

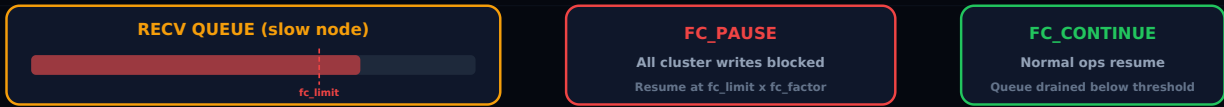
Galera Flow Control

Sylvain ARBAUDIE · 2024-10-28

- GALERA
- MARIADB
- FLOW-CONTROL
- CLUSTERING
- TUNING

GALERA FLOW CONTROL — THE CLUSTER HANDBRAKE

recv queue > fc_limit → FC_PAUSE → cluster writes blocked



EXPERT TUNING RULES

- slave_threads = 2 x CPU**
Drain recv queue faster
- fc_limit = 5 x threads**
Enough room for parallelism
- fc_factor = 0.8**
Progressive resume, no oscillation

wsrep_flow_control_paused > 0.10 = serious problem — investigate immediately

Flow Control is the guardian of Galera consistency — understand it, tune it, monitor it



Galera writeset
recv queue
Flow Control

Flow Control

Flow Control `gcs.fc_limit`

Galera `recv_queue` `FC_PAUSE`

`FC_PAUSE` — `FC_CONTINUE`

`recv_queue` * `gcs.fc_factor` `FC_CONTINUE`

5 wsrep

Flow Control

```
SHOW GLOBAL STATUS WHERE Variable_name IN (
    'wsrep_local_recv_queue',
    'wsrep_local_recv_queue_avg',
    'wsrep_flow_control_paused',
    'wsrep_flow_control_paused_ns',
    'wsrep_flow_control_sent'
);
```

wsrep_local_recv_queue

recv gcs.fc_limit 0

wsrep_local_recv_queue_avg

recv 0.5

wsrep_flow_control_paused

Flow Control 0 1 0.1 10%

wsrep_flow_control_paused_ns

Flow Control

wsrep_flow_control_sent

FC_PAUSE FC_PAUSE

6

gcs.fc_limit

Flow Control 16

gcs.fc_factor

Of 5 gcs.fc_factor Flow Control fc_limit=100 fc_factor=0.8 FC 80

wsrep_slave_threads

= recv 2 x CPU

wsrep_cert_deps_distance

```
wsrep_slave_threads
```

gcs.recv_q_hard_limit

recv OOM swap

gcs.max_throttle

Flow Control 0 1 0.25 FC 25% 0 FC

Galera

Slave

```
wsrep_slave_threads = 2 * CPU_CORES
```

```
wsrep_slave_threads = 16 wsrep_cert_deps_distance -- slave
```

Slave fc_limit

```
gcs.fc_limit = 5 * wsrep_slave_threads
```

16 slaves fc_limit = 80 FC

fc_factor

```
gcs.fc_factor = 0.8
```

0.8 fc_factor 0.5 FC_PAUSE / FC_CONTINUE

```
gcs.recv_q_hard_limit = HALF_RAM_PLUS_SWAP
```

32 GB 16 GB swap recv_q_hard_limit = 24G OOM

