

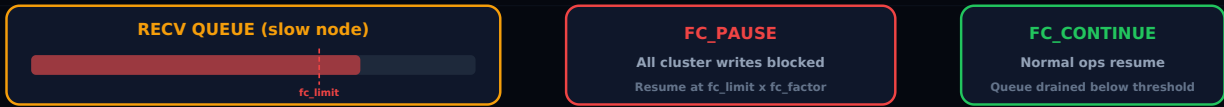
# Galera Flow Control

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_fc_limit * 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

