBJECTAGG, JSON_ARRAYAGG
- and retrieving
 - JSON_VALUE, JSON_QUERY, JSON_TABLE,
IS [NOT] JSON, JSON_EXISTS
- Extensions: more methods, JSONB op JSONPATH

SQL/JSON

```
{
  "address": {
    "city": "Moscow",
    "street": "Ulyanova, 7A"
  },
  "lift": false,
  "floor": [
    {
      "level": 1,
      "apt": [
        {"no": 1, "area": 40, "rooms": 1},
        {"no": 2, "area": 80, "rooms": 3},
        {"no": 3, "area": 50, "rooms": 2}
      ]
    },
    {
      "level": 2,
      "apt": [
        {"no": 4, "area": 100, "rooms": 3},
        {"no": 5, "area": 60, "rooms": 2}
      ]
    }
  ]
}
```

```

SELECT apt.*
FROM
  house,
  JSON_TABLE(js, '$.floor[0, 1]' COLUMNS (
    level int,
    NESTED PATH '$.apt[1 to last]' COLUMNS (
      no int,
      area int,
      rooms int
    )
  )) apt;

```

level	no	area	rooms
1	2	80	3
1	3	50	2
2	5	60	2

(3 rows)

SQL/JSON

- GIN indexing
- JSONBC
- Generic/JSONB subscripting
`UPDATE test_table SET JS['a']['b'] = '2'::jsonb;`

Nikita Glukhov, Fedor Sigaev,
Alexander Korotkov, Oleg Bartunov,
Ildus Kurbangaliev, Dmitry Dolgov

Параллелизм

- Ускорение parallel seqscan 2x
David Rowley
- Parallel hash join
Thomas Munro
- Parallel Append implementation
Amit Khandekar
- Parallel tuplesort (up to 3x)
Peter Geoghegan, Rushabh Lathia

Партиционирование

- Advanced partition matching algorithm (join)
Ashutosh Bapat
- Cross-partition update
Amit Khandekar
- Hash partitioning
Amul Sul
- Constraint exclusion on partitioned tables
Ashutosh Bapat
- Local indexes for partitioned table
Álvaro Herrera, Maksim Milyutin

Процедуры

- SQL procedures
Peter Eisentraut
- Transactions in PL/pgSQL, PL/Perl, PL/Python,
PL/Tcl
Peter Eisentraut

Мелкое, но полезное

- toast_tuple_target
Simon Riggs
- VACUUM table1, table2 ...
Nathan Bossart
- Atomics ops optimization
Yura Sokolov
- Compress method for SP-GiST & polygons
Heikki Linnakangas, Alexander Korotkov,
Fedor Sigaev, Nikita Glukhov

Мелкое, но полезное

- Initdb —with-wal-segsize
Beena Emerson
- Default namespace for Xpath
Pavel Stehule
- citext_pattern_ops
Alexey Chernyshov
- pg_prewarm
Mithun Cy, Robert Haas
- hook for session
Fabrízio de Royes Mello, Yugo Nagata

Мелкое, но полезное

- Pushdown limit
Douglas Doole
- Bitmap index scan only count
Alexander Kuzmenkov
- Distinguish selectivity of $<$ from \leq & $>$ from \geq
Tom Lane
- Постраничная оптимизация hash-индекса
Ashutosh Sharma
- Производительность агрегатов
Andres Freund

Мелкое, но полезное

- Wait for LSN
Ivan Kartashov
- Improve compactify_tuples and PageRepairFragmentation
Yura Sokolov
- Fix LWLock degradation on NUMA
Yura Sokolov
- Lazy hash table for snapshot's "xid in progress"
Yura Sokolov

Pgbench/psql

- Log, pow, <, = WHEN/CASE/THEN/END, etc
Fabien Coelho, Raúl Marín
- random_zipfian
Alik Khilazhev
- Alternative commands quit and exit for psql
Everaldo Canuto
- Улучшения на очереди

В работе

- allow usage of more than 1GB of work mem
Claudio Freire
- Convert join OR clauses into UNION queries
Tom Lane
- Incremental sort
Alexander Korotkov
- recheck_on_update functional indexes
Konstantin Knizhnik
- Partition-wise aggregation/grouping
Jeevan Chalke

В работе

- Generic WAL compression
Oleg Ivanov
- Range Merge join
Jeff Janes
- Precalculated stable functions
Alexander Korotkov, Arthur Zakirov,
Marina Polyakova
- Jsonb transform for PL/Python and PL/Perl
Anthony Bykov
- 64-bit transaction identifiers
Alexander Korotkov

В работе

- Logical decoding of two-phase transactions
Stas Kelvich, Nikhil Sontakke
- Predicate locking in GIST/GIN/HASH index
Shubham Barai
- Custom compression methods
Ildus Kurbangaliev
- Covering B-tree indexes (aka INCLUDE)
Anastasia Lubennikova
- Delete pages during GiST VACUUM
Constantine Kuznetsov, Andrey Borodin

В работе

- Tsquery, around, configuration
Victor Drobny, Aleksandr Parfenov
- Shared ispell dictionaries
Artur Zakirov
- Merge join on non-equal operators
Alexender Kuzmenkov
- Global indexes (unique too)
Alvaro Herrera, Ildar Musin
- Generated columns
Peter Eisentraut
- Parameters of opclasses
Nikita Glukhov

Будущее, но уже в работе

- Pluggable Storage API
Alexander Korotkov, Alvaro Herrera,
- JIT
Andres Freund
- Merge
Simon Riggs
- Threading, pooling
Konstantin Knizhnik
- TM API to community
Konstantin Knizhnik, Stas Kelvich
- Autoprepere
Konstantin Knizhnik



Спасибо за внимание!

teodor@postgrespro.ru

o.bartunov@postgrespro.ru

a.korotkov@postgrespro.ru