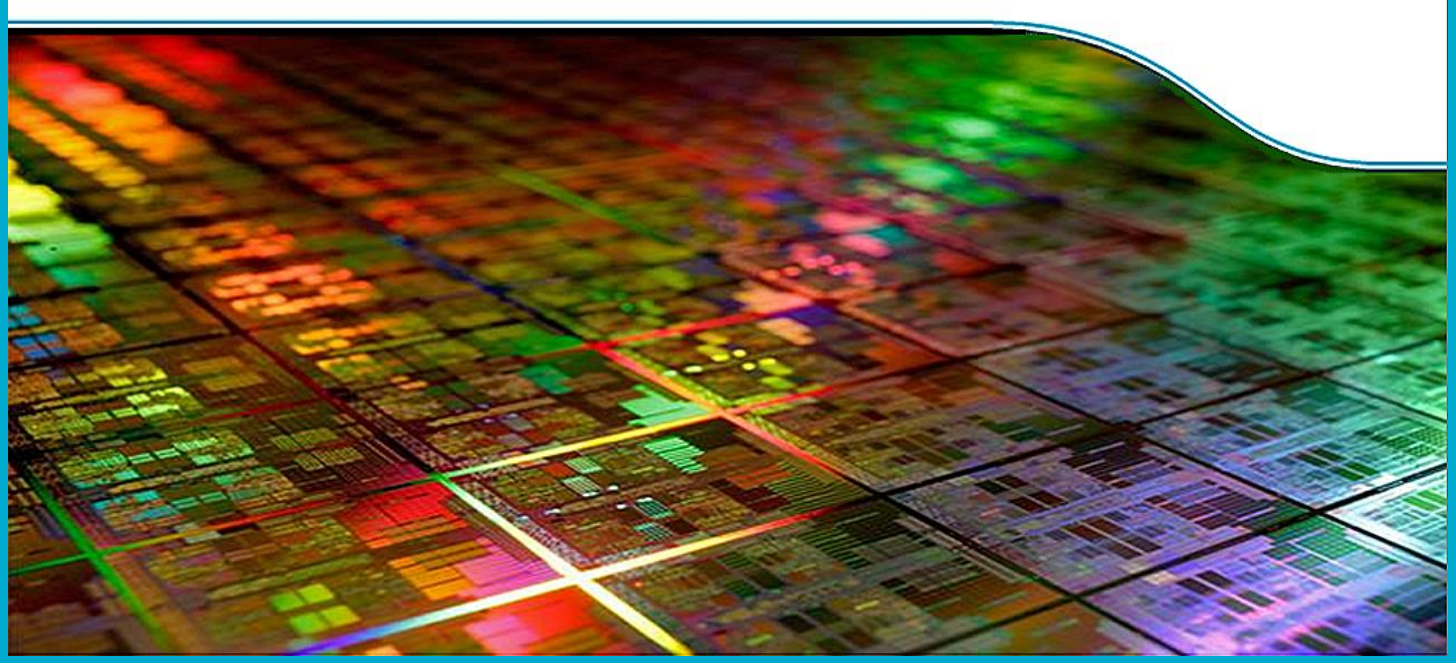

From CSIRAC to Cray (again) and onwards

- Robert C. Bell – CSIRO and Bureau of Meteorology

Summary

- The past:
  - history of computing support for science in Australia
- The present:
  - the Bureau's Cray system
- The future:
  - computing for science

