openlab tests Oracle VM for database service virtualization

Why database virtualization

There is a combination of factors that leads to increased interest for virtualization of database servers: number of servers, power usage, cost, new platforms with multi/many cores, high availability and ease of management. Database servers often do not use the full CPU/server capacity. This can be for reasons such as separate test-and-development servers or because spare capacity has to be kept available in case of increased load or for high availability.

In addition, several database services can often not be merged onto a single database or physical "machine" as they need to be on different versions of the database software or operating system. We also need to "protect" each of the database instances from the others so that memory can be guaranteed for a given instance.

Having a smaller number of servers will obviously reduce the cost of hardware, maintenance and electrical power. The power usage is especially important due to the constraints on "critical power" in computer centres as well as the impact on a cooling system.

We believe that consolidation with virtualization will become more and more interesting, in particular given that processors will continue to have an increasing number of cores in the medium term future. Another attractive opportunity with virtualization is the possibility to relocate and distribute virtual machines on different nodes depending on the load and intervention needs. This leads to higher availability while having fewer constraints on the operations to be performed.

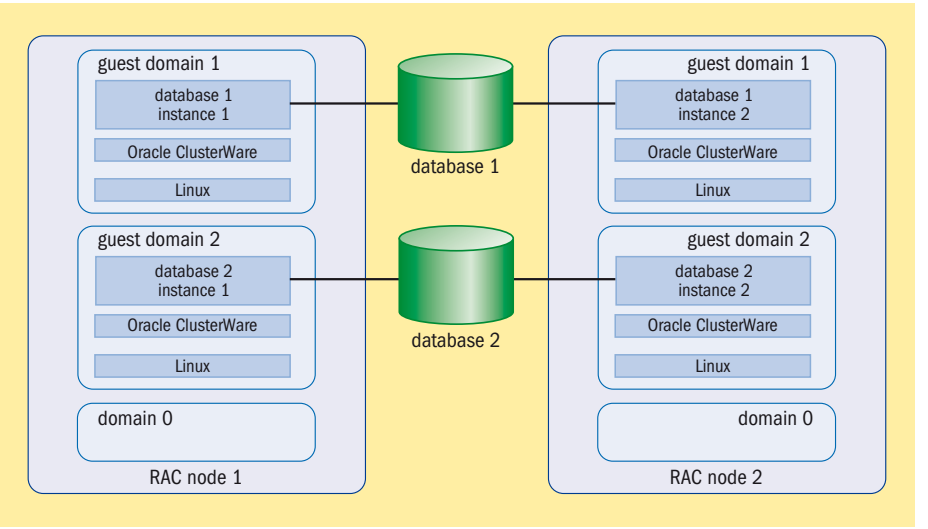

Fig. 1. Fighting underutilization by combining several clustered databases on two servers.

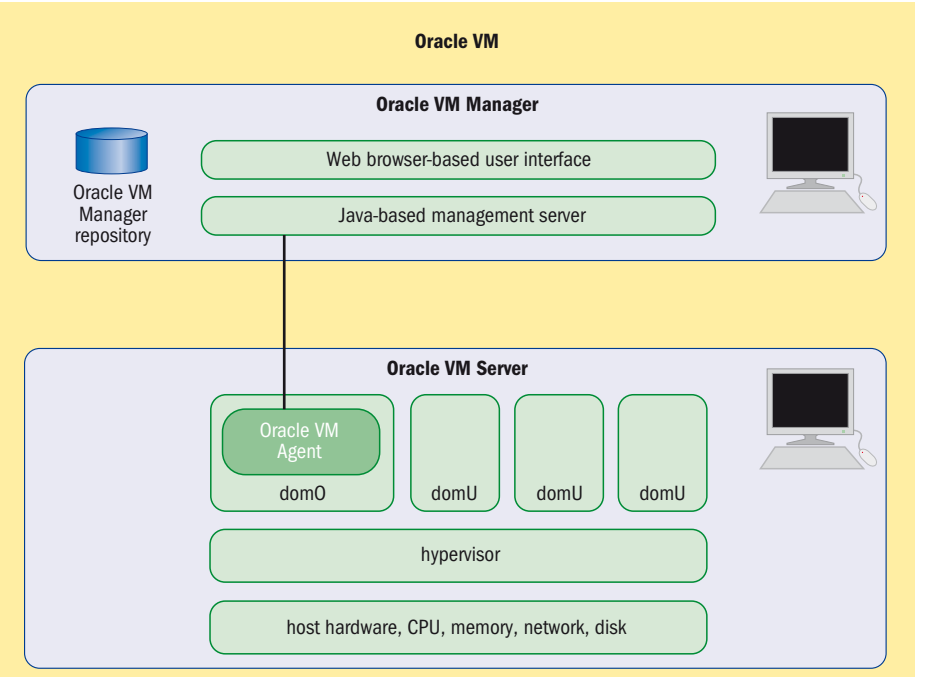

Fig. 2. Architecture: Oracle VM Manager and Oracle VM Server are installed on separate servers with communication provided by the management server and Oracle VM Agent.

