

Managing multi-PB data: Perspectives from Earth systems

- Claire Trenham, NCI

A. Hotan (CASS), C. Allen, K. Druken, A. Steer, B. Evans, J. Smillie,
J. Wang and L. Wyborn (NCI)

- ASKAP Early Science is happening!
 - *But how to store and share the data so other scientists can make use of CASDA?*
- Estimated archive volume is 5PB/yr
(http://www.atnf.csiro.au/projects/askap/news_computing_05112015.html)
- That's huge for a single dataset, but not crazy for a data collection. Hooray for post-processing to archival volumes :)

- *“The new wide-field radio telescopes, such as: ASKAP, MWA, and SKA; will produce spectral-imaging data-cubes (SIDC) of unprecedented volume. This requires new approaches to managing and servicing the data to the end-user.”*
 - Kitaeff et al 2012, <http://skuareview.net/wp-content/uploads/2016/05/astro04-kitaeff.pdf>
- *“... At the same time, other research and development communities, such as: remote sensing, geographic information systems, medical imaging, have indeed developed interesting techniques which could solve many problems which radio astronomy is about to face with extremely large size imaging data”*
 - Kitaeff et al 2012, <http://skuareview.net/wp-content/uploads/2016/05/astro04-kitaeff.pdf>

- NCI – National Computational Infrastructure

- Highly integrated peak machine

- Raijin: 1.2PFlops, >57k cores, Infiniband

- data store


- >30PB disk, ~10PB tape, 56Gb FDR Infiniband & 10GigE

- research clouds

- NeCTAR public cloud; Tenjin private cloud with Virtual Labs and access to 10+PB National Research Data Collection

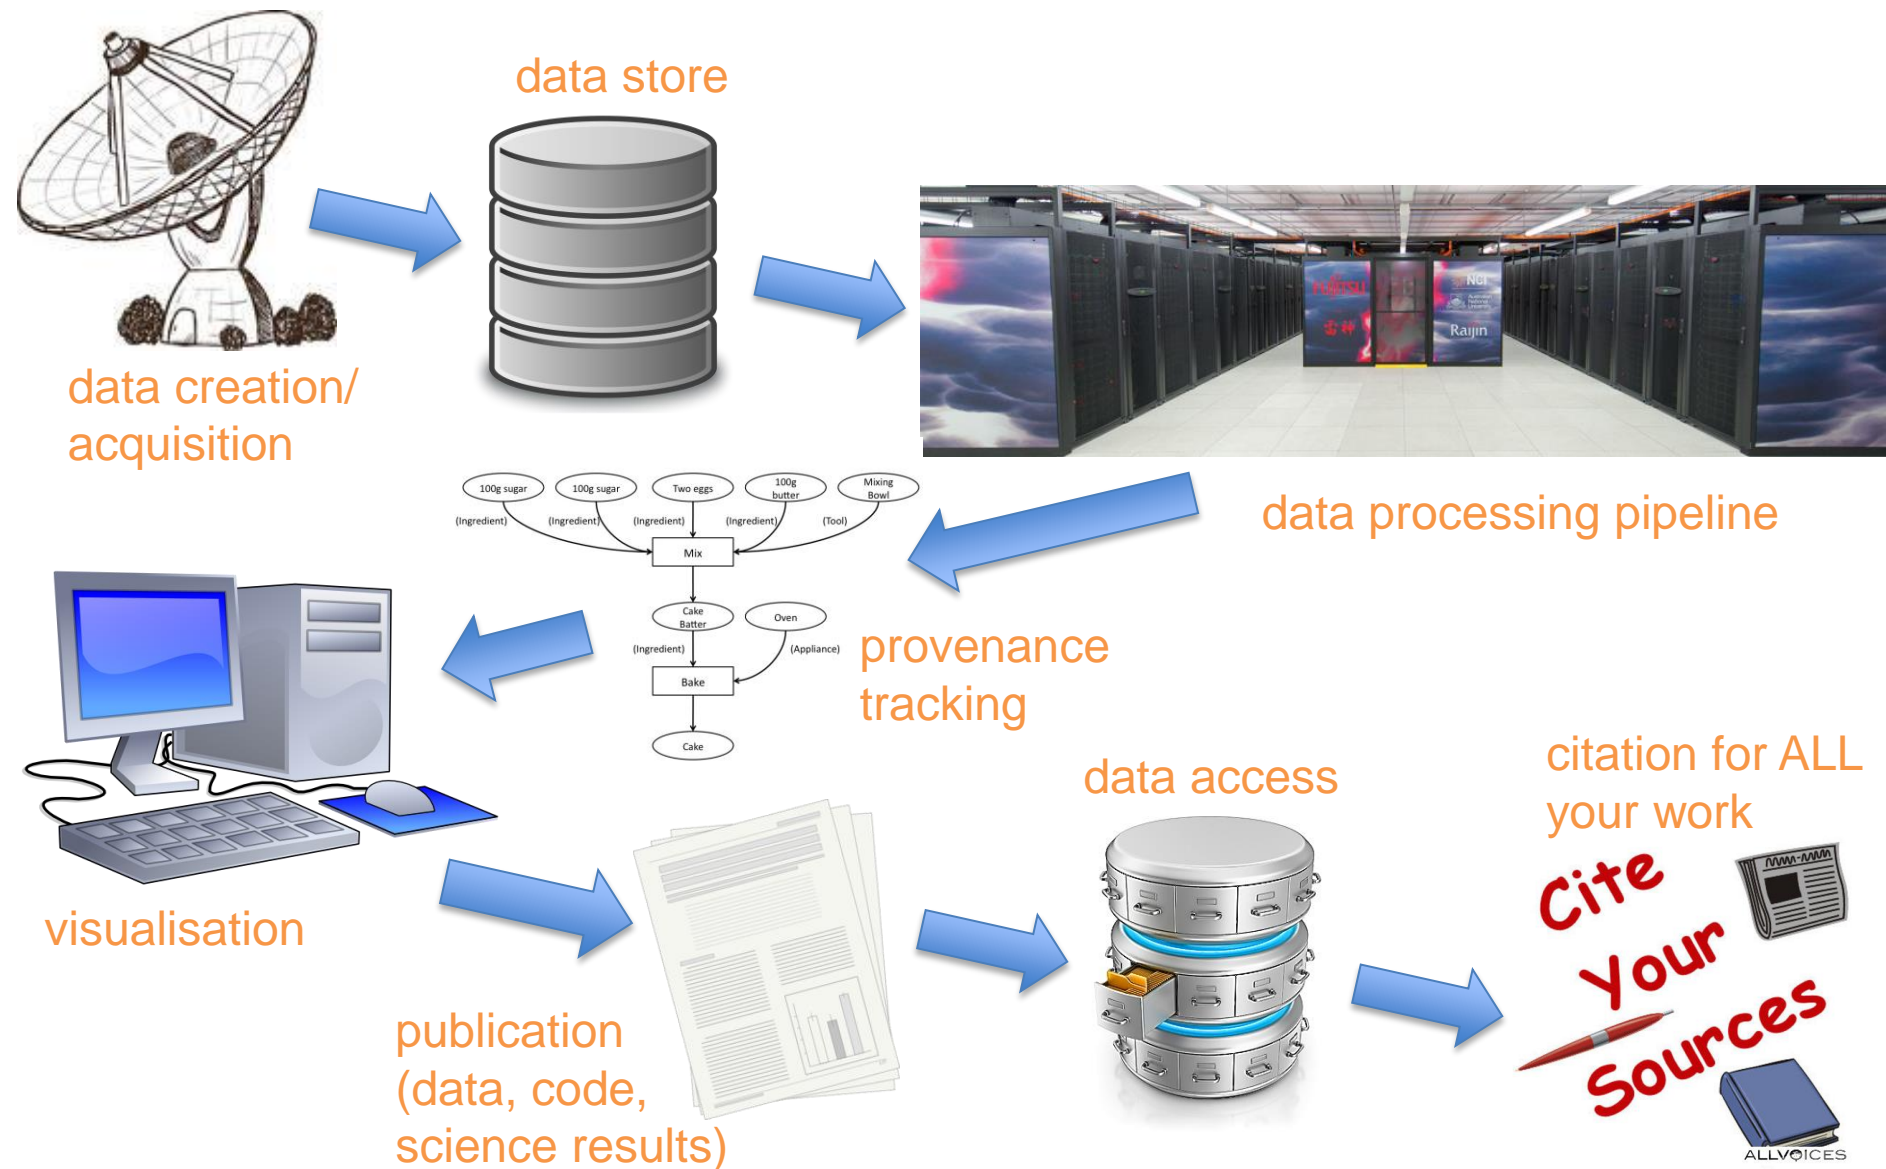
- Services

- Academic consultants provide user support; scientific visualization; virtual laboratories; application optimization



- [RDS\(I\)](#) funding provided to nodes around Australia for the storage of nationally significant data collections.
- NCI focus on the National Environmental Research Data Collection, comprising a range of fields including: climate, weather, Earth observations, ecology & land use, geophysics, geoscience, and astronomy; as well as data holding in social sciences, and bioinformatics.
- Over 10PB ingested and made available to community.
- [Earth Systems Grid Federation](#) primary node (climate models); [Copernicus Hub](#) for ESA data.

