NCI

NATIONAL COMPUTATIONAL INFRASTRUCTURE

Integrating Big Geoscience Data into the Petascale National Environmental Research Interoperability Platform (NERDIP): Successes and Unforeseen challenges

Lesley Wyborn and Ben Evans.

Australian Government Department of Education
Australian National University
Australian Government Bureau of Meteorology
CSIRO
Australian Government Geoscience Australia
Australian Government Australian Research Council

nci.org.au
@NCInews

The 2006-2015 Australian Research Funding Schemes

- Two main tranches of funding:
  - National Collaborative Research Infrastructure Strategy (NCRIS)
    - $542M for 2006-2011 ($75 M for cyberinfrastructure)
  - Super Science Initiative
    - $901 million for 2009-2013 ($347M for cyberinfrastructure)
  - Annual Maintenance funding of around $180M pa since 2014-2015
- All programmes were designed ensure that Australian research continues to be competitive and rank highly on an international scale.

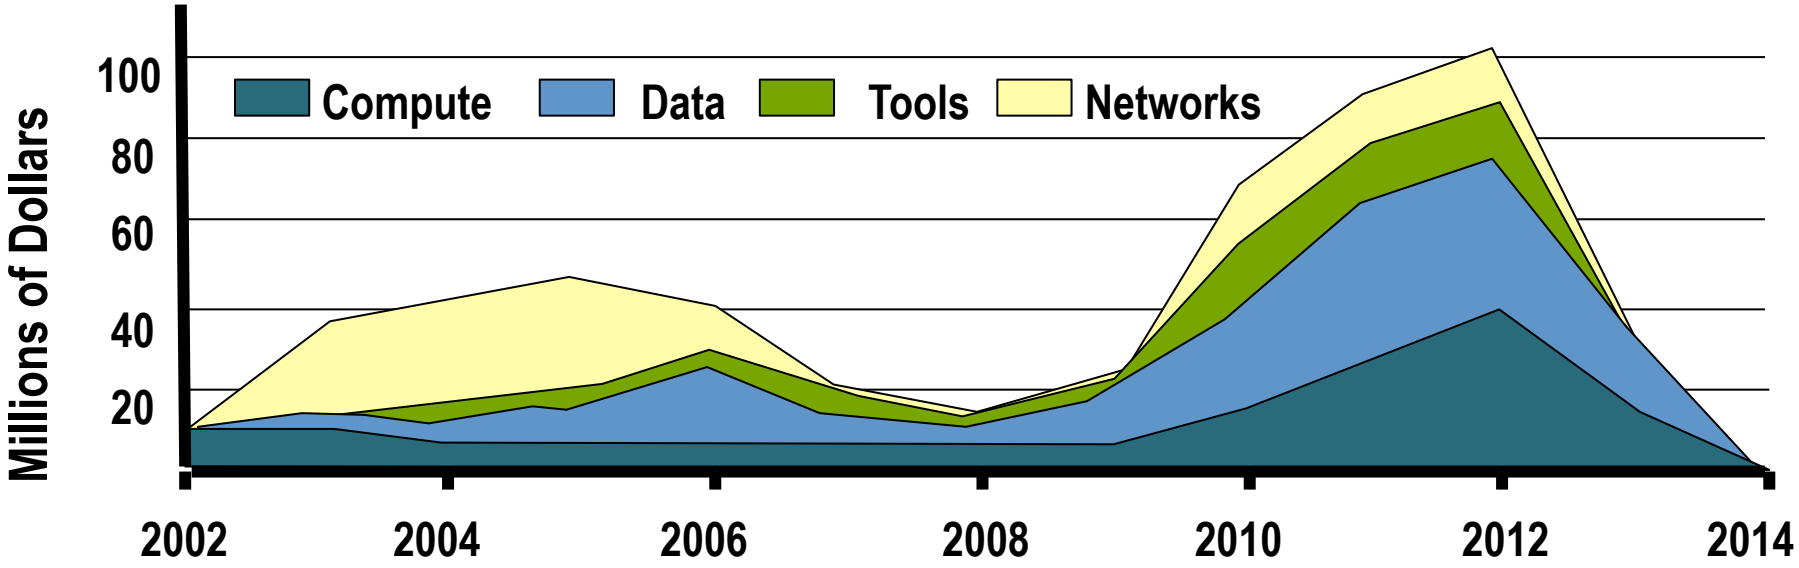

We went from nothing…..to……..

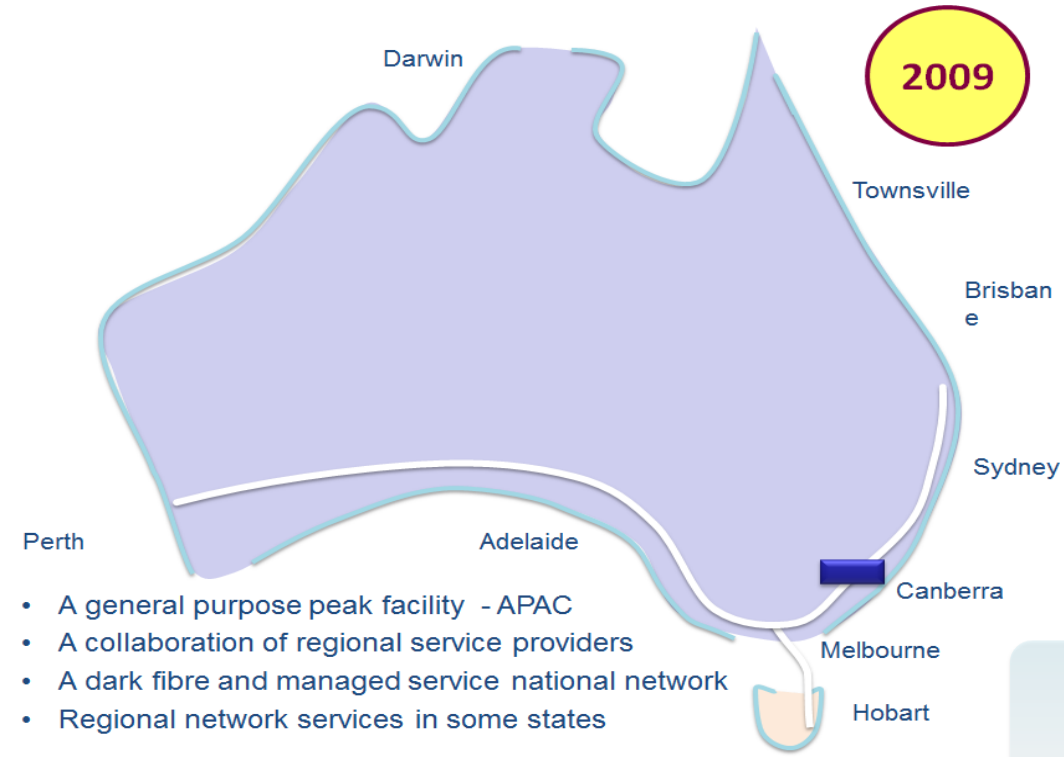

- A general purpose peak facility - APAC
- A collaboration of regional service providers
- A dark fibre and managed service national network
- Regional network services in some states

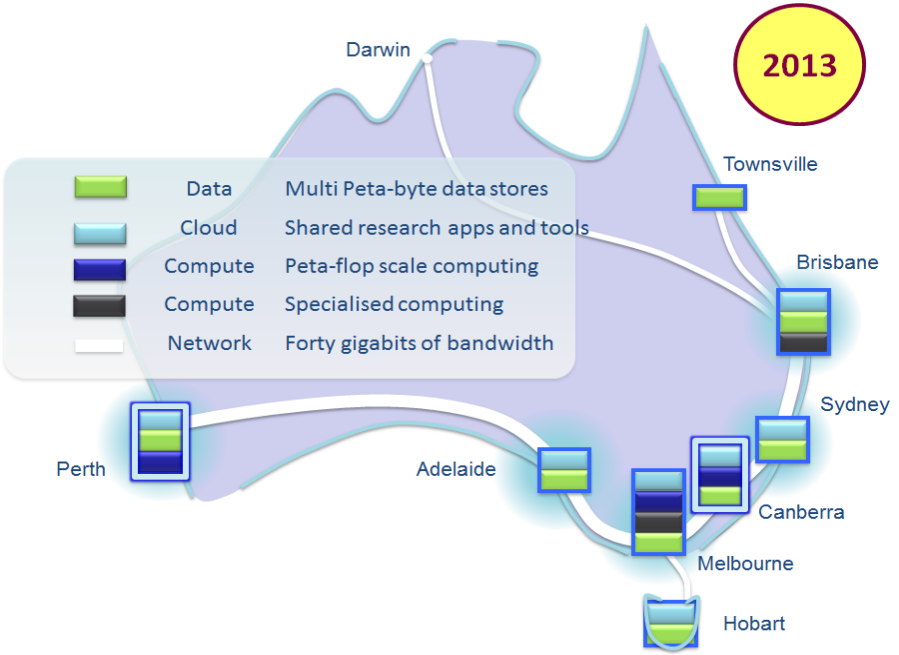

The Research Data Storage Infrastructure

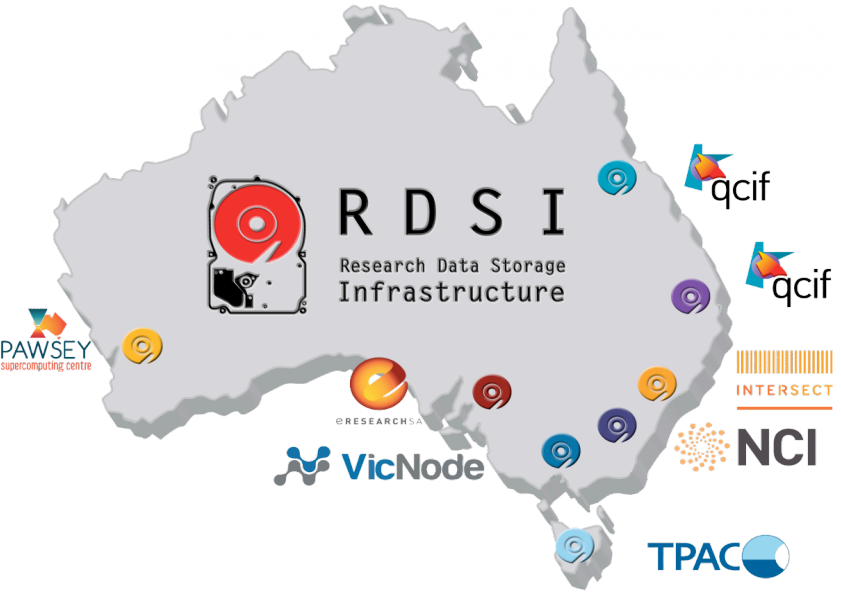

Progress on Data Ingest as of 16 October, 2015:
~43 Petabytes in 8 distributed nodes

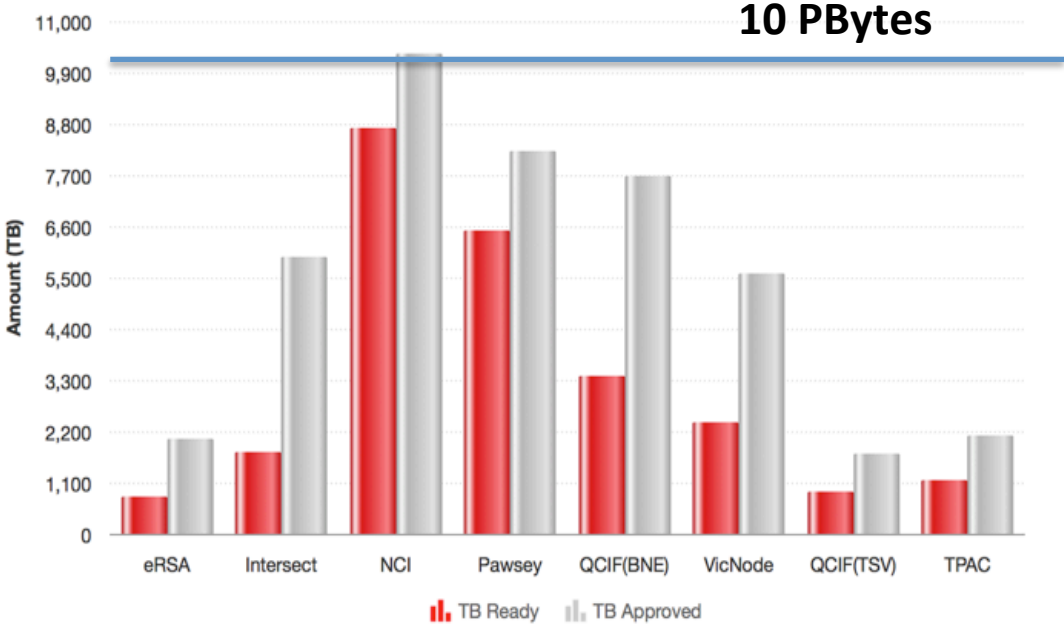

eRSA	Intersect	NCI	Pawsey	QCIF(BNE)	VicNode	QCIF(TSV)	TPAC
785	1743	8709	6504	3406	2404	907	1160
2049	5957	10296	8202	7699	5575	1729	2105

