

MySQL and DRBD Guide

Abstract

This is the MySQL and DRBD extract from the MySQL Reference Manual.

For legal information, see the [Legal Notices](#).

For help with using MySQL, please visit either the [MySQL Forums](#) or [MySQL Mailing Lists](#), where you can discuss your issues with other MySQL users.

For additional documentation on MySQL products, including translations of the documentation into other languages, and downloadable versions in variety of formats, including HTML and PDF formats, see the [MySQL Documentation Library](#).

Document generated on: 2017-04-28 (revision: 51931)

Table of Contents

Preface and Legal Notices	v
1 Overview of MySQL with DRBD/Pacemaker/Corosync/Oracle Linux	1

Preface and Legal Notices

This is the MySQL and DRBD extract from the MySQL Reference Manual.

Legal Notices

Copyright © 1997, 2017, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the

documentation in a manner similar to how Oracle disseminates it (that is, electronically for download on a Web site with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Oracle. Oracle and/or its affiliates reserve any and all rights to this documentation not expressly granted above.

Chapter 1 Overview of MySQL with DRBD/Pacemaker/Corosync/Oracle Linux

DRBD (Distributed Replication Block Device) is one of the leading solutions for MySQL HA (High Availability). When combined with Pacemaker and Corosync, users have:

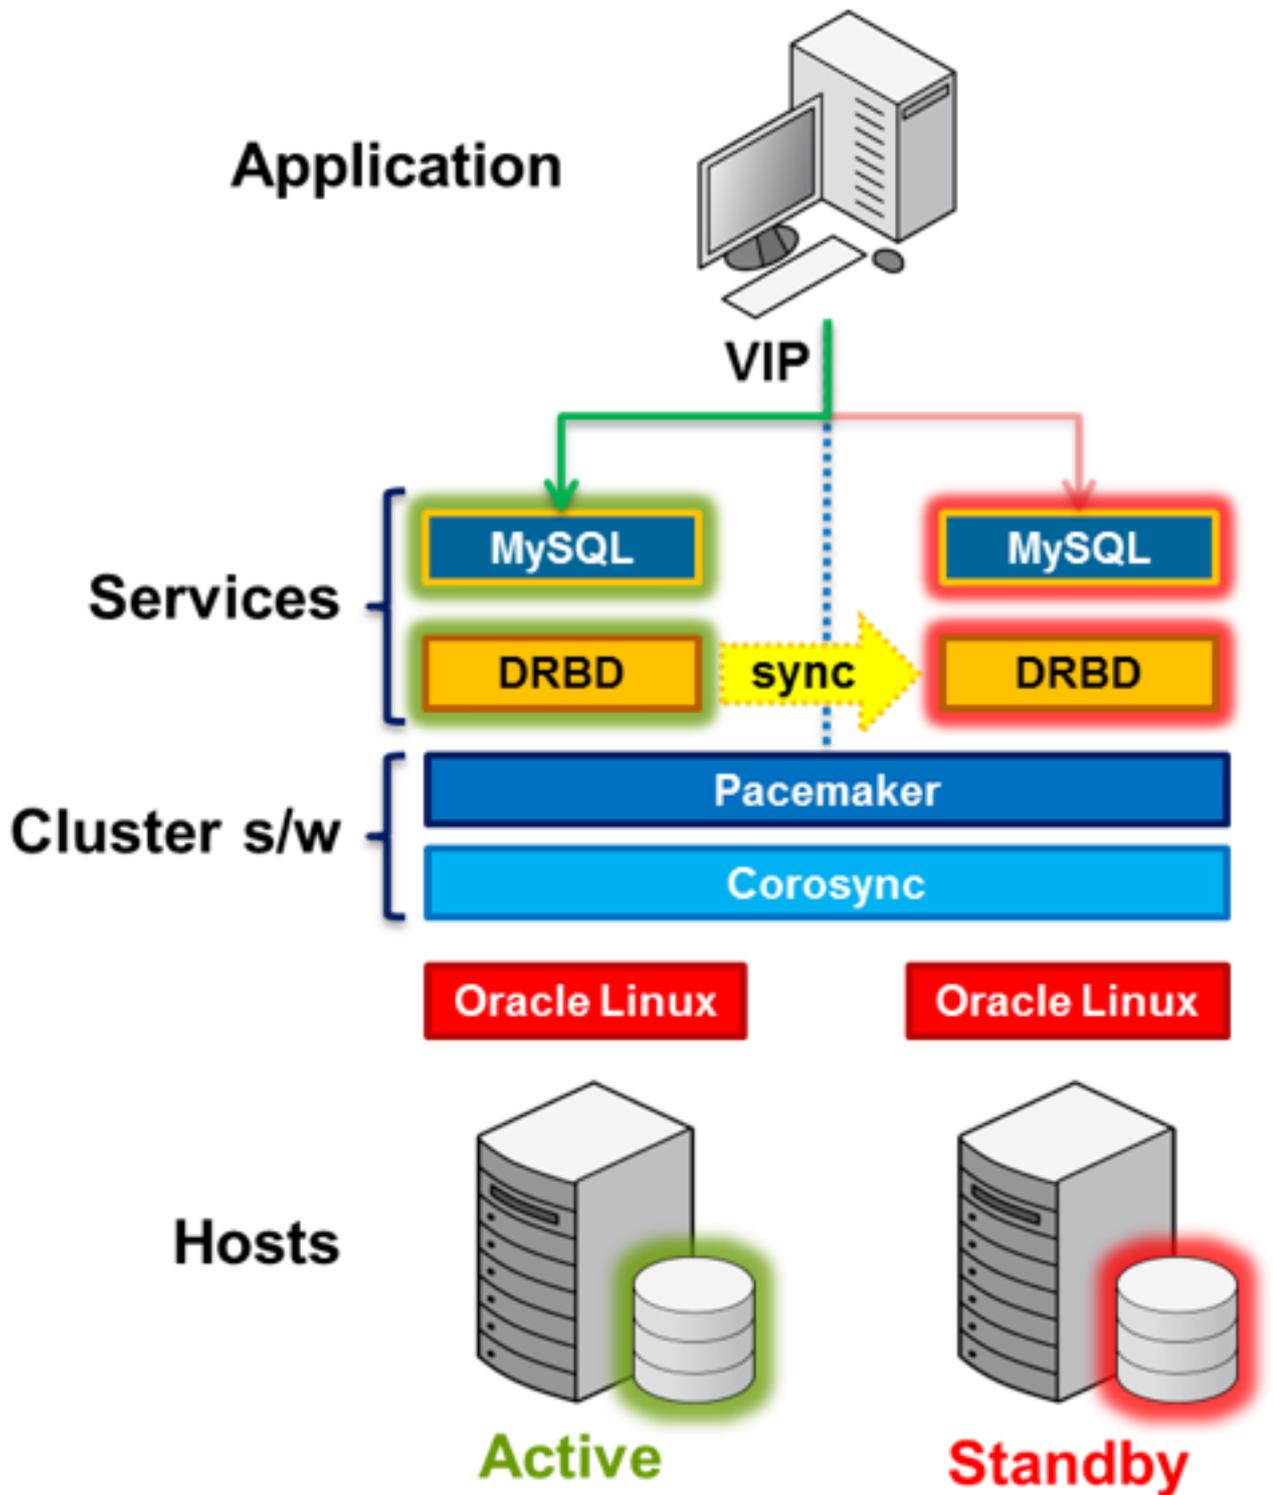
- An end-to-end, integrated stack of mature and proven open source technologies, fully supported by Oracle (as part of MySQL Enterprise Edition).
- Automatic failover and recovery for service continuity.
- Mirroring, via synchronous replication, to ensure failover between nodes without the risk of losing committed transactions.
- Building of HA clusters from commodity hardware, without the requirement for shared-storage.