- Data is not just a single concept of 1s and 0s
- There are many steps to go from observing an area on the sky to making sources readily searchable and useful data accessible to research astronomers



- Many of these steps are well accounted for or can readily be managed.
- High Performance Computing (HPC) can be leveraged for highly parallel problems to rapidly process data into manageable quantities
 - e.g., producing images from raw visibilities.
- Have an end product (ASKAP Science Data Archive) that requires significant infrastructure to store and analyse.
- Interacting with ASKAP's High Performance Data (HPD) archive involves activities which are often not well suited to an HPC environment, but require access to HPD.



Preparation

- Data management plan, Group access, License and product descriptions, Directory structure, update/back up plan




Data ingest

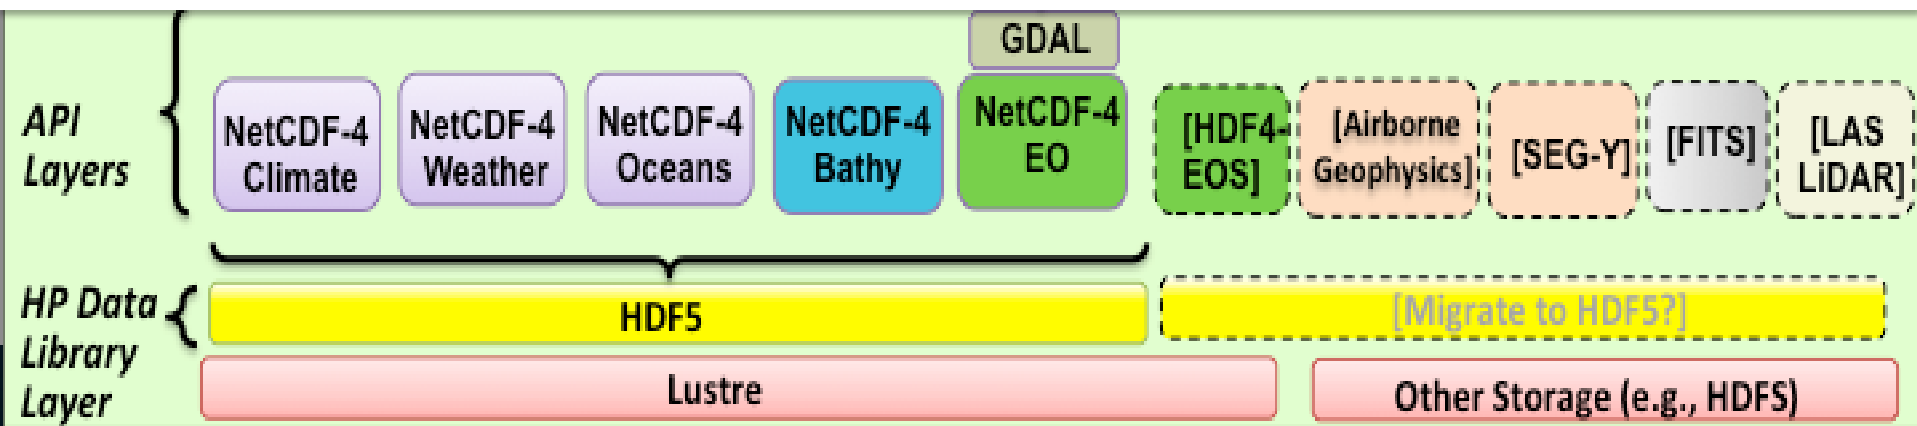
- Data curation and replication, Metadata catalogue