Primary Nodes | Additional Nodes

Total: TB Ready	25,617	Total: TB Approved	43,611

Source: https://www.rds.edu.au/

Integrated World-class Scientific Computing Environment at NCI

Puppet DevOps Framework

3000 Core Cloud

openstack

Integrated HPC-HPD Environment

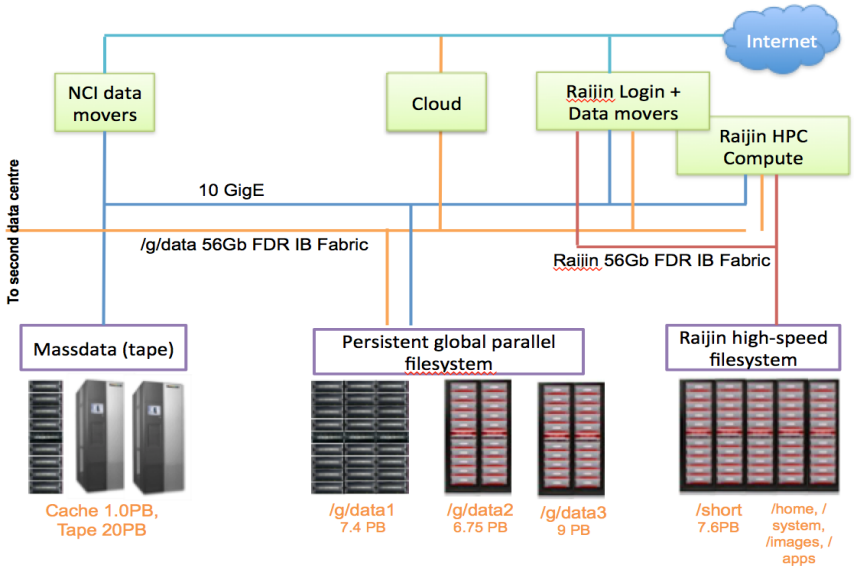

Data Services

THREDDS ESGF Earth System Grid Federation

OPeNDAP GeoServer

Server-side analysis and visualization

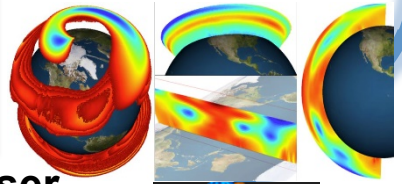

UV CDAT

VDI: Cloud scale user desktops on data

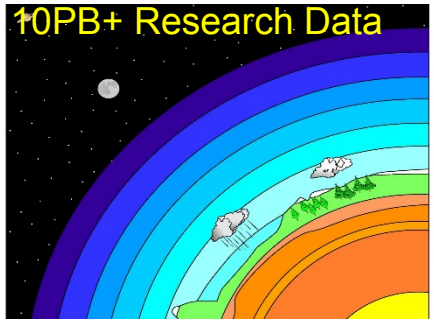

Web-time analytics software

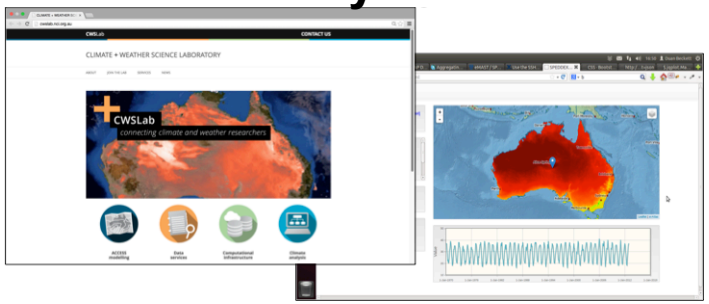

National Environment Research Data Collections (NERDC)

1. Climate/ESS Model Assets and Data Products
2. Earth and Marine Observations and Data Products
3. Geoscience Collections
4. Terrestrial Ecosystems Collections
5. Water Management and Hydrology Collections

Data Collections	Approx. Capacity
CMIP5, CORDEX	2 Pbytes
ACCESS products	3.3 Pbytes
LANDSAT, MODIS, VIIRS, AVHRR, INSAR, MERIS	2 Pbytes
Digital Elevation, Bathymetry, Onshore Geophysics	400 Tbytes
Seasonal Climate	600 Tbytes
Bureau of Meteorology Observations	400 Tbytes
Bureau of Meteorology Ocean-Marine	220 Tbytes
Terrestrial Ecosystem	290 Tbytes
Reanalysis products	175 Tbytes

10+ PB of Data for Interdisciplinary Science

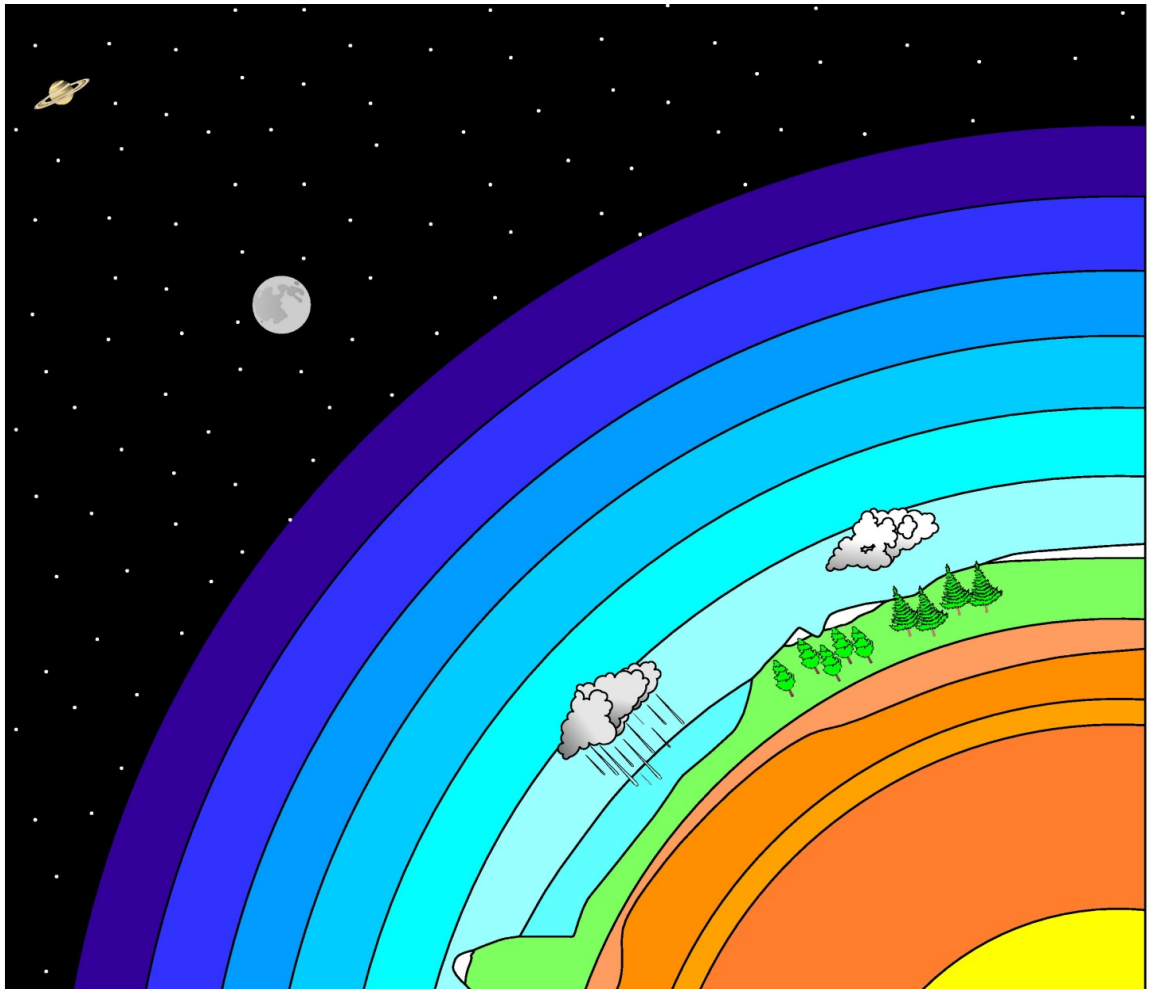

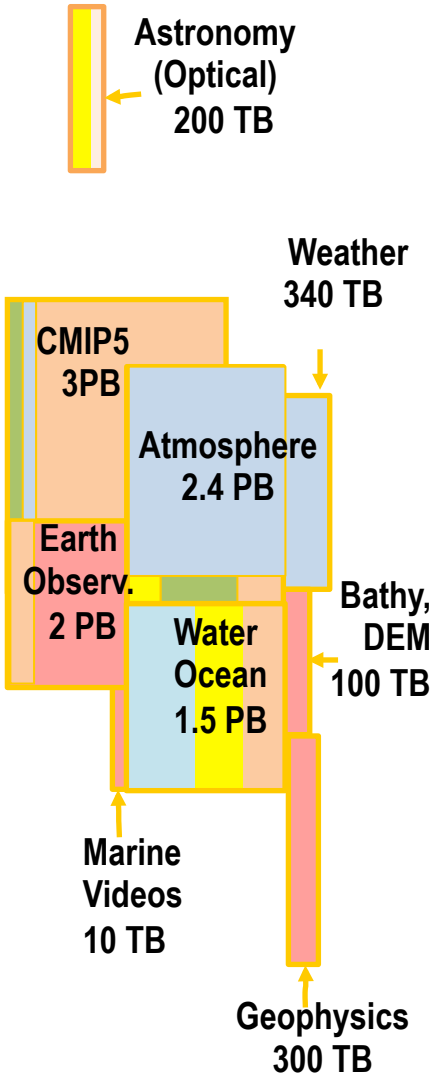

GA | CSIRO | ANU | Other National | Inter-national

Managing 10+ PB of Data for Scalable In-situ Access

- Combined and integrated, the NCI collections are too large to move
  - bandwidth limits the capacity to move them easily
  - the data transfers are too slow, complicated and too expensive
  - even if our data can be moved, few can afford to store 10 PB on spinning disk
- We need to change our focus to:
  - moving users to the data (for sophisticated analysis)
  - moving processing to data
  - having online applications to process the data in-situ
  - Improving the sophistication of users – with our help

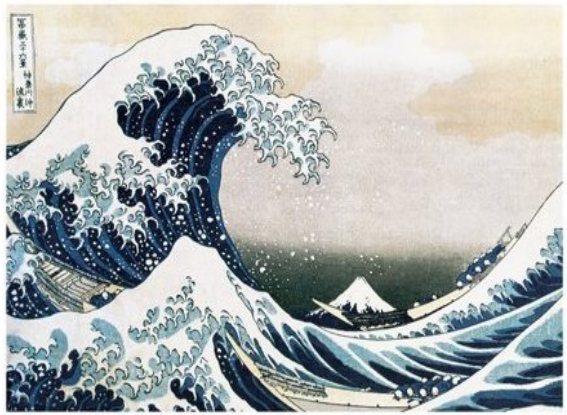

- We called for a new form of system design where:
  - storage and various types of computation are co-located
  - systems are programmed and operated to allow users to interactively invoke different forms of analysis in-situ over integrated large-scale data collections

Rethinking Hardware Architectures for Data-intensive Science

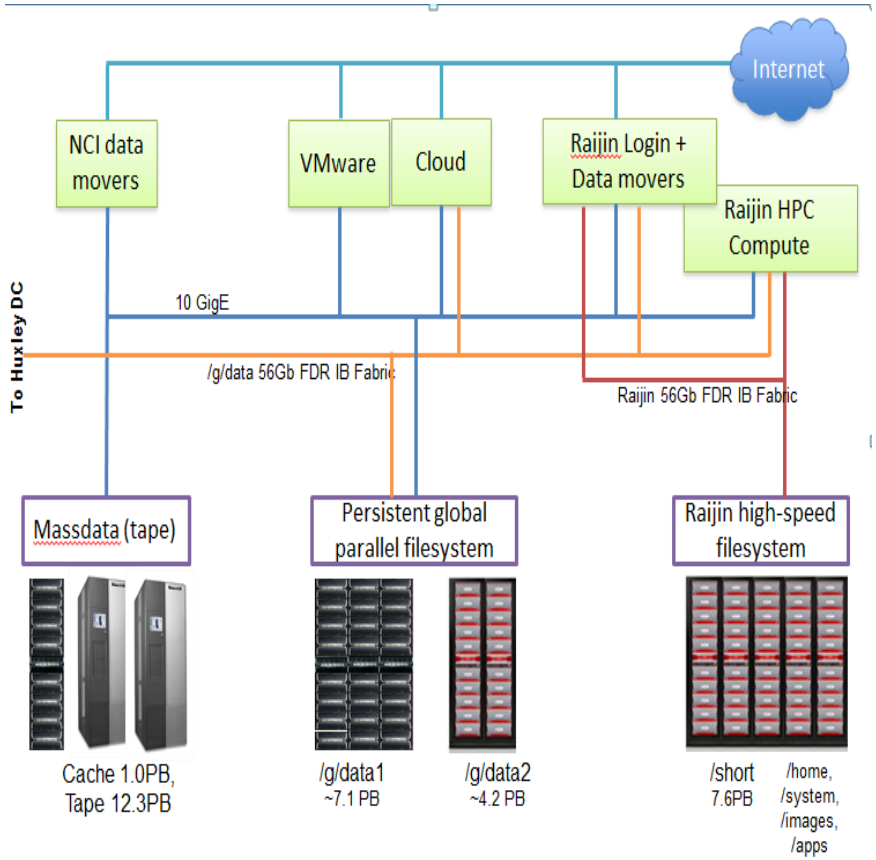

NCI's Integrated High Performance Environment

- Work at NCI has also highlighted the need for balanced systems to enable Data-intensive Science including:
  - Interconnecting processes and high throughput to reduce inefficiencies
  - The need to really care about placement of data resources
  - Better communications between the nodes
  - I/O capability to match the computational power
  - Close coupling of cluster, cloud and storage

My take is that 'Big Data' is not just about the "V's"

1. Volume: data at rest
2. Velocity: data in motion (streaming)
3. Variety: many types, forms and structures (or no structures)
4. Veracity: trustworthiness, provenance, lineage, quality
5. Validity: data that is correct
6. Visualization: data in patterns
7. Vulnerability: data at risk
8. Value: data that is meaningful
9. V?????
10. V?????

