yb_stats

Snapshot all available operational data

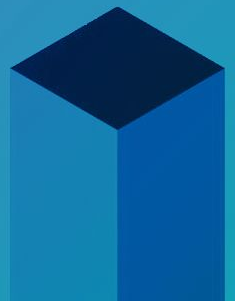

YFTT YUGABYTEDB FRIDAY TECH TALKS

yugabyteDB

What is yb_stats?

A utility to extract all available operational data from a YugabyteDB cluster.

Why do I need yb_stats?

- To obtain all available **facts** and **performance** details.
- For YugabyteDB clusters of any kind, any version in any kind of situation.
- One stop information source: everything that can be obtained is there.
- In an easy transportable, open format (UTF8 CSV).
- yb_stats results can be shared between people, even from clusters which don't have internet access. ('airgapped')
- Allows the collection of facts to base investigations on, instead of guesses and advise based on best practices which might or might not apply.
- Get mathematical exact facts, which allows performance and capacity calculations.
- Get an accurate **before** and **after** situation.

What does yb_stats NOT do?

- It is not a replacement for YugabyteDB Anywhere's metrics component.
- It does not interpret data, it provides raw technical facts.

- Currently, yb_stats is not part of a default YugabyteDB install.

Resource usage of yb_stats

- A common question: does running yb_stats have a performance impact?
  - All yb_stats (and YugabyteDB Anywhere) data sources:
    - are explicit metadata endpoints that do not obstruct database processing.
    - provide a relatively low amount of data.

- In exceptional cases cases the threads view might cause noticeable impact.
  - (the view asks each thread to provide it's current stack).
  - For these cases there is the '--disable-threads' switch.

- Running yb_stats from a YugabyteDB cluster node has insignificant performance impact.

Setting up yb_stats

- yb_stats needs the ip addresses for the endpoints.
  - Pretty much like the 'yb-admin' tool.
- It stores the endpoint ip addresses and ports in the file '.env'
  - In the current working directory.
  - Automatically reused, so one-time setup!
- Automatically probes all set ip addresses for set ports and endpoints.

Using yb_stats

- yb_stats has got two modes:
  - 'ad hoc' mode:
    - show performance data difference only ("investigation mode").
    - nothing is stored, only performance data is obtained and shown.
  - 'snapshot' mode:
    - all data is gathered and stored.
    - CLI output is snapshot number only.
    - snapshot output via '--snapshot-diff' (performance) or '--print-*' switches.

Demo - Basic usage

- Using --help
- https://github.com/fritshoogland-yugabyte/yb_stats#readme

- Perform snapshots
  - Performance data (--snapshot-diff)
  - Master status (--print-masters *N*)
  - Versions (--print-version *N*)
  - Gflags (--print-gflags *N*)
  - Logfile contents, mem-tracker, entities (table and tablet metadata)

Demo - Investigate work

- Scenario: get an impression of performed work in YSQL/postgres.

- Typical steps for workload investigation:
  - snapshot before situation.
  - perform workload.
  - snapshot after situation.

Installing yb_stats

yb_stats can be installed via yum in the following way:

EL7 (Centos 7)

```
sudo yum install -y
https://github.com/fritshoogland-yugabyte/yb_stats/releases/download/v0.8.5/yb_stats-
0.8.5-1.el7.x86_64.rpm
```

EL8 (Alma 8)

```
sudo yum install -y
https://github.com/fritshoogland-yugabyte/yb_stats/releases/download/v0.8.5/yb_stats-
0.8.5-1.el8.x86_64.rpm
```

Installing yb_stats

For other platforms:

- Install cargo, git clone yb_stats, compile yb_stats.
  - https://github.com/fritshoogland-yugabyte/yb_stats/blob/main/BUILD.md

- I am thinking about: OSX/brew, Ubuntu/apt..
  - Please request or ask to indicate demand!

yb_stats useful indicators

node exporter: dmi: host "hardware":

```
--hostname-match 9300 --gauges-enable --stat-name-match dmi_info
```

node exporter: cpu: number of CPUs:

```
--hostname-match 9300 --stat-name-match cpu_seconds
```

node exporter: memory: total size, available, free:

```
--hostname-match 9300 --gauges-enable --stat-name-match memory_Mem
```

node exporter: disk: number and size:

```
--hostname-match 9300 --gauges-enable --stat-name-match filesystem_size
```

node exporter: network: NICs and MTU:

```
--hostname-match 9300 --gauges-enable --stat-name-match network_mtu
```

yb_stats useful indicators

tablet servers: YSQL layer info (pggate, pg_stat_statements):

```
--hostname-match 13000
```

tablet servers: YCQL layer info:

```
--hostname-match 12000
```

tablet servers: generic engine/DocDB:

```
--hostname-match 9000
```

master servers: generic:

```
--hostname-match 7000
```

yb_stats useful indicators

tablet server and master:

Memory assessment:

```
--snapshot-diff --gauges-enable --stat-name-match
'(tcmalloc|generic|mem_tracker)'
```

- Difference between master (metadata) and tablet server (data).
- Add `--details-enable` for looking at tablet level statistics details.
- `--snapshot-diff` allows seeing difference between two points in time.

yb_stats useful indicators

tablet server and master:

IO assessment:

```
--snapshot-diff --hostname-match '(7000|9000)' --stat-name-match
'(log_append_latency|log_sync_latency|rocksdb_sst_read_micros|rocksdb_w
rite_raw_blocks)'
```

- Should mainly show tablet server IO, unless DDL is performed.
- Add `--details-enable` for looking at tablet level statistics details.
- Log IO is in critical codepath, memtable/SST file IO is mainly independent.

yb_stats useful indicators

tablet server and master:

CPU assessment:

```
--snapshot-diff --hostname-match '(7000|9000)' --stat-name-match
'(cpu|context|schedstat)'
```

- CPU is divided between 'stime' and 'utime' (system/kernel mode and user mode).
- Compare CPU time with available CPU, context switches.
- Significant schedstat_waiting indicates CPU shortage.

Thank You

Join us on Slack: yugabyte.com/slack (#yftt channel)

Star us on Github: github.com/yugabyte/yugabyte-db

YFTT YUGABYTEDB FRIDAY TECH TALKS

yugabyteDB