



# The Exciting MySQL 5.7 Replication Enhancements

**Luís Soares (luis.soares@oracle.com)**  
**Principal Software Engineer, MySQL Replication Team Lead**

# Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

# Program Agenda

# Program Agenda

- 1 ➤ Background
- 2 ➤ MySQL 5.7 Replication Generally Available
- 3 ➤ The New Replication Features in MySQL 5.7
- 4 ➤ Development Sneak Peek
- 5 ➤ Roadmap
- 6 ➤ Summary

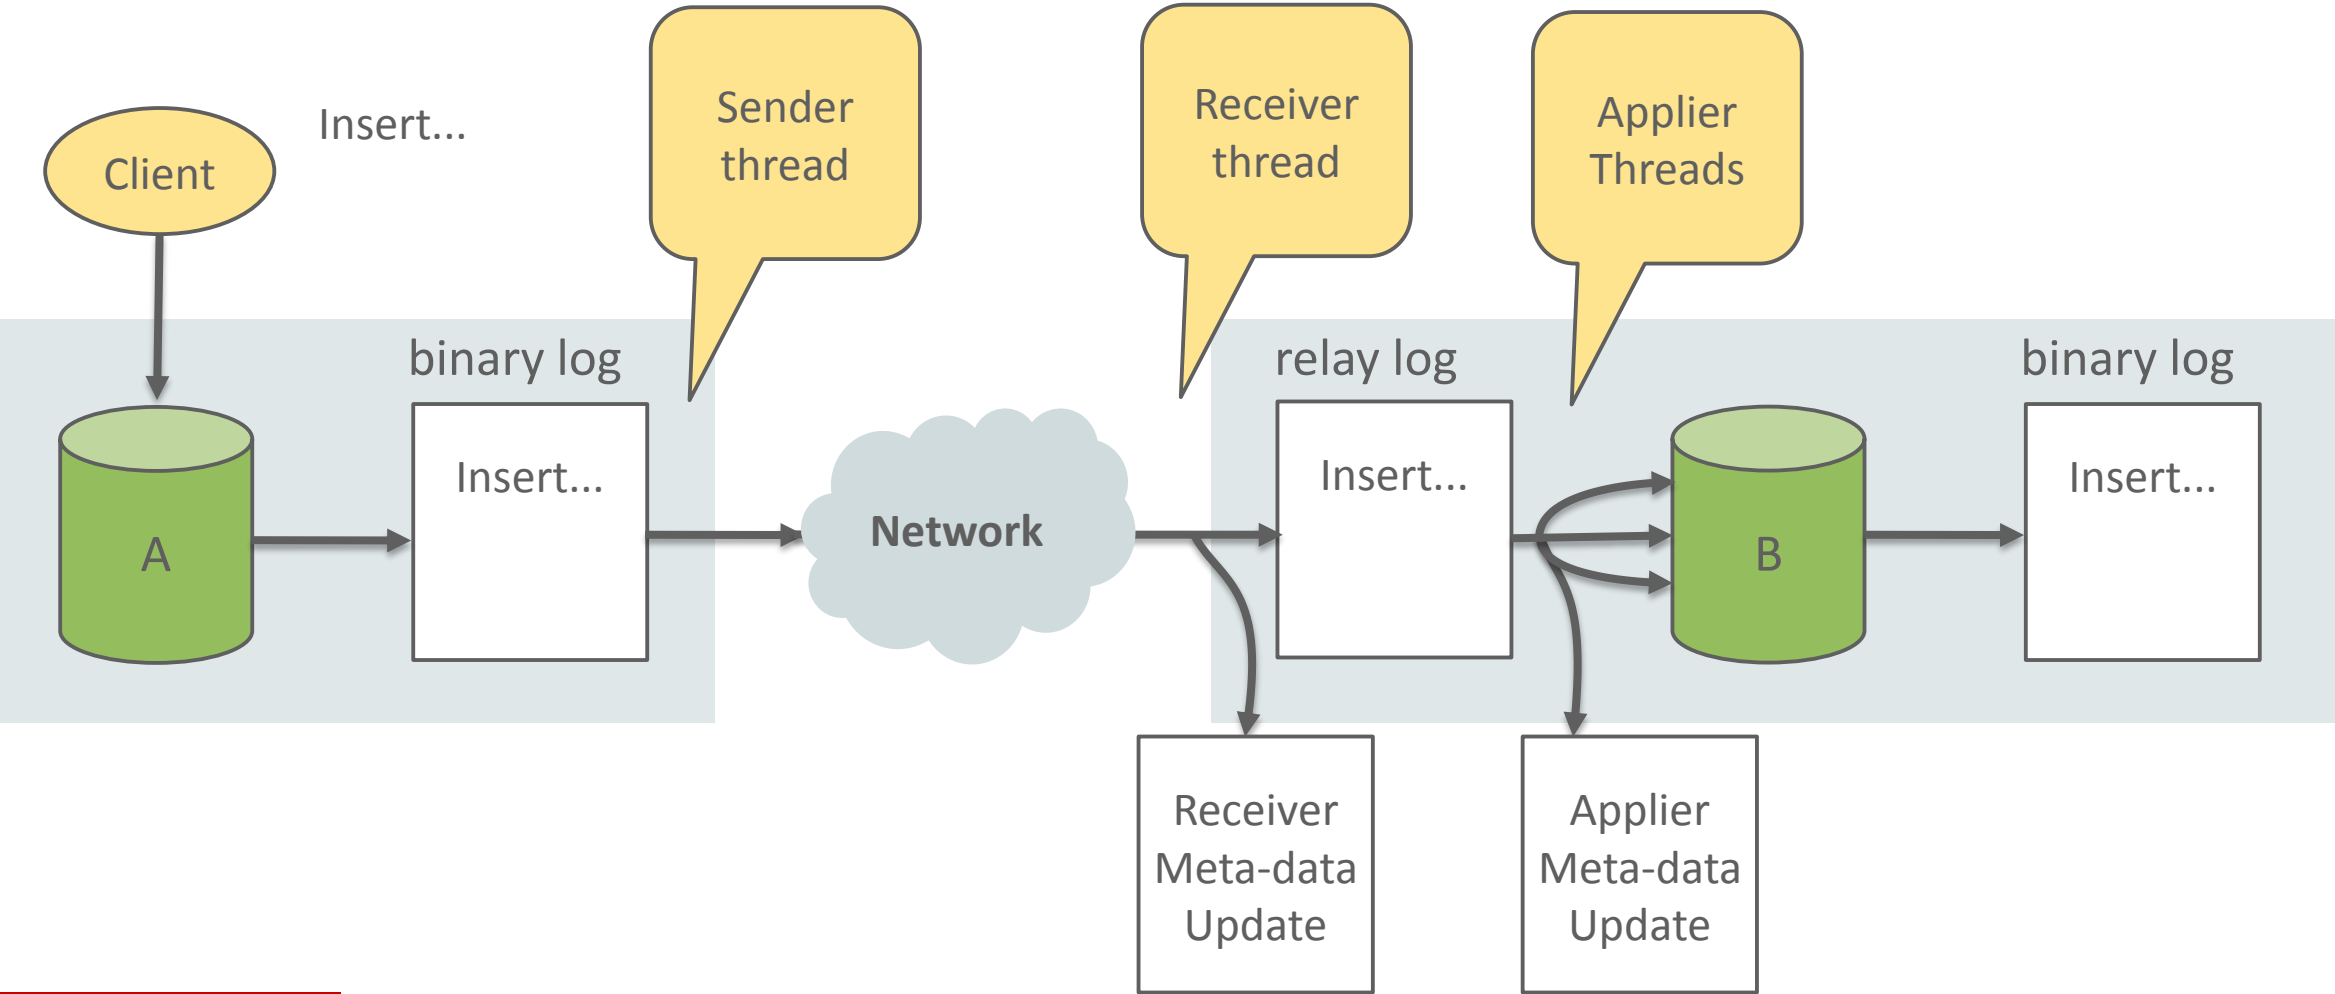
# 1 MySQL 5.7 Replication Generally Available

# Interesting Stats Related to MySQL 5.7

- **40 replication worklogs** pushed to MySQL 5.7
  - 29 related to on MySQL replication Core
  - 11 are ground work for MySQL Group Replication.
- **8 contributions** merged into MySQL 5.7.
- **19 major enhancements** to the replication core.
- **14 refactoring and modularization** related changes to the code.

## 2 Background

# Background: Replication Components

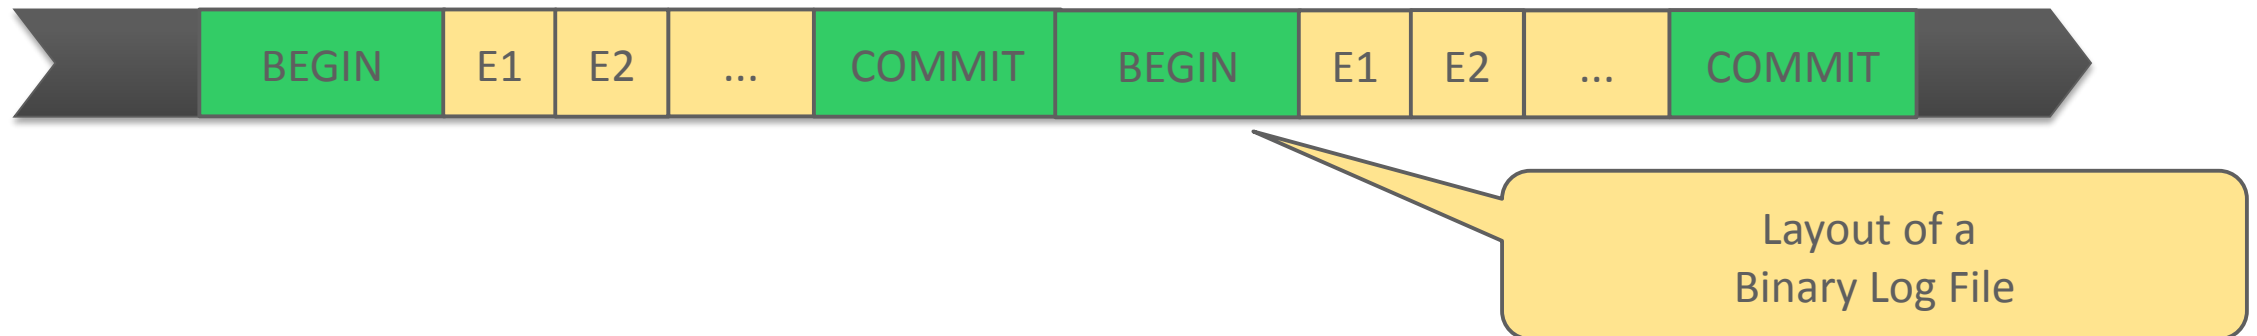




# Background: Replication Components

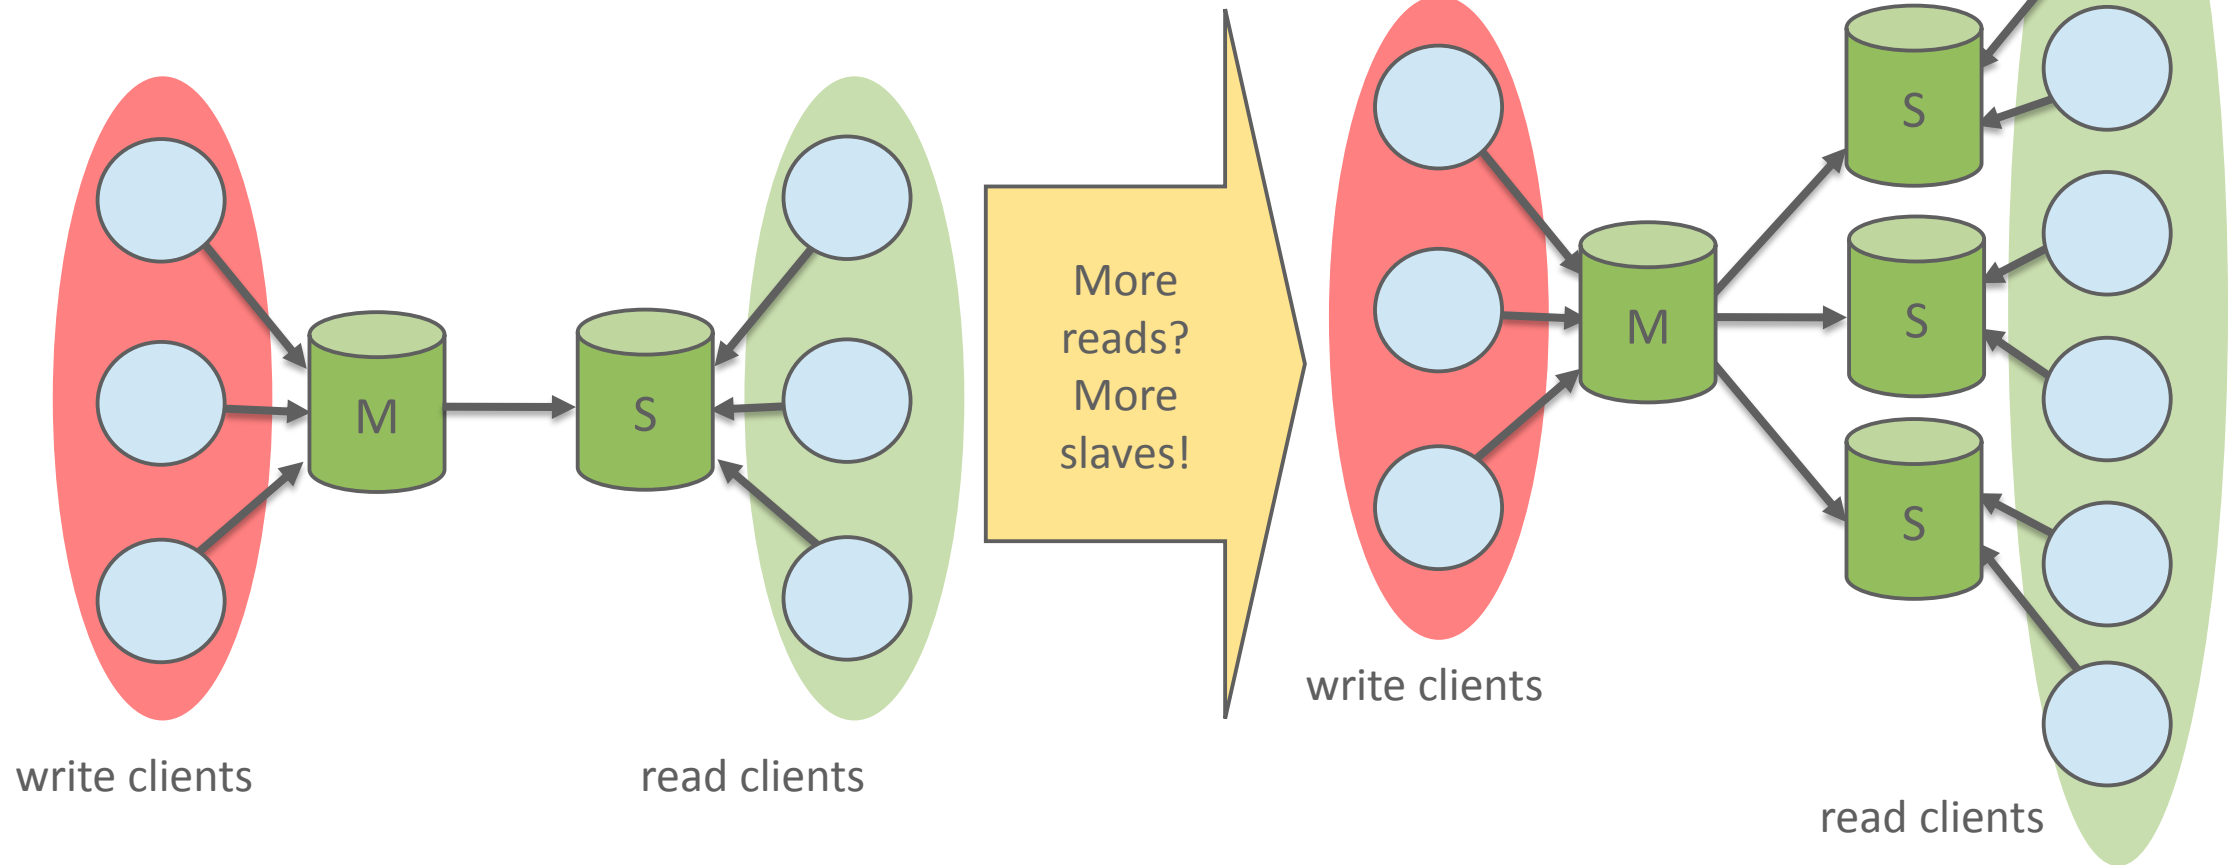
- **Binary Log**

- File based logical log that records the changes on the master.
- Statement or Row based format (may be intermixed).
- Each transaction is split into groups of events.
- Control events: Rotate, Format Description, Gtid, ...