‘Big Data’ vs High Performance Data

- Big Data is a relative term where the volume, velocity and variety of data exceed an organisations storage or compute capacity for accurate and timely decision making

- We define High Performance Data (HPD) as data that is carefully prepared, standardised and structured so that it can be used in Data-Intensive Science on HPC (Evans et al., 2015)

- To get on top of the Data Tsunami, we need to convert ‘Big data’ collections into HPD by
  - Aggregating data into seamless ‘pre-processed’ data products
  - Creating hyper-cubes and self describing data arrays

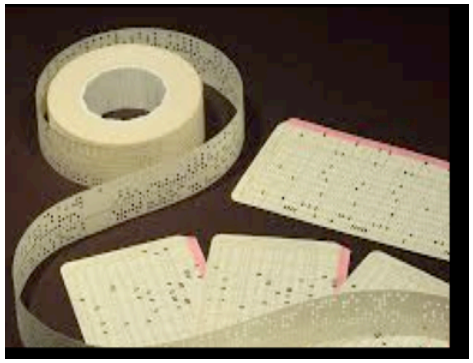

1964: 1KB = 2m of tape or ~ 20 cards

2014: a 4 GB Thumb drive =
~8000 Km of Tape
or ~83 million cards

2014: 20 PB of modern storage =
~ 32 trillion metres of tape
~ 320 trillion cards

http://www.sas.com/content/dam/SAS/en_us/doc/whitepaper1/big-data-meets-big-data-analytics-105777.pdf

Creating HPD collections: eg the Landsat Cube

- The Landsat cube arranges 636,000 Landsat Source scenes spatially and temporally, to allow flexible but efficient large-scale in-situ analysis
- The data is partitioned into spatially-regular, time-stamped, band-aggregated tiles which are presented as temporal stacks.

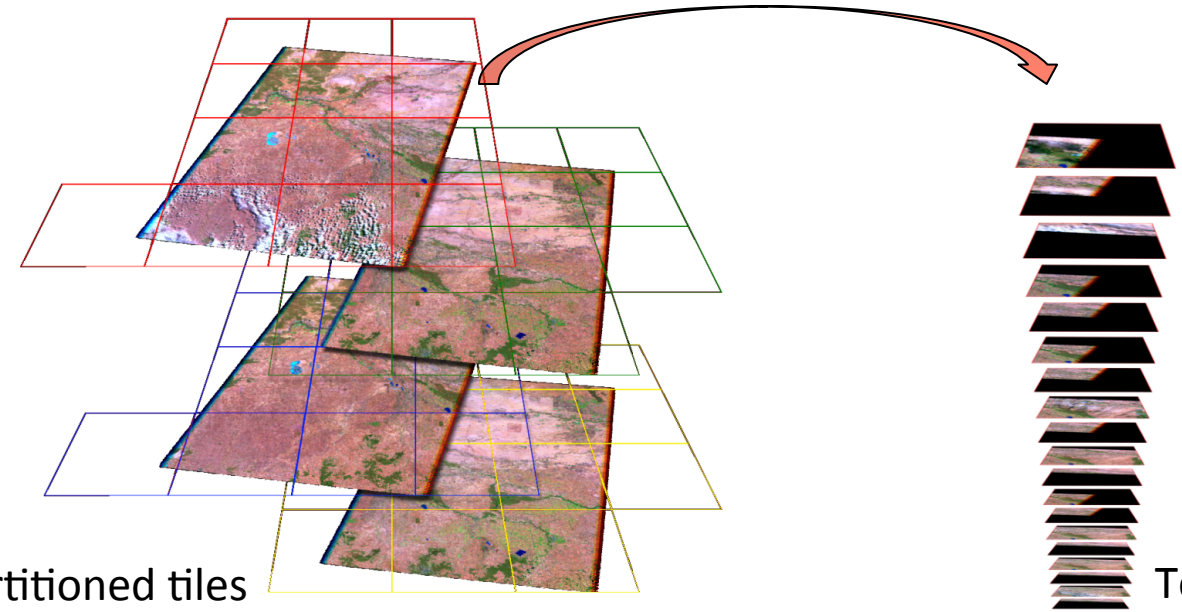

Spatially partitioned tiles

Temporal Stack

Current Landsat Holdings Reformatted as HPD

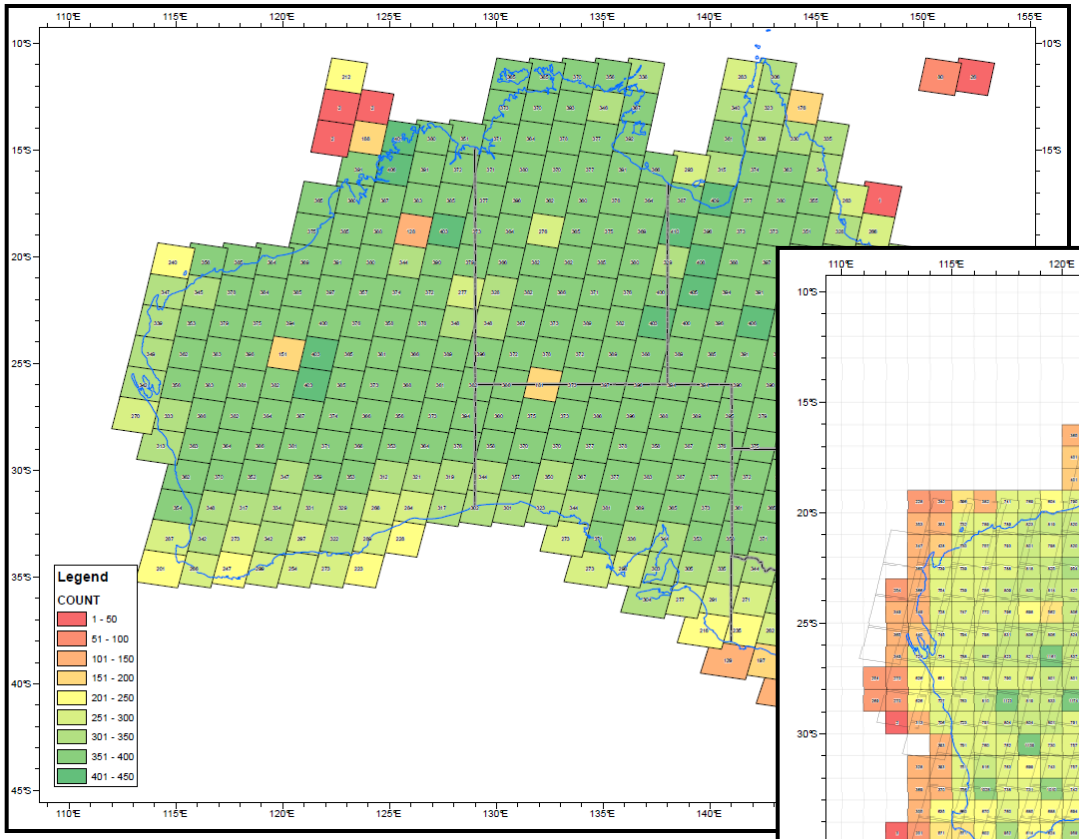

636,000 Landsat Source Scene Datasets (~52 x 10^{12} Pixels)

4M Spatially-Regular Time-Stamped Tiles (0.5 PB)

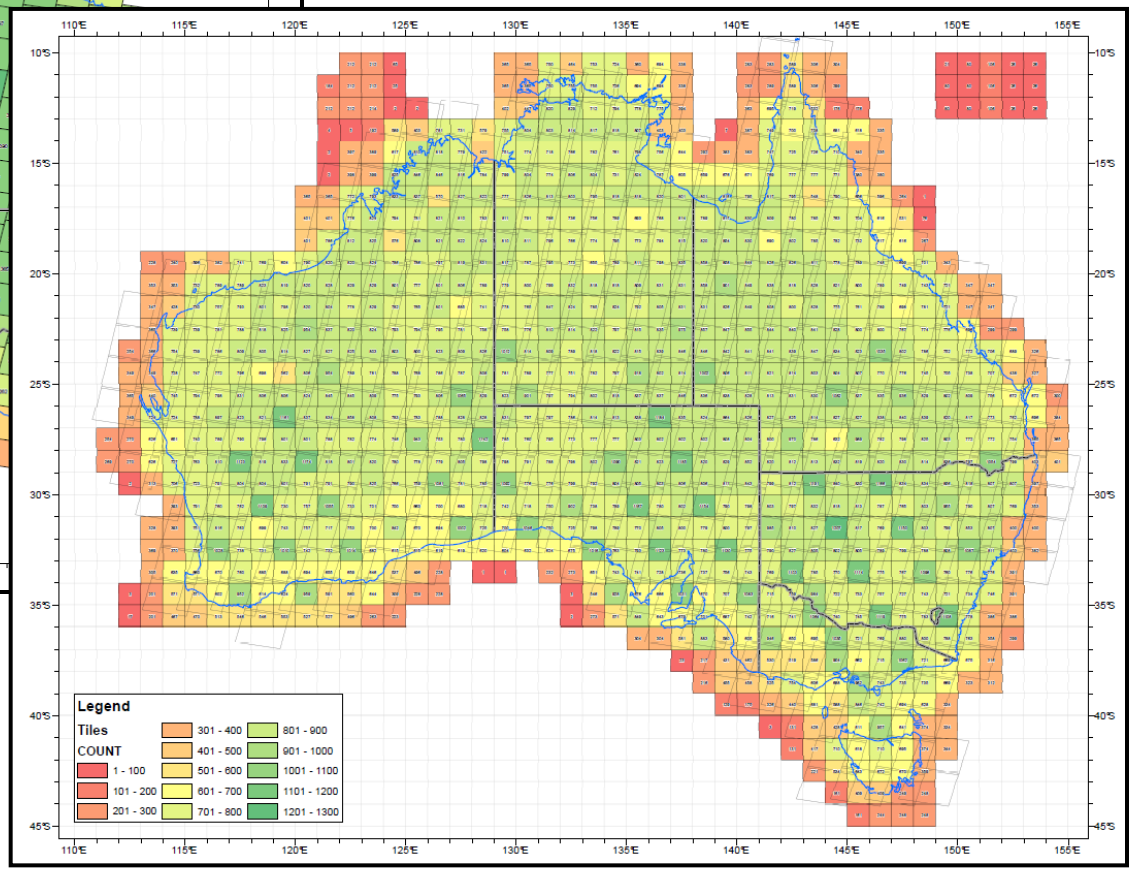

Water Detection from Space

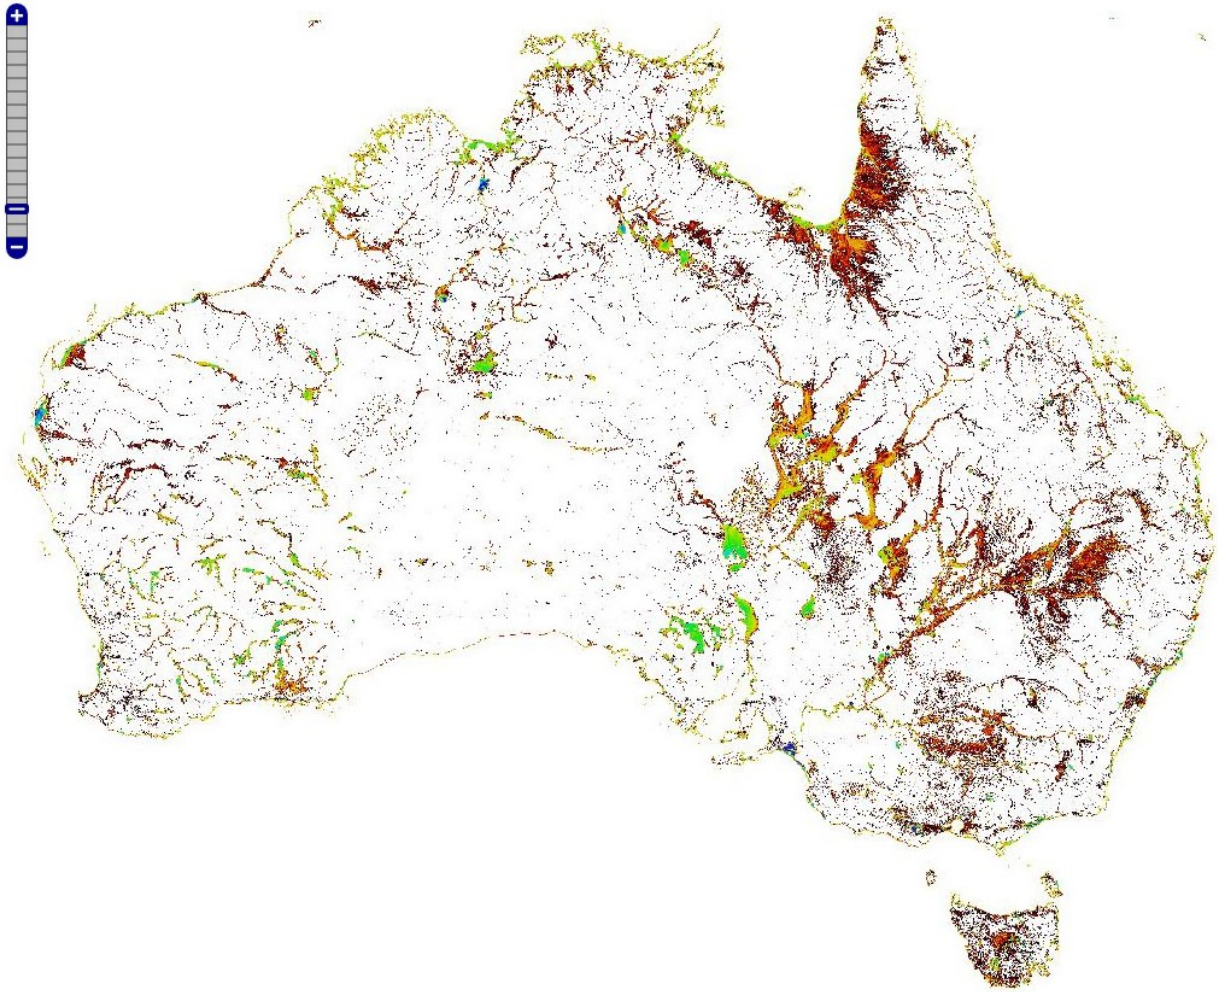

- 15 Years of data from LS5 & LS7(1998-2012)
- 25m Nominal Pixel Resolution
- Approx. 133,000 individual source scenes in approx. 12,400 passes
- Entire archive of 1,312,087 ARG25 tiles => $21x10^{12}$ pixels can be processed in ~8 hours

