iVEC and The Pawsey Centre

Neil Stringfellow, iVEC Executive Director

Tuesday 3rd December 2013

iVEC and the Pawsey Centre

- iVEC is an unincorporated joint venture of four public universities in Western Australia and the CSIRO
  - iVEC has been supported by the state government of Western Australia since its founding at the turn of the century
- In 2009 iVEC was selected to operate national supercomputer facilities as part of the Australian Government's Super Science scheme
  - The new centre was named the Pawsey Centre after the radio astronomer Joseph Pawsey
  - 25% of resources of the Pawsey Centre are for operational radio astronomy

Australian Government

iVEC Facilities and Expertise

40+ staff across five facilities, around Perth

- **CSIRO Pawsey Centre/ ARRC** – uptake, supercomputing , data, visualisation
- **Curtin University** – uptake and visualisation
- **Edith Cowan University** – uptake and visualisation
- **Murdoch University** – supercomputing
- **University of Western Australia** – supercomputing, uptake, visualisation

Sites linked together by dedicated high-speed network, and to rest of world via AARNet

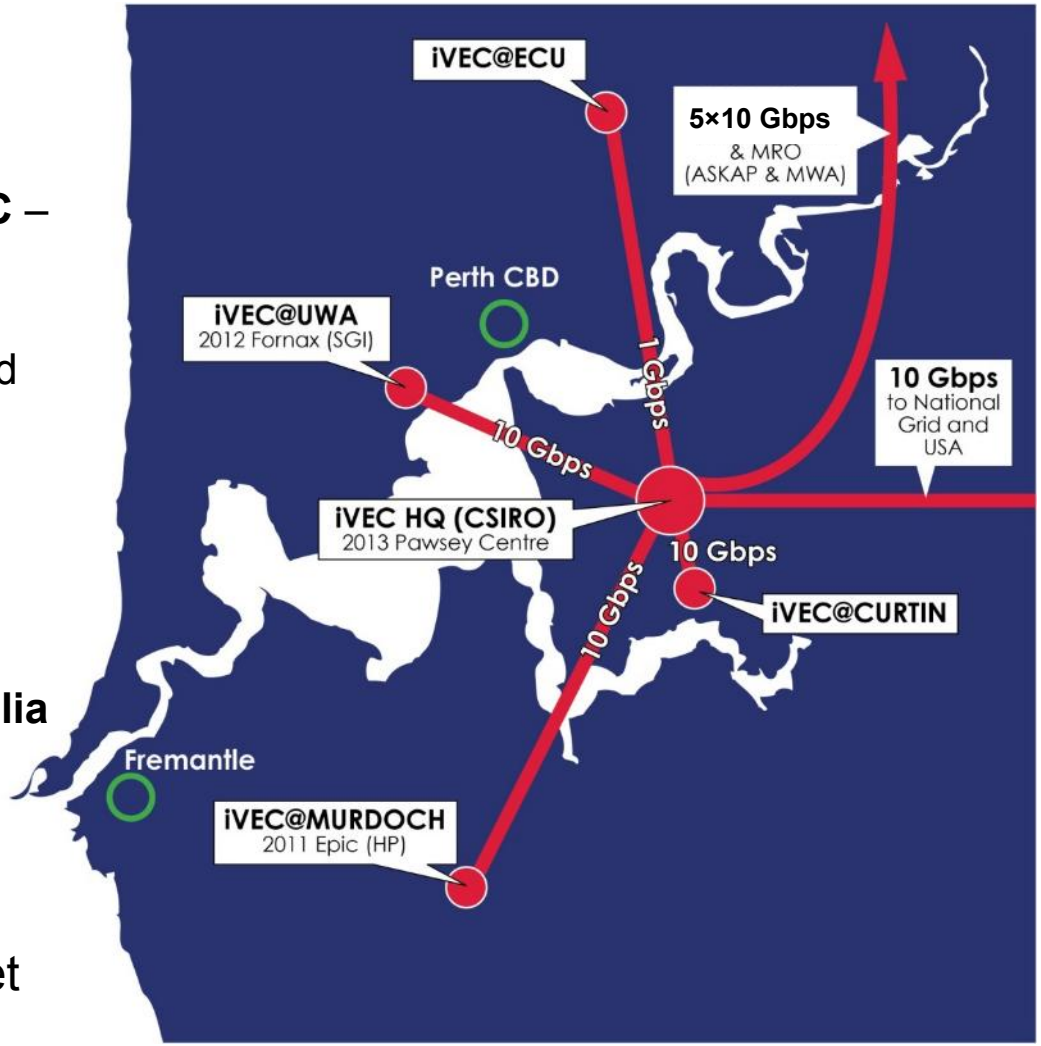

iVEC

Super Science Initiative (May 2009)

Substantial investment, by Australian Government, in eResearch infrastructure – $312M AUD over 4 years:

- $80M supercomputer in Perth (Pawsey Centre)
- $50M supercomputer in Canberra (NCI)
- $97M data management/ cloud computing (RDSI/ NeCTAR)
- $48M data tooling and management (ANDS)
- $37M research and education network

Operational and staffing costs from partners, regional agencies, etc.

The Pawsey Project

- In May 2009, Australian Government chose iVEC to establish and manage $80M AUD Pawsey Supercomputing Centre
- World-class petascale facility
  - towards enormous challenges of computing and data-processing for Square Kilometre Array (SKA)
- Hub for supercomputing supports high-end research in:
  - radio astronomy, geosciences, chemistry, nanotechnology, biotechnology, engineering, physics, ...