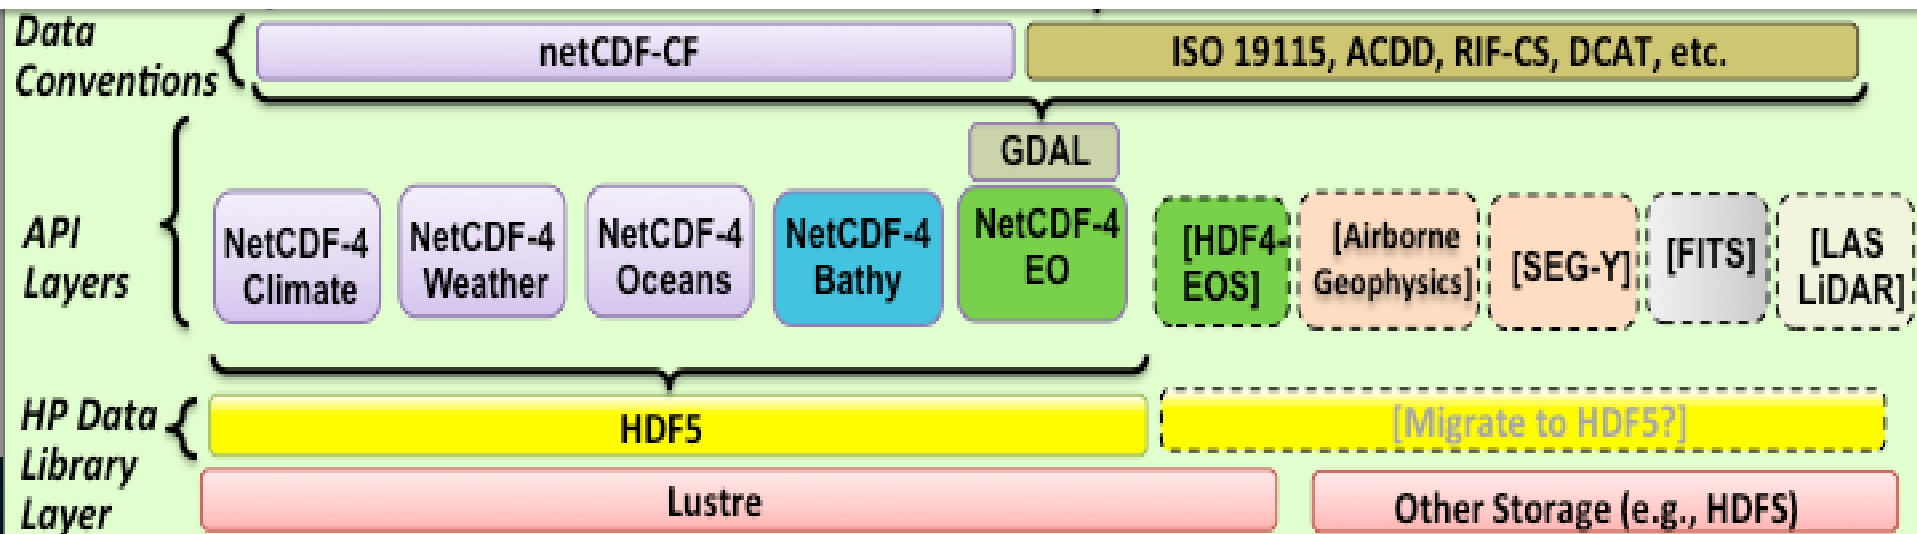


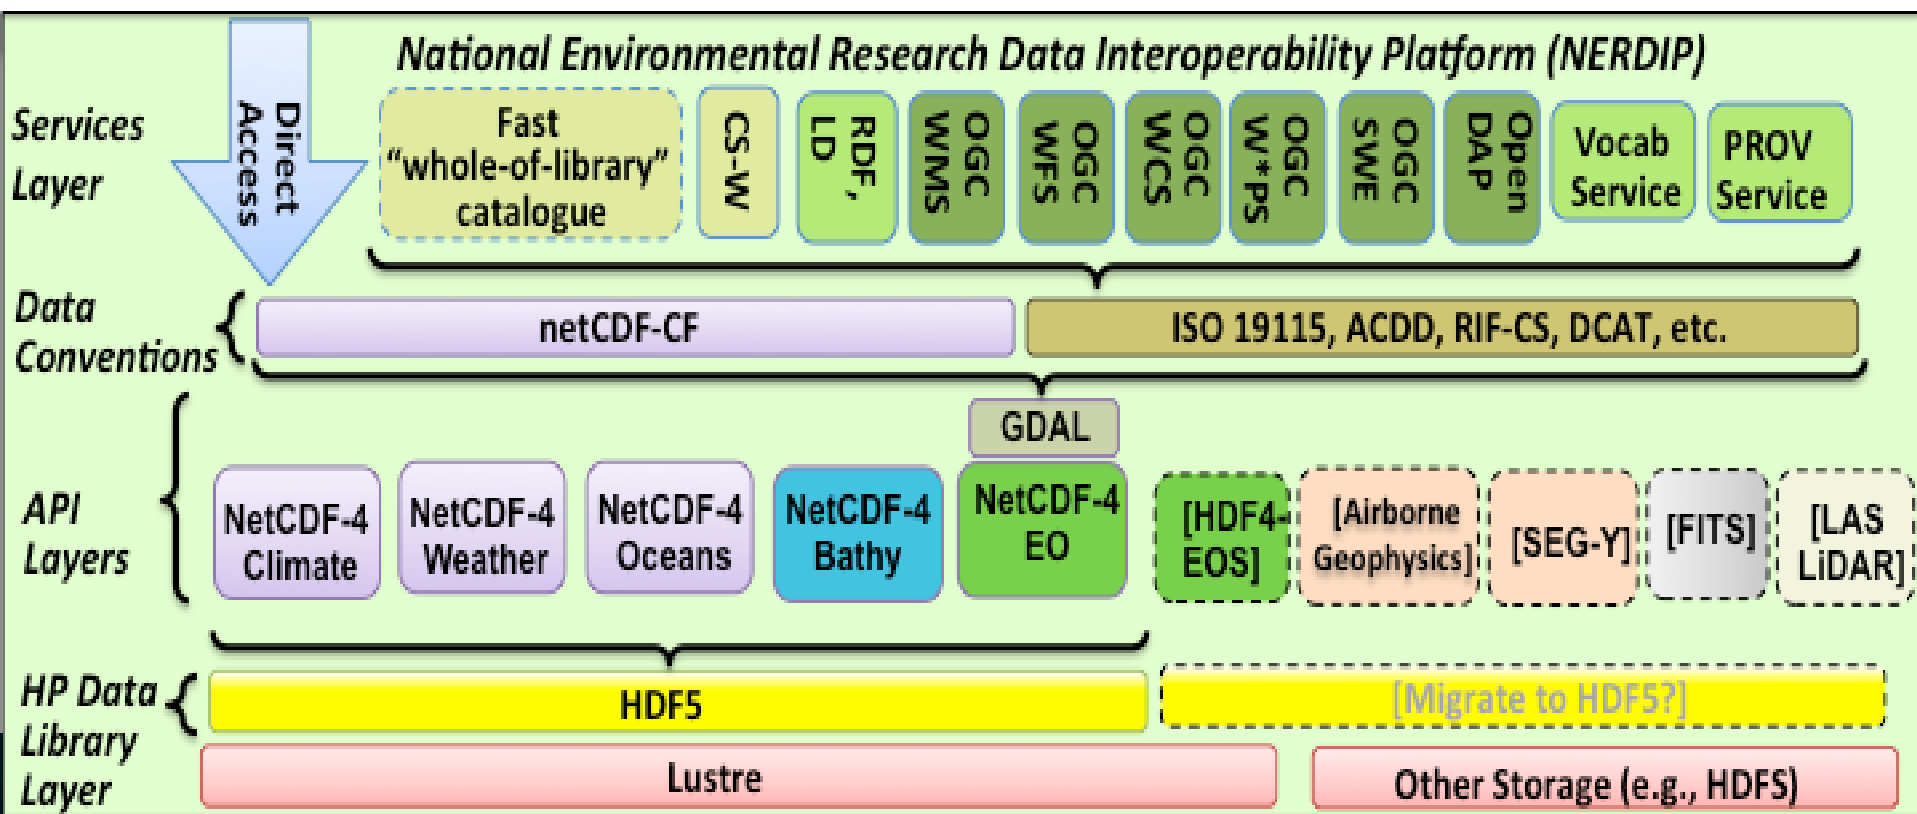
Data
Publishing

