Review of CSIRO's Capability in Advanced Scientific Computing

Brennan Arrold
Justin Baker
Marcus Foster
Tad Matuszkiewicz
John Morrissey
David Needham
Alf Uhlherr

7 November 2007

CSIRO Information Management & Technology

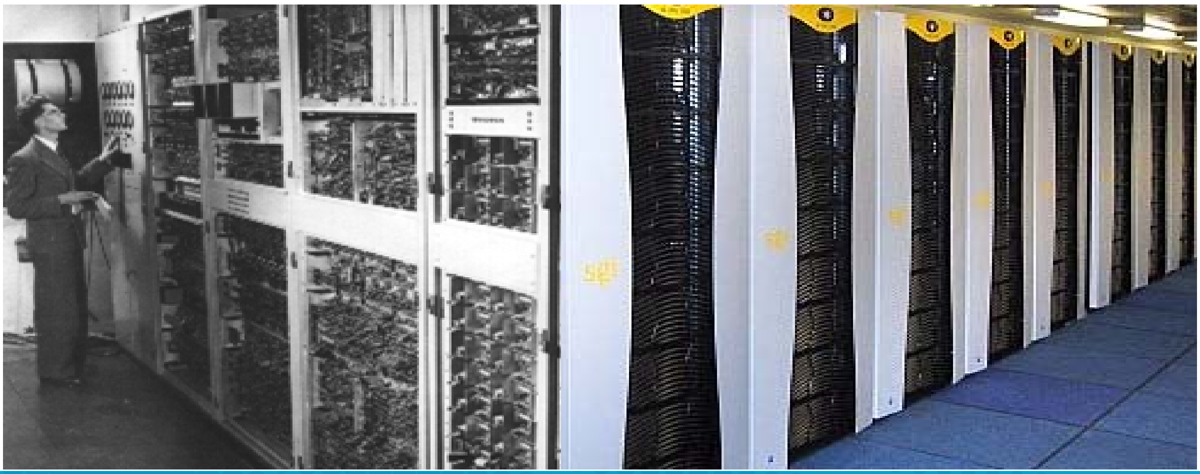

Version 1.50A

Executive Summary ☒ Main Report ☒ Appendices ☒

Enquiries should be addressed to:

David Toll
Director, Information Property and Services
CSIRO
PO Box 225
Dickson ACT 2062
(02) 6214 2933
David.Toll@csiro.au

Copyright and Disclaimer

CONTENTS

EXECUTIVE SUMMARY 6

Background 6

A Call to Action 6

External and Internal Influences 7

Benefits of an Enhanced Program for Advanced Scientific Computing 7

Options Considered 8

Recommendations 8

INTRODUCTION 10

What is Advanced Scientific Computing? 10

Why Is Advanced Scientific Computing Important? 10

Scientific Computing Interrelationships 11

Types of Facilities and Services 13

Review Process 14

ADVANCED SCIENTIFIC COMPUTING IN CSIRO 15

Current Computational Science Activities 15

Current Scientific Computing Services 15

Current Shared Facilities 16

Current Local Facilities 16

Current CSIRO Use of Partner Facilities 17

REVIEW FINDINGS 18

Comparison with Peer Organisations 18

Comparison with National Partners 21

Diversity of Requirements 21

Adoption of Local and Centralised Facilities 22

Projected CSIRO Demand 23

Constraints on Uptake 24

Opportunities for CSIRO Science 25

RECOMMENDATIONS 27

Options 27

Proposed Capability 29

Structure and Governance 30

Provision of Services 30

Developing a Community of Practice 32

Securing a Financial Basis for Growth 33

APPENDICES 35

Appendix 1: Glossary 35

Appendix 2: History of High Performance Computing in CSIRO 43

Appendix 3: Technology Trends 47

Appendix 4: Cost of Ownership 54

Appendix 5: Funding and Allocation Models 58

Appendix 6: Utilisation of Advanced Scientific Computing in CSIRO 61

Appendix 7: Survey Results 68

Appendix 8: Workshop Summary 76

Appendix 9: External Organisational Benchmarks 80

Appendix 10: Sources 91

Appendix 11: List of Contributors 93

List of Tables

Table I. Investment in Advanced Scientific Computing by CSIRO and peer organisations. . 18
Table II. Service structure and delivery for APAC partners. 21
Table III. CSIRO investment options for Advanced Scientific Computing. 27
Table IV. Proposed fund sources for scientific computing services. 33
Table V. Utilisation economies for shared HPSC systems. 52
Table VI. Compute and Infrastructure Costs. 54
Table VII. Storage Costs for 2007. 55

List of Figures

Figure 1. Initiatives associated with scientific computing in CSIRO. 12
Figure 2. Hierarchy of scientific computing systems. 13
Figure 3. Comparison of hardware performance between CSIRO and peer organisations. . 19
Figure 4. Apportioning of computing budget according to science discipline. 20
Figure 5. Projected opportunities for large scale scientific computing (NCCS). 26
Figure 6. Proposed Structure for CSIRO's ASC capability. 29
Figure 7. CPU utilisation of project-dedicated and shared access clusters. 52
Figure 8. Utilisation of HPSC SGI Altix shared memory system. 64
Figure 9. Utilisation of HPSC/HPCCC NEC SX-6 parallel vector system. 65
Figure 10. Utilisation of HPSC IBM BladeCenter cluster system. 65
Figure 11. Utilisation of CSIRO Data Store HFS system. 66
Figure 12. CSIRO utilisation of iVEC SGI Altix shared memory system. 66
Figure 13. Utilisation of CSIRO e-Science Condor desktop cycle harvesting system. 67
Figure 14. Predicted balance of CPU requirements. 73
Figure 15. Predicted balance of storage needs and annual growth, with & without CMAR. . 74

EXECUTIVE SUMMARY

AT A GLANCE

- Advanced Scientific Computing is critical to the future of CSIRO science.
- CSIRO investment in Advanced Scientific Computing has stagnated during the last decade, with detrimental impact to science outcomes.
- This review recommends an incremental growth strategy, leading to an approximate doubling of the current annual expenditure over the next three years.
- The recommended focus is on people, processes and competencies, with external and internal partnerships being used to develop computing infrastructure, information management and collaborative research tools.

Background

Computational science is an essential and growing component of all fields of scientific and industrial research. CSIRO recognizes this through the recent selection of computational and simulation science as one of four transformational capability platforms for increased science investment in the *CSIRO Strategic Plan for 2007-2011*.

This science is heavily reliant on Advanced Scientific Computing (ASC) facilities and services, beyond those available as part of standard desktop PC infrastructure. Such facilities currently form the "high end" of the portfolio of services delivered by CSIRO Information Management & Technology (IM&T). ASC includes compute servers, "cluster" systems, high performance workstations, large/fast data storage and visualisation systems, managed locally within Business Units, at shared facilities or through partnerships within the national innovation system. Within this portfolio, ASC services derive the most benefit from the spectacular improvements in price-performance that underpin the Information Technology revolution, require the highest levels of technical agility, proficiency and user engagement, and contribute the highest level of benefit to specific research outcomes.

Advanced Scientific Computing is especially important to CSIRO as a research enterprise, because:

- It is now, or will become, the principal means for characterising and tackling the biggest technological challenges facing Australian and global society, such as climate change, sustainability, human health, biodiversity and security
- It is the one category of specialised research tool that is useful to all scientists in CSIRO, spanning and uniting the enterprise's unique geographical and discipline diversity
- Current compute-intensive research fields, such as earth systems, astronomy, minerals and genetics, are areas of CSIRO strength and particular national significance
- ASC collaborative "e-Research" platforms are becoming vital to the operation of large agile interdisciplinary teams, which are an essential component of CSIRO's differentiation advantage.

A Call to Action

Recent comparisons with analogous facilities for peer organisations in the Global Research Alliance show the following trends:

- CSIRO's commitment to ASC has been relatively stagnant over the past 2 decades, with investment levels remaining constant, while those of peer organisations have increased substantially
- Utilisation of ASC is uneven across CSIRO's research portfolio, with much lower levels of adoption in materials, energy, chemical and life sciences than in comparable peer organisations
- CSIRO staffing levels for support and outreach activities have remained at a comparatively low level, which has constrained ASC adoption and demand for the existing facilities to a sub-optimal level.

