

ENGLISH




MEETUP ONLINE '21



WHAT TAKES DISK SPACE

AIGARS KADIKIS
TECHNICAL SUPPORT ENGINEER

ZABBIX



WHAT TAKES DISK SPACE - AGENDA

- ✓ Tables
- ✓ Data types
- ✓ Items
- ✓ Hosts
- ✓ Tips



BIGGEST TABLES

history
history_uint

history_str
history_text
history_log

events

MASURE SIZE OF TABLES (MYSQL)

```
SELECT table_name,  
       table_rows,  
       data_length,  
       index_length,  
       round(((data_length + index_length) / 1024 / 1024 / 1024),2) "Size in GB"  
FROM information_schema.tables  
WHERE table_schema = "zabbix"  
ORDER BY round(((data_length + index_length) / 1024 / 1024 / 1024),2) DESC  
LIMIT 8;
```

TABLE_NAME	TABLE_ROWS	DATA_LENGTH	INDEX_LENGTH	Size in GB
history	39287666	2276048896	1389445120	3.41
history_uint	37597109	2179383296	1393065984	3.33
history_text	381535	871006208	22708224	0.83
trends_uint	2257508	163282944	0	0.15
trends	1149013	85590016	0	0.08
alerts	8196	4734976	1376256	0.01
items	7924	3686400	2424832	0.01
history_str	47559	5783552	1589248	0.01


MEASURE SIZE OF TABLES (POSTGRESQL)

```
SELECT *, pg_size_pretty(total_bytes) AS total , pg_size_pretty(index_bytes) AS index ,
         pg_size_pretty(toast_bytes) AS toast , pg_size_pretty(table_bytes) AS table
FROM (SELECT *, total_bytes-index_bytes-coalesce(toast_bytes, 0) AS table_bytes
      FROM (SELECT c.oid,
                  nspname AS table_schema,
                  relname AS table_name ,
                  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 = 'r' ) a) a;
```

In case if TimescaleDB extension in place, it will tell the biggest hypertables (AKA chunks)

BUILD IN WAY TO TRACK INCOMING FLOW

<input type="checkbox"/>	Wizard	Name ▲	Triggers	Key	Interval	History	Trends	Type	Applications	Status	Info
<input type="checkbox"/>	...	Template App Zabbix Server: Number of processed character values per second		zabbix[wcache,values,str]	6s	5d	365d	Zabbix internal	Zabbix write cache	Enabled	
<input type="checkbox"/>	...	Template App Zabbix Server: Number of processed log values per second		zabbix[wcache,values,log]	6s	5d	365d	Zabbix internal	Zabbix write cache	Enabled	
<input type="checkbox"/>	...	Template App Zabbix Server: Number of processed numeric (float) values per second		zabbix[wcache,values,float]	6s	5d	365d	Zabbix internal	Zabbix write cache	Enabled	
<input type="checkbox"/>	...	Template App Zabbix Server: Number of processed numeric (unsigned) values per second		zabbix[wcache,values,uint]	6s	5d	365d	Zabbix internal	Zabbix write cache	Enabled	
<input type="checkbox"/>	...	Template App Zabbix Server: Number of processed text values per second		zabbix[wcache,values,text]	6s	5d	365d	Zabbix internal	Zabbix write cache	Enabled	



BIGGEST DATA COMING TO INSTANCE RIGHT NOW

```
SELECT hosts.host,items.itemid,items.key_,
COUNT(history_log.itemid) AS 'count', AVG(LENGTH(history_log.value)) AS 'avg size',
(COUNT(history_log.itemid) * AVG(LENGTH(history_log.value))) AS 'Count x AVG'
FROM history_log
JOIN items ON (items.itemid=history_log.itemid)
JOIN hosts ON (hosts.hostid=items.hostid)
WHERE clock > UNIX_TIMESTAMP(NOW() - INTERVAL 30 MINUTE)
GROUP BY hosts.host,history_log.itemid
ORDER BY 6 DESC
LIMIT 1\G
```

```
***** 1. row *****
      host: VMware host
     itemid: 123780
      key_: vmware.eventlog[{$VMWARE.URL},skip]
     count: 474
   avg size: 93.3586
Count x AVG: 44252.0000
```

