

BEUTH HOCHSCHULE FÜR TECHNIK BERLIN University of Applied Sciences



#### **PRO & CONTRA**

#### Felix Kunde and Petra Sauer



BEUTH HOCHSCHULE FÜR TECHNIK BERLIN

University of Applied Sciences







wird aus dem Europäischen Fonds für Regionale Entwicklung (EFRE) gefördert.

In the GIS scene more and more users are interested in data curation. Data creation has been too cost-intensive to just throw it away. Plus, a view into the past is desirable to measure effects of location-based decision making

## MOTIVATION

- Track all data changes in the database
- Revisit previous data versions
- Undo changes of certain write operations
- Work against multiple branches of a database

### HOW TO AUDIT? IN A RELATIONAL WAY



## GOOD ...

, **BUT** ...

- Easy to setup
- Easy to query the past
- Does not break apps

- What to do with DDL changes?
- Rely on timestamp fields?
- Store complete tuples?

ТооІ	Method	Log type	Revision
timetravel	Audit Trail	Extra Columns	Timestamps
temporal_tables	Audit Trail	Shadow Tables	Timestamps
table_version	Audit Trail	Shadow Tables	UD revision
table_log	Audit Trail	Shadow Tables	Trigger seq
audit_trigger	Audit Trail	Generic (hstore)	Transactions
pgMemento	Audit Trail	Generic (jsonb)	Transactions
CyanAudit	Audit Trail	Generic (pivot)	Transactions
pgVersion	Version Control	Extra Columns	UD revision
QGIS Versioning	Version Control	Extra Columns	UD revision
GeoGig	Version Control	External (binary)	UD revision
Flyway	Migration	External (SQL)	UD revision
Liquibase	Migration	External (XML)	UD revision

FOSS4G 2017 Talk: How to version my spatial database? > http://slides.com/fxku/foss4g17\_dbversion

# WHY AUDIT IN JSONB?

- Stop care about DDL changes
- Can be indexed, so queries are fast
- Store your JSON somewhere else
- \*Everybody loves JSONB!!\*

# pgmemens

## WHAT IS DIFFERENT?

- Relies on transaction IDs, not timestamps
- Stores only deltas in the logs
- Has a powerful undo feature

## LOG TABLES



- When using JSONB everything could be stored in one log table
- pgMemento stores transaction and table event metadata in separate tables to facilitate the lookup for historic actions
- Less redundancy vs. higher logging overhead

## DML-AUDITING

more columns) audit\_id id type geom POINT(1 1)2 1 txid = 2800000car SET type = 'moto' 2 bike POINT(12)23 3 train POINT(2 2) 42 insert BEFORE statement-level trigger transaction\_id op\_id Table\_operation table\_relid id txid id ... ... ... ••• ... ... ... ... 10 2800000 50 2800000 2005030 UPDATE 4 • • • TRANSACTION LOG TABLE\_EVENT\_LOG insert

Surrogate Key

(as PK can cover

## DML-AUDITING



## **DDL-AUDITING**



## **DDL-AUDITING**



## FAST QUERIES POWERED BY GIN INDEX

*For which transactions column 'type'* exists in the logs?

```
SELECT DISTINCT
  e.transaction_id
FROM
  pgmemento.table_event_log e
JOIN
  pgmemento.row_log r
  ON r.event_id = e.id
WHERE
  r.audit_id = 23
AND (r.changes ? 'type');
```

Which tuples once contained certain combinations of key(s) and value(s)?

SELECT DISTINCT
 audit\_id
FROM
 pgmemento.row\_log
WHERE
 changes @> '{"type": "bike"}'::jsonb;

Imagine the whole tuple is stored every time. More overhead on data processing.

## TIME FOR PAIN

### RESTORE PREVIOUS VERSIONS OF TUPLES

## THE LIFE OF A TUPLE

	id	name	type	Status	Audit_id
BIRTH			NULL		10
GROWING UP				{"status": <i>"</i> hello world"}	10
MARRIAGE		{"name":"foo",		"status":"alone"}	10
WISDOM			{"type":"phd"}		10
DEATH	{Id:1,	"name":"bar",	"type":"prof",	"status":"happy"}	10

### HOW WAS LIFE BEFORE MARRIAGE?

	id	name	type	Status	Audit_id
BIRTH			NULL		10
GROWING UP Uh oh, where is the log?! MARRIAGE				{"status": <i>"</i> hello world"}	10
		{"name":"foo" <i>,</i>		"status":"alone"}	10
WISDOM			{"type":"phd"}		10
DEATH	{Id:1,	"name":"bar",	"type":"prof",	"status":"happy"}	10

### STRATEGY 1: ROLLING BACK

- Concat all JSONB logs in reverse order
- Starting from recent state (or delete event) until requested point in time
- JSONB trick: duplicate keys get overwritten
- Too much overhead, if history is long

### STRATEGY 2: JSONB QUERIES

- Check audit\_column\_log
- For each column find the first entry in logs after requested point in time
- Feed result to jsonb\_build\_object
- Produces giant queries, but still quite fast

## RESTORE – PART 3

```
SELECT
  p.*
FROM
  generate_log_entries(1,2800000,'my_table') entries
LATERAL (
  SELECT
    *
  FROM
    jsonb_populate_record(
                                         Works only with a template. Could be the
        null::my table,
                                          actual table, but to be correct in case of
        entries
                                          any DDL changes, a temporary template
                                         can be created on the fly with information
) p;
                                         from audit_column_log.
```





- ... why would you use JSONB anyway then?
- Relational is faster and easier
- Takes up more space on your disk
- Your app might not like it

## REVERT

- Query all changes and referenced events for a given txid (or range of txids) in reverse order
- Loop over result set and perform the opposite event
- Consider dependencies between tables in order to avoid foreign key violations

## REVERT

op_id	Event	Reverse Event	Log Content
1	CREATE TABLE	DROP TABLE	-
2	ALTER TABLE ADD COLUMN	ALTER TABLE DROP COLUMN	-
3	INSERT	DELETE	NULL
4	UPDATE	UPDATE	Changed fields of changed rows
5	ALTER TABLE ALTER COLUMN	ALTER TABLE ALTER COLUMN	All rows of altered columns
6	ALTER TABLE DROP COLUMN	ALTER TABLE ADD COLUMN	All rows of deleted columns
7	DELETE	INSERT	All fields of deleted rows
8	TRUNCATE	INSERT	All fields of table
9	DROP TABLE	CREATE TABLE	All fields of table (logged as truncate)

# JSONB > QUERIES

- Updates are pretty easy to setup (ok we could produce deltas also here and not during trigger phase)
- Could lead to a branching concept ...
- What about long running transactions? (typical GIS workflows – many edits, commit once)

## TO DOs

- Branching concept
- Log tables for more DB objects
- Extending the test suite
- Maybe: Logical decoding instead of triggers

## PERFORMANCE



Kunde F., Sauer P. (2017) pgMemento – A Generic Transaction-Based Audit Trail for Spatial Databases. In: Gertz M. et al. (eds) Advances in Spatial and Temporal Databases. SSTD 2017. Lecture Notes in Computer Science, vol 10411. Springer, Cham

## TECHNICAL DETAILS

- Written entirely in PL/pgSQL
- Requires at least PostgreSQL 9.5
- Repo: github.com/pgmemento
- LGPL v3 Licence



#### BEUTH HOCHSCHULE FÜR TECHNIK BERLIN University of Applied Sciences





#### Felix Kunde

fkunde[at]beuth-hochschule.de

#### **Petra Sauer**

Sauer[at]beuth-hochschule.de





Bundesministerium für Wirtschaft und Energie

