



# What's new in Percona Xtradb Cluster 5.6

Jay Janssen  
Lead Consultant  
February 5th, 2014

# Overview

PXC 5.6 is the aggregation of

- Percona Server 5.6

- Codership MySQL 5.6 patches

- Galera 3.x

## Agenda

- Major new features

- Async replication and PXC 5.6

- 5.6 features that play nicely (or not) with Galera

- Future features

- Questions

# What is Percona XtraDB Cluster?

Read and Write anywhere Innodb cluster

Highly available quorum-based failover

Synchronous replication for data consistency

Cluster responsible for node state

PXC product 2 years in GA with 5.5, 13 releases total, over 150k downloads

Lots of development effort in the test suite

Easy to migrate from and integrated with MySQL  
async replication

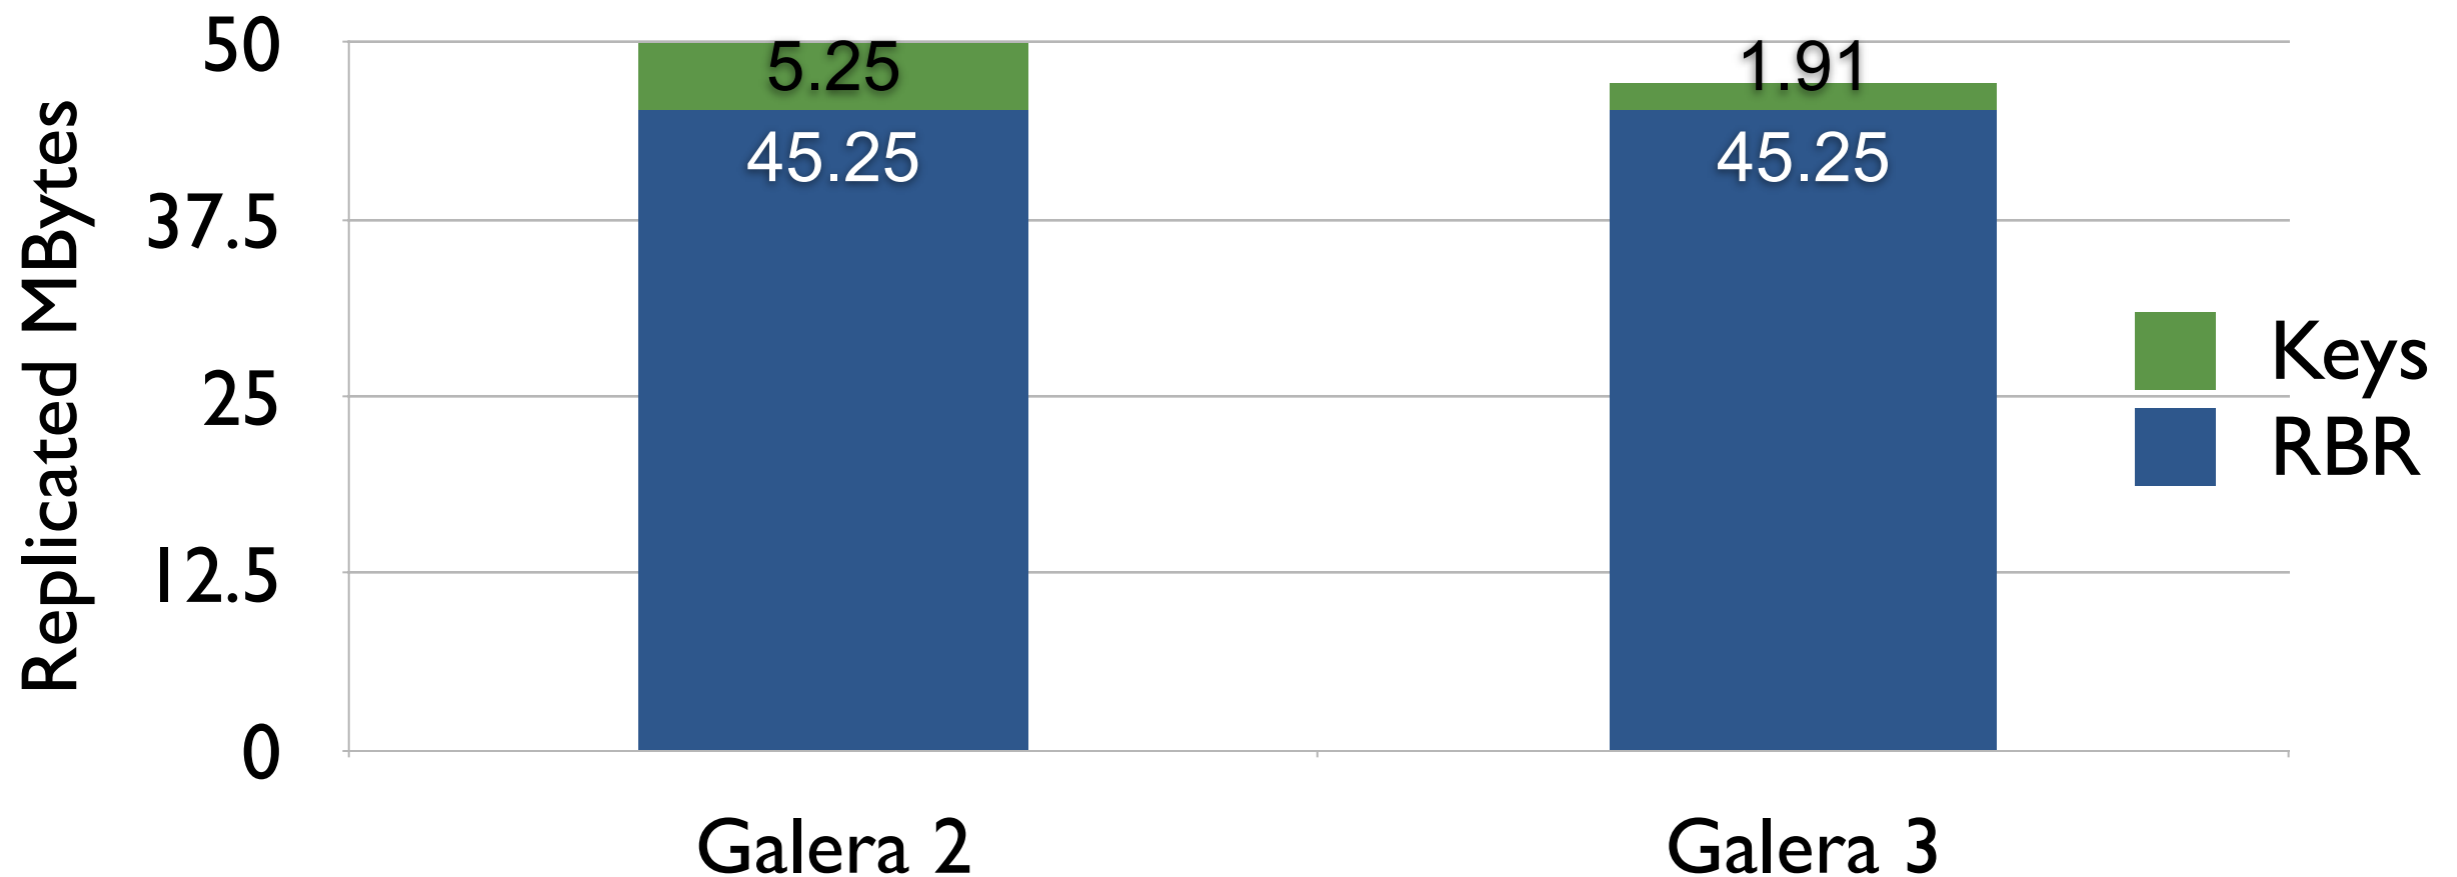


# Galera Replication Enhancements

# Replication overhead for Keys

Keys track all Schemas, Tables, and PK/UKs/FKs of rows being modified in the writeset

