



Top Filter

Overview

The top filter is a filter module for MaxScale that monitors every SQL statement that passes through the filter. It measures the duration of that statement, the time between the statement being sent and the first result being returned. The top N times are kept, along with the SQL text itself and a list sorted on the execution times of the query is written to a file upon closure of the client session.

Configuration

The configuration block for the TOP filter requires the minimal filter options in it's section within the MaxScale.cnf file, stored in \$MAXSCALE_HOME/etc/MaxScale.cnf.

```
[MyLogFilter]
type=filter
module=topfilter
```

Filter Options

The top filter does not support any filter options currently.

Filter Parameters

The top filter accepts a number of optional parameters.

Filebase

The basename of the output file created for each session. A session index is added to the filename for each file written.

```
filebase=/tmp/SqlQueryLog
```

The filebase may also be set as the filter, the mechanism to set the filebase via the filter option is superseded by the parameter. If both are set the parameter setting will be used and the filter option ignored.

Count

The number of SQL statements to store and report upon.

```
count=30
```

The default value for the number of statements recorded is 10.

Match

An optional parameter that can be used to limit the queries that will be logged by the top filter. The parameter value is a regular expression that is used to match against the SQL text. Only SQL statements that matches the text passed as the value of this parameter will be logged.

```
match=select.*from.*customer.*where
```

All regular expressions are evaluated with the option to ignore the case of the text, therefore a match option of select will match both `select`, `SELECT` and any form of the word with upper or lowercase characters.

Exclude

An optional parameter that can be used to limit the queries that will be logged by the top filter. The parameter value is a regular expression that is used to match against the SQL text. SQL statements that match the text passed as the value of this parameter will be excluded from the log output.

```
exclude=where
```

All regular expressions are evaluated with the option to ignore the case of the text, therefore an exclude option of select will exclude statements that contain both `select`, `SELECT` or any form of the word with upper or lowercase characters.

Source

The optional source parameter defines an address that is used to match against the address from which the client connection to MaxScale originates. Only sessions that originate from this address will be logged.

```
source=127.0.0.1
```

Examples

Example 1 - Heavily Contended Table

You have an order system and believe the updates of the `PRODUCTS` table is causing some performance issues for the rest of your application. You would like to know which of the many updates in your application is causing the issue.

Add a filter with the following definition;

```
[ProductsUpdateTop20]  
type=filter  
module=topfilter  
count=20
```

```
match=UPDATE.*PRODUCTS.*WHERE
exclude=UPDATE.*PRODUCTS_STOCK.*WHERE
filebase=/var/logs/top/ProductsUpdate
```

Note the exclude entry, this is to prevent updates to the PRODUCTS_STOCK table from being included in the report.

Example 2 - One Application Server is Slow

One of your applications servers is slower than the rest, you believe it is related to database access but you not not sure what is taking the time.

Add a filter with the following definition;

```
[SlowAppServer]
type=filter
module=topfilter
count=20
source=192.168.0.32
filebase=/var/logs/top/SlowAppServer
```

In order to produce a comparison with an unaffected application server you can also add a second filter as a control.

```
[ControlAppServer]
type=filter
module=topfilter
count=20
source=192.168.0.42
filebase=/var/logs/top/ControlAppServer
```

In the router definition add both filters

```
filters=SlowAppServer | ControlAppServer
```

You will then have two sets of logs files written, one which profiles the top 20 queries of the slow application server and another that gives you the top 20 queries of your control application server. These two sets of files can then be compared to determine what if anything is different between the two.