

Monitoring PostgreSQL

PostgreSQL monitoring

- Nezbytná služba pro DBA k zajištění provozu
 - místo na svazcích
 - replikace
- Zabbix, HP Openview, ...
- Nástroj zpětné vazby pro vývojáře
 - neefektivní dotazy
 - velikost tabulek a indexů
 - počty transakcí
 - ...

Zdroje informací

- Operační systém
- Interní statistiky PostgreSQL a LOG soubor instance
- Rozšíření v contributed balíčcích
 - pg_stat_statements
- Ostatní rozšíření/extenze
 - pg_stat_kcache
- Externí nástroje
 - [check_postgres](#)
 - [pgBadger](#)

Data z operačního systému

- Processor
 - dedikovaný DB server ?
 - user, system, IO wait
- Paměť
 - used, buffered, cached, swap
- IO
 - reads, writes, Bytes, IO time
- síť

Kam s ním?

- Kam se všemi těmi daty?
- **InfluxDB & Grafana**
 - **Telegraf** jako agent pro sběr dat
 - InfluxDB vstupní metody
 - „řádkový protokol“
 - Json
 - **Python API**
 - InfluxDB výstupní metody
 - Poskytuje REST API
 - InfluxDB **Continuous Queries**
 - Historizace dat

Interní statistiky

- **Monitoring Database Activity** - dokumentace
 - Dynamické
 - Kumulativní
- Globální informace na úrovni instance
 - pg_stat_activity
 - pg_stat_bgwriter
 - pg_stat_database
 - ...
- Lokální informace v rámci připojené databáze
 - pg_stat_user_tables
 - ...

Konfigurace pro monitoring

- CPU - rychlost dotazu na aktuální čas
- pg_test_timing

```
postgres$ /usr/pgsql-9.4/bin/pg_test_timing
Testing timing overhead for 3 seconds.
Per loop time including overhead: 41.79 nsec
Histogram of timing durations:
< usec    % of total    count
   1      95.84044   68798574
   2       4.15628   2983562
   4       0.00242     1734
   8       0.00046      328
  16       0.00037      263
  32       0.00000        3
  64       0.00000        0
 128       0.00000        1
 256       0.00000        1
 512       0.00003      25
```

Konfigurace pro monitoring

```
# - Kernel Resource Usage -
shared_preload_libraries = 'pg_stat_statements'

# - When to Log -
log_min_duration_statement = 100

# - What to Log -
log_checkpoints = on
log_connections = on
log_disconnections = on
log_line_prefix = '%t %p %c %l %r %d %u %x %v %a %i:'
log_lock_waits = on
log_statement = 'ddl'
log_temp_files = 2048

# - Query/Index Statistics Collector -
track_activities = on # information on the currently executing command
track_counts = on # statistics on database activity - potrebuje autovacuum
track_io_timing = on # timing of database I/O calls
track_functions = all
track_activity_query_size = 8192 # default 1024
# AUTOVACUUM PARAMETERS

log_autovacuum_min_duration = 250
```


Nastavení pro perf. test analýzu

- Pro běžný provoz by byl log rychle narůstal
- pg_badger „perfctest“ nastavení

```
# - What to Log -  
log_duration = on  
log_min_duration_statement = 0  
log_statement = 'all'  
log_temp_files = 0  
  
log_autovacuum_min_duration = 0
```

pgbadger

- parsing logů, generuje HTML report
 - podporuje **inkrementální** zpracování
 - paralelizace
 - čtení logů přes ssh
 - zpracuje komprimované gzip logy (logrotate)
 - reporty za časové okno (v týdnech)

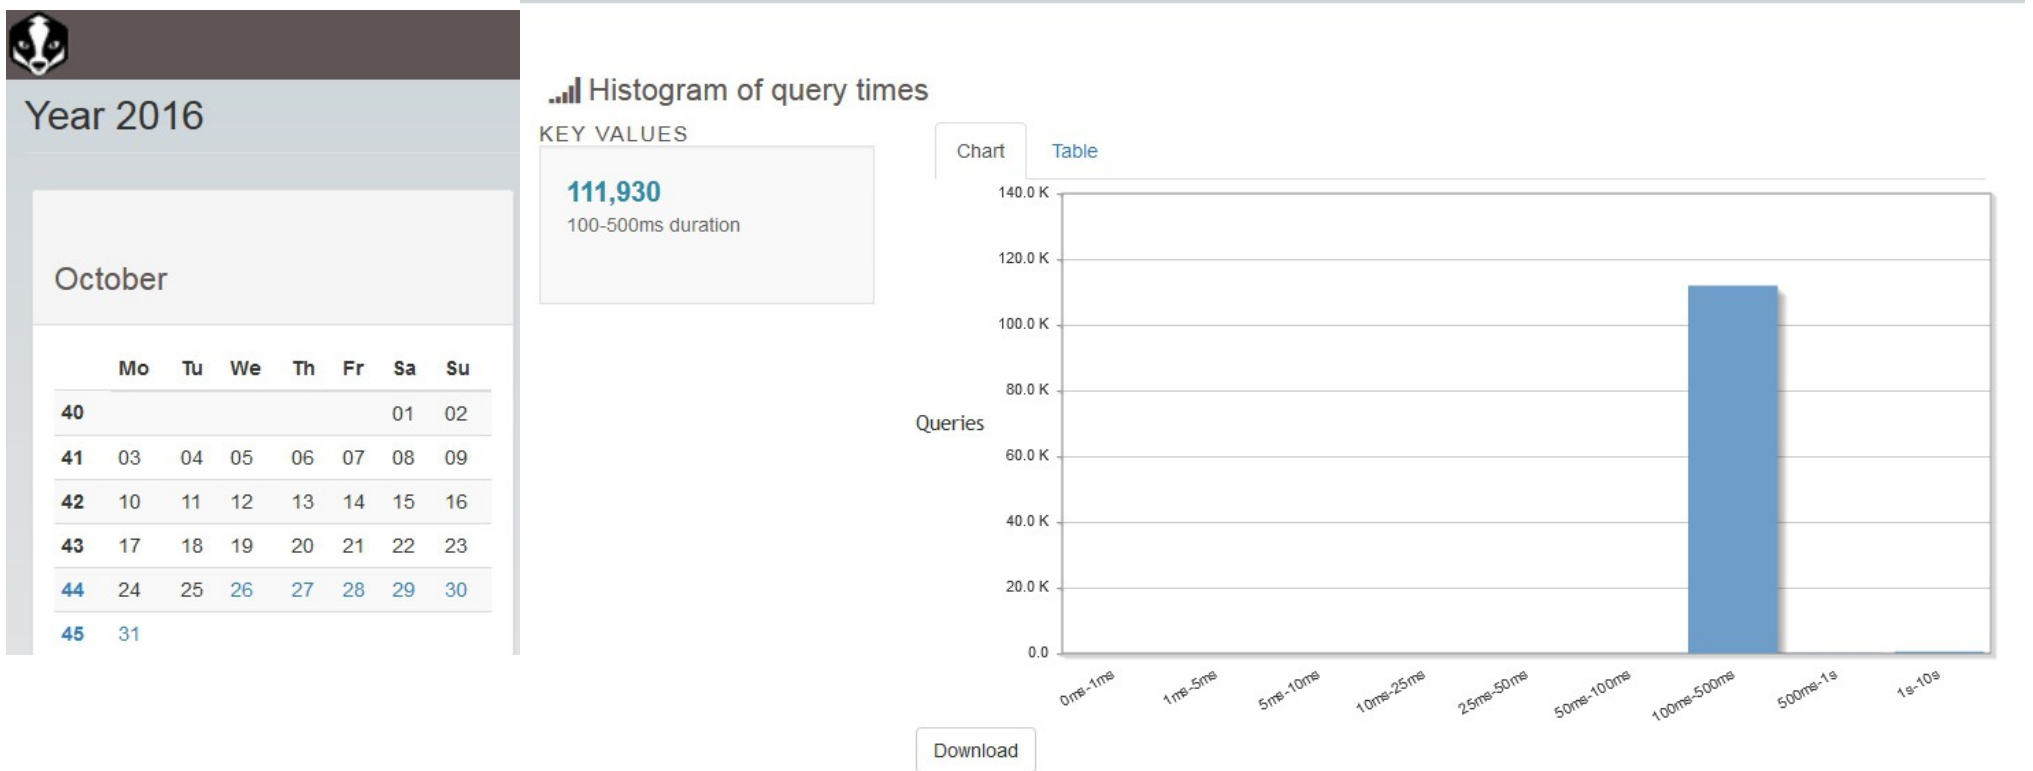
```
pgbadger -I -R 4 -j 2 -J 4 -f stderr -start-monday --ssh-option -q \  
-r cosi.kdesi.net --prefix %t %p %c %l %r %d %u %x %v %a %i: \  
/var/lib/pgsql/data/pg_log/postgresql.log*.gz \  
-o /var/lib/pgreports/my_cluster_name
```

