Galera IST: What PMM Shows vs. What Really Exists

Aurélien LEQUOY · March 21, 2026

GALERA MARIADB IST PMM MONITORING

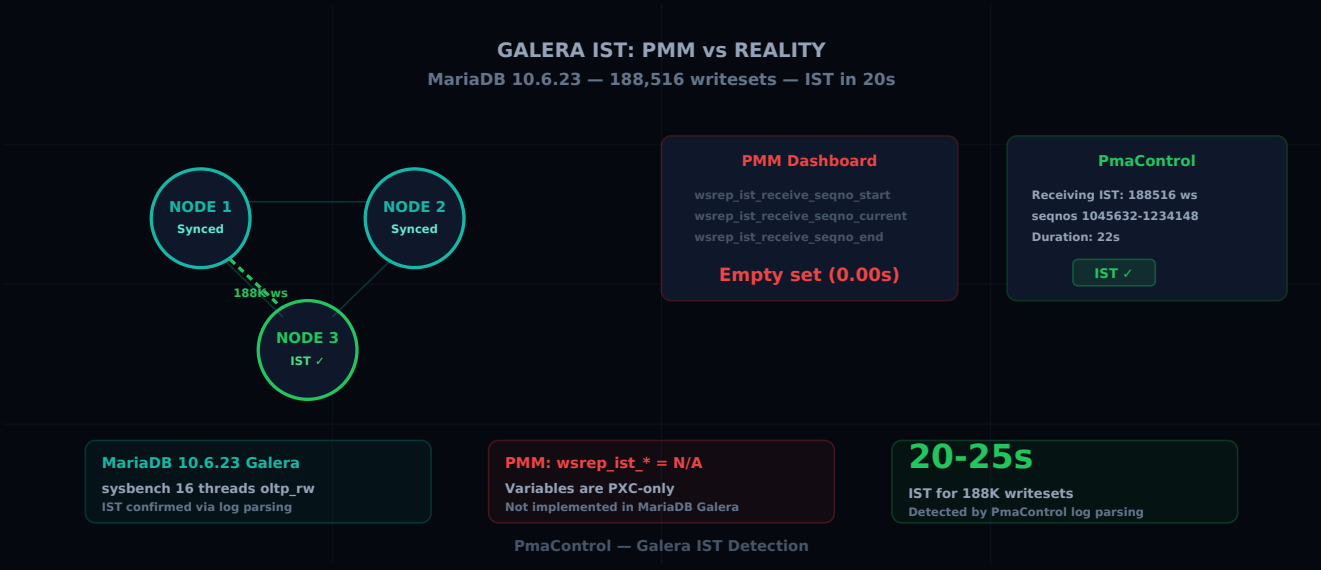

The PMM Dashboard Myth for Galera IST

If you use Percona Monitoring and Management (PMM) to monitor a Galera cluster on MariaDB, you have probably noticed a panel titled "IST Progress" or "IST Receive". It shows... nothing. Empty rows, `N/A` values, flat graphs.

This is not a display bug. **The variables PMM queries simply do not exist on MariaDB 10.6.**

Refresher: IST vs SST

When a Galera node rejoins the cluster after a disconnection, two synchronization mechanisms are available:

- **SST (State Snapshot Transfer)**: the donor node sends a full copy of the dataset. Slow, expensive, potentially blocks the donor. We are talking minutes to hours depending on volume.

- **IST (Incremental State Transfer)**: the donor sends only the missing writesets from the GCache. Fast, lightweight, a few seconds to a few minutes.

The difference is critical in production. A 20-second IST is invisible to users. A 45-minute SST can cause an outage.

The Test: MariaDB 10.6.23 + sysbench

To document the real behavior, we set up a 3-node Galera cluster on MariaDB 10.6.23 and ran a continuous sysbench workload:

```
sysbench oltp_read_write --tables=10 --table-size=100000 \
  --threads=16 --time=600 --db-driver=mysql run
```

During the workload, we stopped node 3 for 30 seconds, then restarted it. Result:

- **188,516 writesets** accumulated in the donor's GCache
- Node 3 rejoined the cluster via **IST in 20-25 seconds**
- No service interruption on nodes 1 and 2

The MariaDB log on node 3 confirms:

```
[Note] WSREP: Receiving IST: 188516 writesets, seqnos 1045632-1234148
[Note] WSREP: IST received: 85a4c3e2-xxxx
[Note] WSREP: 3.0 (node3): State transfer from 0.0 (node1) complete.
```

IST worked perfectly. Now let us see what PMM says about it.