Consequently, CSIRO's investment in ASC is approximately half of the expected level. Lack of action will cause this gap to increase in future, making it ever more difficult to bridge.

External and Internal Influences

This review of CSIRO's commitment to ASC coincides with the following events:

- Initiation of the NCRIS 5.16 Platforms for Collaboration capability, which provides opportunities for leveraging investments in computing infrastructure, data management and collaborative services, as well as increasing engagement with the Australian research community
- Renewal of the High Performance Computing and Communications Centre (HPCCC) joint venture with the Bureau of Meteorology, which forms the cornerstone of the new Centre for Australian Weather and Climate Research (CAWCR) joint research organisation
- Initiation of a number of complementary activities within CSIRO, notably the computational & simulation science platform (methodology and applications), the Terabyte Science and e-Research Themes (underpinning technology), and the e-Science Information Management framework (policy and implementation)
- Requirement for a sustainable hardware refresh program, at a time when the majority of ASC infrastructure in Divisional and partner facilities is approaching end of life.

Benefits of an Enhanced Program for Advanced Scientific Computing

Increased investment in ASC will not only increase CSIRO's overall impact, it will lower overall operating costs, by increasing efficiency and reducing demand for other capital expenditure on non-computational research infrastructure.

Specifically, the main benefits of growing ASC are that this would:

- Maintain CSIRO competitiveness in established areas of strength
- Enhance CSIRO's reputation and competitiveness as a whole, by increasing the focus on a growth area that places maximum emphasis on the technical expertise of the researchers
- Create opportunities for new science and new areas of scientific leadership
- Facilitate growth of computational science as a capability platform for CSIRO
- Provide greater value from existing data sets, collections and acquisitions, to complement and selectively replace capital investment in laboratory instrumentation
- Create infrastructure and environments for increased collaborative and interdisciplinary research, both internally and with the wider research community, in Australia and internationally.

Options Considered

The principal options for CSIRO's future investment in ASC are:

1. Status Quo – maintain current structure and investment levels

2. Incremental growth –increase in investment staged over a three year period, with greater emphasis on partnering, collaboration and science support

3. Rapid growth – significant increase in investment to bring CSIRO up to par with peer organisations immediately.

Option 1 is the easiest and least expensive to implement. However, this course of action fails to meet existing demand, nor addresses latent demand from new application areas, leading to ongoing and potentially disastrous erosion of CSIRO's competitiveness and impact.

Option 3 is highly proactive and represents a strong commitment to growth in ASC. However, this option is expensive, disruptive, and runs the risk of wasting investment on underutilised resources if CSIRO researchers are unable to adapt rapidly to a large increase in computing capability.

Option 2 offers an optimum position between the above two extremes, and forms the main basis for the detailed recommendations from this review.

Recommendations

The recommendations in this review may be summarised as follows:

Structure and Governance:

- Position ASC as a CSIRO enterprise capability with hardware capacity, e-Research fabric, user support and science applications levels, where the base levels of this capability are located primarily within CSIRO IM&T
- Establish an ASC Steering Committee as a cross-business authority to set directions and priorities.

Provision of Services:

- Leverage partnership opportunities, including participation in the NCRIS National Computing Infrastructure initiative and renewal of the HPCCC partnership agreement
- Develop internal shared cluster and cycle harvesting capability
- Expand data storage capacity
- Develop specialist staff to support ASC across CSIRO.

Community of Practice:

- Develop a significantly expanded outreach program
- Foster an Advanced Scientific Computing culture across CSIRO.

Securing a Financial Basis for Growth:

- Grow ASC capability over the period 2007-2011 through incremental investment growth
- Develop procedures for research groups to specify dedicated large-scale computing requirements for funding through the Science Investment Process
- Implement annual rolling capacity and workforce plans.

The proposed level of investment to meet these requirements is as follows:

Current year	Year 1	Year 2	Year 3
$4.8M	$7.2M	$8.6M	$9.5M

The major benefits of this alteration in CSIRO’s investment balance will be to allow a larger proportion of CSIRO to adopt cost-effective computational science solutions, and to increase our ability to undertake large-scale collaborative research projects through internal, national and international partnerships.

INTRODUCTION

What is Advanced Scientific Computing?

Computational science and scientific computing are complementary concepts. *Computational science* is simply science performed on computing systems, or more generally using information technology (IT) systems. Implicit in this definition is that what is happening on the computer is novel. For example, routine processing operations performed on experimental or instrumental data, typically using a commercial software package, do not come under the definition, whereas developing new data processing methods does.

Computational science is sometimes referred to as *e-Science* (or more broadly as a form of *e-Research*), although these terms more accurately refer to the use of IT platforms to enable large scale collaborative projects, typically involving substantial computational and/or data requirements. The labels computational science and e-Science are used to distinguish from "wet science", where the new information is entirely derived from laboratory or field measurements.

IT has enabled the rise of numerical experiment, or *simulation*, to complement theory and observation (laboratory experiment) to form the three basic components of scientific research. Computation is also often crucial in evaluating complex theories, in extracting knowledge from large datasets (*informatics*) and in evaluating, correlating and combining these different forms of scientific investigation.

Scientific computing refers to the computing facilities, environment and services needed to do computational science and e-Research. These facilities may be largely the same as those used in a standard desktop for word-processing etc; a modern PC is more powerful than the world's fastest computer of only 15 years ago. Where the requirements exceed the capabilities of a commodity PC, they fall under the heading of *Advanced Scientific Computing (ASC)*.

As described in Appendix 3: Technology Trends, a key feature of computational science is that it is inherently *scalable*. In cases where an individual calculation takes a given amount of time – say 1 CPU second – it is typically possible to increase the scale of the calculation by an arbitrarily large factor (say 1 million) to obtain an equivalently large increase in science value, at only a moderate increase in cost or researcher effort. This scalability is frequently obtained by such approaches as *parameter sweeping* (or "what if" scenarios), parameter optimisation, *ensemble* calculations or changes in model resolution. Effective scalability allows computational research activities to expand to the size of the available resource. Thus, any scalable science application is by definition an application of advanced scientific computing. Equivalently, a scalable application can be termed a *high throughput* computation, in analogy with the notion of high throughput experimentation (HTE).

Why Is Advanced Scientific Computing Important?

The importance of scientific computing to the wider community is reflected by the following observations:

1. Modern Information Technology is frequently identified[1] as *the most important characteristic defining modern society*

[1] C. Livingston, CSIRO Chair, address to staff 2006.

2. The IT revolution is now widely recognised *as the most significant transformation in the history of scientific research*, as described for example in a recent special edition of *Nature*[2]

3. Based on the essential role of computational science[3] in identifying and rectifying current global challenges (climate change, sustainability, biodiversity, biosecurity, human health, artificial intelligence etc), one may conclude that *scientific computing will have a dramatic impact on society as a whole.*

This is why Computational and Simulation Science is now featured in Objective 2.1.2 in the CSIRO Strategic Plan 2007-2011, as one of four areas for Building Transformational Capability Platforms.

The performance of IT systems per unit cost continues to grow at an exponential rate. This continually increases the attractiveness of computational science in complementing and replacing wet science. Typical advantages include:

- Greater ability to access small (nanoscopic) and large (astronomical) size and time scales
- Greater level of detail accessible, especially regarding the interrelationships and mechanisms underlying observable quantities
- Frequently greater cost-effectiveness that is increasing more rapidly
- Greater scalability and throughput, enabling greater science benefit per researcher
- Fewer occupational health, safety and environmental hazards.

This growing utility and impact defines Advanced Scientific Computing as a "smart" choice for tackling current scientific challenges, where the level of impact is largely governed by the skill, creativity and focus of the organisation's individual researchers.