Dinosaurs

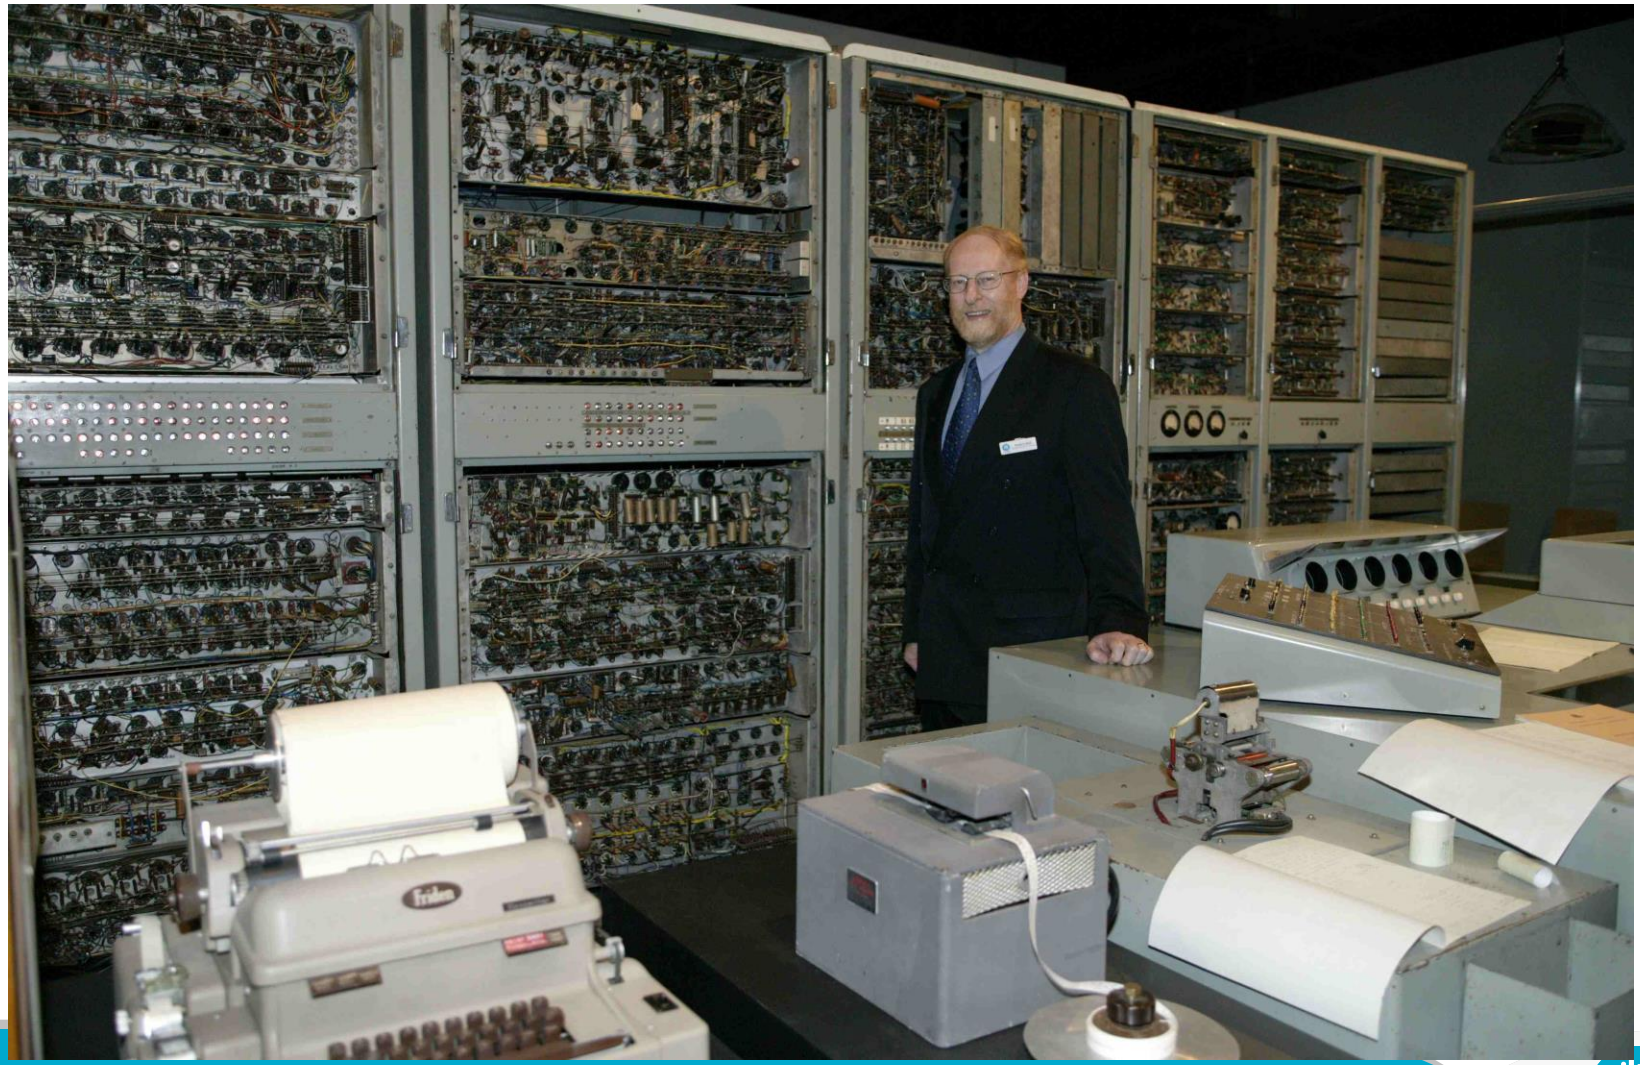


CRAY
CRAY
CRAY
CRAY
CRAY
CRAY
CRAY
CRAY
N US
CRAY

History of Compute performance in CSIRO

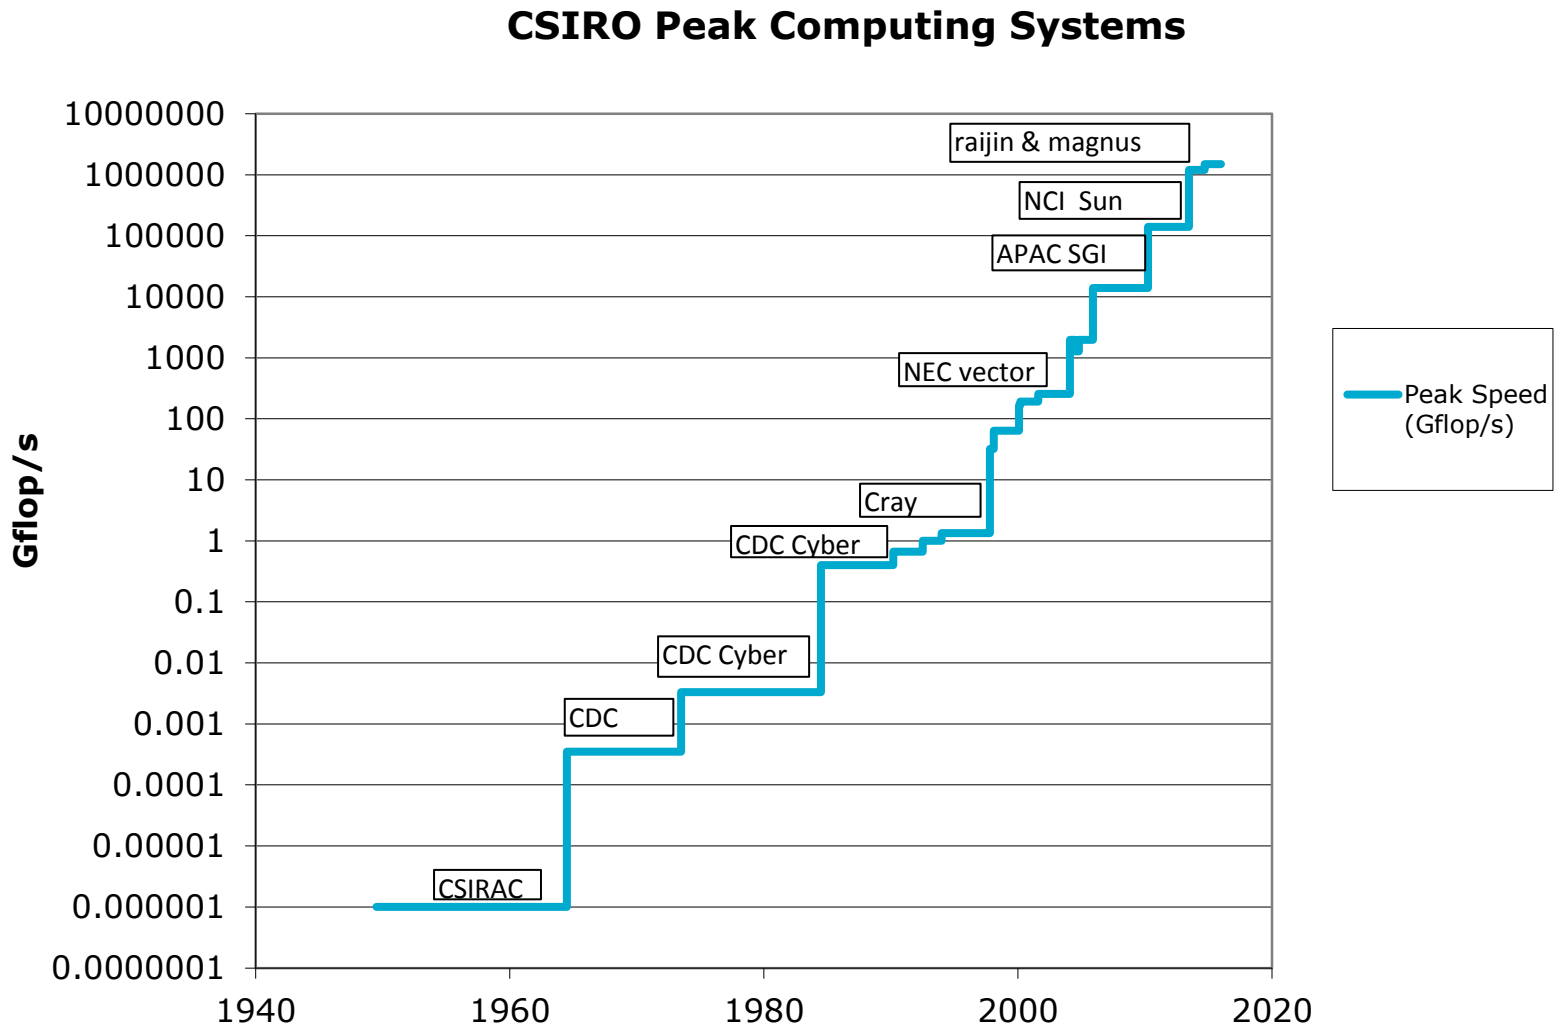

History of Computing in the Bureau

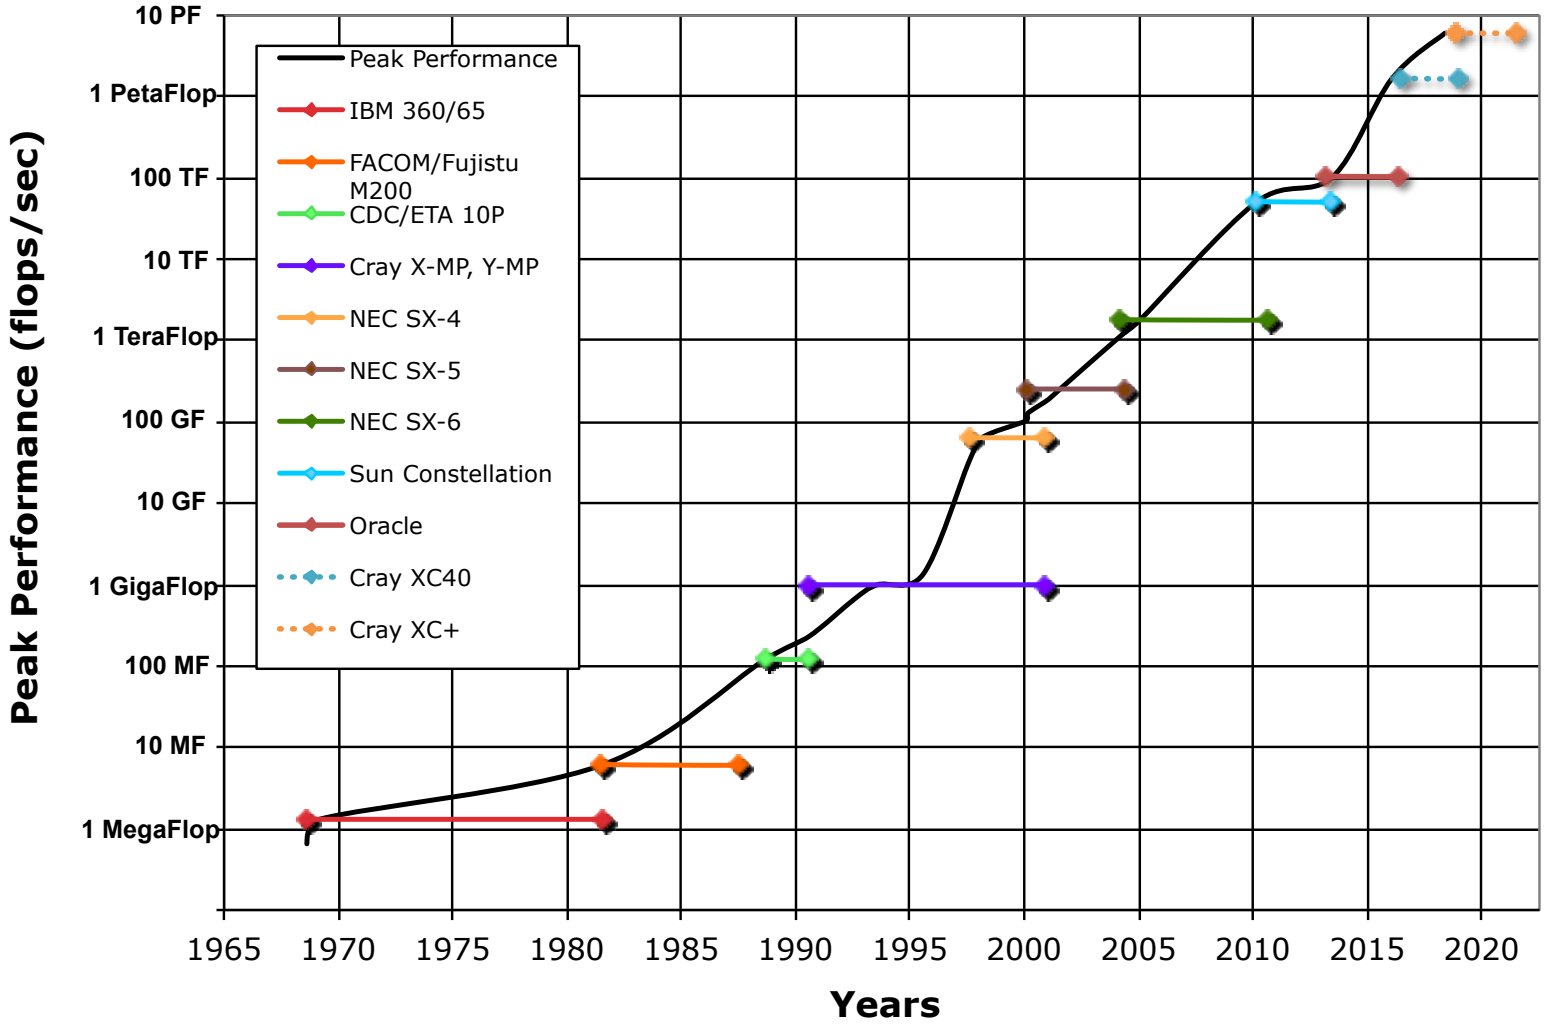

Early computing

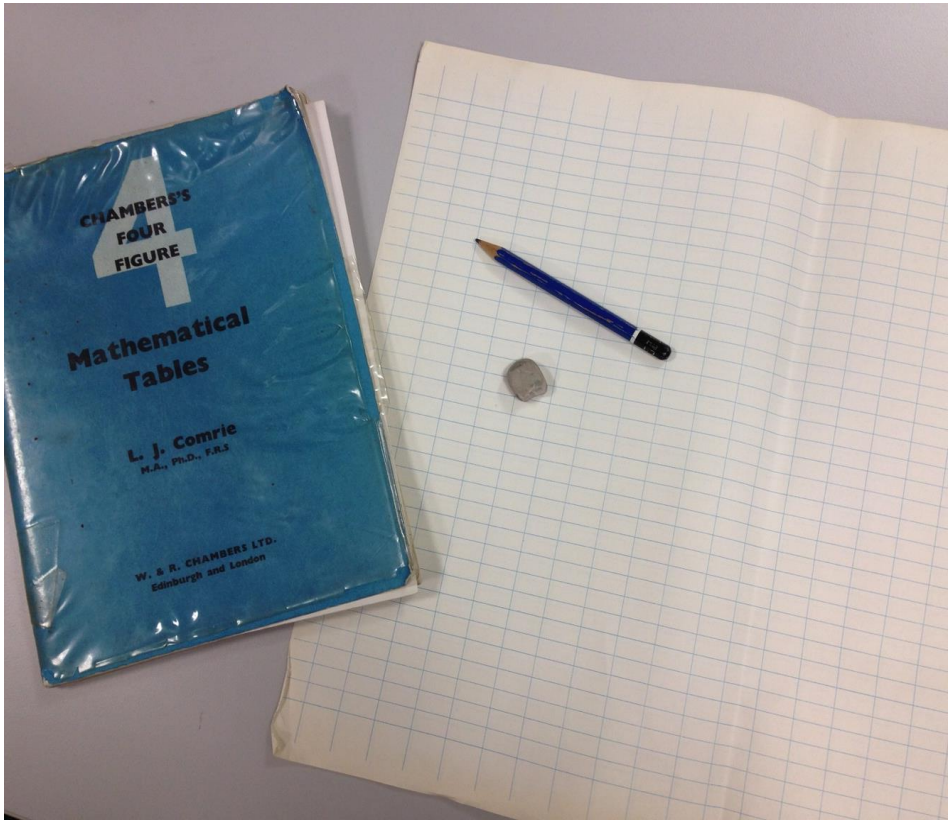

Courtesy Nigel Tout (www.VintageCalculators.com)

MADAS 20BZS circa 1950s-1960s

Early history

- CSIRAC in CSIR Radiophysics from 1949
  - Moved to University of Melbourne in 1956
  - Dick Jensen ran first numerical forecasts in Australia on CSIRAC
- CSIRO Computing Research – CDC 3600 (used by Bureau until IBM 360/65 systems in 1968)
- CSIRO: Csironet: CDC Cybers, then post-Csironet Crays
- Bureau: FACOMs, ETA, Crays
- HPCCC: NEC