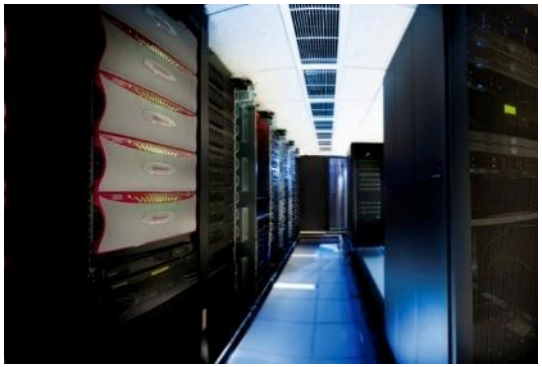

The Pawsey Centre – completed April 2013

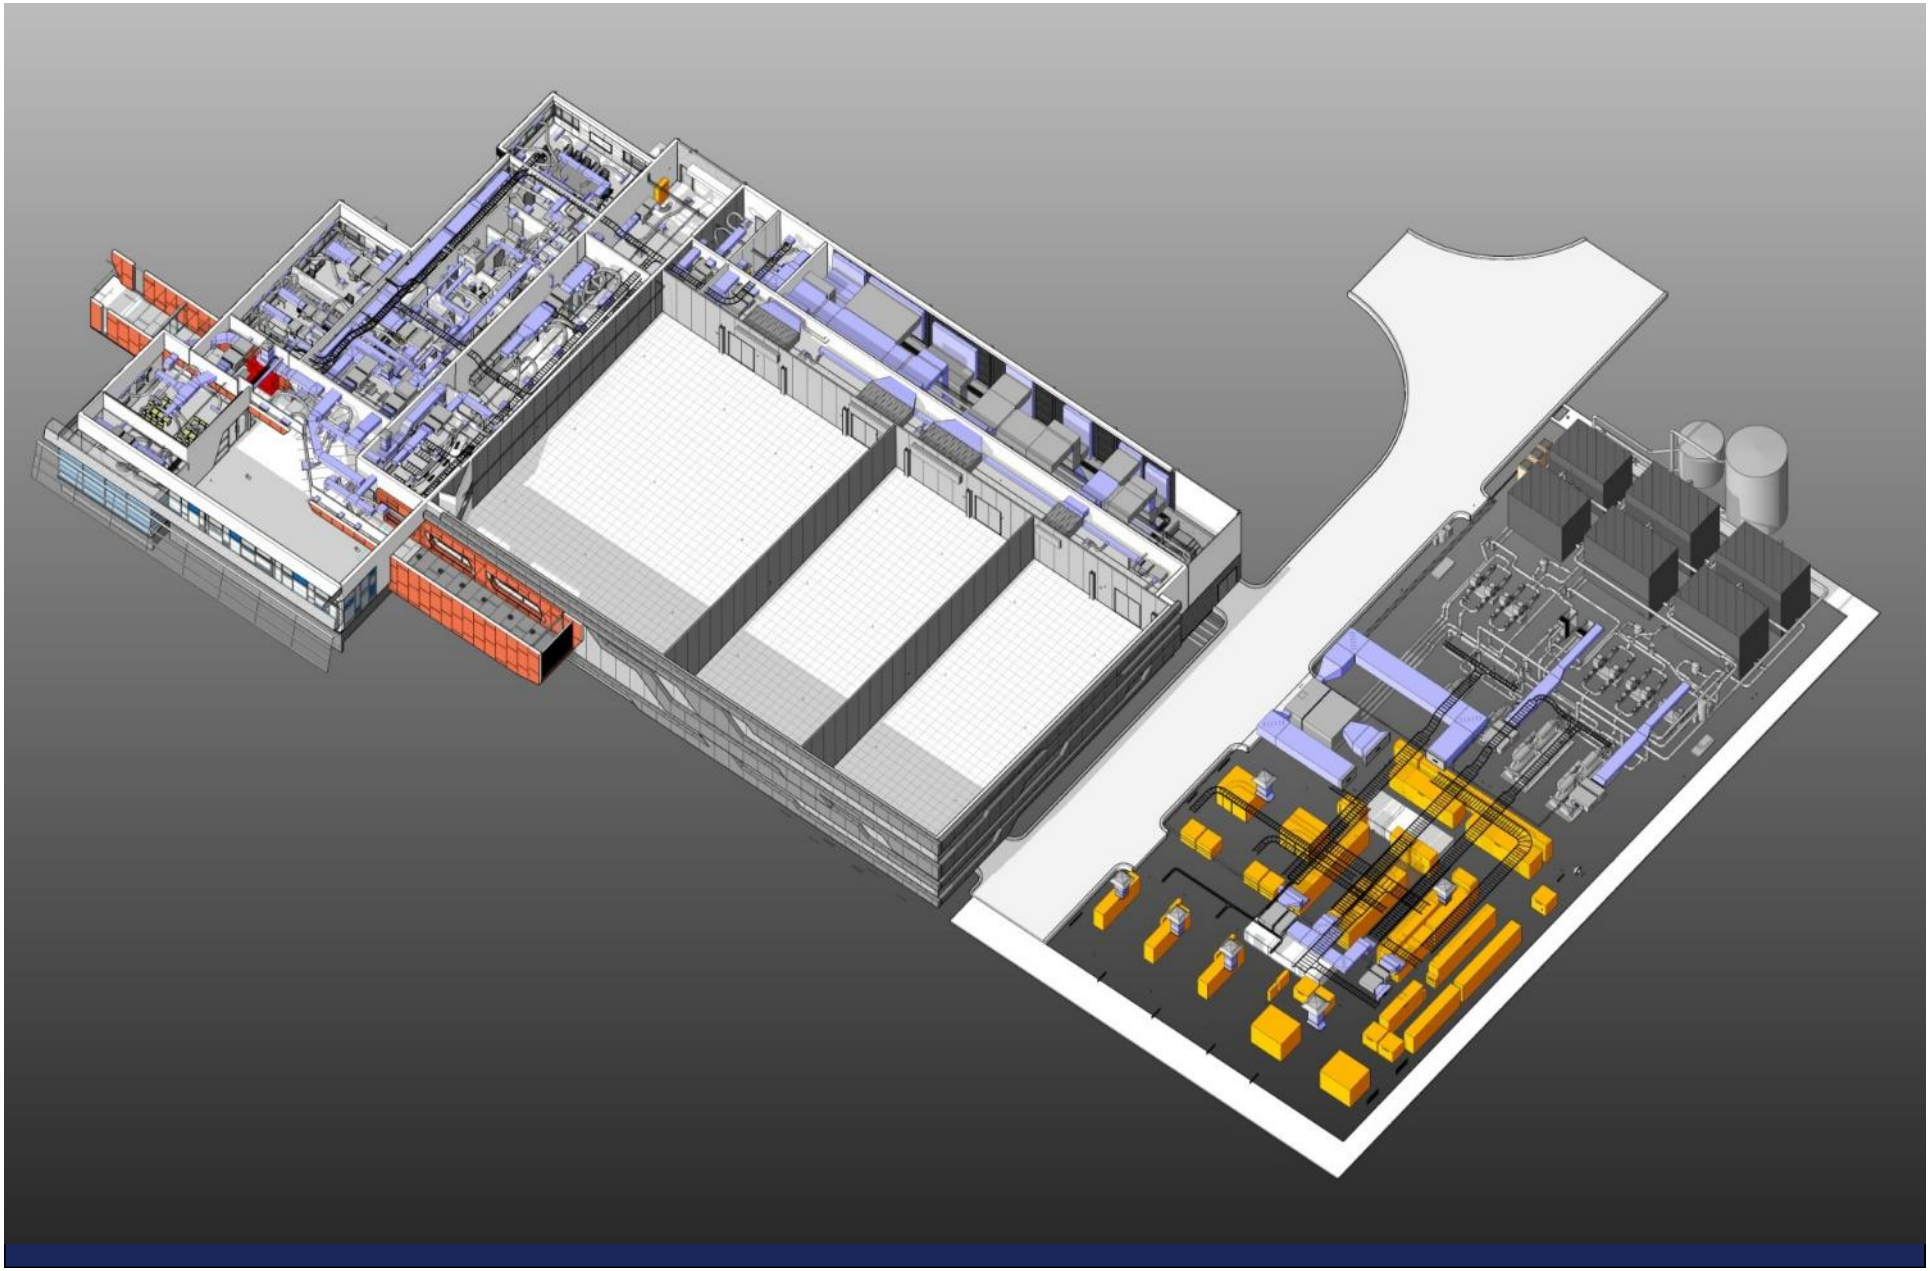

Pawsey Project Current Systems

Epic at Murdoch University

- 800 nodes each with 2 x 6-core Intel Westmere processors
  - 9,600 Intel Westmere cores
- Most nodes have 24 Gigabytes of memory
- General purpose machine for “classical” HPC
- Used by ASKAP for data processing

