NATIONAL COMPUTATIONAL INFRASTRUCTURE

High-performance computing and data services for Australian research and innovation — challenges and opportunities: A perspective from NCI

QUESTnet 2016

Lindsay Botten
Director, NCI

6 J uly 2016

Overview

- A retrospective to characterise the national eResearch future
  - Recap of the Australian eResearch landscape 2007–16
  - What lessons have we learned?
  - Where to next?

- Advanced computing — its importance in national and international R&D
  - A global perspective
  - HPC in Australia — achievements, drivers and evolution; NCI as an example
  - The need for fit-for purpose, integrated, production-quality, expert, innovative services
    - International competitiveness
    - National benefits for research and innovation
  - The "end of the free lunch" — urgent need for
    - Investments in skills, education and training
    - Establishment of a national computational science capability

NCI

NATIONAL COMPUTATIONAL INFRASTRUCTURE

Australian e-Infrastructure 2007-16

Recap of the Australian eResearch Infrastructure Landscape

- 2007/08: NCRIS: ARCS, ANDS, AAF, NCI
- 2009-11: EIF: ANDS, NeCTAR, RDSI/RDS, Pawsey, NCI
- 2011 Roadmap never implemented
- Strategy and Implementation not refreshed/reconsidered until
  - eResearch Framework 2015-16 and NCRIS Roadmap 2016

- Characteristics of the past:
  - Structured through a technology lens, rather than research/outcome lens
  - Understandably—no mature understanding of research requirements then
  - No explicit references to service accountability, one-stop shop, integration, etc.

- What has happened during the past 8-10 years?
  - Growing maturity: eResearch methods are now integral rather than adjunct
  - Intensity and integration: research increasingly ambitious, data-intense, collaborative
  - Research challenges: increasingly global, collaborative and competitive
  - Advanced economies R&D: increasingly "*to out-compute is to out-compete*"

What lessons have we learned?

- Research-directed planning — rather than technology-directed
  - Need for stronger alignment with NCRIS, national priorities, etc.
  - Co-planned/directed by communities and institutions that “care”
  - Need for a national e-infrastructure strategy, e.g., UK e-infrastructure leadership council

- National vs Institutional planning emphasis
  - NCRIS: for national-scale initiatives beyond institutional remit
  - Not for a national collection of institutional-scale infrastructures (cost-shifting)
    - But an imperative for institutional engagement

- Holistic and balanced investments of infrastructure and expertise
  - Infrastructure spend without adequate skills/expertise emphasis dooms leverage
  - Need for a *national computational science capability (expertise)* to ensure Australia can compete and engage in the world
  - Mainstreamed investments — the opposite of grant-like provisioning of today
    - For sustainability, quality, preservation/advancement of skills/expertise

What lessons have we learned?

- Integrated, whole-of-solution structures cf. fragmentation
  - *Joined-up* solutions (hardware/software) focussed on outcomes
  - *Friction-free* environments to maximise productive research practice
  - *Accountability* for the effective delivery of outcomes is vital
  - *One phone call rather than three*

- Importance of aggregation
  - *Reduce infrastructure service costs yet optimise operational quality/innovation* — enabling greater spend on research engagement/initiative functions
  - Infrastructure aggregation with distributed researcher-facing expertise is ideal
  - Plus critical mass of expertise for national-scale and long-range priorities

- Shared responsibility — government and institutions
  - *Planning/funding* of national e-infrastructure needs to be a *shared responsibility* of national government and institutions which are beneficiaries
  - Whole-of-government approach to national priorities and responsibilities
  - Co-investment as a proxy for value

Where to from here?

- NCRIS
  - NCRIS recurrent—$1.5B over 2017-2026, NISA
  - Capital funding still to be found (Roadmap?)
- National Research Infrastructure Roadmap
  - Expert group chaired by Chief Scientist
  - Six capability working groups: environment and natural resources; advanced physics, maths and materials; health and medical research; culture and communities; national benefits and security; *underpinning infrastructure*
  - Report and recommendations Nov/Dec 2016 as (possible) input to next year's budget
- Issues confronting the Roadmap Expert Group
  - Urgent capital refresh of "hard" e-infrastructure
  - NCI peak system — 4-years old at end of 2016
  - RDSI storage/NeCTAR cloud approaching end-of-life
  - Scale of investments
    - Necessitating a whole-of-gov't approach?
    - Difficult prioritising decisions?

Australian Government

NATIONAL INNOVATION & SCIENCE AGENDA

WELCOME TO THE IDEAS BOOM

innovation.gov.au

Building world-class national research infrastructure

The initiative

The Australian Government will invest $2.3 billion over the next 10 years in cutting-edge national research infrastructure:

- $1.5 billion for the National Collaborative Research Infrastructure Strategy (NCRIS)
- $520 million for the Australian Synchrotron
- $294 million for Australia's commitment to the Square Kilometre Array (SKA)

This initiative delivers funding certainty for our research infrastructure, which is critical to maintaining our world-class research capability and capacity for innovation in our industries. In 2016, Australia's Chief Scientist will chair an expert group to undertake a road-mapping process to identify specific future research infrastructure capability requirements.

The NCRIS funded Australian Plant Phenomics Facility providesniversity the art infrastructure to address food security, nutrition, sustainable agriculture and biofuel feed-stocks. Source: Australian Plant Phenomics Facility, 2015

Advance warning of extreme weather events

High quality, cutting-edge research requires advanced, computational and data-intensive methods and infrastructure. The NCRIS-funded National Computational Infrastructure (NCI), launched in 2007, gives Australian researchers this capacity.

NCI, with CSIRO, the Bureau of Meteorology and ARC Centre of Excellence for Climate System Science, has improved the weather and climate model for Australia (ACCESS) which provides information on weather forecasts, tropical cyclones, fire weather forecasting, flood warnings and climate information. Investment in ACCESS is already paying off. It is now possible to forecast the weather three days ahead with the same level of accuracy as for two-day forecasts previously. Improved forecasting of tropical cyclone tracks allows for better planning of emergency events and decision making. Predicting the path of tropical cyclone Yasi in 2011 allowed communities, industry and emergency management agencies to make decisions and plan several days ahead of the event.

NCI has also worked with researchers to provide geological modelling and a virtual astronomical observatory, as well as a simulation of blood flow around the human body.

Why this is important

NCRIS drives research excellence and collaboration between 35,000 researchers, government and industry to deliver practical outcomes. Ongoing funding will enable continued operation of critical super-computer capacity and world-class research infrastructure in areas such as nanofabrication, food production, health, environment and sustainable cities.

The SKA is a unique opportunity to co-host a globally significant mega-science project in Australia. The SKA will be the largest and most capable radio telescope ever constructed. It will deliver significant economic, science and technology benefits and advance human capital.

Around 4,000 researchers access the Australian Synchrotron — an accelerator technology platform that reveals the innermost structures of materials. The Synchrotron delivers benefits to diverse scientific and industrial communities, with applications in health, energy manufacturing, food, environment and bio-security.

Implementation

Synchrotron and SKA funding commences in 2016–17. NCRIS funding is ongoing from 2017–18.

Where to from here?

- eResearch Framework 2015-16 (Roadmap input) addresses most "lessons learned"
- Structure — two principal (accountable) delivery entities + software + connectivity
  - National Computation System — deriving from peak facilities—NCI, Pawsey
    - *Deploying computational power to extend boundaries of enquiry*
    - Peak compute, big data serving high-impact or merit-based/prioritised needs
    - Coalescence into a single capability inevitable: single governance/administration
  - Australian Research Data System — genesis in ANDS, RDS and NeCTAR
    - *Connecting and using data assets to better enable new kinds of research*
    - Improving accessibility, value and use of quality research data
    - Implemented as a co-operative system using cloud (but with aggregation)
  - Research Community Platforms — genesis in NeCTAR VL program
    - *Creating research applications that exploit innovative*
    - Software investment, implemented through other infrastructures
  - Connectivity — deriving from AARNet and AAF as now
    - Frictionless connectivity enable collaboration,
    - Implemented as service providers

NATIONAL COMPUTATIONAL INFRASTRUCTURE

Focus on Advanced Computing

The current state of Advanced Computing — a global perspective

- HPC —a 4-6 year advantage over what is available in the commodity marketplace
- Ubiquitous technology on which all science/tech. research & innovation depend
- **Today, to out-compute is to outcompete** (IDC, US Council for Competition)
  - IDC report: 97% of US companies could no longer survive or be competitive
  - UK (Tildesley Report, 2011) deeply concerned about "loosing touch"
- Massive investments—US, China, Japan, EU, UK
  - US Department of Energy largely drives international HPC hardware/software agenda
  - National Strategic Computing Initiative—President Obama (July 2015)
    - HPC technologies for economic competitiveness and scientific discovery
    - Public-private partnership; whole-of-government approach
    - Accelerate exascale delivery; establish enduring ecosystem; DoE/DoD/ NSF
    - Current DoE invest. (10-200PF): CORAL $300M, NERSC Cori $74M, Trinity $174M
- Australian Context (2016): National Research Infrastructure Roadmap
- Future: a National Computation System for Australia
  - About the nexus of HPC simulation and big data
  - Need for major investments in the ecosystems of skills — incl. computational sciences
  - About maximising return on investment — academic research, agencies, industry R&D

NCI

Summary of 100 peak international systems (Top500.org, November 2015)

Country	Number	Total Rmax (Pflops)	Top500 Ranks	GDP-nominal (US$ Billion)	Tflop/ GDP($B)	Popul'n (M)	Tflop/ Pop'n(M)
USA	35	115.31	2,3,5,6,10,12,13,14I,15,16,18,20,21,27,38,40,42/43,45I,51,54/55,58,63,65,72,73,79,82,87,94I,97,99I	18,527	6.22	323.1	356.9
China	11	46.77	1,26,34,57,66/67,76,81,90/91,98	10,356	4.52	1375.5	34.0
Japan	12	28.12	4,22,25,31,35/36,60,61,62,75,78,84,85M,89	4,602	6.11	126.8	221.8
Germany	10	23.65	8,11,23/24,50,56,64M,77,83,100	3,874	6.10	81.5	290.2
United Kingdom	6	8.20	41,46/47M,49,71,88	2,950	2.78	65.1	126.0
Saudi Arabia	2	7.79	9,32I	746	10.44	32.2	241.8
Switzerland	2	7.21	7,92	703	10.26	8.3	868.8
France	5	7.20	33I,44,53,70,74	2,833	2.54	64.5	111.6
Italy	2	4.97	19I,37	2,147	2.31	60.7	81.9
South Korea	2	4.80	29/30M	1,410	3.40	51.6	93.0
Russia	2	2.80	36,95	1,860	1.51	146.5	19.1
Poland	2	2.68	39,80	547	4.90	38.5	69.6
Australia	**2**	**2.08**	**68,86**	**1,442**	**1.44**	**24**	**86.7**
Czech Republic	1	1.46	48,	205	7.12	10.5	139.0
Sweden	1	1.40	52,	570	2.46	9.8	142.9
Finland	1	1.25	59,	272	4.60	5.5	227.3
Netherlands	1	1.09	69,	880	1.24	17	64.1
Spain	1	0.93	93,	1,406	0.66	46.4	20.0
India	1	0.90	96,	2,051	0.44	1286.2	0.7
Totals	**99**	**268.61**					