# Background: What is Replication Used For?

## Read scale-out



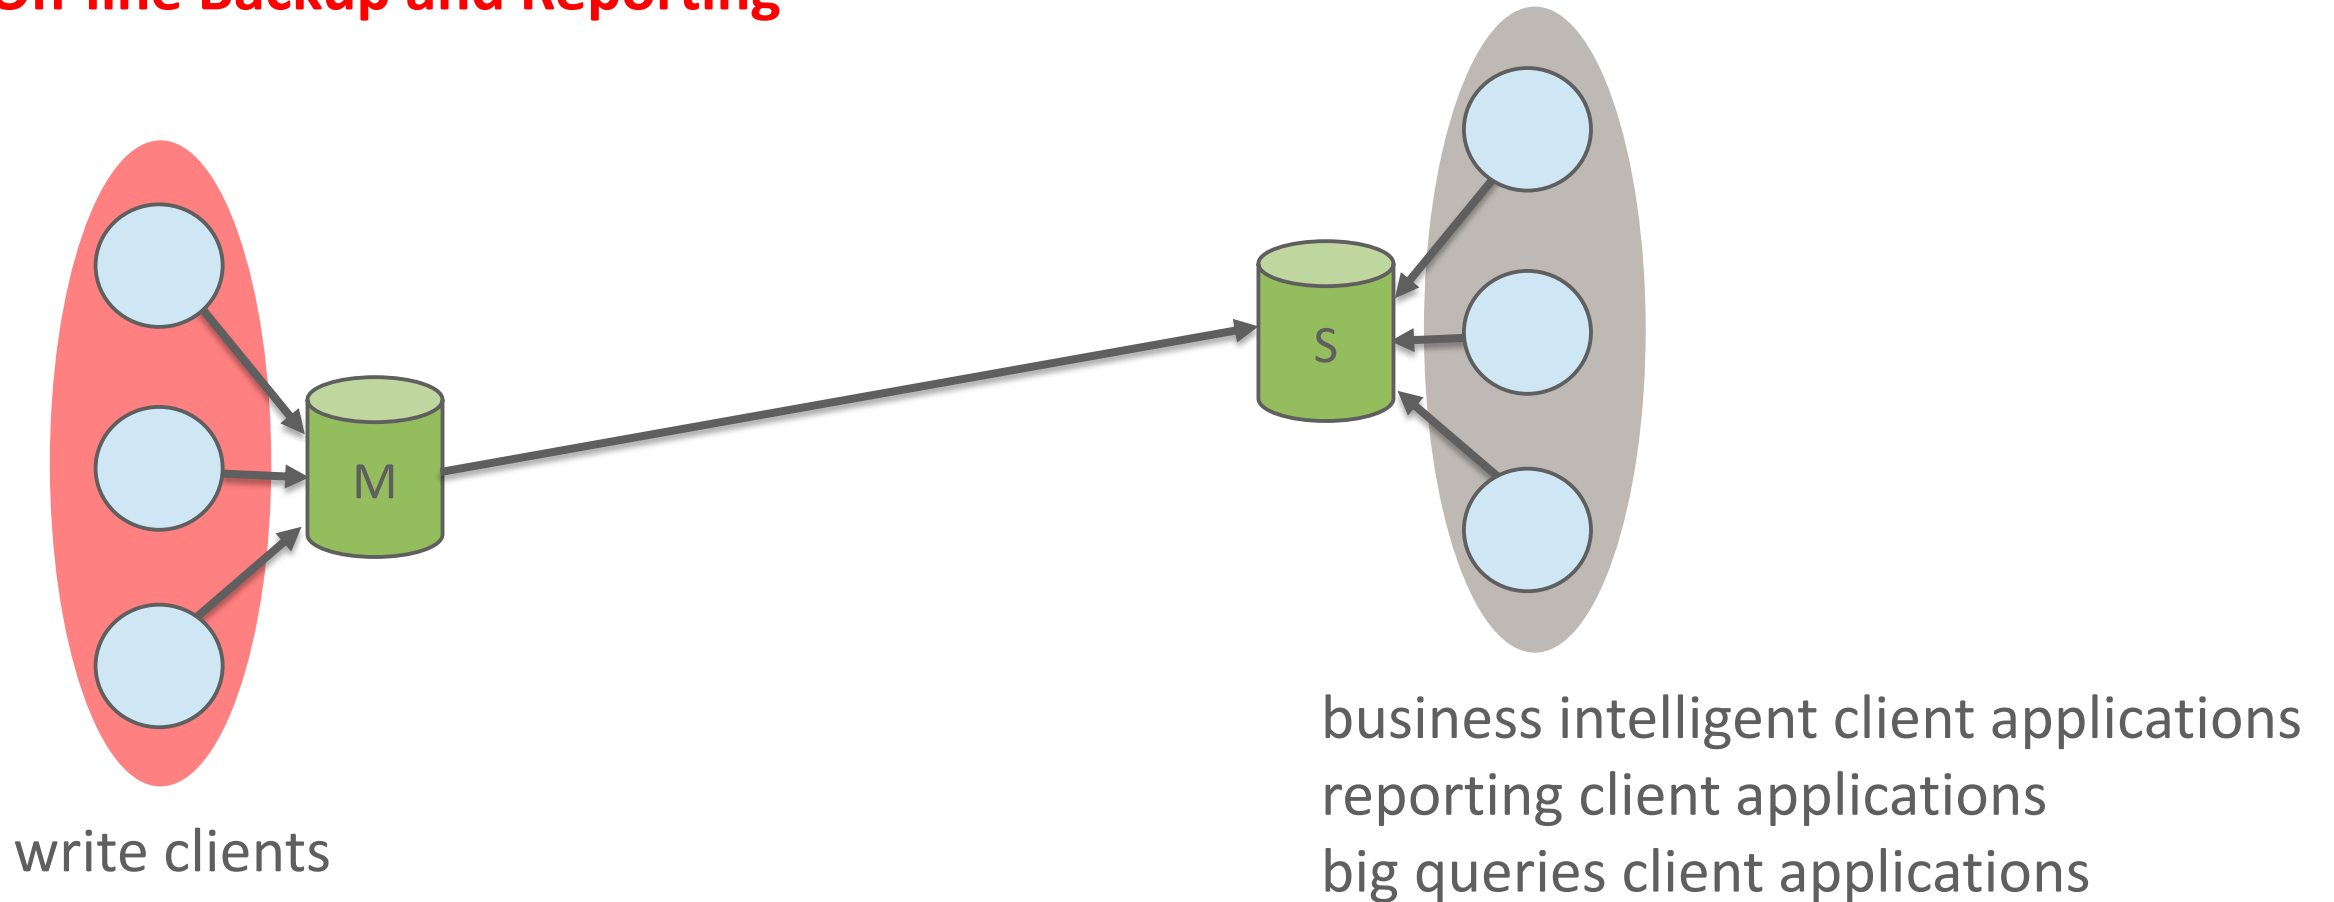
# Background: What is Replication Used For?

**Redundancy:** If master crashes, **promote** slave to master

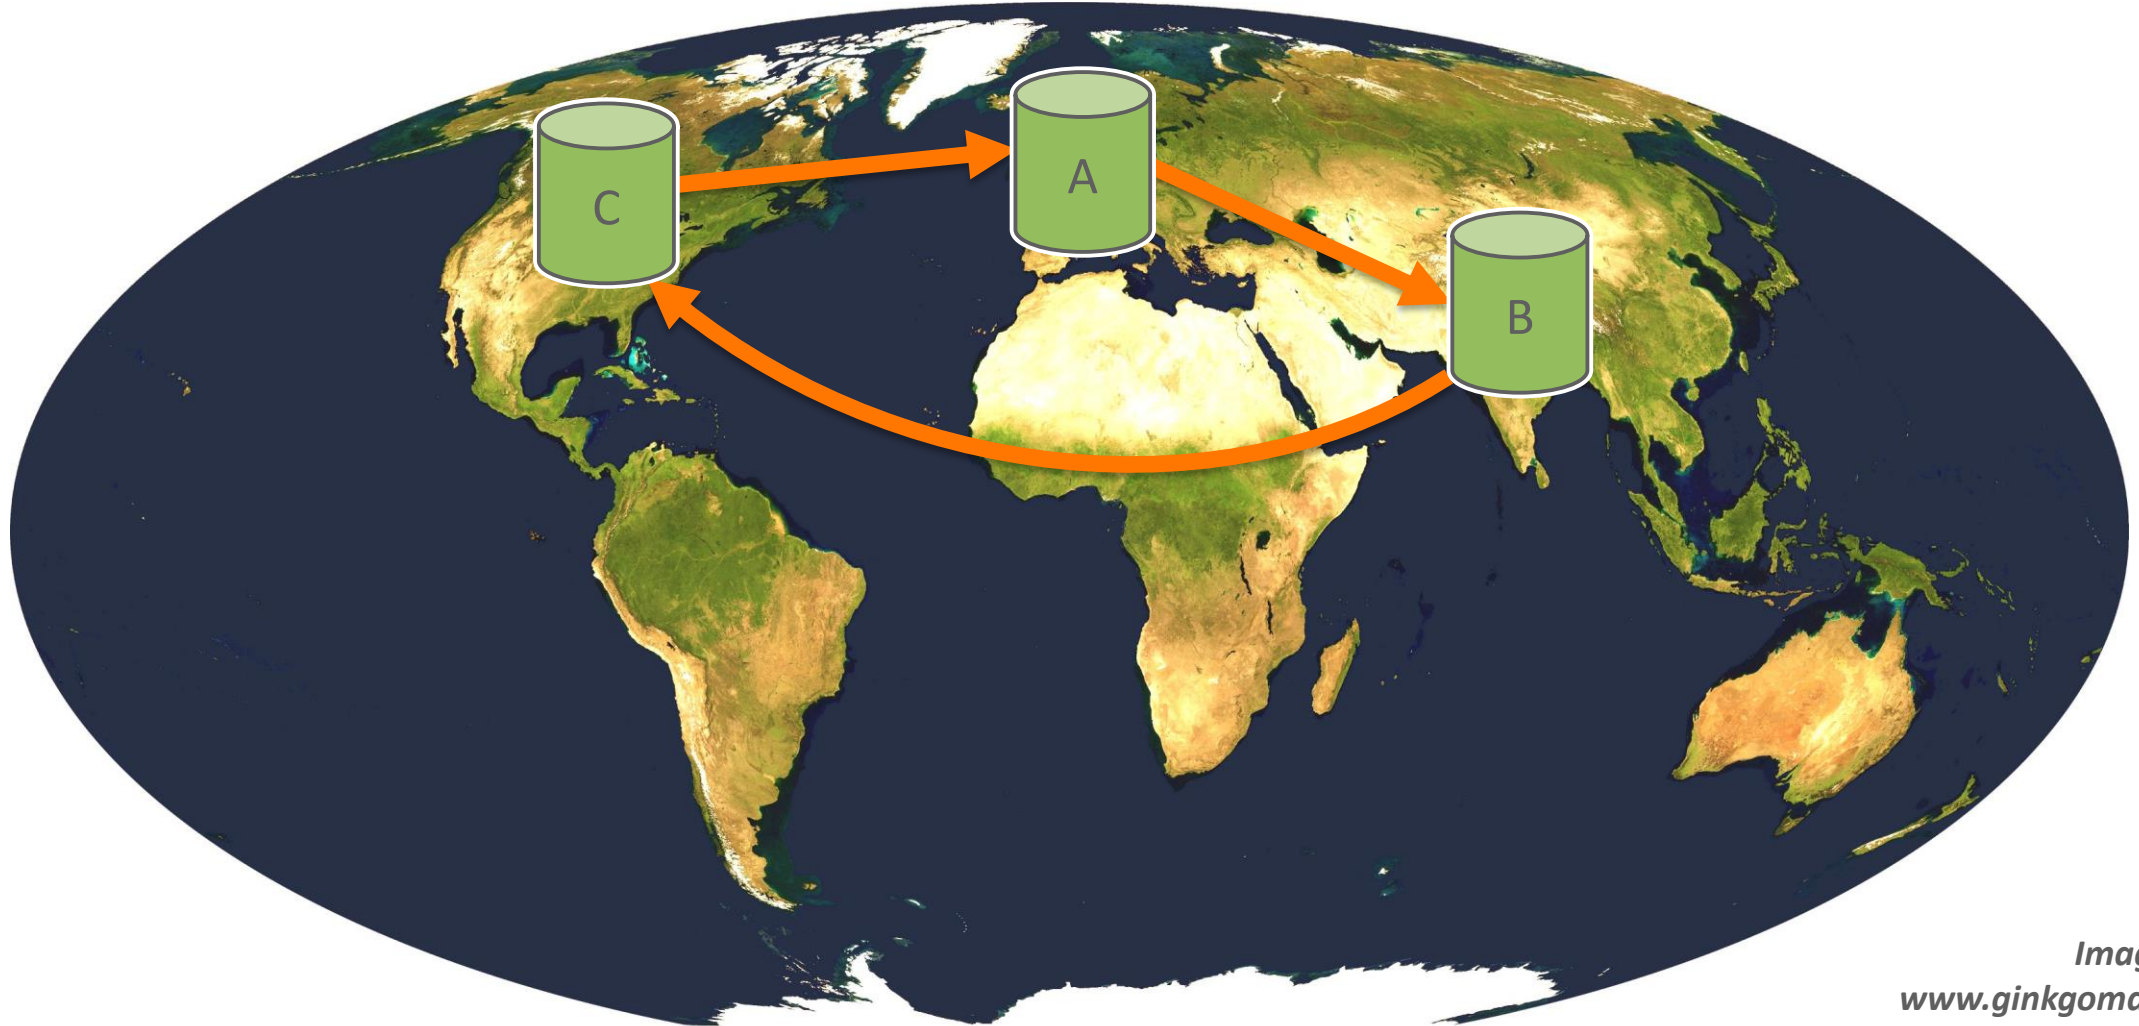


# Background: What is Replication Used For?

## On-line Backup and Reporting



# Background: What is Replication Used For?



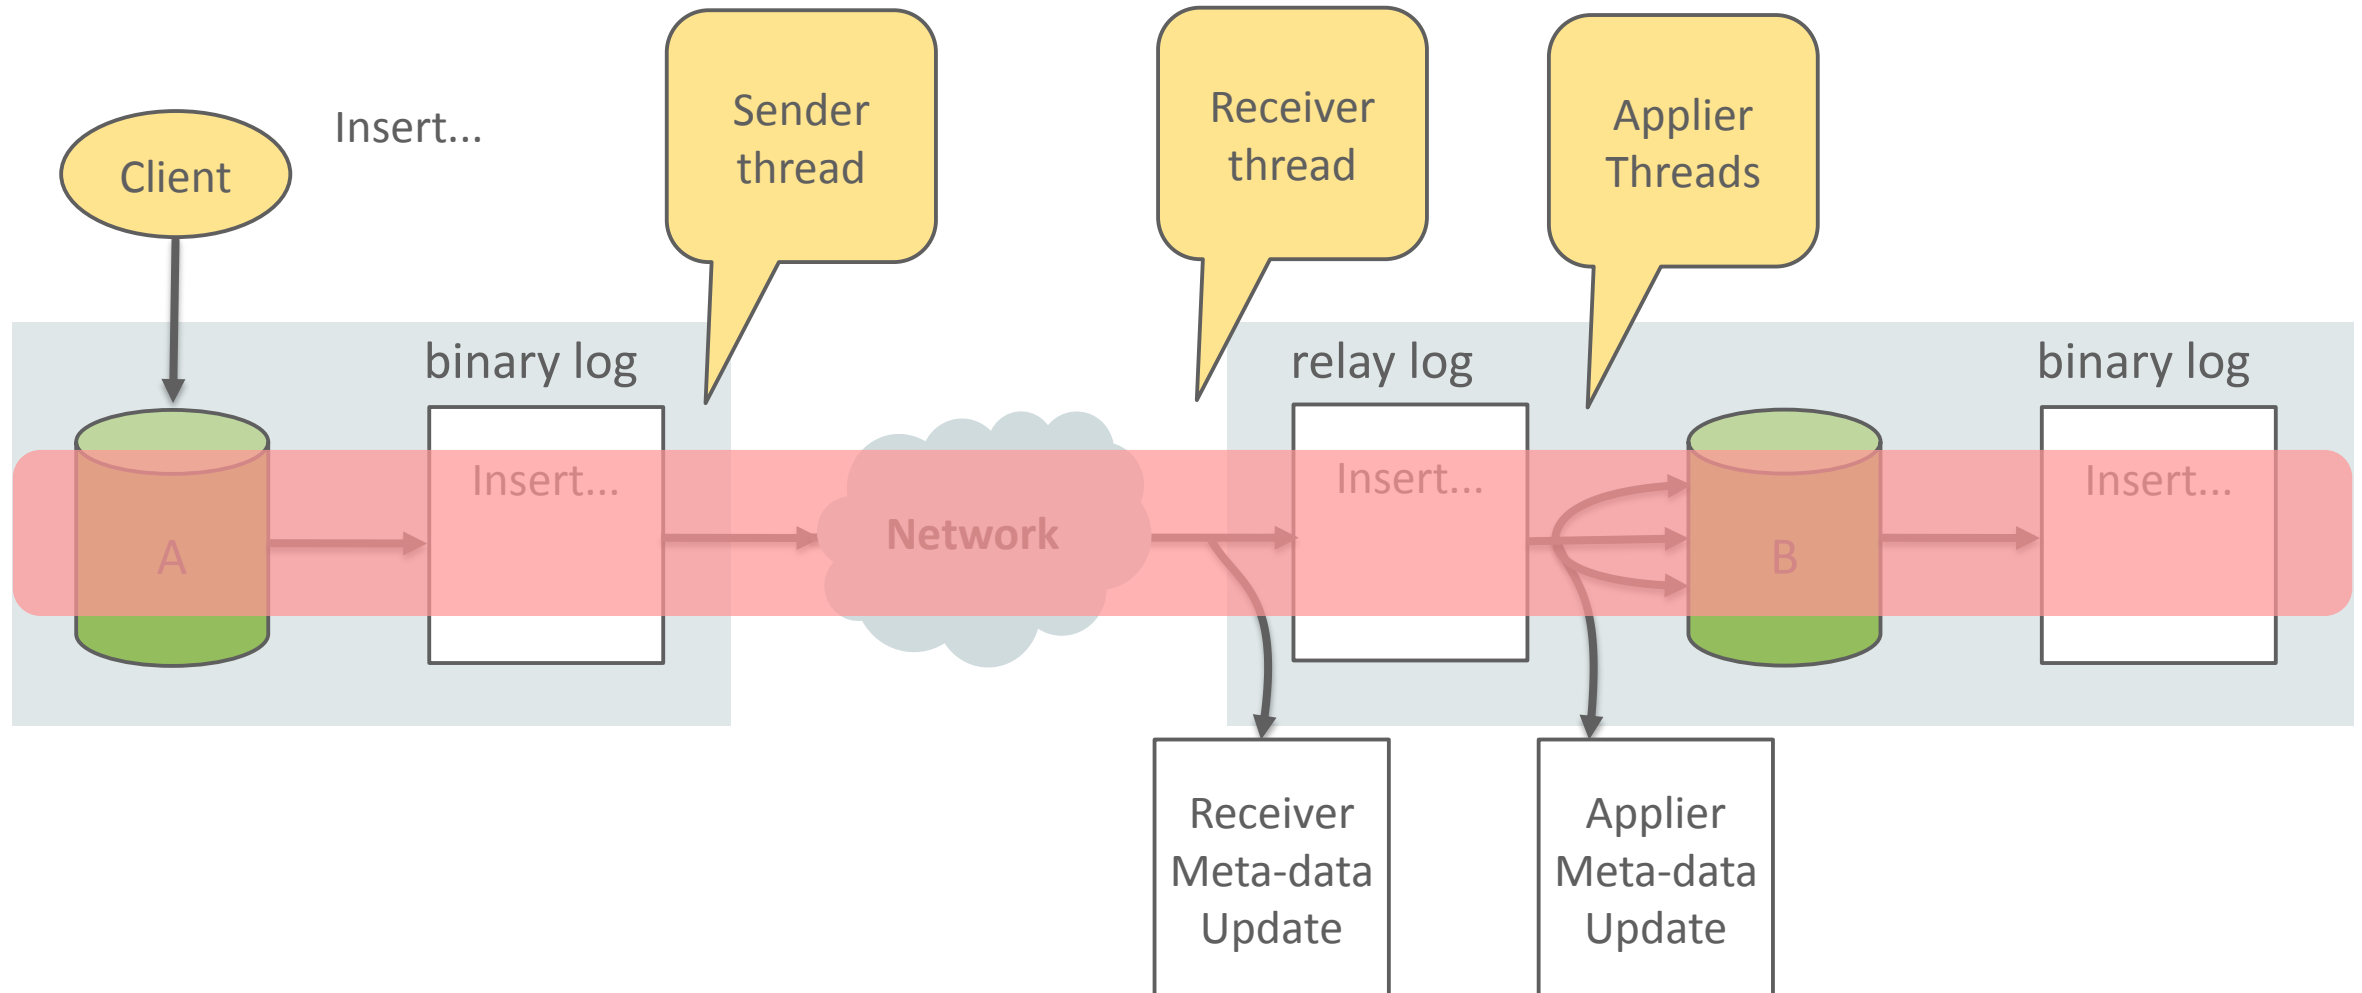
*Image from  
[www.ginkgomaps.com](http://www.ginkgomaps.com)*

# 3 The New Replication Features in MySQL 5.7

# 3 The New Replication Features in MySQL 5.7

## 3.1 Usability / Online

# Online Reconfiguration of Global Transaction Identifiers





# Online Reconfiguration of Global Transaction Identifiers

- The procedure is online.
  - Both reads and writes are allowed while global transaction identifiers are being turned on or off in the entire replication cluster.
- No need to synchronize servers, ever.
- No need to restart servers.
- No need to change replication topology.
  - Works in arbitrary topologies.
- Should anything happen in the middle of the procedure, you can always roll back.