pgbadger

- Přehledný [manuál](#)
- denně generované reporty jsou přístupné vývojářům a aplikační podpoře přes web
- Kontrola po nasazení nové verze, či revize průběhu výkonnostních testů

pbadger – query times

Pozor na `log_min_duration_statement` histogram pak může na první pohled vypadat zvláštně – pro aplikační podporu či vývojáře je to rychlá informace, zda se neobjevil podezřele pomalý dotaz.



pgbadger – sql traffic



pgBadger

Overview ▾

Connections ▾

Sessions ▾

Checkpoints ▾

Temp Files ▾

Vacuums ▾

Locks ▾

Queries ▾

Top ▾

Events ▾



SQL Traffic


KEY VALUES

41 queries/s

Query Peak

2016-10-31 03:22:34

Date



Pozor na `log_min_duration_statement`.

pgbadger – pomalé dotazy

🌀 Slowest individual queries

Rank Duration Query

1 1m56s `DELETE FROM XXXXXXXXXX WHERE ctid IN (SELECT ctid FROM XXXXXXXXXX WHERE created < '2016-08-01 05:01:41.366' LIMIT 10000);`

[Date: 2016-10-31 05:03:37 - Database: ~~XXXXXXXXXX~~ - User: ~~XXXXXXXXXX~~_app - Remote: 10.~~XXXXXXXXXX~~.~~XXXXXXXXXX~~.~~XXXXXXXXXX~~ - Application: [unknown] - Bind query: yes]

2 1m9s `DELETE FROM XXXXXXXXXX WHERE ctid IN (SELECT ctid FROM XXXXXXXXXX WHERE created < '2016-08-01 20:00:59.582' LIMIT 10000);`

[Date: 2016-10-31 20:02:08 - Database: ~~XXXXXXXXXX~~ - User: ~~XXXXXXXXXX~~_app - Remote: 10.~~XXXXXXXXXX~~.~~XXXXXXXXXX~~.~~XXXXXXXXXX~~ - Application: [unknown] - Bind query: yes]

📈 Size of temporary files

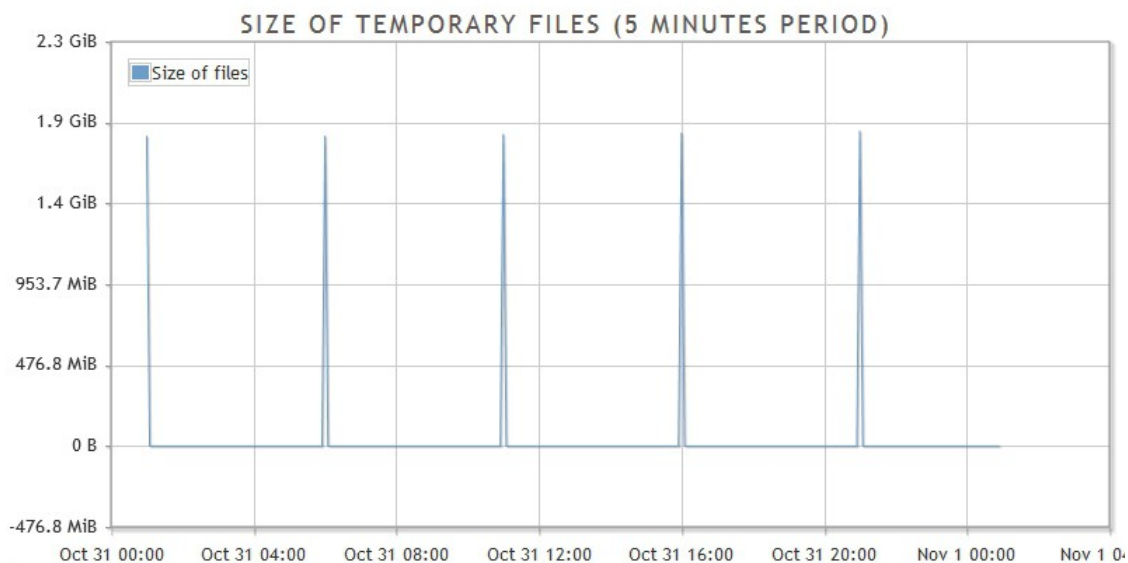
KEY VALUES

601.54 MiB

Temp Files size Peak

2016-10-31 20:02:08

Date



Download

pgbadger - deadlock

2

125

Details

```
LOG: process ... still waiting for ShareLock on transaction ... after ... ms
```

Examples

3

56

Details

```
LOG: process ... still waiting for ExclusiveLock on tuple (...) of relation ... of database ... after ... ms
```

Examples

4

27

Details

```
ERROR: deadlock detected
```

Examples

ERROR: deadlock detected

Detail: Process 16851 waits for ShareLock on transaction 144602509; blocked by process 35447. Process 35447 waits for ShareLock on transaction 144676444; blocked by process 16851. Process 16851: DELETE FROM J WHERE id = \$1 AND # = \$2 Process 35447: UPDATE jms_audit_data SET created_by = NULL, modified = CURRENT_TIMESTAMP WHERE c IS NOT NULL AND CURRENT_TIMESTAMP - (\$1 * '1 minute' :: INTERVAL) > next_process

Context: while deleting tuple (5,36) in relation "jms_audit_data"

Hint: See server log for query details.