250k row inserts ala sysbench prepare

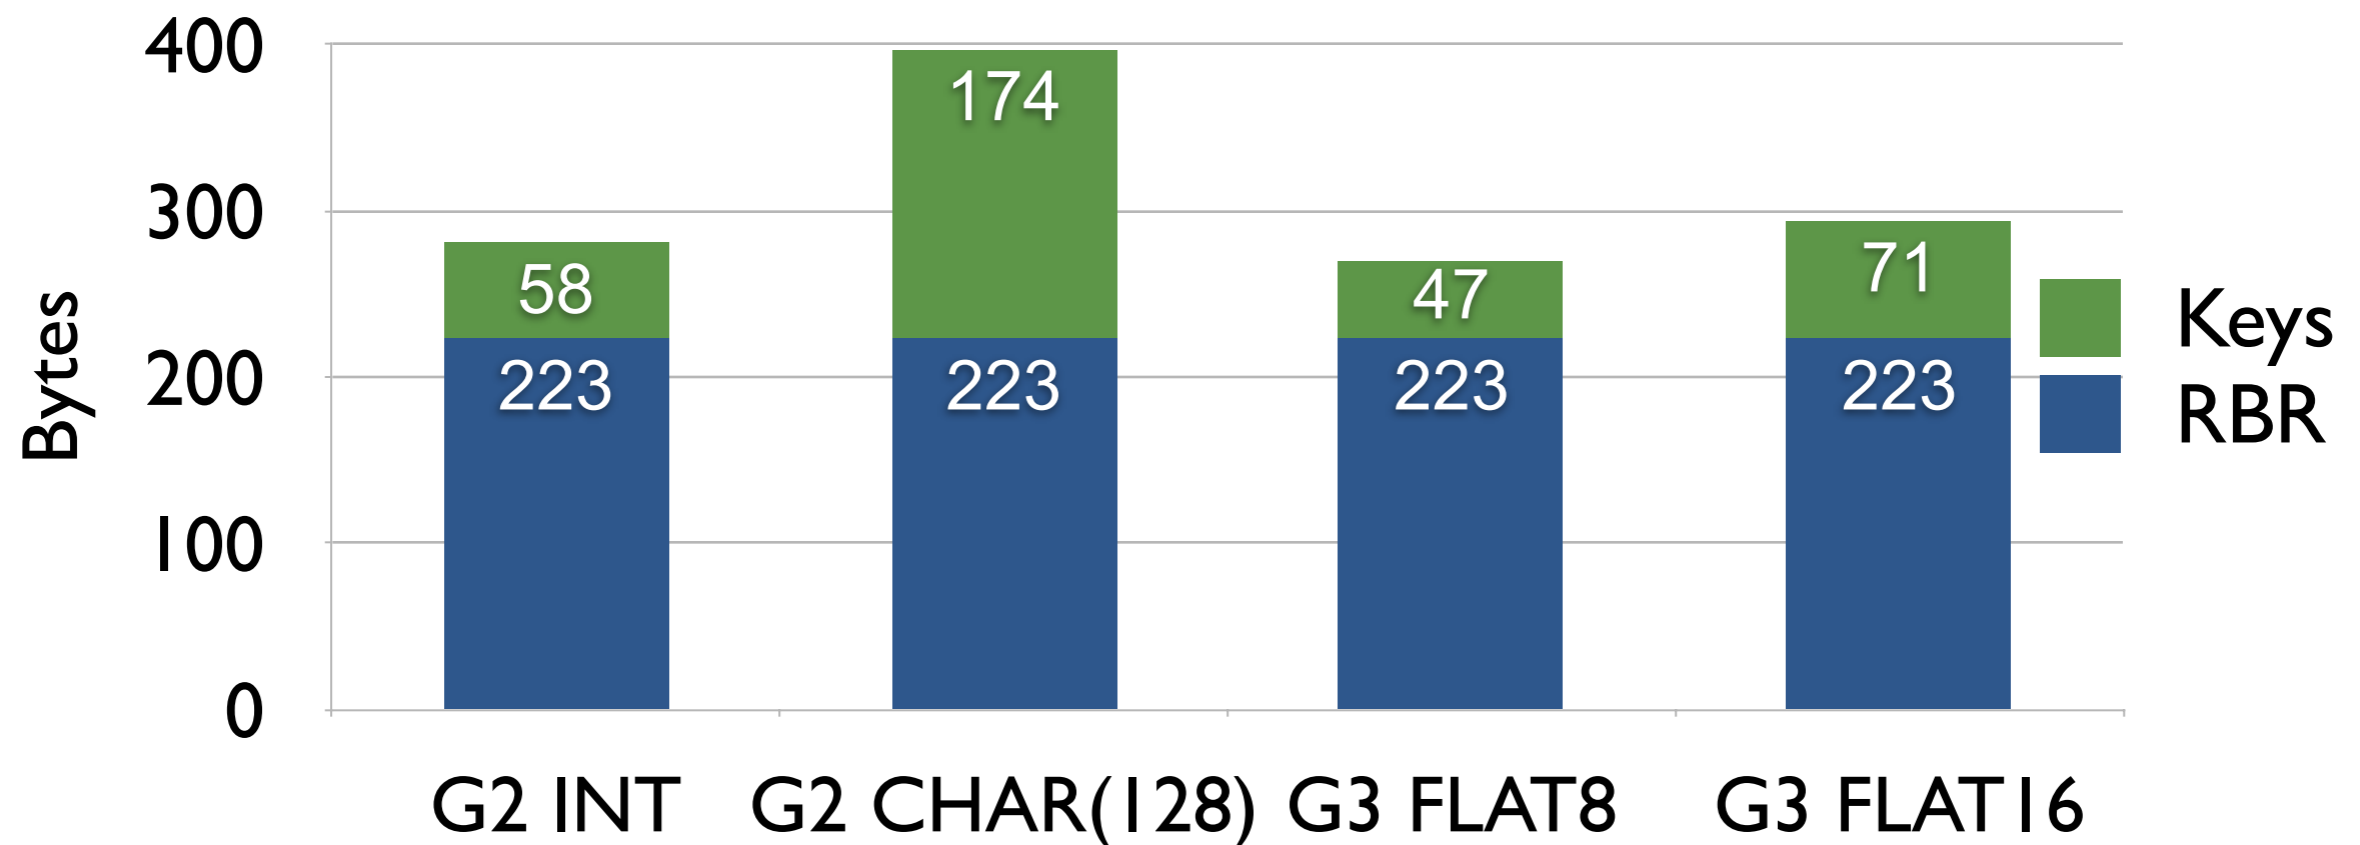


# Keys are now hashed

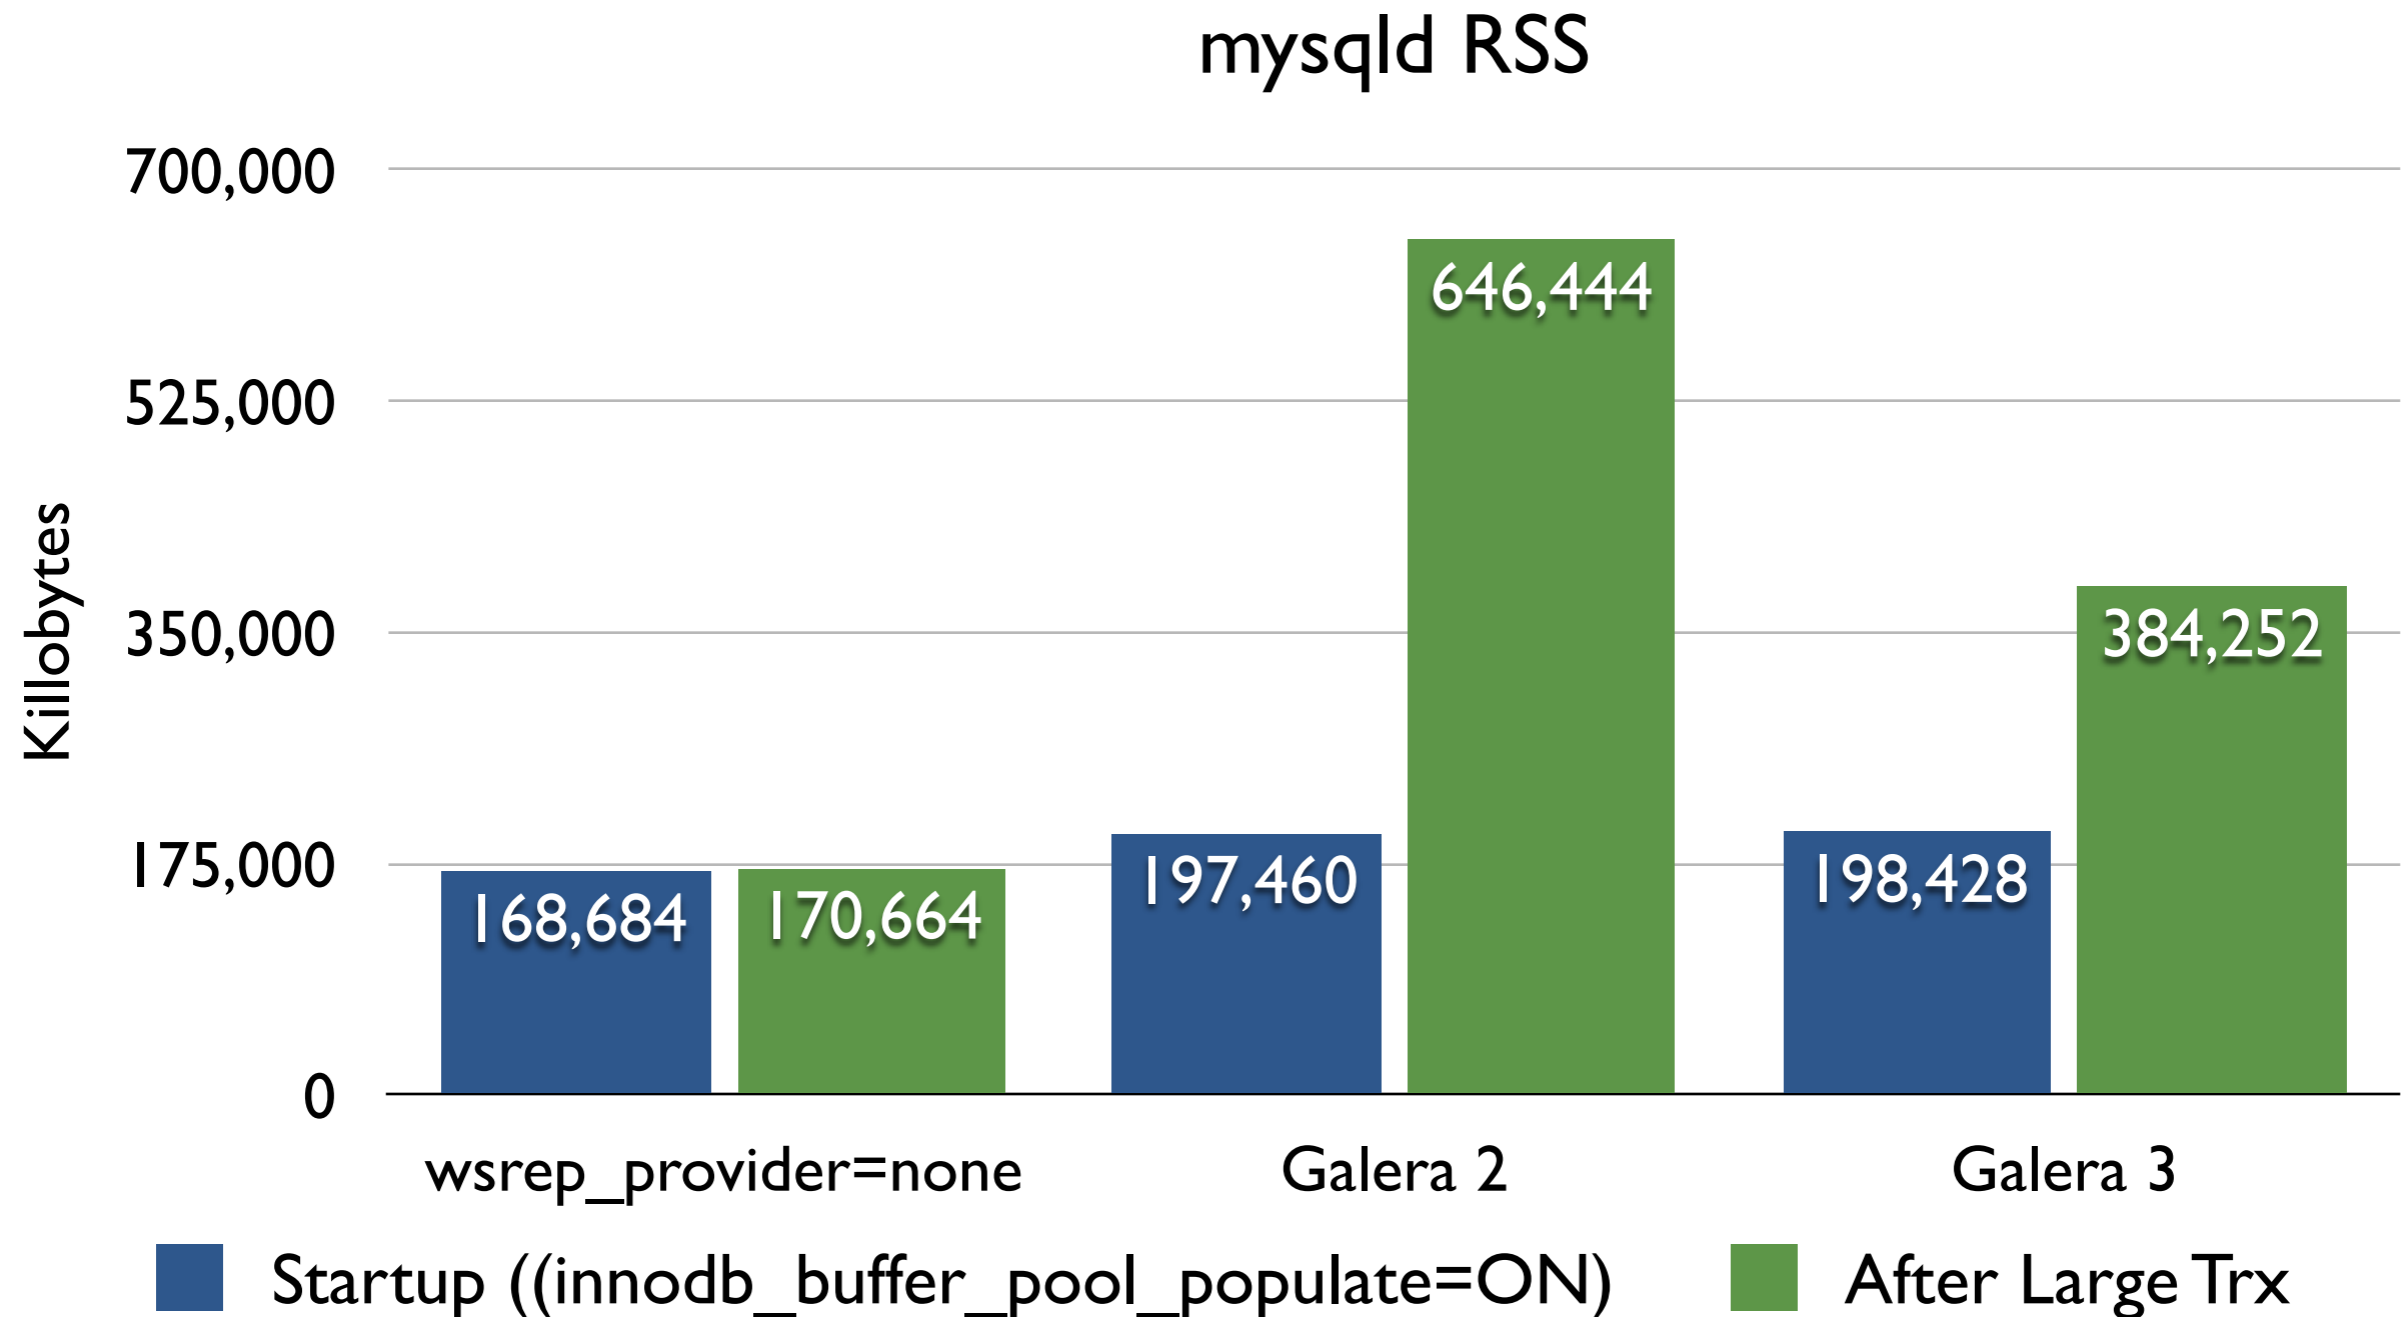
Hashed keys reduces writeset sizes

Especially for large key columns

Single INSERT writeset size