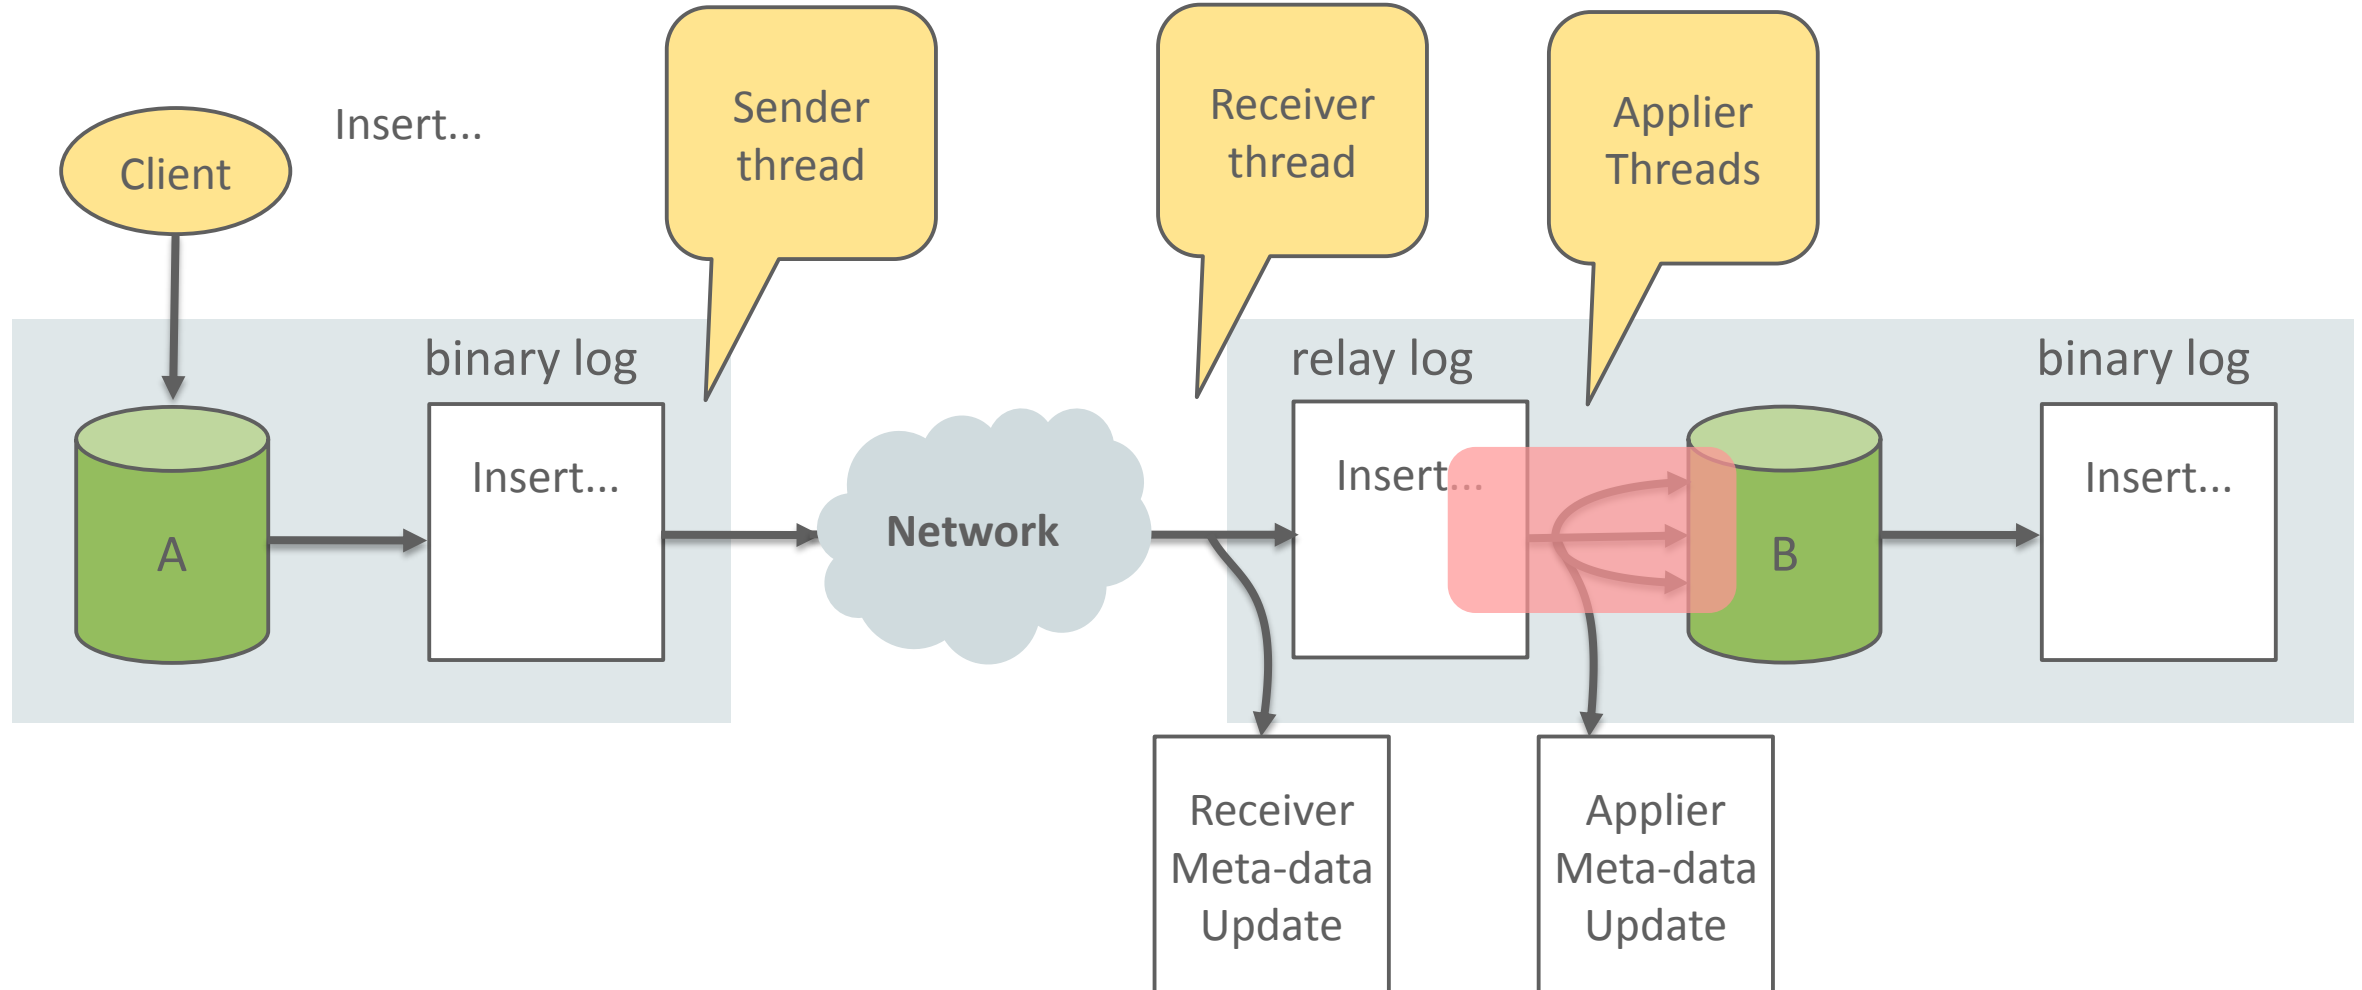
# Online Reconfiguration of Global Transaction Identifiers

- Over simplified procedure to turn on global transaction identifiers:

- 1) SET @@GLOBAL.GTID\_MODE = OFF\_PERMISSIVE; (on every server)
- 2) SET @@GLOBAL.GTID\_MODE = ON\_PERMISSIVE; (on every server)
- 3) wait for a bit longer than your replication lag
- 4) SET @@GLOBAL.GTID\_MODE = ON; (on every server)

- Details: <http://dev.mysql.com/doc/refman/5.7/en/replication-mode-change-online-enable-gtids.html>

# Online Reconfiguration of Replication Filters

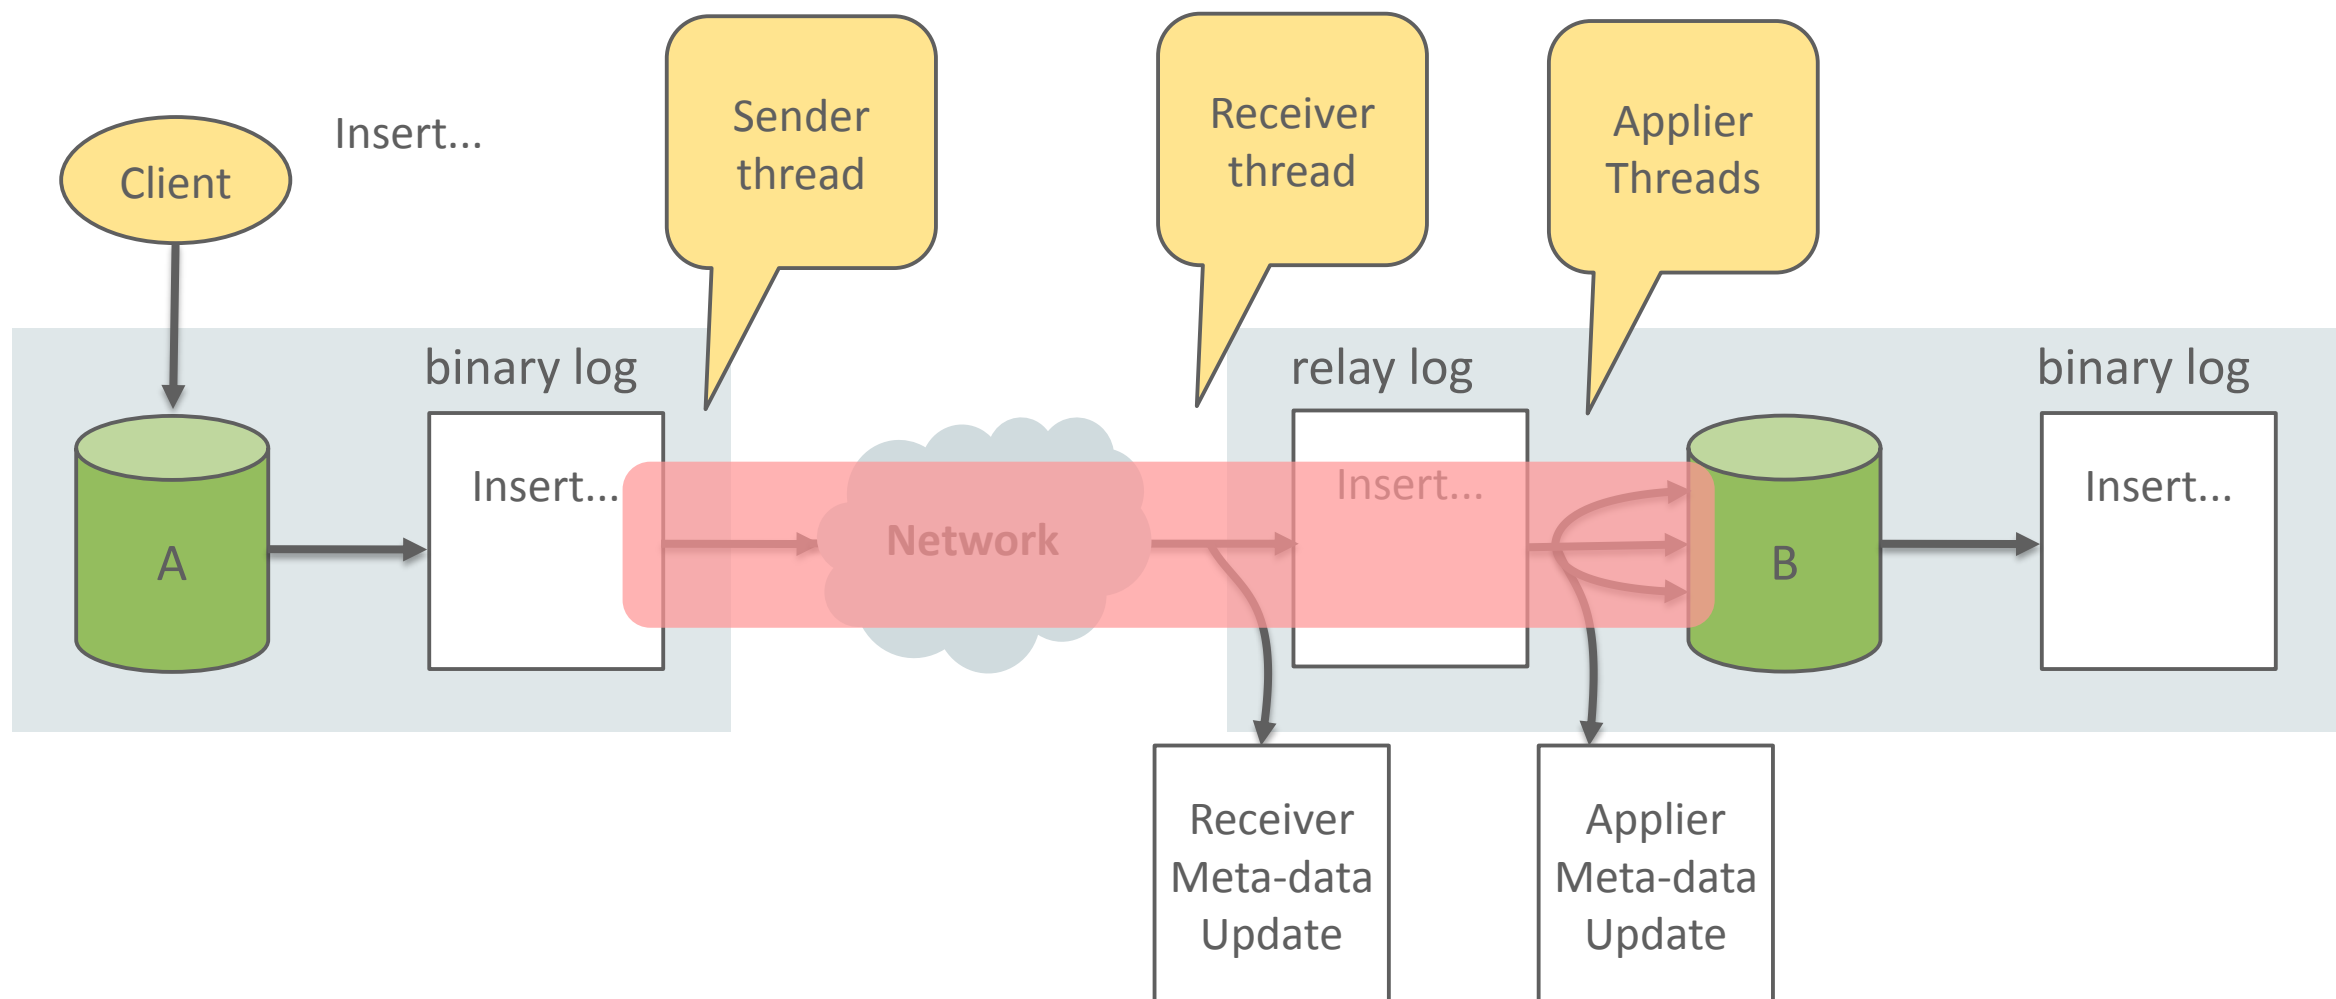


# Online Reconfiguration of Replication Filters

- Change Slave's Replication Filters dynamically.
  - No need to stop and restart the slave server for setting new replication filtering rules.
  - All slave filters are supported.
  - Values can be input in various character sets.


```
mysql> CHANGE REPLICATION FILTER REPLICATE_DO_DB= (db1, db2)
```

# Online Reconfiguration of Replication Receiver/Applier



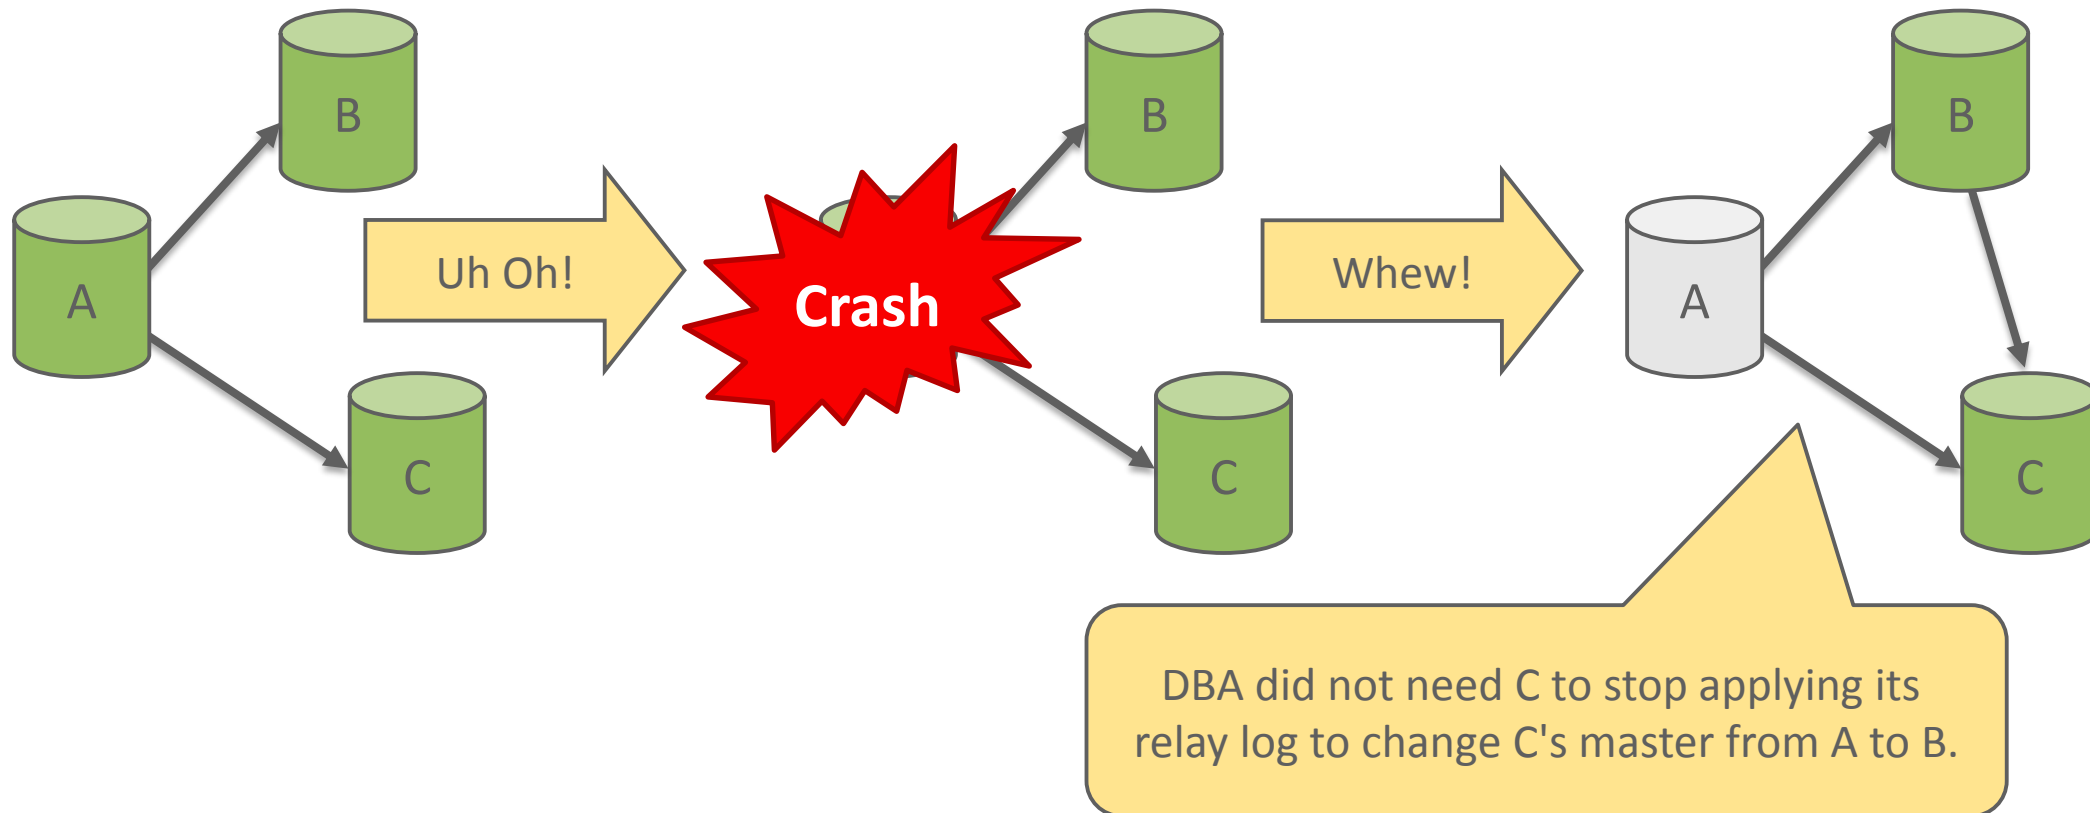
# Online Reconfiguration of Replication Receiver/Applier

Change master from A to B without stopping the applier threads.