CSIRO HPC Systems

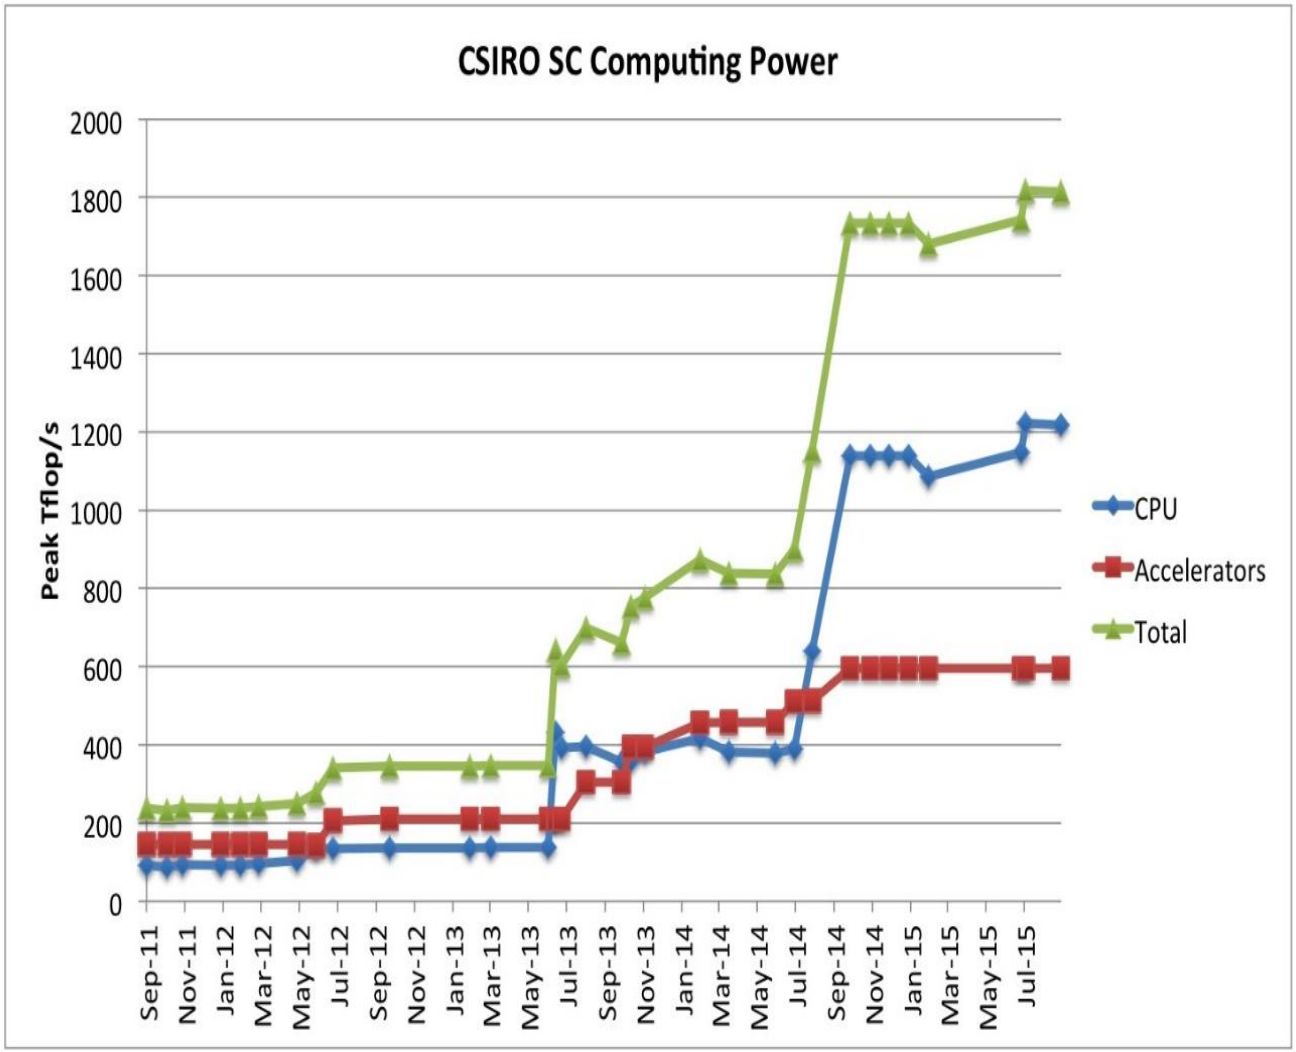

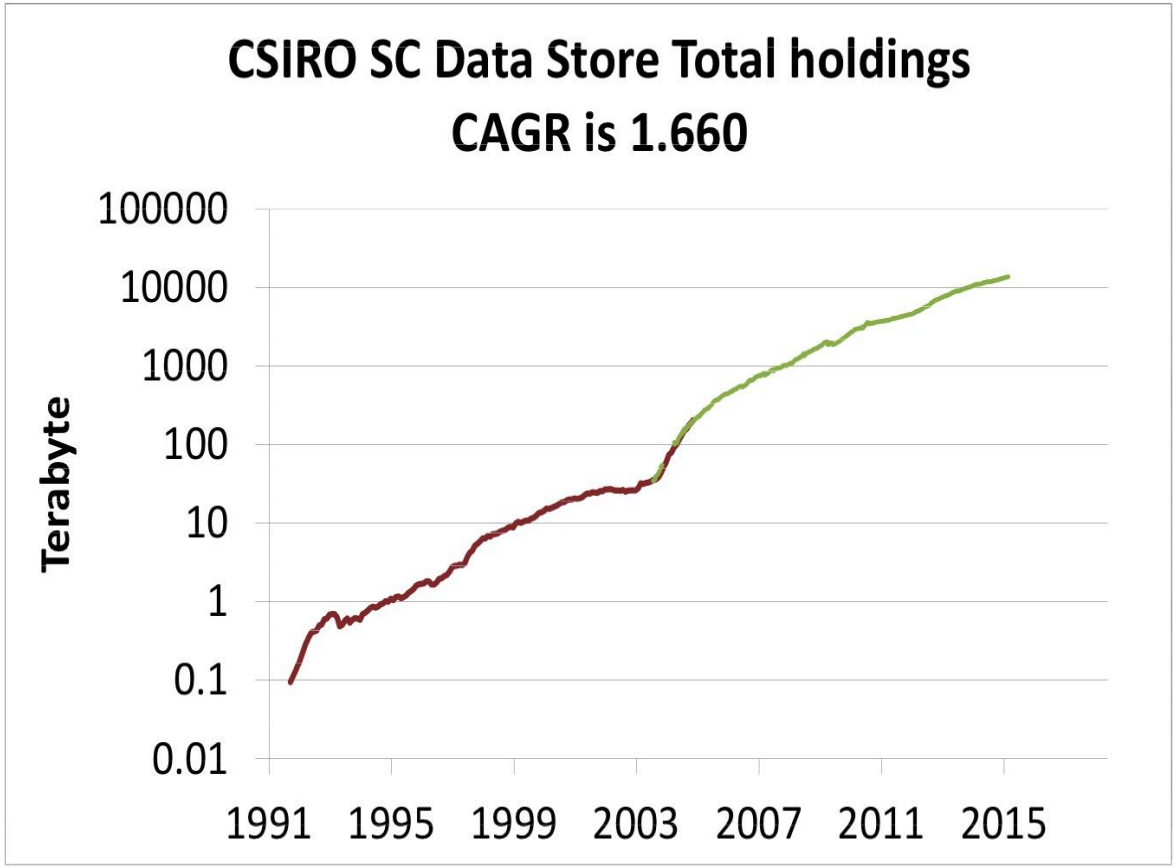
CSIRO SC Data Store Total holdings
CAGR is 1.660
Terabyte
100000
10000
1000
100
10
1
0.1
0.01
1991
1995
1999
2003
2007
2011
2015

History

Provision of:

- Systems – compute
- Storage – HSM
- Software – Open Source treasury
- Services – Help Desk, backups, etc
- Support – people with knowledge

The present: Bureau's Cray system

- New data centre
- Cray XC40 – widely adopted in the community
- Features:
  - Intel Xeon Haswell processors:
    - FMA3: d=round(a x b + c) ; Variable clock speed
  - Dense packaging: 384 processors per cabinet: 90 kW
  - Special compute-node Linux version: low jitter
  - Water cooled
  - Two partitions: for resiliency: compute and storage
  - Aries interconnect:
    - 48-port routers, 500 Gbyte/s; dragonfly topology

Cray XC Rank1 Network

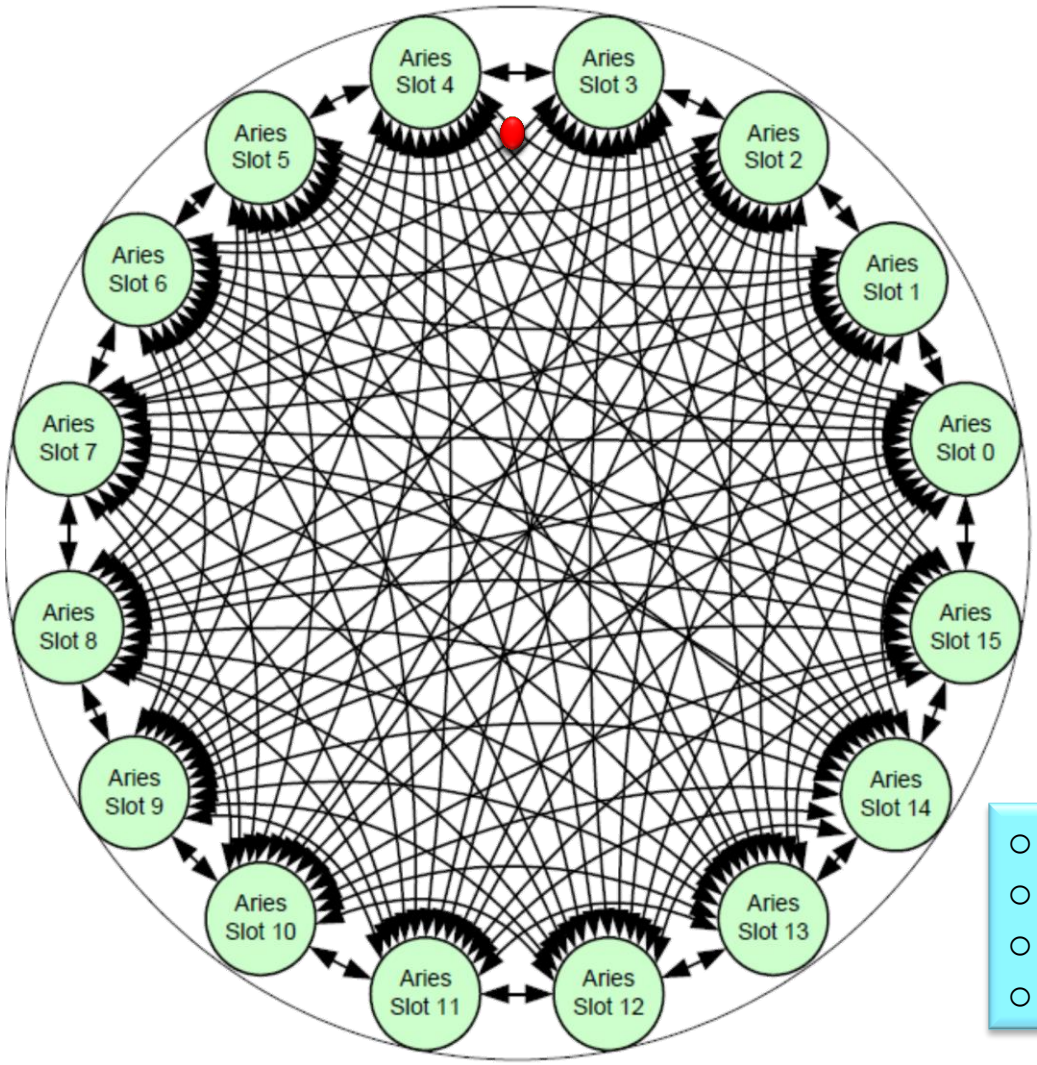

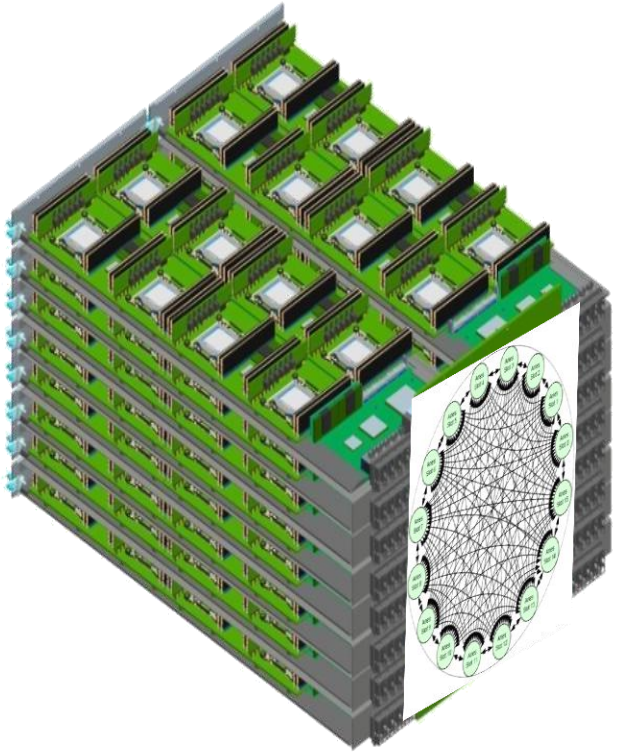

- Chassis with 16 compute blades
- 128 Sockets
- Inter-Aries communication over backplane
- Per-Packet adaptive Routing