Note:

- Striking correlations between GDP / supercomputer power (Tflops) and number of systems in production in research and industry
- List will change markedly in the coming year with the surge of new investments (China, US, UK, EU)
- Australian HPC capability dwindling with time
  - NCI #24 in Nov 2012
  - Pawsey #41 Aug 2014

Australian Peak HPC (Research) Systems in 2016

- **NCI—Raijin** (November 2012) — NCRIS Super Science
  - Fujitsu Primergy (x86, Sandy Bridge)
  - Max performance: 1.05 PF (peak 1.20 PF)
  - Sustained performance: 1,600 kSFRB
  - 57,472 Intel Xeon (Sandy Bridge) cores; 160 TB memory
  - Filesystem: 10 PB (150 GB/sec); Interconnect: 200+ Tbit/sec

- **Pawsey Centre—Magnus** (August 2014) — NCRIS Super Science
  - Cray XC40 (x86, Haswell)
  - Max performance: 1.10 PF (peak 1.50 PF)
  - Sustained performance: 1,140 kSFRB
  - 35,712 Intel Xeon (Haswell) cores; 93 TB memory
  - Filesystem: 3 PB (70 GB/sec);

- **VLSCI—Avoca** (June 2012) —Victorian Government funded
  - IBM Blue Gene/Q
  - Max performance: 0.72 PF (peak 0.84 PF)
  - 65,536 cores; 64 TB memory
  - Likely to be decommissioned in Dec. 2016

Other Australian HPC Systems

- Other Peak Systems
  - BoM Australis
    - 2016 — Cray XC40—1.6 Pflop, 1,600 kSFRB (sustained performance: 1.0 x Raijin)
    - 2018—major upgrade to 7/8 Pflop total, 4,500 kSFRB (sustained performance 2.6 x Raijin)
- Mid-range systems
  - CSIRO
    - Pearcey (Dell x86 cluster, ~5,000 cores)
    - Bragg (GPU and Xeon Phi cluster)
    - Ruby (SGI UV 3000 — shared memory system)
    - Bowen (High-performance cloud)
  - FlashLite
    - UQ, special large memory system running ScaleMP, 1,500 cores
  - Massive
    - Monash/Synchrotron, ~2,000 cores + GPUs, visualisation, real-time imaging
  - Other institutions
    - Numerous 1,500 – 4,000 core facilities in universities—USyd, UNSW, UWA, Adel, Melb, N'cle, …
  - NeCTAR Cloud
    - ~30,000 cores, ~20% HPC standard, some running virtual clusters

NATIONAL COMPUTATIONAL INFRASTRUCTURE

What does Advanced Computing contribute to national R&D?

HPC supporting the full gamut of research

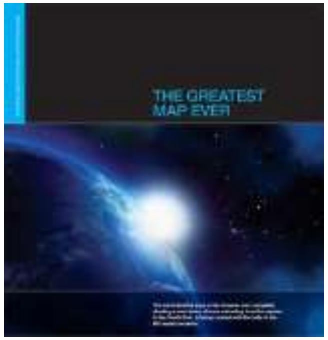

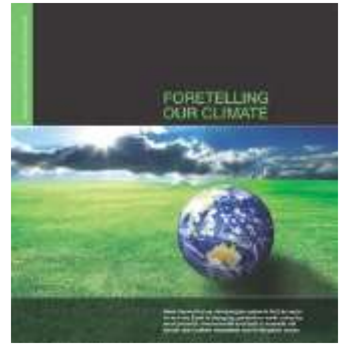

pure

- Fundamental sciences
- Mathematics, physics, chemistry, astronomy,
- ARC Centres of Excellence (ARCCSS, CAASTRO, CUDOS, CNBP)

strategic

- Research with an intended strategic outcome
- Environmental, medical, geoscientific, etc.
- e.g. food security, cancer genomics, geosciences, energy

applied

- Informing public policy; demonstrable economic impacts
- Climate variation, Next-gen weather, Seasonal forecasting, Earth observation services (GA, BoM, CSIRO, ARCCSS)

industry

- Supporting industry and innovation
- e.g., ANU/UNSW start-up, Lithicon, sold for $76M to US company FEI in 2014;
- AMSA/DHI large vessel drifting

NCI at a glance in 2015-16

Profile

- Serves: 31 x universities, 5 x science agencies,
  8 x NCRIS capabilities; 3 x MRIs; industry
- Highly-regarded: *Net Promoter Score* --- 4.6/5 (Survey 2015)

Scale (2015)

- 1.2 petaflop (57,472 core) supercomputer;
- Supercomputer-class private cloud (3,200 cores)
- ~450 Million compute hours; ~5.5 million jobs p.a.
- 36 PB Lustre storage (up to 150 GB/sec b/w); ~430M files
- 500+ projects; ~4,000 users, 1,500 with logins
- Shares: 55% univ./MRIs; 44% Gov't agencies; 1+% industry
- Access on merit (excellence) and priority (significance)
- 60 expert staff: operations (30), engagement/initiatives (20+)
- Australia's most extensive research s/ware library (~$300K pa)
- 1.8 MW of electricity; $2.5M pa. utilities

Collaborations (International and Industry)

- ESGF, UK MetOffice, RIKEN, NOAA, NASA, USGS, EUMETSAT, …
- Industry: FEI,DHI,ResMed,Fujitsu,Mellanox,SGI,DDN,Amazon,..

Excellence dimensions

- Supports 200+ ARC/NHMRC-funded CoEs, projects, fellowships, etc. — ~$58M pa, $270M aggr. value (2016)
- About one-third of recent AAS Fellows are NCI users
- 500+ journal articles p.a. acknowledge NCI support

National Benefits

- Via Government agencies (CSIRO, BoM, GA, AMSA)
- Holding internationally significant data collections
  - CMIP5 Climate Collection
  - Landsat (GA)
  - Australasian hub for EU Sentinel earth obs. (GA)
- Supporting high-profile medical research (MRIs)
  - National partner in NH&MRC AGHA (Genomics)
- Developing future industry services —govt. agencies
- Strong alignment with 7/9 national priorities

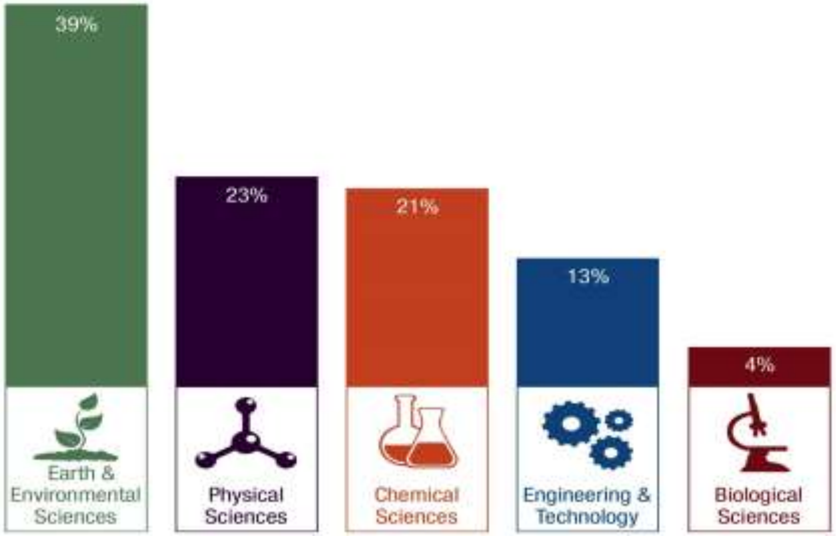

The greatest map ever made

Led by Nobel Laureate, Professor Brian Schmidt, Australian astronomers are using NCI to carry our the most detailed optical survey yet of the southern sky. The project involves processing and storing of many terabytes of optical telescopic images, and has led to the discovery of the oldest star in the universe.

Designing better engines for renewable biofuels

ARC Future Fellow Professor Evatt Hawkes and his team at UNSW are using a highly scaling code that uses up to half of the NCI supercomputer to run the world's first simulation of hydrocarbon combustion. They aim to develop a cheap model to be used by industry to design diesel engines that use less fuel and produce fewer emissions .

Ocean modelling breakthroughs

Using the high performance computing capabilities of the NCI, Professor Matthew England (UNSW) and Associate Professor Andy Hogg (ANU) from the ARC Centre of Excellence for Climate System Science are discovering how previously unknown small ocean eddies drive vast ocean circulation processes in the Southern Ocean around Antarctica.

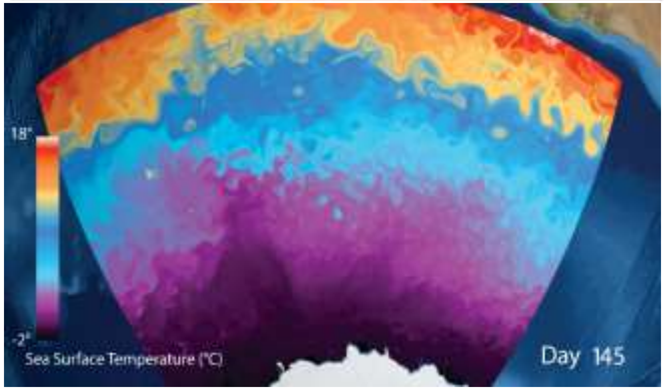

Understanding the and structure / stability of nanostructures

Using the power of NCI's Raijin, CSIRO's Dr Amanda Barnard, winner of the Foresight Institute Feynman Prize in 2015, has spearheaded the understanding of the self-organising properties of diamond nano-particles using computational methods. Such discoveries have wide application in self-cleaning surfaces, fuel cells, and the development of a brain tumour chemotherapy treatment in use at UCLA.

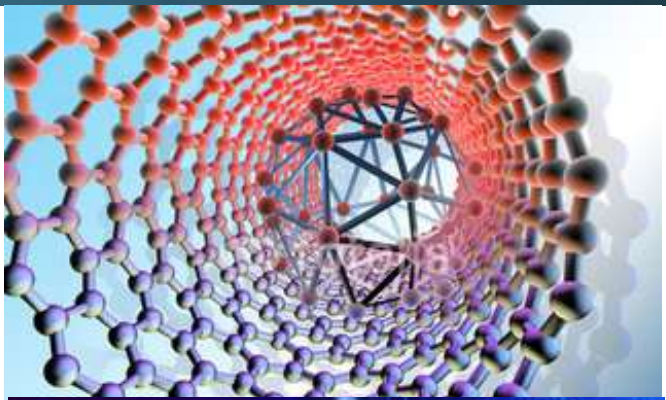

Producing a world leading genome database

Comprising data from thousands of genomes from healthy elderly Australians, sequenced by Garvan, the new Medical Genome Reference Bank ($10M from NSW Health) will be processed, analysed and stored within NCI's HPC environment in 2016-17. The MGRB will give doctors and scientists a comprehensive database of genetic data, for diagnosis, treatment and research in the fight against disease.