# Online Reconfiguration of Replication Receiver/Applier

Change master from A to B without stopping the applier threads.



# Online Reconfiguration of Replication Receiver/Applier

- Enables more online operations during fail-over:

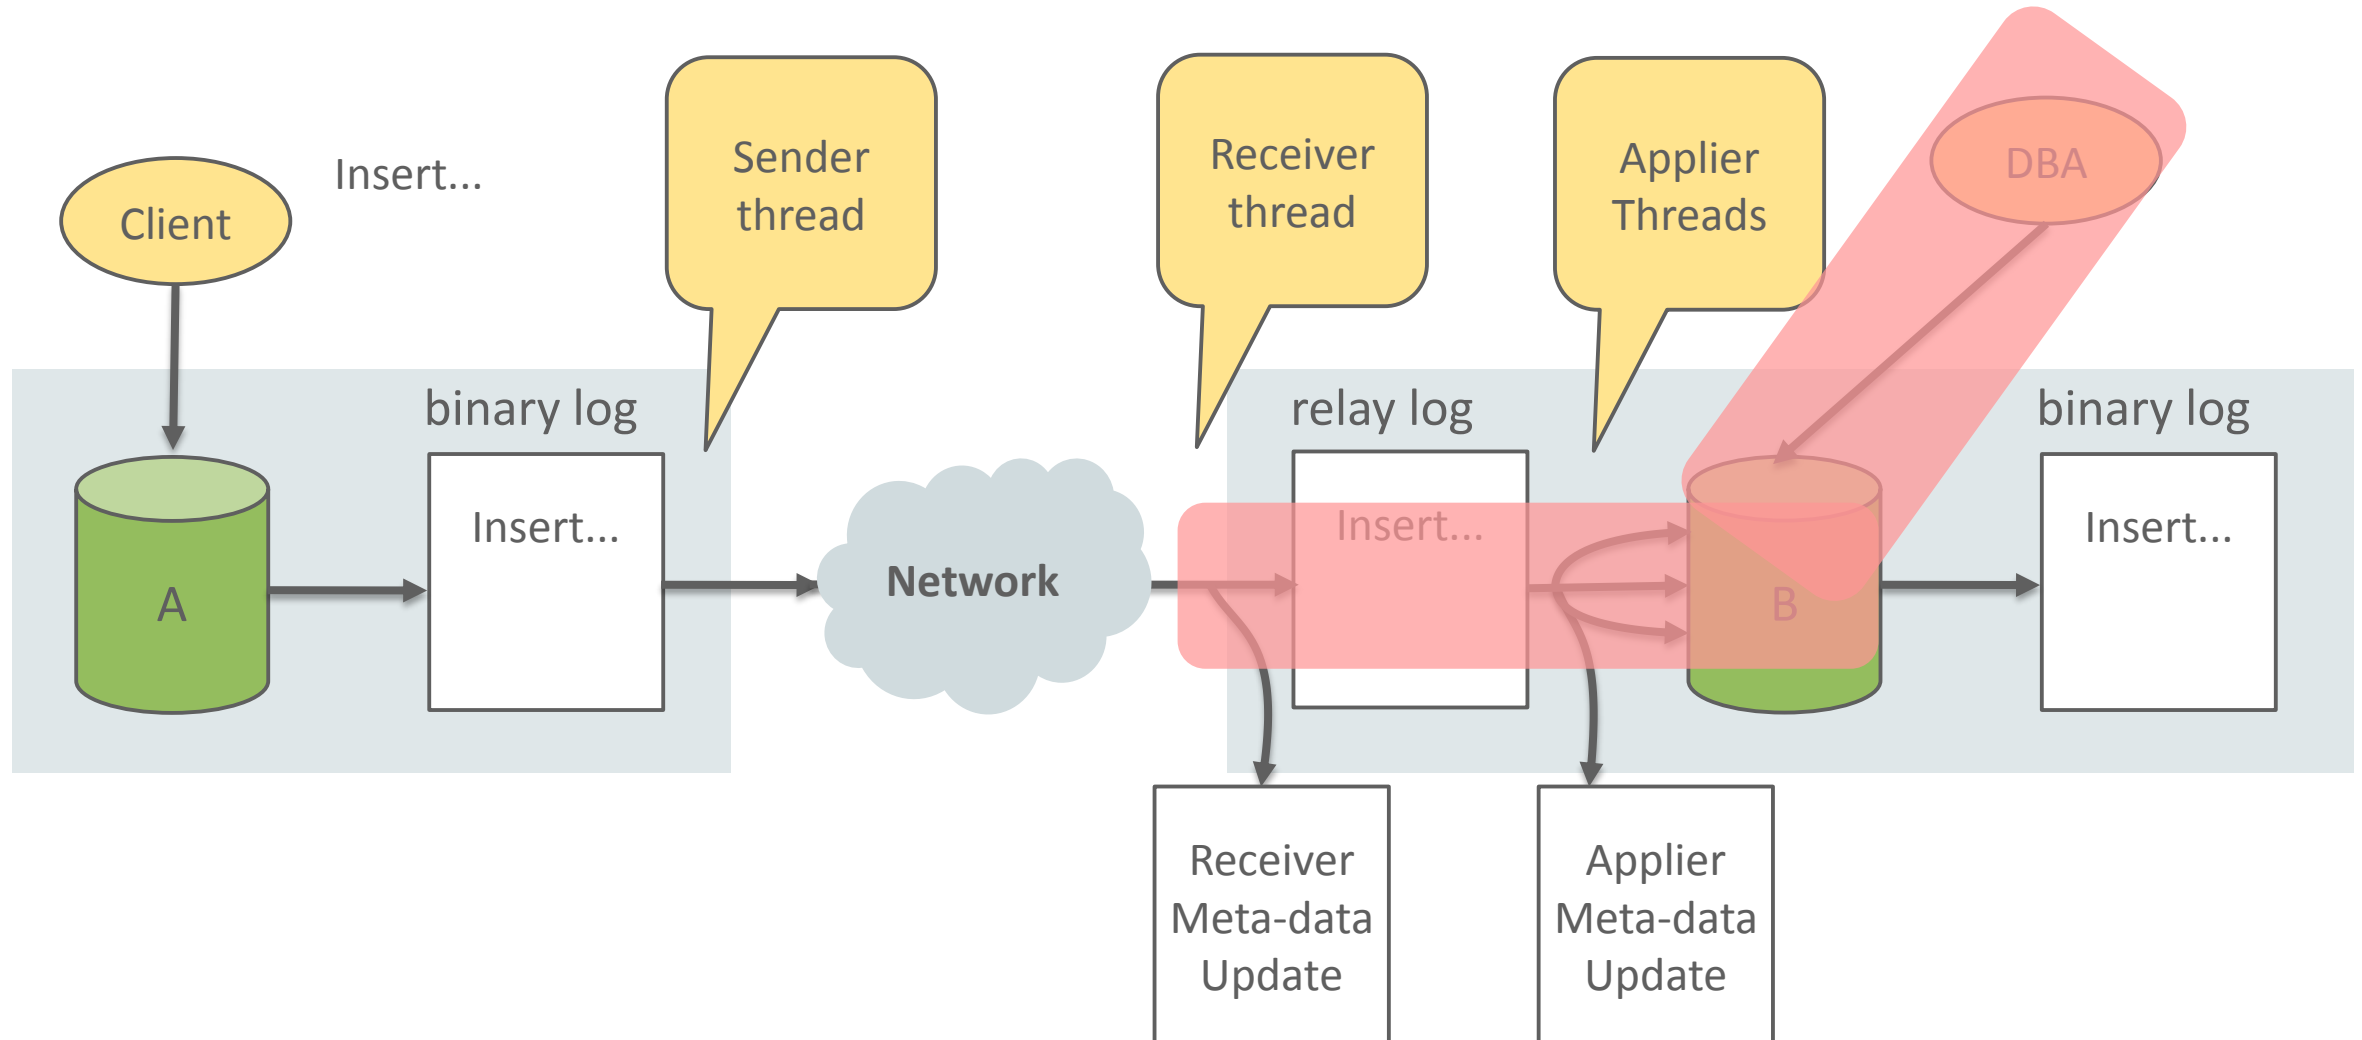
```
mysql> STOP SLAVE IO_THREAD;  
mysql> CHANGE MASTER TO MASTER_HOST='master2', ...;  
mysql> START SLAVE IO_THREAD;
```

- Stopping, changing master and restarting the receiver thread are all done while the applier thread is running.
- Change applier properties while the receiver thread is working:

```
mysql> STOP SLAVE SQL_THREAD;  
mysql> CHANGE MASTER TO MASTER_DELAY=3600, ...;  
mysql> START SLAVE SQL_THREAD;
```



# Improved Replication Monitoring

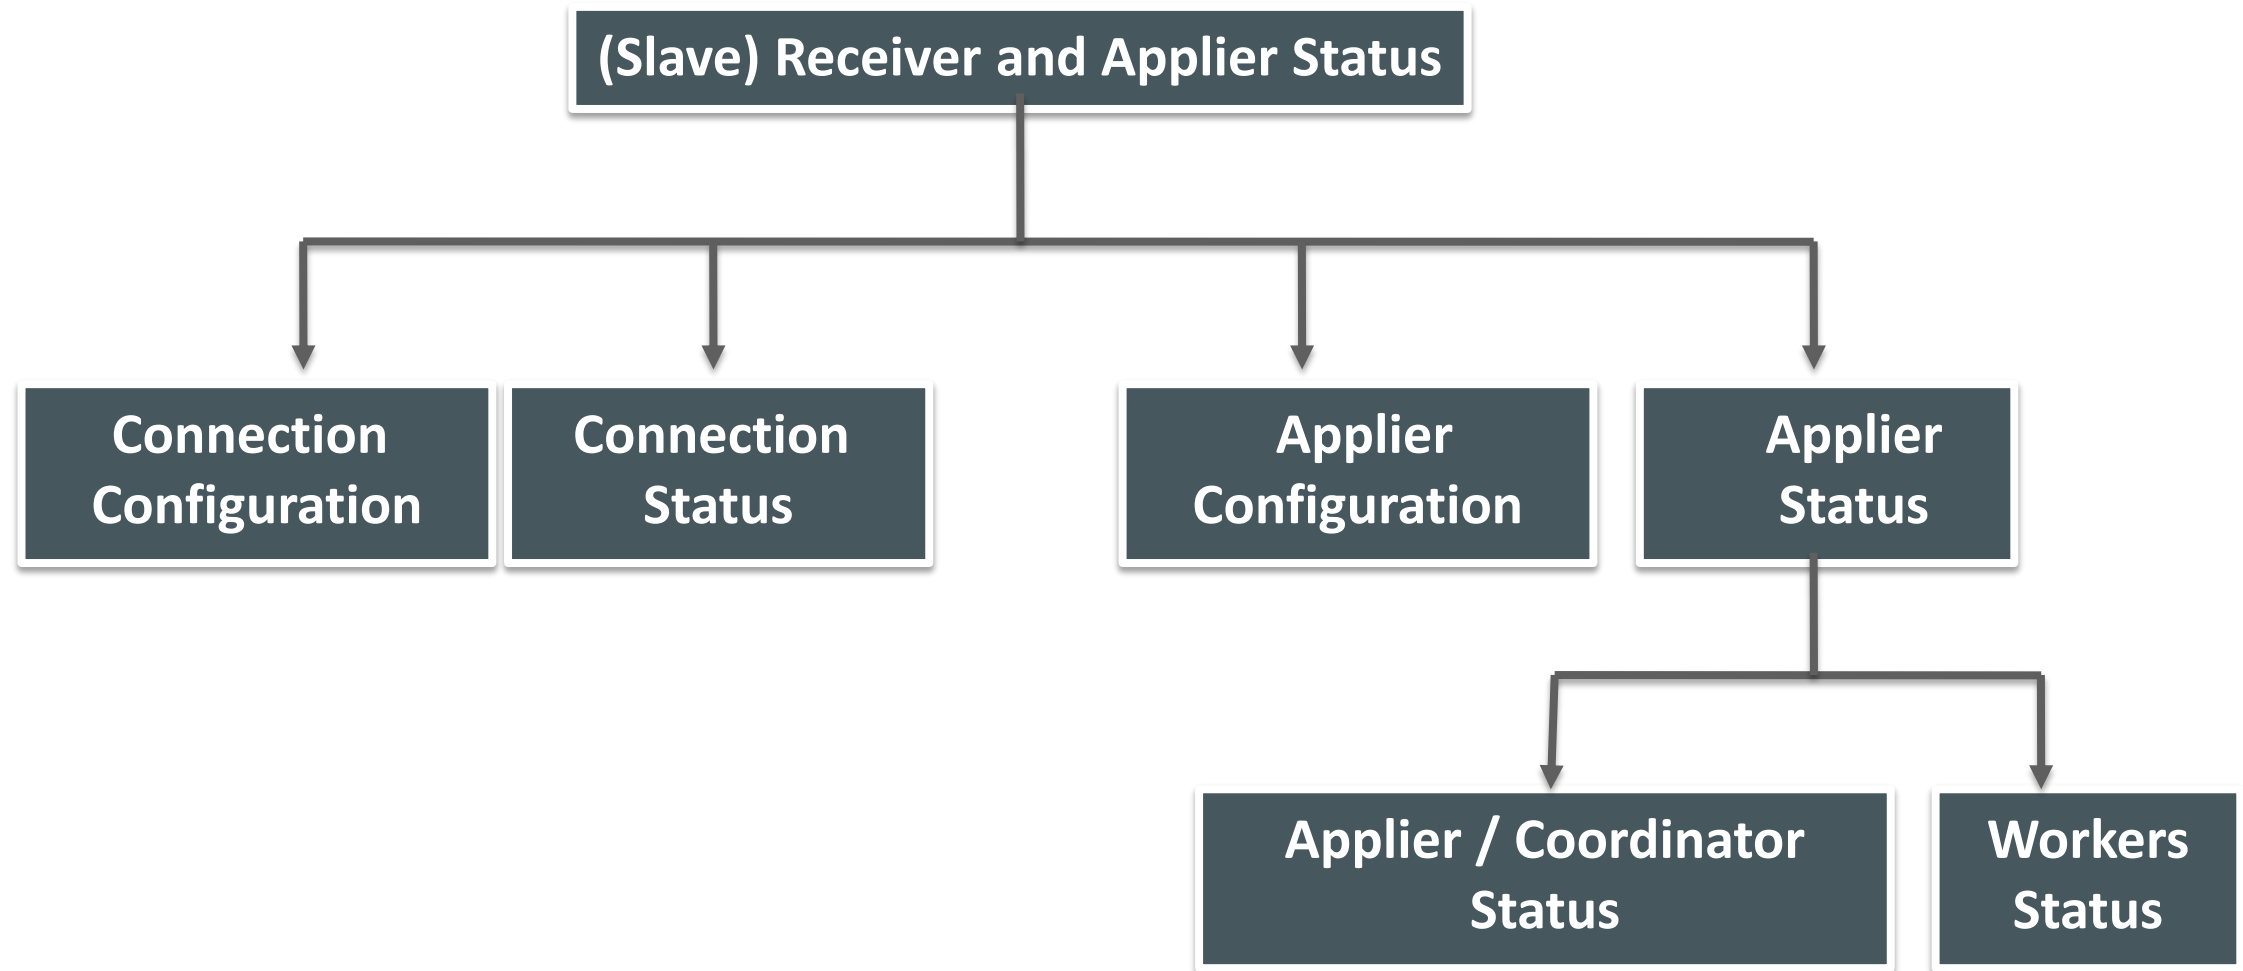


# Improved Replication Monitoring

## Performance Schema Replication Tables

- Monitoring through SQL.
- Grouping logically related information.
- Extensible, can be adapted to new features.
- More accurate and consistent identifier names.
- Master-slave, Multi-source, Group Replication support.

