

## **High Availability Messaging Solution** **Using the AXIGEN Mail Server, Heartbeat and DRBD**



July 08, 2008



## 1. Introduction

In today's business environments, we often hear the term 'high-availability'. Whether a hardware or software solution, we aim for redundancy at all levels, so as to maximize the availability of mission-critical services and operations.

As always, the hardware redundancy is more costly to achieve. Implementing a certain level of redundancy at the hardware layer is usually over ten times more expensive than obtaining the same redundancy only through the software layer. So whenever possible to substitute high-availability at the hardware level with a similar availability at the software level, one should definitely go ahead with it, given the cost-effectiveness of the solution, and especially if the process does not involve any performance drawbacks.

This white paper will discuss a high-availability solution for the AXIGEN Mail Server, using the Heartbeat package with DRBD. AXIGEN is a proprietary messaging solution while both Heartbeat and DRBD are open source software released under the GNU Public License (GPL). These software packages can be used together to build scalable and highly available integrated cluster messaging applications on the Linux operating system.

## 2. Products Overview

- **AXIGEN** (<http://www.axigen.com/mail-server/isp/>) – Leveraging on proprietary technologies SmartProcessing™, UltraStorage™ and GrowSecure™, the AXIGEN Mail Server delivers redundancy at the mail protocol level of a messaging solution, as well as for the actual message container storage.

AXIGEN is a carrier-class messaging solution with advanced email filtering and routing capabilities, also featuring groupware and advanced collaboration functionalities, and backed up at the same time by around the clock, highly effective technical support.

- **Heartbeat** (<http://www.linux-ha.org/Heartbeat>) – A fundamental part of the High-Availability Linux project, Heartbeat provides core cluster management services, including membership, communication, resource monitoring and management services, IP address takeover etc. Heartbeat version 1.2.3 supports multiple IP addresses and a simple two node primary/secondary model.

When used with AXIGEN in a cluster environment, multiple Heartbeat pairs, each of them consisting of two nodes, can be configured in order to support larger clusters with the additional help of a layer 4-7 load balancer, a LDAP centralized user database and, of course, AXIGEN proxy nodes. With the newly-released version 2.0.0 of Heartbeat, the two node size limit of the cluster is raised.

- **DRBD** (<http://www.drbd.org/>) – It is a block device or disk replication technology that can be viewed as network RAID-1. It can be used effectively to mirror a whole block device via a network onto another block device. Thus, DRBD involves two block devices, one labeled as primary (local) and the other labeled as secondary (remote/backup/standby).

Every write operation to the primary local device is written to disk and also sent to the other host across the network to be written to the secondary device. The remote host (secondary) writes the data to its configured disk. If the primary node fails, then the secondary node can take over in a typical failover scenario.

### 3. Solution Architecture & Implementation

In the solution examined, DRBD can be viewed as a resource that is controlled by Heartbeat. This 'node management' layer is transparent from AXIGEN's point of view. DRBD uses the underlying physical disk as a virtual device. This virtual device is essentially a DRBD resource that is used by AXIGEN to access the disk. When the DRBD resource is made 'active' on a particular node, it means that the disk configured as a DRBD resource is now accessible and ready for I/O operations.