Statement: DELETE FROM J WHERE id = \$1 AND # = \$2

Date: 2016-11-16 03:35:04

Database: ~~XXXXXXXXXX~~

Application: [unknown]

User: ~~XXXXXXXXXX~~

Remote: ~~192.168.1.100~~

Statistické pohledy




- aktuální stav (Dynamic Statistics Views)
 - pg_stat_activity
 - pg_stat_replication
 - pg_stat_ssl
- Kumulativní (Collected Statistics Views) – v rozsahu instance
 - Lze vynulovat funkcí pg_stat_clear_snapshot()
 - pg_stat_database
 - ...
- Kumulativní – v rozsahu připojené databáze
 - pg_stat_all_tables
 - ...

Instance: pg_stat_activity

- Přehled aktuálně připojených sessions (backend procesy)
 - **datname** – kam je proces připojen
 - `age(now(), query_start)` – trvání dotazu
 - `age(now(), xact_start)` – trvání transakce
 - **state** – stav session (active, idle, idle in transaction...)
 - **waiting** (změna v 9.6 – `wait_event_type`, `wait_event`)

pg_stat_activity

- idle in transaction
- využití max_sessions
- čekající sessions



Instance: pg_stat_archiver

- **archived_count**

- last_archived_time

- age(now() - last_archived_time) -- interval
- extract(epoch from age(now(), last_archived_time))
as arch_age_sec -- integer

- **failed_count**

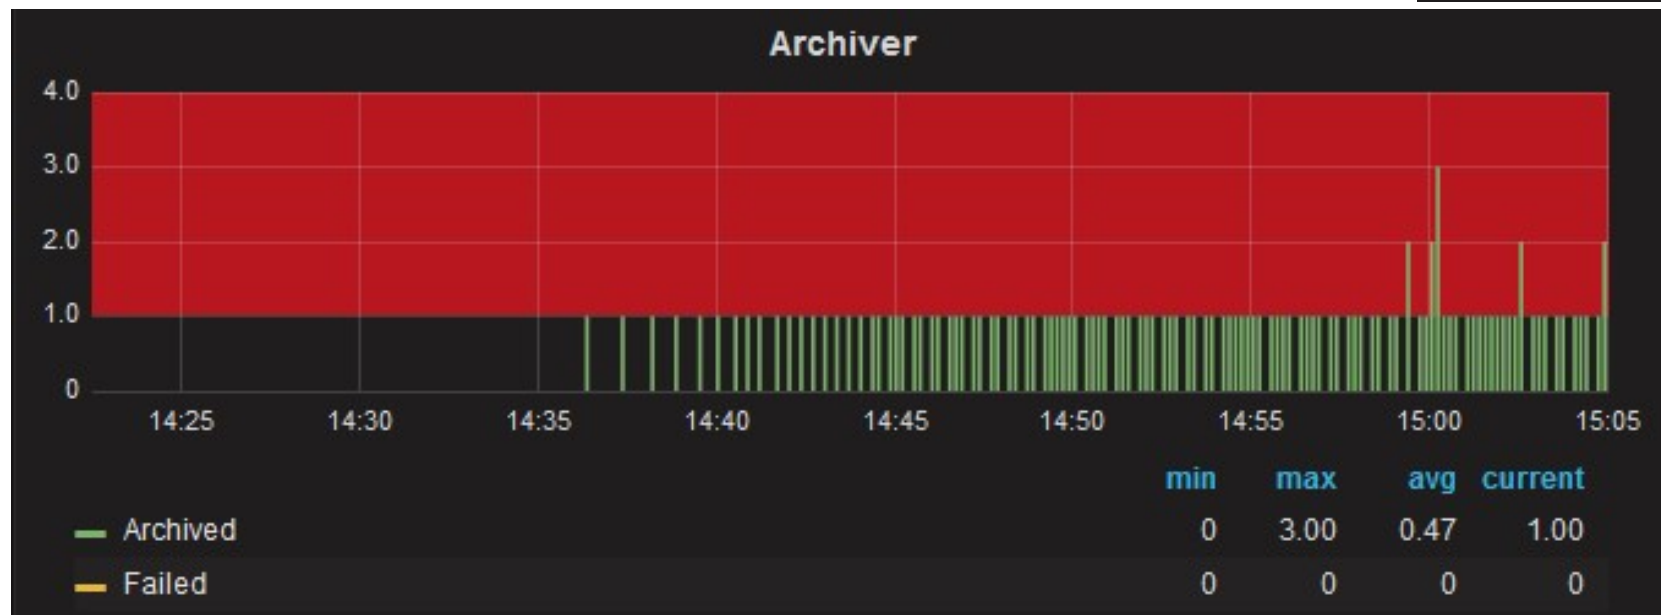

- last_failed_time

- Lze přidat objem transakčních logů

```
select pg_xlog_location_diff(pg_current_xlog_location(),  
'0/00000000'::pg_lsn) as xlog_volume;
```

pg_stat_archiver

- četnost přepínání xlog segmentů
 - Vizuální kontrola zda je nastaven `archive_timeout`
- Četnost selhání `archive_command`



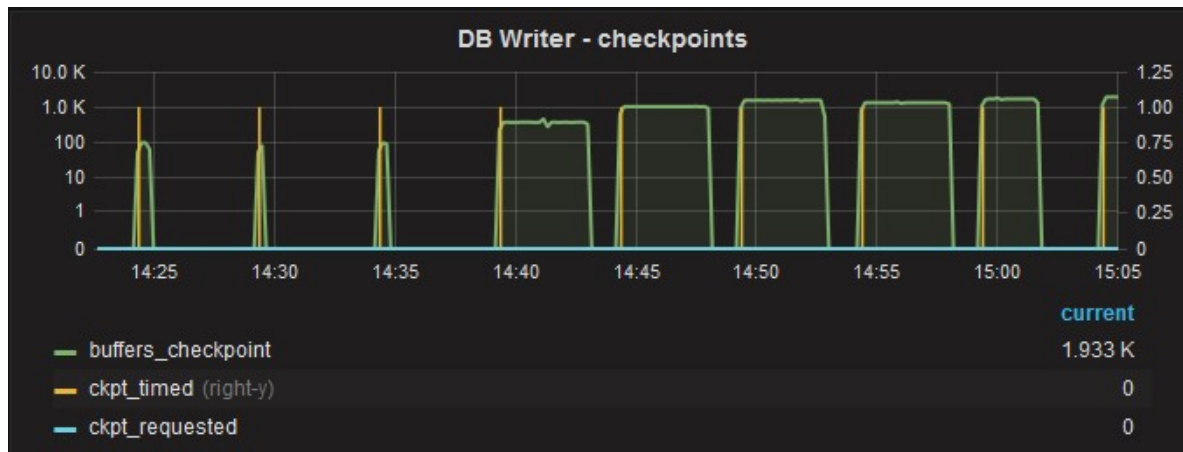


Instance: pg_stat_bgwriter

- **checkpoints_timed**
- **checkpoints_req**
- **buffers_checkpoint**
- **buffers_clean** – buffery zapsané bgwriter procesem
- **buffers_backend** – buffery zapsané backend procesy přímo
- **buffers_alloc** – alokované buffery