# Improved Replication Monitoring



# Improved Replication Monitoring

```
mysql> select * from performance_schema.replication_applier_status_by_worker\G
```

```
***** 1. row *****
```

```
CHANNEL_NAME:
```

```
WORKER_ID: 1
```

```
THREAD_ID: 35
```

```
SERVICE_STATE: ON
```

```
LAST_SEEN_TRANSACTION: 4ba0eb86-63c3-11e4-92ba-28b2bd168d07:2368
```

```
LAST_ERROR_NUMBER: 0
```

```
LAST_ERROR_MESSAGE:
```

```
LAST_ERROR_TIMESTAMP: 0000-00-00 00:00:00
```

```
***** 2. row *****
```

```
CHANNEL_NAME:
```

```
WORKER_ID: 2
```

```
THREAD_ID: 36
```

```
SERVICE_STATE: ON
```

```
LAST_SEEN_TRANSACTION: 4ba0eb86-63c3-11e4-92ba-28b2bd168d07:2367
```

```
LAST_ERROR_NUMBER: 0
```

```
LAST_ERROR_MESSAGE:
```

```
LAST_ERROR_TIMESTAMP: 0000-00-00 00:00:00
```

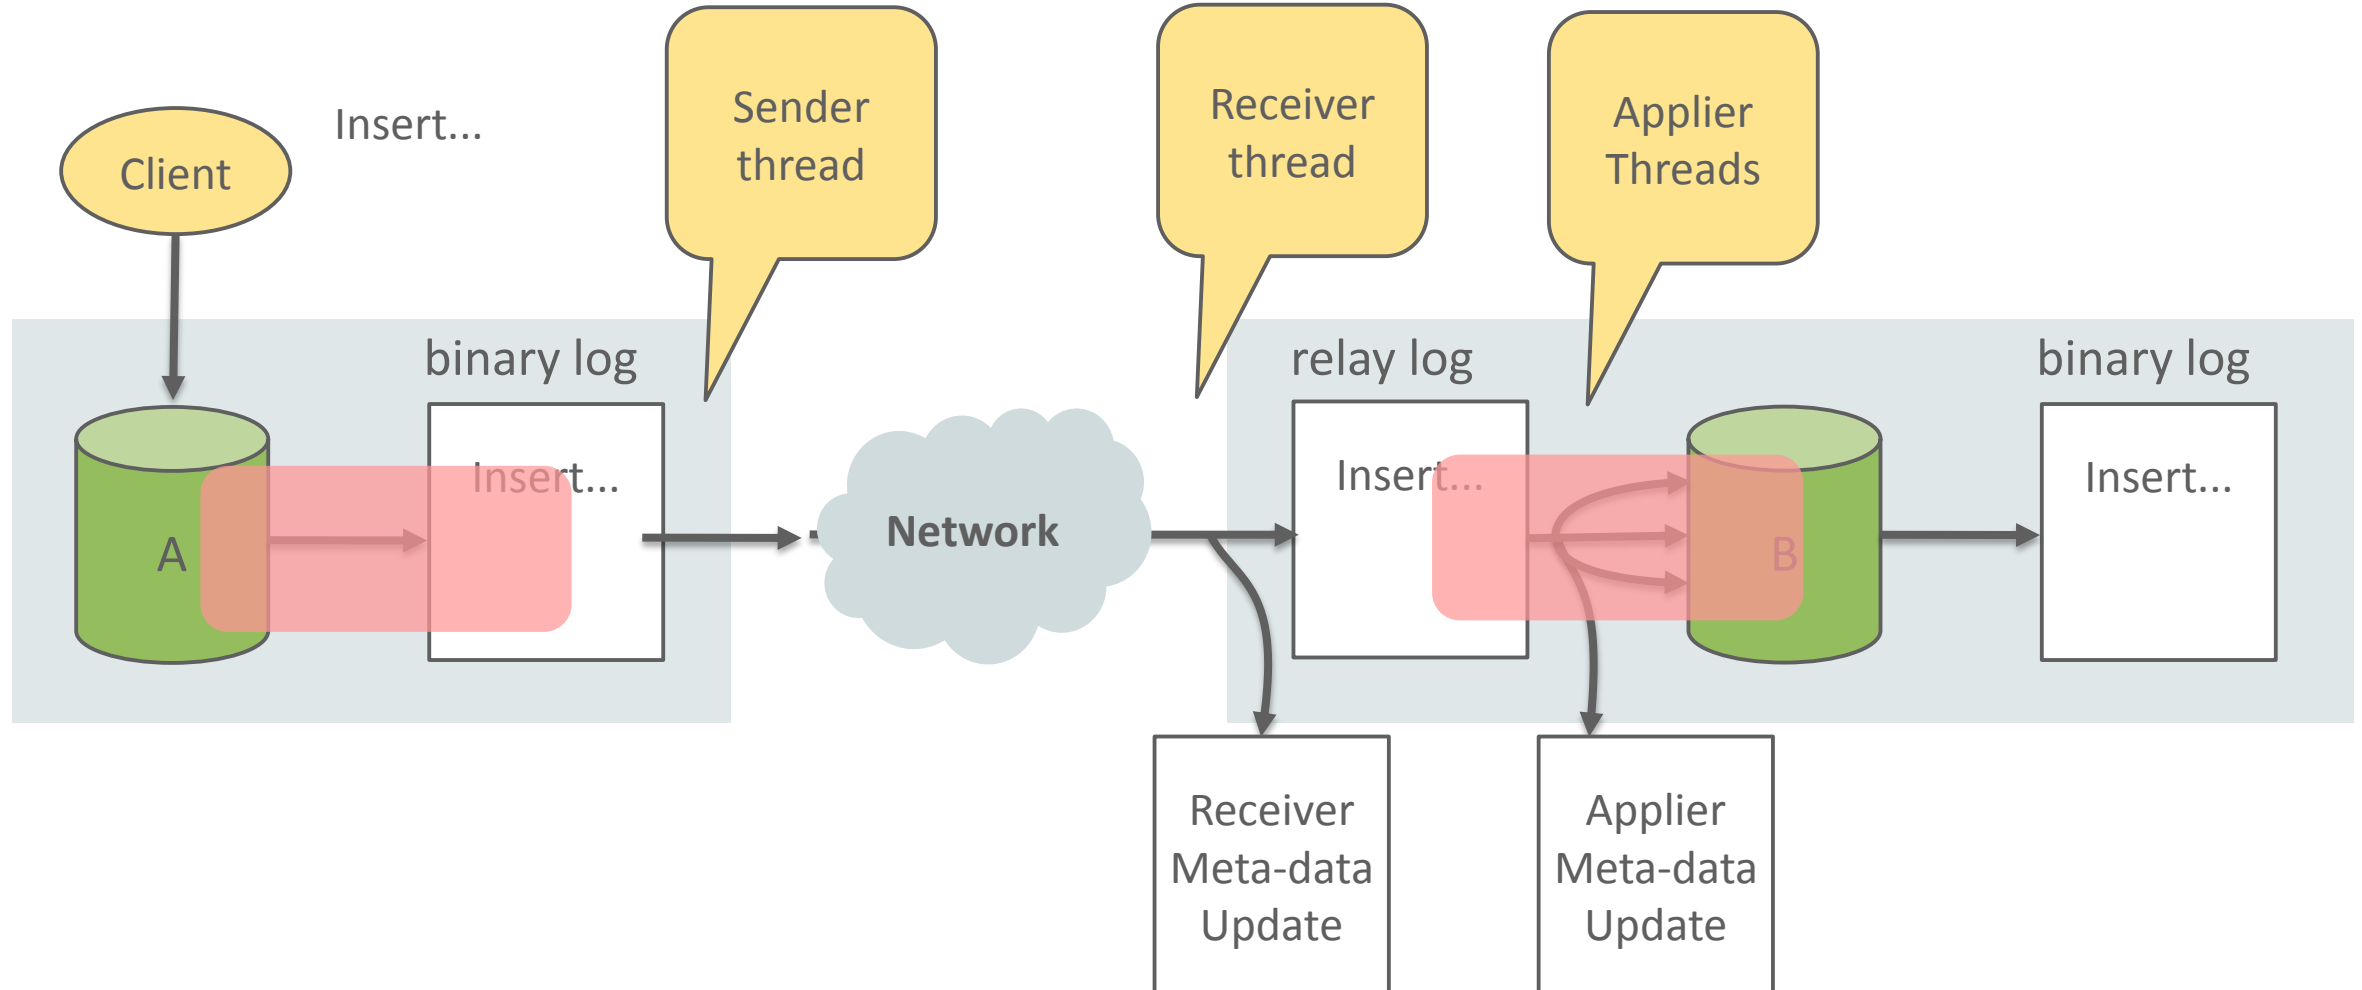
```
2 rows in set (0,00 sec)
```

# 3 The New Replication Features in MySQL 5.7


3.1 Usability / Online

3.2 Performance

# Improved Applier Throughput – Locking-based Parallelism



# Improved Applier Throughput – Locking-based Parallelism

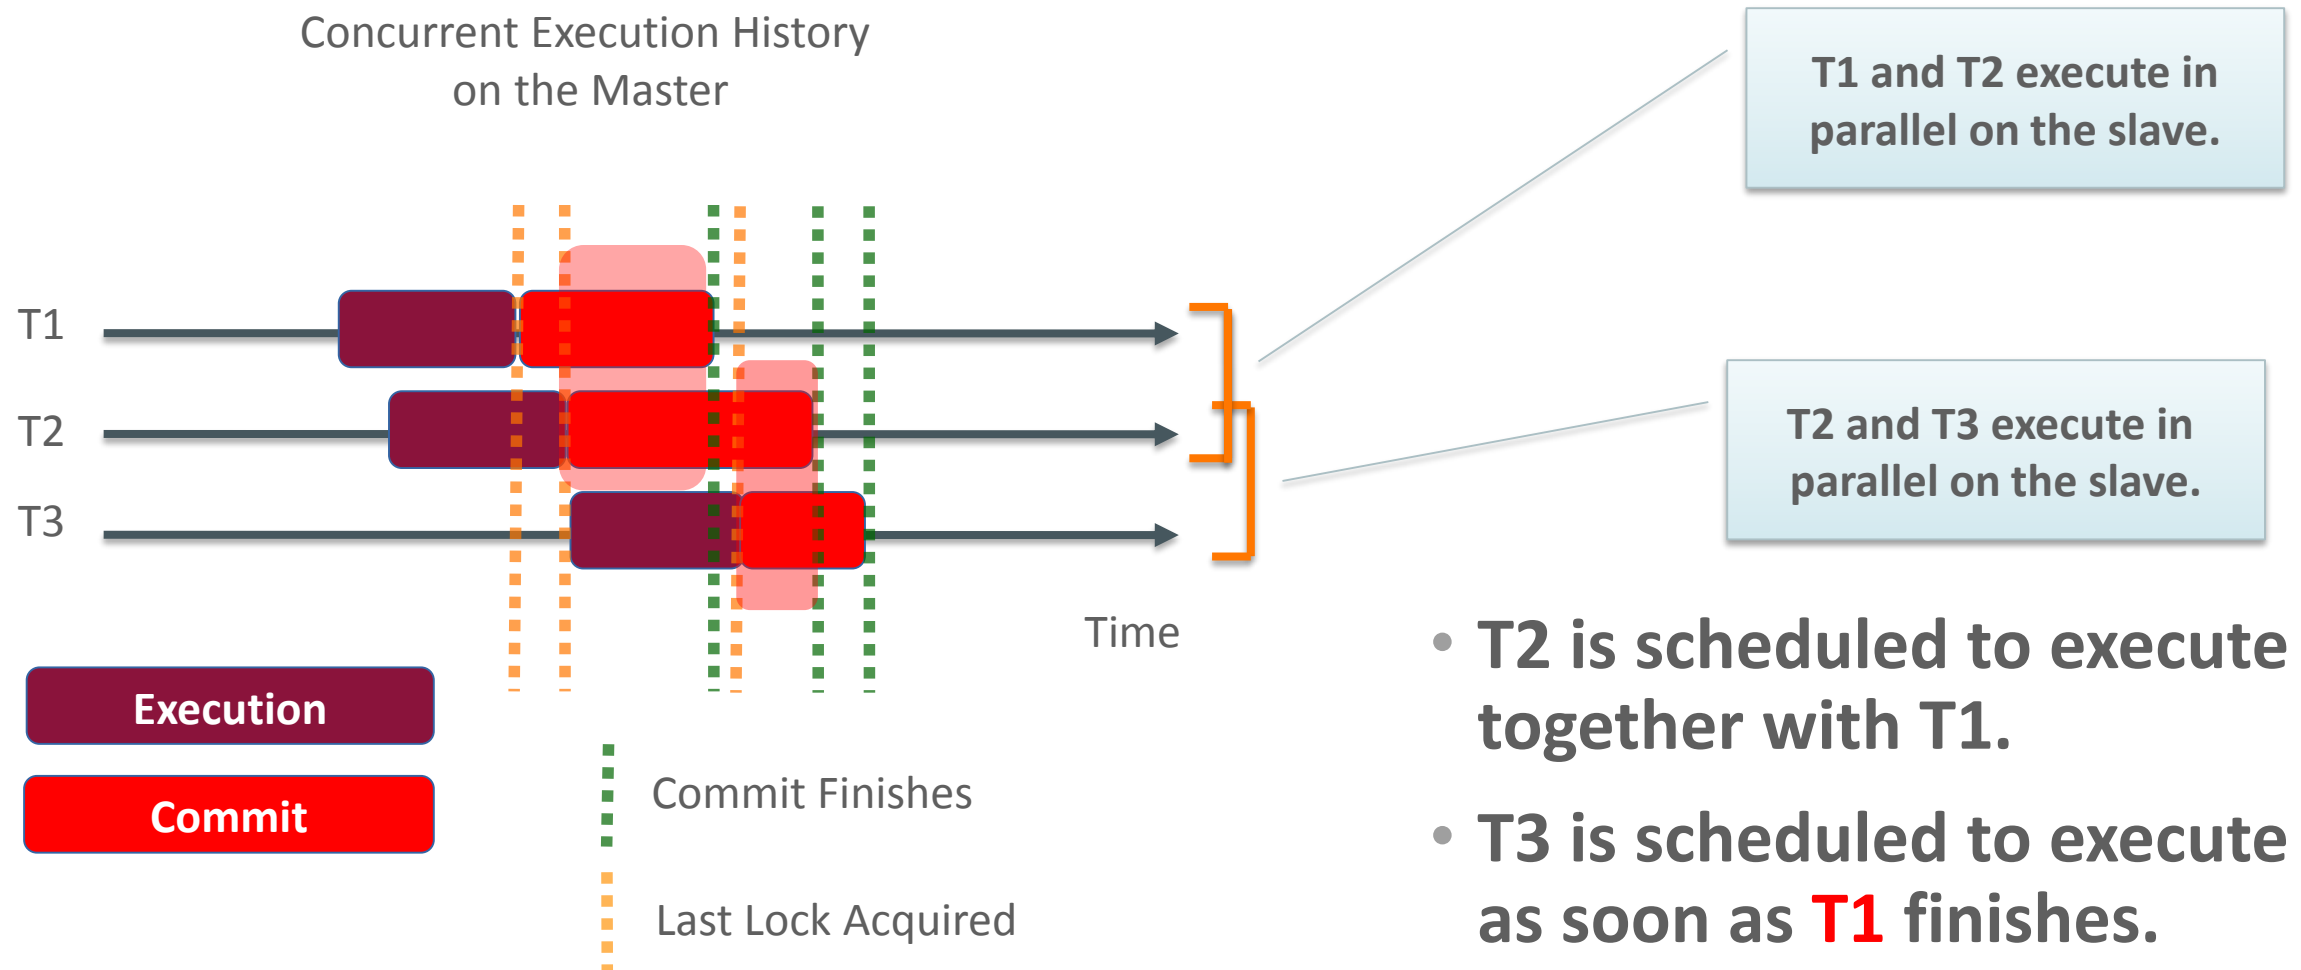


# Improved Applier Throughput – Locking-based Parallelism

- Leverage parallelization information from master execution:
  - Concurrent transactions, which **have not blocked each other** on the master, are marked as non-contending in the binary log.
- Meanwhile, at the slave:
  - Non-contending transactions are scheduled in **parallel**;
  - Concurrent transactions commit independently, thus **no waiting** involved.



# Improved Applier Throughput – Locking-based Parallelism



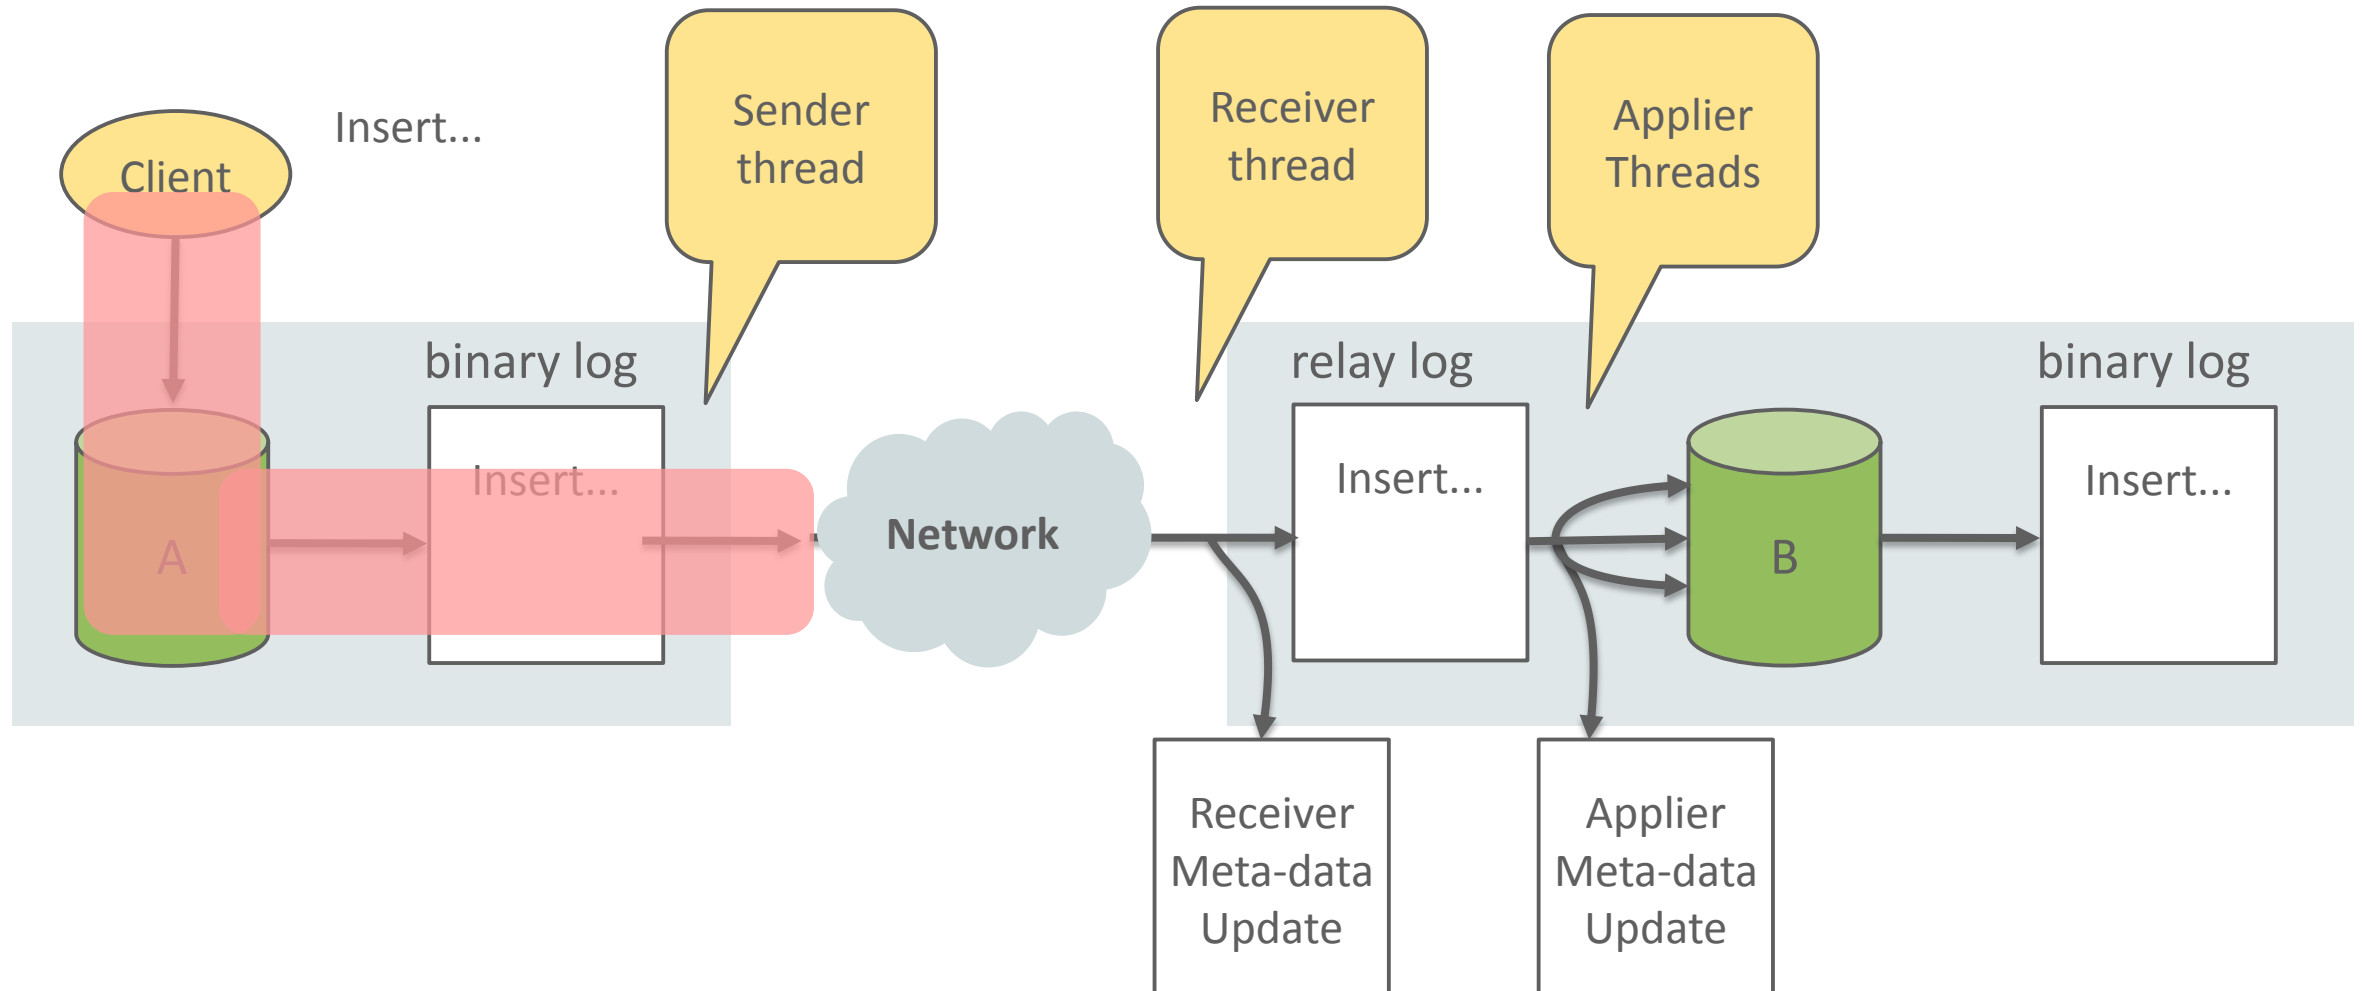
# Improved Applier Throughput – Locking-based Parallelism