Advanced Scientific Computing is especially important to CSIRO as a research enterprise, for the following reasons:

- It is the one category of specialised research tools that is useful to all scientists in CSIRO, spanning and uniting the enterprise's unique geographical and discipline diversity
- Compute-intensive research fields such as climate, astronomy, minerals and genetics are areas of CSIRO strength and particular national significance, providing justification for substantial portions of CSIRO's budget
- Collaborative e-Research platforms are becoming vital in facilitating the kinds of large agile interdisciplinary teams (e.g. within Flagships) that define CSIRO's differentiation advantage.

Hence, it is difficult to overstate the importance of Advanced Scientific Computing to CSIRO's immediate and longer-term future.

Scientific Computing Interrelationships

Advanced Scientific Computing is a component of an interrelated set of capabilities and associated funding initiatives within CSIRO, which are illustrated in Figure 1. These include the following:

---

[2] "2020 Vision: How computers will change the face of science", *Nature* **440** (2006) 398 *et seq*.
[3] Benioff, M. and Lazowska, E. 2005. *Computational Science: Ensuring America's Competitiveness*; President's Information Technology Advisory Committee (PITAC) Report.

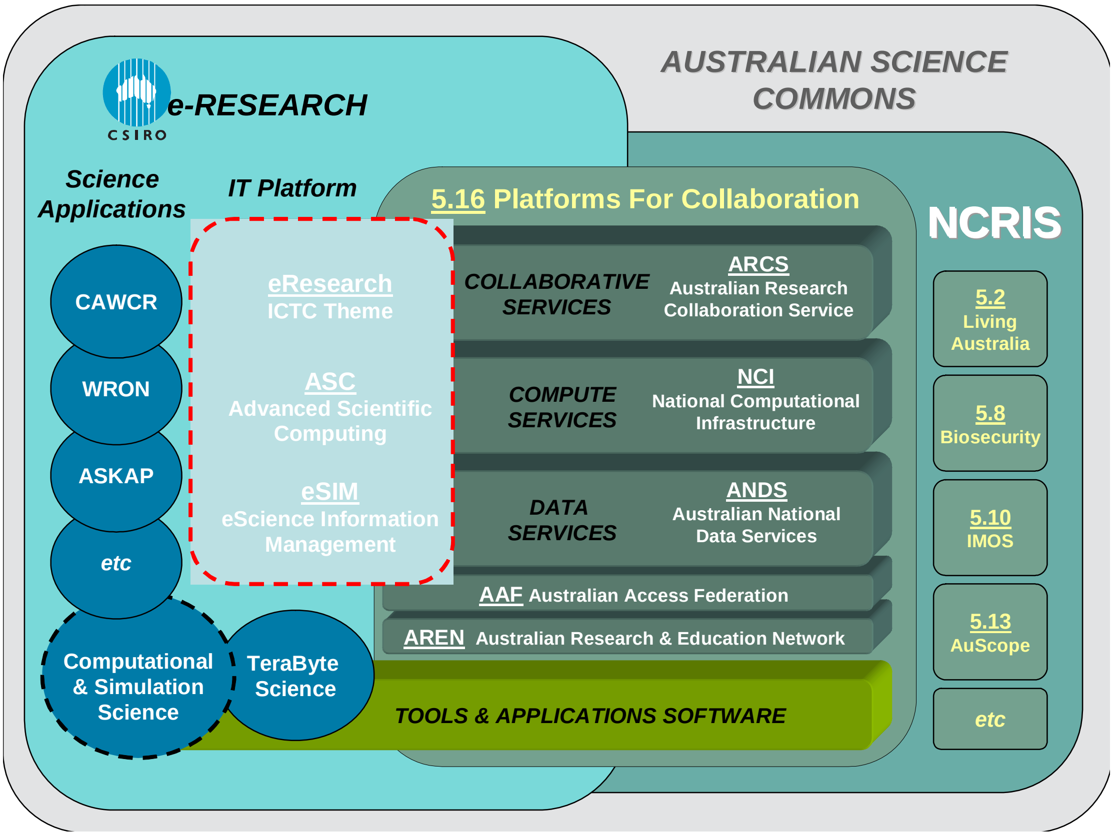

Figure 1. Initiatives associated with scientific computing in CSIRO.

- Computational & Simulation Science, now defined as one of CSIRO's four Transformational Capability Platforms for science investment in the CSIRO Strategic Plan 2007-2011
- CSIRO's engagement with Platforms for Collaboration (PfC)[4], which forms Initiative 5.16 in the DEST NCRIS program[5]. PfC is funded for $75M over 5 years to develop functions for:
  - National Computing Infrastructure (NCI)
  - Australian National Data Service (ANDS)
  - Australian Research Collaboration Service (ARCS)[6]
  - Australian Research and Education Network (AREN)
  - Australian Access Federation (AAF)[7]

  NCI provides opportunities for CSIRO to co-invest in large scale national "peak" and regional "shoulder" computing facilities suitable for CSIRO researchers. ARCS and ANDS in particular provide additional opportunities for CSIRO ICTC and IM&T to develop technologies and services of value to the national research community, which feature in CSIRO's submission to PfC. Note that ARCS is the joint venture agreement formed from the PfC Interoperation and Collaboration Infrastructure (ICI) project
- Other NCRIS initiatives that also require collaborative and IT infrastructure technology including 5.1 Evolving Bio-Molecular Platforms and Informatics, 5.2 Integrated Biological

---

4 http://www.pfc.org.au//cgi-bin/twiki/view.
5 http://www.ncris.dest.gov.au.
6 http://www.arcs.org.au/.
7 http://www.aaf.edu.au/.

Systems, 5.3 Characterisation, 5.8 Networked Biosecurity Framework, 5.10 Optical and Radio Astronomy, 5.12 Integrated Marine Observation System and 5.13 Structure and Evolution of the Australian Continent

- The CSIRO e-Science Information Management (eSIM) strategy, which is currently being initiated through the eSIM pilot projects
- The CSIRO IM&T Foundation Program, which is engaged in renewing much of CSIRO's server, storage and network infrastructure.

An e-Research partnership is currently being developed between CSIRO IM&T and the Information and Communications Science and Technology (ICST) Group, whose Divisions are actively developing a number of related enabling technologies, notably through the Themes in e-Research, Terabyte Science and the Australian Square Kilometer Array Pathfinder (ASKAP).

Types of Facilities and Services

Advanced Scientific Computing typically encompasses the following components:

- Compute servers
- Data storage
- Networking, communications and visualisation infrastructure
- Systems, applications and development software
- Support staff.

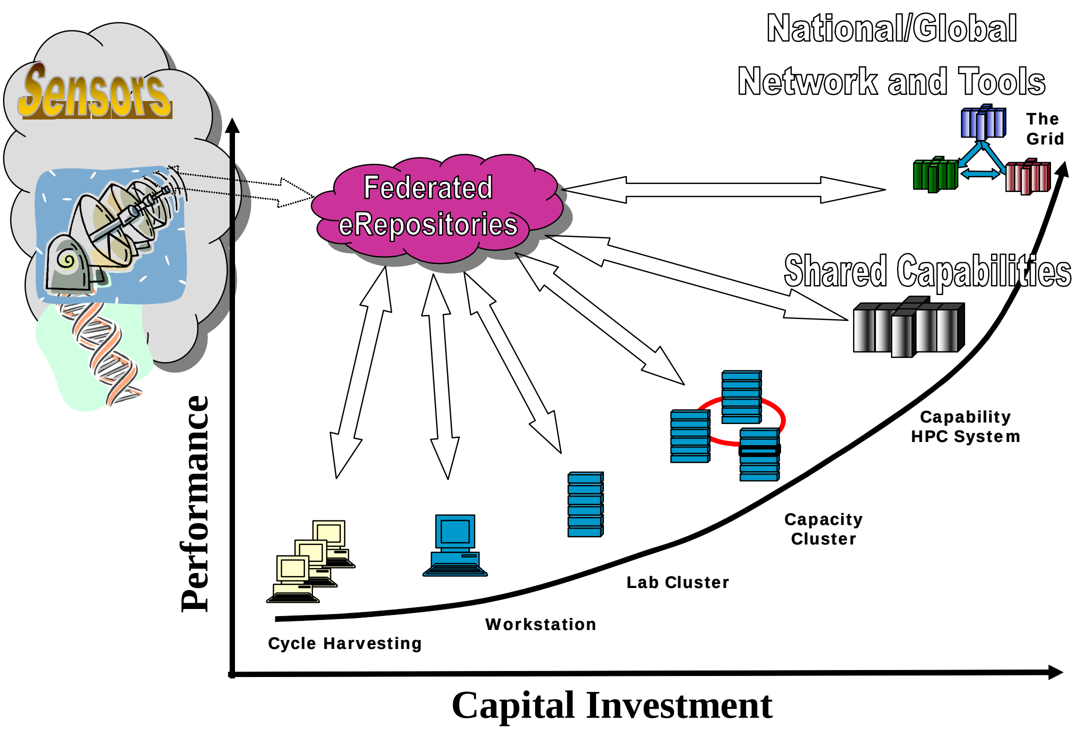

Figure 2. Hierarchy of scientific computing systems.