New tools to diagnose auto-immune diseases

NCI is providing the computational power for Professor Carola Vinuesa from the Centre for Personalised Immunology at ANU—a NH&MRC Centre of Excellence— to store, analyse and compare the genomes of patients with auto-immune diseases. Diseases like lupus vary significantly at the genetic level, and discovering those variants opens the door to much more targeted treatments in the future.

Creating supercharged crops

NCI is supporting Prof. Jill Gready (ANU John Curtin School of Medical Research) in industry-funded research that develops 'super-charged' crops that can perform Earth's most important chemical reaction—photosynthesis— with far greater efficiency. Prof. Gready uses the supercomputer to study how Rubisco, the key enzyme in photo-synthesis, can be made more efficient to yield greater food security.

Doubling solar cell efficiency

Professor Michelle Coote from ANU and the ARC Centre of Excellence for Electromaterials Science used NCI's supercomputer to investigate how particular coloured molecules used in solar cells might be stabilised while maintaining a high efficiency. Professor Coote found five compounds which led to the manufacture of a solar cell with a two-fold increase in efficiency.

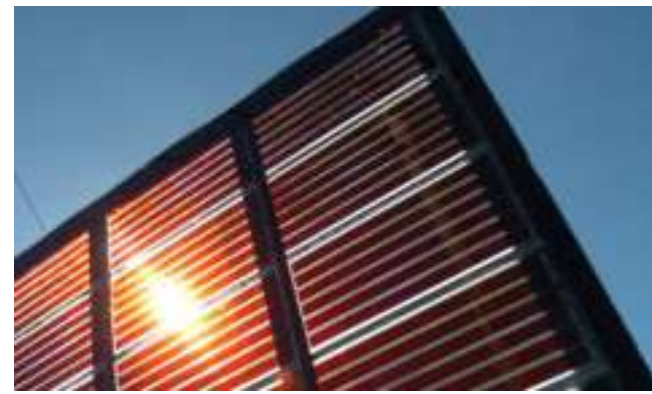

4D Earth modelling for industry

To extract resources from between grains of sand within sedimentary rocks, industry needs reliable data and models about how these sedimentary systems work. NCI provides essential HPC resources for the $5.4M ARC Industry Transformation Research Hub for Basin Geodynamics and Evolution of Sedimentary Systems (GENESIS), led by Professor Dietmar Muller of the University of Sydney.

Research Highlights/Impact

Predicting the unpredictable

Australia's weather and future climate are predicted using the ACCESS model—developed by BoM, CSIRO, and ARCCSS—operating on time scales from hours/days, to centuries. Collaborations with NCI and Fujitsu are leading to performance increases of up to 40% and model scalability to 1000s of cores for more accurate predictions (extreme weather, seasonal prediction) on BoM's next-gen production system.

Unlocking the Landsat Archive

NCI is enabling researchers at Geoscience Australia to 'unlock' decades of Landsat earth observation satellite images of Australia since 1979. A one petabyte *data cube* has been generated by processing and analysing hundreds of thousands of images, yielding important insights for water/land management decision-making and policy, with benefits for the environment and agriculture.

Delving for natural resources

NCI supports the Australian digital rock start-up company, Lithicon, sold for $76M to a US multi-national (FEI) in 2014. The technology, which is at the forefront of the efficient recovery of oil and gas resources worldwide, is a fusion of a new generation of micro-CT scanner, and advanced analysis and simulations, using ANU/UNSW IP based on composite material theory, that require supercomputer performance.

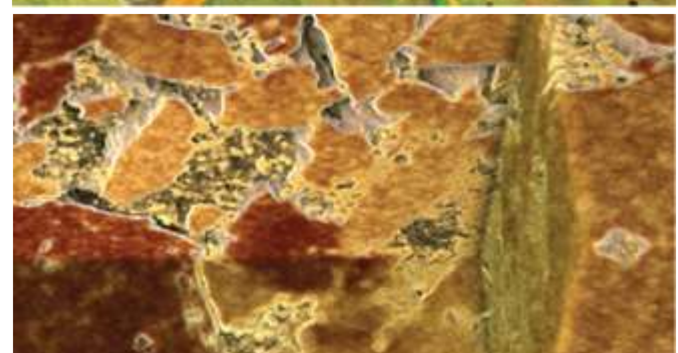

ACCESS
The Australian Community Climate and Earth-System Simulator

Earth system model

Economic benefits:

- Extreme weather/climate (5% of GDP)
- Reduce agriculture losses: ~$2 billion)
- Improved seasonal prediction (>$1 billion p.a.)