- Supports statement-based or row-based formats.
- Scheduling policy controlled through:

```
mysql> SET slave_parallel_type= [logical_clock|database]
```

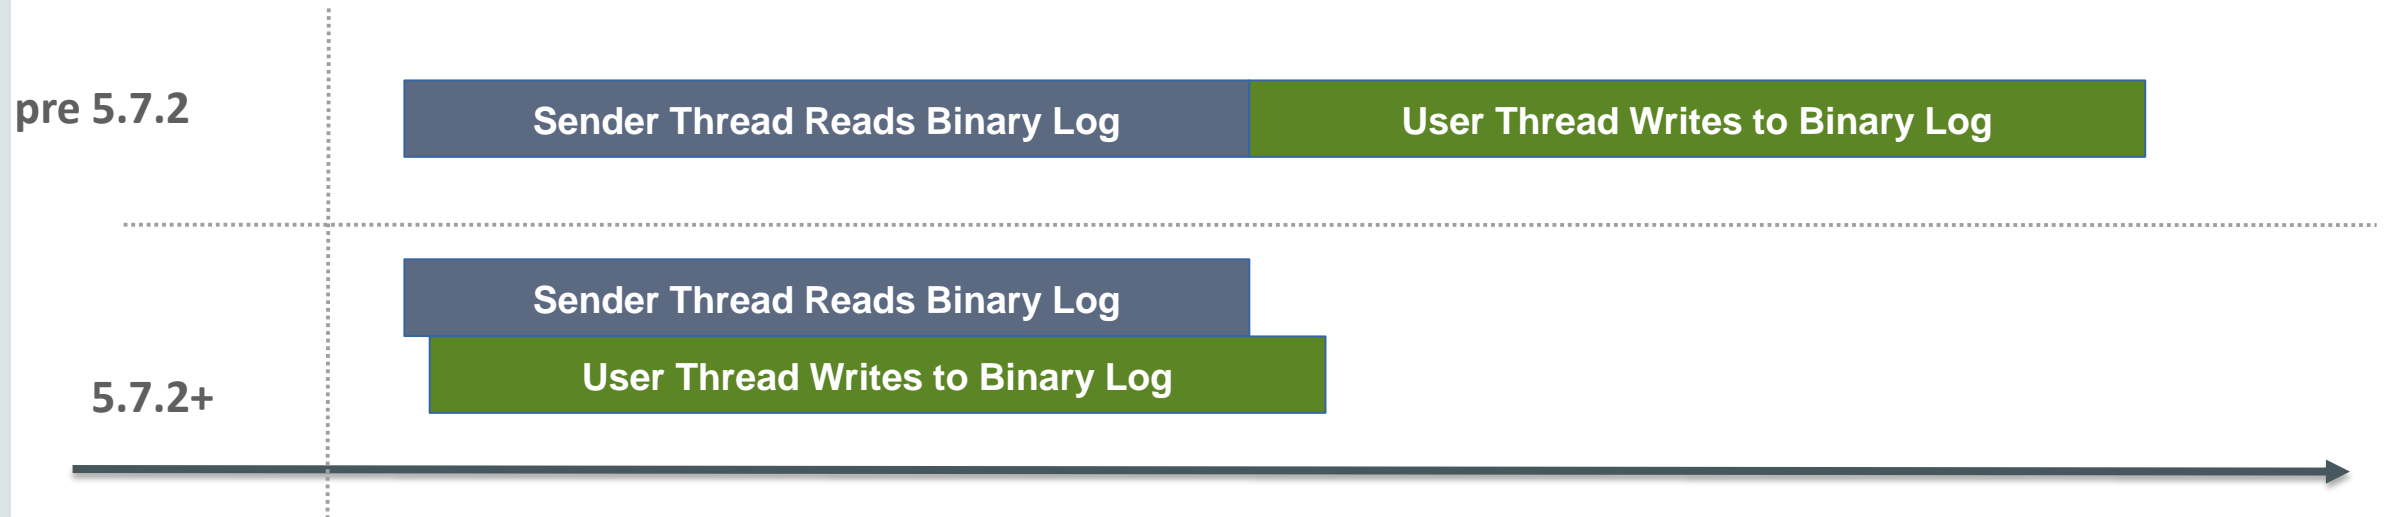
- ***logical\_clock*** - means schedule based on the locking interval timestamps.
- ***database*** - the scheduling policy from 5.6 (concurrency control done per database).
- Work to improve slave scalability continues, does not stop here.

# Improved User Threads-Sender Threads Synchronization

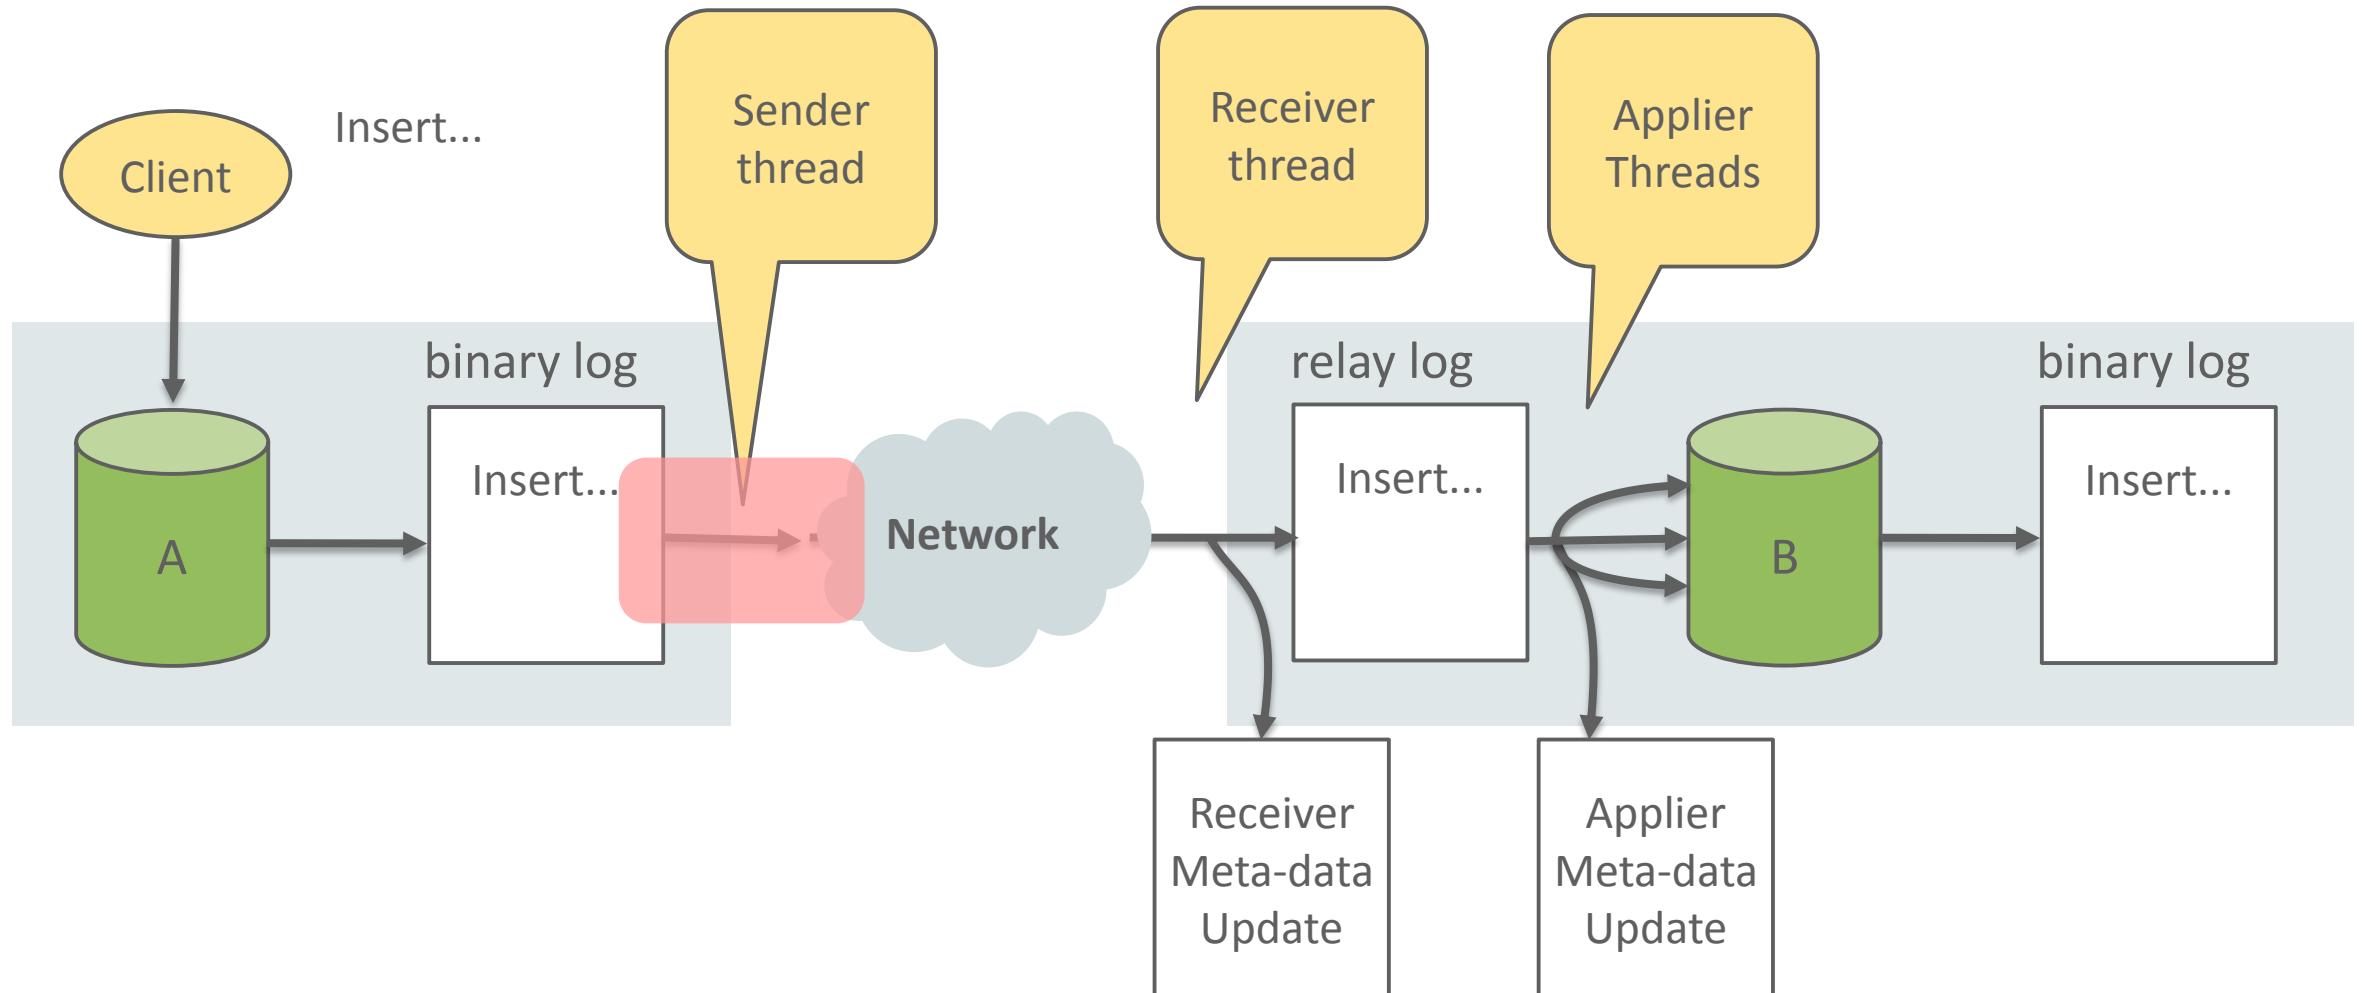


# Improved User Threads-Sender Threads Synchronization

- Concurrent reads by the sender thread with ongoing writes from user threads.
  - Sender thread does not block user sessions more than necessary.
  - Higher throughput for both sender threads and user sessions.




# Faster Sender Thread



# Faster Sender Thread

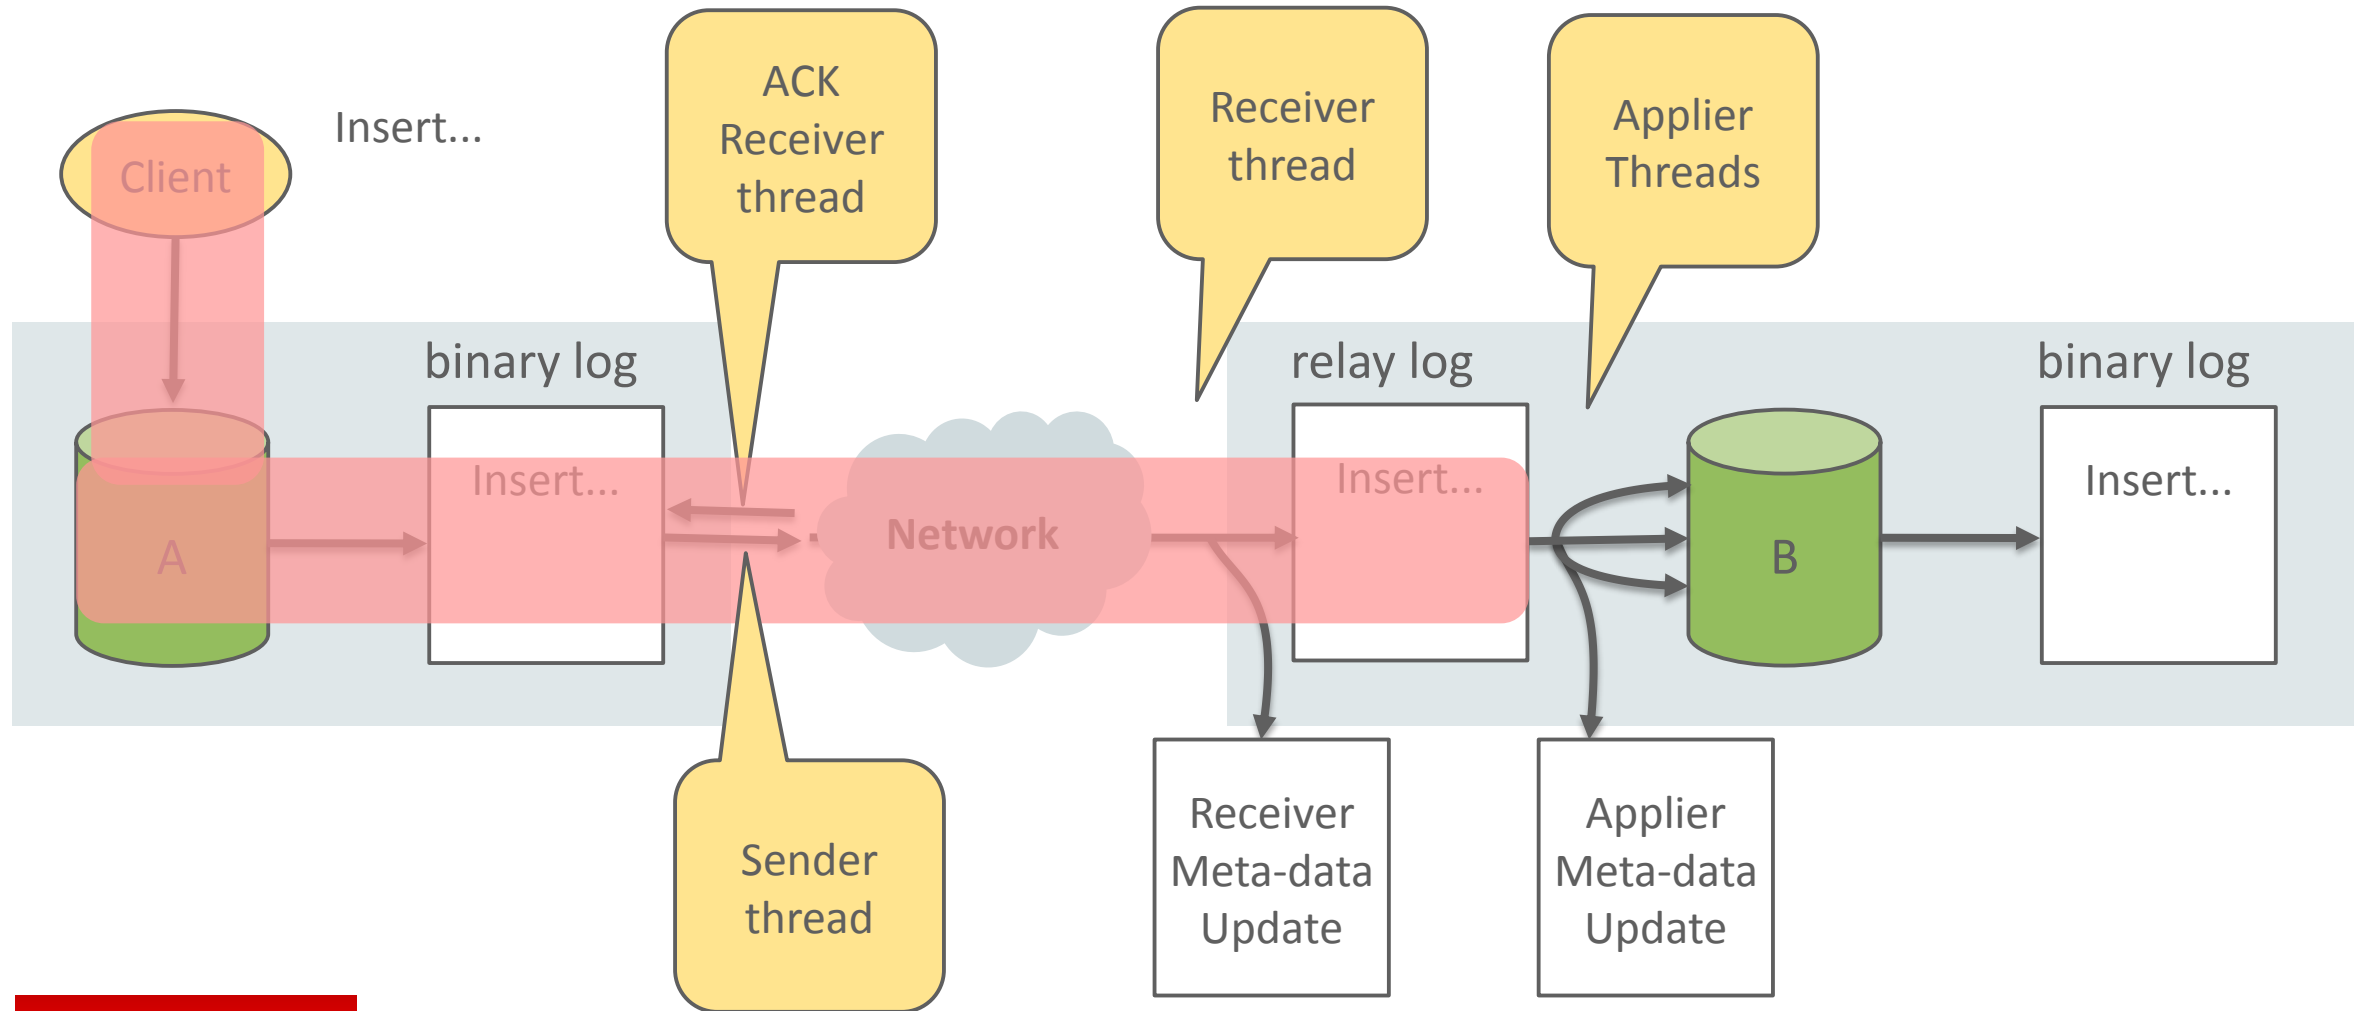
- Send buffer is not allocated and freed every time an event is sent.
- When sender threads require larger send buffer - buffer grows as needed.
- When buffer is larger than needed - buffer shrinks dynamically.
- Together with the sender thread enhancements released on MySQL 5.7.2:
  - increases master scalability;
  - reduces resource consumption (CPU);
  - Master, with dump threads attached, copes better with peak loads.