pg_stat_bgwriter

- Vizuální kontrola checkpoint-ů a checkpoint_completion_target

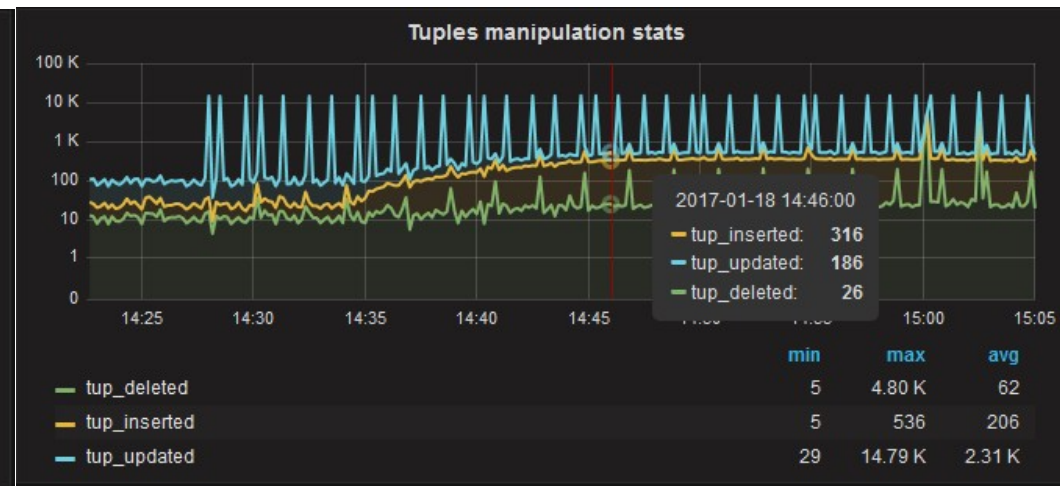





Instance: pg_stat_database

- jeden řádek pro každou DB, kromě numbackends jsou statistiky kumulativní
- datname
- **numbackends** – aktuální hodnota
- **xact_commit, xact_rollback** – kumulativní sledovat poměr commit/rollback
- **blks_read, blks_hit** – efektivita buffer cache
- temp_files, temp_bytes
- **deadlocks**
- tup_% statistiky pro řádky (return, fetch, ins, udp, del)
- blk_read_time, blk_write_time – čas strávený backendy na IO – efektivita cache na úrovni OS

pg_stat_database

- úspěšnost buffer cache
- DML statistiky
- transakce



DB: pg_stat_all_tables

- pg_stat_sys_tables
- pg_stat_user_tables
- seq_scan,
seq_tup_read/seq_scan
 - Kolikrát se tabulka četla a **kolik řádek bylo vráceno na jedno čtení** – nechybí index ?
- n_tup_upd/n_tup_hot_upd
 - Potenciální kandidát na změnu FILLFACTOR
- schemaname,
relname...
- **seq_scan**
- seq_tup_read
- **idx_scan**
- idx_tup_fetch
- n_tup_upd
- n_tup_hot_upd
- autovacuum_count
- autoanalyze_count