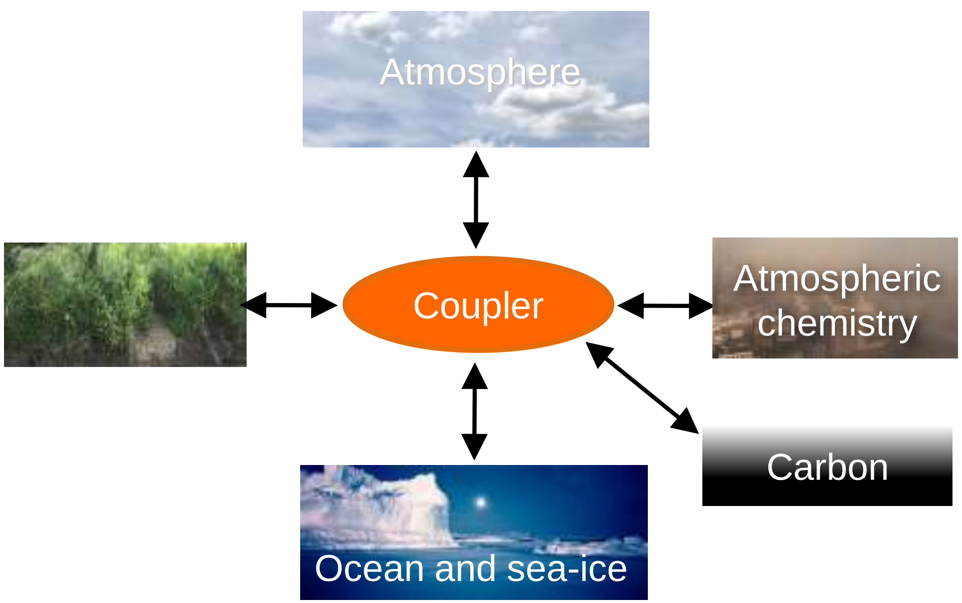

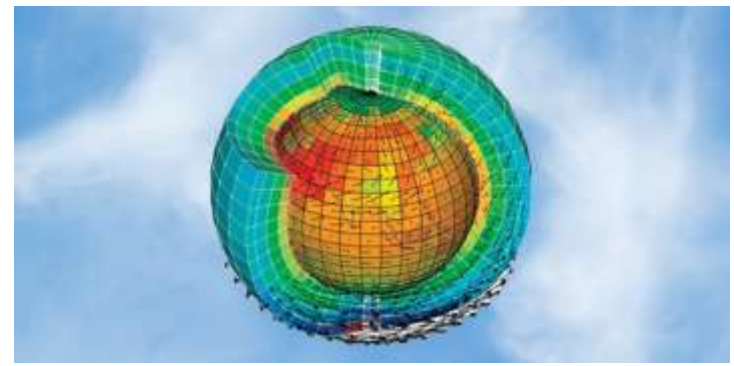

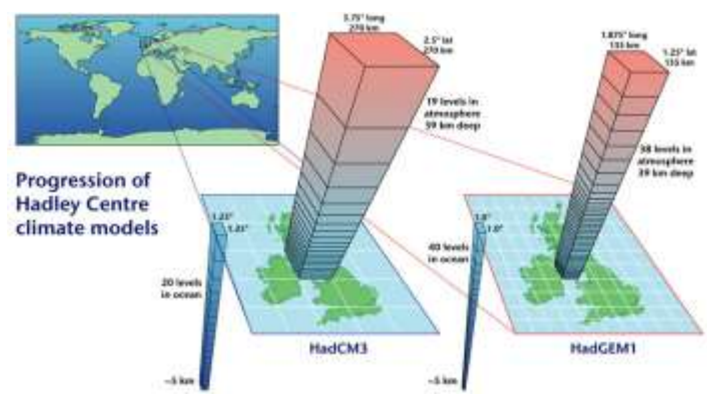

HPC driver: Model resolution—Weather/Climate Prediction

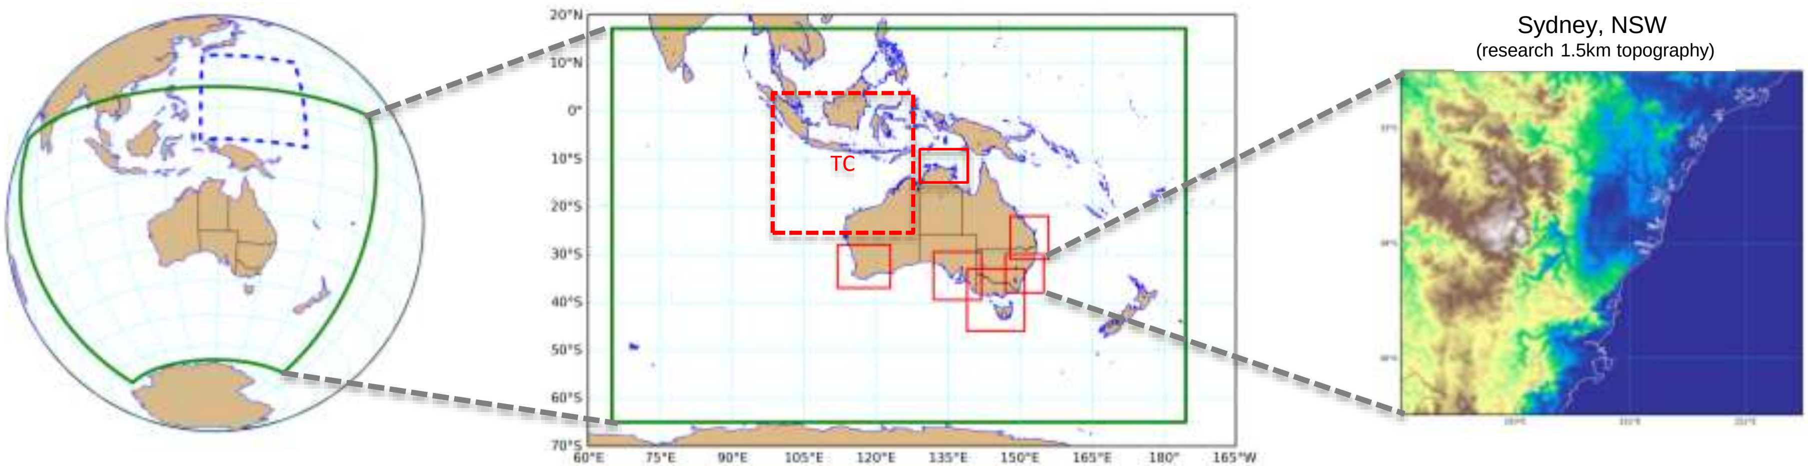

2 x daily 10-day & 3-day forecast
40km Global Model

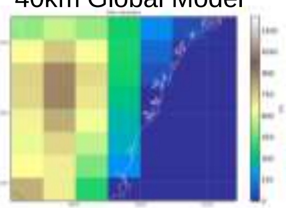

Increasing model resolution
for improved local information

4 x daily 3-day forecast
12km Regional Model

Future model ensembles for likelihood of significant weather

4 x daily 36-hour forecast
4km City/State Model

2013

Model Topography of Sydney, NSW

2020

2 x daily 10-day & 3-day forecast
12km Global Model

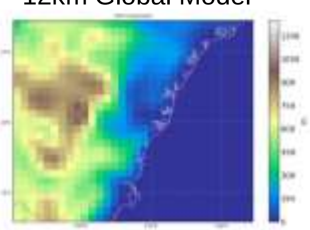

8 x daily 3-day forecast
5km Regional Model

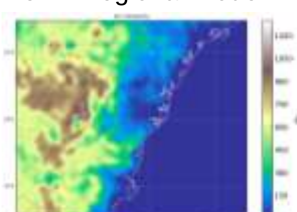

24 x daily 18h or 36h forecast
1.0km City/State Model

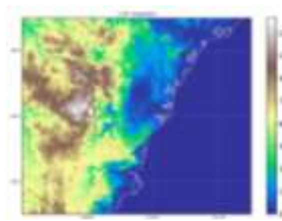

Projection of Nominal Modelling Resolutions for Future Computing Systems

Subject to sufficient investments in High Performance Computing (Courtesy: BoM)

Optimisation/scaling: faster results, better RoI, more ambitious research

- Weather codes optimisation:
  - UKMO, BoM, STFC Daresbury, NCI Collaboration
  - Joint work also with BoM and Fujitsu
  - UM: 40% performance improvement with parallel I/O taken up by UKMO

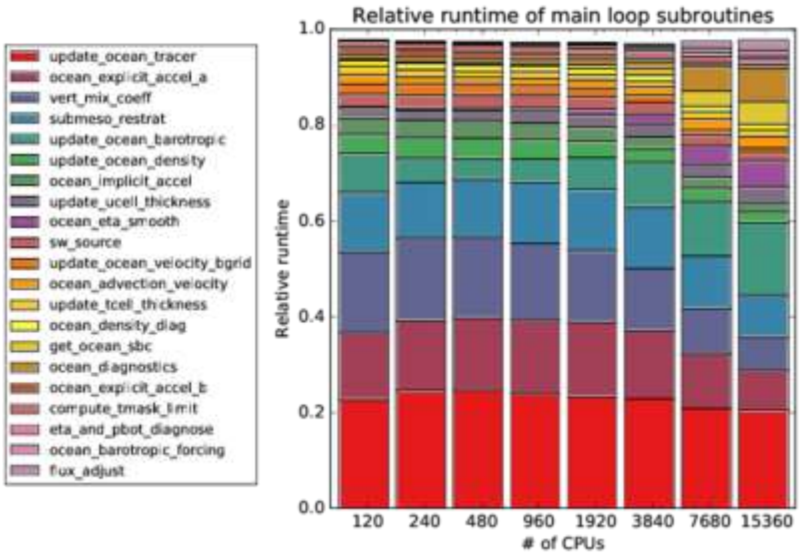

- Earth System Science Optimisation
  - 30 % performance gains in ocean model taken up by NOAA
  - Code now runs well at up to 15, 360 cores, previously performing poorly at 960 cores.
  - Resolution improvement to 0.1 degree (from 0.25 degree) enables more ambitious science — eddies resolved

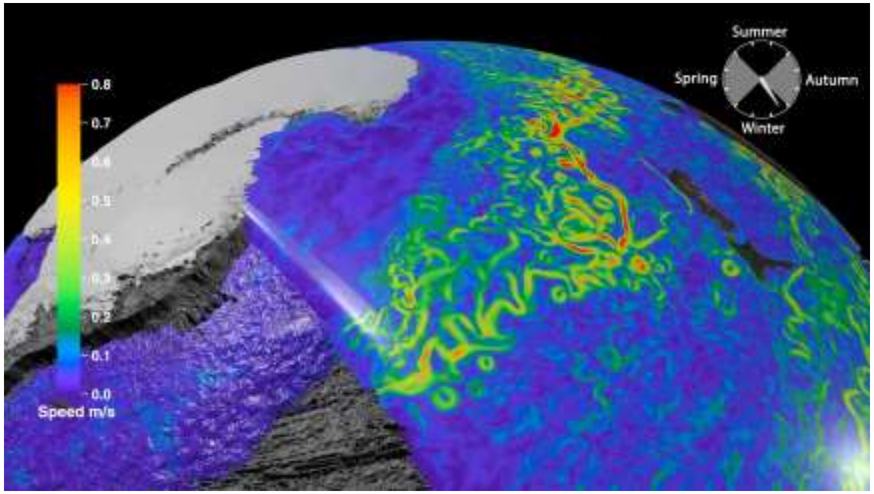

Data-intensive Scientific Computing Environments

Library of Congress: 10 TB

10 PB = 10,000 TB

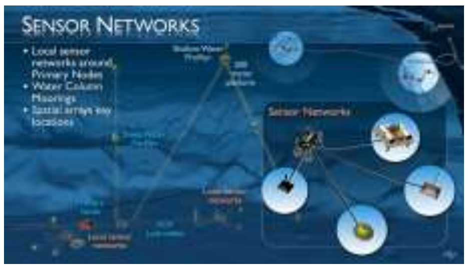

Big Data vs High-Performance Data Analytics

- HPDA: "big data" meets HPC
- Driving HPDA: data from massive instruments and sensor networks, massive simulations, need to quantify of uncertainty, etc…
- Focus on <u>content</u> and <u>HPC analytics</u> on content
- NCI's focus in on HPD/HPDA

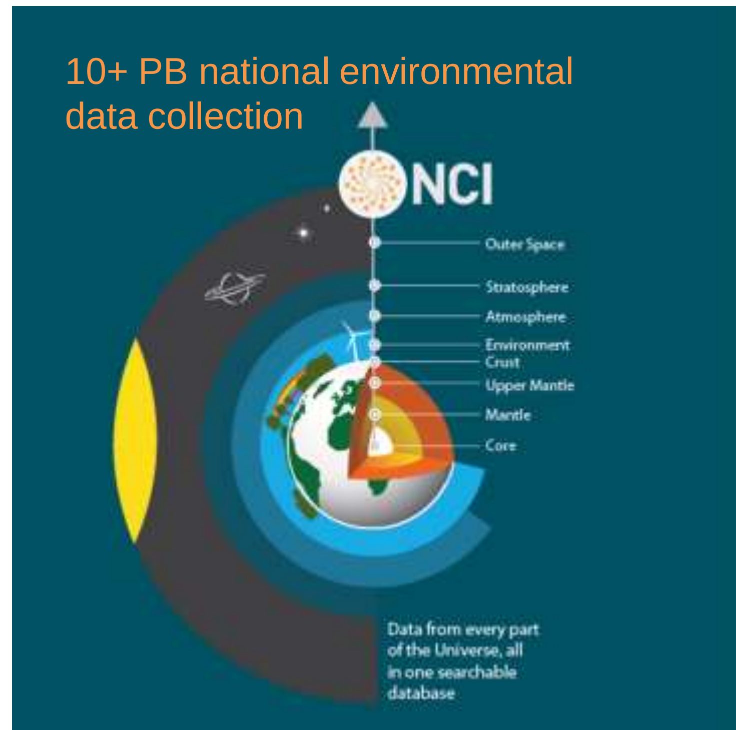

NCI National Environment Research Collections

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

http://geonetwork.nci.org.au

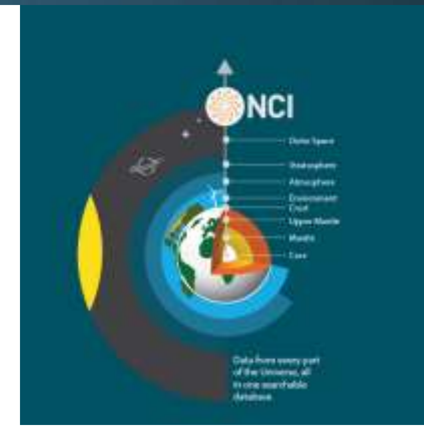

Data Collections	Approx. Capacity
CMIP5, CORDEX, ACCESS Models	5 Pbytes
Earth Obs: LANDSAT, Sentinel, MODIS, INSAR	2 Pbytes
Digital Elevation, Bathymetry, Onshore/Offshore Geophysics	1 Pbytes
Seasonal Climate	700 Tbytes
Bureau of Meteorology Observations	350 Tbytes
Bureau of Meteorology Ocean-Marine	350 Tbytes
Terrestrial Ecosystem	290 Tbytes
Reanalysis products	100 Tbytes

Australian Geoscience Data Cube

- Data cube: 300,000 Landsat scenes (spatial/temporal) allowing flexible, efficient, large-scale in-situ analysis
- Spatially-regular, time-stamped, band-aggregated tiles presented as temporal stacks.