These components may further be based *locally*, i.e. collocated with the users, or consolidated into managed *central* facilities. They may also be shared with research or corporate partners, on a project, state, national or international scale. The label *High Performance Computing* (HPC) is widely used for large-scale facilities.

The options for these components can be considered as a *hierarchy*, as outlined in Figure 2. As a diverse research enterprise, CSIRO has different users operating across all levels of this hierarchy. At times users may need to move their activities up or down these levels.

CSIRO has requirements for both capability and capacity HPC systems. Capability systems are designed to tackle a comparatively small number of tasks that are large, irreducible, high priority and frequently time-critical; examples include weather and seismic forecasting. Such systems typically have high bandwidth shared memory between processors. Capacity systems are typically clusters of commodity processors, which are designed to maximise throughput for a large number of distinct tasks and/or users. Capability systems can also handle capacity tasks, but their use of specialised hardware generally makes them more expensive than a capacity system.

Review Process

This document is intended to provide a blueprint for Advanced Scientific Computing in CSIRO over the period 2007 to 2010 and beyond.

The document was developed using the following process:

(i) The ASC review process was developed by the CSIRO IM&T Consulting Group and endorsed by the CSIRO IT Advisory Committee (ITAC) meeting on 26 June 2007

(ii) A call for input from CSIRO staff into the review was broadcast through Monday Mail #243 (4 June 2007), IM&T Update #2/2007 (13 June 2007) and HPC Bulletin #169 (30 May 2007)

(iii) Heads of all Business Units (BUs) were invited to nominate representatives to participate in the planning process

(iv) An outline document was circulated to the BU representatives and other interested parties on 4 July 2007

(v) A survey of BU activities, projected demand and issues was circulated to the BU representatives on 23 July 2007

(vi) An ASC workshop was held on 14-15 August 2007 at CSIRO Ian Wark Laboratory in Clayton

(vii) Input was solicited from peer overseas research organisations in the Global Research Alliance and from national partners (the "PACS") in the Australian Partnership for Advanced Computing (APAC)

(viii) Additional input was received from CSIRO executive, scientific and IM&T staff, via meetings, videoconference, telephone and email

(ix) The resultant material was collated and analysed to develop the content and recommendations in this document.

ADVANCED SCIENTIFIC COMPUTING IN CSIRO

Current Computational Science Activities

CSIRO currently has a diverse portfolio of computational science research. The following summary presents a subset of these research areas with significant Advanced Scientific Computing requirements or overlaps.

- Australian Square Kilometre Array Pathfinder “ASKAP”, pulsar evolution (ATNF)
- Molecular modelling of energy generation and storage materials (CET, ETF)
- Geomechanical modelling of ore deposition, predictive minerals discovery, modelling mine excavation impact (CEM, MDU)
- Modelling pedigree gene regulatory networks (FF)
- Modelling food processes and unit operations for dryers, extruders, ovens, high pressure vessels, refrigeration, storage and transport systems (FSA)
- e-Research platforms, National Collaboration Network, trusted systems (ICTC)
- Hydrological modelling of groundwater systems, catchments and aquifers (CLW)
- Coupled climate-earth system modelling “ACCESS”, ecosystem modelling, biogeochemical modelling, ensemble forecasting, complex systems science (CMAR, CCA, CSE)
- Engineering and thermofluids, microfluidics, materials theory & modelling, corrosion modelling, image reconstruction (CMSE, NMF)
- Fluid and particle dynamics for industrial applications and visual animations, geospatial mapping and monitoring, quantitative risk analysis, terabyte dataset processing (CMIS)
- Fluid dynamics and multiphase flows in mineral processing and power generation, instrument design for Online Analysis and Control (Minerals)
- Computer aided drug discovery, structural biology, materials informatics (CMHT)
- Stratigraphic forward modelling, geophysics, reservoir characterisation, underground CO_2 storage, coupling of subsurface flow with geomechanical and geochemical processes (CPR)
- Gene annotation, classification, prediction and expression, microarray analysis, comparative genomics, transcriptomics, genotyping, proteomics, phylogenomics, metagenomics, crop management, biosecurity (CPI, CLI, Ento)
- Modelling of neurodegenerative diseases (pHealth)
- Water Resources Observational Network “WRON” (WfHC)
- National ocean forecasting system “BlueLink” (WfO)

A more extensive list is given in Appendix 6: Utilisation of Advanced Scientific Computing in CSIRO.

Current Scientific Computing Services

Scientific computing in CSIRO is currently supported principally through the High Performance Scientific Computing (HPSC) and e-Science teams. Both are part of IM&T, but HPSC is funded as a separate theme. HPSC focuses mainly on operating shared systems and access to partner facilities, while e-Science focuses mostly on supporting local clusters, Unix-based workstation systems and cycle harvesting.

Current Shared Facilities

The majority of CSIRO's shared ASC facilities are collocated with the Bureau of Meteorology (BoM) at their head office in Docklands, Victoria. These facilities are managed by the High Performance Computing and Communications Centre (HPCCC) and include:

- five nodes (40 CPUs) of the shared NEC SX-6 parallel vector supercomputer, and associated TX7 disk server systems
- an SGI Altix 3000 NUMA server with 128 IA64 CPUs and 224GB shared memory
- an IBM BladeCenter 1350 cluster with 136 dual Xeon CPU nodes
- the CSIRO Data Store, a hierarchical file store (HFS) consisting of 31 TB staged high- and low-performance disk, and a StorageTek 9310 tape silo, with current duplicate holdings of over 700 TB of near-line storage
- a portfolio of software licenses including compilers, development tools and applications[8].

The NEC capability system is currently used predominantly by CMAR researchers. The Altix and IBM cluster systems currently have diverse and varying usage patterns from 12 divisions, with CMIS, CEM, CMHT, CET, CMAR, CLW and CLI being some of the more prominent users. Utilisation of these machines is depicted in more detail in Appendix 6: Utilisation of Advanced Scientific Computing in CSIRO.

The other main shared facility is the CSIRO Bioinformatics Facility (CBF). The CBF consists of a Dell cluster with 66 dual CPU compute nodes and five server nodes, where the selection of architecture and software is designed for performing genomics and proteomics applications. The CBF was initially funded through the Emerging Science Initiative in Biotechnology, and has users from nine different CSIRO divisions, with the biggest users being CPI, CLI and CMIS.

The majority of these shared systems have been in service for four or more years and are due for replacement. The Data Store currently has a comprehensive backup policy, but no offsite backup in the event of a fire or other destructive event at Docklands. CSIRO currently does not have comparable storage capacity elsewhere to provide automated offsite on-line or near-line backup.

CSIRO IM&T e-Science runs a pilot cycle harvesting facility, consisting of the Condor scheduler accessing a pool of 700 desktop PCs. The system currently has significant utilisation from a small set of users, largely within Minerals and CEM. The system can be expanded to utilise CSIRO's 6000+ available desktops and cater for a broader range of suitable applications and users.

Current Local Facilities

In addition to using high-specification PCs as standard, CSIRO currently has over 1300 fast scientific workstations. Typically, these desktop or deskside machines are dedicated to single users, and have single or dual high-specification commodity CPUs under a single system image. About 60% or these machines run under Windows or Windows Server, with the remainder utilising Linux, Solaris or other Unix operating systems. In 2007 IM&T established a standard supported Linux environment (SLES

[8] http://intranet.csiro.au/intranet/it/slas/catalogue/concepts/SWHPSC.asp.