# Faster Sender Thread



- Ad hoc microbenchmark using mysqlslap
- 1M queries, concurrency=200, commit=1
- N slaves attached (fake slaves using remote mysqlbinlogs)
- 48 cores HT / 512 GB RAM / HDD

# Faster Semi-sync Replication – ACK Receiver Thread





# Faster Semi-sync Replication – ACK Receiver Thread

- Consecutive transactions do not block each other while waiting for ACKs
  - Transaction t1 and t2 are sent immediately to the slave by the sender thread
  - ACKs are received only by a special thread
  - Transaction t2 does not include t1 round trip in its semi-sync overall latency
- Thread starts when semi-sync is activated


```
mysql> SET GLOBAL rpl_semi_sync_master_enabled= ON
```

- Thread stops when semi-sync is deactivated

```
mysql> SET GLOBAL rpl_semi_sync_master_enabled= OFF
```

# Faster Semi-sync: Durability over the Network

Using Semi-synchronous Replication for Durability over the Network



Orange bar **percentage** shows the throughput gain (orange) when compared to the corresponding blue bar. Blue and stacked bars show throughput ratio computed against the best blue bar.

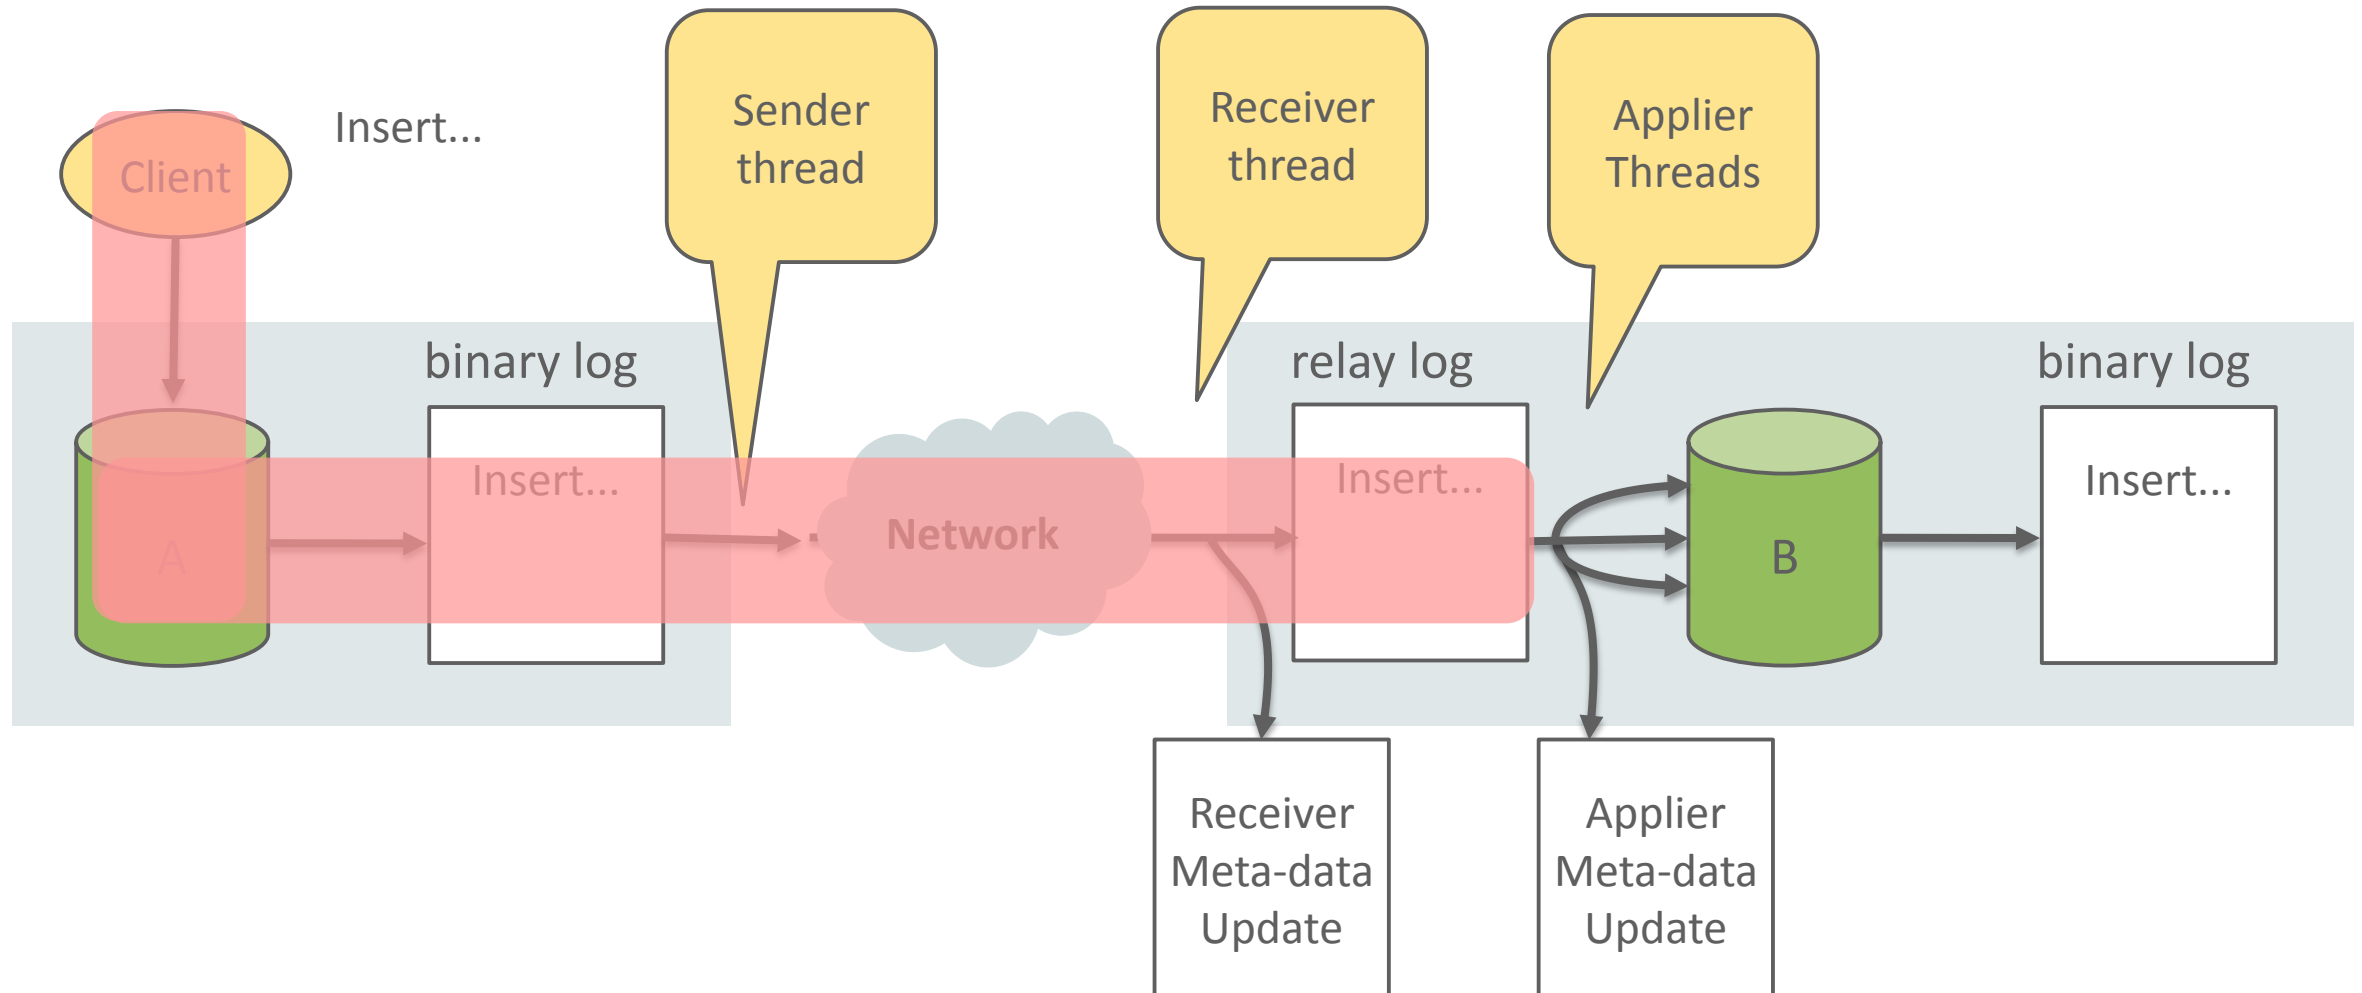
# 3 The New Replication Features in MySQL 5.7

3.1 Usability / Online

3.2 Performance

3.3 Dependability

# Loss-less Semi-sync Replication

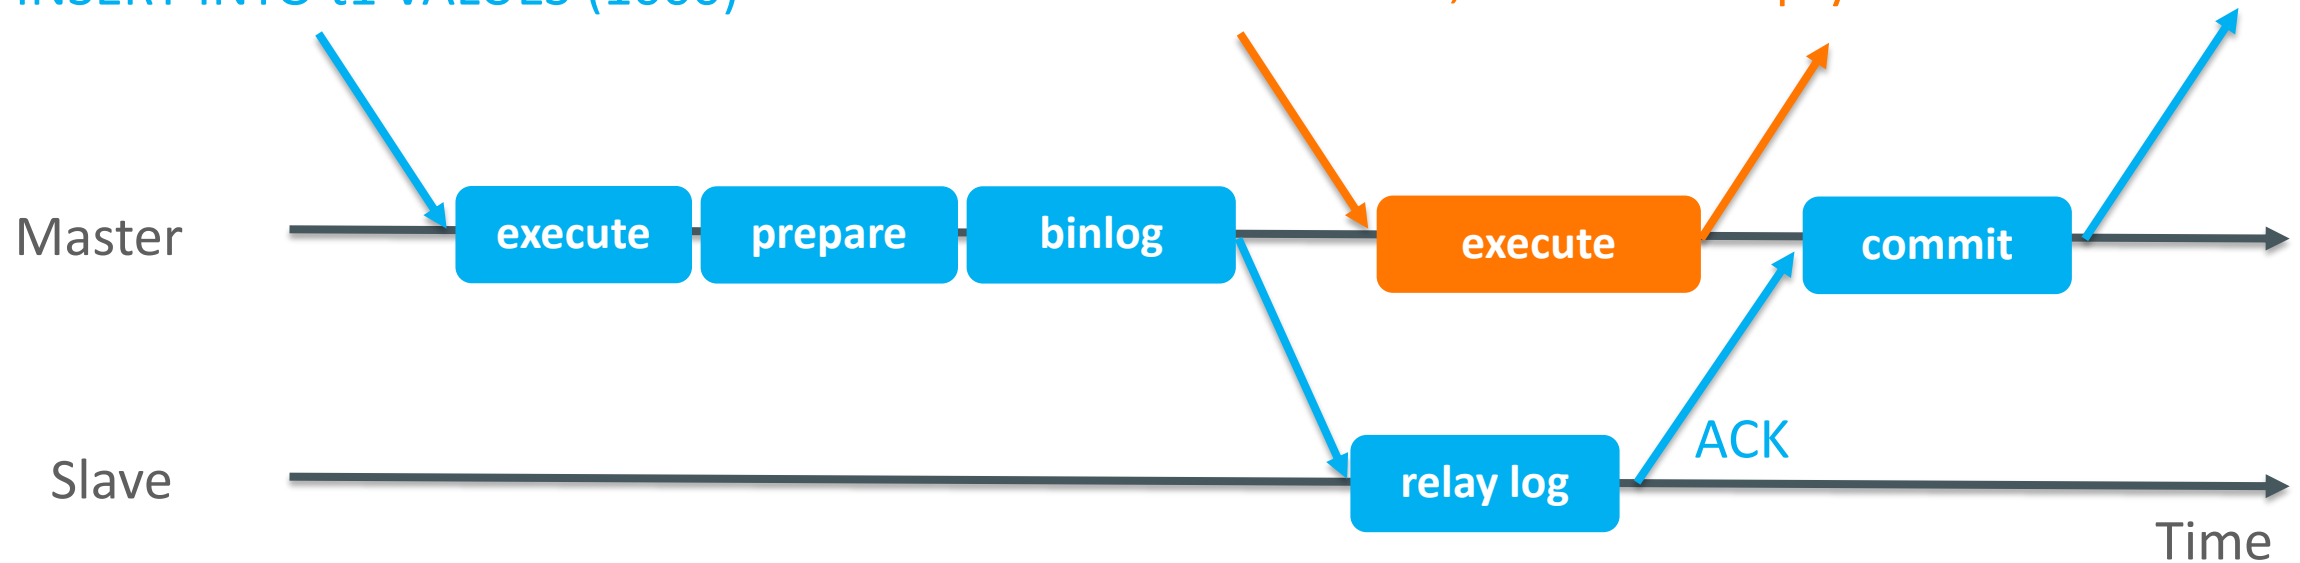


# Loss-less Semi-sync Replication

T1: INSERT INTO t1 VALUES (1000)

T2: SELECT \* FROM t1;

empty set

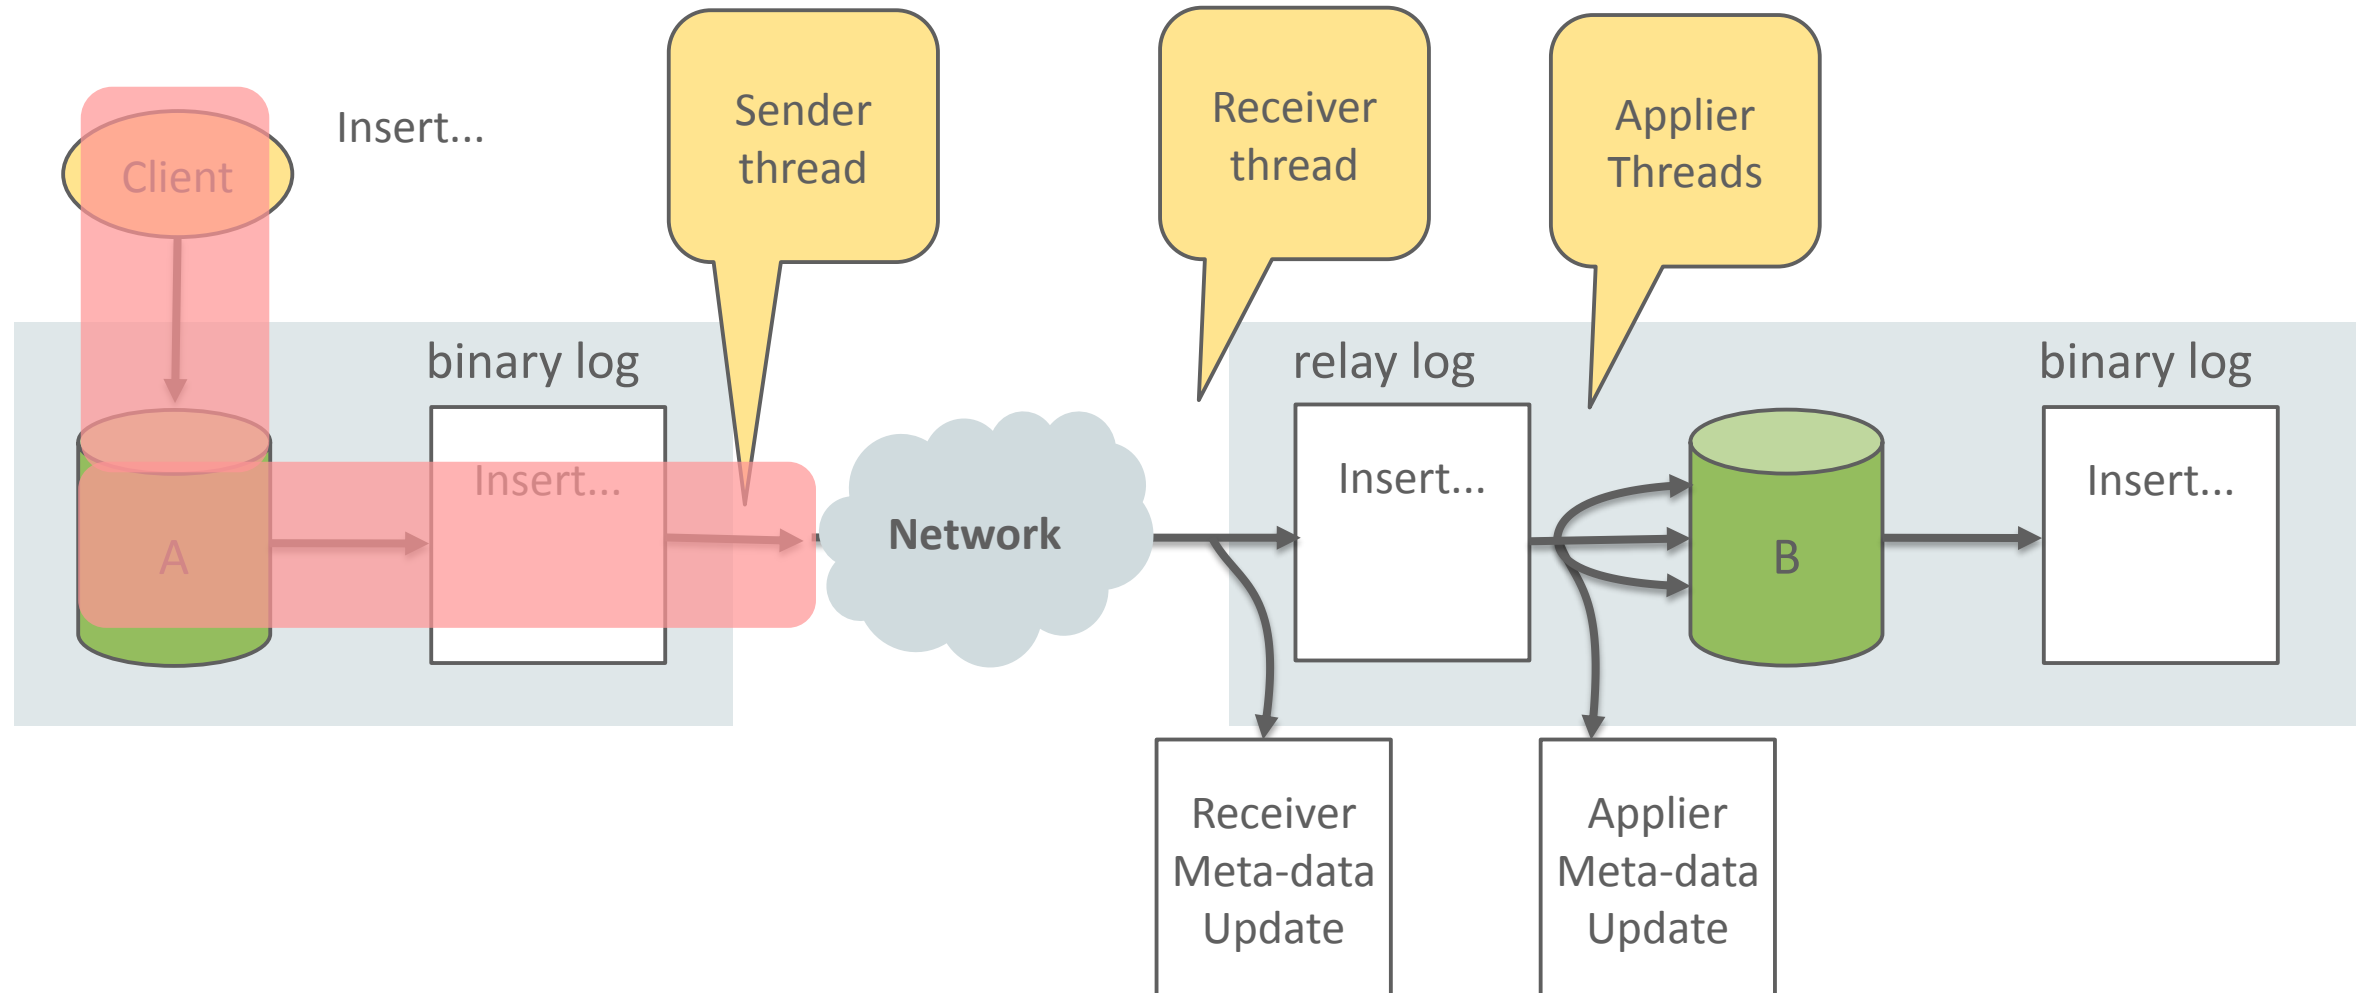


# Loss-less Semi-sync Replication

- Master waits for slave's ACK **before** committing (as opposed to: master waits for slave's ACK **after** committing).
  - Therefore, concurrent transactions do not see changes while this transaction waits for ack.
- Should a master fail, then any transaction that it may have externalized is also persisted on a slave.
- User can choose between the original semi-sync behavior and the new

```
mysql> SET rpl_semi_sync_master_wait_point= [AFTER_SYNC|AFTER_COMMIT]
```