Fornax at UWA

- SGI system of 96 nodes each with 2 x 6-core Intel Westmere processors
- Most nodes have 72 Gigabytes of memory
- Every node has a Nvidia “Fermi” GPU
- Used for data-intensive science
- Used by MWA for data processing of operational observations

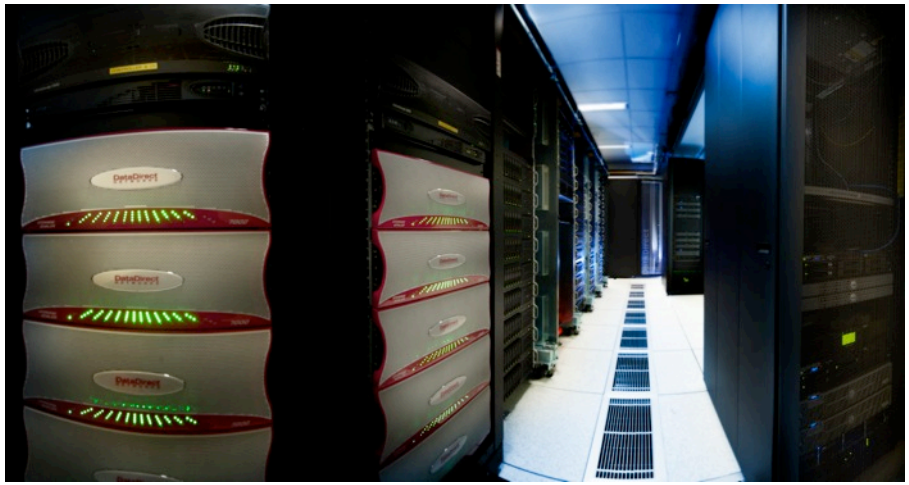

iVEC in the Petascale Era

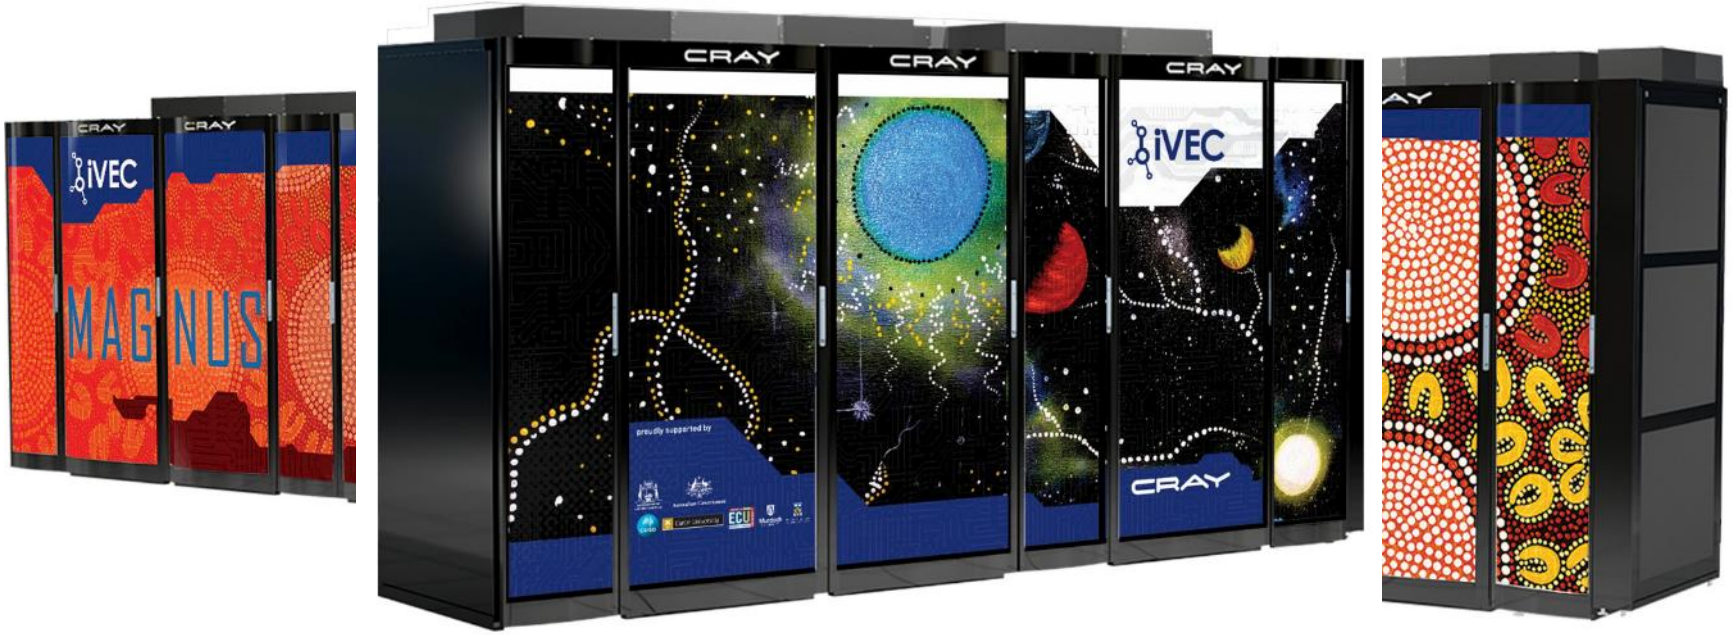

- Two Cray XC30 ‘Cascade’ supercomputers
  - RTC – 200TFlops (Ivy Bridge) real-time computer for radio-astronomy operations
  - Magnus – 1+ Pflops system for wider scientific community
    - Initial deployment of 2 cabinets in 2013
    - Final deployment in mid-2014

Petascale Storage Facility

- Hierarchical Storage Manager (SGI and HDS/Spectra)
- Capacity for 100 PB (tape)
  - Initial media for 2×20PB
- Front-end 6PB disk cache, plus specialised D/B and fast-disk
- FDR Infiniband across infrastructure
- First two years of operation for SKA pathfinders (16+ PB)