9/10), with variant operating systems (e.g. Debian, Red Hat) available for self-support where required. Additionally, it is estimated that about 30% of CSIRO's Dell PowerEdge servers are used for research project work, generally under Windows Server, and in some cases serving as local cluster nodes.

These scientific workstation systems are only moderately more expensive to purchase or maintain than standard desktop PCs. Hence, they provide a cost-effective platform for moderate scale scientific computing, especially work that is interactive and/or non-scalable in nature. Growing and streamlining this as a mature commodity service is important not only to support established science activities, but also as an entry path for raising the level of computational science expertise and deployment across CSIRO.

Several CSIRO sites have significant cluster facilities, which are typically dedicated to selected projects or groups. The cluster systems in Clayton (Minerals, CPR), Parkville (CMHT) and QCAT (CEM) are typical examples. In addition, HPSC operates portions of its IBM cluster that are dedicated to specific CMIS and CMAR groups.

Current CSIRO Use of Partner Facilities

CSIRO is a partner in the Australian Partnership for Advanced Computing (APAC). Note that APAC is currently being replaced by a new structure within PfC, as outlined earlier in Scientific Computing Interrelationships, with NCI responsible for providing major national computing infrastructure.

CSIRO is also a participant in the iVEC and TPAC regional partnerships. The relationship with iVEC is facilitated by being collocated with CSIRO at the Australian Resources Research Centre (ARRC) in Perth, with CSIRO users being allocated a significant share of the facility's compute and storage systems. CMAR make use of TPAC compute and storage facilities, which are based at the University of Tasmania in Hobart.

The APAC National Facility, managed by the ANU Supercomputer Facility, operates a 1920-processor SGI Altix, which is the largest computing system in Australia. CSIRO has made significant use of the National Facility for a number of years. In the past, this has been predominantly through a partner share at a cost of $200k per annum. As of 2007, this share was reduced to $40k per annum with recognition that CSIRO investigators were eligible to apply for Merit Allocation Scheme (MAS) grants. For 2007, this MAS share has a value of $217k. The user base has been small but diverse, and currently has about 12 users from eight different Business Units. A number of APAC grants have also been awarded to CSIRO staff with university affiliations, or to university chief investigators with CSIRO partners.

Several CSIRO groups also engage in research partnerships, which use external computing resources in Australia and overseas. For example, the large new cluster system at the Swinburne Centre for Astrophysics and Supercomputing[9] will be heavily utilised in collaborative analysis of Australia Telescope data. Commercial HPC facilities are also available[10], although these do not offer the same collaboration and leverage benefits that are available through the PAC partners.

[9] http://astronomy.swinburne.edu.au/supercomputing/.
[10] e.g. Amazon EC2 (http://www.amazon.com/gp/browse.html?node=201590011).

REVIEW FINDINGS

Comparison with Peer Organisations

The four international research agencies contacted by the CEO to provide external "benchmarks" were Battelle (USA), Fraunhofer-Gesellschaft (FhG, Germany), TNO (Netherlands) and VTT (Finland). All are participants in the Global Research Alliance (GRA). Inquiries to these agencies were in turn referred to four specialist HPC providers, namely NCCS, ITWM, NCF and CSC respectively. Note that the first two are subsidiaries of their respective agencies, while the latter two are partners. NCCS is notable for establishing a "Leadership Computing" facility at Oak Ridge National Laboratories (ORNL). This includes an 11700-processor Cray XT3/XT4 system, which is currently ranked the second fastest computer in the world[11]. More detailed information on these organisations and their interrelationships is given in Appendix 9: External Organisational Benchmarks.

Table I. Investment in Advanced Scientific Computing by CSIRO and peer organisations[12].

Research Institute	**Fraunhofer Gesellschaft**	**VTT**	**Battelle Institute**	**TNO**	**CSIRO**
Country	Germany	Finland	US	Netherlands	Australia
Institute Budget A$M	1980	359	4400	943	1020
Staff	12500	2800	4200	4600	6600
ASC Supplier	**ITWM**	**CSC**	**NCCS**	**NCF**	**HPSC**
ASC Staff	150	150	~180	70	10
Budget A$M	17.4	25.6	385	20	4.4
ASC Facility Fund Source	Fraunhofer	Ministry of Education	Battelle/ DoE	NWO	CSIRO
Fund Source Science Budget A$M	1980	430	4400	730	970
% Budget ASC	0.9%	6.0%	8.8%	2.7%	0.4%
Peak capability TFlop/s	10.2	11.3	119	14.6	1.1
Peak aggregated TFlop/s	11.3	22.0	137	23.0	3.0
Max capability TFlop/s	7.9	8.9	102	11.5	0.6
Max aggregated TFlop/s	8.5	17.1	115	~15	2.0

[11] http://www.top500.org.
[12] "Peak" and "Max" refer respectively to theoretical and measured (Linpack benchmark) maximum performance in teraflop/s. "Capability" denotes the largest single system, while "aggregate" gives the sum performance of large systems operated by the facility.

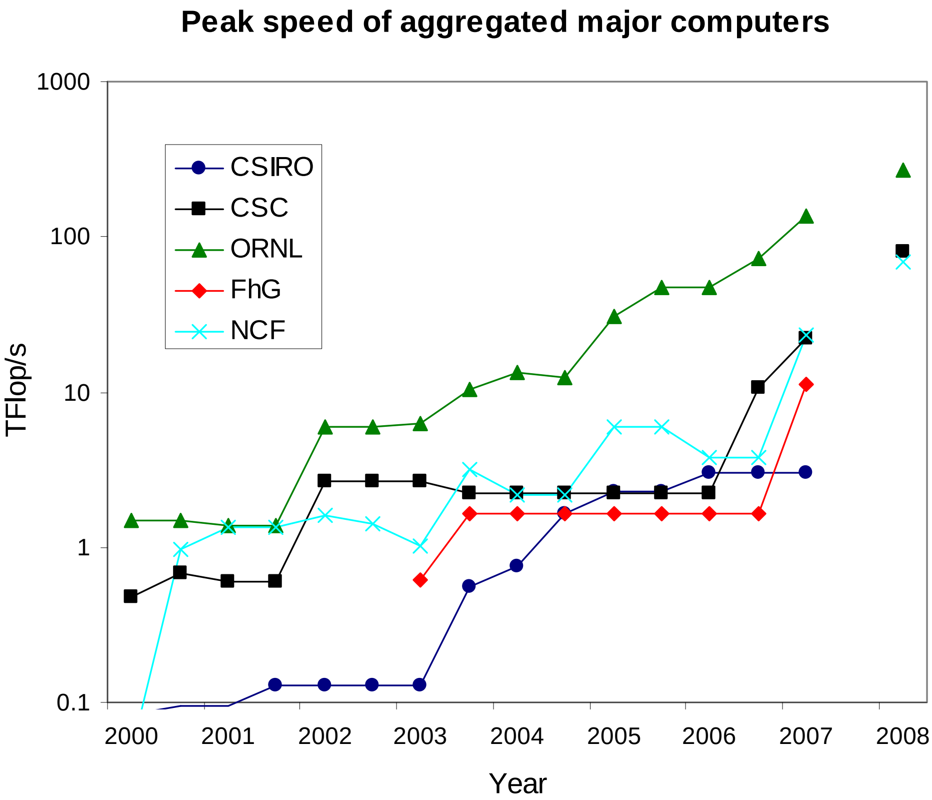

Figure 3. Comparison of hardware performance between CSIRO and peer organisations.