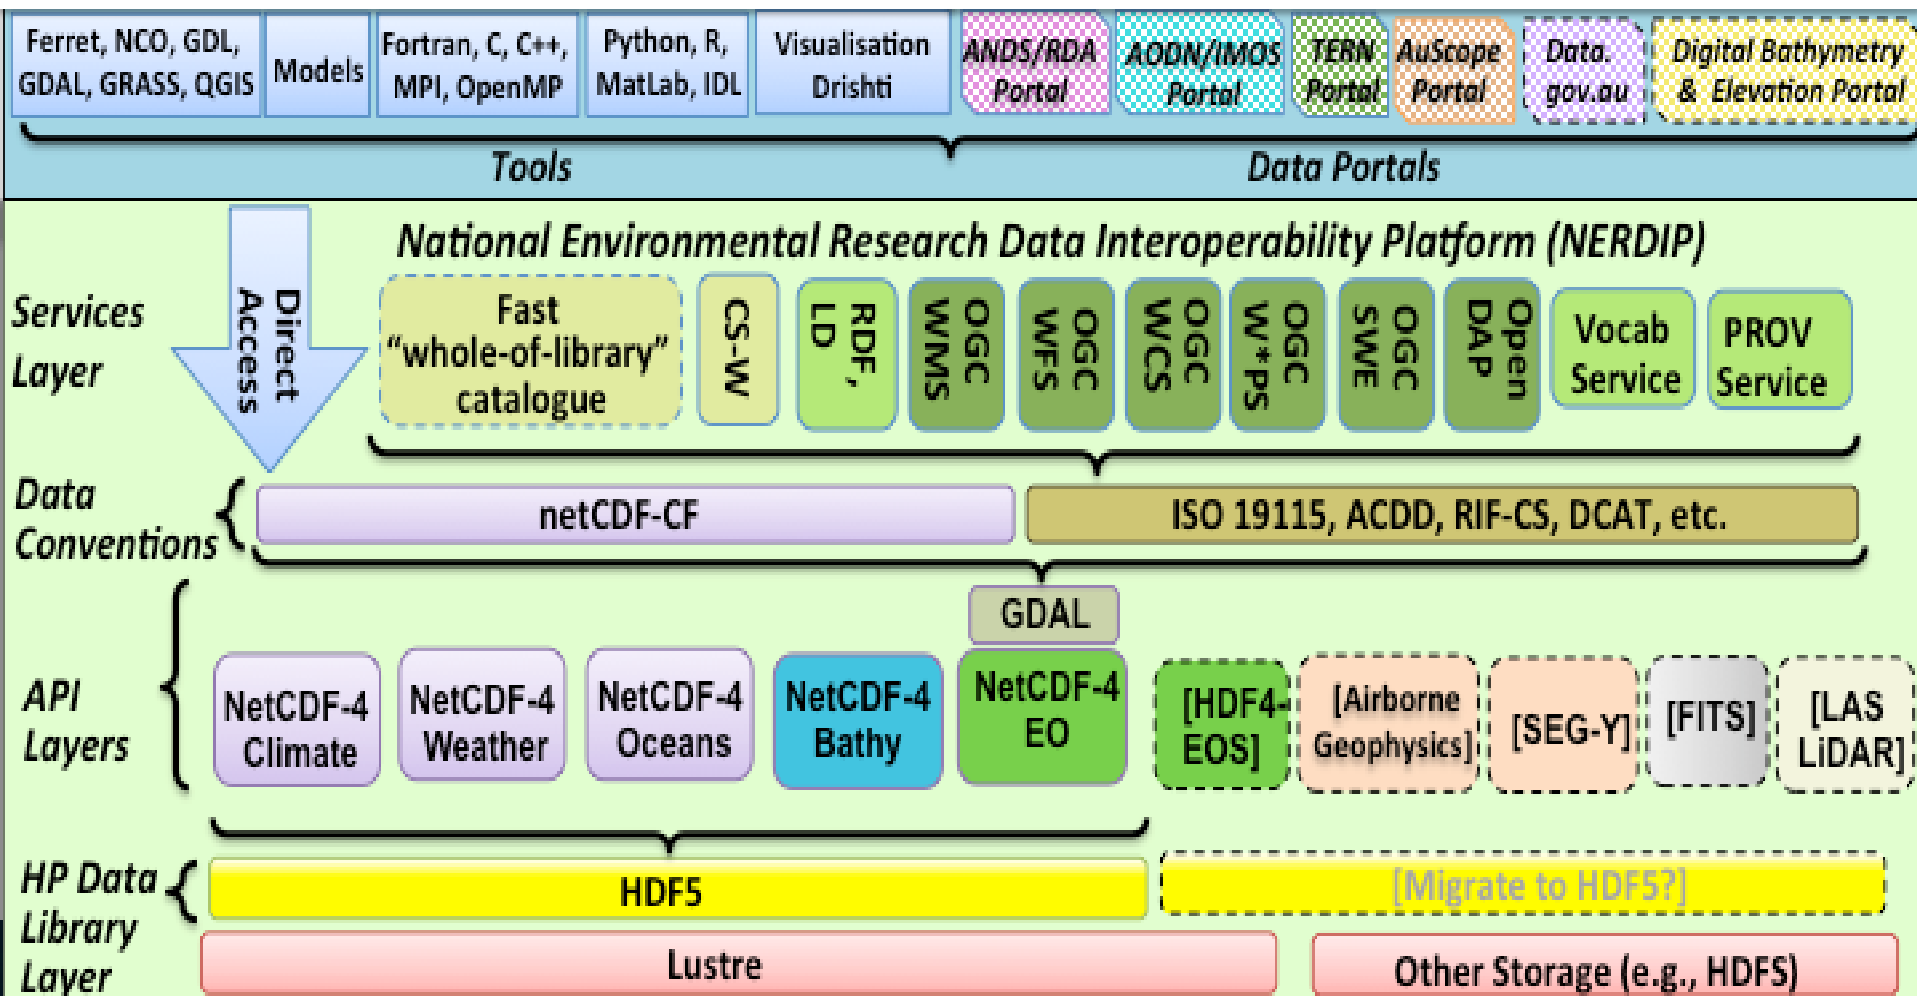
- DOI minting, publish data through THREDDS, or other data services

