

Part of the vertex of an Expressed Sequence Tag (EST), taken from the visualisation of EST clusters showing the relationship between the different fragments.
Acknowledgement: Dane Kennedy, University of the Witwatersrand.



Contact details

- General enquiries: info@meraka.org.za
- Telephone number: +27(0)12 841-3028
- Technical enquiries: support@c4.csir.co.za
- Mailing list subscription page: <http://lists.meraka.csir.co.za/mailman/listinfo/c4-users>
- C4 website: <http://c4.csir.co.za>
- Meraka HPC web page: <http://www.meraka.org.za/HPC>

CSIR Cluster Computing Centre
Shared and Open HPC Infrastructure
c4.csir.co.za | www.csir.co.za

CSIR
our future through science

What is HPC?

High performance computing (HPC) can be thought of as the Formula 1 of the computing world. From a technology perspective, it is the specialist domain where cutting-edge computing equipment and associated software components – predominantly open source in nature – come together to push the boundaries of what can be achieved in terms of processing, networking, storage and visualisation. From an application perspective, HPC enables researchers to probe the unknown areas of their domains by means of modelling and simulation at levels of complexity and resolutions, which far exceed that which is attainable by conventional techniques.


Background

In the middle of 2006, the CSIR invested in the establishment of a new HPC platform to be used as a research tool into the technologies of HPC itself, as well as an enabler of computational science in the organisation and beyond. The platform has subsequently grown through various additions and enhancements made possible through the efforts of several research groups and projects, especially that of the Meraka Institute at the CSIR.

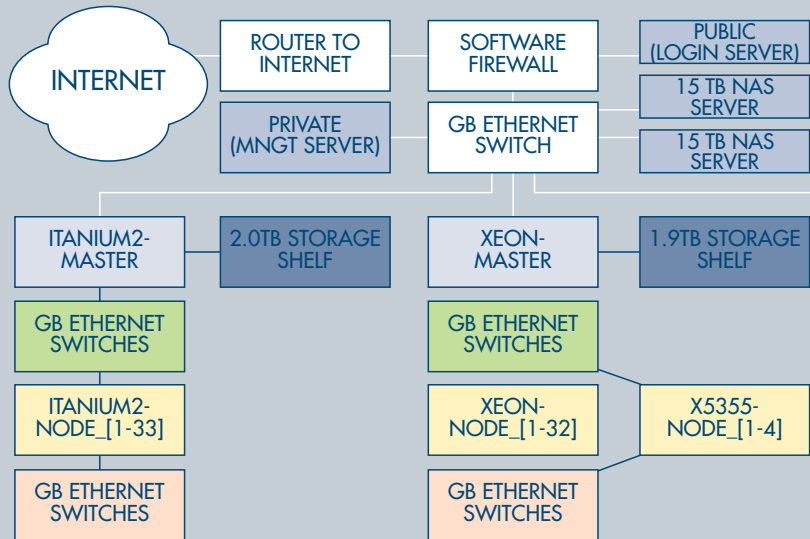
C4 is among the most substantial organisational HPC resources in South Africa, and offers computing capacity to the wider scientific community in addition to its core users in the CSIR. There are currently in the order of 100 registered scientific users, a number that is constantly growing. A significant proportion depend directly on C4 for the advancement of their various research efforts.

Getting Started

- The resources of C4 include three computing clusters, named opteron, xeon and itanium2 (April 2008)
- Usage of C4 is at no cost to scientific users
- Users are categorised as primary or secondary, depending on alignment with the mandate of the particular resource
- C4 is a shared resource and employs flexible rules regarding resource allocation
- An *ad hoc* peer-review mechanism for policy modification underpins C4's 'infrastructure commons' philosophy



This pyramid illustrates the hierarchy of HPC facilities, with C4 being a good example of an organisational class facility that is typified by moderate scale and high levels of user diversity.

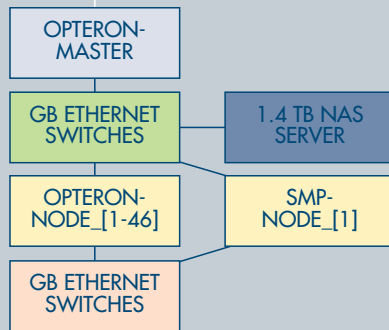


- 1 headnode providing 4 x single core Itanium®2 1.6 GHz and 16GB of RAM
- 33 nodes providing 2 x single core Itanium®2 1.6GHz and 4GB of RAM

This diagram shows the various components of C4 and their relative positions in the high-level network architecture.

