

openQA Infrastructure - action #93922

grafana dashboard for "approximate result size by job group" fails to render any data with "InfluxDB Error: unsupported mean iterator type: *query.stringInterruptIterator"

2021-06-11 21:47 - okurz

Status:	Resolved	Start date:	2021-06-11
Priority:	Normal	Due date:	2021-06-29
Assignee:	mkittler	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Ready		
Description			
Observation			
https://monitor.qa.suse.de/d/nRDab3Jiz/openqa-jobs-test?viewPanel=19&orgId=1 shows no data and an error message "InfluxDB Error: unsupported mean iterator type: *query.stringInterruptIterator"			
Related issues:			
Related to openQA Infrastructure - action #94456: no data from any arm host o...		Resolved	2021-06-22

History

#1 - 2021-06-14 11:56 - mkittler

- Assignee set to mkittler

I'll have a look.

#2 - 2021-06-14 12:33 - mkittler

Looks like the result size is stored as a string (at least at some point) but needed to be stored as an integer:

```
> show field keys from postgresql;
name: postgresql
fieldKey          fieldType
-----
...
result_size       integer
result_size       string
```

Apparently telegraf already passes the value as string:

```
sudo telegraf --test --debug -config /etc/telegraf/telegraf.conf -config-directory /etc/telegraf/telegraf.d |
grep -i result_size
...
2021-06-14T12:28:48Z D! postgresql_extensible: column: group_id = int64: 380
2021-06-14T12:28:48Z D! postgresql_extensible: column: group_name = string: Containers/SLE Base Image Updates
15-SP1
2021-06-14T12:28:48Z D! postgresql_extensible: column: result_size = string: 91464511
2021-06-14T12:28:48Z D! postgresql_extensible: column: group_id = int64: 381
2021-06-14T12:28:48Z D! postgresql_extensible: column: group_name = string: Containers/SLE Base Image Updates
15 GA
2021-06-14T12:28:48Z D! postgresql_extensible: column: result_size = string: 88512145
> postgresql,db=postgres,group_id=380,group_name=Containers/SLE\ Base\ Image\ Updates\ 15-SP1,host=openqa,server=dbname=openqa\ host=localhost\ user=telegraf result_size="91464511" 1623673729000000000
> postgresql,db=postgres,group_id=381,group_name=Containers/SLE\ Base\ Image\ Updates\ 15\ GA,host=openqa,server=dbname=openqa\ host=localhost\ user=telegraf result_size="88512145" 1623673729000000000
...
```

It looks like one cannot easily change the data type on InfluxDB level (<https://docs.influxdata.com/influxdb/v1.8/troubleshooting/frequently-asked-questions/#can-i-change-a-fields-data-type>¹) so I'll check on telegraf-level. Likely we cannot easily fix the values which are already present then.

¹ I've also tested it by changing the query to `SELECT moving_average(mean("result_size"::integer), 10) / 1024 / 1024 / 1024 FROM "postgresql" WHERE $timeFilter GROUP BY time($__interval), "group_name"`. The error goes away but no data is returned anymore, likely because the cast only returns a null value or something.

#3 - 2021-06-14 13:07 - mkittler

- Status changed from New to In Progress

SR: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/505

This should fix the cause of the problem. Note that casting to numeric wouldn't work here; it would be end up being passes as string in my tests. In fact, I believe that we're getting a numeric when computing the sum of bigints (see <https://www.postgresql.org/docs/13/functions-aggregate.html>) is actually the problem here. However, I'm still wondering why it worked at some point before. (The documentation for older PostgreSQL versions shows that it was always like this. Maybe a change in the other tools introduced the problem.)

I'm also wondering what to do about the string values which are now in the database. They still prevent displaying the graph. I see the toInt() function only showing up in InfluxDB's documentation for v2 but we're still at v1.8¹.

¹ It looks like [upgrading to v2](https://docs.influxdata.com/influxdb/v2.0/tools/grafana) would be possible with Grafana (<https://docs.influxdata.com/influxdb/v2.0/tools/grafana>).

#4 - 2021-06-14 14:46 - mkittler

Now I remember: This column wasn't always bigint. I've been changing it not too long ago (see <https://github.com/os-autoinst/openQA/commit/69998774a13bc337c02140b7b0eec6cd704ef63f>). If I knew that before I would of course had created the SR before but now we're stuck with the string values which have already been added and there seems no easy way to cope with them (see previous comments) or to get rid of them (see https://docs.influxdata.com/influxdb/v1.8/query_language/manage-database and <https://stackoverflow.com/questions/39685114/delete-points-with-unwanted-field-values-from-influxdb-measurement>).

#5 - 2021-06-14 18:13 - okurz

Your MR was merged but <https://monitor.qa.suse.de/d/nRDab3Jiz/openqa-jobs-test?viewPanel=19&orgId=1&from=now-15m&to=now> doesn't show something yet.

#6 - 2021-06-15 04:09 - openqa_review

- Due date set to 2021-06-29

Setting due date based on mean cycle time of SUSE QE Tools

#7 - 2021-06-15 14:39 - mkittler

Maybe it wasn't quite clear from the previous comments: The SR fixes the cause of the problem but doesn't repair the damage which has already been done.