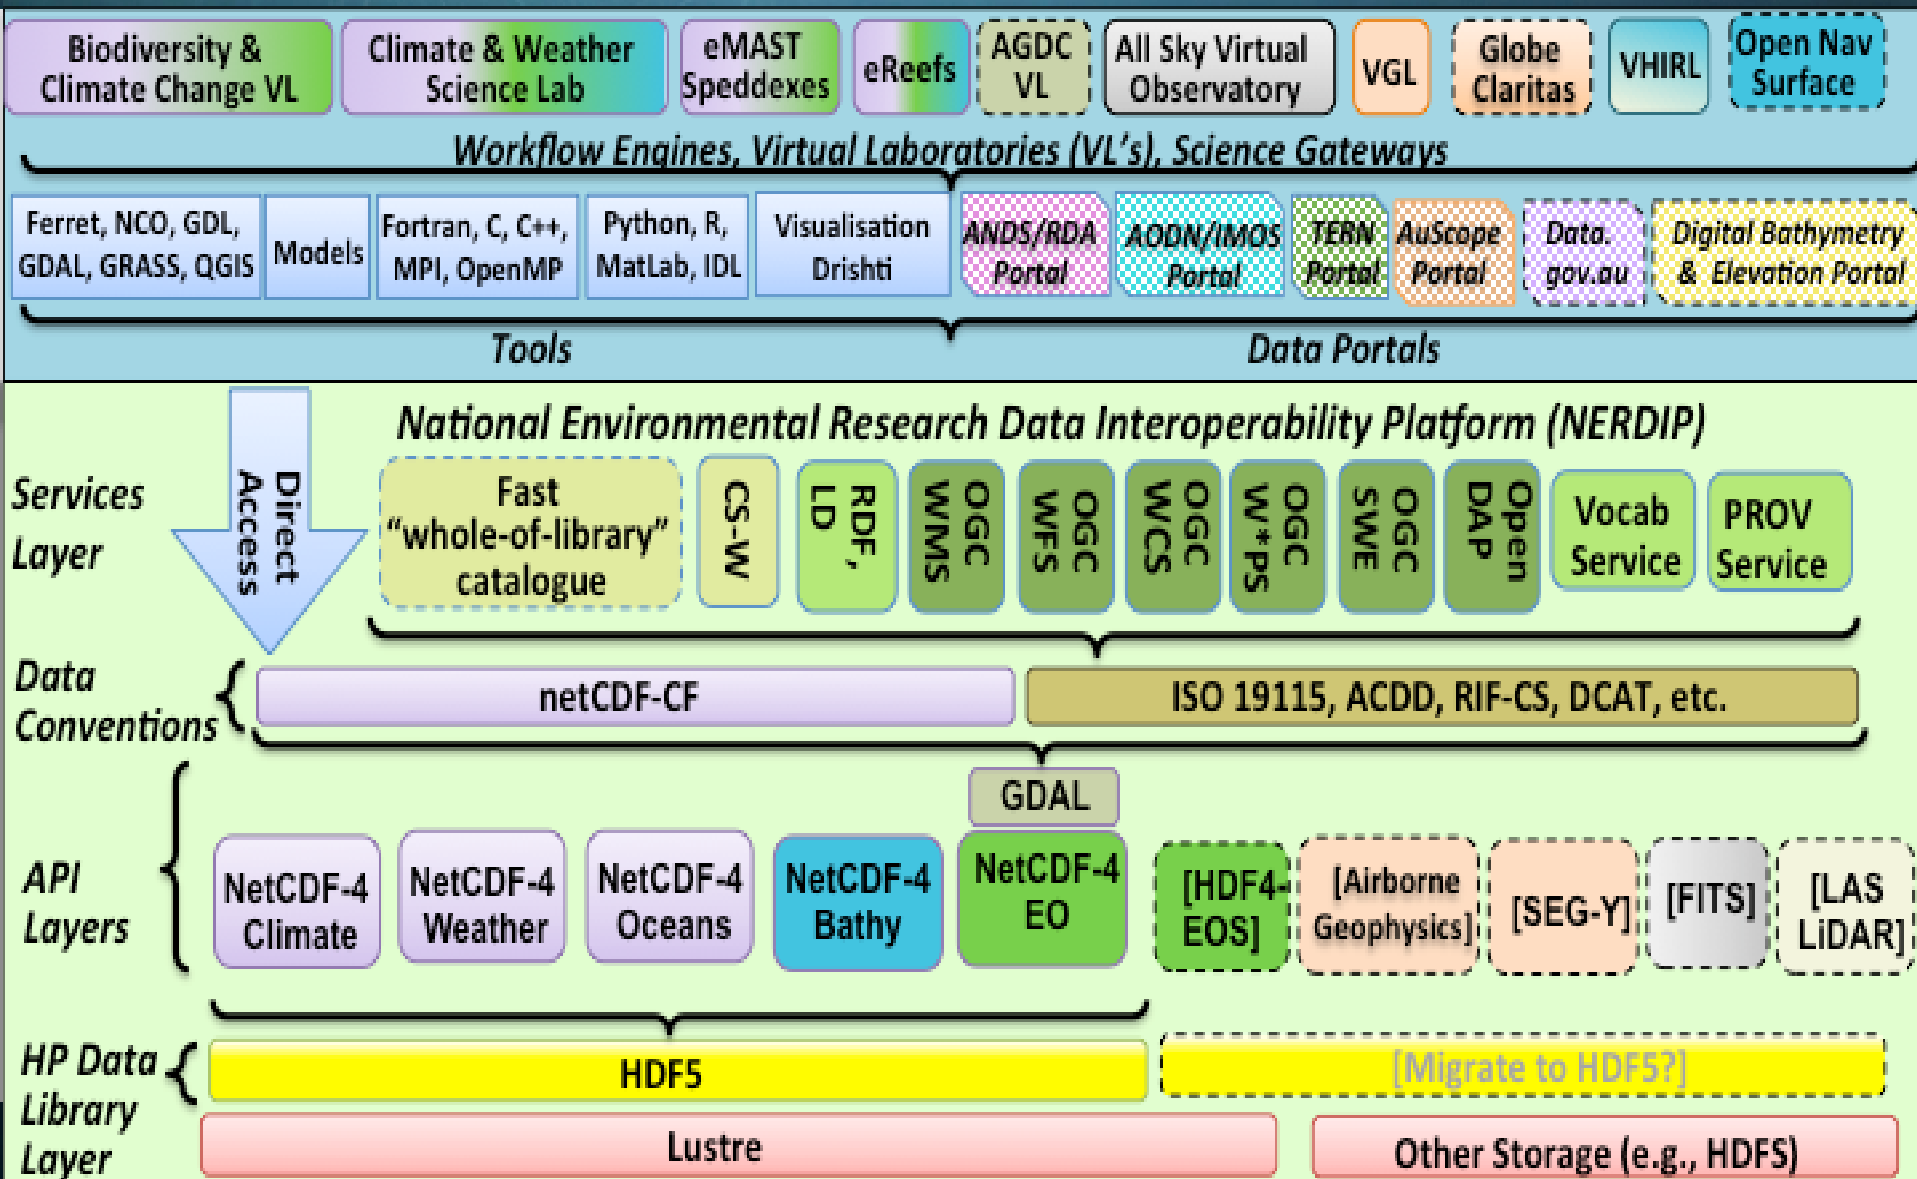






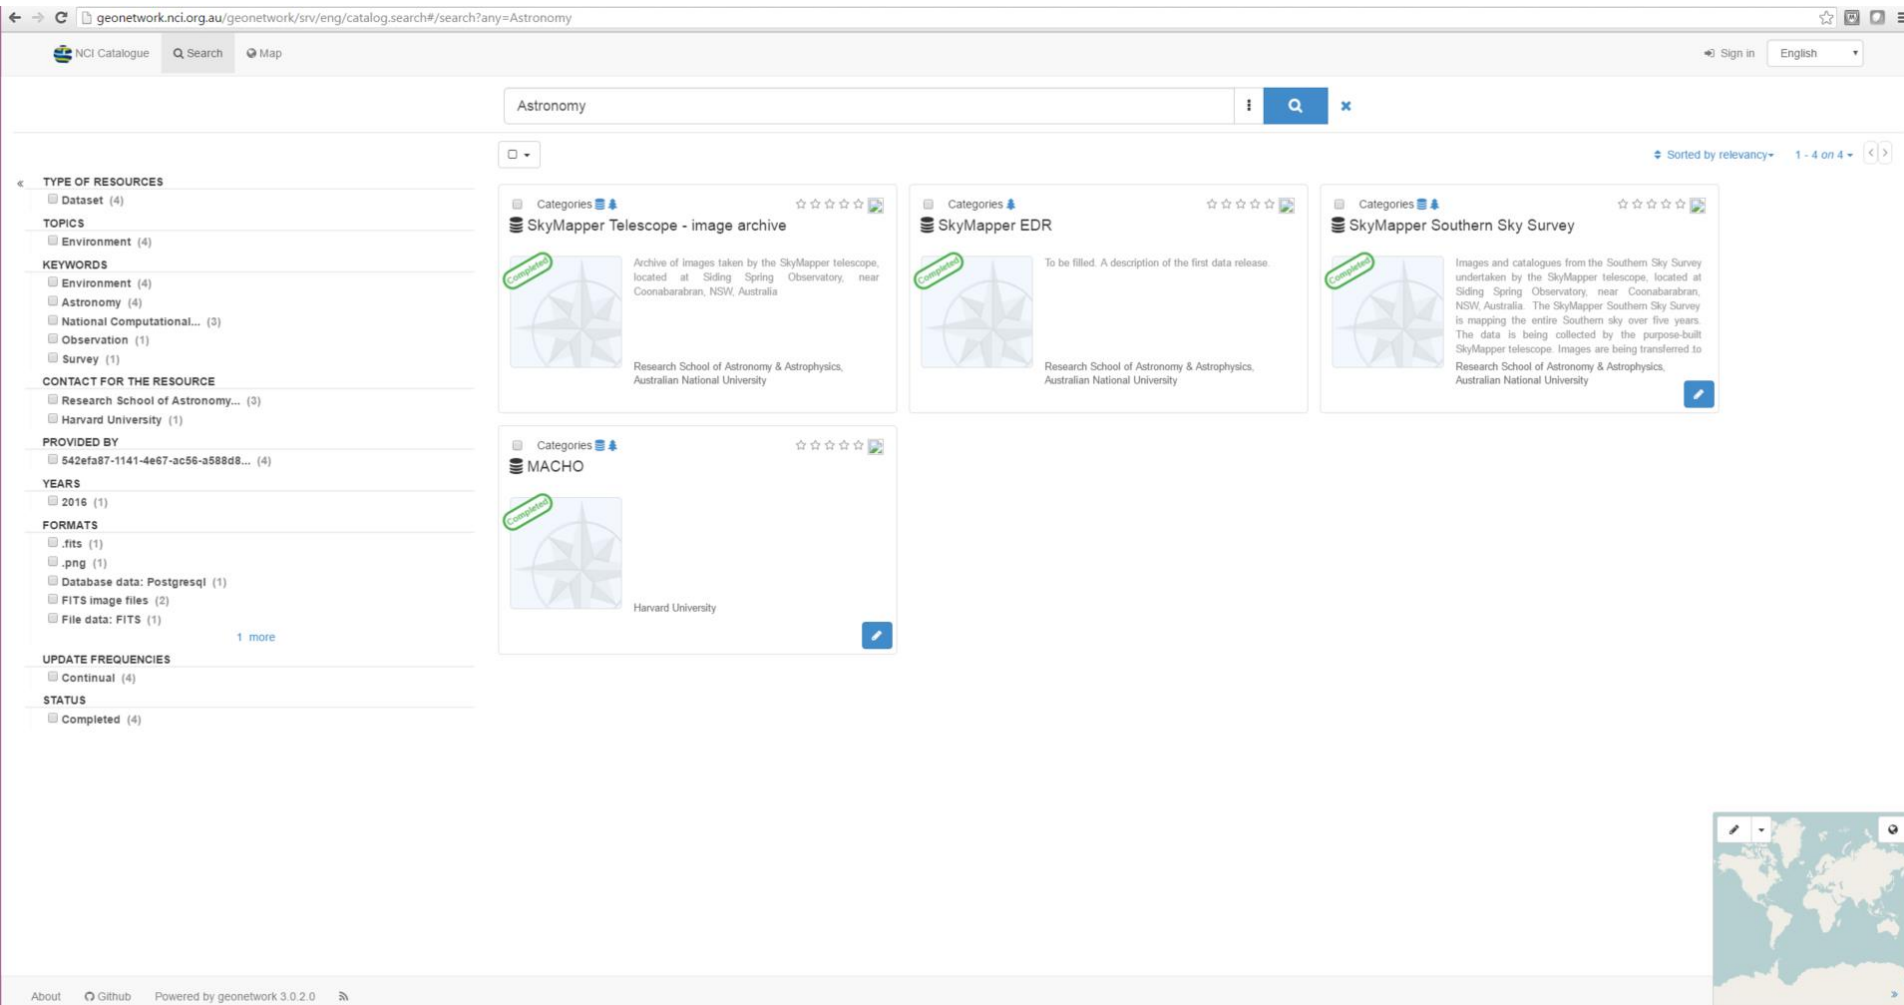






- *Great, so you have lots of data, now what?*
- ASKAP has CASDA
 - <https://confluence.csiro.au/display/CASDA/CASDA+Project+Wiki>
- NCI have a multi-element system for metadata catalogues and data services
 - GeoNetwork: Find metadata records (akin to CSIRO DAP)
 - THREDDS Data Service: download or remotely access or view data
 - Geoserver, ERDDAP, Hyrax, others... and filesystem
 - PROMS (provenance), DOI minting (citation)