The principal findings of the comparison are summarised in Table I for the 2006-7 financial year, while the variation in facility performance over time is compared on a logarithmic scale in Figure 3. A number of pertinent conclusions can be drawn from these findings, as follows:

1. CSIRO has inferior computing facilities compared to these peer organisations. Furthermore, the gap is currently widening, because all four organisations have launched significant upgrades to their facilities within the past year

2. CSIRO invests a substantially smaller fraction of its overall science budget on ASC than any of these peers

3. During this decade, CSIRO and its external peers have increased their diversity of service beyond capability HPC systems, to include such elements as large capacity clusters, large data stores and grid services. A significant difference is that CSIRO has done so by subdivision of a relatively constant budget, while the other organisations have done so through corresponding additions to their budget

4. CSIRO spends a significantly smaller fraction of its ASC budget on support staff than its peers

5. The peer organisations with in-house ASC capability have these embedded in a wider computational science or e-Research research structure, which has been missing in CSIRO.

Figure 4 displays how the 2006-2007 budgets of different Advanced Scientific Computing facilities were apportioned across different science disciplines, according to the following classification:

- Information and mathematical sciences
- Physics, including nuclear, particle and astrophysics
- Chemical sciences
- Biological sciences, including molecular biology and ecology
- Materials, condensed matter and nanosciences
- Geosciences, earth systems and environmental sciences
- Engineering, including CAD, fluids and process modelling

It is worth noting that these are commonly used classifications, which do not discriminate between earth sciences and marine & atmospheric sciences.

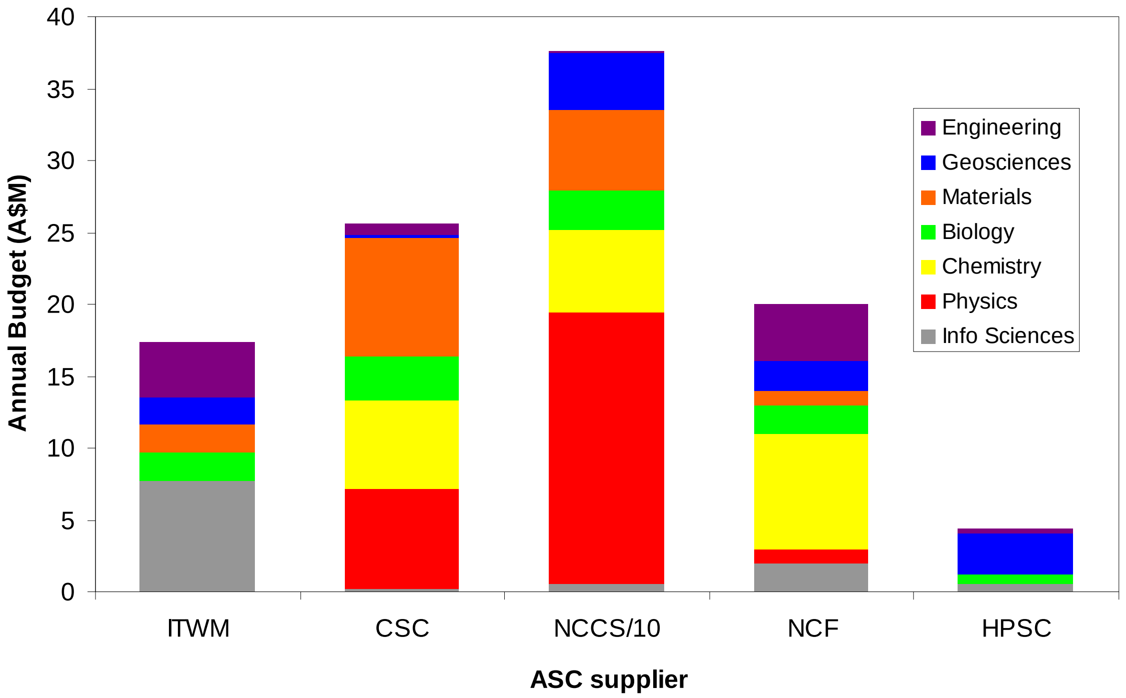

Figure 4. Apportioning of computing budget according to science discipline[13].

The other organisations display significant variations in the relative proportions, but in each case the balance of disciplines is roughly commensurate with the organisation's overall research profile. For CSIRO it is not. Computational research activity in the physics, chemistry and materials domains in CSIRO is too small to be measurable on the scale of this figure. Clearly the concern with CSIRO is not that the geoscience activities are too large – particularly bearing in mind their national significance – but rather that most of the other activities are too small.

[13] NCCS figure has been scaled down by a factor of 10. CSIRO figure includes HPSC and CBF.

This comparison shows that *CSIRO would need to at least double its ASC investment in order to be comparable with the nearest peer organisation* (Fraunhofer), where "nearest" means both "most analogous" in structure and purpose, and "next lowest" in proportional investment. A minimum factor of 2 discrepancy is consistent with earlier findings[14] and also reflects the need to double the uptake across CSIRO's range of science application areas, as recommended in the Broad Directions Setting process[15].

Comparison with National Partners

As well as international peers, information regarding services and service delivery was solicited from the APAC regional partners. The findings are summarised in Table II. The responses generally reflect the nature of the regional partners as independent and specialist providers of ASC services, rather than as an integrated component of a broader IT function.

Table II. Service structure and delivery for APAC partners.

PAC	VPAC	SAPAC	iVEC	TPAC	QCIF	HPSC
Staff:Infrastruct. $	1 : 1	1 : 1	1.8 : 1	0.6 : 1	0.5 : 1	0.4 : 1
Embedded staff	Yes	Partial	Yes	Yes	Yes	Partial
IT Framework	ISO 9000	No	No	No	ITIL aspects	No
Tracking	TR	Request Tracker	Email lists	Helpdesk	ITSM	Wreq

Diversity of Requirements

The following is a summary of responses by Business Units and individual scientists to the planning process, which are described in more detail in Appendix 7: Survey Results.

If there is a single message to draw from the responses, it is that CSIRO researchers have a wide range of different requirements, as befitting the wide range of disciplines, skills, partners and drivers associated with their research. Where commodity products typically present a range of options with comparable performance, for ASC the choice of software and hardware platforms in particular are influenced by subtle factors that can and do have order of magnitude influences on project productivity. This is illustrated for example by the very different divisional utilisation patterns for HPSC systems with different architectures. These performance issues frequently override smaller scale economies, such as those associated with preferred suppliers, common environments etc. They are a clear reflection of the near-unanimous emphasis of respondents that a "one size fits all" philosophy is inappropriate for ASC.

Among the requirements that have been flagged as significant are:

14 http://intra.hpsc.csiro.au/user/general_information/HPSC_Comparative_Review.pdf.
15 I. Elsum, Computational Science: SIP2 Broad Directions Setting, August 2006.

- Increased demand for Linux cluster resources in particular
- Desire to grow Windows cluster facilities
- Increased demand for other compute systems, e.g. shared memory and cycle harvesting
- Significant increases in data storage, access and retention requirements
- Enhanced data management for re-use, knowledge extraction, development of associated processing techniques, curation, regulatory requirements etc
- Enhanced code and software management.

The need for more Windows cluster systems is a relatively recent requirement, noting that Windows is currently better suited as an environment for multiple independent calculations (e.g. parameter sweep, ensemble) rather than coupled parallel codes. Hence, some of this demand can be borne by small local blade servers, while other users have larger scale but intermittent requirements that would be better served by a shared facility.

A few respondents highlighted the potential performance and cost benefits of custom hardware platforms such as field-programmable gate array (FPGA) or cell processor systems for their specific applications, with the ASKAP project being a notable example.

Adoption of Local and Centralised Facilities

Objective 3.1.2 of the CSIRO Strategic Plan, for "Working Effectively and Efficiently in our Enterprise", is to "Utilise common systems, structures and improved processes… to optimise the use of our facilities, equipment and information assets."

IM&T's consolidation programs for servers, printers etc are driven by the following benefits:

- improved risk management, notably through easier backup and greater security
- better capacity planning
- economies of scale and utilisation obtained by aggregating smaller systems.

For the IM&T Server Consolidation Project, the economy is estimated[16] at 25%, which is a relatively small benefit compared to, for example, the reduced risks.

For ASC the economies of utilisation predominate over other economies. For example, analysis of the usage of the HPSC cluster and Altix systems indicate that this economy is a factor of 8-9, i.e. *individual users would require an order of magnitude more resource for the same level of responsiveness*. The more these individual users are aggregated, the greater the economy becomes. Details of this analysis are given in Appendix 3: Technology Trends. The other principal advantage of aggregation for scientific computing systems is that it increases the *capability* of the system, i.e. it increases the number of connected CPUs that are available to individual tasks and users, enabling larger tasks to be tackled.