- 1 headnode providing 4 x single core Intel® Xeon® 3.6 GHz (model Irwindale) and 16GB of RAM (xeon)
- 31 nodes providing 2 x single core Intel®2 Xeon® 3.6 GHz (model Irwindale) and 4GB of RAM (xeon-node_1-xeon-node_32)
- 1 node providing 4 x single core Intel® Xeon® 3.6 GHz (model Irwindale) and 4GB of RAM (xeon-node_27)
- 4 nodes providing 2 x quad core Intel® Xeon® 2.6 GHz (model Cloverton) and 2GB of RAM (x5355-node_1 - x5355-node_4)

C4 at a glance



- 1 headnode providing 2 x dual core AMD® Operton® 2.6 GHz (model 285) and 4GB of RAM (operton)
- 46 nodes providing 2 x dual core AMD® Operton® 2.6 GHz (model 285) and 4GB of RAM (operton-node_1 - operon-node_46)
- 1 node providing 4 x dual core AMD® Operton® 2.6 GHz (model 885) and 64GB of RAM (smp-node_1)

Getting a user account

Subscribe to c4-users mailing list (see URL on back page)

The sys admin sends an email with account details and login instructions

Use ssh or PuTTY (on Windows) to log into C4 at c4.csir.co.za

Once logged in the first time, use 'passwd' to change your password

Questions?

Read the README file located in your home directory

Search for a solution in the online archives of the c4-users mailing list

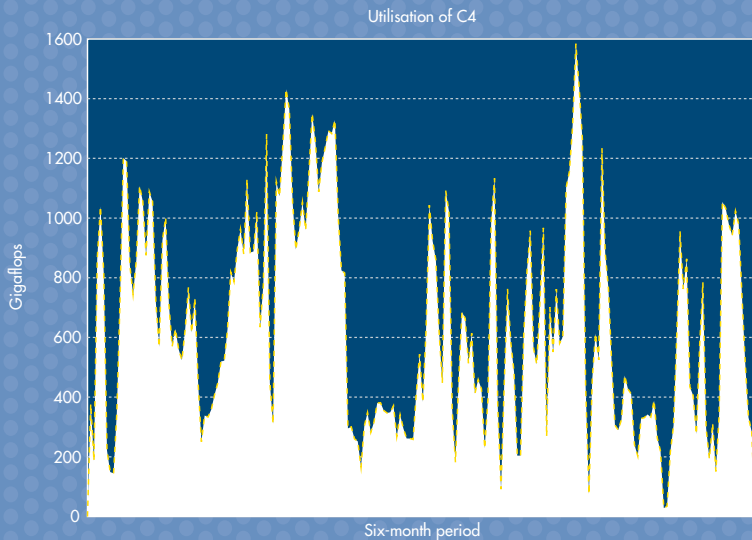
Search for a solution on the web

Post a question to the c4-users mailing list

Send an email to support@c4.csir.co.za

Application examples

There are several examples of application areas that benefit from HPC in general, and also many of researchers that utilise the resources of the C4 in particular. These range from the modelling of DNA substitution in HIV, to experiments in artificial intelligence, the simulation of stresses and strains below road surfaces, the development of improved radar processing techniques, refinement of data compression algorithms, and beyond.



BioMedical Informatics Centre

There are a number of definitions for biomedical informatics, which converge around the interface between disciplines such as traditional bioinformatics, structural biology and population dynamics. Its exclusive objective is the understanding and solution of health threats by employing the techniques of modern informatics.

The BioMedical Informatics Centre (BMIC) is located at the Pretoria offices of the Meraka Institute. It is an HPC service to the national biomedical informatics community, and was largely made possible through donations made by Intel. The BMIC was launched by the chairman of Intel, Dr Craig Barrett, on 11 December 2006.



Albert Gazendam (right): HPC technical leader at the Meraka Institute, with Dr Barrett (second from right) and representatives from Intel Corporation after the BMIC launch