<http://geonetwork.nci.org.au/geonetwork>



The screenshot shows a web browser window displaying the geonetwork.nci.org.au website. The search bar contains the word "Astronomy". The left sidebar shows filters for "TYPE OF RESOURCES", "TOPICS", "KEYWORDS", "CONTACT FOR THE RESOURCE", "PROVIDED BY", "YEARS", "FORMATS", "UPDATE FREQUENCIES", and "STATUS". The main content area displays three search results:

- SkyMapper Telescope - image archive**: Archive of images taken by the SkyMapper telescope, located at Siding Spring Observatory, near Coonabarabran, NSW, Australia. Research School of Astronomy & Astrophysics, Australian National University.
- SkyMapper EDR**: To be filled. A description of the first data release. Research School of Astronomy & Astrophysics, Australian National University.
- SkyMapper Southern Sky Survey**: Images and catalogues from the Southern Sky Survey undertaken by the SkyMapper telescope, located at Siding Spring Observatory, near Coonabarabran, NSW, Australia. The SkyMapper Southern Sky Survey is mapping the entire Southern sky over five years. The data is being collected by the purpose-built SkyMapper telescope. Images are being transferred to Research School of Astronomy & Astrophysics, Australian National University.

Below these results is a fourth result:

- MACHO**: Harvard University.

The bottom right corner of the page features a world map. The footer contains the text "About", "GitHub", and "Powered by geonetwork 3.0.2.0".


Geological Survey of Western Australia, 2013 gravity merged 400m grid of Western Australia - test 1


Updated: 29 minutes ago


This grid combines onshore Spherical Cap Bouguer gravity anomalies with offshore Free-air anomalies for Western Australia. Data over the continental region has been gridded from the Australian National Gravity Database (ANGD) and includes GA/GSWA gravity surveys which were acquired during 2013. Data over the marine region is sourced from the Geoscience Australia June 2009 offshore-onshore gravity grid, which uses offshore data derived exclusively from satellite altimetry. The grid is supplied in ERMMapper format. The horizontal datum and projection are GDA94, GEODETTIC. The grid cell size is 0.00416666 degrees (approx 400 metres). The data are in units of micrometres per second squared (μms^{-2}) which is equivalent to 0.1 milligals. The recommended reference for this grid is "Geological Survey of Western Australia, 2013 gravity merged 400m grid of Western Australia".


Completed


Download and links

- 

[OPeNDAP for data sub-setting](#)
OPeNDAP at NCI
Open link
- 

[OPeNDAP for remote data access from local tools](#)
OPeNDAP at NCI
Open link
- 

[netCDF Subset Service to extract file contents by spatial extents](#)
NetCDF Subset Service at NCI
Open link
- 

[Direct file download to local machine](#)
HTTP Download from NCI
Open link
- 

[Web Coverage Service](#)
This dataset is published in the download service (WCS) available at http://dapds00.nci.org.au/thredds/wcs/r11/GSWA_Geophysics/WA_Gravity_Grids/WA_400m_Grav_Merge_v1_2016.nc with layer name WCS Service at NCI.

This is neat

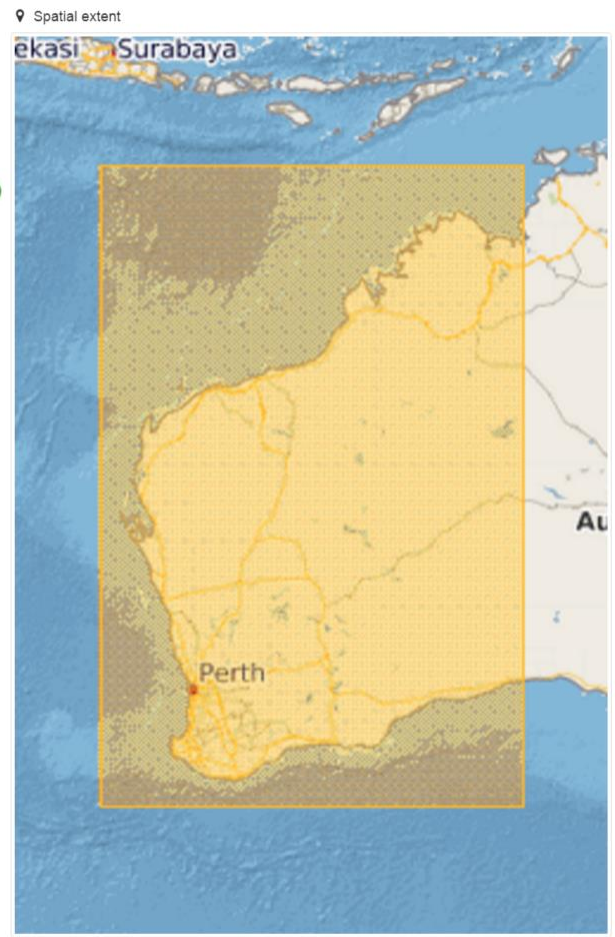
Associated resources

- 

[Geological Survey of Western Australia Geophysics Data Collection](#)
Geophysical datasets published by GSWA and available via the Australian National Virtual Geophysical Laboratory portal presently at...
Parent record

About this resource

Categories	
Keywords	• Environment
Language	• English
Resource Identifier	• 111



- OPeNDAP – Network Data Access Protocol
 - Subset HDF5/netCDF4 data
 - only bring little bits of the data to you instead of downloading whole file
- <http://dap.nci.org.au> THREDDS server
 - OPeNDAP is one of the protocols served, permits subsetting and remote access to files
 - Other protocols include HTTP download and Open Geospatial Consortium Web Services to stream JPEG, TIFF etc.

- Web interface to OPeNDAP allows subset selection and retrieval as ASCII (or other formats using an alternate implementation)
- Can access files directly from tools (Python etc) by dropping the .html from URL.
- Only works with netCDF/HDF

← → ↻ dapds00.nci.org.au/thredds/dodsC/r11/GSWA_Geophysics/WA_Gravity_Grids ☆ ☰

OPeNDAP Dataset Access Form

Action:

Data URL:

Global Attributes:

```
Conventions: CF-1.5
GDAL: GDAL 1.11.3, released 2015/09/16
history: Wed Apr 13 09:32:08 2016: ncrename -v
Band1,gravity_merged
./GSWA_Geophysics/WA_Gravity_Grids/WA_400m_Grav_Merge_v1_2016.nc
```

Variables:

CRS: String

crs =

```
grid_mapping_name: latitude_longitude
longitude_of_prime_meridian: 0.0
semi_major_axis: 6378137.0
inverse_flattening: 298.257222101
```

lat: Array of 64 bit Reals [lat = 0.5773]

lat:

```
standard_name: latitude
long_name: latitude
units: degrees_north
```

lon: Array of 64 bit Reals [lon = 0.4224]

lon:

```
standard_name: longitude
long_name: longitude
units: degrees_east
```

gravity_merged: Grid

lat: lon:

```
long_name: WA State grid merge of Gravity data
_fillValue: -99999.0
grid_mapping: crs
units: um/s^2
```