# Improved Memory usage (certification index)



# Overall improvements

---

Writesets have a new format

- faster certification

- less memory

- checksums

Writeset keys stored as hashes

- FLAT8/FLAT16

Socket Checksums

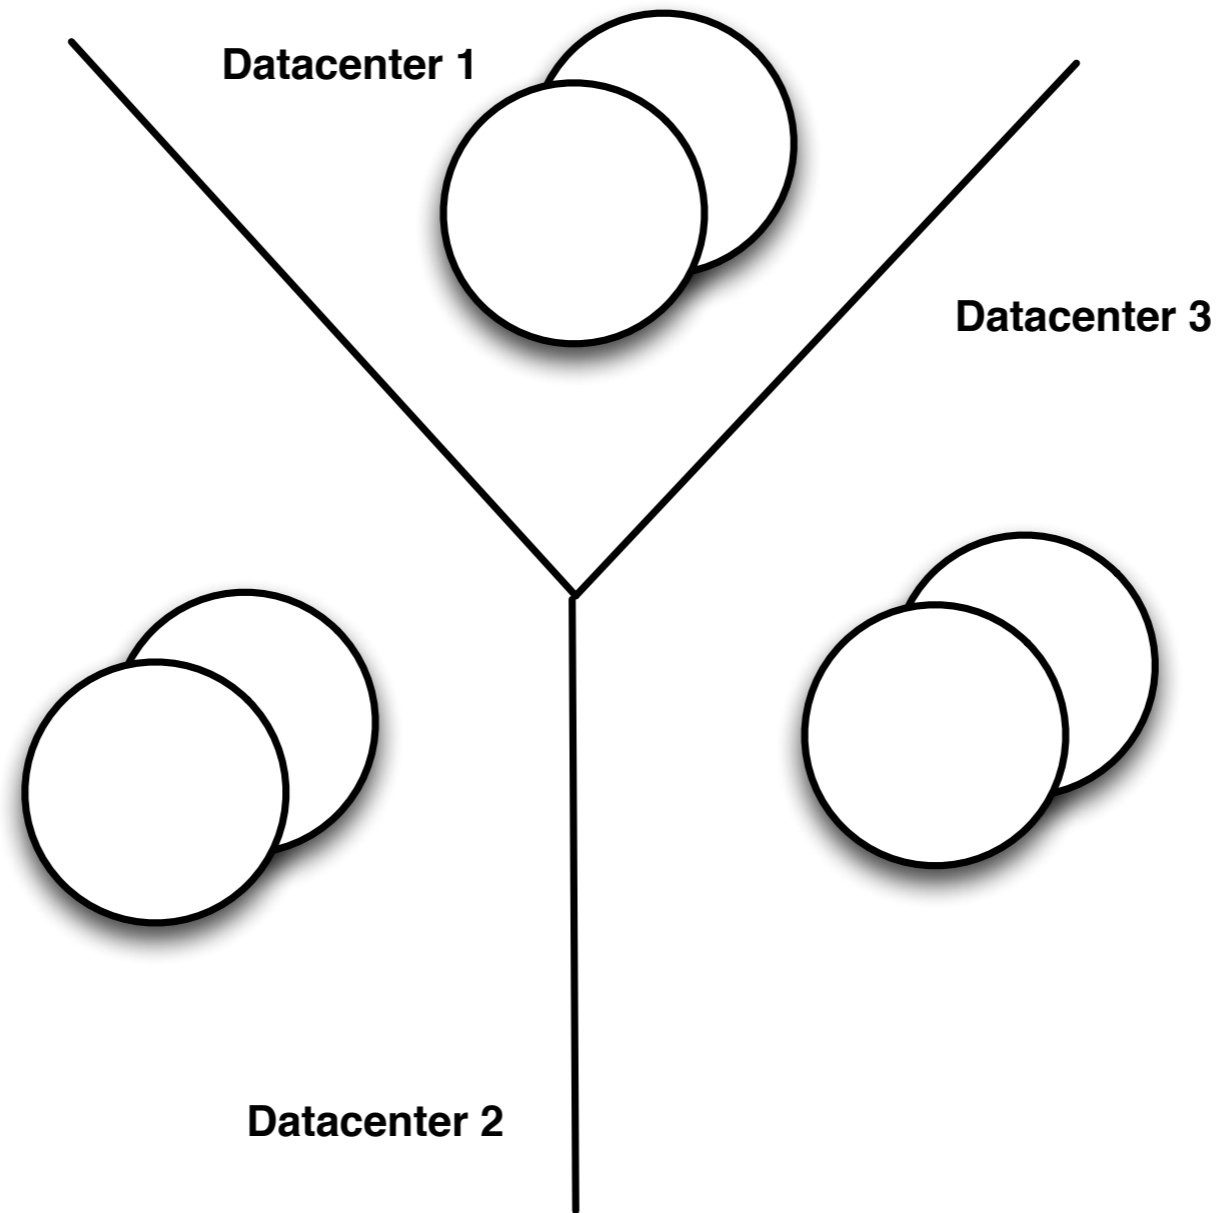
- CRC32-C (hw accel where supported)