# Semi-sync Replication – Wait for Multiple ACKs



# Semi-sync Replication – Wait for Multiple ACKs

- Master does **not commit** transaction until it receives **N** ACKs from **N** slaves.
- Dynamically settable:

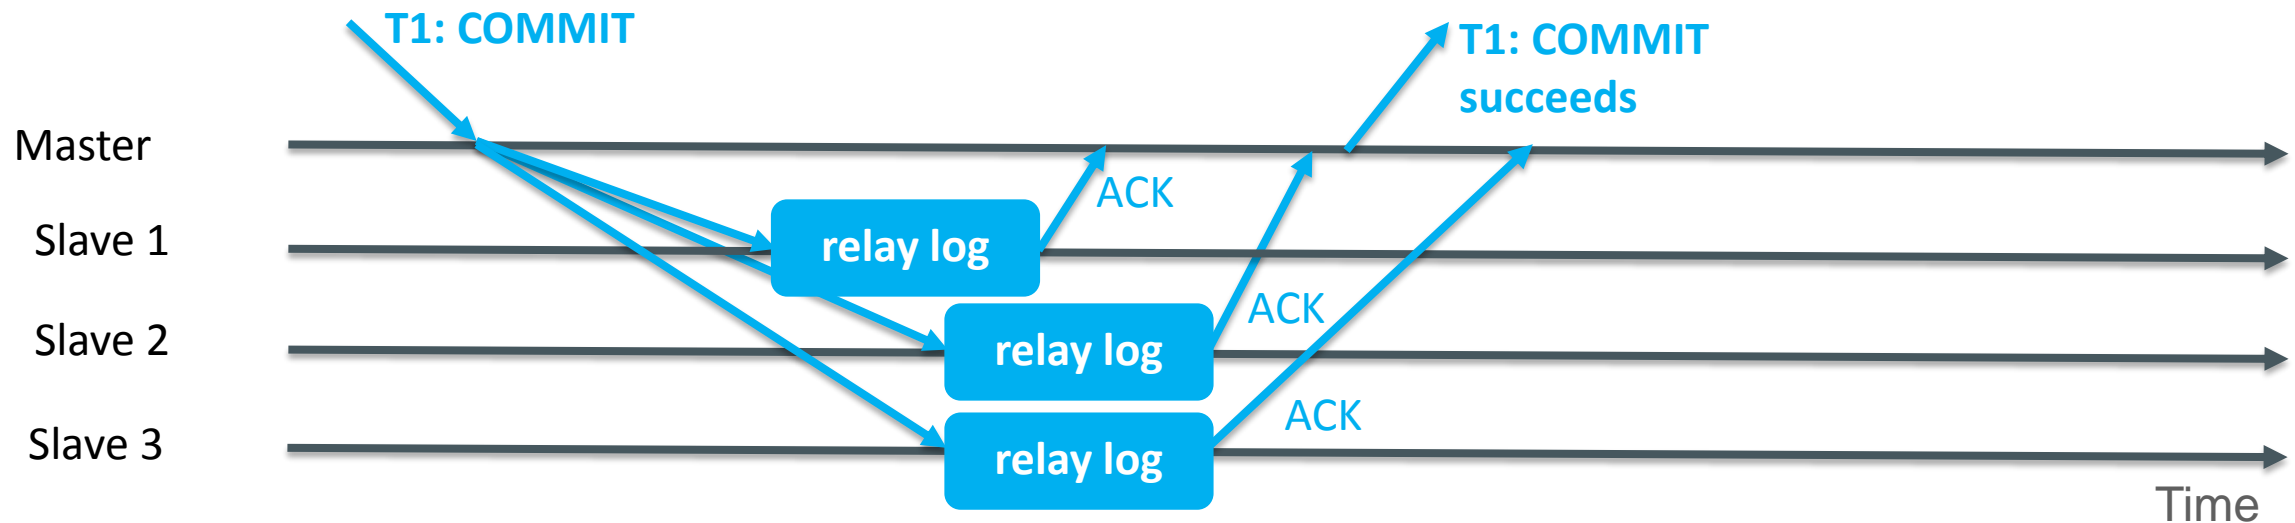
```
mysql> SET GLOBAL rpl_semi_sync_master_wait_for_slave_count= N
```



# Semi-sync Replication – Wait for Multiple ACKs

- Master does **not commit** transaction until it receives **N** ACKs from **N** slaves.
- Dynamically settable:

```
mysql> SET GLOBAL rpl_semi_sync_master_wait_for_slave_count= N
```



```
mysql> SET GLOBAL rpl_semi_sync_master_wait_for_slave_count= 2
```

# 3 The New Replication Features in MySQL 5.7

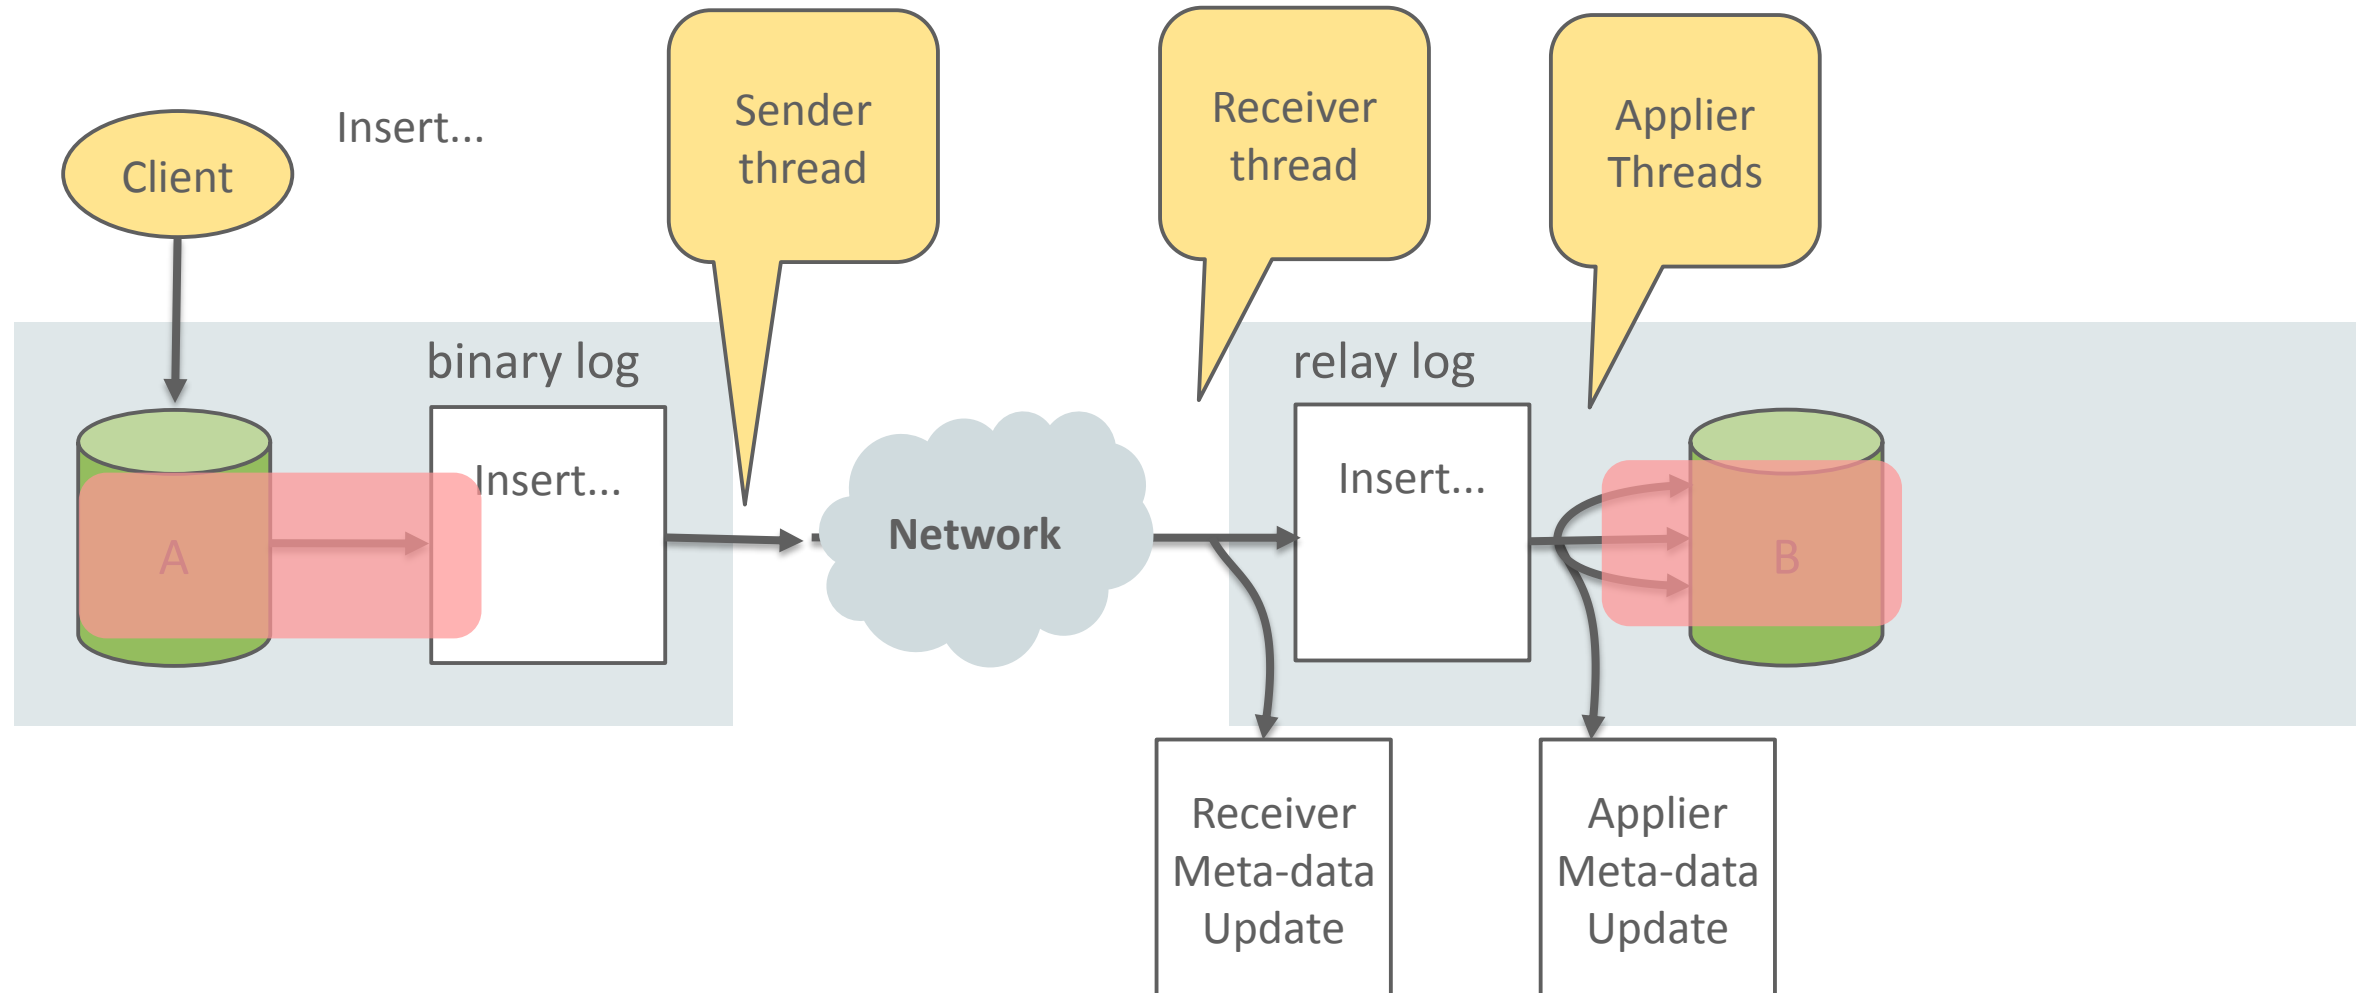
3.1 Usability / Online

3.2 Performance

3.3 Dependability

3.4 Flexibility

# Storing Global Transaction Identifiers in a Table



# Storing Global Transaction Identifiers in a Table

## Use Cases

- Slaves can use GTIDs with binary logs disabled.
  - Slaves that are never candidates to become a master can still use GTIDs for auto positioning.

# Storing Global Transaction Identifiers in a Table Mechanics

- GTIDs are saved into a mysql system table (range-compressed).

```
CREATE TABLE gtid_executed(  
    source_uuid CHAR(36) NOT NULL,  
    interval_start BIGINT NOT NULL,  
    interval_end BIGINT NOT NULL,  
    PRIMARY KEY(source_uuid, interval_start)  
);
```


- GTIDs are **inserted** into the table as transactions commit.
- A compression thread runs periodically and compresses the table into ranges.
  - New global variable controls the period: gtid\_executed\_compression\_period.

```
mysql> SET GLOBAL gtid_executed_compression_period= N; (N - number of transactions)
```

# Storing Global Transaction Identifiers in a Table Mechanics

- Binary Log Enabled
  - Store GTID in binary log (transactionally).
  - Copy GTID set to table on binary log rotation.
- Binary Log Disabled
  - Store GTID in table (transactionally).
  - New transactions do not get an automatic GTID assigned.
- Always store GTIDs transactionally.

# Multi-Source Replication



Slave **can** have more than **one master**.

The need for gathering data in a central server:

- Integrated backup;
- Complex queries for analytics purposes;
- Data HUB for inter-cluster replication.

# Multi-Source Replication

- A server (slave) can replicate from multiple sources (masters).
- Multiple channels (channel: receiver thread, relay logs, applier threads) that can be stopped started individually.
- Integrated with Multi-threaded Slave (each channel has its own multi-threaded applier set of threads).
- Integrated with the new P\_S tables.
  - replication\_applier\_status\_by\_coordinator shows multiple entries, one per channel/source coordinator.
  - replication\_applier\_status\_by\_worker shows multiple entries, multiple entries per channel
  - replication\_connection\_status shows multiple entries, one per connection to different sources.



# Multi-Source Replication

- Integrated with GTIDs.
- Integrated with crash-safe tables.
  - Progress state is stored in these tables for recoverability purposes.
- Virtually no limit on the number of sources (capped to 256, but can be changed if server binary is rebuilt).
- Able to manage each source separately.

# 3 The New Replication Features in MySQL 5.7

3.1 Usability / Online

3.2 Performance

3.3 Dependability

3.4 Flexibility


3.5 Misc

# “Smaller”, yet interesting, enhancements!

- **Multi-threaded applier** is able to **retry failed transactions**.
- Option to make the **Multi-threaded applier** **preserve commit order**.
- **SSL** options for **mysqlbinlog** tool.
- **Rewrite DB** rules for the **mysqlbinlog** tool.
- **Function to wait for transactions to be applied**, regardless of the replication stream they come from.
- Options to **track global transaction identifiers** in the **return packet** of the mysql protocol. Useful for tracking GTID session state and enabler for session consistency.
- Support for **XA transactions** when the **binary log is enabled**.
- Change in the defaults. E.g., **binlog\_format=ROW** and **sync\_binlog=1**.
- Options to **fine tune** the binary log **group commit** procedure.

# 4 Development Sneak Peek

# MySQL Group Replication Plugin



- Multi-master update everywhere
- **Automatic group membership management and failure detection.**
- **No need for server fail-over.**
- **Automatic reconfiguration.**
- No single point of failure.
- Shared-nothing state machine replication.
- **InnoDB compliant.** Off-the-shelf hardware friendly.

# 5 Roadmap

# What is Next?

- **MySQL Group Replication**
  - Rapid releases.
  - Improve performance, stability and usability.
- **MySQL Replication Usability**
  - Instrument even more replication and extend replication P\_S tables
  - Simpler administrative commands, more online operations
- **MySQL Replication Performance**
  - Continue to improve replication stream multi-threaded (slave) applier
  - Continue to improve replication stream pipeline performance and efficiency

# 6 Summary



# Summary

- Replication feature set in MySQL 5.7 shows many improvements:
  - Semi-sync is more flexible and faster.
  - Less contention on the binary log.
  - Slave is faster and provides an inter-transaction parallelization scheme.
  - More online reconfiguration: Global Transaction Identifiers, filters, configuration.
  - Improved monitoring.
  - Even more flexible replication with multi-source enabling new deployment scenarios.
  - Enhancements all over: transaction retries, sequential commits, XA support, new replication functions, security.

# Where to go from here?

- Packages
  - <http://dev.mysql.com>
  - <http://labs.mysql.com>
- Reference Documentation
  - <http://dev.mysql.com/doc/refman/5.7/en/replication.html>
- Blogs from the Engineers (news, quirks, technical information and much more)
  - <http://mysqlhighavailability.com>

ORACLE®