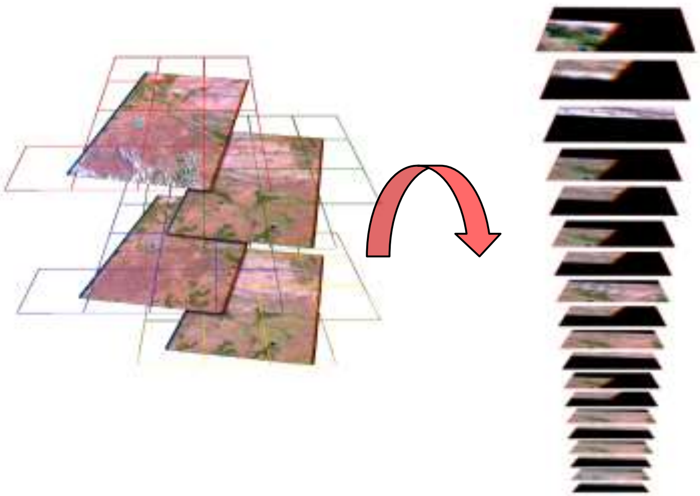

Spatially partitioned tiles   Temporal Stack

Continental-Scale Water Observations from Space

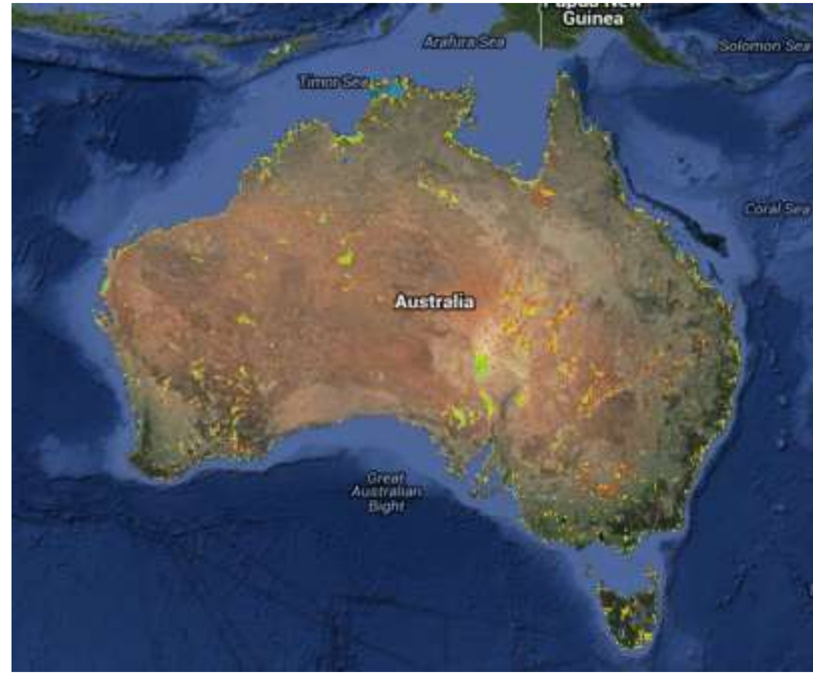

WOFS water detection

- **27 Years** of data from LS5 & LS7(1987-2014)
- **25m** Nominal Pixel Resolution
- Approx. *300,000* individual source ARG-25 scenes in approx. 20,000 passes
- Entire 27 years of 1,312,087 ARG25 tiles => **$93x10^{12}$ pixels** visited
- **0.75 PB** of data
- **3 hrs** at NCI (elapsed time) to compute.

Australian Government Geoscience Australia   CSIRO   NCI NATIONAL COMPUTATIONAL INFRASTRUCTURE

EU Copernicus Sentinel Earth Observation

- Six families of satellites: Sentinels 1-6 progressively from 2014
  - Monitoring of land, ocean, vegetation, soil, altimetry, etc.
- Australia to provide the *regional data access and analysis hub*
- Consortium: GA, CSIRO, State Govt. agencies (WA, NSW, Qld)
- NCI as the Australian Sentinel Hub
- Implementation Partners: NCI, aarnet (Australia's Academic and Research Network)
- Exploitation Collaborators: BoM, RDS, CRC-SI

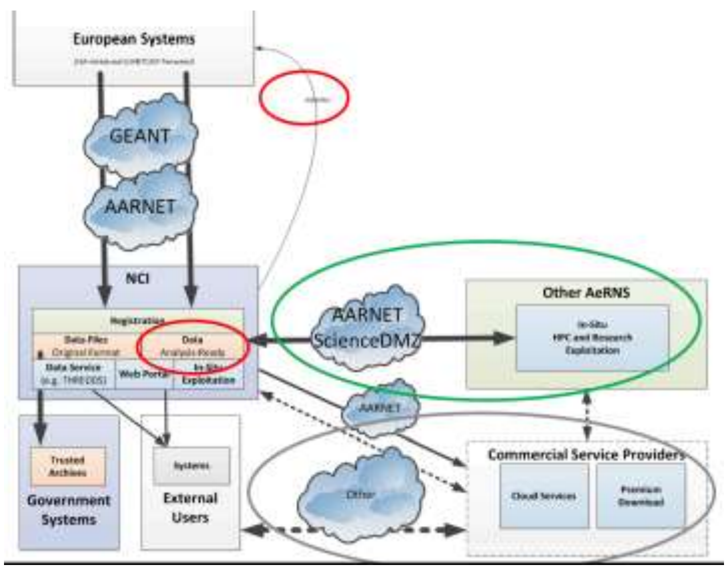

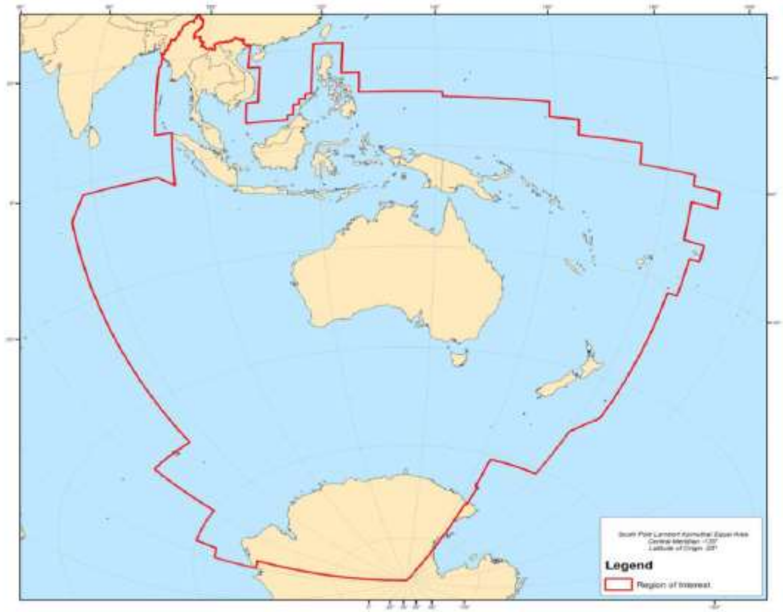

Industry Engagement

Direct Engagement ⟷ **Indirect Engagement**

Commercial Access	(IT) Industry Collaboration	R&D Collaboration	Infrastructure Parent of Innovation
FEI	FUJITSU	FUJITSU	Australian Government Bureau of Meteorology
DHI	DataDirect NETWORKS IN MOTION	Linux	ESGF
	sgi	lustre	Australian Government Geoscience Australia
	Mellanox TECHNOLOGIES	intel	LITHICON
	amazon web services		Australian Government Geoscience Australia
			CSIRO








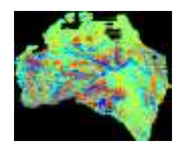



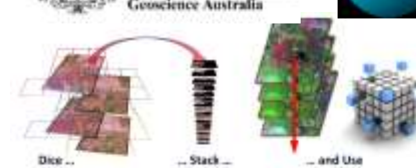

NATIONAL COMPUTATIONAL INFRASTRUCTURE

The architecture of the future HPC environment

Service Portfolio of a contemporary HPC facility— NCI as it is today

Services and Technologies portfolio

- Comprehensive/integrated to meet all needs
- High-performance: HPC, Cloud, storage; user services
- Deep engagement and collaboration with vendors
- Expert environment tailored for specialised needs
- World-class operational capability for Australia
  - e.g., greatest number of single institution contributions to international open source Lustre parallel filesystem software

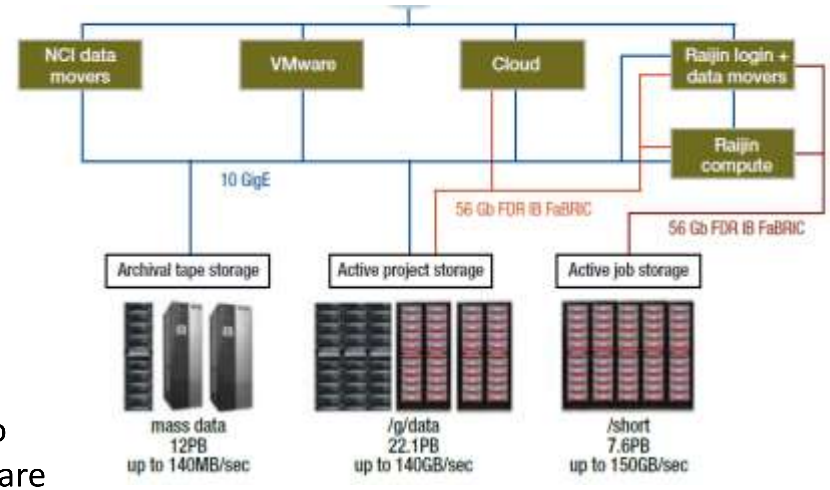

Research Engagement/Initiatives portfolio

- HPC/HPD innovation; virtual environments (incl. virtual desktops); data collection management and services; visualisation
- Provides “special sauce”/“glue” for valued, friction-free services
- General capabilities evolve from specialised requirements
- Shaped by requirements of NCRIS, partners, MRIs industry
- Recognised in international vanguard of data-intensive computation

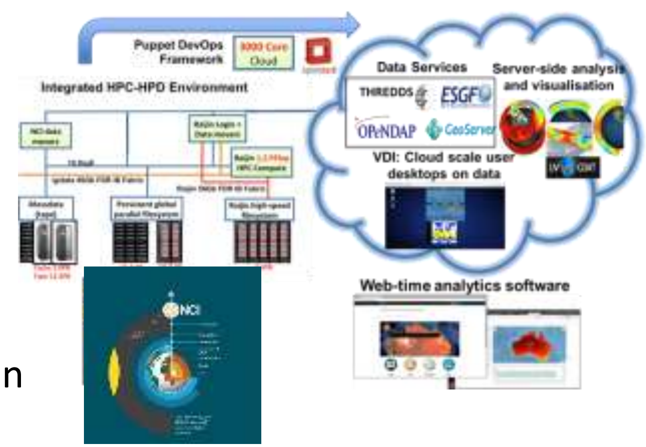

NCI: Major Infrastructure

- **Supercomputer—Raijin** (November 2012) — NCRIS Super Science
  - Fujitsu Primergy (x86 cluster,)
  - Max. perf.: 1.05 PF (peak 1.20 PF); Sustained Perf.: 1,600 kSFRB
  - 57,472 Intel Xeon (Sandy Bridge) cores; 160 TB memory
  - Filesystem: 10 PB (150 GB/sec); Interconnect: 200+ TB/sec
- **HPC Cloud—NeCTAR and Tenjin**
  - Dell OpenStack Cloud (x86, Xeon Sandy Bridge)
  - Max. perf.: 67 TF
  - 3,200 Xeon (Sandy Bridge) cores; 25 TB memory; 160 TB of SSD
- **Global integrated storage (highest performance in Australia)**
  - 36 PB of Lustre disk (3 generations: 120/70/50 TB/sec)
  - 40 PB if tape for archive purposes
- **Data Centre — purpose build 900 sq. m. (November 2012)**
  - Power capacity (after 2014-15 upgrades)
    - 4.5 MW capacity raw; 1 MW UPS; 2 x 1.1 MVA generators
  - Cooling in two loops
    - Server:
      - 2 x 1.8 MW chillers; 3 x 0.8 MW “free cooling” heat exchangers
      - 18^o C; 75 l/sec pumping rate
    - Data: 3 x 0.5 MW Carrier chillers; 15^o C