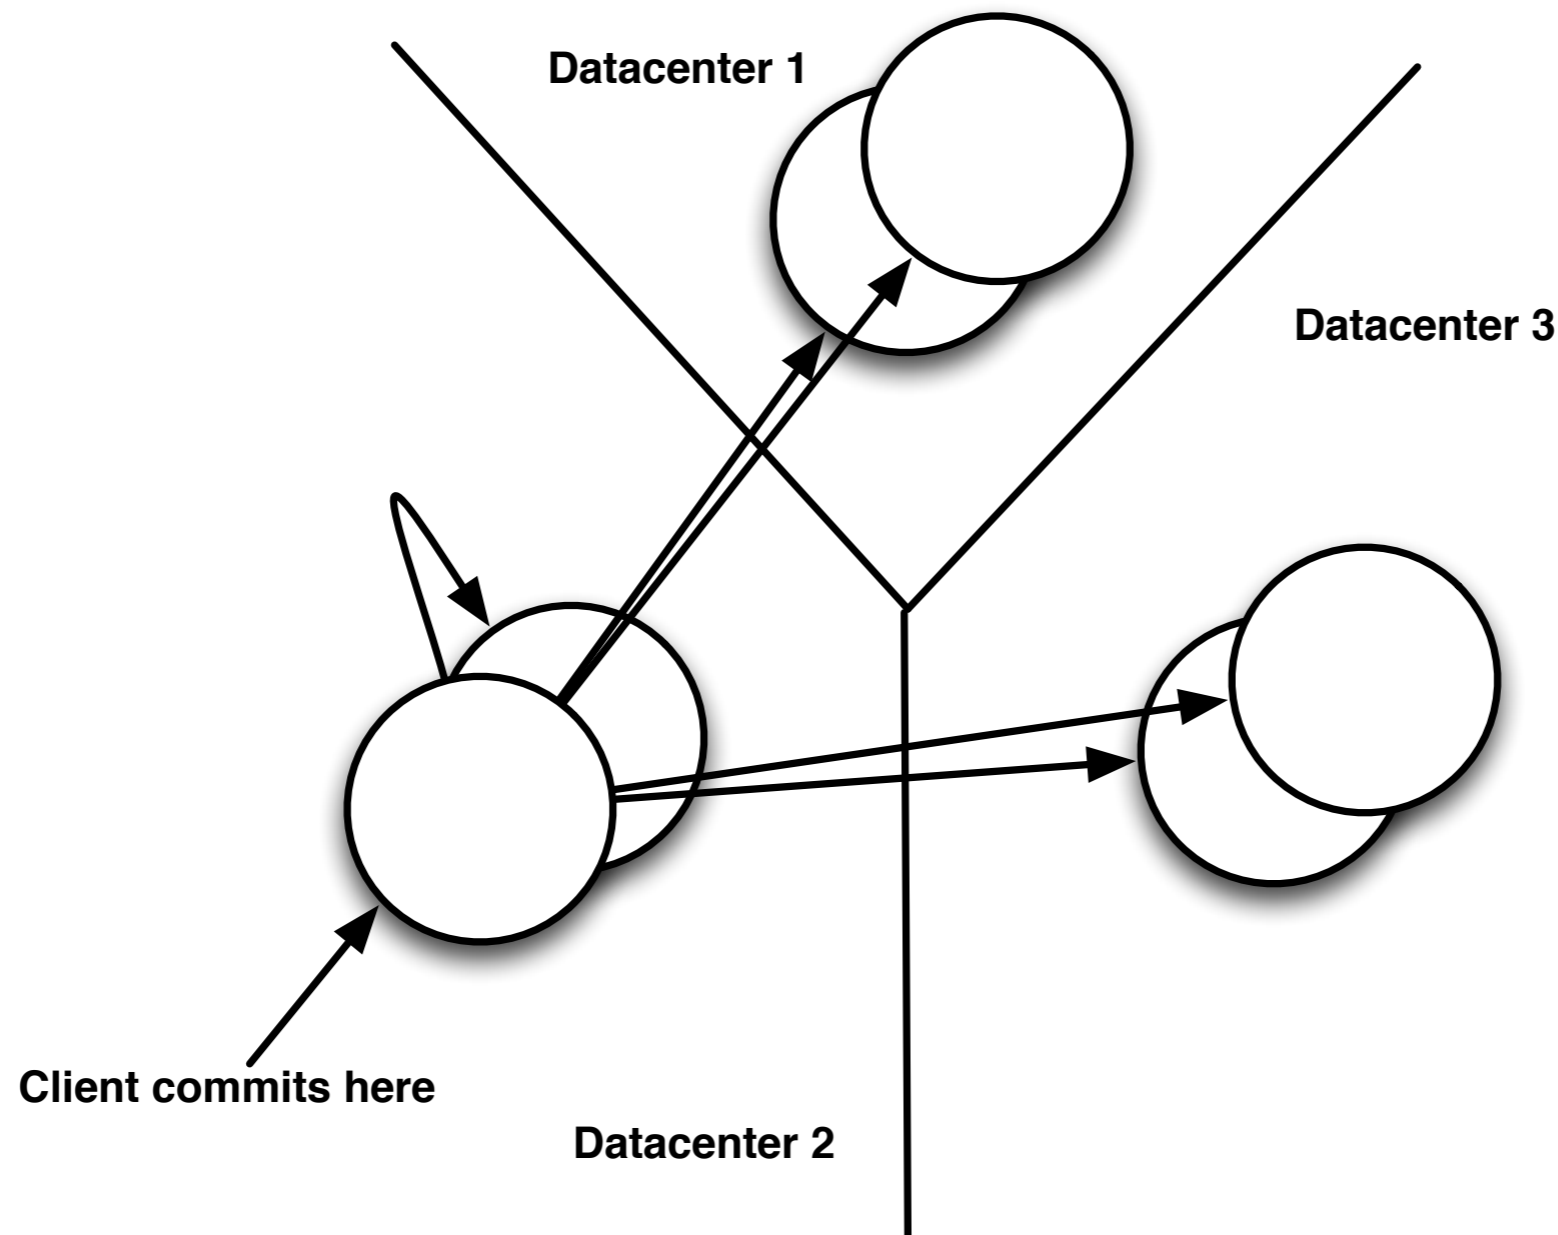


# WAN Segments

# Multiple Datacenter Cluster



# Replication without Segments



# Cluster segments

Set `gmcast.segment` distinct *per-location*

All nodes in colo1 have `gmcast.segment=1`

All nodes in colo2 have `gmcast.segment=2`

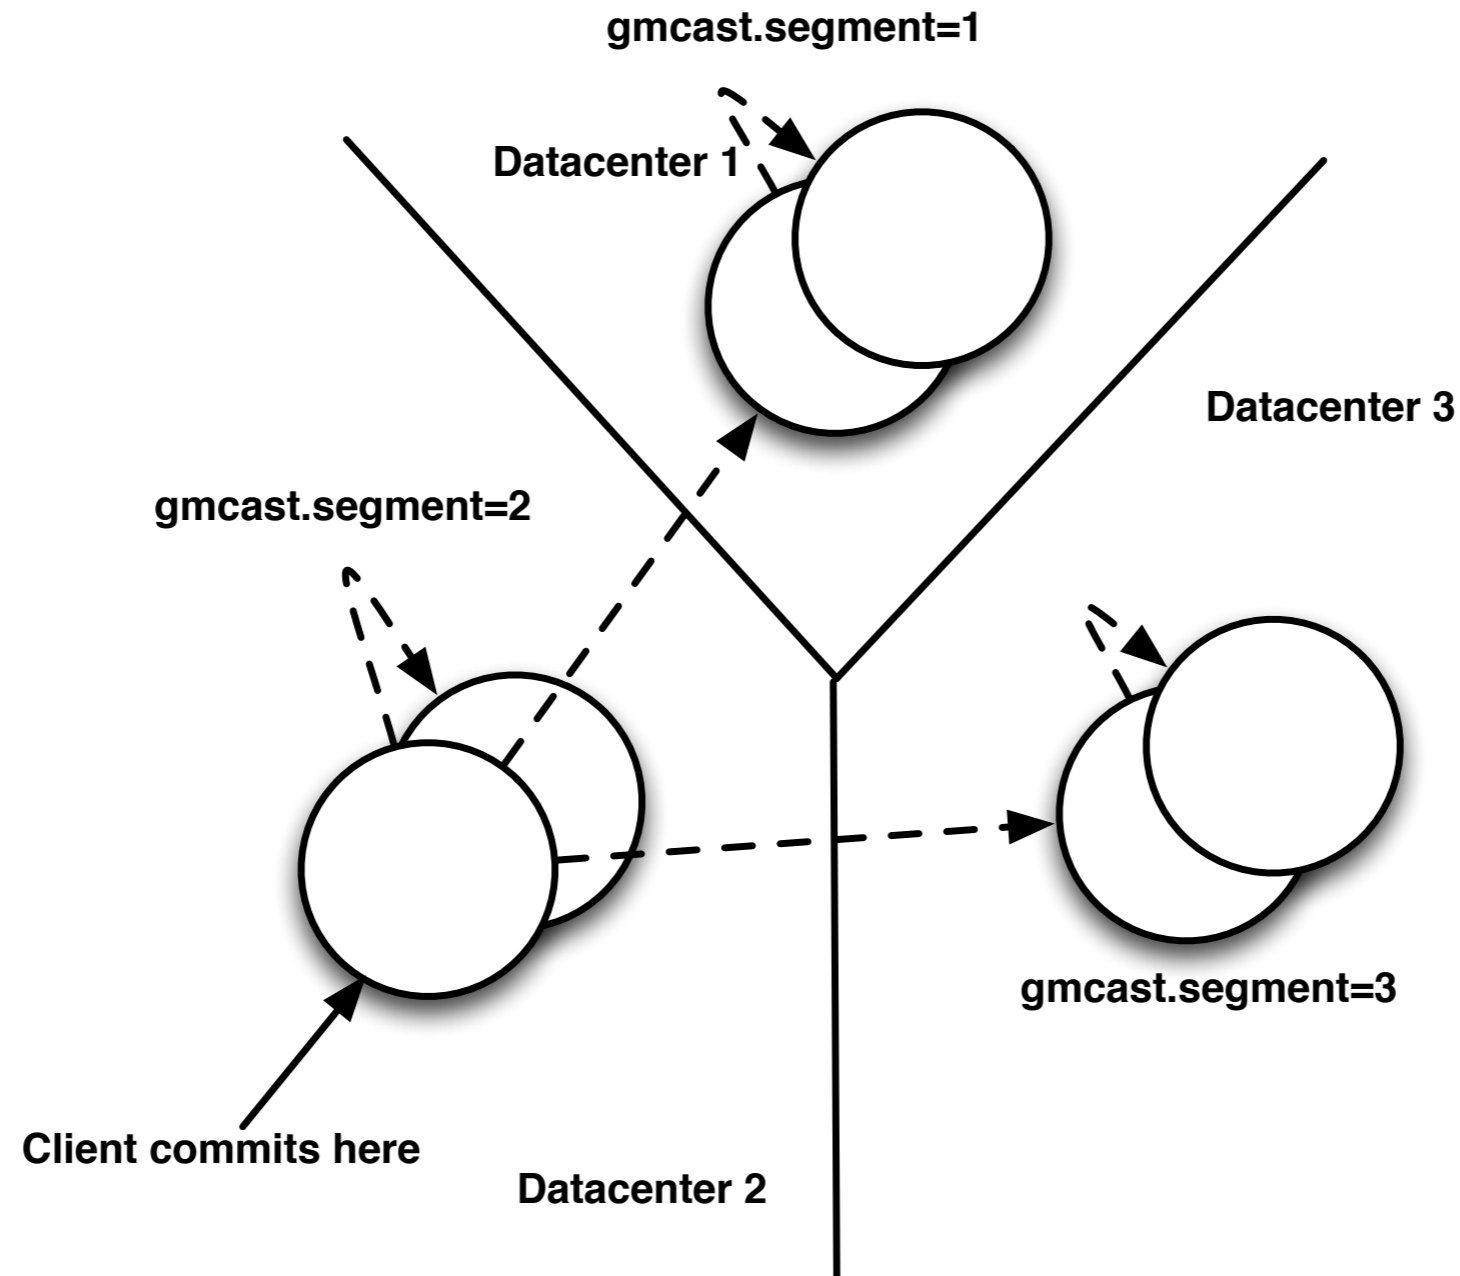
etc.

