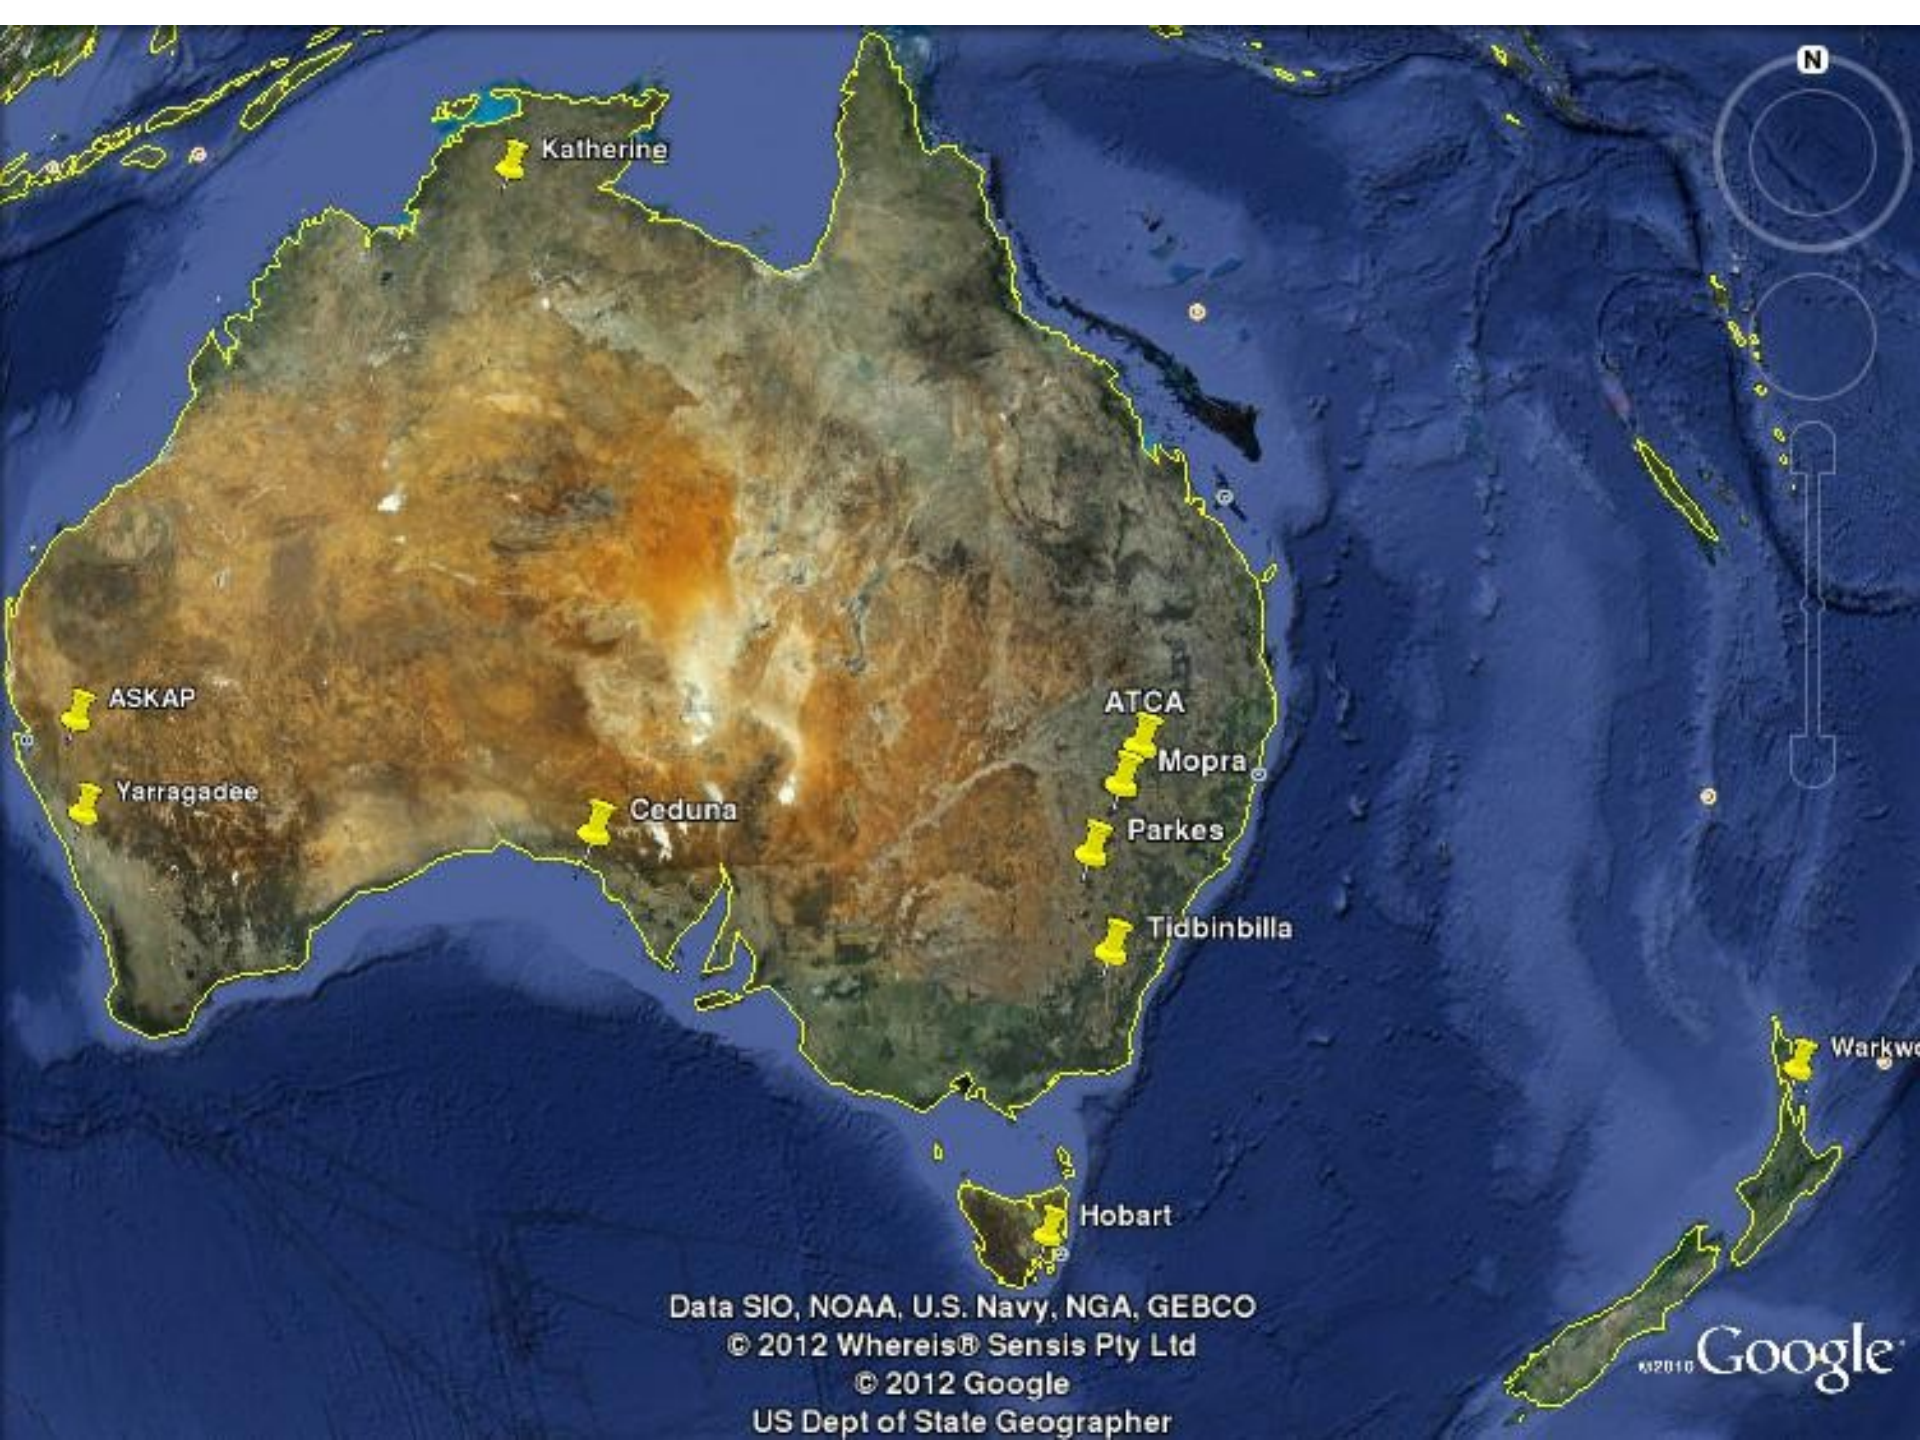


ICRAR-Curtin Correlator Facility

- Primary correlator for the Australian Long Baseline Array
 - 4 ATNF Antennas (Parkes, ATCA, Mopra, ASKAP)
 - 2 UTAS antennas (Hobart, Ceduna)
 - 3 AuScope Antennas (Yarragadee, Hobart, Katherine - UTAS)
 - Warkworth antenna (AUT, NZ)
 - +Tidbinbilla, Hartebeesthoek, Tigo, O'Higgins, Shanghai...
- CUPPA
 - 20 node (160 core) beowulf cluster
 - Gb ethernet
 - >150 TB attached storage
 - Runs DiFX under Espresso