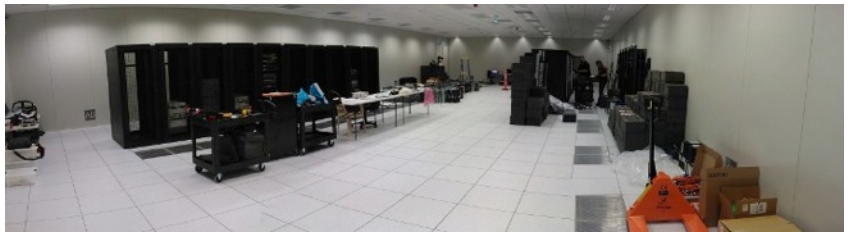

Pawsey Centre Architecture

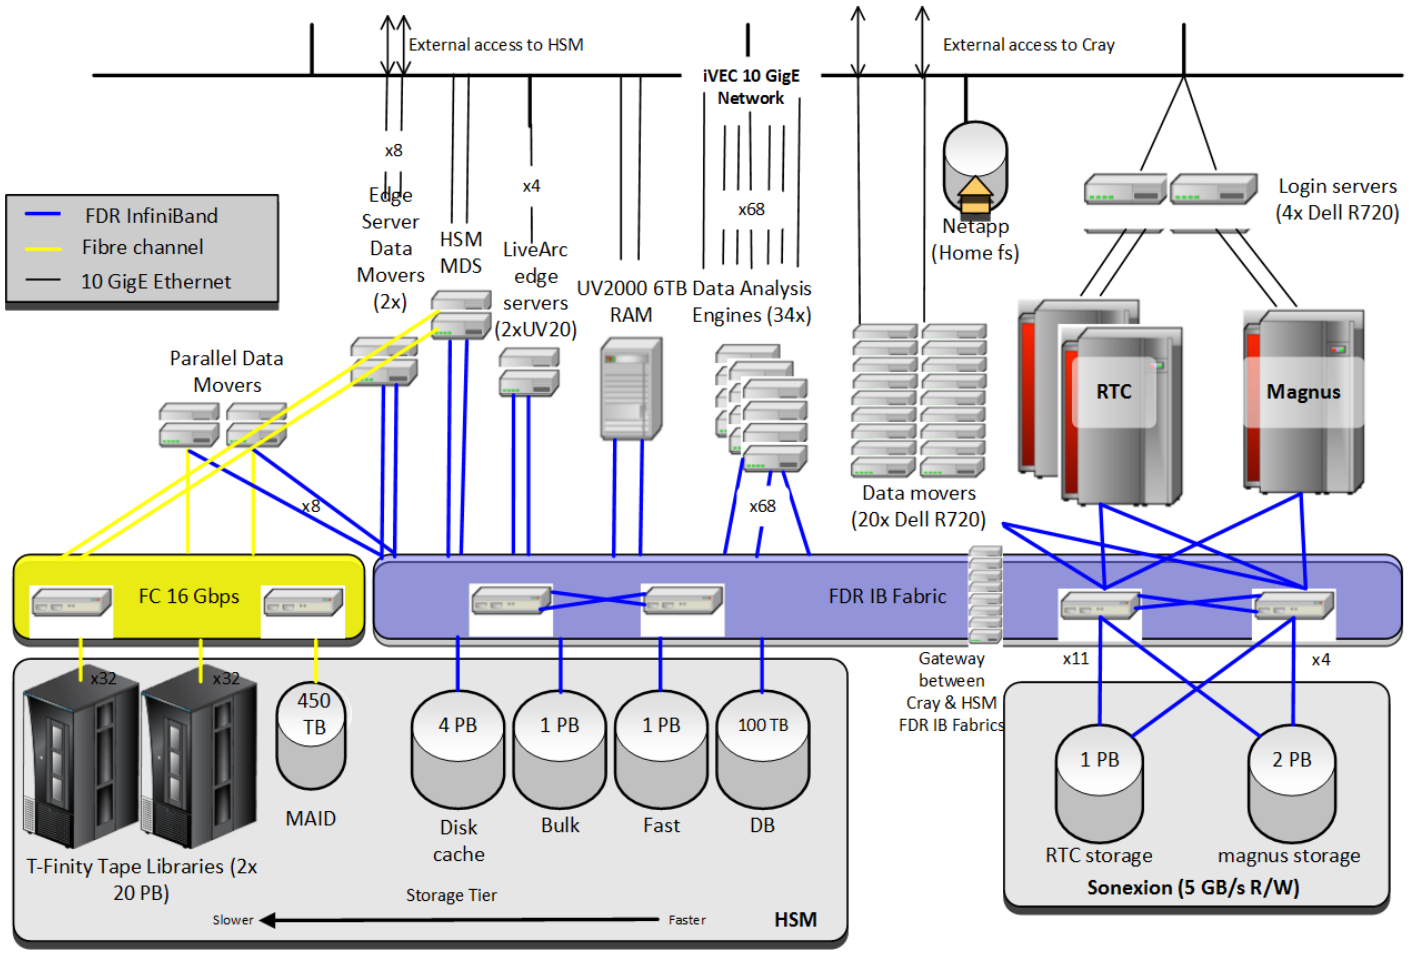

The Pawsey Centre and Sustainability

- Supercomputers are power hungry - the supercomputers and related equipment in the Pawsey Centre will consume ***megawatts*** of power
  - That power needs to be dissipated away
- The Pawsey Centre uses innovative groundwater cooling to take away the waste heat
- The pumps for the groundwater cooling are powered by solar power from photovoltaic cells