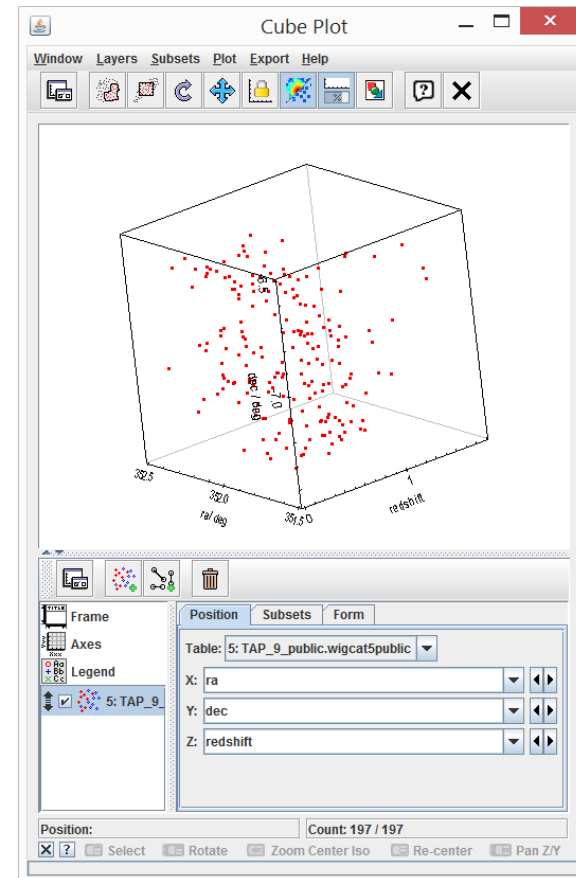
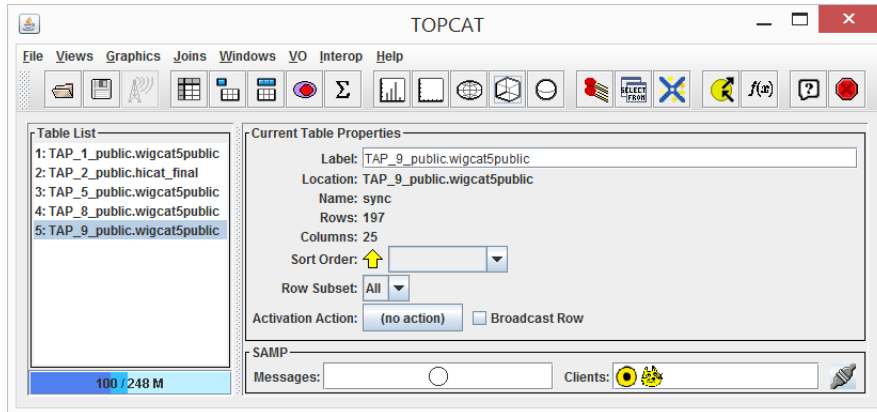
For questions or comments about this dataset, contact the administrator of this server [Support] at: help@nci.org.au

For questions or comments about OPeNDAP, email OPeNDAP support at: support@opendap.org

DDS:


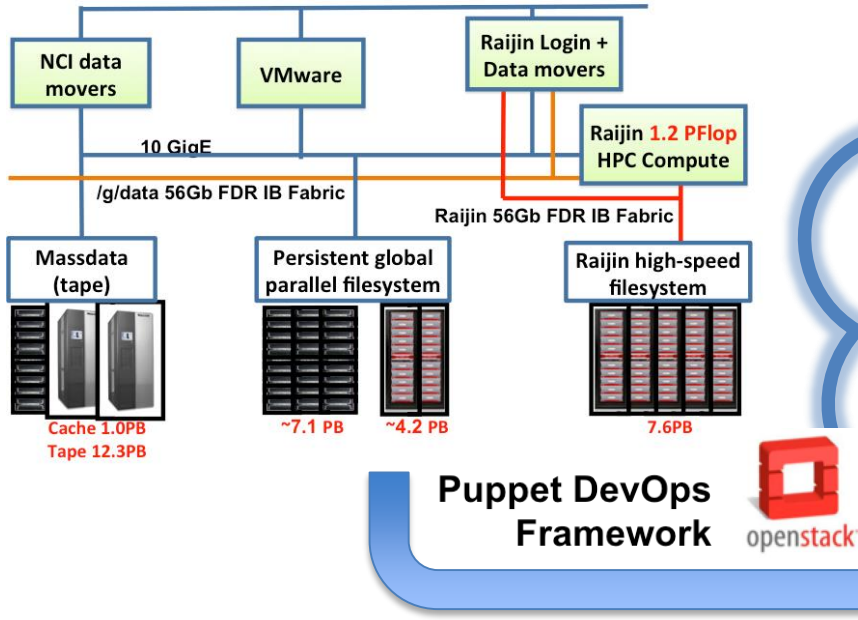
```
Dataset {
  String crs;
  Float64 lat[lat = 5774];
  Float64 lon[lon = 4225];
  Grid {
    ARRAY:
      Float32 gravity_merged[lat = 5774][lon = 4225];
    MAPS:
      Float64 lat[lat = 5774];
      Float64 lon[lon = 4225];
  } gravity_merged;
} r11/GSWA_geophysics/WA_Gravity_Grids/WA_400m_Grav_Merge_v1_2016.nc;
```

- <http://skymapper.anu.edu.au/news/early-data-release-live/>



Bring the
scientists TO the
data!