Data Platforms of today need to scale down to small users

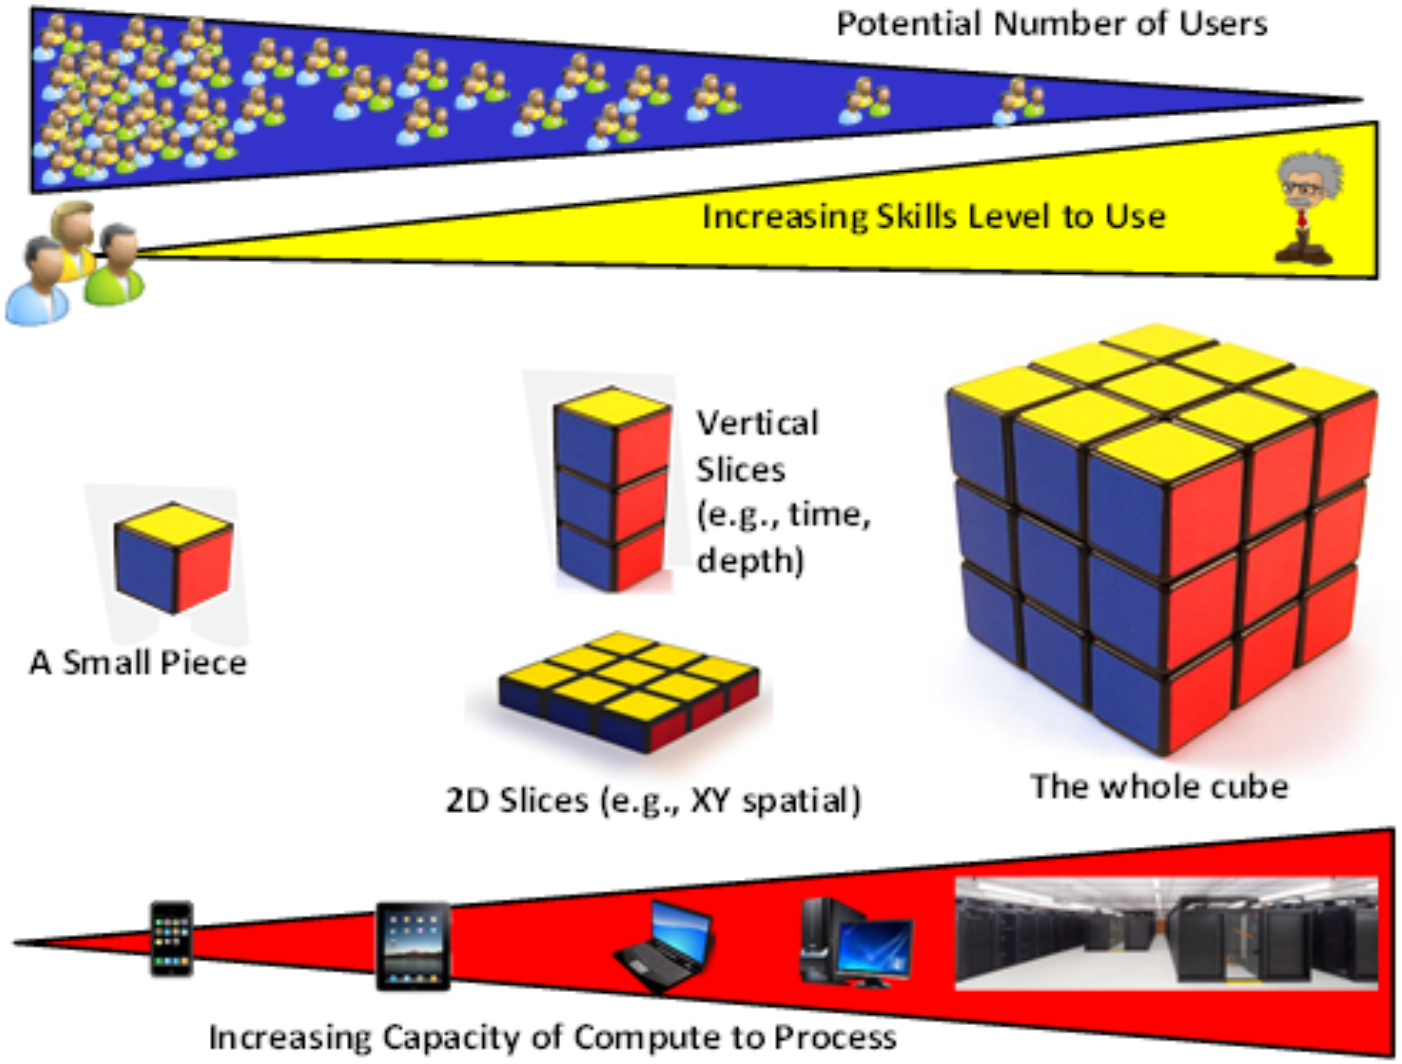

Scaling down to the smaller users

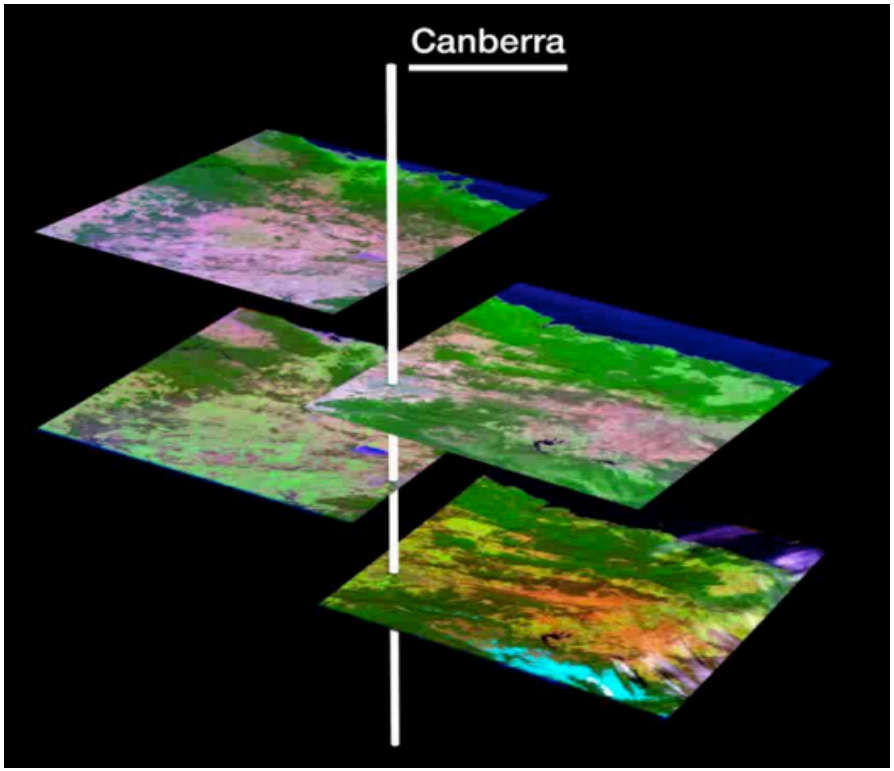

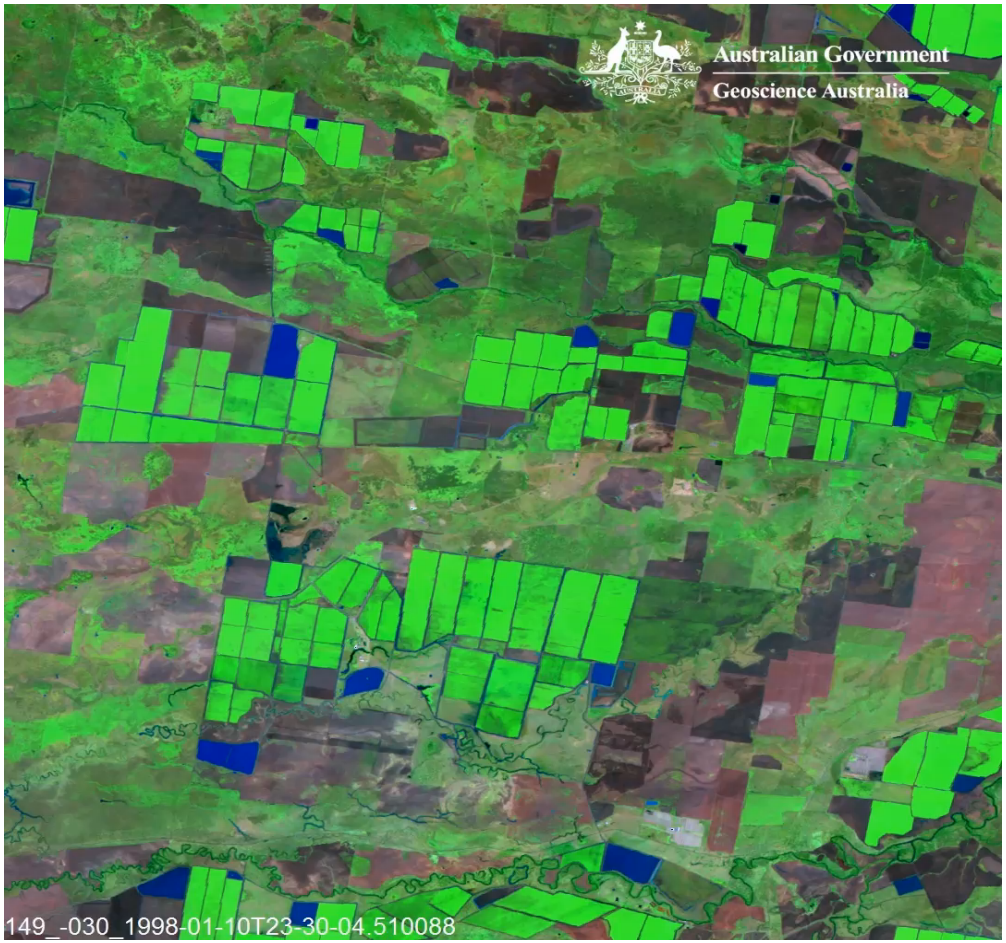

Do we enable individual scenes to be downloaded for locally hosted small scale analysis?
Or do we facilitate small scale analysis, in-situ on data sets that are dynamically updated?

Introducing the National Environmental Data Interoperability Research Platform (NERDIP)