Pawsey Centre Green Initiatives

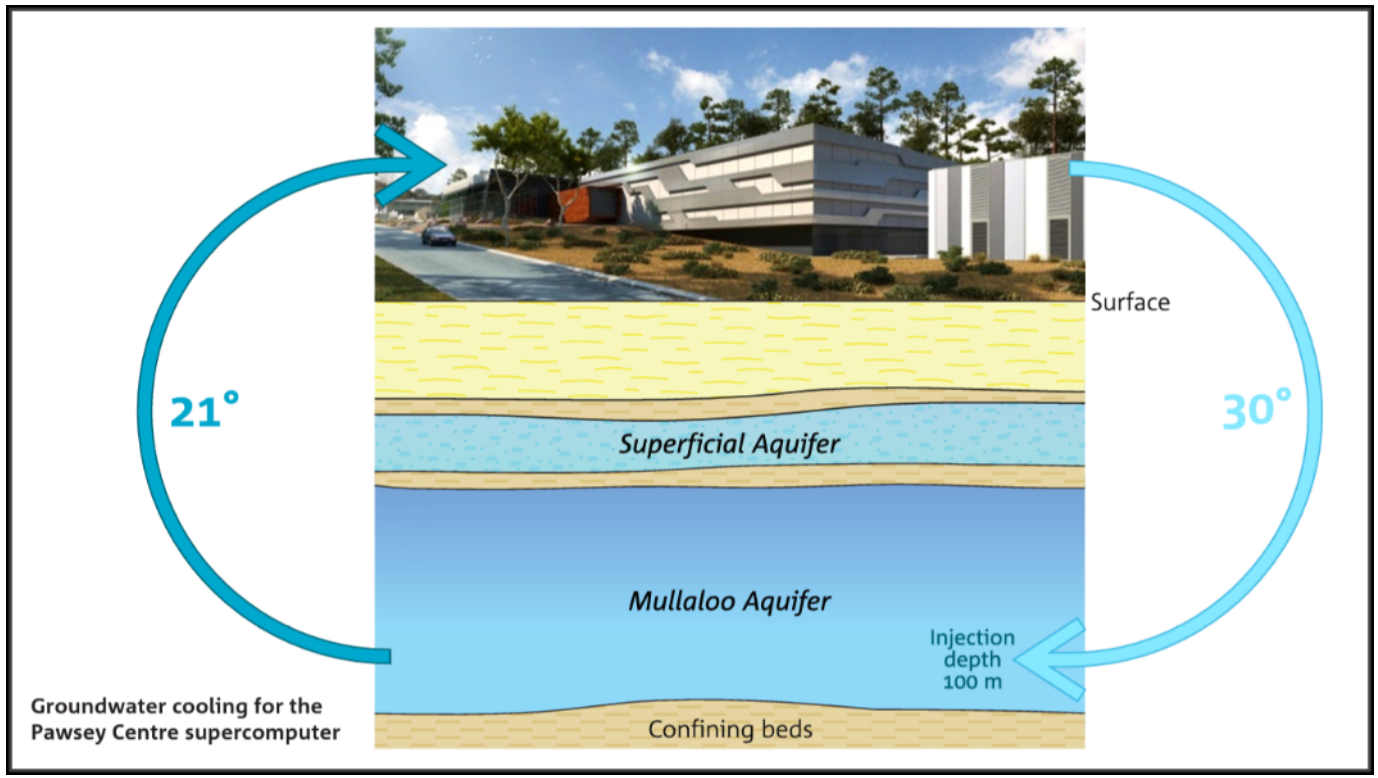

- Groundwater Cooling
- Closed-loop circuit provides cooling for supercomputers without need for chillers or cooling towers
- Estimated saving of 38M litres potable water per year

Pawsey Centre Green Initiatives

- 140 kWatt photovoltaic cells use plentiful Western Australian sun to drive cooling systems
- Aim is to produce a zero-cost cooling system for supercomputers

iVEC

Large Scale Science Needs Heavy Computing

- Major multi-billion dollar projects rely heavily on supercomputing or large-scale number-crunching resources
- Examples of such projects include the Large Hadron Collider, ITER, and the Square Kilometer Array telescope
- ITER relies on large-scale supercomputer simulations to model turbulence in Tokomaks
- The Large Hadron Collider (LHC) makes extensive use of a distributed grid of computers for data analysis

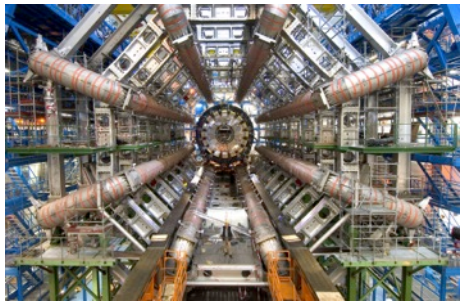

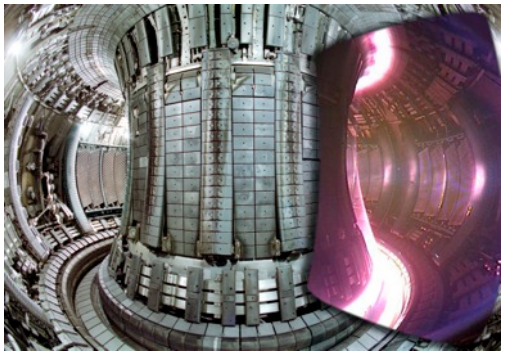

iVEC

Large Scale Science in Australia – the SKA

SKA SQUARE KILOMETRE ARRAY

- The Square Kilometer Array (SKA) will be the world's largest radio telescope when completed
  - The collecting area will be one square kilometer
- South Africa and Australia have been jointly selected to host parts of the telescope
- The SKA1 telescopes are due to begin early science in 2020
- SKA2 telescopes should be delivering data from 2024

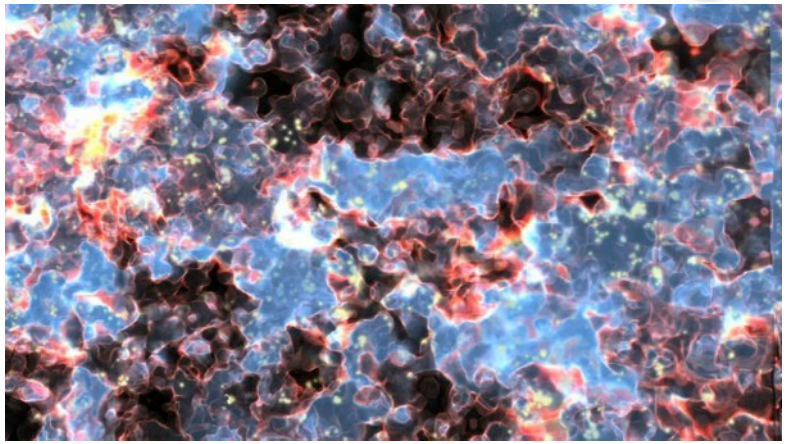

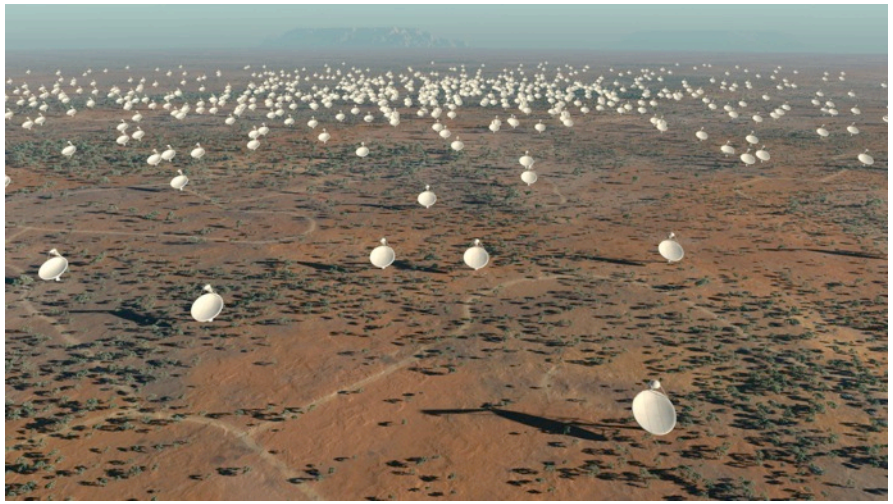

SKA Pathfinders

- 3 SKA “pathfinder” projects are being established

MeerKAT in South Africa

Murchison Widefield Array (MWA) in Australia

ASKAP in Australia

SKA Location in Australia

- The Australian pathfinder telescopes are located at the Murchison Radio Observatory (MRO) in the Shire of Murchison
  - The Shire of Murchison is radio-quiet
  - It has an area of 41,172 km^2
  - The population is never more than 120 people
  - For comparison …
    - Switzerland - area 41,285 km^2, pop: 8,014,000
    - Netherlands - area 41,543 km^2, pop: 16,789,000
- The MRO is in a desert region
  - Limited infrastructure for power, cooling, people
  - Have to strike a balance between *essential* data processing on-site and shipping data elsewhere