#8 - 2021-06-18 11:03 - mkittler

I've been removing all the string values. Now `curl -G 'http://localhost:8086/query?db=telegraf&epoch=ns' --data-urlencode 'q=SELECT "result_size"::string FROM "postgresql"'` returns no more results.

However, the query still fails. Adding `::integer` helps. Like mentioned before, this won't actually cast any string values to ints but at least acts as a filter and changes the iterator type. The graph then shows no data but that's because there's no *recent* data. When changing the time parameters older data can be shown. Annoyingly, new data cannot be inserted as well:

```
Jun 18 13:00:08 openqa telegraf[15959]: 2021-06-18T11:00:08Z E! [outputs.influxdb] when writing to [http://openqa-monitor.qa.suse.de:8086]: received error partial write: field type conflict: input field "result_size" on measurement "postgresql" is type integer, already exists as type string dropped=35; discarding points
```

This is apparently just how InfluxDB works: https://docs.influxdata.com/influxdb/v1.8/guides/write_data/#schemaless-design

I'm wondering how the switch from integer to string could happen then, though.

I was actually hoping these errors would go away when deleting all string values but that's apparently not the case. Any ideas how to fix that? So far the only solution which comes to my mind is using a new input field (e.g. "result_size2"). It had the disadvantage that old data wouldn't be shown anymore in the graph.

```
(Query to reproduce write error: curl -i -X POST 'http://localhost:8086/write?db=telegraf' --data-binary 'postgresql,db=postgres,group_id=378,group_name=Containers/SLE\ Base\ Image\ Updates\ 15-SP3,host=openqa,server=dbname\=openqa\ host\=localhost\ user\=telegraf result_size=253051198i 1624014635000000000')
```

Maybe one could also:

1. [export](#) everything
2. alter the export file via sed
3. drop postgresql from InfluxDB
4. re-insert the exported values

#9 - 2021-06-21 09:36 - mkittler

Thinking about it further, likely the last option (export/import) is the only way. Of course it means that all data is temporarily not accessible. Deleting all the string values took almost one day. Not sure how fast the full export/import will be. Since all the string values have already been deleted point 2. will likely be a trivial check whether that's actually the case. Has anybody already done this or other ideas? (A full export/import seems to be the last resort.)

#10 - 2021-06-21 12:34 - mkittler

I'm trying the export now: `sudo -u influxdb bash -c 'influx_inspect export -database telegraf -out - -datadir /var/lib/influxdb/data -waldir /var/lib/influxdb/wal | grep "postgresql," > /var/lib/influxdb/dumps/postgresql'`

It only allows to export the whole database and not just the postgresql measurement, hence the grep command. (One can however drop a specific measurement, so DROP MEASUREMENT postgresql should work later on.)

I'll have an eye on it but it'll likely take a while. Feel free to abort the export if necessary, e.g. if we're running out of disk space.

#11 - 2021-06-22 07:09 - okurz

- Related to action #94456: no data from any arm host on <https://stats.openqa-monitor.qa.suse.de/d/1bNU0StZz/automatic-actions?orgId=1> added

#12 - 2021-06-22 11:41 - mkittler

After dropping the measurement, new data shows up again:

<https://stats.openqa-monitor.qa.suse.de/d/nRDab3Jiz/openqa-jobs-test?viewPanel=19&orgId=1>

I'm currently still in the process of importing the exported data. Unfortunately `influx -import -path=/var/lib/influxdb/dumps/postgresql` doesn't work because it always expects the db name at the beginning of the file (which I've unfortunately greped way thinking it wasn't necessary). So I have been resorting to `while read -r LINE; do curl -i -X POST 'http://localhost:8086/write?db=telegraf' --data-binary "$LINE"; done < /var/lib/influxdb/dumps/postgresql`.

#13 - 2021-06-24 12:45 - mkittler

- Status changed from *In Progress* to *Feedback*

The import is still running. I've stopped and masked `rebootmgr.service` on `monitor.suse.de` to avoid any interruptions because it likely take another few days to finish.

By the way, one can see the imported data here:

<https://stats.openqa-monitor.qa.suse.de/d/nRDab3Jiz/openqa-jobs-test?viewPanel=19&orgId=1&from=now-1y&to=now>

I'm marking the issue as "feedback" because I don't have to do anything actively at this point. (For upgrading InfluxDB and setting retention policies I've created additional tickets.)

#14 - 2021-06-29 08:42 - mkittler

The import has been aborted because the machine has been rebooted despite masking `rebootmgr`. Should I try to continue it? (It won't be that easy because I don't know where it left off.)

#15 - 2021-06-29 12:06 - okurz

mkittler wrote:

Should I try to continue it?

nah, leave it

#16 - 2021-06-30 08:38 - mkittler

- Status changed from *Feedback* to *Resolved*

Just for the record, if we wanted to continue again for some reason or free up the space, the dump is still under `/var/lib/influxdb/dumps/postgresql` on the monitoring host.