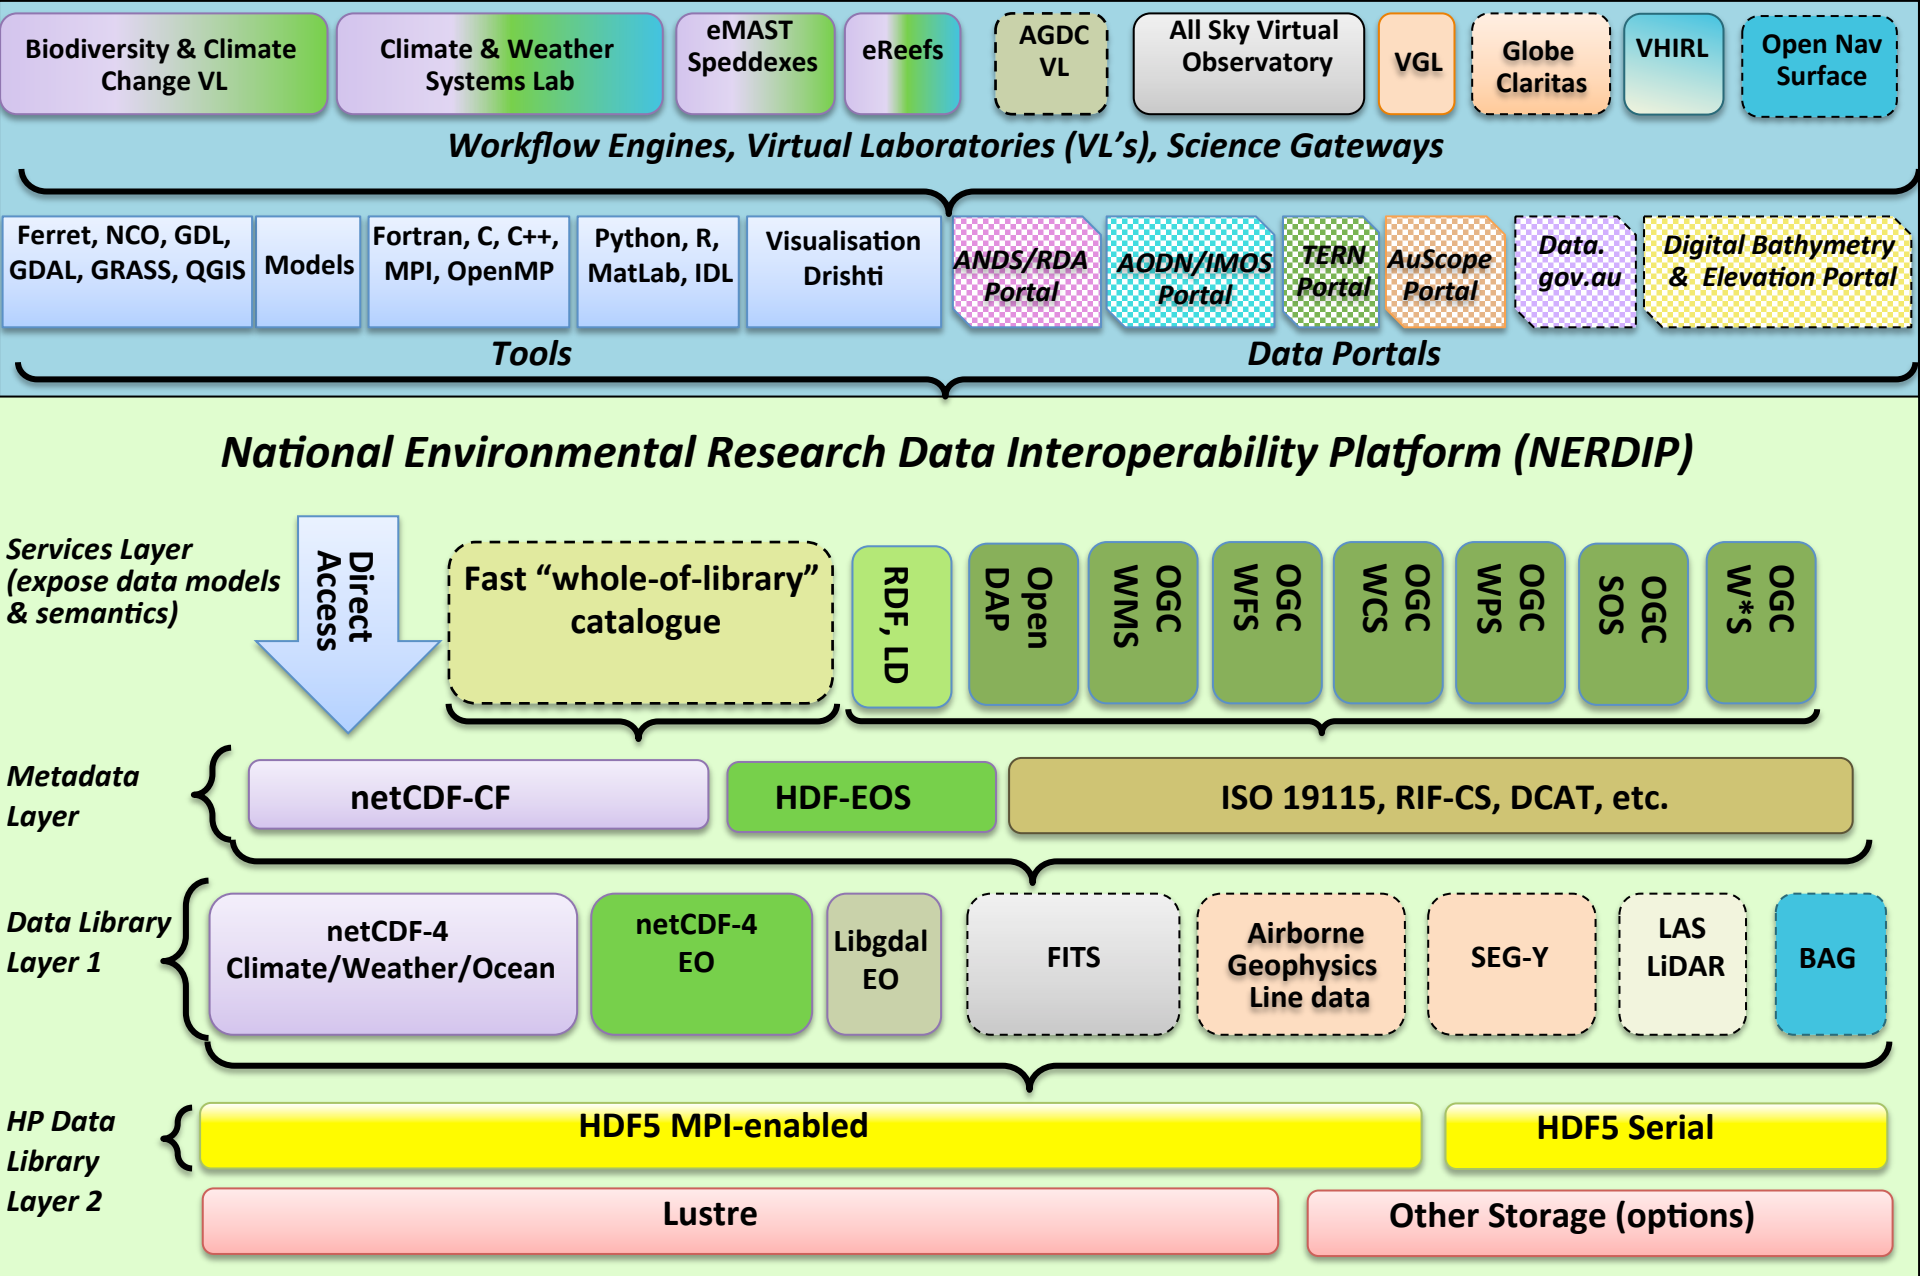

NERDIP: Enabling Multiple Ways to Interact with the Data

Infrastructure to Lower Barriers to Entry

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

Ace Users

Tools

Data Discovery

Data Portals

Data Platform

NERDIP: Enabling Multiple Ways to Interact with the Data

Infrastructure to Lower Barriers to Entry

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

Ace Users

Tools

Ferret, NCO, GDL, GDAL, GRASS, QGIS | Models | Fortran, C, C++, MPI | Python, R | Visualisation Drishti

Data Portals

Data Portals

ANDS/RDA Portal | AODN/IMOS Portal | Digital Bathymetry & Elevation Portal

Data Platform

National Environmental Research Data Interoperability Platform (NERDIP)

Services Layer (expose data models & semantics)

Direct Access

Fast "whole-of-library" catalogue

RDF, LD | Open DAP | OGC WMS | OGC WFS | OGC WCS | OGC WPS | OGC SOS | OGC W*S

Metadata Layer

netCDF-CF | HDF-EOS | ISO 19115, RIF-CS, DCAT, etc.

Data Library Layer 1

netCDF-4 Climate/Weather/Ocean | netCDF-4 EO | Libgdal EO | FITS | Airborne Geophysics Line data | SEG-Y | LAS LiDAR | BAG

HP Data Library Layer 2

HDF5 MPI-enabled | HDF5 Serial | Lustre | Other Storage (options)

Platforms Free Data from the "Prison of the Portals"

- **Portals are for visiting, platforms are for building on**
- Portals present aggregated content in a way that invites exploration, but the experience is pre-determined by a set of decisions by the builder about what is necessary, relevant and useful.
- Platforms put design decisions into the hands of users: there are innumerable ways of interacting with the data
- Platforms offer many more opportunities for innovation: new interfaces can be built, new visualisations framed, ultimately new science rapidly emerges

Tim Sherratt http://www.nla.gov.au/our-publications/staff-papers/from-portal-to-platform