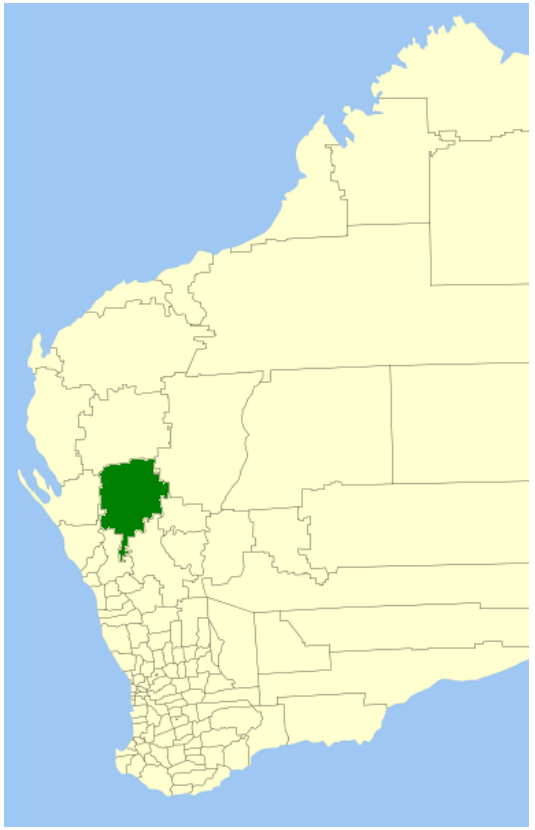

LHC as an Example

- The LHC uses a distributed grid of computers and storage facilities arranged in tiers
- CERN is the tier-0 site and has about 50 petabytes of raw data from experiments
- The 12 tier-1 and around 140 tier-2 sites deliver the bulk of the resources
  - 200 petabytes of disk storage and 200 petabytes of tape storage
  - Over 500,000 standard compute processors