Australis HPC

Parameter	Oracle HPC System	2015 System	2018 System	Relative Increase
Processor	Intel Xeon Sandy Bridge 6-core 2.5GHz	**Intel Xeon Haswell E5-2690v3 12-core 2.6GHz**	**Intel Xeon Haswell + Skylake**	**Increase relative to Oracle HPC System Ngamai**
Nodes	576	**2,160**	**4,112**	**2015: 3.8x 2018: 7.1x**
Cores	6,912	**51,840**	**129,920**	**2015: 7.5x 2018: 18.8x**
Aggregate Memory	36.9 TB	**276 TB**	**651 TB**	**2015: 7.5x 2018: 17.7x**
Global Filesystem Technology	Oracle Lustre 1.8.8	**Sonexion Lustre 2.5.1+**	**Sonexion Lustre 2.x**	
Usable Storage	214 TB	**4,320 TB**	**8,640 TB**	**2015: 20.3x 2018: 60.5x**
Storage Bandwidth	16 GB/s	**135 GB/s**	**306 GB/s**	
Compute Interconnect	Mellanox Infiniband QDR 40Gb/s	**Cray Aries 93 – 157Gb/s**	**Cray Aries 93 – 157Gb/s**	
Typical Power Use (kiloWatts)	200 kW	**865 kW**	**1,648 kW**	**2015: 4.3x 2018: 8.2x**
Sustain system performance (SSP)	**16**	**253**	**618**	**2015: 15.6x 2018: 38.1x**
Top 500 R_{max} Linpack (TF)	104	**1663**	**5000+**	**2015: 16.0x 2018: 50.0x**

Operational Capacity
(Half of total capacity)

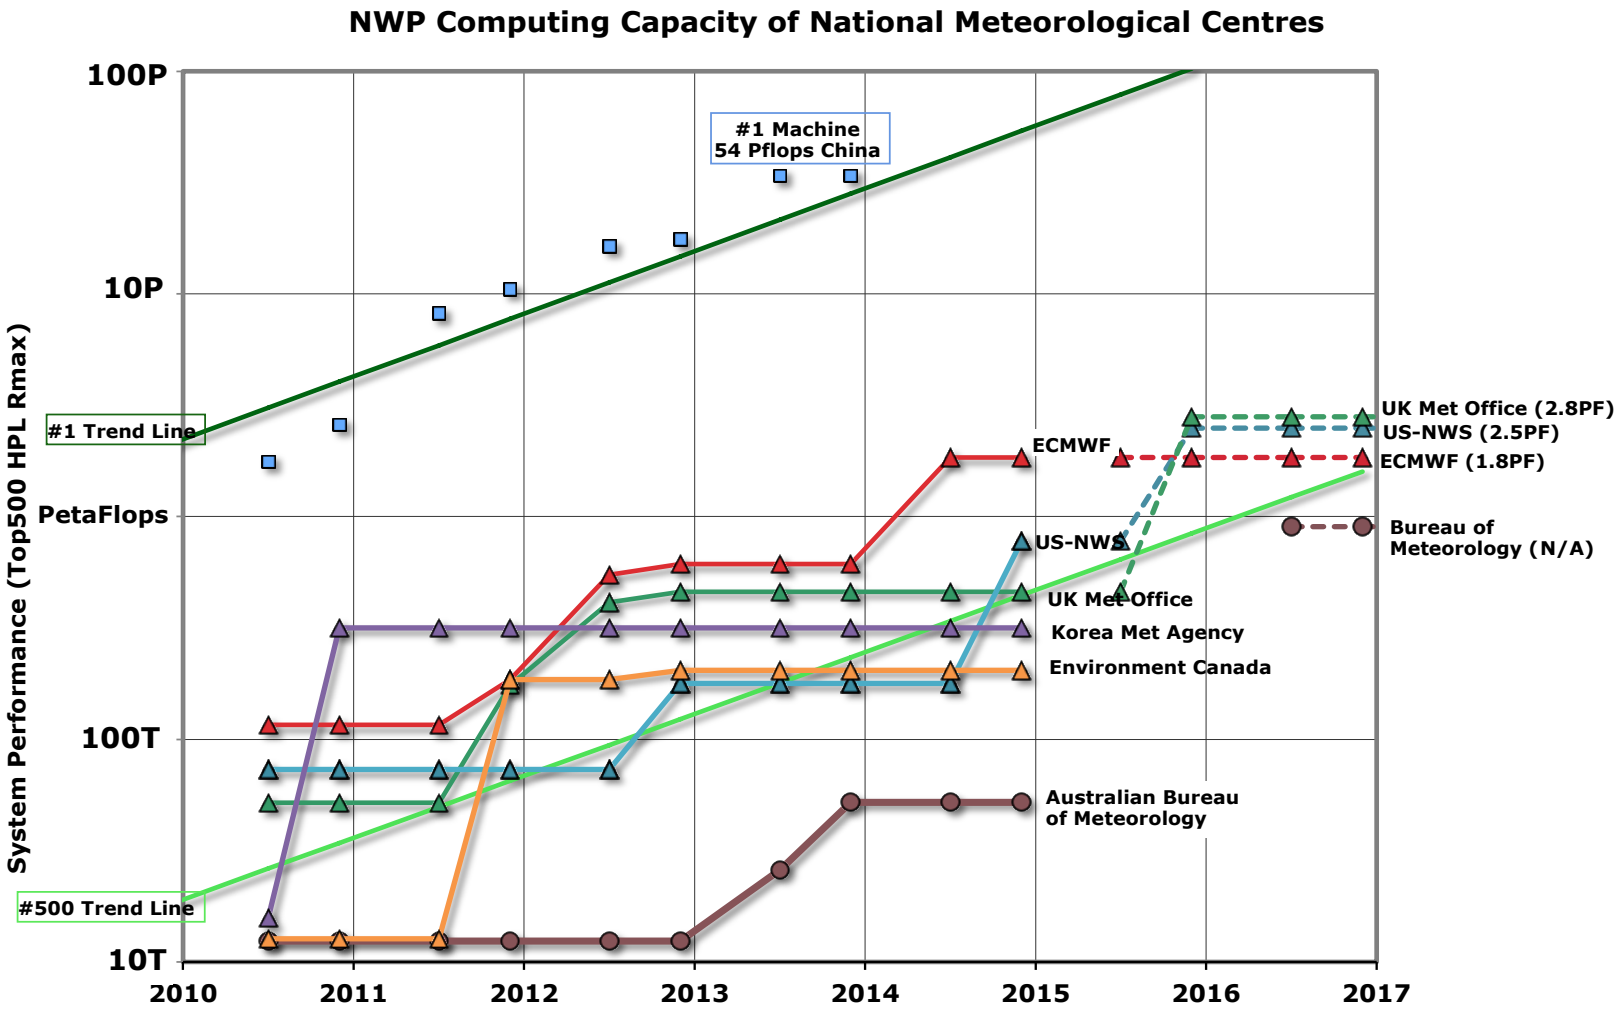

Future possibilities

- Systems – no more easy gains?
- Storage – cloud?
- Software – narrowing of options because of architectures
- Services – computer science, collaboration, visualisation
- Support – people with knowledge