NERDIP: Enabling Ace Users to Interact with the Data

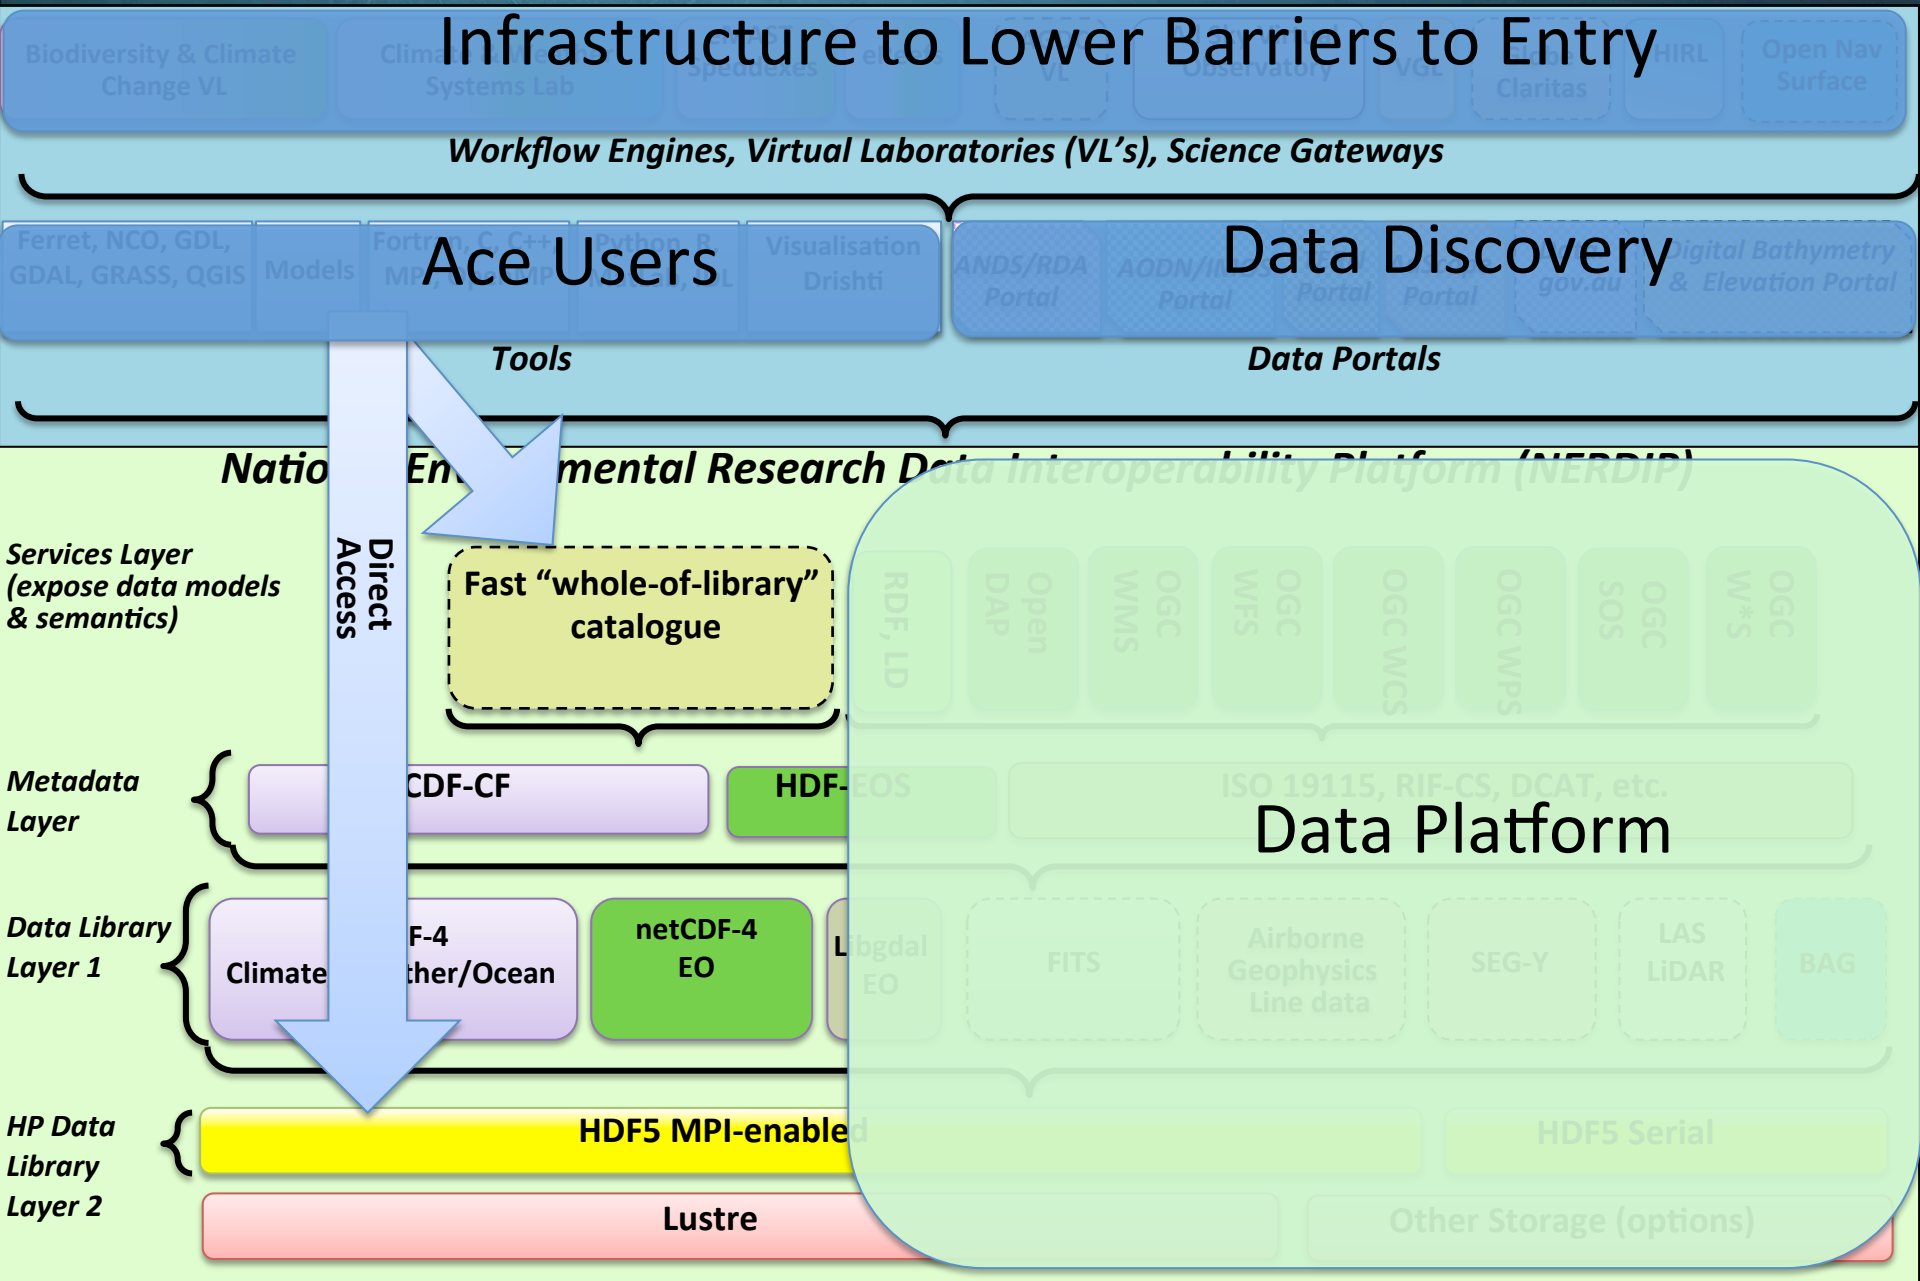

NERDIP: Enabling Ace Users to Interact with the Data

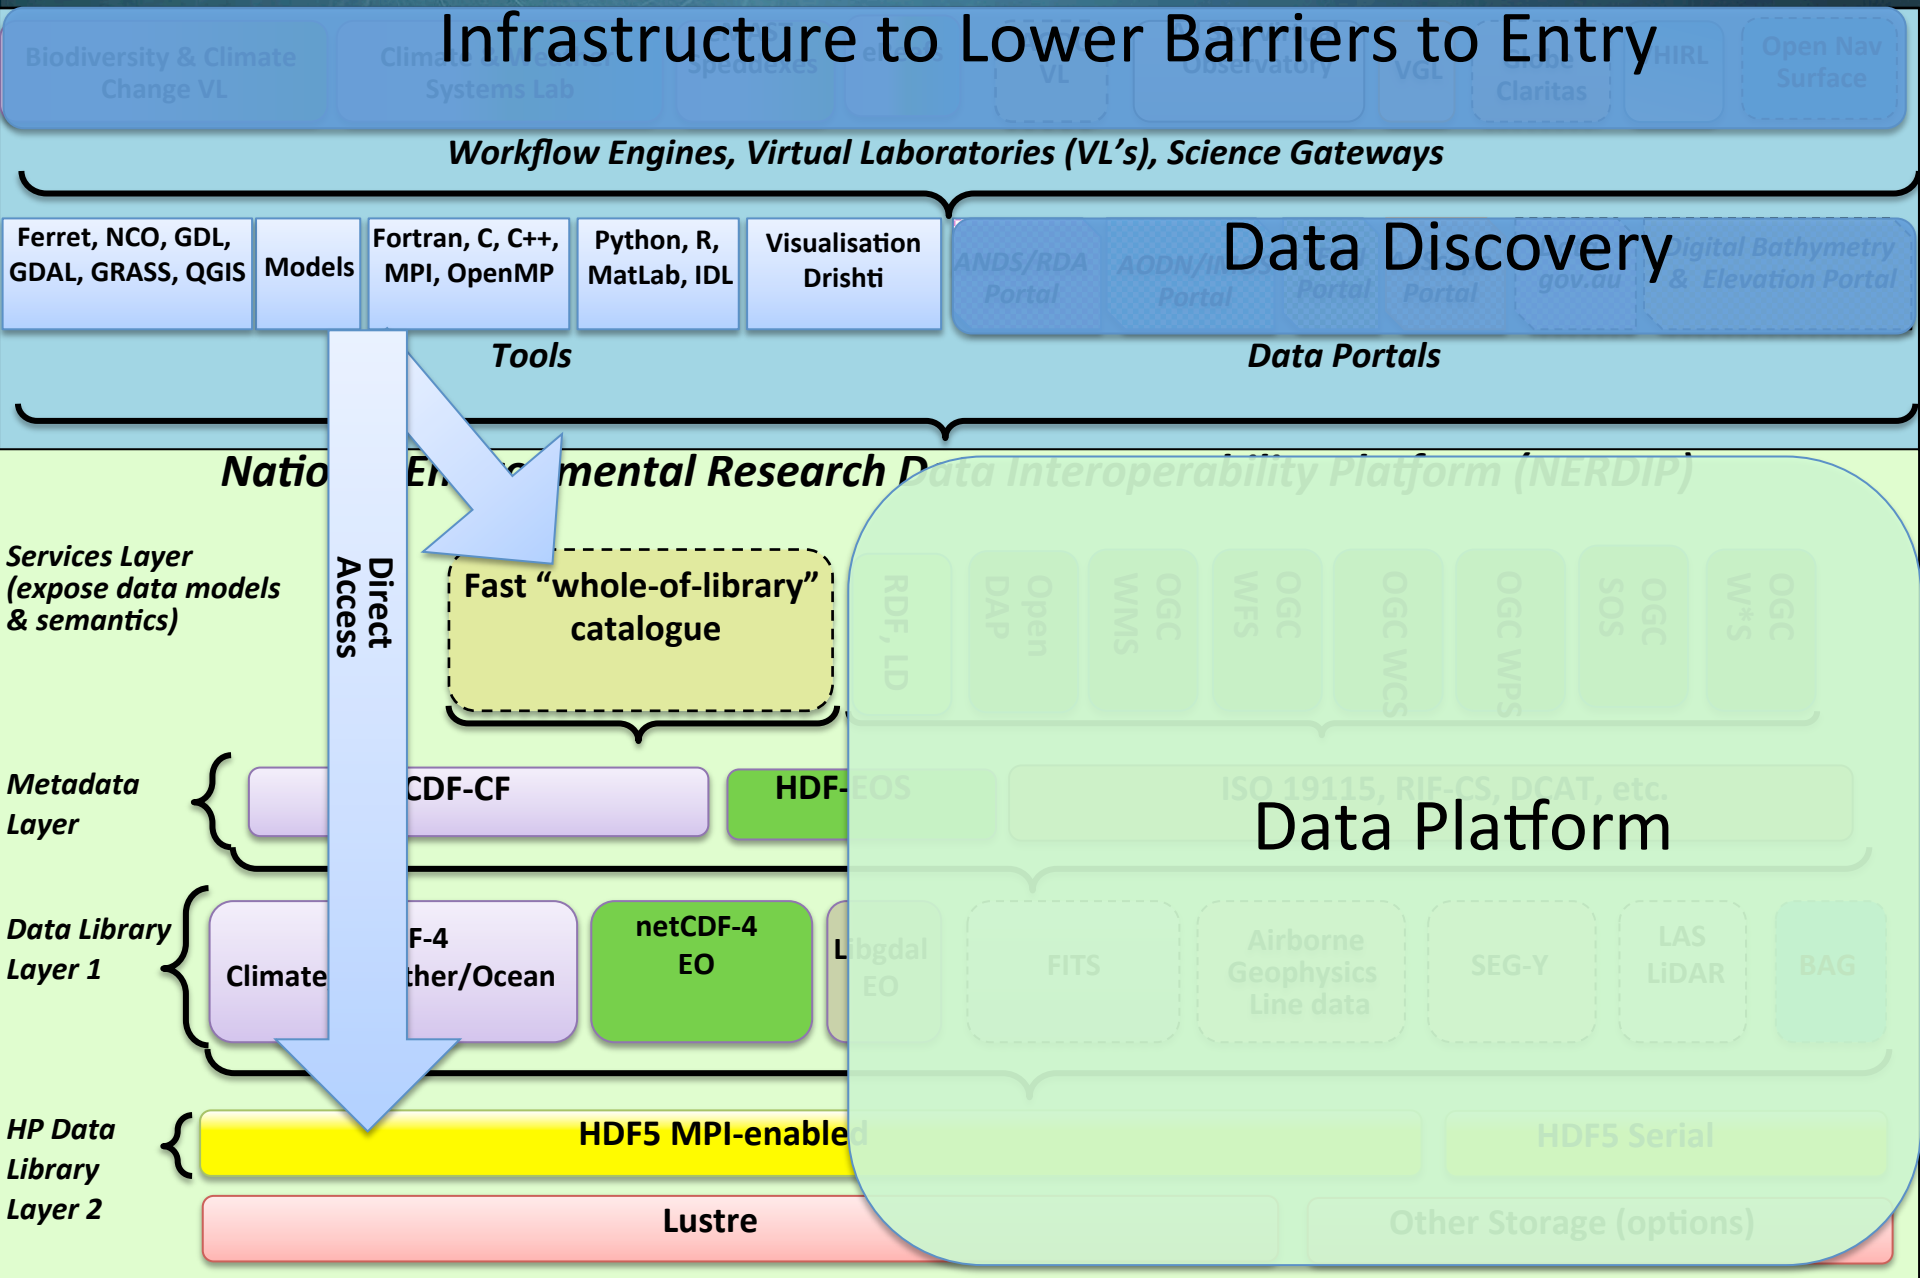

Infrastructure to Lower Barriers to Entry

APPLICATION

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

FOCUSSED DEVELOPERS

Tools | ***Data Portals***

National Environmental Research Data Interoperability Platform (NERDIP)