## Benefits

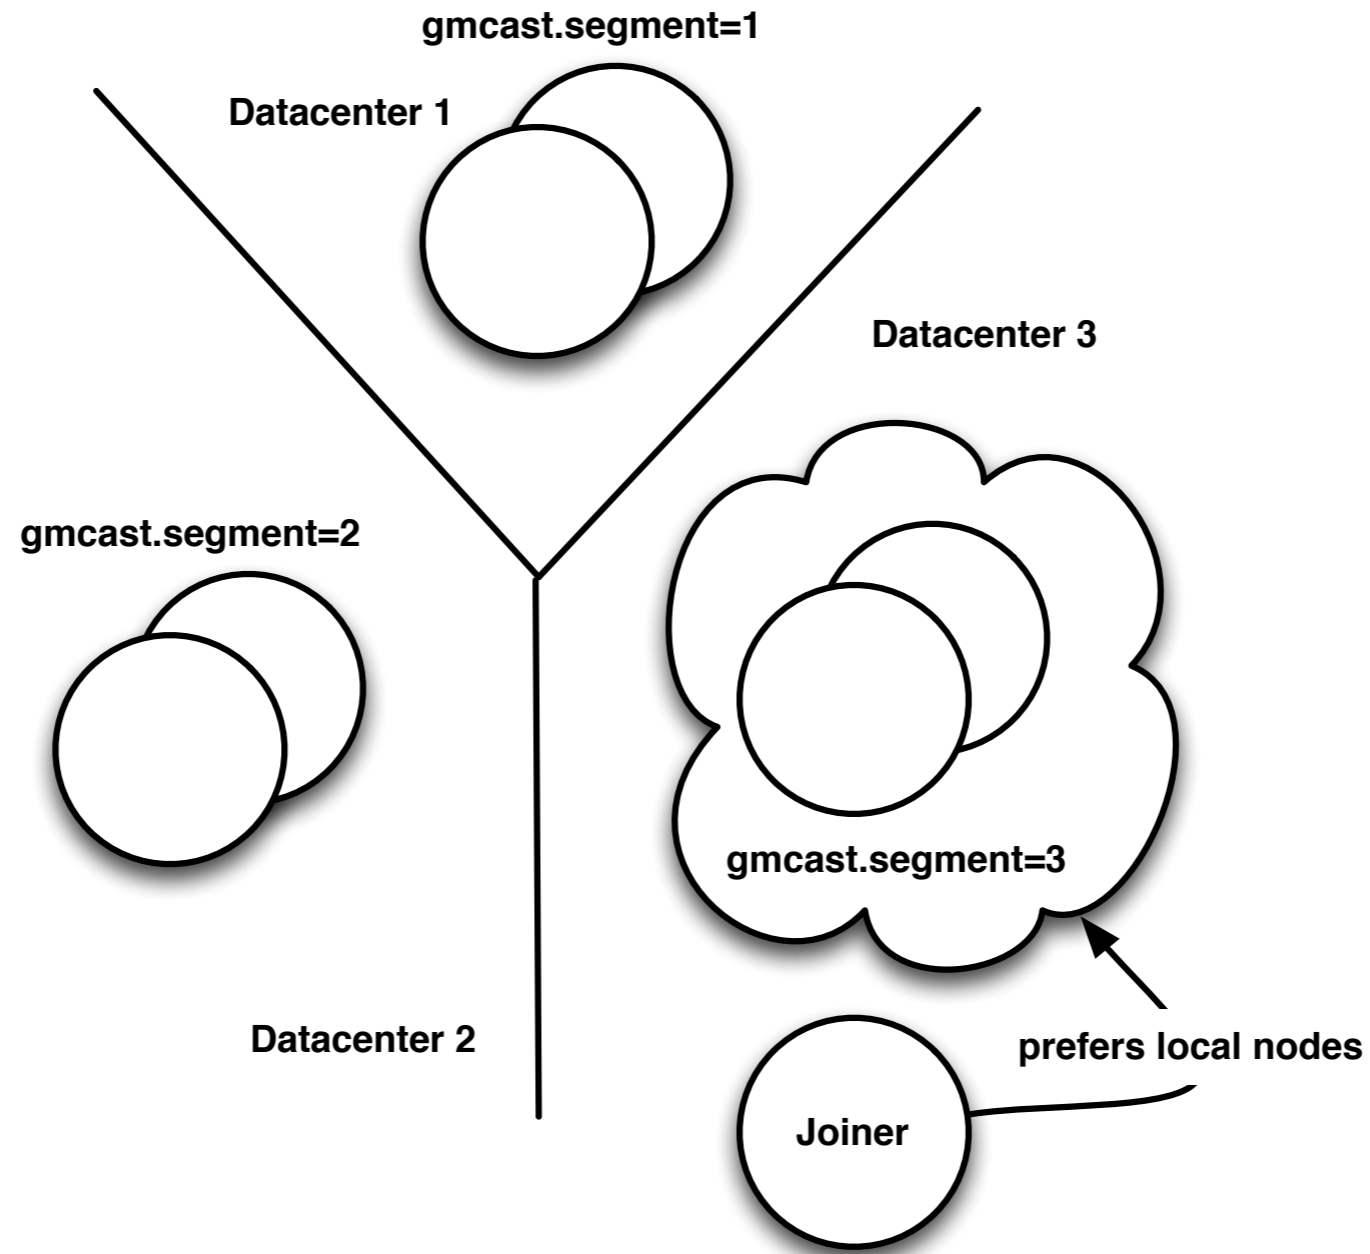
Replication traffic is minimized between segments