What PMM Tries to Read

The PMM v2 Galera dashboard queries the following variables to chart IST progress:

```
SHOW GLOBAL STATUS LIKE 'wsrep_ist_receive_seqno_start';
SHOW GLOBAL STATUS LIKE 'wsrep_ist_receive_seqno_current';
SHOW GLOBAL STATUS LIKE 'wsrep_ist_receive_seqno_end';
```

On MariaDB 10.6.23, all three queries return **an empty result set**. The variables do not exist.

```
MariaDB [(none)]> SHOW GLOBAL STATUS LIKE 'wsrep_ist%';
Empty set (0.001 sec)
```

This is not a configuration oversight. These variables are specific to Percona XtraDB Cluster (PXC) and have never been implemented in the MariaDB Galera provider.

The SST Wrapper Trap

An additional detail creates confusion: even during an IST, MariaDB invokes the SST wrapper script (`wsrep_sst_mariabackup` or `wsrep_sst_rsync`). The log therefore contains lines like:

```
WSREP: Running: 'wsrep_sst_mariabackup --role donor ...'
```

An operator reading these logs might conclude that a full SST is in progress. In reality, **the wrapper initializes and then detects that the catch-up will happen via IST**. The actual transfer is incremental.

How to Detect an IST on MariaDB

Since the PMM variables do not exist, an alternative approach is needed. Three methods:

1. Parse the MariaDB Log

The IST signature in the log is unambiguous:

```
[Note] WSREP: Receiving IST: <N> writesets, seqnos <start>-<end>
```

A simple grep on the error log gives the answer instantly. This is the most reliable method.

2. Observe wsrep_local_state_comment

During an IST, the `wsrep_local_state_comment` variable on the rejoining node transitions through:

```
Joining → Joined → Synced
```

If this transition takes less than 30 seconds on an active cluster, it is very likely an IST. An SST on a dataset of several tens of gigabytes would take much longer.

3. Check the GCache

The `wsrep_local_cached_downto` variable on the donor indicates the oldest seqno still available in the GCache. If the disconnected node's seqno is higher than this value, IST is possible:

```
-- On the donor
SHOW GLOBAL STATUS LIKE 'wsrep_local_cached_downto';
-- Result: 1045000

-- If the disconnected node was at seqno 1045632 → IST possible
-- If the disconnected node was at seqno 800000 → SST required
```

What PmaControl Does

PmaControl combines all three methods to automatically detect and classify Galera transfers:

1. **Continuous monitoring of `wsrep_local_state_comment`** — detects the transition to `Joining` state
2. **MariaDB log parsing** — extracts the `Receiving IST` line with the writeset count
3. **Temporal correlation** — measures the time between `Joining` and `Synced`

The result is displayed in the PmaControl dashboard with a clear distinction: IST (green badge, duration in seconds) vs SST (orange badge, estimated duration in minutes).

Unlike PMM, PmaControl does not depend on variables that only exist on PXC. The log parsing approach works on **all versions of MariaDB Galera** since 10.1.

Key Numbers

Metric	Observed Value
Tested version	MariaDB 10.6.23 Galera
sysbench workload	16 threads, oltp_read_write
Accumulated writesets	188,516
IST duration	20-25 seconds
PMM wsrep*ist** variables	**Do not exist**
PmaControl detection	Automatic via log parsing

Recommendations

1. **Do not rely on the PMM Galera dashboard** for IST tracking if you run MariaDB — the panels will remain empty
2. **Size the GCache generously** (`gcache.size=2G` minimum) to maximize IST chances after a brief disconnection
3. **Centralize MariaDB logs** — they are the source of truth for Galera transfers
4. **Use PmaControl** for Galera monitoring that actually works on MariaDB, not only on PXC

Conclusion

The PMM dashboard for Galera IST is designed for Percona XtraDB Cluster. On MariaDB Galera, it shows nothing — not because IST does not work, but because the variables it queries do not exist.

IST works perfectly on MariaDB 10.6. You just need to know where to look: in the logs, not in phantom variables.