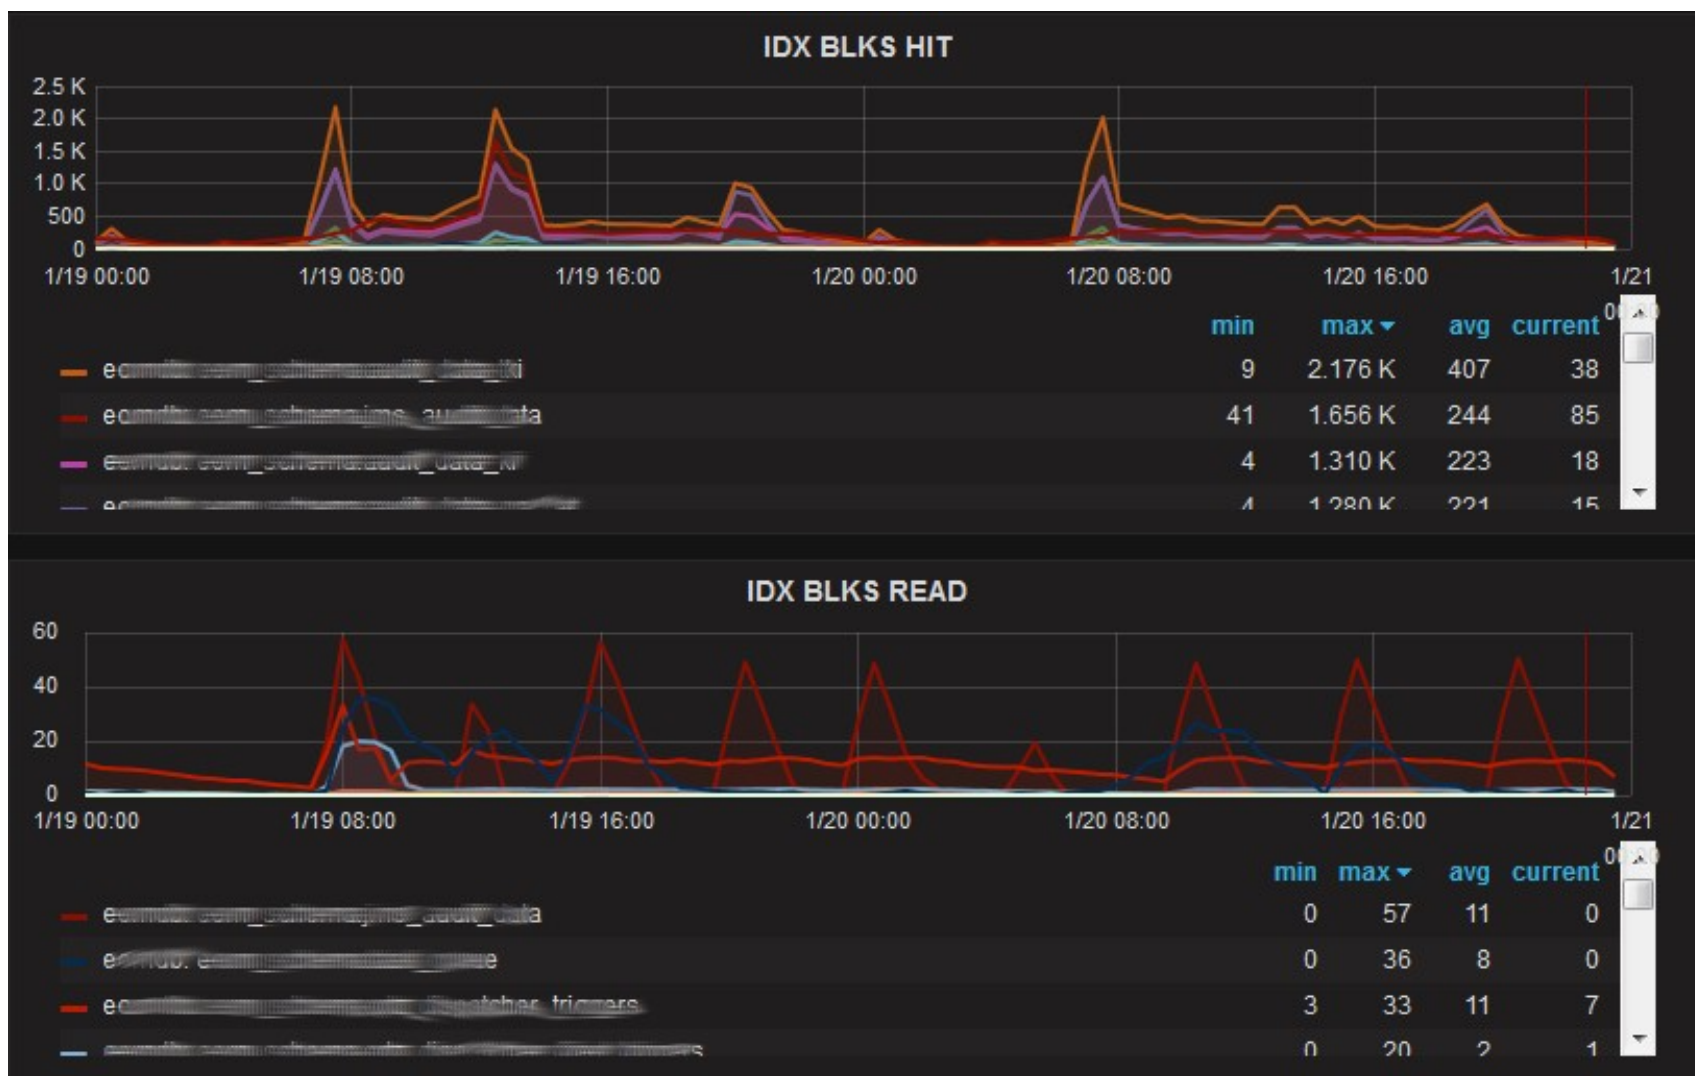
pg_stat_user_tables



DB: pg_stat_all_indexes

- pg_stat_sys_indexes
- pg_stat_user_indexes
- idx_scan
 - používá se index ?
 - Pozor na časový úsek, za který data vyhodnocujeme (měsíční zpracování..)
- schemaname,
relname,
indexrelname...
- **idx_scan**
- idx_tup_read
- idx_tup_fetch

pg_stat_user_indexes



DB: pg_statio_all_tables

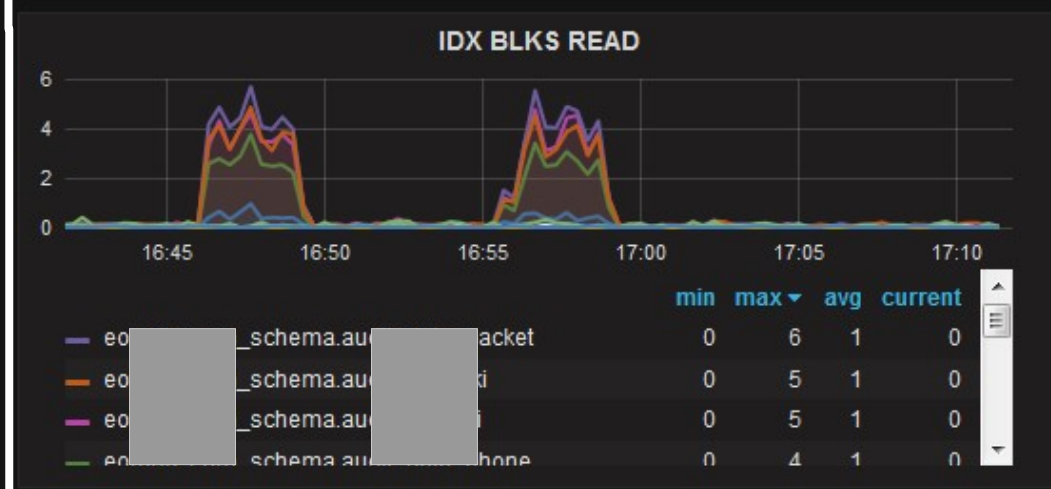

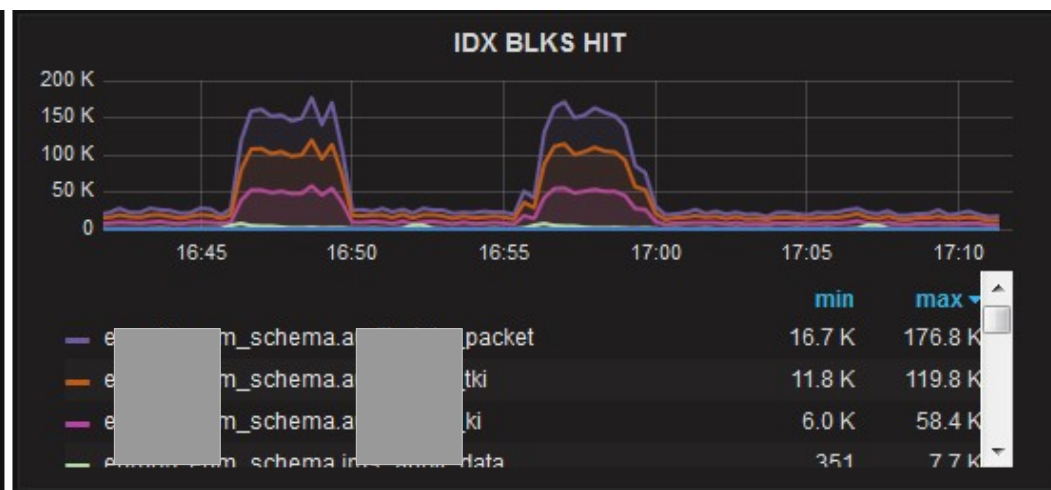

- pg_statio_sys_tables
- pg_statio_user_tables

$\frac{\text{heap_blks_hit}}{\text{heap_blks_hit} + \text{heap_blks_read}}$

- efektivita buffer cache
- fyzické čtení nemusí být problém, pokud dobře funguje cache OS, viz `pg_stat_database.blk_read_time`

- schemaname, relname
- heap_blks_read
- heap_blks_hit
- **idx_blks_read**
- **idx_blks_hit**
- Toast_..., tidx_...

pg_statio_user_tables



DB: pg_statio_all_indexes

- pg_statio_sys_indexes
- pg_statio_user_indexes

idx_blks_hit/
(idx_blks_hit+idx_blks_read)