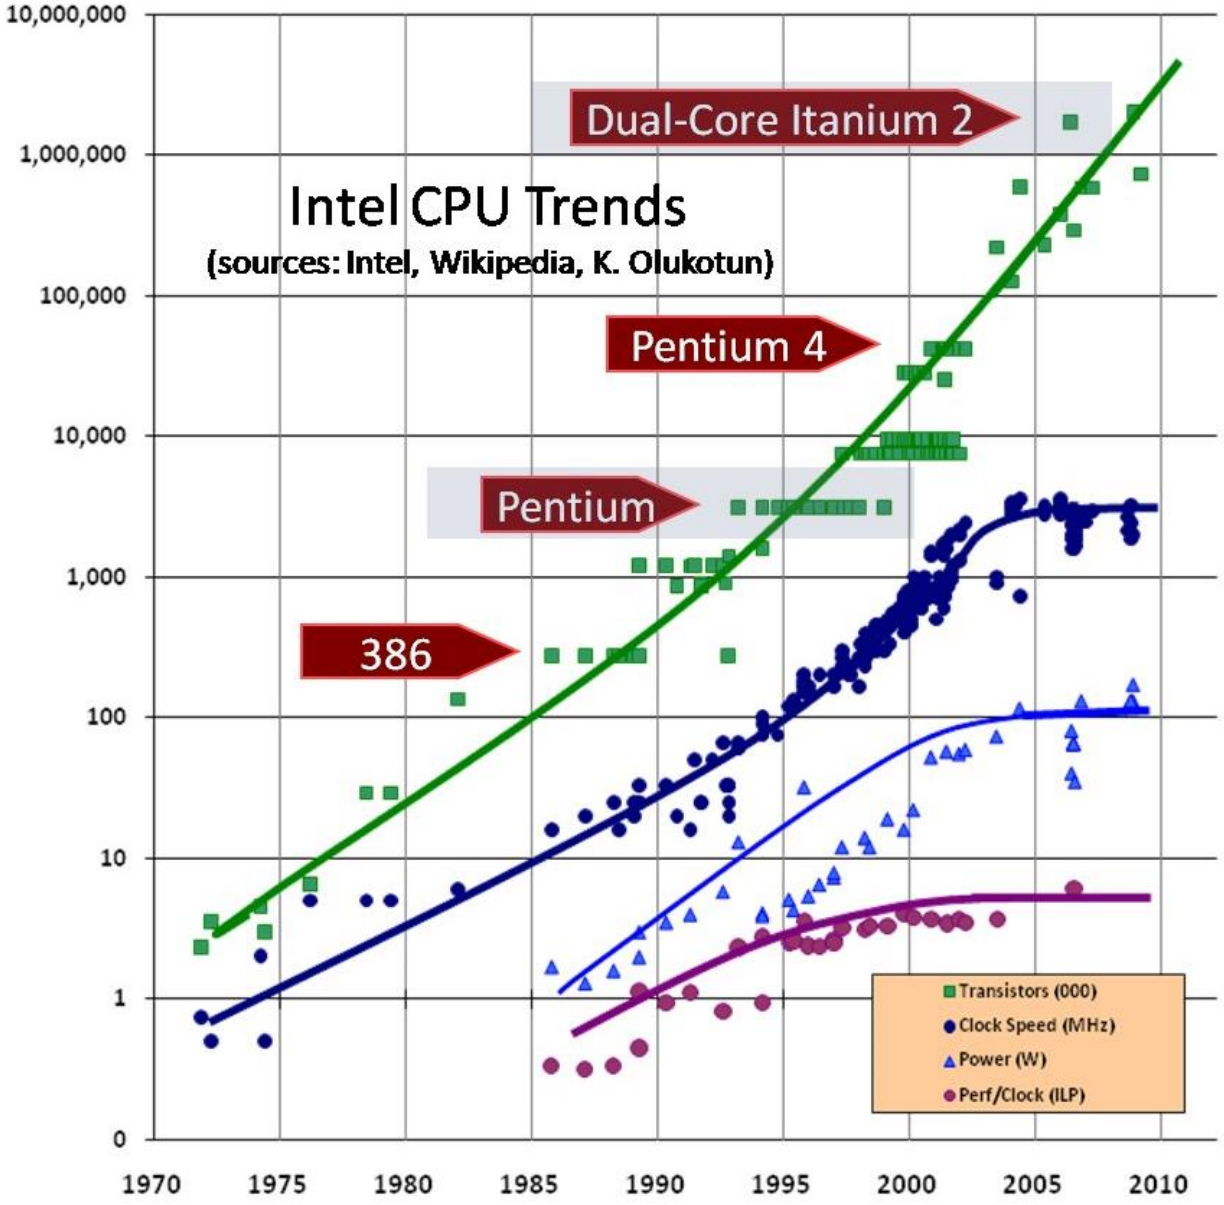
Intel CPU Trends
(sources: Intel, Wikipedia, K. Olukotun)
Dual-Core Itanium 2
Pentium 4
Pentium
386
Transistors (000)
Clock Speed (MHz)
Power (W)
Perf/Clock (ILP)
10,000,000
1,000,000
100,000
10,000
1,000
100
10
1
0
1970
1975
1980
1985
1990
1995
2000
2005
2010

Stream2 - 2004

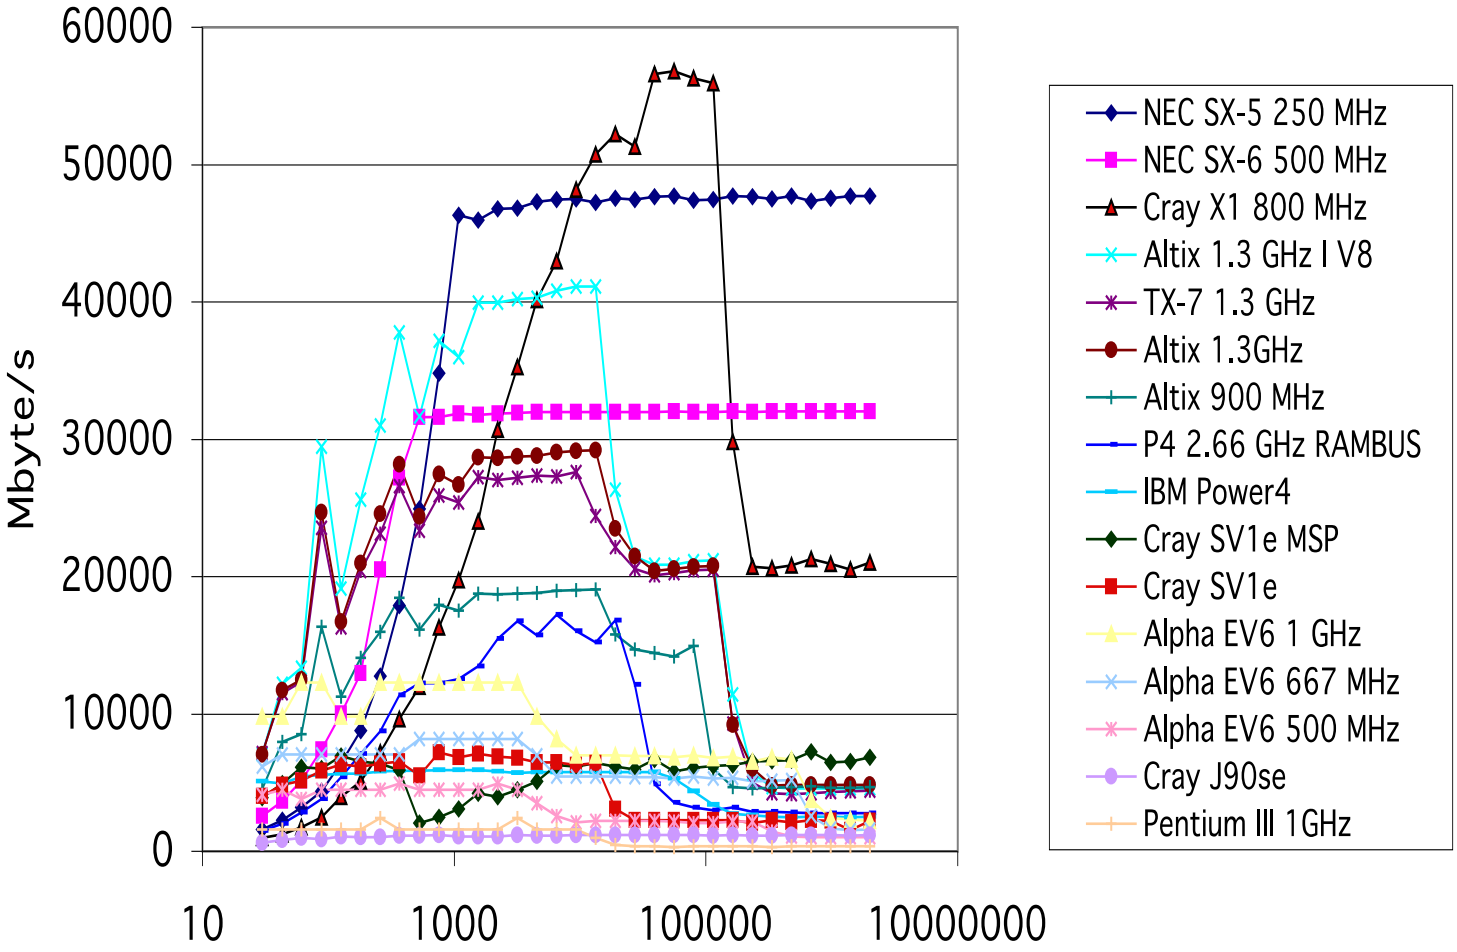

Future systems

- Systems – more cores, not much faster?
- Limited bandwidth gains
- Accelerators: for some applications
  - CSIRO accelerated computing program – GPU systems – some success
- Commodity market – ARM processors?