Costs of operating and maintaining clusters are typically approximately proportional to the number of nodes – and hence CPUs – in the cluster. These costs are considered in more detail in Appendix 4: Cost of Ownership. Note that *the purchase price is usually a minority component of the total cost of ownership to the organisation*. Similar ratios are also commonly observed for server systems[16], while

[16] R. Knudson, CSIRO IM&T Server Consolidation Project presentation, September 2007. http://intranet.csiro.au/intranet/it/about/PO/FP/serverConsolidation.htm#Details.

for desktop PCs the high levels of standardisation mean that the purchase price is a marginally higher fraction (approximately 50-60%) of the total cost of ownership[17]. Another significant factor is that vendors' increasing reliance on multicore architectures for performance improvements, and increasing costs of electricity due to environmental pressures, which are anticipated to lead to dramatic increases in energy costs as a proportion of total cost of ownership over the next 3 years[18].

Maintaining effective utilisation of small, dedicated clusters is more difficult than for larger shared systems. Low levels of utilisation give poor value, while conversely prolonged utilisation levels close to 100% imply that the resource is not large enough.

Hence, use of local clusters and other HPC systems is most appropriate in cases where the economies of utilisation are either achieved via local critical mass or are limited by other factors. Typically, these might include:

- Commercial software requirements/costs that exceed or govern the hardware requirements
- High data throughput from non-computational source, e.g. collocated scientific instrument
- Requirement for specialised or novel hardware platforms e.g. FPGA
- Requirement for specialised software development environments
- Requirement for maximum predictability of throughput, rather than maximum throughput, e.g. to meet client deadlines
- External requirement or perception of dedicated access, security or confidentiality
- Low bandwidth network access, e.g. at remote sites
- High visualisation or interactivity requirements that can lead to poor scalability (see ASC definition).

Note that less than a decade ago, clusters themselves constituted novel technology, such that the costs of implementing, administering and running the system could be considered as part of the researcher's training.

Unlike compute servers, data storage facilities have no real economies of utilisation but significant economies of scale, which depend on the size and format of the storage medium.

Projected CSIRO Demand

During the present review, *all* responding Business Units anticipated a growth in demand for Advanced Scientific Computing facilities. Estimating this growth in demand for compute and storage resources is easiest at the highest level of aggregation, where historical trends can be used to smooth out the vagaries of individual projects and research groups.

With CSIRO now looking to increase focus on computational science, a number of new large-scale projects will increase the historical growth rate of demand for computational resources. These projects include the Minerals Down Under, Climate Change Adaptation and Niche Manufacturing flagships, the Australian Square Kilometre Array Pathfinder (ASKAP), the Centre for Australian Weather and Climate Research (CAWCR), the Water Resources Observation Network (WRON), new themes in Terabyte Science and Predictive Minerals Discovery, a proposed computational &

[17] M. Sheppard, CSIRO IM&T Technical Services Team – Desktops, September 2007.
[18] http://www.uptimeinstitute.org.

simulation science capability theme (and its attendant large scale applications), growing activity in energy generation and storage materials, and significantly increased requirements in livestock and plant genomics and systems biology.

These projects alone will require an approximate doubling of the present level of resourcing for cluster and shared memory systems within 2-3 years, not including ASKAP, which will require a large (20 Teraflop/s) dedicated HPC facility. Hence, the combined requirement for 2010 is estimated as at least 40 Teraflop/s.

Meeting this growth in demand will require a substantial increase in computing capacity that will be difficult for CSIRO to attain on its own. The leveraged funding, science support and collaborative environments available through NCI and other partnerships can be of significant benefit to many of these project areas.

At time of writing, CSIRO stores approximately 1 Petabyte of science data. The growth rate is approximately a factor of 2000 over the past 13 years, corresponding to a storage requirement that has been virtually doubling every 12 months. With data-intensive research projects now becoming prevalent in CSIRO, this historical growth rate may be exceeded over the next few years, at least until the "data avalanche" is curtailed by the growing proportion of research costs associated with retaining and managing this data. A conservative estimate for CSIRO's storage requirement for 2010 would be 10 Petabytes (not including ASKAP).

Constraints on Uptake

Currently, access to shared CSIRO ASC facilities is readily available. CSIRO scientists can gain effectively unrestricted access to the IM&T HPSC systems overnight simply by obtaining a signature from their next level manager. Access to APAC partner share resources has been similarly straightforward.

The fact that these resources are smaller than those of comparable peers, but have generally not been oversubscribed in the past, reflects several issues:

- The proportion of CSIRO researchers who are able to implement truly scalable applications or workflows appears to be relatively low
- It appears that some staff who do possess the necessary skills are simply unaware of the ease of access to these facilities
- Computational science is frequently more technically challenging than experimental science, and many BUs have restricted computational research activity, due to their inability to locate, allocate or develop researchers with suitable skills and expertise
- Parts of CSIRO do not have a strong scientific computing culture, i.e. they do not always identify existing computing infrastructure as part of their portfolio of instruments and services for tackling client problems
- Compared to the academic sector, CSIRO research costs are more typically limited by staffing costs rather than equipment costs
- Some CSIRO researchers use alternative facilities because of previous technical or service issues with HPSC, other parts of IM&T or non-CSIRO equivalents

- These facilities are largely limited to Unix-based operating systems, which are the standard platform for large-scale computations, rather than Windows, which is frequently more suitable for desktop scientific computing applications.

The IM&T e-Science team has had intermittent outreach activities to address these issues via its science strategy and road show roles. The intermittency is necessary because these activities generally result in the creation of new projects and user requests, which stretch the capacity of the hardware and support staff currently available.

Opportunities for CSIRO Science

Advanced Scientific Computing will obviously form an essential component in supporting the growth of CSIRO’s capability in Computational and Simulation Science. It also has a clear role in supporting the other current priority areas for science investment, namely Advanced Materials, Transformational Biology, and Sensors and Sensor Networks.

Increased computational capability can play a key role in achieving many of CSIRO’s most significant goals. These include:

- Predicting the ecological and socioeconomic impact of climate change
- Discovery and extraction of new minerals and fossil fuel resources
- Development of novel materials for energy generation and storage, sequestration, desalination, biomedical, sensor, transport and infrastructure applications
- Developments of new processes and treatments for metals and mineral feedstocks
- Diagnosis and treatment of neurological diseases
- Rational design of bioactive compounds and drug delivery systems
- Improved breeds and management of crops and livestock
- Successful bid to host the Square Kilometre Array telescope.

An NCCS perspective of opportunities for computational science[19] is given in Figure 5. What can be seen is that in each of the disciplines, progress requires large-scale computations using coupled mechanistic simulations across multiple scales. Depending on the field, such approaches have been given labels such as “atoms to aircraft”, “quantum to planet” and “molecules to ecosystems”. Clearly, they define commonalities across these fields, which offer profound opportunities for synergistic benefits in a multidisciplinary research enterprise.

[19] Verastegui, B. 2005. Computing and Computational Sciences Directorate. Workshop Presentation to ORNL Research Alliance in Math and Science.