- efektivita buffer cache
- Informace na úrovni konkrétních indexů

- schemaname,
relname,
indexrelname
- **idx_blks_read**
- **idx_blks_hit**

DB: pg_statio_all_sequences

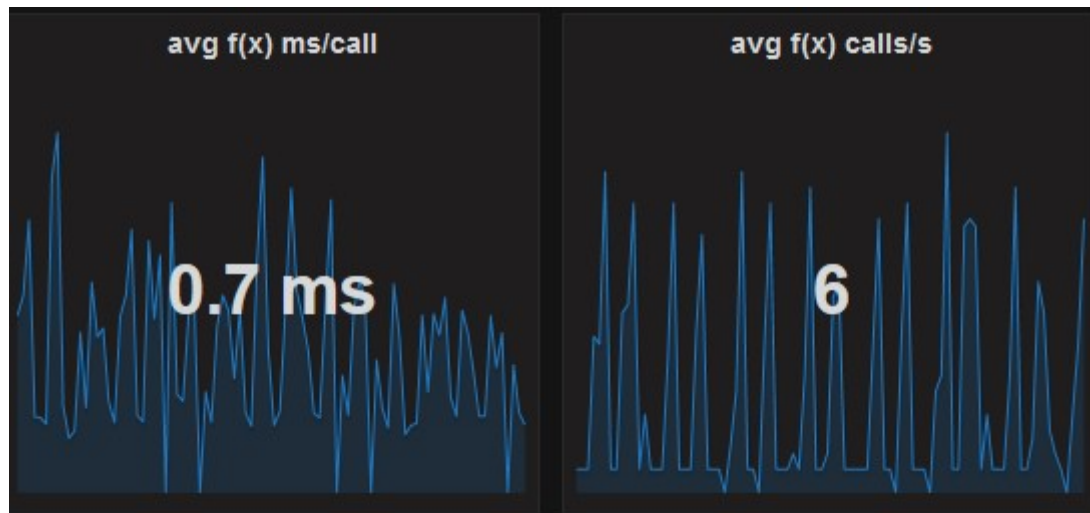
- pg_statio_sys_sequences
- pg_statio_user_sequences

- Nezaznamenali jsme žádný problém či důvod k systematickému sledování

- rename, schemaname, blks_read, blks_hit

DB: pg_stat_user_functions

- Zjištění prodlužujících se časů na jednotlivé volání
- ne / lineární vztah k objemu dat
- **calls**
- **total_time**
- **self_time**



instance: pg_stat_progress_vacuum

- od verze 9.6
- online progress reporting
- neukládáme do monitorovací DB

Vlastní dotazy


- velikost databází
- velikost jednotlivých relací (tabulky, indexy, materializované pohledy, TOAST tabulky)
- blokující session
- objem WAL záznamů – přidali jsme k archiveru
- ne / povolené autovacuum nad tabulkami
- využití max_sessions
- autovacuum threshold podle počtu řádek tabulky
- ...

instance: DB size

```
SELECT d.datname,  
       CASE  
         WHEN pg_catalog.has_database_privilege(d.datname,  
         'CONNECT') THEN  
           pg_catalog.pg_database_size(d.datname)  
         ELSE  
           -1  
         END as size  
FROM pg_catalog.pg_database d;
```

datname	size
template0	6513156
postgres	6631452
hibernate	25325596
jackrabbit	22327324
quartz	7209476
template1	6521348

(6 rows)



DB: relation_size – 9.1

```
SELECT current_database() as datname, a.schemaname, a.relation_name,
a.relation_kind,
a.relation_persistence, a.row_estimate, a.total_bytes, a.index_bytes,
total_bytes-index_bytes AS relation_bytes
FROM
( SELECT nspname AS schemaname, relname AS relation_name ,
  (CASE
    WHEN c.relkind = 'r' THEN 'table'
    when c.relkind = 'i' then 'index'
    when c.relkind = 'm' then 'materialized view'
    when c.relkind = 't' then 'TOAST table'
    ELSE 'other'
  END) as relation_kind ,
  (case
    when c.relpersistence = 'p' then 'permanent'
    when c.relpersistence = 'u' then 'unlogged'
    when c.relpersistence = 't' then 'temporary'
  END) as relation_persistence,
  c.reltuples AS row_estimate,
  pg_total_relation_size(c.oid) AS total_bytes,
  pg_indexes_size(c.oid) AS index_bytes
FROM pg_class c LEFT JOIN
  pg_namespace n ON n.oid = c.relnamespace
WHERE relkind in ('r', 'i', 'm', 't' )
) a;
```

DB: relation_size – 9.2 – 9.6

```
SELECT current_database() as datname, a.schemaname, a.relation_name,
a.relation_kind,
a.relation_persistence, a.row_estimate, a.total_bytes, a.index_bytes,
a.toast_bytes, total_bytes-index_bytes-COALESCE(toast_bytes,0) AS
relation_bytes FROM
( SELECT nspname AS schemaname, relname AS relation_name ,
(CASE
WHEN c.relkind = 'r' THEN 'table'
when c.relkind = 'i' then 'index'
when c.relkind = 'm' then 'materialized view'
when c.relkind = 't' then 'TOAST table'
ELSE 'other'
END) as relation_kind ,
(case
when c.relpersistence = 'p' then 'permanent'
when c.relpersistence = 'u' then 'unlogged'
when c.relpersistence = 't' then 'temporary'
END) as relation_persistence,
c.reltuples AS row_estimate,
pg_total_relation_size(c.oid) AS total_bytes,
pg_indexes_size(c.oid) AS index_bytes ,
pg_total_relation_size(reltoastrelid) AS toast_bytes
FROM pg_class c LEFT JOIN
pg_namespace n ON n.oid = c.relnamespace
WHERE relkind in ('r', 'i', 'm', 't' )
) a;
```

instance: blocking sessions – 9.2 – 9.6

[Lock Monitoring](#) - vhodné příklady, přidat trvání blokování dotazu – threshold, grafy

– zdroj postgresql wiki...

```
SELECT a.datname AS db,  
       kl.pid AS blocking_pid,  
       ka.username AS blocking_user,  
       ka.query AS blocking_query,  
       bl.pid AS blocked_pid,  
       a.username AS blocked_user,  
       a.query AS blocked_query,  
       extract( epoch from age(now(), a.query_start)) as age_sec,  
       to_char(age(now(), a.query_start), 'HH24h:MIm:SSs'::text) AS age  
FROM pg_locks bl  
     JOIN pg_stat_activity a ON bl.pid = a.pid  
     JOIN pg_locks kl ON bl.locktype = kl.locktype AND NOT bl.database IS  
DISTINCT FROM kl.database AND NOT bl.relation IS DISTINCT FROM kl.relation  
AND NOT bl.page IS DISTINCT FROM kl.page AND NOT bl.tuple IS DISTINCT FROM  
kl.tuple AND NOT bl.virtualxid IS DISTINCT FROM kl.virtualxid AND NOT  
bl.transactionid IS DISTINCT FROM kl.transactionid AND NOT bl.classid IS  
DISTINCT FROM kl.classid AND NOT bl.objid IS DISTINCT FROM kl.objid AND  
NOT bl.objsubid IS DISTINCT FROM kl.objsubid AND bl.pid <> kl.pid  
     JOIN pg_stat_activity ka ON kl.pid = ka.pid  
WHERE kl.granted AND NOT bl.granted  
ORDER BY a.query_start;
```

