

Sleepy Carnarvon to host supercomputer?

ITWeb (www.itweb.co.za)

24 April 2008

Leon Engelbrecht

Carnarvon, a small town of about 10 000 people in the Northern Cape, could be home to SA's next centre for high-performance computing, says Square Kilometre Array (SKA) project leader Dr Bernie Fanaroff.

The SKA is a proposed 3 000 to 5 000 dish radio telescope that is to be built either in SA or Australia from 2013.

The cost, to be funded by SA and a consortium of other nations, is now estimated at 1.5 billion euro (R18 billion), up from an earlier projection of one billion euro (R12 billion).

Construction on a SKA "pathfinder", MeerKAT, gets under way this October.

Fanaroff says the initial hope was to place MeerKAT and SKA in a remote location, away from radio frequency interference, and pipe the data to Cape Town for processing. But that is not proving cost-effective.

"We are going to do most of the data-crunching on site because it will be too expensive to pipe the raw data down to Cape Town," he says. "So we will have our science and control centre in Cape Town, but the data-crunching will take place on site."

Our own CERN

MeerKAT project leader Anita Loots says the SKA will have a data requirement similar to the Large Hadron Collider experiment at the European Organisation for Nuclear Research, better known as CERN.

Loots says the core of the 3 000 to 5 000 dish SKA alone will require two computers capable of handling 10 petaflops each.

"As far as my knowledge goes, no proven petaflops machines exist in the world today, although the Japanese say they want to build one. So there's still a lot to be done just to build the machines [we need to make SKA work]."

This data volume also dictates that the SKA have its own high-performance computing centre. There was initial talk of it using the Meraka Institute's Cape-based Centre for High Performance Computing, but "we will flood their systems, their facility is slightly too small for us," Loots says.

Fanaroff says SKA and MeerKAT will have similar computing architectures. "There are two different computing centres. The first is a correlator. The correlator multiplies the signal of each dish against all the others by producing an inverse transform that produces a picture of the sky.

"But then you have to calibrate that picture, you have to modify the picture, clean it up, take out the noise and that's high-performance computing. We were originally planning to have the correlator on site and the high performance computing centre in Cape Town, but it is too expensive to send the raw data; so what we are now going to do is 'slice-and-dice' the data on site until we are down to the calibrated clean picture that we will then send down on a 10, 20GB line to Cape Town."

Fanaroff says this may be increased to 100GB, "later even more" depending on the cost of the associated infrastructure, which is being supplied by Infracore and Neotel.

"In Cape Town, we will have the science centre, the control centre, so the whole computer on site and the telescope will be controlled from Cape Town," Fanaroff continues. "So you will have telemetry going up and monitoring data as well as telescope data come down. The scientists will get the clean calibrated data to do analysis on."