The following figure illustrates the stack that can be used to deliver a level of High Availability for the MySQL service.

At the lowest level, 2 hosts are required in order to provide physical redundancy; if using a virtual environment, those 2 hosts should be on different physical machines. It is an important feature that no shared storage is required. At any point in time, the services will be active on one host and in standby mode on the other.

Pacemaker and Corosync combine to provide the clustering layer that sits between the services and the underlying hosts and operating systems. Pacemaker is responsible for starting and stopping services, ensuring that they are running on exactly one host, thus delivering high availability and avoiding data corruption. Corosync provides the underlying messaging infrastructure between the nodes that enables Pacemaker to do its job; it also handles the nodes membership within the cluster and informs Pacemaker of any changes.

Figure 1.1 MySQL, DRBD, Pacemaker, and Corosync Stack



The core Pacemaker process does not have built-in knowledge of the specific services to be managed; instead, it uses agents that provide a wrapper for the service-specific actions. For example, in this solution we use agents for Virtual IP Addresses, MySQL and DRBD: these are all existing agents and come packaged with Pacemaker.

The essential services managed by Pacemaker in this configuration are DRBD, MySQL and the Virtual IP Address that applications use to connect to the active MySQL service.

DRBD synchronizes data at the block device (typically a spinning or solid state disk) – transparent to the application, database and even the file system. DRBD requires the use of a journaling file system such as [ext3](#) or [ext4](#). For this solution, it acts in an active-standby mode: at any point in time, the directories being managed by DRBD are accessible for reads and writes on exactly one of the two hosts and inaccessible (even for reads) on the other. Any changes made on the active host are synchronously replicated to the standby host by DRBD.

Download the following guide for detailed instructions on installing, configuring, provisioning and testing the complete MySQL and DRBD stack, including:

- MySQL Database.
- DRBD kernel module and userland utilities.
- Pacemaker and Corosync cluster messaging and management processes.
- Oracle Linux operating system.

Download the guide at: <http://www.mysql.com/why-mysql/white-papers/mysql-high-availability-drbd-configuration-deployment-guide/>.

Support for DRBD

The complete DRBD stack for MySQL has been certified by Oracle. Commercial support, which provides a single point of contact for the entire stack, whether issues relate to the operating system, DRBD, clustering software or MySQL, is available to those who have both [MySQL Enterprise Edition](#) and [Oracle Linux Premier Support](#) contracts.