DB: autovacuum=on

- Ve skutečnosti je zajímavé, zda některá tabulka nemá **off**

```
with relav as (  
    select cropt.oid, cropt.ropt[2]::boolean from  
        ( SELECT c.oid, string_to_array(unnest(c.reloptions), '=') as ropt  
          FROM pg_class c  
        ) cropt  
    where cropt.ropt[1] = 'autovacuum_enabled'  
)  
select current_database() as datname, nspname AS schemaname,  
       c.relname as table_name,  
       coalesce(relav.ropt, current_setting('autovacuum')::boolean)  
       as autovacuum_enabled  
from pg_class c  
     left join relav on c.oid = relav.oid  
     LEFT JOIN pg_namespace n ON n.oid = c.relnamespace  
WHERE c.relkind IN ('r', 'm', 't');
```

datname	schemaname	table_name	autovacuum_enabled
postgres	pg_catalog	pg_statistic	t
postgres	pg_catalog	pg_type	t
postgres	pg_catalog	pg_authid	t
postgres	pg_catalog	pg_proc	t
postgres	pg_catalog	pg_class	t

DB: autovacuum threshold

- Pro velké tabulky může být výchozí autovacuum_vacuum_scale_factor příliš vysoký

```
with relav as (  
select cropt.oid, cropt.ropt[2]::real from  
  ( SELECT c.oid, string_to_array(unnest(c.reloptions), '=') as  
ropt FROM pg_class c ) cropt  
  where cropt.ropt[1] = 'autovacuum_vacuum_scale_factor'  
)  
select current_database() as datname, nspname AS schemaname,  
  c.relname as table_name, c.reltuples::int as row_estimate,  
  coalesce(relav.ropt,  
current_setting('autovacuum_vacuum_scale_factor')::real) avsf,  
  current_setting('autovacuum_vacuum_threshold')::int +  
  c.reltuples*coalesce(relav.ropt,  
current_setting('autovacuum_vacuum_scale_factor')::real)  
  as av_tuples_threshold  
from pg_class c  
left join relav on c.oid = relav.oid  
LEFT JOIN pg_namespace n ON n.oid = c.relnamespace  
where --c.relkind in ('r', 'm', 't')  
c.relkind = ANY ('{r,m,t}'::char[])  
order by av_tuples_threshold desc;
```

DB: autovacuum threshold

- Alternativa téhož dotazu

```
WITH relopt AS (  
    select OID, (pg_options_to_table(reloptions)).option_name,  
           (pg_options_to_table(reloptions)).option_value  
    from pg_class c  
),  
relav as (  
    select ro.oid, ro.option_value::real as avsf from relopt ro where  
    ro.option_name = 'autovacuum_vacuum_scale_factor'  
)  
select current_database() as datname, nspname AS schemaname,  
       c.relname as table_name,  
       c.reltuples::int as row_estimate,  
       relav.avsf,  
       coalesce(relav.avsf,  
current_setting('autovacuum_vacuum_scale_factor')::real) avsf,  
       current_setting('autovacuum_vacuum_threshold')::int +  
       c.reltuples*coalesce(relav.avsf,  
current_setting('autovacuum_vacuum_scale_factor')::real)  
       as av_tuples_threshold  
from pg_class c  
left join relav on c.oid = relav.oid  
LEFT JOIN pg_namespace n ON n.oid = c.relnamespace  
where --c.relkind in ('r', 'm', 't')  
c.relkind = ANY ('{r,m,t}'::char[])  
order by av_tuples_threshold desc limit 5;
```

DB: autovacuum threshold

table_name	row_estimate	avsf	av_tuples_threshold
measurement_mid	43656168	0.2	8731284
measurement_high	491057280	0.01	4910622
measurement_low	7307218	0.2	1461494
blmeasurements	958449	0.2	191740
...			
ac_data	152	0.2	80
pg_aggregate	133	0.2	77
...			
pg_toast_2606	0	0.2	50
pg_toast_1255	0	0.2	50
pg_toast_2620	0	0.2	50

Instance: wraparound

- [Dokumentace](#)
- The maximum time that a table can go unvacuumed is two billion transactions minus the `vacuum_freeze_min_age` value **at the time of** the last aggressive vacuum.
- uvedená kontrola *předpokládá*, že se `vacuum_freeze_min_age` neměnilo

```
SELECT datname, age(datfrozenxid) as age,  
(  
    age(datfrozenxid)/(2*10^9-current_setting('vacuum_freeze_min_age')::int)  
)::real as pct_to_wraparound  
FROM pg_database;
```

datname	age	pct_to_wraparound
template0	64573	3.31144e-05
postgres	64573	3.31144e-05
hibernate	64573	3.31144e-05

DB: wraparound – table level

```
with relafma as(
select cropt.oid, cropt.ropt[2]::int from
  ( SELECT c.oid, string_to_array(unnest(c.reloptions), '=') as ropt FROM    pg_class c ) cropt
  where cropt.ropt[1] ilike 'autovacuum_freeze_max_age'
),
relvfma as(
select cropt.oid, cropt.ropt[2]::int from
  ( SELECT c.oid, string_to_array(unnest(c.reloptions), '=') as ropt FROM    pg_class c ) cropt
  where cropt.ropt[1] ilike 'vacuum_freeze_min_age'
),
pgcl as ( select c.oid, c.relnamespace, c.relkind,
  age(c.relfrozenxid) as tbl_age,
  age(t.relfrozenxid) as toast_age,
  greatest(age(c.relfrozenxid),age(t.relfrozenxid)) as age,
  coalesce(relvfma.ropt, current_setting('vacuum_freeze_min_age')::int) as tbl_vfma,
  coalesce(relafma.ropt, current_setting('autovacuum_freeze_max_age')::int) as tbl_afma
FROM pg_class c
LEFT JOIN pg_class t ON c.reltoastrelid = t.oid
left join relafma on c.oid = relafma.oid
left join relvfma on c.oid = relvfma.oid
)
  SELECT current_database() as datname, nspname AS schemaname,
    c.oid::regclass::text as table_name,
    c.relkind,
    c.tbl_age,
    c.toast_age,
    c.age,
    (c.age/(2*10^9 - c.tbl_vfma))::real as ptc_to_tbl_wraparound,
    (c.age / (least(c.tbl_afma, 2*10^9)))::real AS pct_to_tbl_aggressive_vacuum
FROM pgcl c
LEFT JOIN pg_namespace n ON n.oid = c.relnamespace
WHERE c.relkind IN ('r', 'm')
order by age desc;
```