Inside NCI's 900 m^2 machine room

Frictionless infrastructure: Compute/Storage Integration

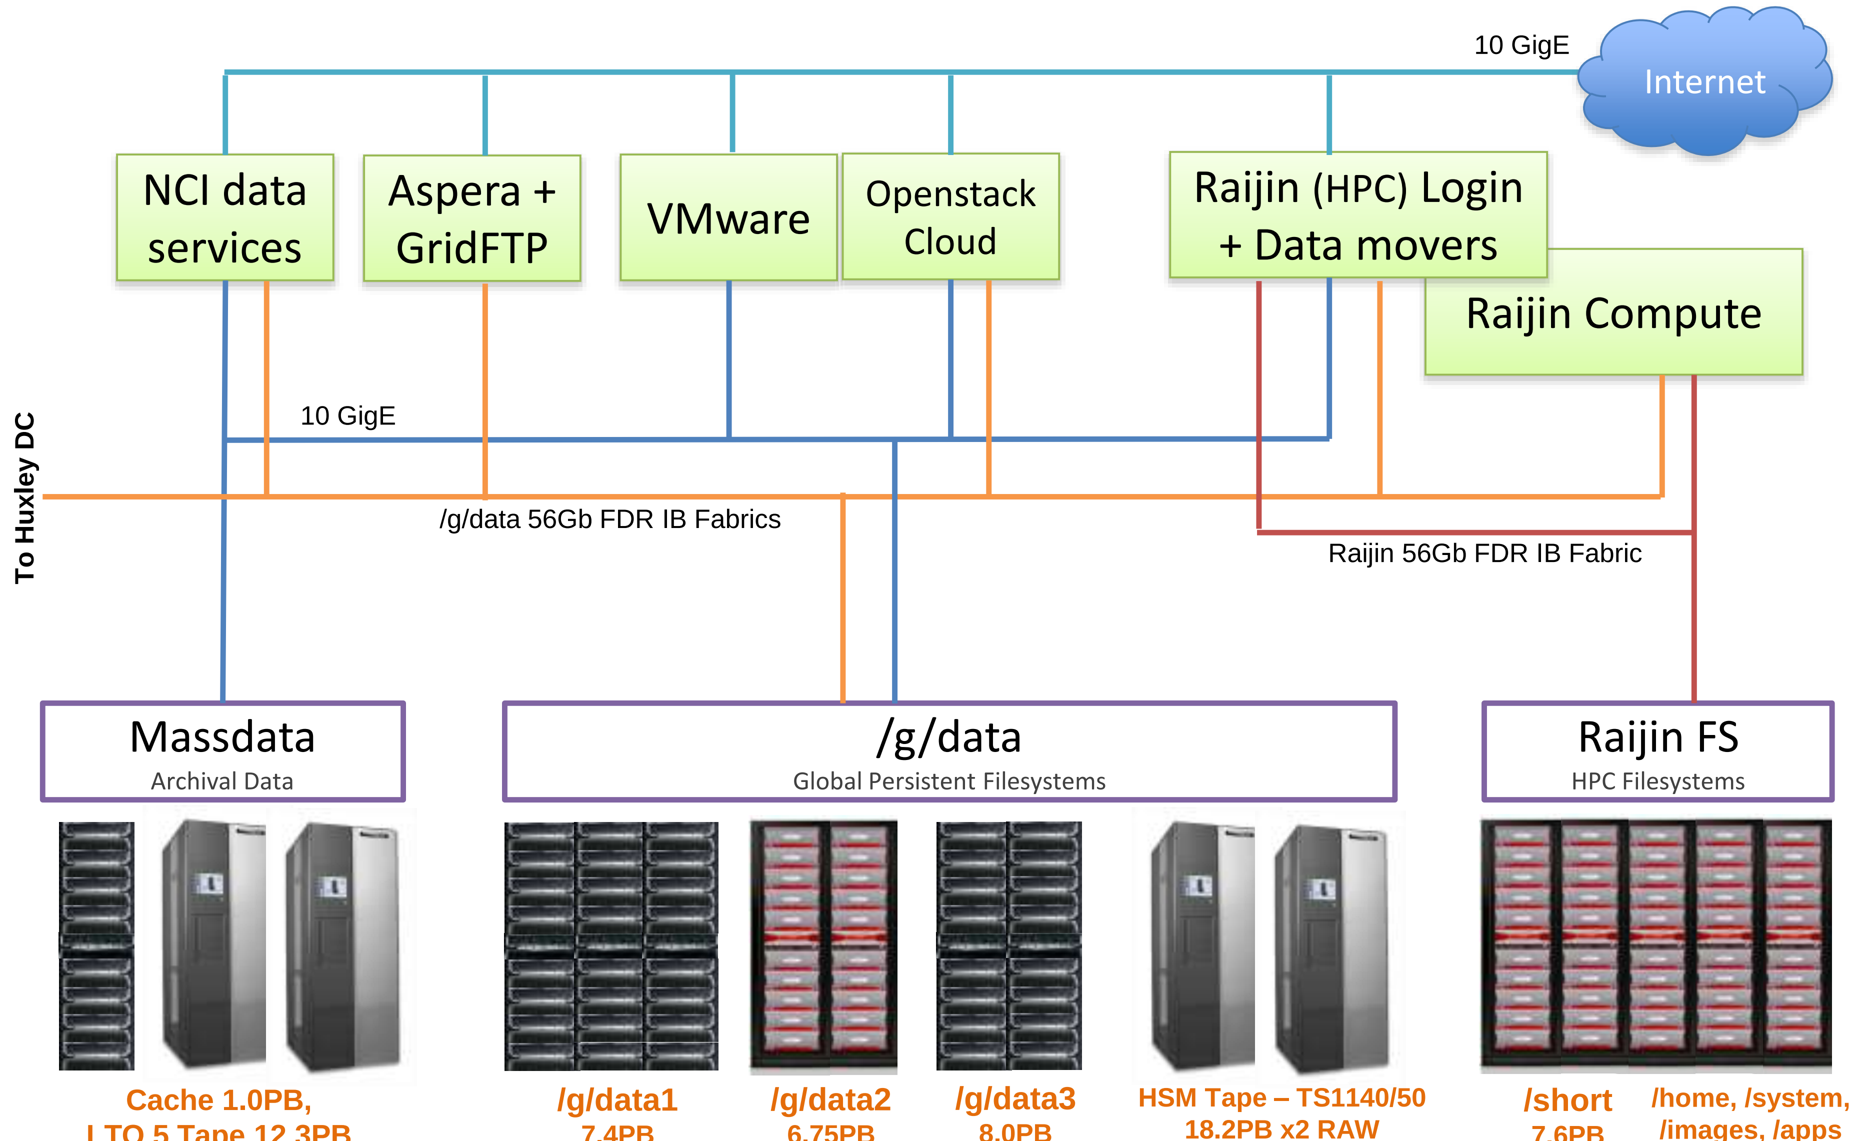

Frictionless software environment

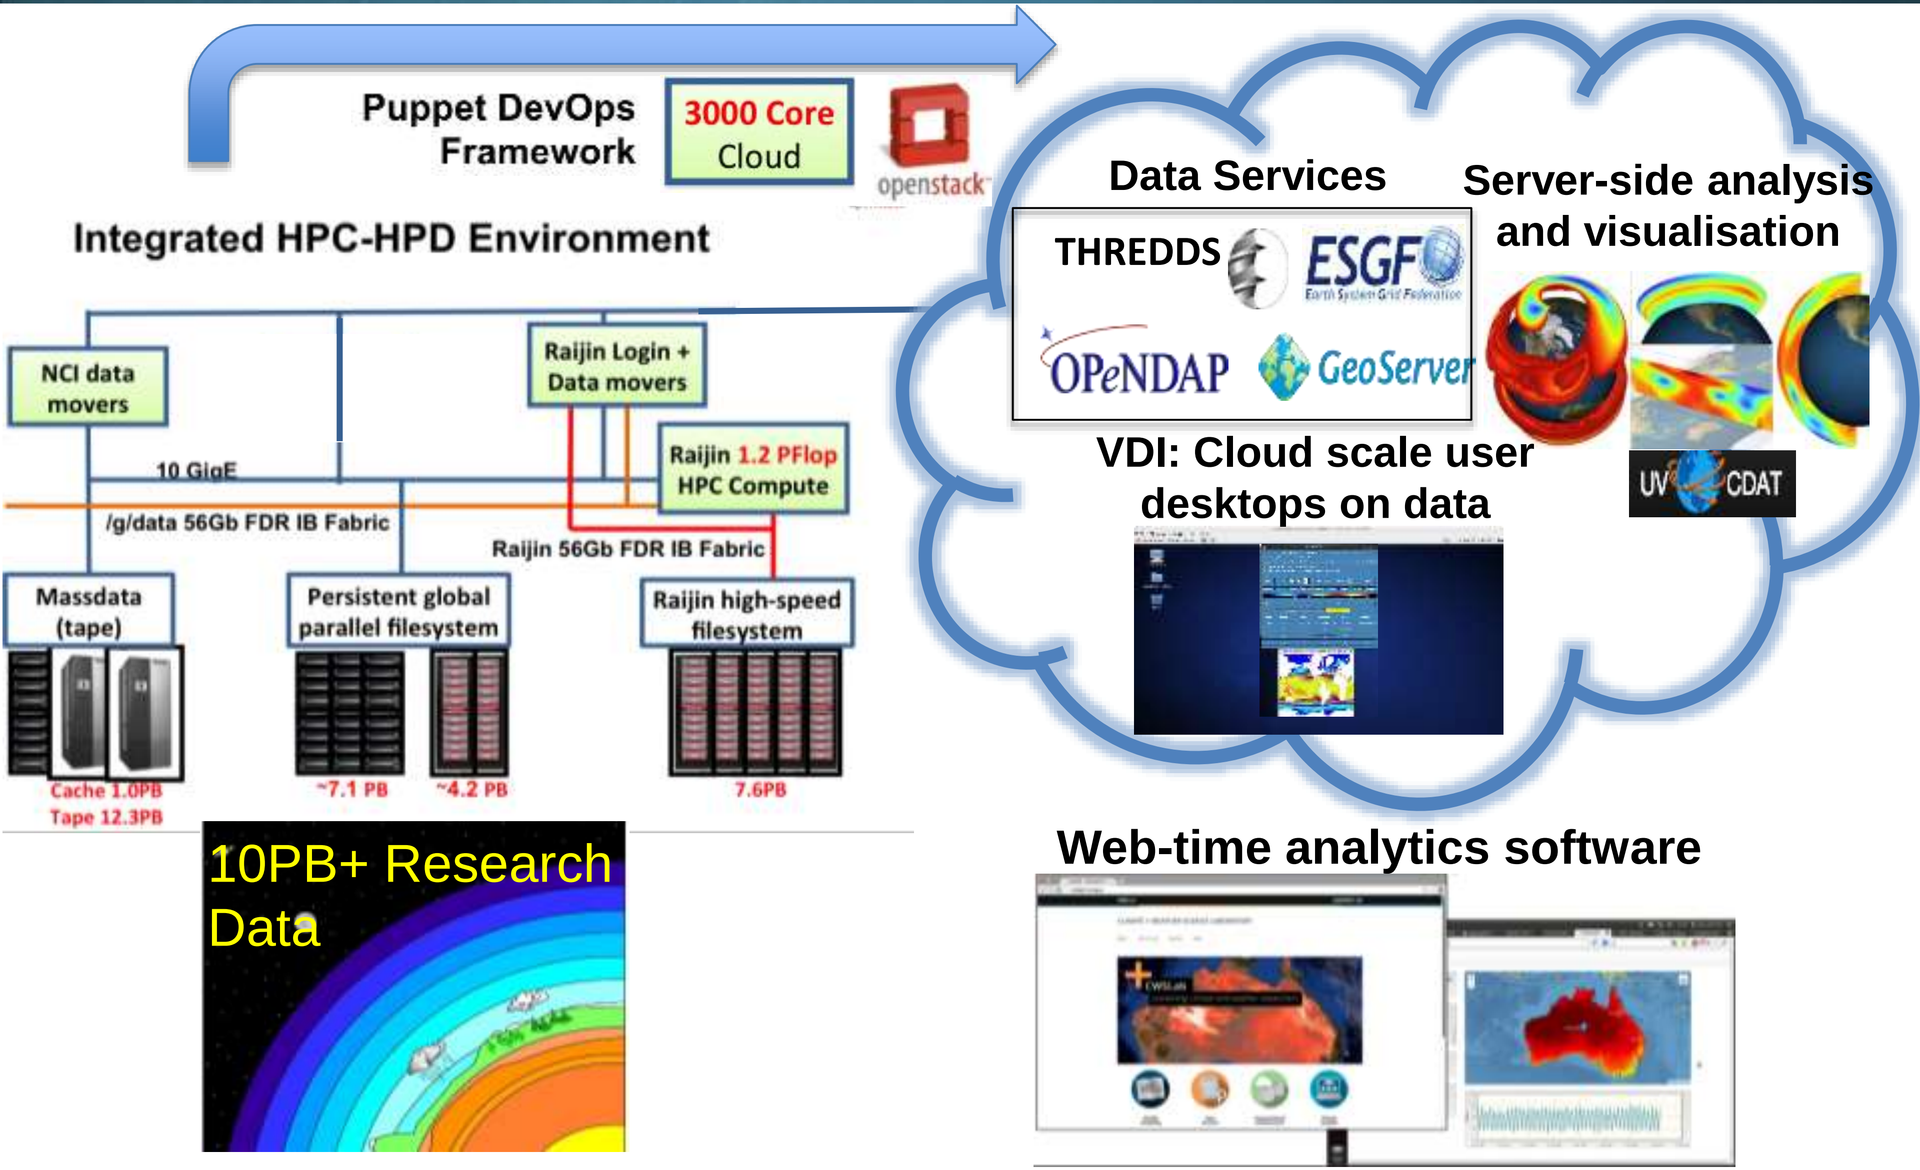

NATIONAL COMPUTATIONAL INFRASTRUCTURE

The realities of HPC today — the end of the “free lunch”

Supply, demand, and workloads at NCI in 2016

- System utilisation at ~95%
  - Maximum expected in a big, multi-user facility — diverse task load (32 cores up to 20,000 cores)

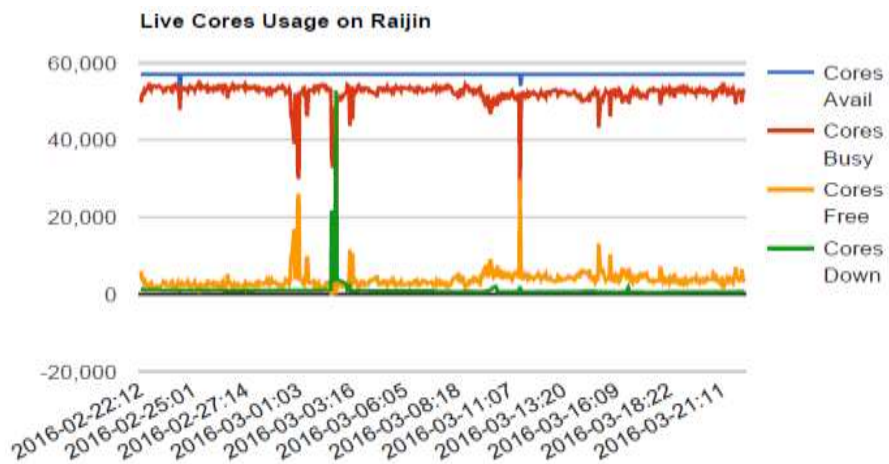

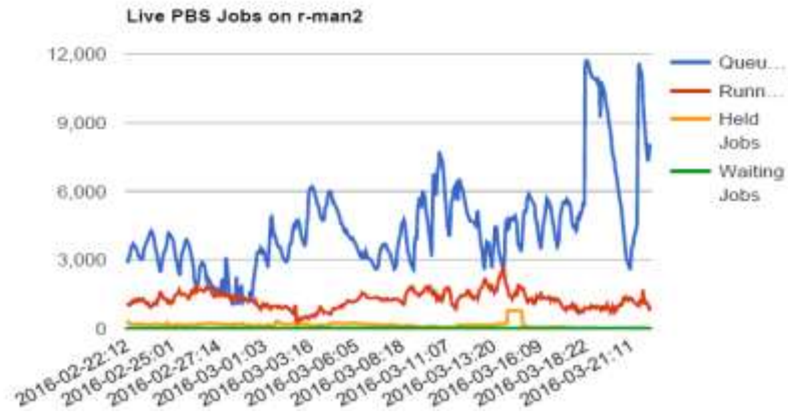

- Demand exceeds supply in both universities and agency environments
- Comparison: For Raijin to be at #24 today, as in 2012, it would be a 3 PF

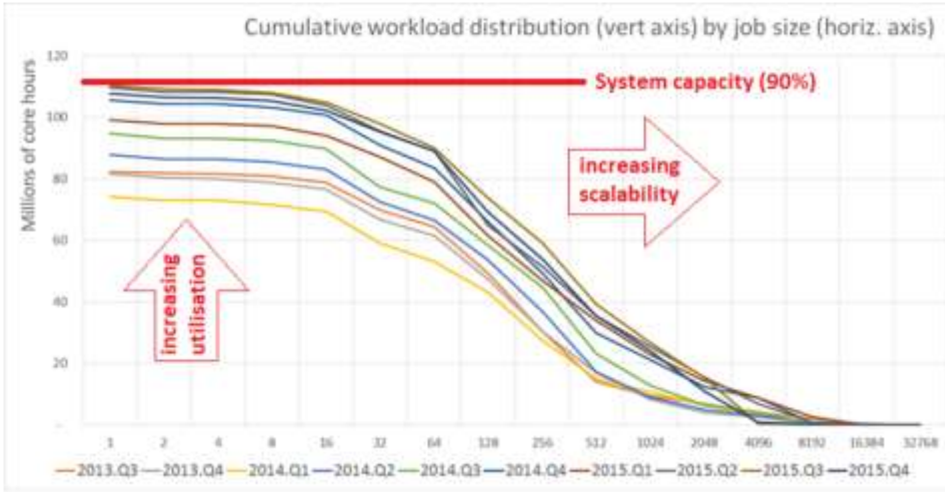

Top 500 Supercomputer list since 1990

http://www.top500.org/statistics/perfdevel/

- What is the next-generation system?
- What is the mix of cores/accelerators

- Not just about an HPC system but a comprehensive and integrated environment