Possible tables to analyze history_text, history_log, history_str

SEE DATA, POSSIBLY DELETE

```
mysql> SELECT hosts.host,items.itemid,items.key_,
-> COUNT(history_log.itemid) AS 'count', AVG(LENGTH(history_log.value)) AS 'avg size',
-> (COUNT(history_log.itemid) * AVG(LENGTH(history_log.value))) AS 'Count x AVG'
-> FROM history_log
-> JOIN items ON (items.itemid=history_log.itemid)
-> JOIN hosts ON (hosts.hostid=items.hostid)
-> WHERE clock > UNIX_TIMESTAMP(NOW() - INTERVAL 30 MINUTE)
-> GROUP BY hosts.host,history_log.itemid
-> ORDER BY 6 DESC
-> LIMIT 1\G
***** 1. row *****
      host: VMware host
      itemid: 123780
      key_: vmware.eventlog[{$VMWARE.URL},skip]
      count: 474
      avg size: 93.3586
      Count x AVG: 44252.0000
```

```
SELECT value FROM history_log WHERE itemid=123780 LIMIT 1\G
```

```
SET SESSION SQL_LOG_BIN=0; DELETE FROM FROM history_log WHERE itemid=123780;
```


BIGGEST DATA PER HOST PER ONE DATA TYPE

```
SELECT ho.hostid, ho.name, count(*) AS records,  
(count(*) * (SELECT AVG_ROW_LENGTH FROM information_schema.tables  
WHERE TABLE_NAME = 'history_text' and TABLE_SCHEMA = 'zabbix')/1024/1024)  
AS 'Total size average (Mb)', sum(length(history_text.value))/1024/1024 +  
sum(length(history_text.clock))/1024/1024 + sum(length(history_text.ns))/1024/1024 +  
sum(length(history_text.itemid))/1024/1024 AS 'history_text Column Size (Mb)'  
FROM history_text  
LEFT OUTER JOIN items i on history_text.itemid = i.itemid  
LEFT OUTER JOIN hosts ho on i.hostid = ho.hostid  
WHERE ho.status IN (0,1)  
AND clock > UNIX_TIMESTAMP(now() - INTERVAL 1 DAY - INTERVAL 6 MINUTE)  
AND clock < UNIX_TIMESTAMP(now() - INTERVAL 1 DAY)  
GROUP BY ho.hostid ORDER BY 4 DESC LIMIT 5;
```

hostid	Name	records	Total size average (Mb)	history_text Column Size (Mb)
11120	Java gateway	2400	0.53787231	0.07963181
11125	Nginx	600	0.13446808	0.16013908

Repeat process with tables «history_log» and «history_str»

TRACK SIZE ON PARTITIONS (MYSQL)

Locate what is the data directory with command:

```
select @@datadir;
```

Most of times it prints '/var/lib/mysql'. In case the database name is 'zabbix', then we need to navigate:

```
cd /var/lib/mysql/zabbix
ls -lh history#*
ls -lh history_uint#*
ls -lh history_str#*
ls -lh history_text#*
ls -lh history_log#*
ls -lh trends#*
ls -lh trends_uint#*
```

```
[root@demo zabbix]# ls -lh history_log#*
-rw-r-----. 1 mysql mysql 44M Jan 22 14:16 history_log#p#p2021_02w.ibd
-rw-r-----. 1 mysql mysql 13M Jan 22 14:56 history_log#p#p2021_03w.ibd
```

ANALYZE A PARTITION (MYSQL)

```
SELECT ho.hostid, ho.name, count(*) AS records,  
(count(*)* (SELECT AVG_ROW_LENGTH FROM information_schema.tables  
WHERE TABLE_NAME = 'history_log' and TABLE_SCHEMA = 'zabbix')/1024/1024)  
AS 'Total size average (Mb)', sum(length(history_log.value))/1024/1024 +  
sum(length(history_log.clock))/1024/1024 + sum(length(history_log.ns))/1024/1024 +  
sum(length(history_log.itemid))/1024/1024 AS 'history_log Column Size (Mb)'  
FROM history_log PARTITION (p2021_02w)  
LEFT OUTER JOIN items i on history_log.itemid = i.itemid  
LEFT OUTER JOIN hosts ho on i.hostid = ho.hostid  
WHERE ho.status IN (0,1)  
GROUP BY ho.hostid ORDER BY 4 DESC LIMIT 10;
```

hostid	name	records	Total size average (Mb)	history_log Column Size (Mb)
10461	VMware host	117938	20.92024612	13.35941887
10084	Zabbix server	41535	7.36762047	7.34180927

Repeat process with tables «history_text» and «history_str»

TO FREE UP SPACE ON MYSQL

It's required to rebuild partitions. This can be a timeconsuming process:

```
SHOW CREATE TABLE history;  
ALTER TABLE history REBUILD PARTITION p202101160000;
```

If there is not enough free disk space, we can crash the DB engine.