DATA MANAGEMENT

Data Platform

FOCUSSED DEVELOPERS

NERDIP Territorial Wars: Application Developers vs Data Managers

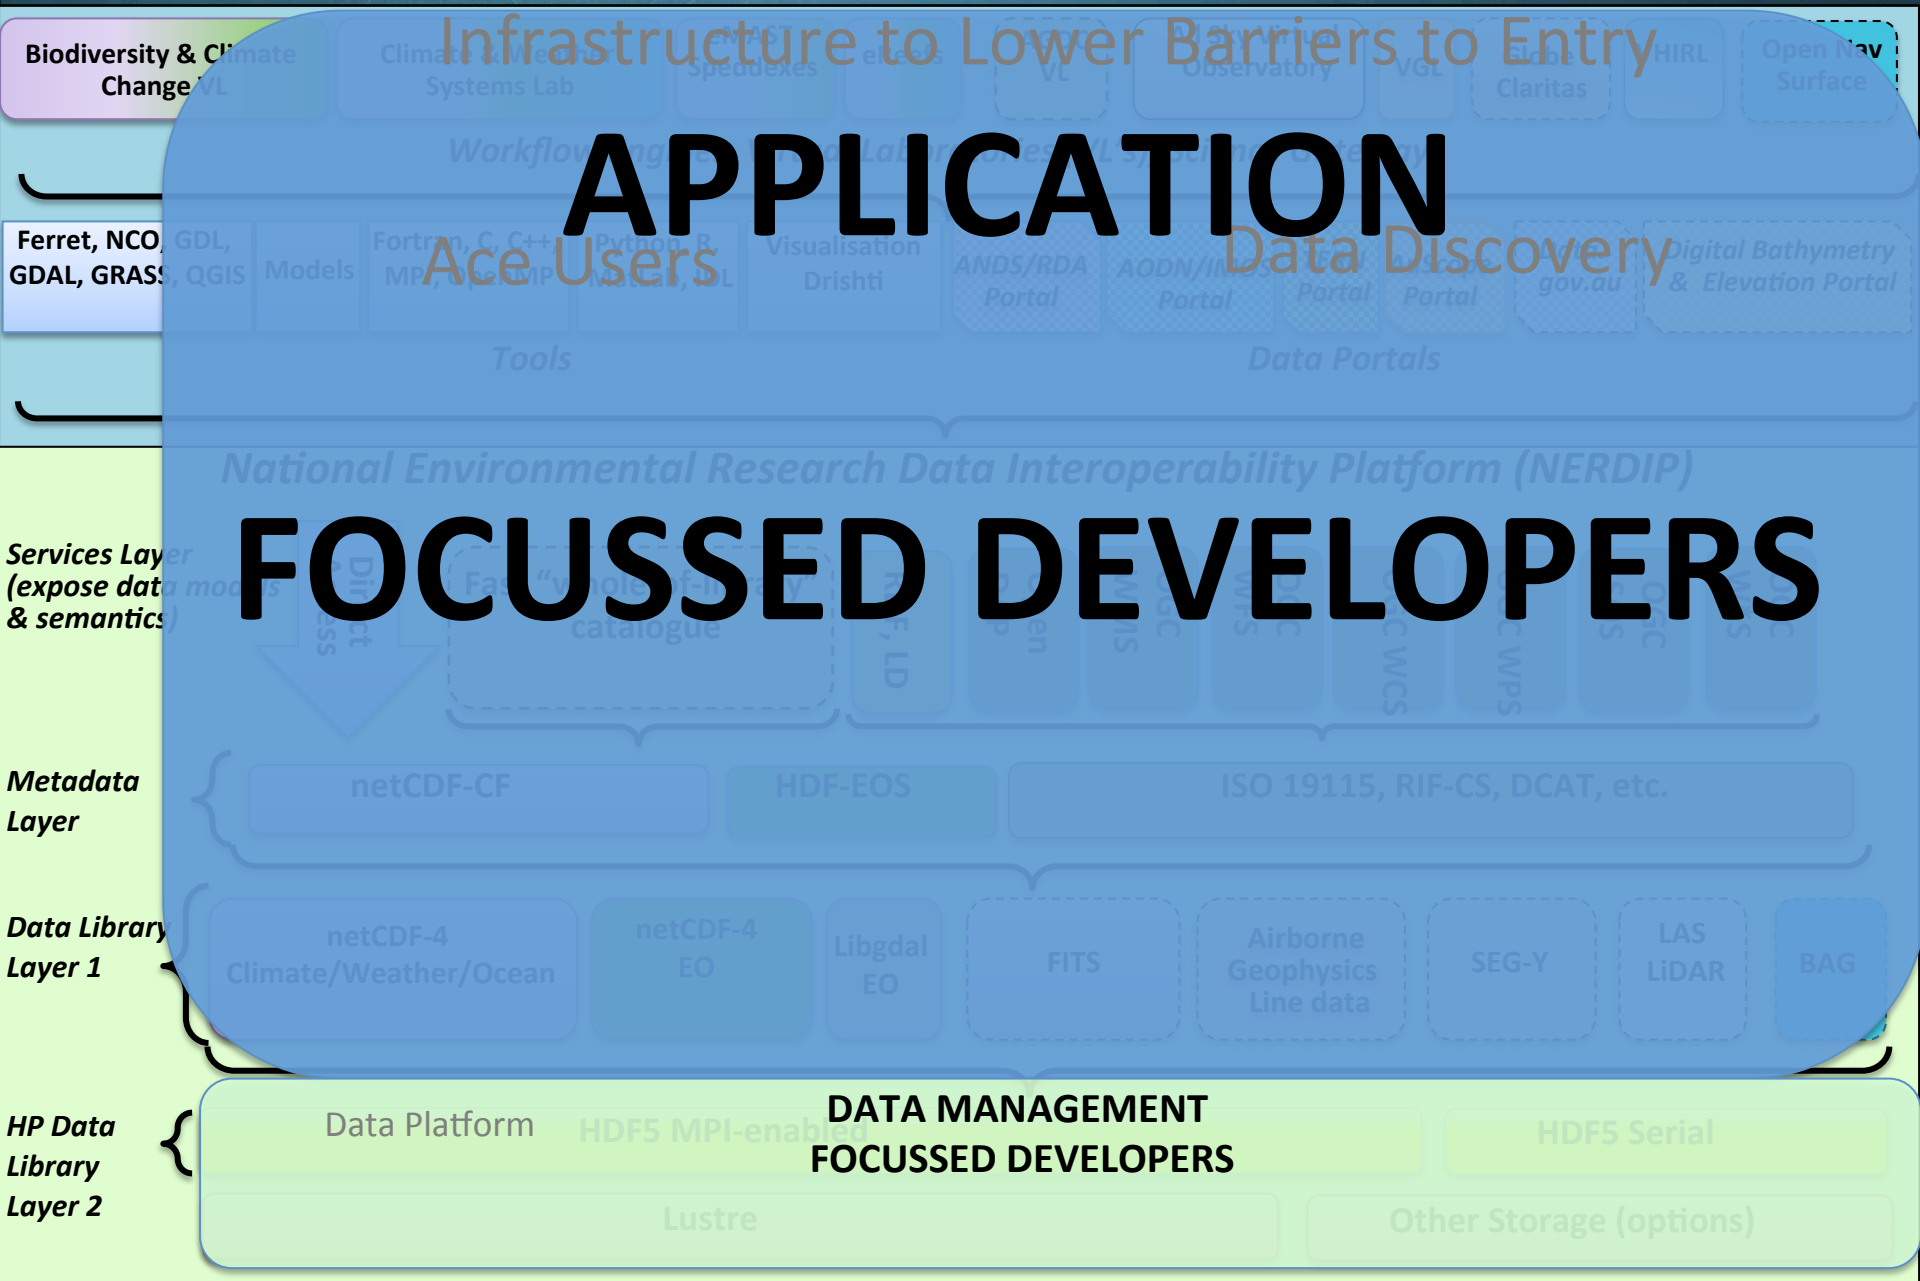

NERDIP: Applications Replicating Ways of Interacting with the Data

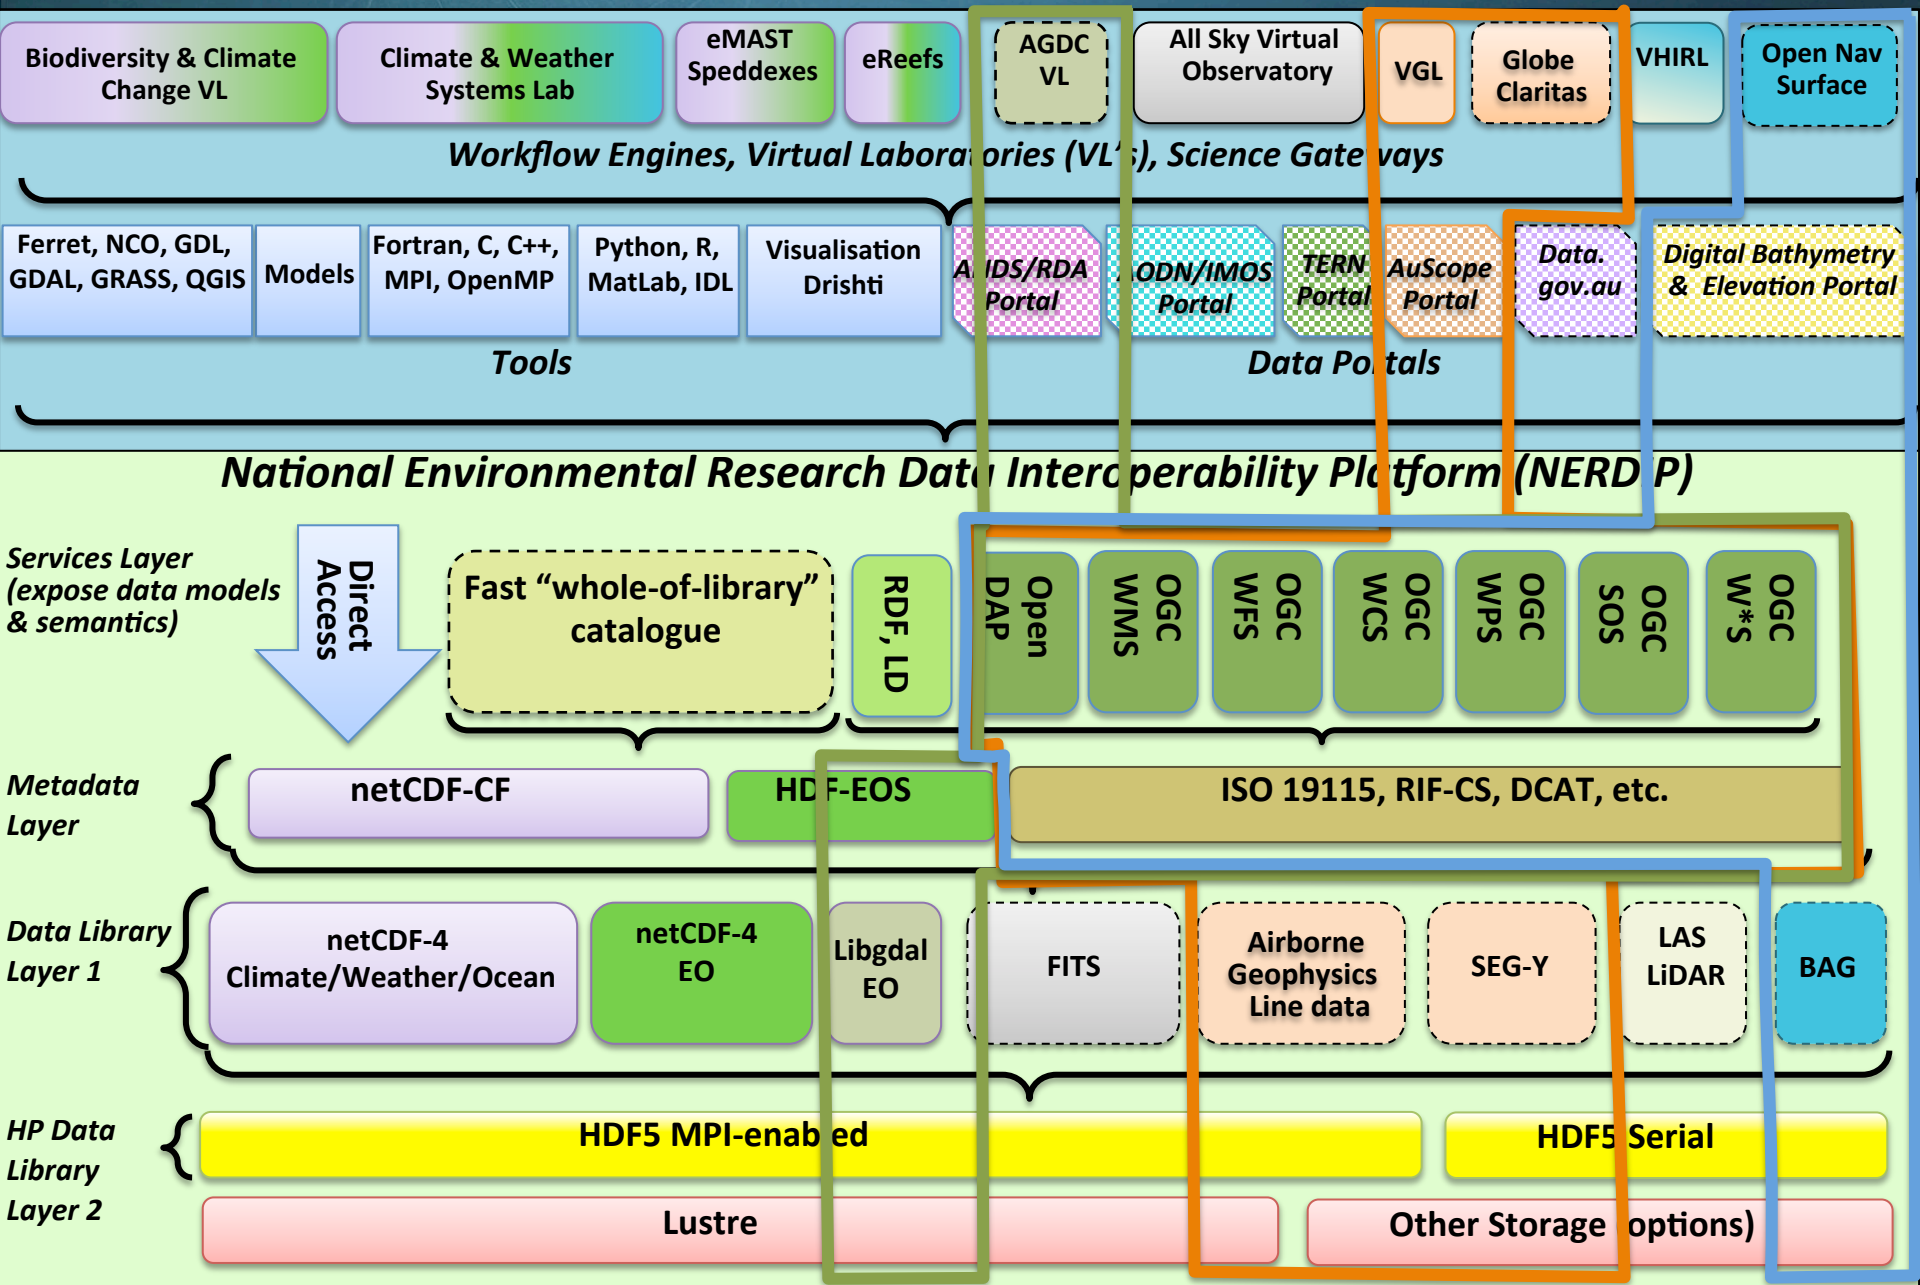

Infrastructure to Lower Barriers to Entry

APPLICATION

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

FOCUSSED DEVELOPERS

Tools ***Data Portals***

SERVICES INTERFACE

DATA MANAGEMENT

Data Platform

FOCUSSED DEVELOPERS

NERDIP: Loosely coupling Applications and Data via a Services Layer

Infrastructure to Lower Barriers to Entry

APPLICATION

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

FOCUSSED DEVELOPERS

Tools | *Data Portals*

National Environmental Research Data Interoperability Platform (NERDIP)

Services Layer (expose data models & semantics)

Direct Access

Fast "whole-of-library" catalogue

RDF, LD | Open DAP | OGC WMS | OGC WFS | OGC WCS | OGC WPS | OGC SOS | OGC W*S

DATA MANAGEMENT

Data Platform

FOCUSSED DEVELOPERS

NERDIP: the Metadata Layer

Infrastructure to Lower Barriers to Entry

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

Ace Users

Tools

Data Discovery

Data Portals

National Environmental Research Data Interoperability Platform (NERDIP)

Services Layer (expose data models & semantics)

Direct Access

Fast "whole-of-library" catalogue

RDF, LD

Open DAP

OGC WMS

OGC WFS

OGC WCS

OGC WPS

OGC SOS

OGC W*S

Metadata Layer

netCDF-CF

HDF-EOS

ISO 19115, RIF-CS, DCAT, etc.

Data Library Layer 1

netCDF-4 Climate/Weather/Ocean

netCDF-4 EO

Libgdal EO

FITS

Airborne Geophysics Line data

SEG-Y

LAS LiDAR

BAG

HP Data Library Layer 2

HDF5 MPI-enabled

HDF5 Serial

Lustre

Other Storage (options)

NERDIP: The Data Layers

Infrastructure to Lower Barriers to Entry

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

Ace Users

Tools

Ferret, NCO, GDL, GDAL, GRASS, QGIS | Models | Fortran, C, C++, MPI | Python, R | Visualisation Drishti

Data Discovery

Data Portals

ANDS/RDA Portal | AODN/IMOS Portal | Portal | Portal | gov.au | Digital Bathymetry & Elevation Portal

National Environmental Research Data Interoperability Platform (NERDIP)

Services Layer (expose data models & semantics)

Direct Access

Fast "whole-of-library" catalogue

RDF, LD | Open DAP | OGC WMS | OGC WFS | OGC WCS | OGC WPS | OGC SOS | OGC W*S

Metadata Layer

netCDF-CF | HDF-EOS | ISO 19115, RIF-CS, DCAT, etc.