Donor selection prefers local segment

# Replication with Segments




# Donor Selection with Segments





# Asynchronous Replication into the cluster

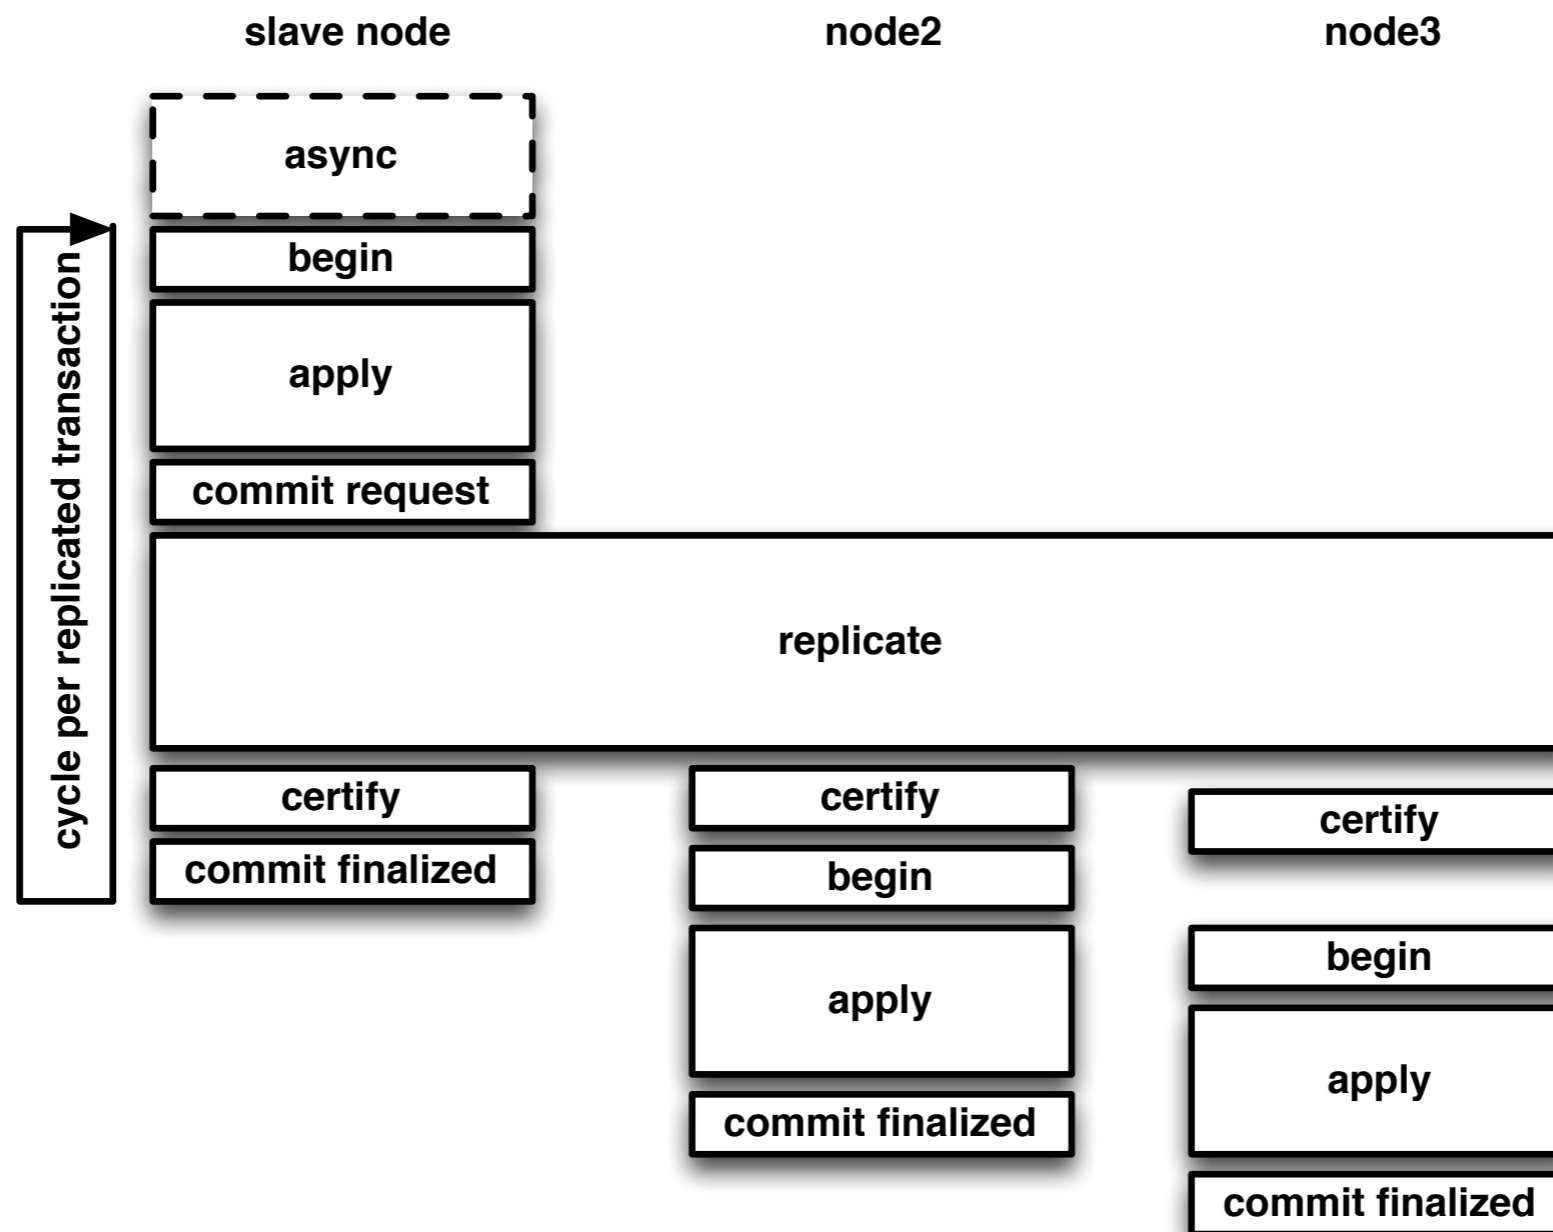
# Any node can be a slave





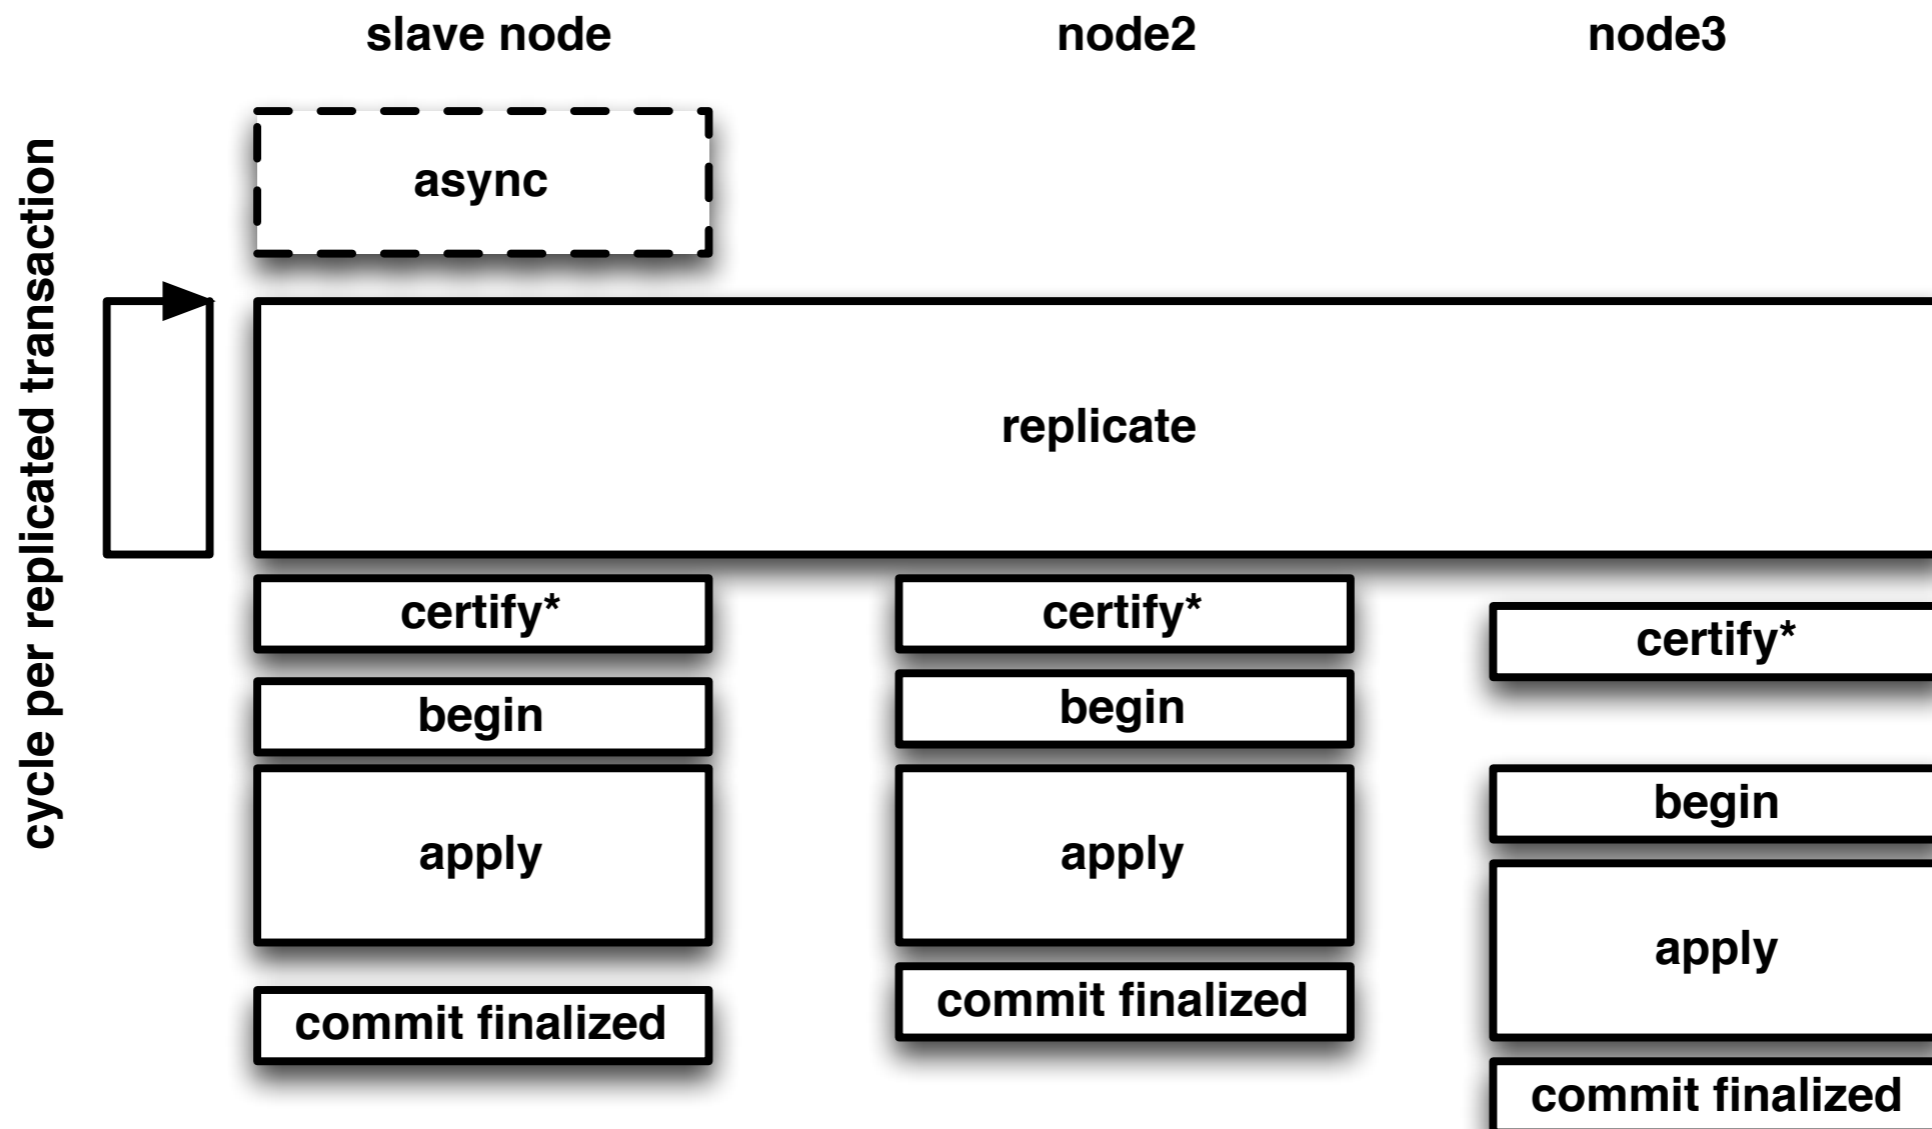
# Galera 2 behavior

Galera <= 2.x or wsrep\_preordered=OFF



# Galera 3 with wsrep\_preordered=ON

## Galera 3+ and wsrep\_preordered=ON



# When to use `wsrep_preordered`

`wsrep_preordered=ON`

Better performance

Does not allow for conflicts with any other writes

No parallel apply for these transactions

Good for Master / Slave to Cluster migration

`wsrep_preordered=OFF`

Detects conflicts with any other writes

Allows parallel apply

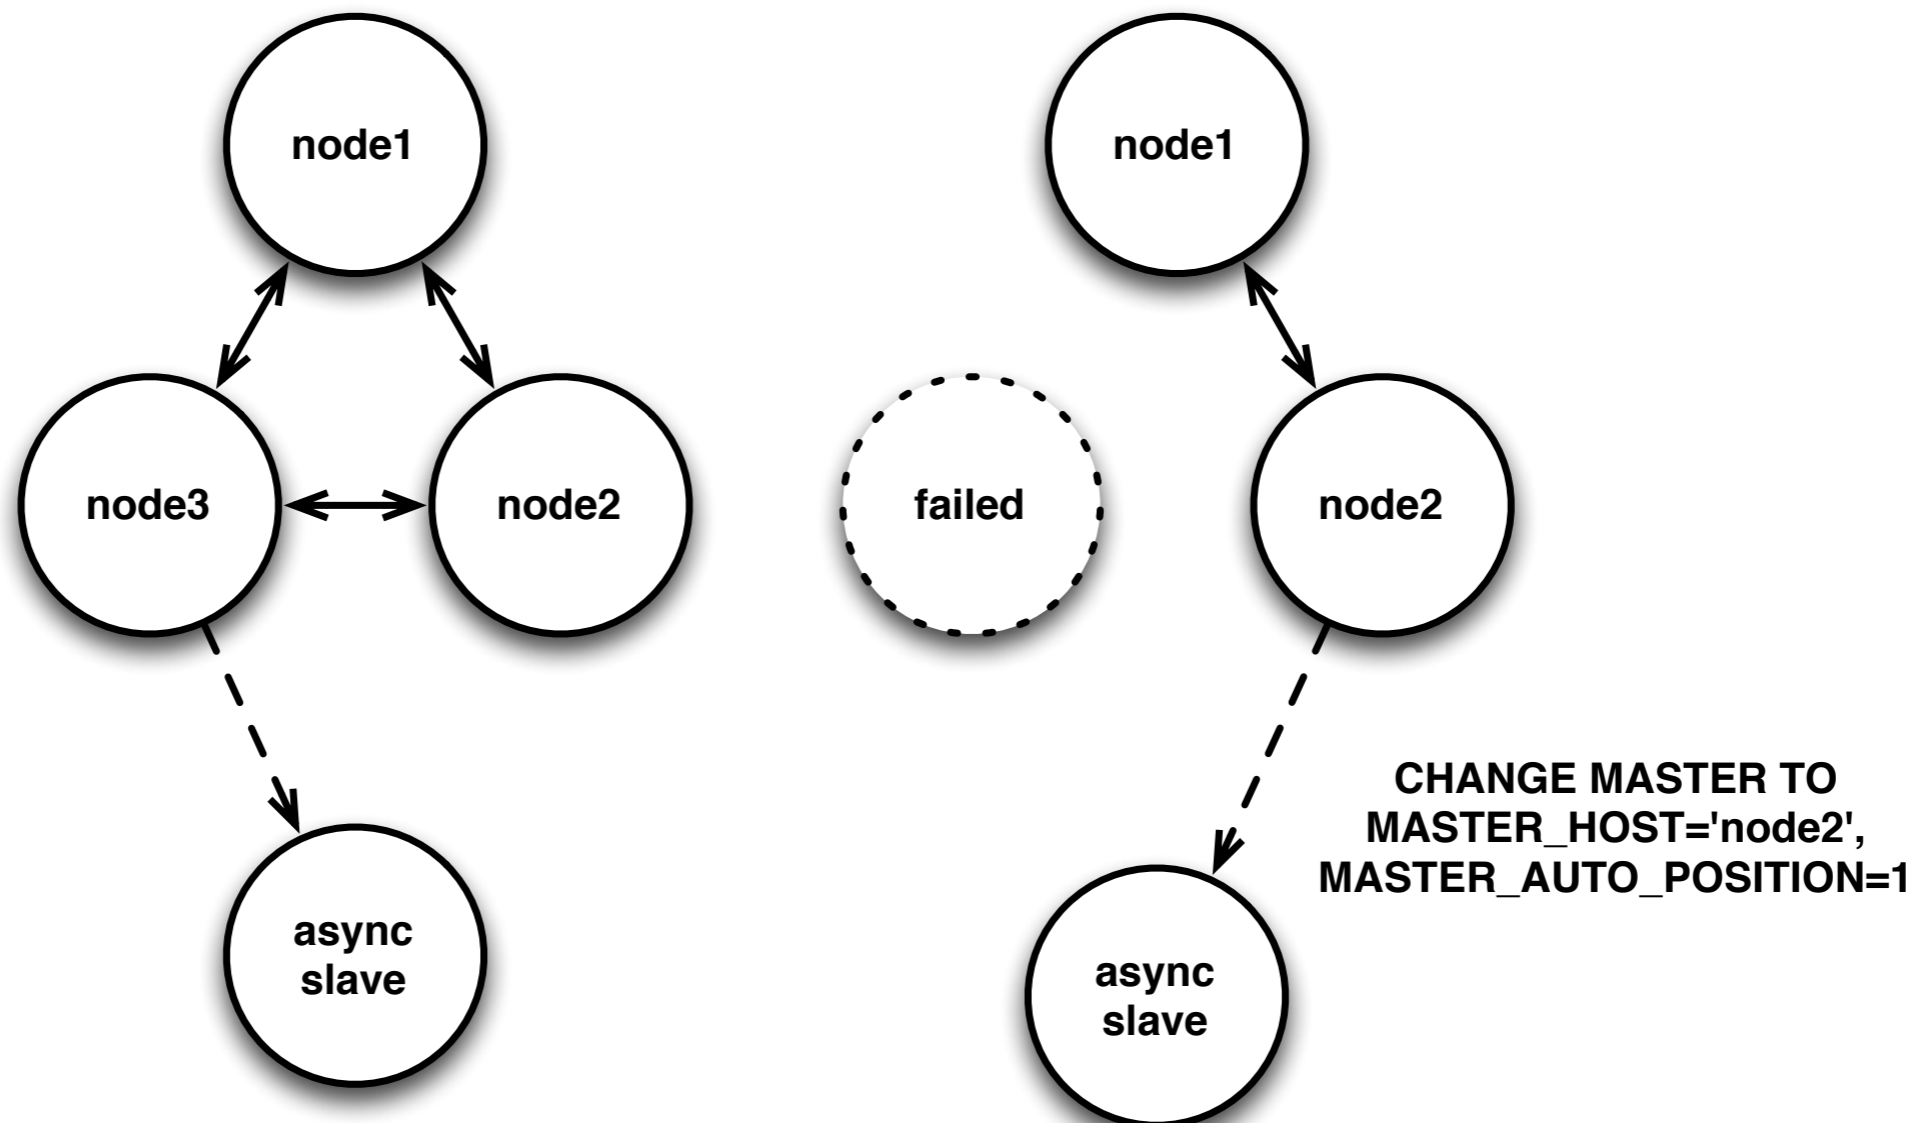
Good for permanent Cluster-is-a-slave



# Asynchronous Replication out of the cluster

# 5.6 async GTID integration

Every node with log-bin will have the same GTID for the same transaction





# General 5.6 Improvements

# Can I use <5.6 feature> in PXC?

Every new feature works fine within a single node

