

Virtualisation, live migration, and data storage project plan

Attila de Groot, Michael van Kleij
attilla.degroot@os3.nl
michael.vankleij@os3.nl

February 24, 2009

Abstract

Virtual machine environments such as Xen or VMware, provide a function to migrate a running virtual machine between two hosts. To migrate a virtual machine the same data store has to be available from both locations. iSCSI is a solution to provide a storage system that is accessible from the two separate locations, however this introduces latency if a vm is migrated to a location far from the iSCSI node.

The goal of this project is to evaluate the possibility of using RAID or DRBD to create a mirrored data store with less latency to the virtual machine.

1 Introduction

Live migration functionality[1] of systems such as Xen and VMware allows an administrator to transfer a virtual machine from one host to the other without downtime. Live migration can help to provide a high availability environment, because services can still be provided while performing maintenance or in case of hardware failure.

When transferring a virtual machine, both need to have access to the same data store for their file system. This creates the problem if a virtual machine needs to be transferred over a WAN[2]. This can be solved by using iSCSI[3] to let virtual machines at both locations have access to the same data store. When a virtual machine is transferred to another continent, this introduces latency to the data store from one of the virtual machines.

To solve the latency problem, research has been done[4] to migrate the data store along with the virtual machine to the other location. It has been shown that a fully functional web server with data store can be migrated over a WAN. With this method a migration still has to be planned and a virtual machine can't be resumed in case of a hardware failure as the data store isn't migrated. The goal of this research is to evaluate the possibility of using RAID or Distributed Replicated Block Device[5] (DRBD) to create a mirrored data store. This would create the possibility to migrate a virtual machine to another location with minimal downtime. Because the data store is continuously mirrored a resume in case of a hardware failure is also possible.

2 Planning

The activities that are defined in the previous section are planned over the next four weeks as shown in the next sections.

2.1 Week 9

Goal: Literature study, Implementation of live migration

- Study available papers on iSCSI, DRBD, and virtualisation migration.
- Implement Xen and live migration on OS3 servers with iSCSI

2.2 Week 10

Goal: Use raid to mirror iSCSI nodes

- Implement raid to create a mirror of two iSCSI nodes from a virtual machine.
- Test live migration between hosts.
- Research latency with iSCSI solution.

2.3 Week 11

Goal: Use DRBD to mirror iSCSI nodes

- Implement DRBD to create a mirror of two iSCSI nodes from a virtual machine.
- Test live migration between hosts
- Research latency with DRBD solution

2.4 Week 12

Goal: Paper on virtualisation and mirroring with raid and DRBD

- Write paper

References

- [1] C. Clark *et al.*, Live migration of virtual machines, 2005, https://www.usenix.org/publications/library/proceedings/nsdi05/tech/full_papers/clark/clark.pdf.
- [2] F. Travostino *et al.*, Seamless live migration of virtual machines over the man/wan, 2006, <http://www.francotravostino.name/papers/VMT.pdf>.
- [3] J. Satran *et al.*, Internet small computer systems interface, 2004, <http://www.ietf.org/rfc/rfc3720.txt>.

- [4] R. Bradford *et al.*, Live wide-area migration of virtual machines including local persistent state, 2007, <http://portal.acm.org/citation.cfm?id=1254810.1254834>.
- [5] P. Reisner, Distributed replicated block device, 2002, http://www.drbd.org/fileadmin/drbd/publications/drbd_lk9.pdf.