

RLS Timetravel

- Possible alternative to "history" tables
- Use RLS to filter based on point in time
- Use PG10 partitioning to keep history separate from current



RLS Timetravel

```
CREATE TABLE timetravel
(
  id int8,
  f1 text not null,
  tr tstzrange not null default tstzrange(clock_timestamp(), 'infinity', '[]')
) PARTITION BY RANGE (upper(tr));
```

```
CREATE TABLE timetravel_current PARTITION OF timetravel
(
  primary key (id, tr) DEFERRABLE
) FOR VALUES FROM ('infinity') TO (MAXVALUE);
CREATE INDEX timetravel_current_tr_idx ON timetravel_current USING GIST (tr);
```

```
CREATE TABLE timetravel_history PARTITION OF timetravel
(
  primary key (id, tr) DEFERRABLE
) FOR VALUES FROM (MINVALUE) TO ('infinity');
CREATE INDEX timetravel_history_tr_idx ON timetravel_history USING GIST (tr);
```



RLS Timetravel

```
CREATE OR REPLACE FUNCTION get_pit() RETURNS timestampz AS $$
  SELECT
    CASE WHEN current_setting('tt.cts', true) IS NULL OR
           current_setting('tt.cts', true) = '' THEN
      clock_timestamp()
    ELSE
      current_setting('tt.cts', true)::timestampz
    END
  $$ LANGUAGE sql;

-- only current rows are visible
-- unless tt.cts is defined, in which case get rows current as of that time
ALTER TABLE timetravel ENABLE ROW LEVEL SECURITY;
CREATE POLICY p1 ON timetravel
USING (tr @> get_pit())
WITH CHECK (/* NEW */tr @> clock_timestamp() OR
            /* OLD */ upper(tr) <= clock_timestamp());
```



RLS Timetravel

```
CREATE OR REPLACE FUNCTION modify_timetravel()
RETURNS TRIGGER AS $$
DECLARE
    ctr timestamptz := clock_timestamp();
BEGIN
    OLD.tr = tstzrange(lower(OLD.tr), ctr, '[]');
    INSERT INTO timetravel VALUES (OLD.*);
    IF (TG_OP = 'UPDATE') THEN
        NEW.tr = tstzrange(ctr, 'infinity', '[]');
        RETURN NEW;
    ELSIF (TG_OP = 'DELETE') THEN
        RETURN OLD;
    END IF;
END;
$$ LANGUAGE plpgsql;

CREATE TRIGGER timetravel_audit BEFORE DELETE OR UPDATE
ON timetravel_current FOR EACH ROW EXECUTE PROCEDURE modify_timetravel();
```



RLS Timetravel

```
GRANT ALL ON timetravel TO dba;  
SET SESSION AUTHORIZATION dba;
```

```
INSERT INTO timetravel(id, f1)  
SELECT g.i, 'row-' || g.i::text  
FROM generate_series(1,1000000) AS g(i);
```

```
RESET SESSION AUTHORIZATION;  
VACUUM FREEZE ANALYZE timetravel;  
SET SESSION AUTHORIZATION dba;
```



RLS Timetravel

-- NOTE: Point In Time (pit) is after update to id 42 and before delete of id 4242

```
UPDATE timetravel SET f1 = 'update number 1' WHERE id = 42 RETURNING lower(tr) AS pit \gset
```

```
DELETE FROM timetravel WHERE id = 4242;
```

```
SELECT * FROM timetravel WHERE id in (42, 4242);
```

id	f1	tr
42	update number 1	["2017-09-03 16:44:22.160023-07",infinity]

RLS Timetravel

```
UPDATE timetravel SET f1 = 'update number 2' WHERE id = 42;  
SELECT * FROM timetravel WHERE id = 42;
```

id	f1	tr
42	update number 2	["2017-09-03 16:44:24.206489-07",infinity]

```
UPDATE timetravel SET f1 = 'update number 3' WHERE id = 42;  
SELECT * FROM timetravel WHERE id = 42;
```

id	f1	tr
42	update number 3	["2017-09-03 16:44:25.270849-07",infinity]

RLS Timetravel

```
SELECT set_config('tt.cts', :'pit', false);  
       set_config
```

```
-----  
2017-09-03 16:44:22.160023-07
```

```
SELECT * FROM timetravel WHERE id in (42, 4242) ORDER BY 1, 3;
```

id	f1	tr
42	update number 1	["2017-09-03 16:44:22.160023-07", "2017-09-03 16:44:24.206489-07"]
4242	row-4242	["2017-09-03 16:44:04.406047-07", "2017-09-03 16:44:23.173052-07"]

RLS Timetravel

```
RESET SESSION AUTHORIZATION;
```

```
SELECT * FROM timetravel WHERE id in (42, 4242) ORDER BY 1, 3;
```

id	f1	tr
42	row-42	["2017-09-03 16:44:04.34601-07", "2017-09-03 16:44:22.160023-07")
42	update number 1	["2017-09-03 16:44:22.160023-07", "2017-09-03 16:44:24.206489-07")
42	update number 2	["2017-09-03 16:44:24.206489-07", "2017-09-03 16:44:25.270849-07")
42	update number 3	["2017-09-03 16:44:25.270849-07", infinity]
4242	row-4242	["2017-09-03 16:44:04.406047-07", "2017-09-03 16:44:23.173052-07")