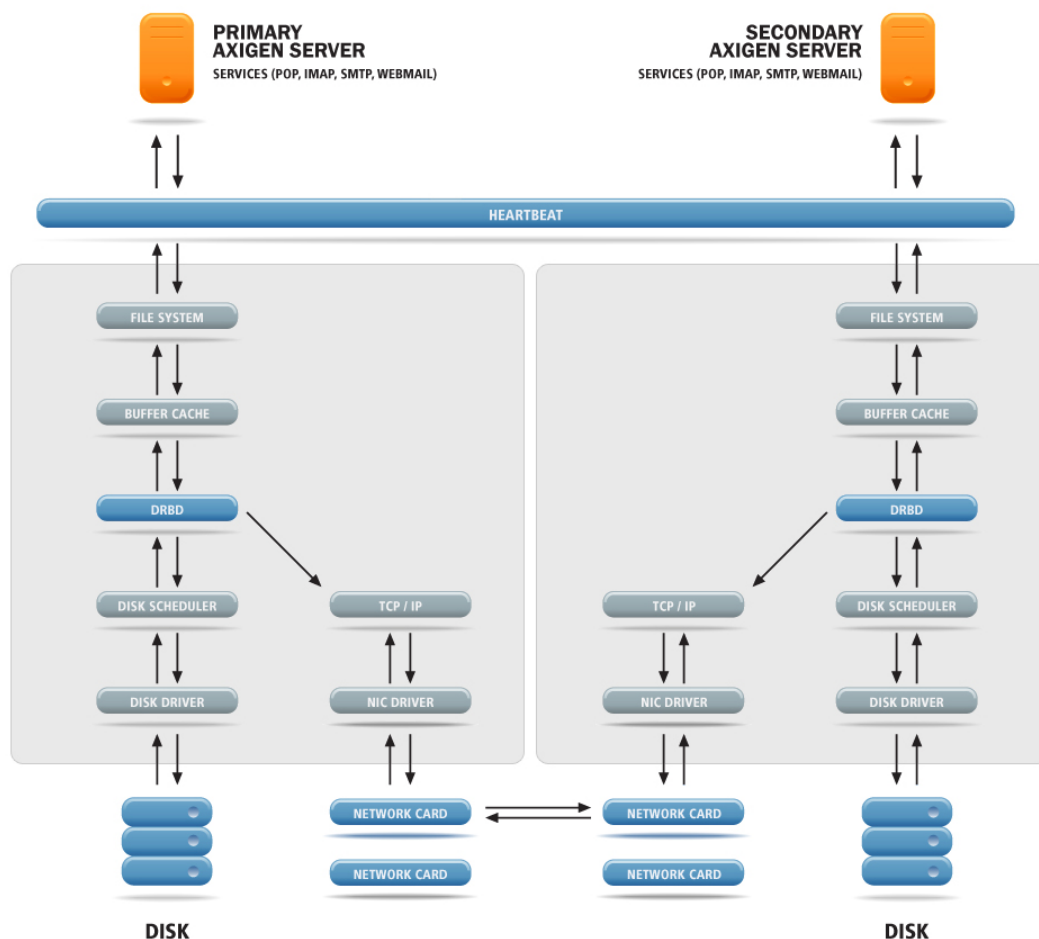


Fig. High Availability Messaging Solution with AXIGEN, Heartbeat and DRBD

### 4. Benefits of Using AXIGEN

The main advantages of using AXIGEN instead of another mailserver solution for this setup would stand in its proprietary UltraStorage™ architecture, with powerful features such as: keeping just a 'single-copy' of the email message and virtually linking it to multiple recipients, the internal caching layer which speeds up the access speed to the actual email from the message store, error checking and recovery logic which can handle unexpected hardware crashes that otherwise would have led to subsequent consistency losses, just to sum up some of the technologies used.

Starting with the current version of DRBD, a disk can be mounted only from the primary node; mounting concurrently in read-only mode from the secondary node is not allowed. This is a limitation by design. If more than one node is concurrently modifying the distributed devices, the process of deciding which part of the device is up-to-date and on which node, or what blocks of the device need to be resynchronized and in which direction, becomes very complex. If the purpose is to allow access to the data from multiple nodes concurrently, one should consider using a shared file system instead.

Due to the dual nature of the high-availability model being used, the 'split-brain' case must be taken in consideration. STONITH is a technique for node fencing, where the errant node which might have run amok with cluster resources is simply 'shot in the head'. This is actually where the STONITH acronym was derived from, as it stands for 'Shoot The Other Node In The Head'. Normally, when a high-availability system declares a node as dead, it is merely speculating that it is dead. STONITH takes that speculation and makes it reality.

AXIGEN ensures that this 'neutralization' of the supposedly 'dead' node is done without notable repercussions, by delivering a proper shutdown of all the messaging resources, be it mails from the queue or message store related mechanisms.

DRBD essentially provides disk replication across a network. Therefore, its performance largely depends on the I/O bandwidth of the physical hard drives that are used and the network bandwidth between the primary and secondary nodes. The main concerns in such a setup are performance and availability (failover), both of them being guided by the disk and network I/O throughput. AXIGEN makes the solution less depending on the aforementioned limitations due to its proprietary message storage which has cached reads and writes.

## 5. Additional Reading

For more information regarding the actual technical implementation details you can read the document at: [http://www.axigen.com/usr/files/ha\\_solution\\_axigen\\_heartbeat\\_drdb.pdf](http://www.axigen.com/usr/files/ha_solution_axigen_heartbeat_drdb.pdf)

### **AXIGEN Copyright © 2008 Gecad Technologies SA [AXIGEN]. All rights reserved.**

This material or parts of the information contained herein cannot be reproduced in any form or by any means without the prior written permission of AXIGEN. The product and the documentation that comes with the product are protected by AXIGEN copyright. AXIGEN reserves the right to revise and modify its products and documentation according to its own necessities, as well as this document content. This material describes a status, as it was in the moment this material was written and may not correctly describe the latest developments. For this reason, we recommend you to periodically check our website, <http://www.axigen.com/>.

AXIGEN cannot be held responsible for any special, collateral or accidental damages, related in any way to the use of this document. AXIGEN does not guarantee either implicitly or explicitly the suitability of this material for your specific needs. This material is provided on an "as-is" basis.