Data Library Layer 1

netCDF-4 Climate/Weather/Ocean | netCDF-4 EO | Libgdal EO | FITS | Airborne Geophysics Line data | SEG-Y | LAS LiDAR | BAG

HP Data Library Layer 2

HDF5 MPI-enabled | HDF5 Serial

Lustre | Other Storage (options)

NERDIP: Infrastructure to Lower Barriers to Entry

Infrastructure to Lower Barriers to Entry

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

Ace Users

Tools

Data Discovery

Data Portals

Data Platform

NERDIP: Infrastructure to Lower Barriers to Entry

Biodiversity & Climate Change VL | Climate & Weather Systems Lab | eMAST Speddexes | eReefs | AGDC VL | All Sky Virtual Observatory | VGL | Globe Claritas | VHIRL | Open Nav Surface

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

Ferret, NCO, GDL, GDAL, GRASS, QGIS | Models | Fortran, C++, MPI, OpenMP | MatLab, IDL | Visualisation Drishti

Ace Users

Tools

ANDS/RDA Portal | AODN/IMOS Portal | Portal | Portal | gov.au | Digital Bathymetry & Elevation Portal

Data Discovery

Data Portals

National Environmental Research Data Interoperability Platform (NERDIP)

Data Platform

Direct Access

Fast "whole-of-library" catalogue

RDF, LD | OGC WMS | OGC WFS | OGC WCS | OGC WPS | OGC SOS | OpenDAP

Services Layer (expose data model & semantics)

Metadata Layer: netCDF-CF | HDF-EOS | ISO 19115, RIF-CS, DCAT, etc.

Data Library Layer 1: NetCDF-4 Climate/Weather/Ocean | NetCDF-4 EO | Libgdal EO | [FITS] | Airborne Geophysics Line data | [SEG-Y] | BAG | LAS LiDAR

HP Data Library Layer 2: HDF5 MPI-enabled | HDF5 Serial | Lustre | Other Storage (options)

NERDIP: Infrastructure to Lower Barriers to Entry

Biodiversity & Climate Change VL | Climate & Weather Systems Lab | eMAST Speddexes | eReefs | AGDC VL | All Sky Virtual Observatory | VGL | Globe Claritas | VHIRL | Open Nav Surface

Workflow Engines, Virtual Laboratories (VL's), Science Gateways

Ace Users — Ferret, NCO, GDL, GDAL, GRASS, QGIS | Models | Fortran, C, C++, Python, R, MPI, OpenMP, CUDA, OpenCL | Visualisation Drishti

Tools

Data Discovery — ANDS/RDA Portal | AODN/IMOS Portal | Portal | Portal | gov.au | Digital Bathymetry & Elevation Portal

Data Portals

National Environmental Research Data Interoperability Platform (NERDIP)

Services Layer (expose data models & semantics)

Direct Access

Fast "whole-of-library" catalogue

RDF, LD | Open DAP | OGC WMS | OGC WFS | OGC WCS | OGC WPS | OGC SOS | OGC W*S

Metadata Layer

netCDF-CF | HDF-EOS | ISO 19115, RIF-CS, DCAT, etc.

Data Library Layer 1

netCDF-4 Climate/Weather/Ocean | netCDF-4 EO | Libgdal EO | FITS | Airborne Geophysics Line data | SEG-Y | LAS LiDAR | BAG

HP Data Library Layer 2

HDF5 MPI-enabled | HDF5 Serial

Lustre | Other Storage (options)

NERDIP: Enabling Multiple Ways to Interact with the Data

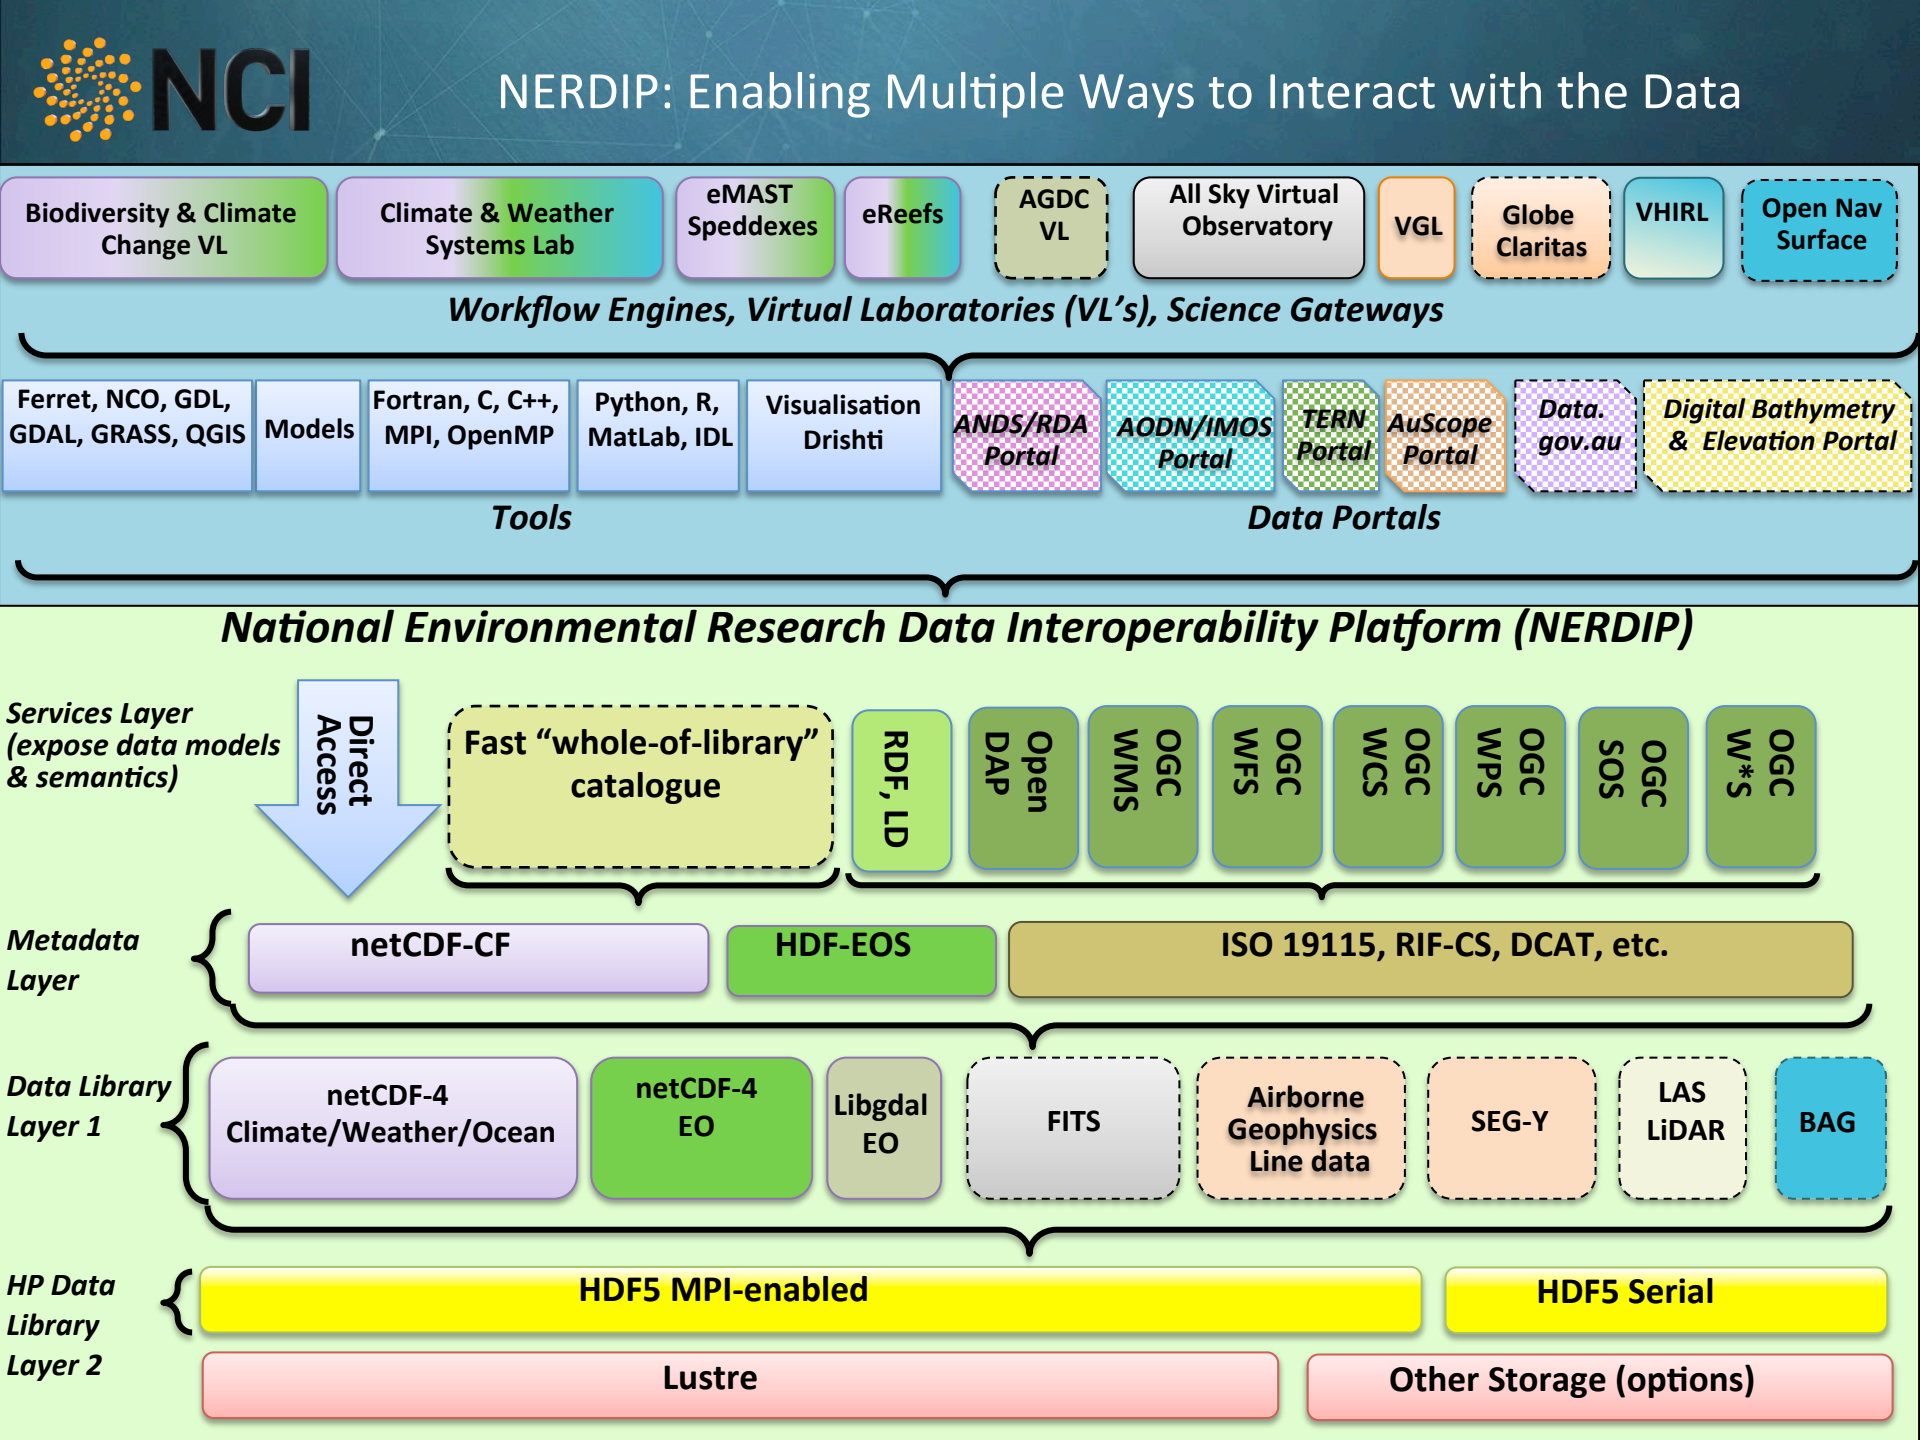

Key Messages on Big Data in the Geosciences

- Data at scales of today have to be built as shared global facilities based around national institutions.
- Domain-neutral international standards for data collections and interoperability are critical for allowing complex interactions in HP environments both within and between HPD collections
- **No one can do it alone**. No one organisation, no one group, no one country has the required resources or the expertise.
- Shared collaborative efforts such as Research Data Alliance, the Earth Systems Grid Federation (ESGF), the Belmont Forum, EarthServer, the Oceans Data Interoperability Platform (ODIP), EarthCube, GEO and OneGeology are needed to realise the full potential of the new data intensive science infrastructures
- For it now takes a 'village of partnerships' to raise a 'HPD data center' in a Big Data World

https://www.sfwa.org/wp-content/uploads/2010/06/iStock_000012734413XSmall.jpg

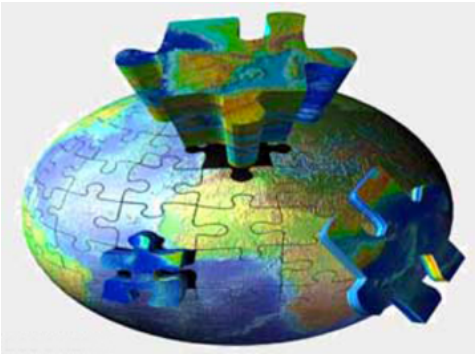

http://www.onegeology.org/