Database virtualization at CERN openlab

In the CERN openlab Database Competence Centre, we started looking at virtualization in early 2006. The first project was aimed at testing the Oracle database in the Xen virtualized environment. The main focus was on the possibility to run a single-instance database in a virtual machine (VM), as well as testing different Xen schedulers to allocate CPU slices to different virtual machines. On successful completion of this project, we decided to expand our tests, and include Oracle Real Application Clusters (RAC) in such an environment.

RAC – an Oracle technology widely used at CERN and in the HEP community – allows several physical machines to work as a single clustered database, thus providing scalability and high availability. High availability is the main driving force behind implementing RAC in less workload-intensive infrastructure databases, which also leads to underutilization of the database nodes. Bearing this in mind, we were attracted to the idea of running RAC in a virtualized environment, thus utilizing more of the hardware without compromising the high availability features (see figure 1).

The concept of RAC on VM was proved to be successful by openlab summer student Maria Leitner in 2007. She created a cluster database inside a single physical host with Red Hat Enterprise Linux 4 as a base operating system. The set-up was quite stable and the performance loss

due to the virtualization layer was at an acceptable level.

In November 2007, Oracle announced its virtualization solution: Oracle VM, which consists of three parts: Oracle VM Server, Oracle VM Manager and Oracle VM Agent (see figure 2).

Based on Red Hat Enterprise Linux (RHEL) and open-source Xen, Oracle VM Server is the base virtualization software to be installed on a bare metal server. The Oracle VM Agent is the application, installed together with the Oracle VM Server and is responsible for communication with Oracle VM Manager Web-based GUI.

Oracle then certified this single instance (November 2007) and RAC (September 2008) databases in the Oracle VM-based virtualized environment, which was the step we were waiting for.

It was decided to continue with our tests and compare RAC performance on Oracle VM versus the native Xen set-up we tested before. Together with openlab summer student Andrei Dumitru in 2008, we created two set-ups on identical hardware, using Oracle VM and pure Xen, included in RHEL 5. The stability of both set-ups proved to be good and the performance of the Oracle VM set-up was 5 to 20% better than on pure Xen, depending on the workload. Tests of live VM migration with an active database showed only a few seconds downtime with no session-state loss. This is an outstanding addition to the high availability features initially provided by RAC.

Following these successful tests, we are now working closely with IT-FIO to integrate this Oracle VM solution into the CERN fabric management system – ELFms.

In summary, these Oracle VM features, their ease of use and the official certification have directed our choice towards Oracle VM as a virtualization solution for our virtualized cluster databases for the future.

Useful links

CERN openlab: http://cern.ch/openlab (with openlab summer student reports in Technical Documents section)
Oracle VM: www.oracle.com/virtualization

Anton Topurov and Eric Grancher, IT-DES

A great success: 4th EGEE User Forum attracts 600 attendees

By the sea, among lemons and sunshine, more than 600 members of the European and international grid community – representing 45 countries – gathered in Catania, Sicily, Italy for the 4th Enabling Grids for E-sciencE (EGEE) User Forum co-located with Open Grid Forum 25 & OGF – Europe's 2nd International Event.

Home base for EGI announced

The most eagerly awaited news to come from the event was the announcement that Amsterdam was selected as the host city at the European Grid Initiative (EGI) policy board meeting on Monday 2 March, ahead of seven other European cities that also bid to host the central organizing body of EGI.

"When I heard the decision I thought, 'Yes! We did it.' This will be a great opportunity for Amsterdam, hosting such an important international organization and we especially hope that the staff of EGI.org will benefit from our beautiful city," said David Groep, NIKHEF, member of the successful bidding team from Amsterdam.

"I'm very grateful to the people who voted for Amsterdam. I know how hard the selection process was. There were other strong contenders that were also excellent options. Of course," he said "I think Amsterdam was the best choice. I hope the location in Amsterdam will help EGI be a great success. We'll do our best to ensure the success of EGI, to keep the users happy so that we get great science."

Some 600 members of the European and international grid community (representing 45 countries) attended the 4th EGEE User Forum. Image courtesy of EGEE.

New memorandum of understanding with BalticGrid-II

On Wednesday 4 March, the project directors of EGEE-III and BalticGrid-II, Bob Jones and Ake Edlund signed a Memorandum of Understanding detailing BalticGrid-II's involvement with EGEE on an infrastructure and application level as the European grid community moves towards the European Grid Initiative.

"More than just an MOU," said Edlund, "this document involves a work plan and concrete targets with names and dates."

"This is an important step for the European grid community," said Jones. "as we work towards a united e-infrastructure in all European states in preparation for EGI."