- A balance between processing power and ability to handle/access data

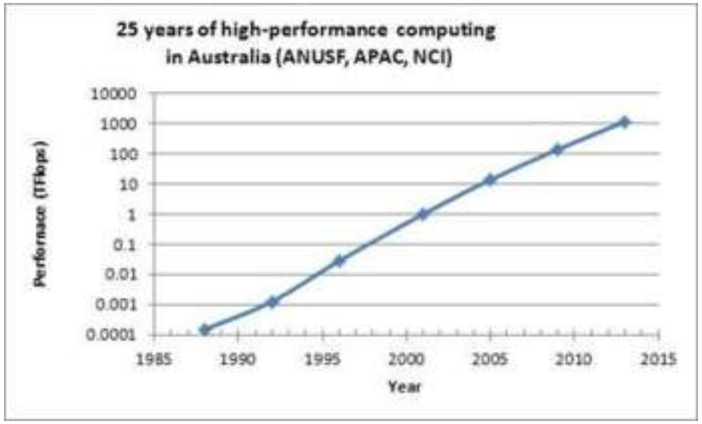

Limits to growth/Changes in approach for HPC

- Current Top 5 on the Top 500 List (June 2016)
  - (1) Sunway TaihuLight (NSC, Wuxi, China) — **93.0 Pflops**; Power: **15.4 MW**; 10,649,600 cores; Sunway SW26010 RISC processors;
  - (2) Tianhe-2 (NUDT, Guangzhou, China) — **33.8 Pflops**; Power: **17.8 MW**; 3,120,000 cores; Intel Xeon E5 and Intel Xeon Phi (accelerators)
  - (3) Titan (DoE Oak Ridge NL, USA) — **17.6 Pflops**; Power: **8.2 MW;** 560,000 cores; Cray XK7 with 18,688 NVIDIA GPUs (accelerators), each with 2688 cores + AMD Opteron
  - (4) Sequoia (DoE Lawrence Livermore NL, USA) — **17.2 Pflops**; Power **7.89 MW**; 1,572,864 cores; IBM Blue Gene/Q
  - (5) K Computer (Riken AICS, Japan) — **10.5 Pflops**; Power: **12.66 MW**; 705,024 cores; Fujitsu SPARC architecture
  - (6) Mira (DoE Argonne NL, USA) — **8.586 Pflops**; Power**: 3.945 MW**; 786,432 cores; IBM Blue Gene/Q

