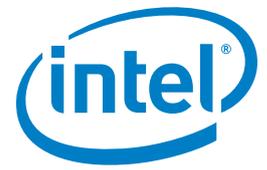


CASE STUDY

Intel® Xeon® processor 5600 series

Enterprise Server

Virtualization



Putting Italy on the cutting edge of scientific computing

Virtualization-based solution developed by INFN-CNAF brings the Grid and Cloud models closer

The Italian National Institute for Nuclear Physics (INFN) operates an organization in Bologna known as CNAF – the National Center for Research and Development in Information and Data-Transmission Technologies. CNAF is responsible for the management and development of the most important information and data transmission services to support INFN's high-energy physics research at a national level. Its research activities are divided into five scientific categories: accelerator physics, astroparticle physics, nuclear physics, theoretical physics and technological development.



“Our data centers support thousands of users and tens of diverse communities. Through efficient use of virtualization technologies, we have been able to expand our offerings and to integrate Grid and Cloud services.”

Davide Salomoni
Computing Research Director, INFN CNAF

CHALLENGES

- **Enhance infrastructure:** Provide the INFN community with a scalable and flexible solution for high-performance scientific computing
- **Guarantee continuity:** Deliver operating system support and scientific data availability for long-term data access at sustainable total cost of ownership (TCO)
- **Expand the customer base:** Offer new and enhanced services

SOLUTIONS

- **Integrated framework:** Implemented on-demand grid/cloud framework for scientific computing, based on open-standard technologies
- **Performance penalties minimized:** Physical and virtual environments have fine-tuned hardware and software solutions and efficient access to large-scale storage systems

IMPACT

- **World-first:** One of the first proven, OS-based implementations to achieve excellent scalability and flexibility in providing shared access to resources and integration between Grids and Clouds – without the need to partition resource pools
- **National use:** The INFN Worker Nodes on Demands Service* (wNoDeS*) framework is the production solution being offered for Grid and Cloud integration by the Italian Grid Initiative (IGI)

Shared Resources for Collaborative Research

For a number of years now, CNAF has worked in close collaboration with the particle research facility at CERN in Geneva, Switzerland, known as the Large Hadron Collider (LHC). The enormous volume of data produced by the LHC led CERN to choose a distributed computing model – the Grid – which entails the use of a number of computing centers across Europe, organized hierarchically. The top-level centers, known as Tier-1, are the main nodes in this infrastructure, and CNAF is Italy's Tier-1 center. There are currently 20 experiments making use of CNAF resources, four of which are being conducted in the LHC.

“The LHC project,” states Mauro Morandin, director of CNAF, “is extremely innovative, not only in terms of physics – in that it will allow us to obtain precise measurements for high-energy physical processes that have never been explored before – but also in terms of computer science. With over 200 data centers operating in 34 countries throughout Europe, Asia, and the Americas, combining to provide over 50,000 CPUs and hundreds of petabytes of mass storage overall, it is the world's leading Grid infrastructure. The Grid model makes it possible to provide access to the disk, tape, and computing resources needed to process huge amounts of data in a much more effective, flexible, and sustainable manner than could be done at one centralized computing center.”

Sharing IT resources is something that people are talking a great deal about but that few organizations, whether public or private, are seriously tackling. One of CNAF's goals is to make its expertise available to a growing number of research fields, which may range, for example, from medicine, to cultural heritage, to the environment. These are all worlds that have specific needs and use varied computer configurations. It is for this reason that the center has developed a framework, wNoDeS, to allow for increased flexibility through its use of virtualization technology.



Italian IT leader brings flexible Grid- and Cloud-based research capabilities to international researchers

Grid and Cloud: Bridging the Gap

Davide Salomoni, computing research director at the Tier-1 INFN center and WNoDeS project manager, explains how CNAF is using virtualization to support a diversifying user base: "We were able to develop software that allows for the creation of on-demand computing resources. These resources can be made available only when necessary, can be fully customized, and accessed through new access channels, thus expanding the user base. In fact, WNoDeS (<http://web.infn.it/wnodes>) creates virtual machines through a dynamic, transparent process, and they are specifically configured to meet the users' demands without the need for static resource partitioning." The virtual machines created by WNoDeS can be defined by the provider depending on the users or groups of users, or they can be set directly by users themselves, using resource-allocation mechanisms via Grid or Cloud interfaces.