Future Storage

- Ten years ago, I wrote:
  - Users typically want every file kept and backed-up, and would be happy to use only one file system, globally visible across all the systems they use, with high-performance everywhere, and infinite capacity!
  - A user added that they want all of the above, at zero cost!
- Storage – cloud? Regulatory, cost and protection issues
- Storage for HPC – Lustre, fast, parallel, but not very reliable.
- Object stores rather than filesystems?
- Tape persists, even in cloud
- CSIRO Cloud storage

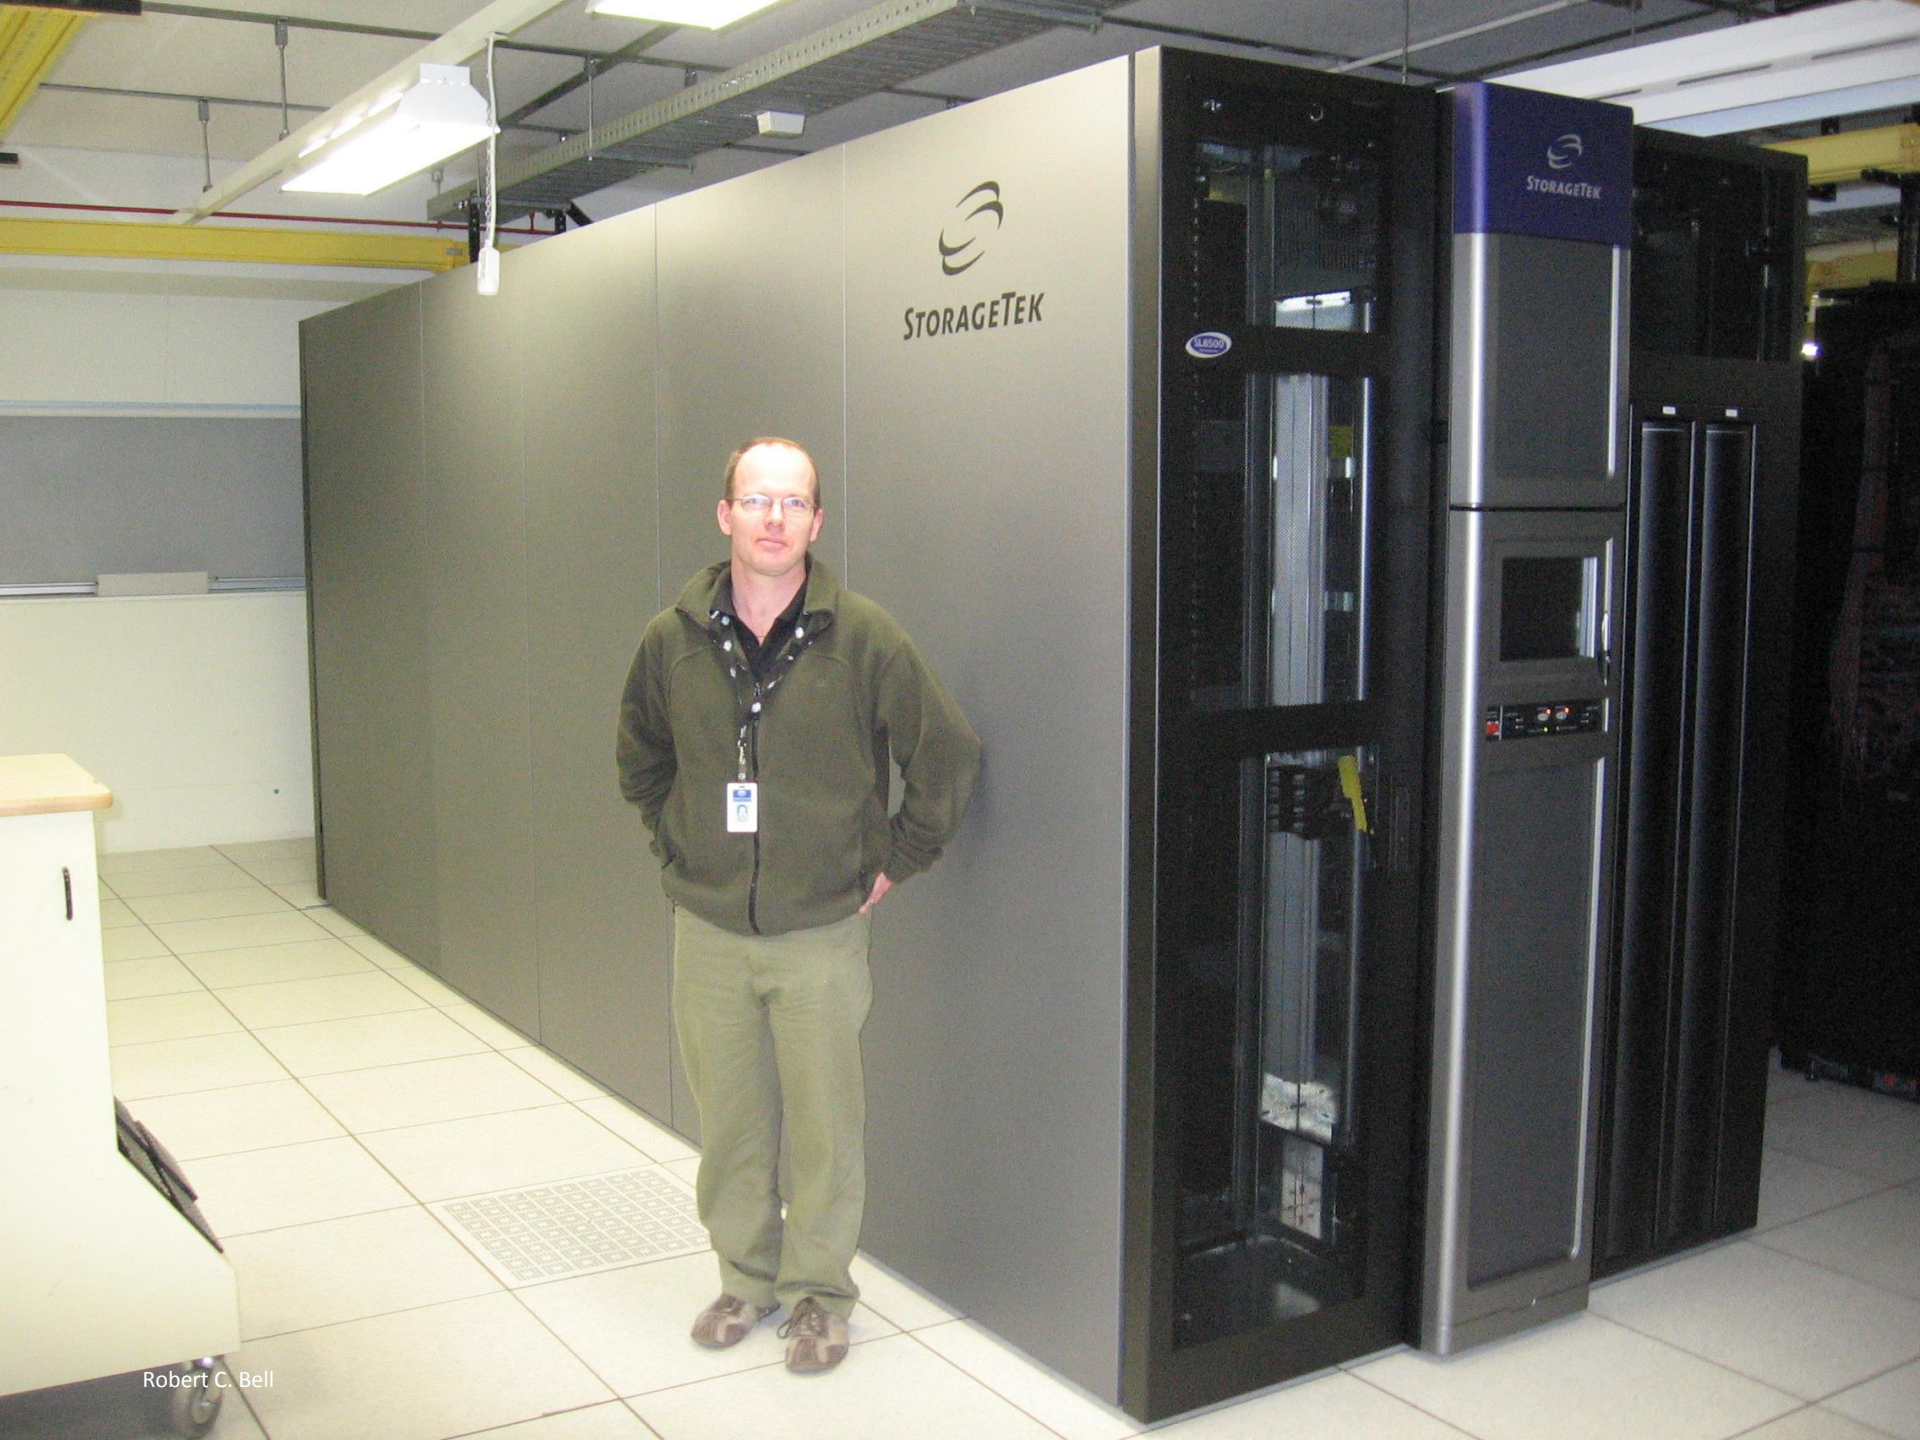

Robert C. Bell


06+
STORAGETEK

Future: Software

- Explosion in consumer software

  but, in HPC, narrowing of options because of architectures
- Massive investment in major models
- Collaboration essential
- CSIRO eResearch model

E-enablement

equipping the organisation to perform activities more efficiently through electronic means. The e-enablement layer includes general computing and storage, networks and office productivity tools both on-premise and remote

- Consolidate infrastructure
- Increase reliability/availability
- Achieve economies of scale
- Virtualise

eResearch

the rapid evolution of research methodologies enabled by information technologies and tools. The eResearch layer includes scientific computing and storage, research data management, visualisation, advanced collaboration technologies and productivity tools.