Integrated HPC-HPD Environment




3000 Core Cloud

Data Services




THREDDS
ESGF Earth System Grid Federation
OPeNDAP
GeoServer

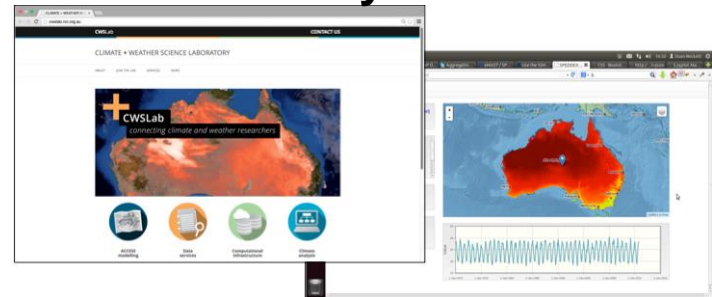
Server-side analysis and visualization




VDI: Cloud scale user desktops on data

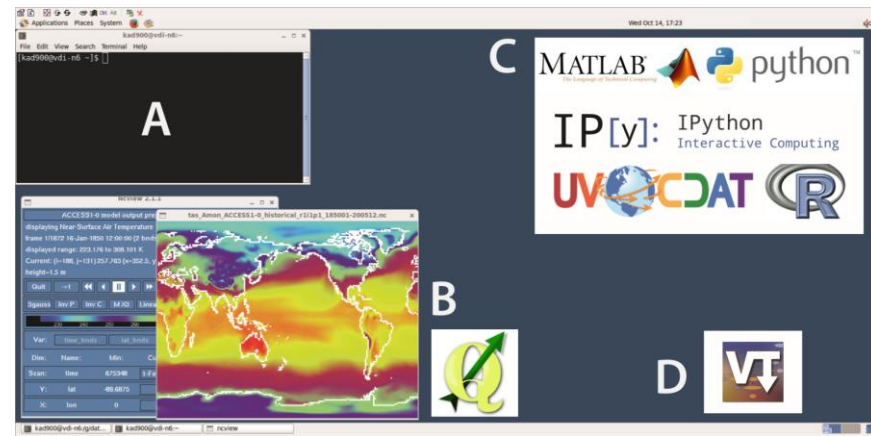


Web-time analytics software

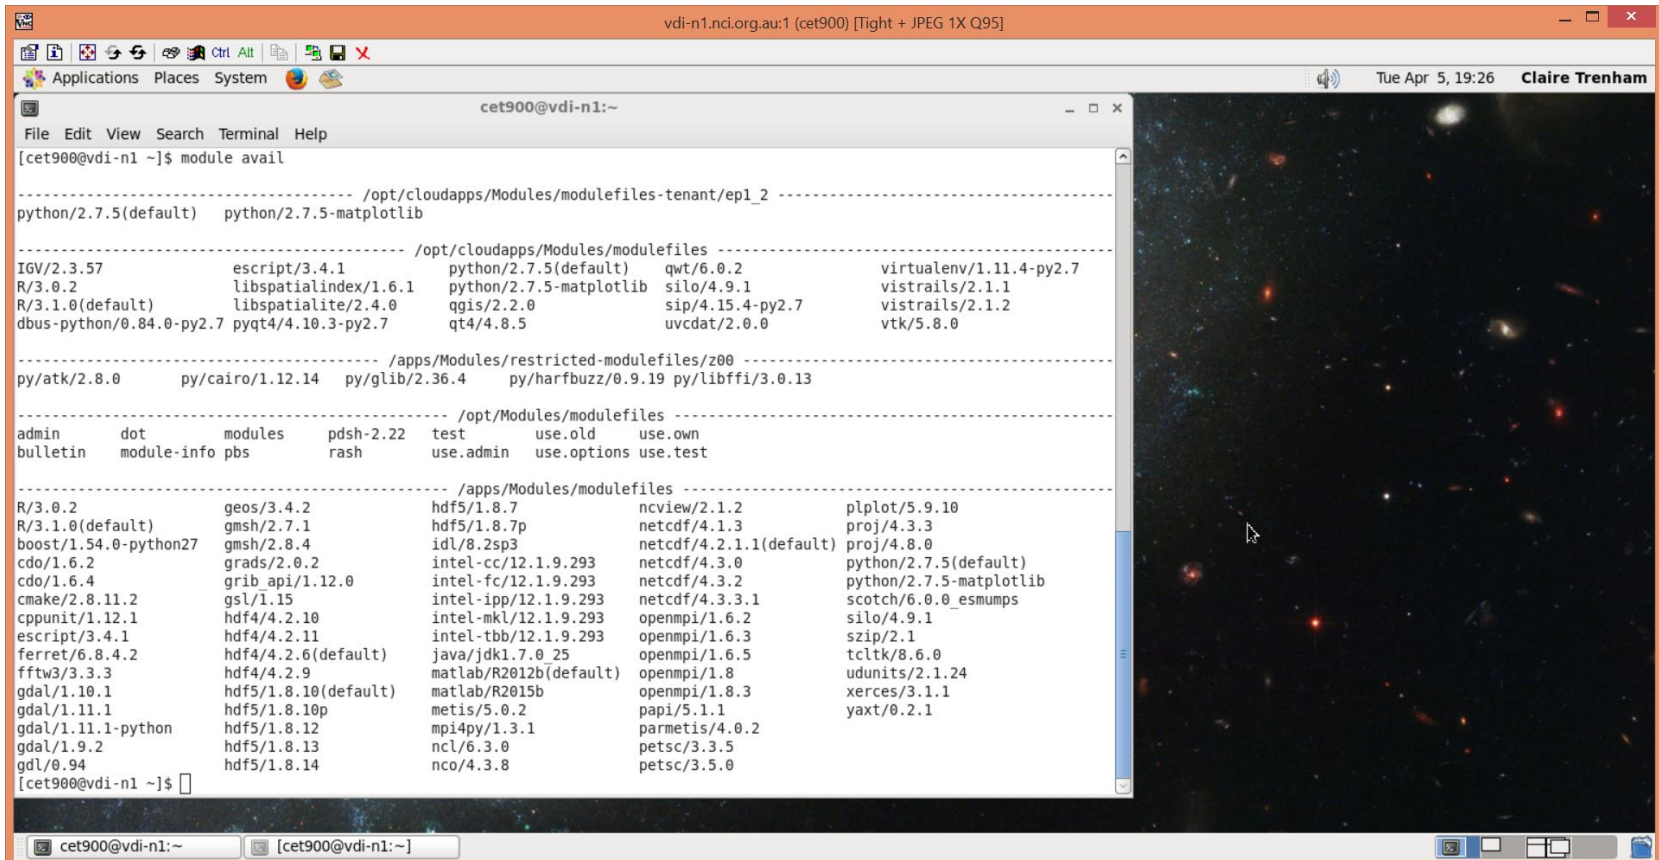




- Tools to support climate data analysis & visualisation
- Virtual laboratory to access, process & analyse data
- Analyses require input data to be consistent format
- Workflow tools allow science community to implement own analyses without dealing directly with filesystems & HPC
- A range of standard software tools available in this environment, connected to the global Lustre filesystem and HPC



- <https://training.nci.org.au/course/view.php?id=3>



```

vdi-n1.nci.org.au:1 (cet900) [Tight + JPEG 1X Q95]
Applications Places System
cet900@vdi-n1:~
File Edit View Search Terminal Help
[cet900@vdi-n1 ~]$ module avail

----- /opt/cloudapps/Modules/modulefiles-tenant/ep1_2 -----
python/2.7.5(default)  python/2.7.5-matplotlib

----- /opt/cloudapps/Modules/modulefiles -----
IGV/2.3.57             escript/3.4.1       python/2.7.5(default)  qwt/6.0.2             virtualenv/1.11.4-py2.7
R/3.0.2               libspatialindex/1.6.1  python/2.7.5-matplotlib  silo/4.9.1           vistrails/2.1.1
R/3.1.0(default)     libspatialite/2.4.0   qgis/2.2.0              sip/4.15.4-py2.7    vistrails/2.1.2
dbus-python/0.84.0-py2.7  pyqt4/4.10.3-py2.7  qt4/4.8.5              uvcdat/2.0.0        vtk/5.8.0

----- /apps/Modules/restricted-modulefiles/z00 -----
py/atk/2.8.0          py/cairo/1.12.14  py/glib/2.36.4          py/harfbuzz/0.9.19  py/libffi/3.0.13

----- /opt/Modules/modulefiles -----
admin      dot      modules  pdsh-2.22  test      use.old  use.own
bulletin   module-info  pbs      rash       use.admin use.options use.test

----- /apps/Modules/modulefiles -----
R/3.0.2          geos/3.4.2          hdf5/1.8.7          ncview/2.1.2          plplot/5.9.10
R/3.1.0(default)  gms/2.7.1          hdf5/1.8.7p        netcdf/4.1.3          proj/4.3.3
boost/1.54.0-python27  gms/2.8.4          idl/8.2sp3         netcdf/4.2.1.1(default)  proj/4.8.0
cdo/1.6.2         grads/2.0.2        intel-cc/12.1.9.293  netcdf/4.3.0         python/2.7.5(default)
cdo/1.6.4         grib_api/1.12.10   intel-fc/12.1.9.293  netcdf/4.3.2         python/2.7.5-matplotlib
cmake/2.8.11.2    gsl/1.15           intel-ipp/12.1.9.293  netcdf/4.3.3.1       scotch/6.0.0_esmumps
cppunit/1.12.1   hdf4/4.2.10        intel-mkl/12.1.9.293  openmpi/1.6.2        silo/4.9.1
escript/3.4.1    hdf4/4.2.11        intel-tbb/12.1.9.293  openmpi/1.6.3        szip/2.1
ferret/6.8.4.2   hdf4/4.2.6(default)  java/jdk1.7.0_25     openmpi/1.6.5        tcltk/8.6.0
fftw3/3.3.3      hdf4/4.2.9         matlab/R2012b(default)  openmpi/1.8          udunits/2.1.24
gdal/1.10.1      hdf5/1.8.10(default)  matlab/R2015b        openmpi/1.8.3        xerces/3.1.1
gdal/1.11.1      hdf5/1.8.10p        metis/5.0.2          papi/5.1.1           yast/0.2.1
gdal/1.11.1-python  hdf5/1.8.12        mpi4py/1.3.1         parmetis/4.0.2
gdal/1.9.2       hdf5/1.8.13        ncl/6.3.0            petsc/3.3.5
gd/0.94          hdf5/1.8.14        nco/4.3.8            petsc/3.5.0
[cet900@vdi-n1 ~]$
  
```

- Data downloading and analysis by many users also has potential risks (apart from the data being too big for this to be feasible!)
 - Versioning of data used in analysis
 - Provenance tracking
 - Errata and Reporting
 - Documentation incorporated in file in case a file gets isolated?
- Bringing scientists to the data can help mitigate these issues by ensuring everyone is working on the same data (with provenance capture?)

- ASKAP data is stored at Pawsey in a different physical architecture than the RDS data at NCI
- ASKAP data search is through CASDA
- FITS format does not support remote subsetting (I think?) – j2k + JPIP?
- But... if you think we might be able to provide advice or ideas, we're happy to talk to you :)



Thanks for listening :)

- Claire Trenham, Research Data Services, NCI
Claire.Trenham@anu.edu.au

- <https://training.nci.org.au>
- <https://github.com/nci/Notebooks>
- <http://nci.org.au/user-support/getting-help/>

- To apply to use NCI facilities
 - Partner Shares (CSIRO, CAASTRO, AAL, uni LIEF)
 - National Computational Merit Allocation Scheme
 - Start-up project
 - <http://nci.org.au/access/getting-access-to-the-national-facility/allocation-schemes/>