

Influx Issue and solutions



Part 4: Plan for development

Marco Roda
CCM meeting
27 March 2024

The issue

- Influxdb has a limit on the number of different values a tag can have
 - Our monitoring system generates entries with 2 tags
 - partition_id
 - source_id
 - `<partition>.<app>.<module>[.<subelement>]`
- The limit (100k) has been reached on the current table used for operation
 - incoming messages with a new tag value are discarded
 - Therefore we start having missing points in the dashboards
- The issues comes from a particular metric
 - TPChannelInfo
 - That is constructed in such a way that every channel has it's own different source_id
 - Even if very few metrics are published, the amount of points of multiple runs and channels piled up

Immediate actions

- this maximum is for us to set, the default 100k can be changed or even removed
 - But all the available documentation suggest that increasing that number can have an impact on performance
- Last week we increased that limit to 120k
 - This should buy some time to complete the new release and test
- I also made 2 PRs to change the metric and its publication
 - Now (from v4.4.0) it will publish the channel as an entry in the structure, not as a tag
 - The plots will not be as nice but I think it's a small price to pay

Plan for the protodune 2 run

- Discussing with Wes, we made a requirements due to the protodune run
 - We would like the data of the monitoring of the run, to be all available in the same table
 - So that people can access the monitoring of the whole run
- My recommendation is that once v4.4.0 is release, we switch to a new table of influxdb
 - And we restore the maximum to 100k
- This should give us plenty of new tags for all possible tests
- Possibility of using influxdb v2
 - That has a default of 1M tags
 - We suspect they made some improvements behind the scenes
 - But Experts are suggesting that at this point is best to wait until the release of v3 to avoid transitioning twice in a short amount of time
 - Pat can comment more on both items

Impact on monitoring revamp v > 5.1

- The existence of that limit will change slightly what we will do for the new monitoring system
 - The tags of the entries will not be redundant
- Tags
 - session
 - app
 - element
 - in the form <module>[.<subelement>]
- In this way we get rid of the multiplicative factor due redundant information in the tags
- I will probably change slightly the OpMonEntry definition (the protobuf schema) to reflect this