- Expand scientific computing
- Embed enterprise-level research data management
- Extend collaboration/e-tools
- Develop visualisation services

IM&T eResearch Supports the Entire Science Data Workflow

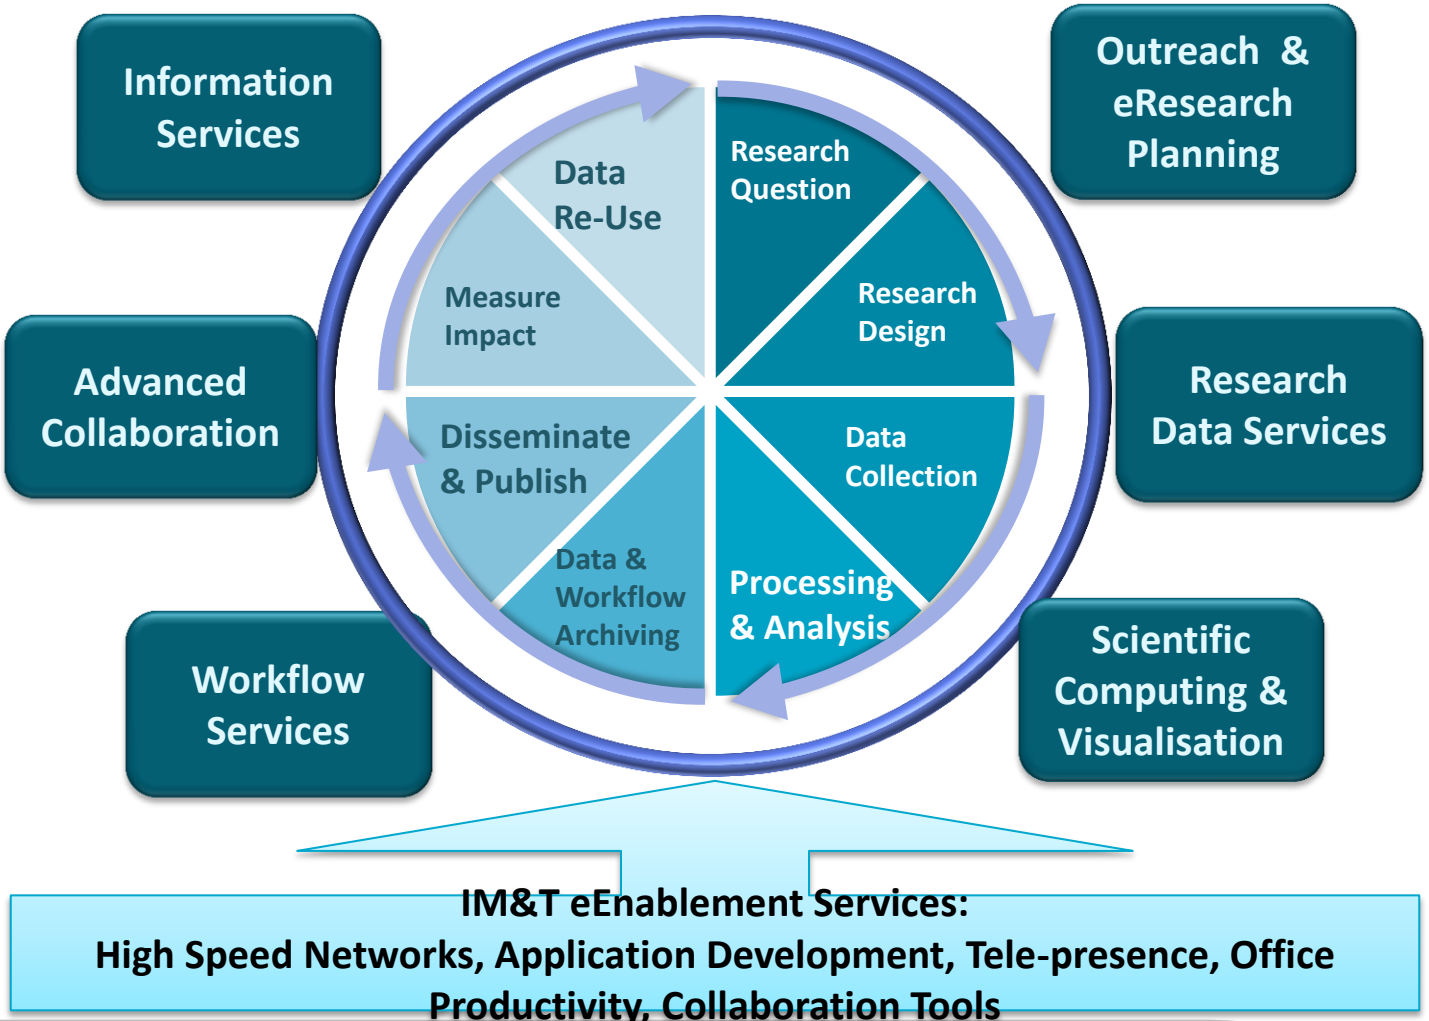

eResearch Project Services

- **Submission History**

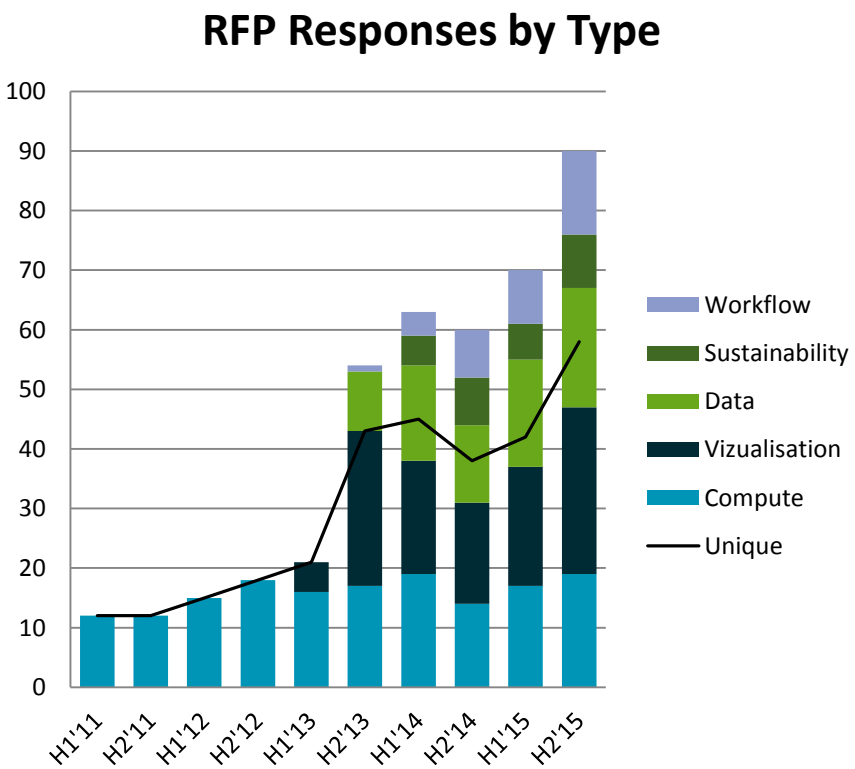

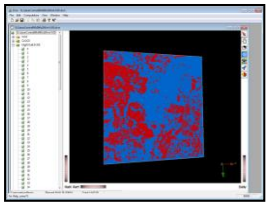

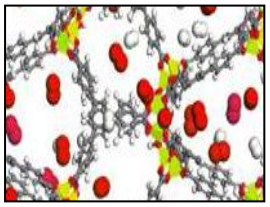

Future: services

- File protection (backup)
- Computer science
- Collaboration
- Visualisation
- Buzz words: e.g. Big Data

Future: support

- People with knowledge
- Critical mass
- Training – little HPC knowledge from undergraduates (Fortran?)
- Staff with commitment to science support
- Bureau: Scientific Computing Services

Conclusion: From CSIRAC to Cray (again) and onwards

- We've come a long way in computing for atmospheric and ocean sciences
- Major acquisition by Bureau for operational computing
- Hard to see future, but we know some dead-ends
- Forecast models running on phones?

Thank you

CSIRO IM&T Scientific Computing Services
Robert C. Bell

t +61 3 9669 8102 | +61 3 9545 2368
e Robert.Bell@csiro.au
w www.hpsc.csiro.au