In the background it will copy the data from one partition to new partition.

It's a lot of I/O operations.

Execute command through the «screen» utility.

TO FREE UP SPACE ON POSTGRES

Postgres is using autovacuum functionality. It is a separate process which cleans up dead tuples:

```
SELECT schemaname, relname, n_live_tup, n_dead_tup, last_autovacuum
FROM pg_stat_all_tables
WHERE n_dead_tup > 0
ORDER BY n_dead_tup DESC;
```

Query will tell how many dead tuples are in each table and when the last autovacuum has occurred.

TO FREE UP SPACE ON POSTGRES

If vacuum has not occurred in last 10 days, it's bad:

schemaname	relname	n_live_tup	n_dead_tup	last_autovacuum
public	history_uint	1228819782	577423316	2020-08-13 04:55:11.54239-07

Must change settings and increase the priority for vacuum process:

```
vacuum_cost_page_miss = 10  
vacuum_cost_page_dirty = 20  
autovacuum_vacuum_threshold = 50  
autovacuum_vacuum_scale_factor = 0.01  
autovacuum_vacuum_cost_delay = 20ms  
autovacuum_vacuum_cost_limit = 3000  
autovacuum_max_workers = 6
```

LOG FILE MONITORING

Log file monitoring has been configured per application.

However, when an application (JMX) encounters an error, it generates a lot of long lines per second.

Solution: use 'log.count' instead of log item and seek the occurrence of patterns.

ZABBIX RAW ITEMS

'Zabbix raw items' works in tandem with dependable items.

A good template solution will use «a master item» and multiple «dependable items».

By default, a master item must have a historical period 0.

Sometimes for troubleshooting purpose we enable «History» for «Zabbix raw items».

It will put an unnecessary content in database

All templates / Oracle by ODBC Applications 2 **Items 73** Triggers 17 Graphs 7 Dashboards Discovery rules 5 Web scenarios Filter

<input type="checkbox"/>	Wizard	Name ▲	Triggers	Key	Interval	History	Trends	Type	Applications	Status	Info
<input type="checkbox"/>	...	Oracle: Get archive log info		db.odbc.get[get_archive_log_stat,"{\$ORACLE.DSN}"]	5m	0		Database monitor	Zabbix raw items	Enabled	
<input type="checkbox"/>	...	Oracle: Get ASM stats		db.odbc.get[get_asm_stat,"{\$ORACLE.DSN}"]	1m	0		Database monitor	Zabbix raw items	Enabled	
<input type="checkbox"/>	...	Oracle: Get CDB and No-CDB info		db.odbc.get[get_cdb_info,"{\$ORACLE.DSN}"]	1m	0		Database monitor	Zabbix raw items	Enabled	
<input type="checkbox"/>	...	Oracle: Get instance state		db.odbc.get[get_instance_state,"{\$ORACLE.DSN}"]	1m	0		Database monitor	Zabbix raw items	Enabled	
<input type="checkbox"/>	...	Oracle: Get PDB info		db.odbc.get[get_pdb_info,"{\$ORACLE.DSN}"]	1m	0		Database monitor	Zabbix raw items	Enabled	
<input type="checkbox"/>	...	Oracle: Get system metrics		db.odbc.get[get_system_metrics,"{\$ORACLE.DSN}"]	0	0		Database monitor	Zabbix raw items	Enabled	
<input type="checkbox"/>	...	Oracle: Get tablespaces stats		db.odbc.get[tablespace_stats,"{\$ORACLE.DSN}"]	1m	0		Database monitor	Zabbix raw items	Enabled	

ENGLISH




MEETUP ONLINE '21

QUESTIONS?

AIGARS KADIKIS
TECHNICAL SUPPORT ENGINEER

ZABBIX



ENGLISH



MEETUP ONLINE '21

THANK YOU!

AIGARS KADIKIS
TECHNICAL SUPPORT ENGINEER

ZABBIX