The way in which these virtual machines are accessed is worth a closer look: in addition to guaranteeing access through a Grid interface, WNoDeS provides access to the same pool of computing resources through a Cloud interface as well. Although the logic behind these two models is the same – to allow access to shared resource pools – the purpose, and often the tools used, may be considerably different. It is no coincidence that Grid infrastructure has been widely adopted by international scientific projects focused on centers like CERN, while its use in the enterprise, which is more interested in the Cloud, remains marginal. The WNoDeS software, developed by CNAF, can provide a bridge between these two worlds, because it allows for complete integration with resource-allocation policies, monitoring, accounting, and security with regard to both Grid and Cloud.

At present, around 2,000 dynamically created virtual machines have been activated at CNAF, and these correspond to about a third of the overall available resources. The goal is to gradually increase the percentage, proceeding very cautiously, to ensure that no existing projects have the slightest interruption. The virtual machines are used to provide custom execution environments supporting the needs of specific scientific collaborations. For example, some collaborations have asked for their job or instantiations to have specific versions of operating systems, custom libraries, or ad-hoc services like distributed database services. These use cases would not have been possible without the flexibility and scalability offered by the WNoDeS framework.

"The key factor convincing us to take up the concept of virtualization was the performance level reached by this technology," continues Salomoni. "We purchase our equipment with public tenders based on objective benchmarks targeted at measuring key factors like density, power consumption, CPU performance and TCO. We were persuaded to deploy our solution when we saw that Intel® processor-powered virtual machines guaranteed performance almost equaling that of physical machines, with an average dip of 3-4 percent in terms of performance levels in CPU-intensive computations. Throughout all phases of the project, from its conception to final testing and all the way to production, the support of Intel, whose architecture is at the base of the vast majority of our systems, has been invaluable. In addition to developing solutions that we found to be cost-effective and very competitive performance-wise, Intel has also allowed us to access both Italian and international resources that are able to meet all of our requests, sharing know-how of the highest quality. The introduction of the latest processors, such as the Intel® Xeon® processor 5600 series, means excellent performance and full functionality for virtualized environments, now that adequate security can be guaranteed throughout all phases of data processing – from start-up to migration to filing. For example, the CNAF now has approximately eight petabytes of disk space available and 10 petabytes of tape storage, and every day it handles something like 50,000 jobs, around 11,000 of which are carried out simultaneously."

Spotlight on CNAF

CNAF is the National Center for Research and Development into the field of information technology for high-energy physics. It is operated by INFN, Italy's National Institute for Nuclear Physics, and is based in Bologna.

In 2005 CNAF became the Italian Tier-1 center for LHC experiments, making it INFN's main center for computer calculus and one of the most important centers for distributed calculus in Italy.

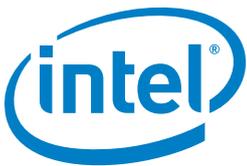
CNAF employees have worked for many years on the development and management of Grid infrastructures at an Italian (GRID.it), European (EGEE and EGI) and International (LHC Computing GRID) level.

Salomoni adds: "In addition, we have just recently started evaluating the advantage of the Intel® solid-state drive (SSD) solutions and the new Single Root I/O Virtualization* (SR-IOV*)-based 10-gigabit Ethernet controllers can bring to our WNoDeS architecture." Maintaining dialogue and open discussion in the field of high-performance distributed computing is absolutely fundamental, and CNAF is actively participating in the work of numerous organizations, through national and international projects, to study developments in this field, paying careful attention to open-source software and interoperability.

Moving Forward

The WNoDeS solution, which makes use of open-source technology and standard interfaces such as the gLite*/EMI* software used within the European Grid Infrastructure (EGI) and Open Cloud Computing Interface (OCCI) for the Cloud, has already been installed at CNAF and at two other big INFN computing centers in Padova and Bari. Additional deployments are in the works. "There are no other solutions, in Italy or any other country, that prove to have the same scalability and flexibility in providing shared access to resources – without the need for partitioning – through the use and integration of both Grid and Cloud interfaces," concludes Salomoni.

Find a solution that is right for your organization. Contact your Intel representative or visit the Reference Room at www.intel.com/itcasestudies



Copyright © 2011, Intel Corporation. All rights reserved. Intel, the Intel logo and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries.

This document and the information given are for the convenience of Intel's customer base and are provided "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel® products are not intended for use in medical, lifesaving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

*Other names and brands may be claimed as the property of others.

0911/JNW/RLC/XX/PDF

326131-001EN