E.g., Innodb Full Text Indexes, Partitioning, Optimizer enhancements, Performance Schema, etc.

Don't expect automatic "cluster" support

E.g., Online DDL improvements

memcached server

<https://bugs.launchpad.net/percona-xtradb-cluster/+bug/1254126>

# Minimal RBR images


Pre-5.6 RBR is full row images for any row change

5.6 adds `binlog_row_image=minimal` (not default)

Seems fine with PXC

RBR image is black-box to Galera

1 Minute Sysbench Update Test (1 col out of 3 modified)







# Upgrading a 5.5 cluster to 5.6

# The easy way

Take the downtime and upgrade it all at once

Check for my.cnf settings that are not 5.6 compatible

Start each node with “wsrep\_provider=none”

Run mysql\_upgrade

Shutdown again

Bootstrap normally

# Rolling upgrade

RBR 5.6 -> 5.5 replication is broken

For each 5.6 upgraded node

Compatibility options with 5.5 must be set

Galera socket.checksum=1

**Don't SST 5.5 to 5.6!**

**Don't write on these nodes!**

Before last node, flip application to 5.6 nodes

Take down remaining 5.5 node(s) for upgrade

Rolling restart to remove 5.5 compat options

<http://bit.ly/pxc56-rolling-upgrade>



# Odds and Ends

# My Favorite Bug Fixes

Adding auto\_increment column to existing table doesn't cause inconsistency with auto\_increment\_control

wsrep\_local\_bf\_aborts catches all BF aborts now

wsrep\_flow\_control\_sent/received now global counters

wsrep\_max\_ws\_size/ws\_rows moving towards being properly enforced now\*

# Future Features

Automatic huge transaction fragmentation/  
streaming support

Non-blocking DDL support

Cluster tolerance to inconsistencies

Intelligent inconsistency handling (e.g., node voting)

Intelligent donor selection (e.g., check gcache)

Performance optimizations (i.e., multi-core)

Multiple provider support



**Questions?**



# PERCONA LIVE: MYSQL CONFERENCE AND EXPO 2014

Learn from leading MySQL experts. *Santa Clara, CA. April 1 - 4, 2014.*

Advanced Rates End March 2<sup>nd</sup> at 11:30pm PST, 2014

Special Discount for Webinar Attendees:  
Use Code **WebinarSC** to receive 10% off of standard  
rates (new registrations only)

[http://www.percona.com/live/mysql-  
conference-2014/](http://www.percona.com/live/mysql-conference-2014/)