Running virtualized Galera instances for fun and profit Linux Conf Au 2014 - Sysadmin Miniconf

Raghavendra Prabhu raghavendra.prabhu@percona.com

Percona

January 6th, 2013



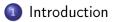
1 / 14

Raghavendra Prabhu (Percona)

Virtualized Galera

→ ◆ 豊 → ◆ 豊 → ■ January 6th, 2013





- 2 Galera and Virtualization
- Coming Together





э

Raghavendra Prabhu (Percona)

Virtualized Galera

January 6th, 2013 2 / 14

(日) (同) (三) (三)

Overview

- Synchronous replication
- Galera and PXC
- Virtualization
- High Availability
- Commonality



3

・ 何 ト ・ ヨ ト ・ ヨ ト

Galera/PXC

- What are they
- How do they work
- MySQL Compatibility
- Synchronous replication
 - Others
 - Google F1/Spanner
 - MySQL NDB
 - Sharding?



э

・ 同 ト ・ ヨ ト ・ ヨ ト

Virtualization

• Popular solutions

- Xen, KVM, OpenStack
- Containers LXC, Docker, ZeroVM
 - Application-isolated virt
- Application
- Constraints
- Advantages
- Data-centers
 - WSCs and clusters



∃ → (∃ →

Warehouse-scale

- Clusters
 - Hadoop, Manta et.al.
- Different from normal data-centers
 - Why?
 - Homogenity of hardware and stacks
 - Demands and constraints



∃ → (∃ →

Current scenario

- Popular Solutions
 - GlusterFS: FUSE
 - Sheepdog: QEMU
- Application level clustering
 - How is this different
 - Can work with them
 - Galera and others
 - Hadoop
 - {No,New,}SQL
 - Mapreduce/Compute



A B M A B M

Scenarios

- Necessity
 - Hardware of today
 - Saturation
 - WSC
 - High Availability
 - Migration
 - WANs
 - Isolation
 - Hosting



3

< 🗗 🕨

Deployment I

- Initialisation
 - Bootstrapping
 - Cloning
 - QCOW zero-copy: I/O
 - Reduced memory footprint: KSM
 - Thin-provisioning: Ballooning Upgrades!
 - Load-balancing galera



э

A B A A B A

< 🗗 🕨

Deployment II

- Operation
 - Migration
 - Live Migration: Streaming
 - Guest Agent
 - FTWRL: Xtrabackup
 - Shared storage
 - COW or COR LAN v/s WAN



10 / 14

э

< A

Deployment III

- Storage
 - Single host
 - VirtFS: share the root! Works well with QCOW
- Network
 - Latency Synchronous replication! Affects pre-commit stage
 - virtio-net/vhost-net



11 / 14

э

Deployment IV

Misc

- Backups
 - Dirty bitmaps
- Roles
 - Throttling
 - Reader Writer
 Upgrades, maintenance
 Non-homogenous hardware
 Async replication support



12 / 14

э

< 行い

Thoughts?

- Questions
- Usage and Experience



3

Raghavendra Prabhu (Percona)

Virtualized Galera

January 6th, 2013 13 / 14

・ロト ・ 日 ・ ・ ヨ ・ ・ ヨ ・

Me

- Raghavendra Prabhu
- http://wnohang.net/about
- Software Engineer at Percona and Product Lead of PXC
 - https://launchpad.net/percona-xtradb-cluster



14 / 14