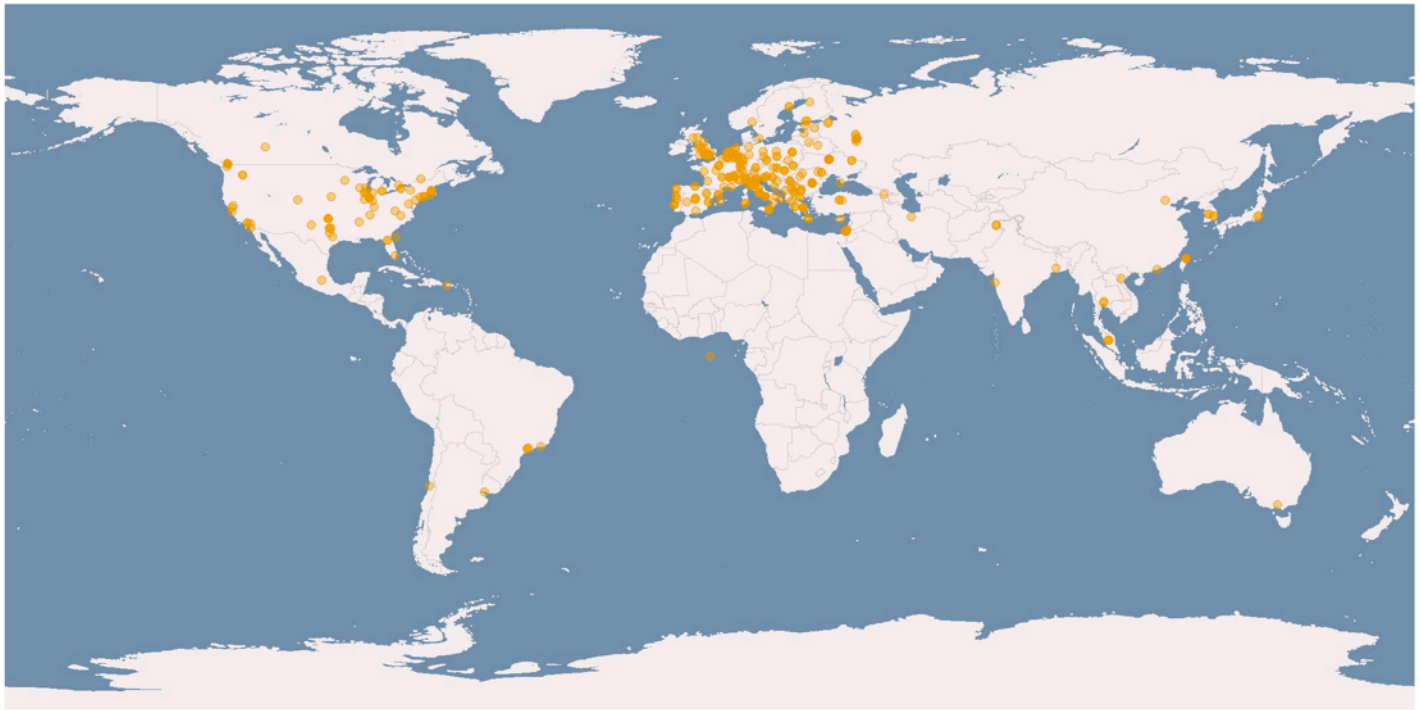

SKA Member Locations

The LHC is located in a region with high connectivity to the scientists

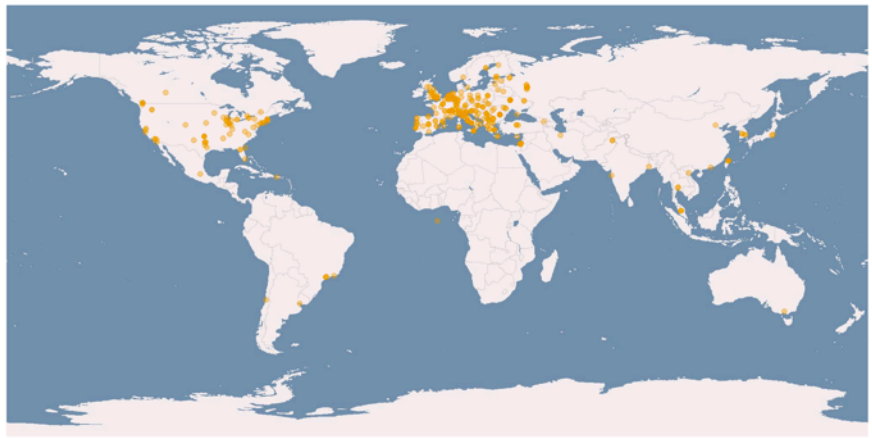

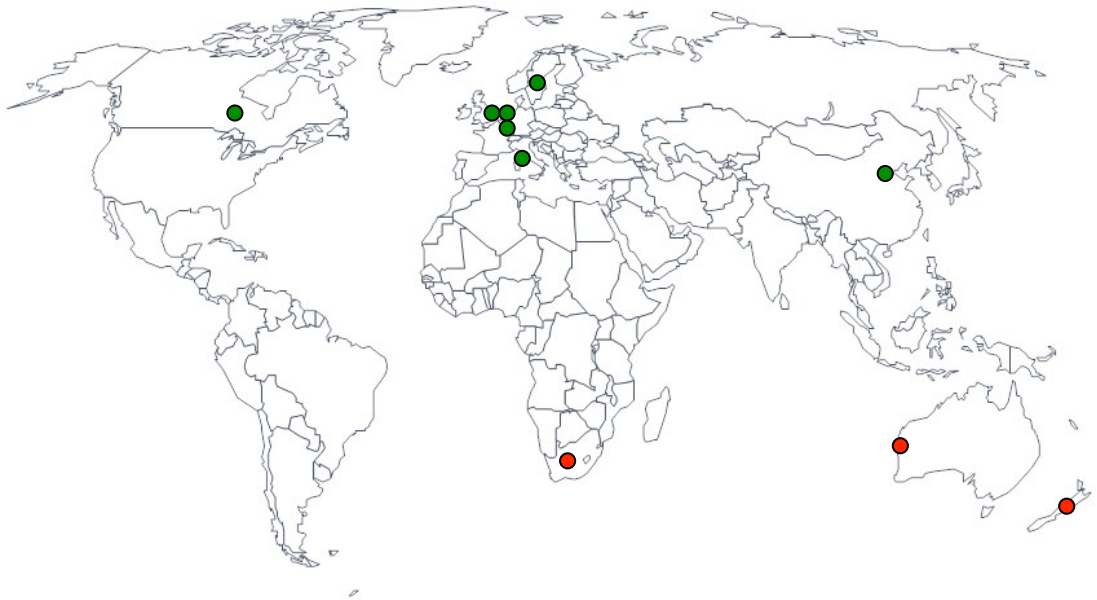

The SKA locations are far removed from most of the SKA member countries

SKA Data Volumes

- ASKAP and MWA will be delivering 5 petabytes and 3 petabytes of data per year to the storage in Perth
  - The raw data volumes from the dishes are much larger and are firsts processed on-site
  - MWA has 128 antennae producing data at a rate of 80 gigabits/s
    - Data is sent to a local GPU cluster on site to process
    - Part-processed data (3.2-6.4 gigabits/s) is sent to Perth on a 10 gigabits/s dedicated link
    - MWA telescope operates only on a 25% duty cycle
  - ASKAP and MWA are just pathfinders, the real SKA will be producing much larger volumes of data
- Future of Production SKA will be much larger
  - The future low-frequency telescope will have about 2.5 million antennae spread over 100 kms
    - The antennae stations will produce data at 10 gigabits/s each
  - SKA is expected to generate an exabyte per day of processed data
- For the LHC many of the tier-1 and tier-2 sites are close to where the data is generated
- For the SKA transferring the data from South Africa and Australia to Europe and Asia will be a major undertaking

ASKAP Data Flow

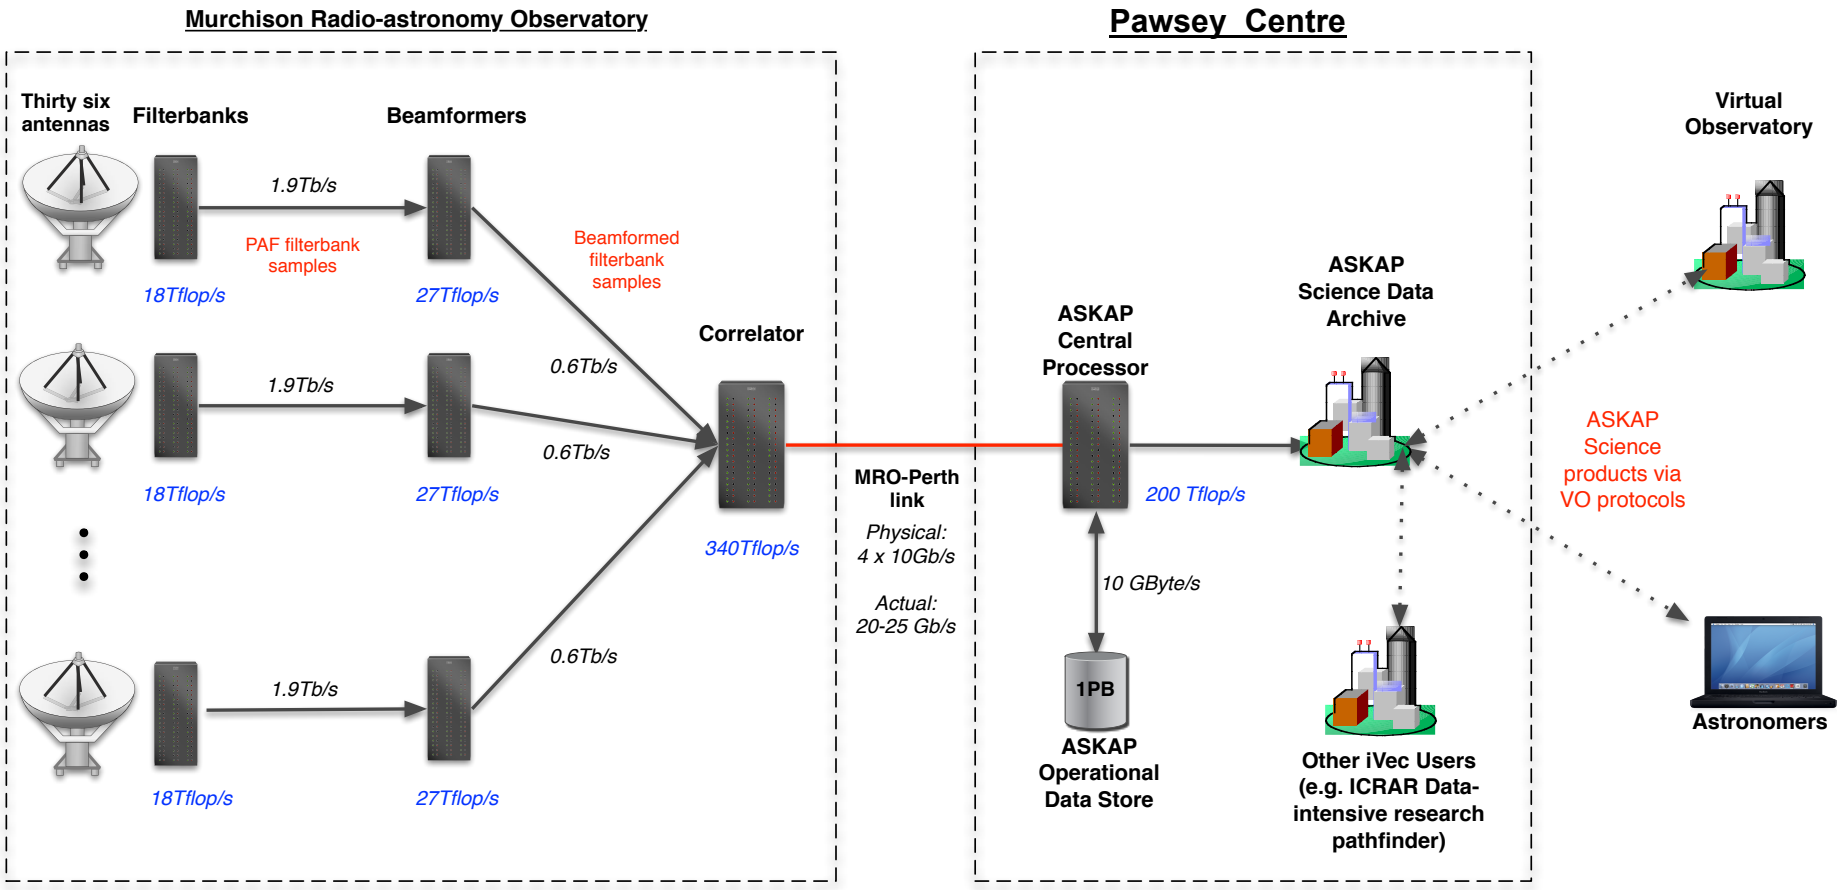

Courtesy of Ben Humphreys, CSIRO

iVEC

What are the Data Products ?