Katherine

ASKAP

Yarragadee

Ceduna

ATCA

Mopra

Parkes

Tidbinbilla

Hobart

Warkwa

Data SIO, NOAA, U.S. Navy, NGA, GEBCO

© 2012 Whereis® Sensis Pty Ltd

© 2012 Google

US Dept of State Geographer

©2010 Google



International
Centre for
Radio
Astronomy
Research

Real-time e-VLBI

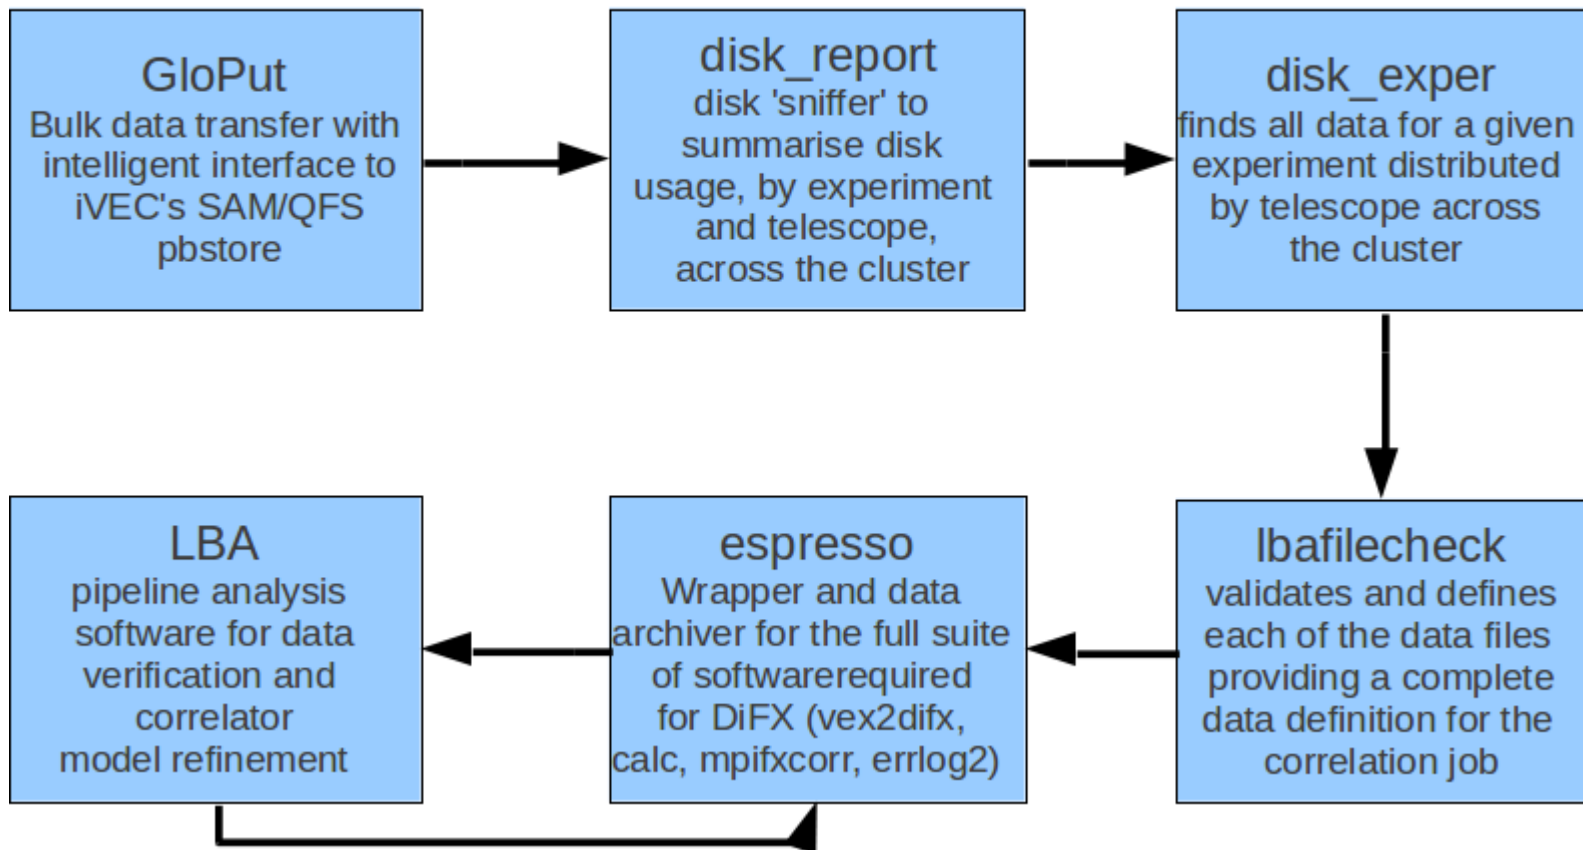


ICRAR-Curtin Correlator Facility

- Data transferred electronically from telescopes (~400 TB per year).
 - Stored on shared facility (iVEC's pbstore)
 - Hybrid disk (~150 TB) and tape (>1 PB) system, SAM/QFS
 - Transferred to CUPPA 'just in time' for correlation
- Approximately 40 separate experiments per year
- All data permanently online – facility is completely remotely operable

ICRAR-Curtin Correlator

The Espresso Processing Pipeline



Espresso

- Distributed with DiFX
- Lightweight, with simple configuration
 - Requires a cluster definition file, a few environment variables and data directory names should follow a simple scheme (project-telescope)
 - Targets file-based correlation. Command line driven..

```
File Edit Tools Syntax Buffers Window Help
node, number compute threads, space separated list of data areas.
cuppa01, 6, /exports/xraid01/l_1 /exports/xraid01/r_1
cuppa02, 6, /exports/xraid02/l_1 /exports/xraid02/r_1
cuppa03, 6, /exports/xraid03/l_1 /exports/xraid03/r_1 /arch/corr/bbdata
cuppa04, 6, /exports/xraid04/l_1 /exports/xraid04/r_1 /arch/corr/bbdata
cuppa05, 6, /exports/xraid05/l_1 /exports/xraid05/r_1 /arch/corr/bbdata
cuppa06, 6
cuppa07, 6
cuppa08, 6
cuppa09, 6
cuppa10, 6
1,1 Top
```

[[difx:espresso]]

Edit this page

Old revisions

Export to ODT

Recent changes

Search

Trace: • start • espresso

You are here: [start](#) » [difx](#) » [espresso](#)

Espresso