  - (2012/2: **24**,27,32,38,51 72,86, 98) Raijin (NCI, Australia) — Peak: **1.05 PFlops**; SPEC: **1,600 kSFRB**; Power: **1.4 MW** ; 57,452 cores; Fujitsu Primergy
  - (2014/2: **41**,58,79) Magnus (Pawsey, Australia)—Peak: **1.10 Pflops**; 35,712 cores; Cray XC40
- Messages
  - Only 2 of the current top 10 (#5,#9,#10), and 8 of top 20, do not have accelerators or specialist hardware
  - *Accelerators (multi-core & multi-threading) needed to maintain Moore's law growth* with feasible power
  - Accelerators are not well suited to all problems; require significant effort/expertise to program
  - A heterogeneous infrastructure will be needed to meet the totality of future requirements

The end of the "free lunch"

- Vayu (2009, core only, 0.140 PFlop)
  - R_{max}: 0.126 PF
  - SFRB: 0.25M
  - Cores: 11,936; Price: $13M;
  - **Power: 0.5 MW; $0.7M p.a.**
- Raijin (2013, core only, 1.2 PFlop)
  - Rmax: 1.05 Pflop (8.3 x Vayu)
  - SFRB: 1.6M (6.4 x Vayu)
  - Cores: 57,542; Price: $26M
  - **Power: 1.4 MW; $2+M p.a.**
- Project to 8 Pflops (2018?, Skylake cores only)
  - Rmax: ~6-7 Pflop (~6-7 x Raijin)
  - SFRB: ~5M; ~3.0 x Raijin
  - Cores: 150K(?); Price: approx. ~$70-80M?
  - **Power: ~3 MW (?); $4+M p.a.**
- Project to 8 Pflops (2018, *cores + accelerators*)
  - *A work in progress*

- Flops available to be extracted, but no gain in sustained performance w/o effort
  - Vayu -> Raijin: $x2, Flops x8, SFRP x6
  - Raijin-> ??: $x3(?), Flops x8, SFRP x3
- *More of the same, but nothing new!*
- ***The free lunch is over***
  - A more cost-effective way is needed
  - Raw infrastructure replaced by:
    - Expertise and know-how
    - Heterogeneous systems, accelerators
  - A major software effort from required

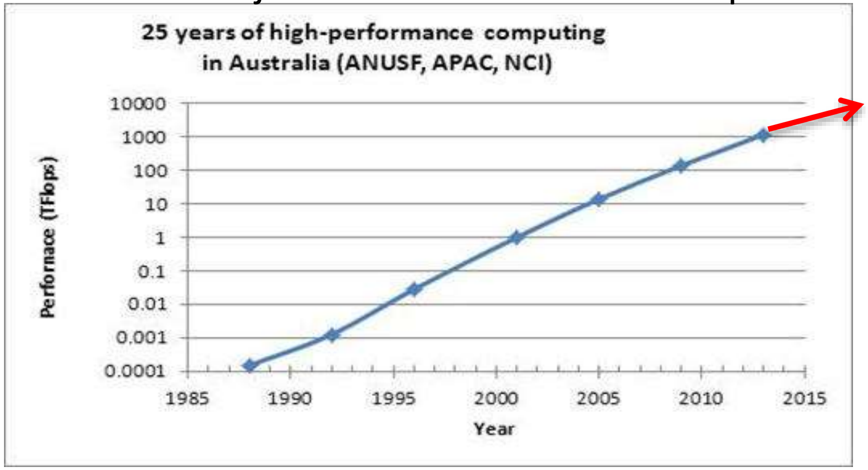

NATIONAL COMPUTATIONAL INFRASTRUCTURE

International trends and drivers — tomorrow is already here

Key trends / science drivers: international and national

- **Model precision and fidelity requiring larger scale computation**
  - Higher-res calculations; Increasing complexity (additional 'physics'); Ensemble modelling
- **Technology changes—current and future**
  - Increasing uptake of accelerator/many core technologies (GPU, XeonPhi, ARM) — energy efficiency
- **More effective methods, tools and algorithms (i.e. expertise)**
  - Big challenges to reengineer codes to realise latent performance gains in next-gen hardware
  - Without S/E effort and new methods: *more of the same but little/nothing new*
  - NCI ex. : 40% gain for atmosphere (parallel i/o etc.—UK Met); 30% perf. gain for ocean (NOAA)
- **Vital investments in people, skills and training (critical)**
  - Critical mass of expert staff creates most powerful computing infrastructures
    - Maximising the effectiveness and efficiency
    - Balanced ecosystem from technology, computational science, mathematics, appl. domains
    - Dissemination of benefits: national research system; industry
  - Exemplars internationally:
    - NCSA; US DoE Labs; RIKEN AICS; PRACE facilities (EU); UK Hartree Centre
- **Importance of long-range investments**
  - Certain for highly-valued staff; alignment with long-term obligations and experiments

Key trends / science drivers: international and national

- **Increasing system integration**
  - Deep integration of compute and data—not HPC with data “bolted on”
  - Transformative shift towards integration for contemporary HPC environments
    - DoE Secretary’s Advisory Board Report:
      - *The need to extract useful information from this explosion of data becomes as important as sheer computational power. …. This shift dictates the need for a balanced ecosystem for high-performance computing with an undergirding infrastructure that supports both computationally-intensive and data-centric computing*
    - National Academies Report to NSF of the USA*:*
      - *For many scientific disciplines, the issue is not whether to use data or simulation, but how the two will be used together*
      - *Co-location of computing and data will be an important aspect of these new environments, and such approaches may work best when the bulk of the data exchange can be kept inside a data centre*
  - In Australia
    - Significant local drivers for HPC/HPD integration
      - Enablement of *system science* by bringing together data from multiple domains
      - Environment: climate science, earth observation (LandSat, Sentinel, etc.)
      - Medical research: genomics, genome/phenome database
    - Pleasing to say that NCI in the world vanguard of compute/active data integration
      - Highly shapes by the institutional collaborations that sustain us

NCI

NATIONAL COMPUTATIONAL INFRASTRUCTURE

A blueprint for the National Computation System

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

National HPC strategy and infrastructure:
Collaboration driving service evolution driving collaborative outcomes

Mission: ***World-class, high-performance computing services for Australian research and innovation***

The vehicle:

- Highly-integrated, HPC/HPD e-infrastructure environment
- Comprehensive/integrated infrastructure and expert service
  - Petaflop-class supercomputer +
    Supercomputer-class private cloud +
    Very high-performance storage systems +
    Internationally-renowned expert support team
  - *This is NCI as it is today*

Values-driven—national and strategic:

- *Enabling high-impact research* that matters/informs policy
- *Research & outcome-driven*—national priorities, excellence
- Designed via *deep engagement*—inst., centres, industry
- To deliver *transformative outcomes / national benefits*
- *Quality/Innovation*—scale/aggregation, experience/expertise
- *Valued*: demonstrated via co-investment, uptake growth
- *Collaboration* driving *service evolution* for *collaborative outcomes* and *research and innovation competitiveness*

Research Goals/Outcomes

Engagement: Communities and Institutions/ Access and Services

Integration: HPC Services Virtual Laboratories/ Data-Intensive Services

Expertise: Support and Development

Compute (HPC/Cloud) Storage/Network Infrastructure

Cash Investments/Co-investments and Costs

Capital (NCRIS)

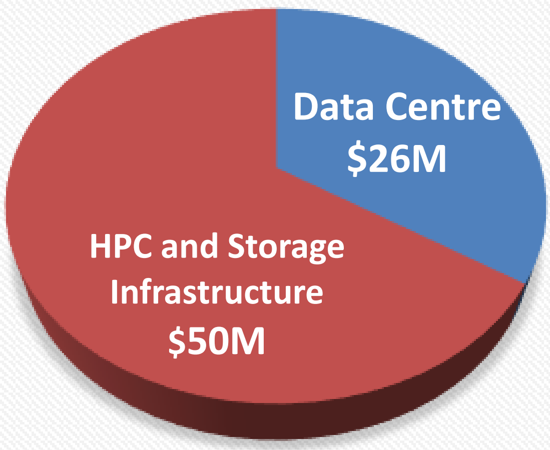

Recurrent (Per Annum)

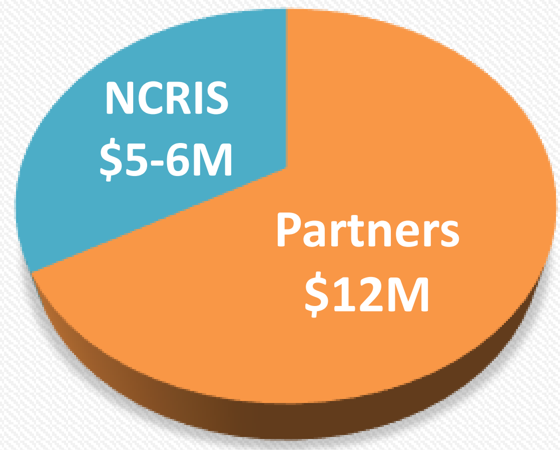

NCRIS investments since 2008:

NCI: $85.5M
Other: ~$8.1M

Partner Co-investment since 2008: $58.9M

Attributes of a future National Advanced Computing Capability

- **National Capability of International Scale**
  - Limited number of peak (leadership-class) facilities for high-performance (HPC/HPD)
    - Aggregated for scale, service quality, service innovation
    - Aspire to an international top 20-30 ranking ($80-100M; ~8 Pflop)
      - 8 x peak performance; data volumes exponential growth (TBD)
  - Integration with (commercial) cloud for high-throughput/lesser requirements
    - Noting limitations of cloud (connectedness, file system performance)
- **Deeply embedded in the National Research Fabric — even more so than today**
  - Preferred, expert and trusted provider to researchers, communities, institutions
- **Converged infrastructure: fit-for-purpose for computational/data-intensive research**
  - Comprehensive/integrated/high-performance supporting:
    inter-/trans-disciplinary research; system science
  - Close integration: HPC and high-performance data — US National Academies paper
- **Emphasis on Expertise, Skills and Education: a national computational sciences capability**
  - Aligned with peak facilities — supporting increased uptake, research ambition
  - Supporting priorities
    - Re-engineering codes of national significance—new processor technologies
    - Increasing use of data-intensive methods analytics, machine learning, etc.
  - Consultancy capability for industries/science agencies
  - Nucleus for a separately-funded research capability (e.g., co-funded CoE)
  - Leadership for a vibrant (u.g./p.g.) education program to skill a future workforce
  - Leadership for Australian participation in international HPC/HPD collaborations
- **Multi-year Investments**
  - Maintain internationally competitive environment, relevant to Australia
  - Retention and development of highly-specialised skills

Optimised Southern Ocean Simulations:

NCI
NATIONAL COMPUTATIONAL INFRASTRUCTURE

Thank You

Questions