Artifact	Size	10 GBit/s (Peak)
Measurement Set 12 hour observation, 4 pols, 16416 spectral channels, 5s integration time	100TB (1)	22.8 hours
Spectral Line Image Cube 4096 x 4096 x 16416 spectral channels Stokes-I only	1TB	13.7 mins
Continuum Image 10240 x 10240 x 3 taylor terms Stokes-I/Q/U/V	4.7 GB	3.76 secs

Examples from ASKAP

Courtesy of Ben Humphreys, CSIRO

Big Data Products - Image Cubes

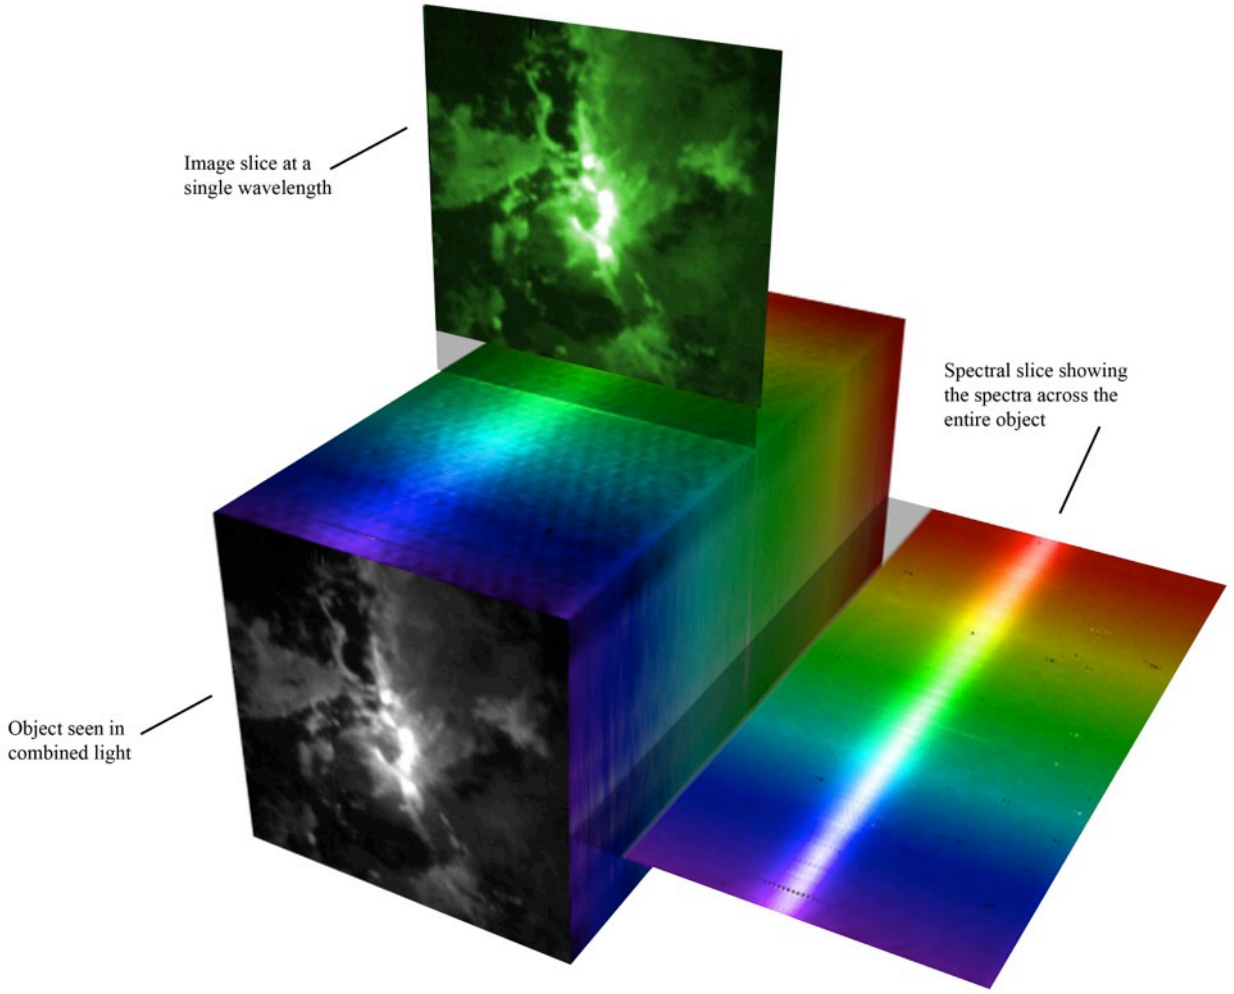

Sizes of Image Cubes

- Image Cubes generated by ASKAP will be of the order of Terabytes
  - Some will be 1 Terabyte
  - Some may be up to 5 Terabytes
- Image cubes from the SKA_2 (full production SKA) could be over 1 Petabyte
- Want to be able to manipulate whole image cubes for data analysis

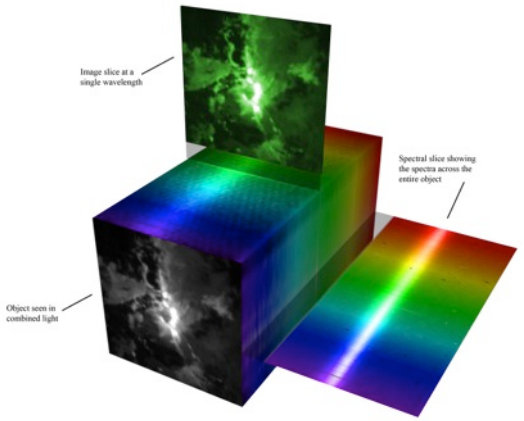

Data Analysis and Storage Systems at Pawsey

- SGI UV with 6 terabytes of shared memory
- 34 data analysis nodes with combinations of CPUs, GPUs and 128 to 256 Gigabytes of memory
- 6 petabytes of disk storage
- 40 petabytes of tape storage
  - *Licences for 100 petabytes of tape storage*
- This the equipment where most of the data analysis will be carried out

iVEC

The Pawsey Centre at Night