DB: wraparound

- pct_to_tbl_wraparound
 - lze najít kerá/é konkrétní tabulky jsou nejstarší
- pct_to_tbl_aggressive_vacuum
 - all-visible but not all-frozen pages are scanned

table_name	relkind	age	pct_to_tbl_wraparound	pct_to_tbl_aggressive_vacuum
pg_type	r	64573	3.31144e-05	0.000322865
pg_authid	r	64573	3.31144e-05	0.000322865
pg_proc	r	64573	3.31144e-05	0.000322865
pg_class	r	64573	3.31144e-05	0.000322865
pg_statistic	r	64573	3.31144e-05	0.000322865

Instance: pg_stat_statements

- „must have“ extenze
- pro interaktivní práci VŽDY seřadit a pracovat jen s nejnáročnějšími dotazy
 - calls, total_time, rows
 - shared_blks_hit, shared_blks_read, temp_blks_read, temp_blks_written, blk_read_time, blk_write_time
 - min_time, max_time, mean_time, stddev_time
- pro monitoring sbírat s rozumnou periodou
 - vyhodnocujeme pak rozdíly za daný časový úsek
 - InfluxDB má pro tento účel fci `non_negative_derivative()`

Dotazy ?

Ukázky dotazů pro jednotlivé verze jsou na dalších stranách

pg_stat_activity 9.1

```
SELECT datid, datname, usesysid, username, application_name,  
client_addr, client_hostname, client_port, backend_start,  
extract(epoch from current_timestamp - xact_start) as  
xact_duration,  
CASE  
    WHEN current_query = '<IDLE>' THEN 'idle'  
    WHEN current_query = '<IDLE> in transaction' THEN 'idle in  
transaction'  
    ELSE 'unknown'  
END as state,  
CASE  
    WHEN current_query = '<IDLE>' THEN null  
    ELSE extract(epoch from current_timestamp - query_start)  
END as query_duration,  
waiting,  
CASE  
    WHEN current_query = '<IDLE>' THEN null  
    ELSE current_query  
END as query_text  
FROM pg_stat_activity;
```

-- InfluxDB tags jsou kurzívou

pg_stat_activity 9.2 – 9.3

```
SELECT datid, datname, pid, usesysid, username, application_name,  
client_addr, client_hostname, client_port, backend_start,  
extract(epoch from current_timestamp - xact_start) as  
xact_duration, state_change, waiting, state,  
CASE  
    WHEN state = 'idle' THEN null  
    ELSE extract(epoch from current_timestamp - query_start)  
END as query_duration,  
CASE state  
    when 'idle' THEN null  
    ELSE query  
END as query_text  
FROM pg_stat_activity;
```

pg_stat_activity 9.4 – 9.5

```
SELECT datid, datname, pid, usesysid, username, application_name,  
client_addr, client_hostname, client_port, backend_start,  
extract(epoch from current_timestamp - xact_start) as  
xact_xact_duration,  
state_change, waiting, state,  
CASE  
    WHEN state = 'idle' THEN null  
    ELSE extract(epoch from current_timestamp - query_start)  
END as query_duration,  
backend_xid, backend_xmin,  
CASE state  
    when 'idle' THEN null  
    ELSE query  
END as query_text  
FROM pg_stat_activity;
```

pg_stat_activity 9.6

```
SELECT datid, datname, pid, usesysid, username, application_name,  
client_addr, client_hostname, client_port, backend_start,  
extract(epoch from current_timestamp - xact_start) as  
xact_duration,  
CASE  
    WHEN state = 'idle' THEN null  
    ELSE extract(epoch from current_timestamp - query_start) END  
as query_duration ,  
state_change,  
CASE  
    WHEN wait_event_type is null THEN false  
    ELSE true  
END as waiting,  
wait_event_type, wait_event, state, backend_xid, backend_xmin,  
CASE  
    when state = 'idle' THEN null  
    ELSE query  
END as query_text  
FROM pg_stat_activity;
```

Pg_stat_archiver 9.4 – 9.6

```
select
  archived_count, last_archived_wal,
  extract(epoch from current_timestamp - last_archived_time)::int as last_arch_sec_age,
  failed_count, last_failed_wal,
  case when (last_failed_wal IS NULL OR last_failed_wal <= last_archived_wal) then
    null
  else
    extract(epoch from current_timestamp - last_failed_time)::int
  end as last_failed_sec_age,
  stats_reset,
  (
    current_setting('archive_mode')::BOOLEAN
    AND ( last_failed_wal IS NULL
          OR
          last_failed_wal <= last_archived_wal
        )
  )AS is_archiving,
  -- „xlog volume appendix“
  pg_xlog_location_diff(
    pg_current_xlog_location(), '0/00000000'::pg_lsn
  ) as xlog_volume
from pg_stat_archiver;
```

-- [Motivace pro "is archiving" z blogu 2nd Quadrant](#)

16.2.2017

Monitoring jako služba pro vývojáře

54

pg_stat_bgwriter

- 9.1

```
SELECT checkpoints_timed, checkpoints_req, buffers_checkpoint,  
buffers_clean, maxwritten_clean, buffers_backend,  
buffers_backend_fsync, buffers_alloc, stats_reset  
FROM pg_stat_bgwriter;
```

- 9.2 – 9.6

```
SELECT  
checkpoints_timed,  
checkpoints_req,  
checkpoint_write_time,  
checkpoint_sync_time,  
buffers_checkpoint,  
buffers_clean,  
maxwritten_clean,  
buffers_backend,  
buffers_backend_fsync,  
buffers_alloc,  
stats_reset  
FROM pg_stat_bgwriter;
```

pg_stat_database

- 9.1

```
SELECT datid, datname, numbackends, xact_commit, xact_rollback,  
blks_read, blks_hit, tup_returned, tup_fetched, tup_inserted,  
tup_updated, tup_deleted, conflicts, stats_reset  
FROM pg_stat_database;
```

- 9.2 – 9.6

```
SELECT  
  datid, datname,  
  numbackends,  
  xact_commit, xact_rollback,  
  blks_read, blks_hit,  
  tup_returned, tup_fetched, tup_inserted,  
  tup_updated, tup_deleted,  
  conflicts,  
  temp_files, temp_bytes,  
  deadlocks,  
  blk_read_time, blk_write_time,  
  stats_reset  
FROM pg_stat_database;
```

pg_stat_user_tables

- 9.1 – 9.6
 - data pro připojenou DB
 - obdobné dotazy pro indexy, funkce, sekvence...

```
select
  current_database() as datname,
  relid, schemaname, relname,
  seq_scan, seq_tup_read,
  idx_scan, idx_tup_fetch,
  n_tup_ins, n_tup_upd, n_tup_del,
  n_tup_hot_upd, n_live_tup, n_dead_tup,
  last_vacuum, last_autovacuum,
  last_analyze, last_autoanalyze,
  vacuum_count, autovacuum_count,
  analyze_count, autoanalyze_count
from pg_stat_user_tables;
```



~~~ definitivní konec presentace ~~~