HPC Cloud Computing

Self Service and Dynamically Scalable High Performance Computing

sara

BiG Grid
the dutch e-science grid

Cloud Computing brings fast compute clusters within the reach of scientists for whom other HPC facilities are not an option. Within “the Cloud” scientists can use a computer environment that is virtually identical to the environment that they have developed for their own internal use – but one which is many times faster. With the BiG Grid HPC Cloud, users get access to Self Service and Dynamically Scalable High Performance Computing Facilities.

Virtual Private HPC Cluster

With the newly developed High Performance Computing (HPC) Cloud environment researchers get access to their own Virtual Private HPC Cluster. It is a virtualized HPC Cluster that users can configure to exactly match their needs, without interfering with the needs of other users. It is flexible, offers self service and is dynamically scalable.

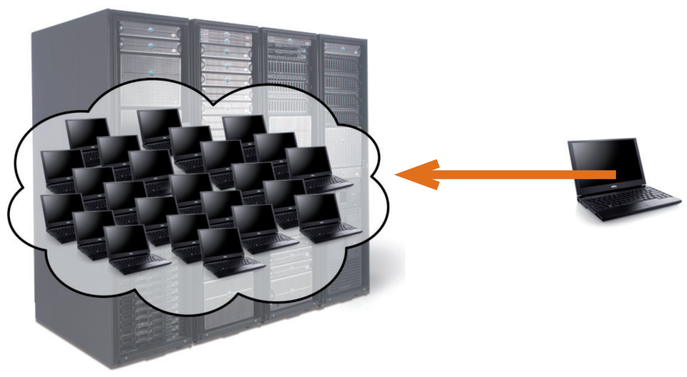

Users can start from existing templates (images), or built their own cluster from scratch. It is even possible to make a copy from their current IT software environment (for example their laptop or desktop pc) and turn that into a HPC cluster in our Cloud. In that way, there will hardly be any difference between their development environment and their production environment. There is no need for an (expensive) rewrite of their software, and scientific challenges can be scaled up very easily from desktop scale to High Performance Compute cluster scale.

Flexibility

The provided flexibility enables new modes of high performance computing. Usage patterns that were hard to realize previously, because of legacy software, security concerns, or source code unavailability, can now be accommodated with a customized cloud solution, tailored to the specific purpose of the scientific end-user. In short, a virtual private HPC cluster gives academic research communities the flexibility they need.

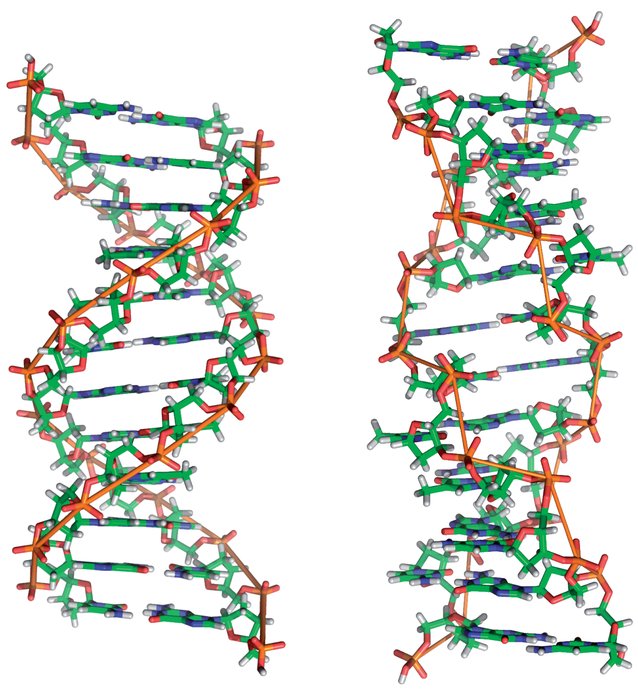

User Participation

SARA has developed the HPC Cloud service for BiG Grid in close collaboration with a number of scientific users, demonstrating real test cases. These test cases, in the areas of geo-ecology, bioinformatics, text processing and named entity recognition, clearly showed the added value in a large variety of scientific research in terms of flexibility, ease of use and productivity.

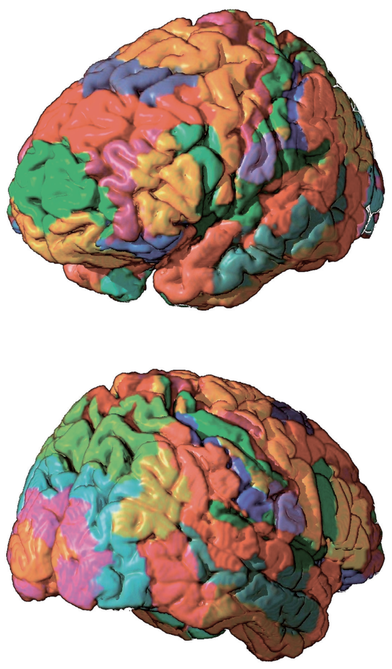

User Experience

One of the participants is Lourens Veen, scientific programmer at the Computational Geo-Ecology group (CGE) at the University of Amsterdam: *“We focused on inverse modeling. With inverse modeling you try to ascertain the optimum settings for a model. We were working with a hydrologic model for identifying precipitation, evaporation and drainage. Our model describes how rainfall influences flow rates. A number of factors play a role here, such as the porosity of the soil. But there are other factors that you can’t always ascertain. In this model, we are trying to discover those hidden parameters. At the same time we are investigating whether the model is tenable.”*

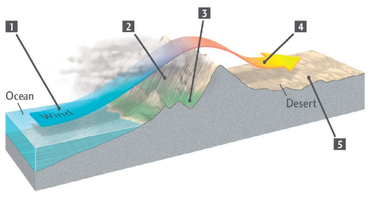

Benefits of customization

Lourens Veen sees huge benefits in Cloud Computing for his team: *“We have all kinds of scientists here who aren’t computational scientists, such as hydrologists, physical geographers, ecologists and ornithologists. Owing to the huge quantities of data and the complex computations, these researchers need to use HPC facilities, but most of these facilities are difficult for these scientists to use. The advantage of Cloud Computing is that we as scientific programmers have full control of the environment. We can build an environment ourselves that our scientists can work with. This allows us to disguise the actual complexity of the system as far as the user is concerned.”*

Graphical User Interface

The cloud environment is accessible through a Management Console, a graphical user interface that end-users can use to deploy and interact with their own Virtual Private HPC Cluster.

A user only needs a web browser with a recent version of Java enabled to access all functionality of the HPC cloud.

Topology

Virtual Machines

- Private network (vlan), or public ip with self service firewall
- Administrative rights inside VM and on private network
- Any Operating System: Linux, Windows, etc
- Any software, including database and webservers

Host Software

- OpenNebula
- Hypervisor: KVM
- Multicore/multiprocess, MPI and OpenMP

SARA Computing and Networking Services, Science Park 121, 1098 XG Amsterdam, tel. +31 (20) 592 3000, www.cloud.sara.nl, e-mail: cloud-support@sara.nl