Espresso is the system developed for managing correlation at Curtin University. It is a lightweight system for managing data on your cluster, automating the correlation process, and providing simple archiving of the outputs. It is designed for correlation from standard linux disks (not direct from Mark5s). Espresso also provides a number of auxiliary scripts which may come in handy during correlation. A typical espresso session, as it is used at Curtin, is available [here](#).

All the scripts will give help if invoked with the `-h` switch.

Table of Contents

- Espresso
 - Installing espresso
 - Running the Espresso Scripts
 - disk_report.py
 - disk_exper.py
 - lbafilecheck.py
 - espresso.py
 - Auxiliary Tools
 - Some Notes on Espresso

Edit

Installing espresso

The scripts come with your DIFX installation (2.0.2 and later), in `$DIFXROOT/applications/espresso`. The included `install.py` script should install them in your DIFX bin directory. To work they need a correlator definition file. See the `corr_hosts.txt` file in `$DIFXROOT/applications/espresso` for an example. For every node in your cluster you need to enter the hostname, maximum number of MPI threads to be run simultaneously on that node (typically \leq number of available cores), and a space separated list of any data areas (directories) on that node where baseband data may be stored.

The environment variable `$CORR_HOSTS` should point at your version of the cluster definition file (`corr_hosts.txt` described above). As this correlator definition file is unrelated to the particular version of DIFX you are using, you probably want to store it in your home directory, or similar.

Espresso allows you to write the output data to a directory other than the one in which the correlation files are stored (this is useful for installations where the NFS disks are too small to store the output data). You should set the environment variable `$CORR_DATA` to point to the