



Contribution ID: 15

Type: Oral

A new approach to High Availability

Wednesday 4 March 2009 11:00 (25 minutes)

High availability has always been one of the main problems for a data center. Till now high availability was achieved by host per host redundancy, a highly expensive method in terms of hardware and human costs. A new approach to the problem can now be offered by virtualization. Using virtualization, it is possible to achieve a redundancy system for all the services running on a data center.

Impact

This new approach to high availability needs only one remote controller, while each one of the physical machines with a virtual environment on board can act as backup server, to recover services running on a crashed physical host. Even the remote controller can run shared among the physical servers of the computing center, to reduce the possibility of a single point of failure in case of disaster.

The system is able to recover in less than 3 minutes a crashed machine (e.g. overloaded), and in less than 7 minutes a critically damaged host (e.g. where the /boot partition has been deleted). To virtualize the main servers of a data center, a new procedure has been developed to migrate physical to virtual hosts. The whole Grid data center SNS-PISA, part of the EGEE/LCG CERN production Grid, is running at the moment in virtual environment under the high availability system.

URL for further information

<http://3rc.sns.it>

Conclusions and Future Work

Exploiting virtualization and ability to install a host from scratch in a completely automatic manner, it is possible to achieve a sort of host on demand, where the start-up of a backup virtual machine is done only when the disaster occurs. As extension of the 3RC architecture, several storage solutions will be tested to store and centralize all the virtual disks, from NAS to SAN, to grant data safety and access from everywhere.

Keywords

High Availability, Virtualization, Host on demand, PXE

Detailed analysis

This new approach to high availability allows to distribute the running virtual machines over the only servers up and running, by exploiting the features of the virtualization layer: start, stop and move virtual machines between physical hosts. The system (3RC) is based on a finite state machine, providing the possibility to restart each virtual machine over any physical host, or reinstall it from scratch. A complete infrastructure has been developed to install operating system and middleware in a few minutes. When the damage concerns a physical host, the hosted virtual machines are automatically moved to other physical hosts, by acting a selection of the lower loaded among the available servers.

Author: Dr CALZOLARI, Federico (Scuola Normale Superiore - INFN Pisa)

Co-authors: Dr CIAMPA, Alberto (INFN Pisa); Dr MAZZONI, Enrico (INFN Pisa); Dr AREZZINI, Silvia (INFN Pisa)

Presenter: Dr CALZOLARI, Federico (Scuola Normale Superiore - INFN Pisa)

Session Classification: Novel Services and Technologies

Track Classification: Emerging Technologies within the EGEE infrastructure