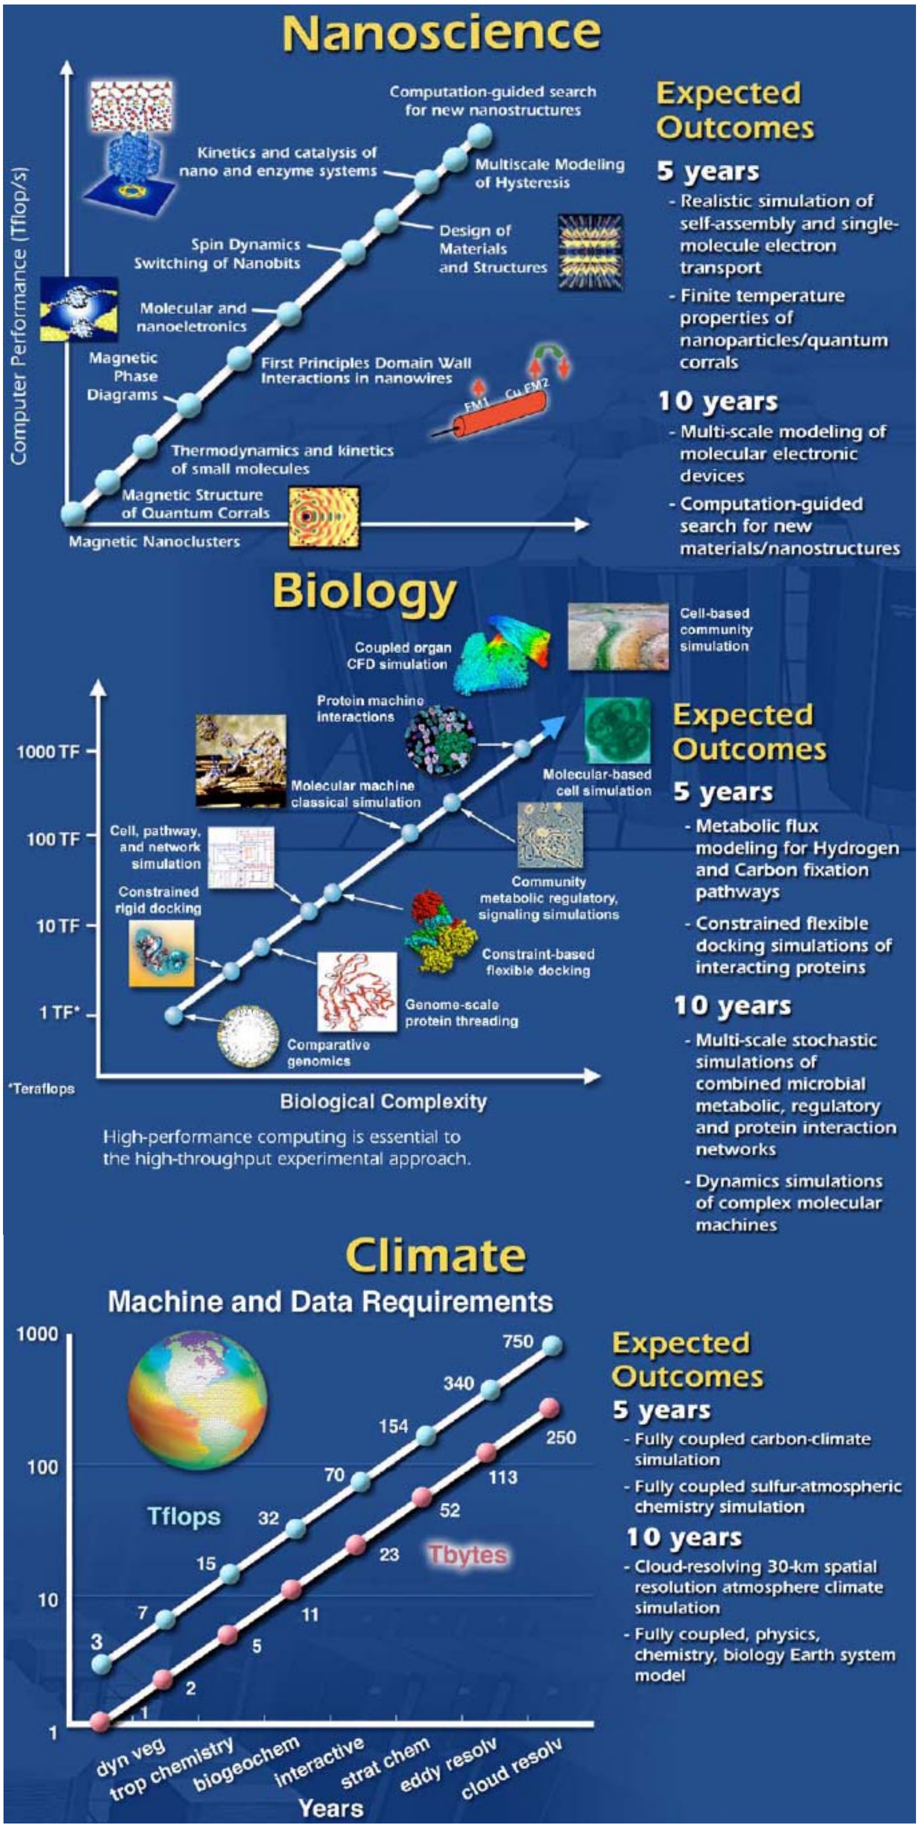

Figure 5. Projected opportunities for large scale scientific computing (NCCS).

RECOMMENDATIONS

Options

The three principal Options available for CSIRO to consider for its Advanced Scientific Computing strategy are outlined in Table III. Based on the relative merits of the different options, Option 2 for Incremental Growth is the recommendation identified through the review process as being optimal for CSIRO. This Option forms the basis for a number of the subsequent recommendations.

Table III. CSIRO investment options for Advanced Scientific Computing.

1. Status Quo

Fragmented approach with mixture of HPSC and BU resourcing

CSIRO approach to ASC stays on current path.

PROS	CONS
• No increase in investment nor staff needed • HPCCC can continue to operate • Modest funding for non-HPCCC advanced computing requirements.	• Reactive approach • Unable to meet current & growing demand • Unrealised science opportunities • Limited ability to engage and collaborate with wider national and international community • CSIRO competitiveness erodes • CSIRO impact diminishes.

.. table continued next page

2. Incremental Growth

Enterprise level investment in ASC is approximately doubled over three years, to match lower end of peer scale

Investment driven by a "Just In Time" approach

Increased focus on people, processes and organisational alignment to computational science outcomes in line with CSIRO Strategic Plan

Emphasis on external partnerships and shared facilities to increase compute power

Maximise value from NCRIS/DEST developed capabilities, i.e. ARCS, ANDS, NCI, AAF, AREN

CSIRO creates an ASC governance structure to manage this investment.

PROS

- Current and backlog demand can be satisfied through partners and shared facilities.
- Implementation is phased in-line with related projects such as eSIM and ARCS.
- Low investment thresholds, no significant new capital quanta required on day one.
- New capital is sought based on detailed demand forecasting and business case development.
- CSIRO ASC initial focus is on people, processes and governance, to be in better position to use capital investment in assets during subsequent years.
- NCRIS compatible with new funding potential.
- CSIRO takes a national leadership role in collaborative computational science.
- People can be recruited and processes created to set the correct directions.

CONS

- Immediate and backlog demands may be under-serviced, leading to suboptimal science outcomes.
- Non-geoscience applications slow to develop.
- Supply lag if shared facilities unable to supply - CSIRO does not have full control of the supply side of the compute power equation.
- Scarcity of suitable recruits for new staff positions.

3. Rapid Growth

Significant immediate increase in enterprise investment for ASC

Major focus is on people, processes and organisational alignment to computational science outcomes in line with CSIRO Strategic Plan, to be achieved as rapidly as possible

CSIRO creates an ASC governance structure to manage this investment.

PROS

- Immediate availability of computational power for any/all researchers
- CSIRO is competitive with peers
- People can be recruited and processes created to set the correct directions.

CONS

- Expensive option with impact on other CSIRO funding opportunities
- Potential for wasted or inappropriate capital funding if facilities are poorly utilised
- Unrealistic in the time-frame - Not ready with services maturity and outreach
- Scarcity of suitable recruits for new staff positions.

Proposed Capability

The focus of the proposed strategy is to create an integrated CSIRO capability for Advanced Scientific Computing. This capability requires commodity and value-added service components, and the skills and flexibility to function in an internal and external partnership role as required. Many of these components have unique, differentiated and leading edge elements.

As seen in Figure 6 the capability contains four integrated layers, each of which has capability development, people, process and governance components:,

- Layer 1 is the hardware capacity and is delivered as a managed service either internally or through partner services
- Layer 2 is the IT fabric for conducting e-Research and involves newly evolving technology
- Layer 3 is the science applications support and development of a community of practice, internally and through partnerships
- Layer 4 is the actual computational science.

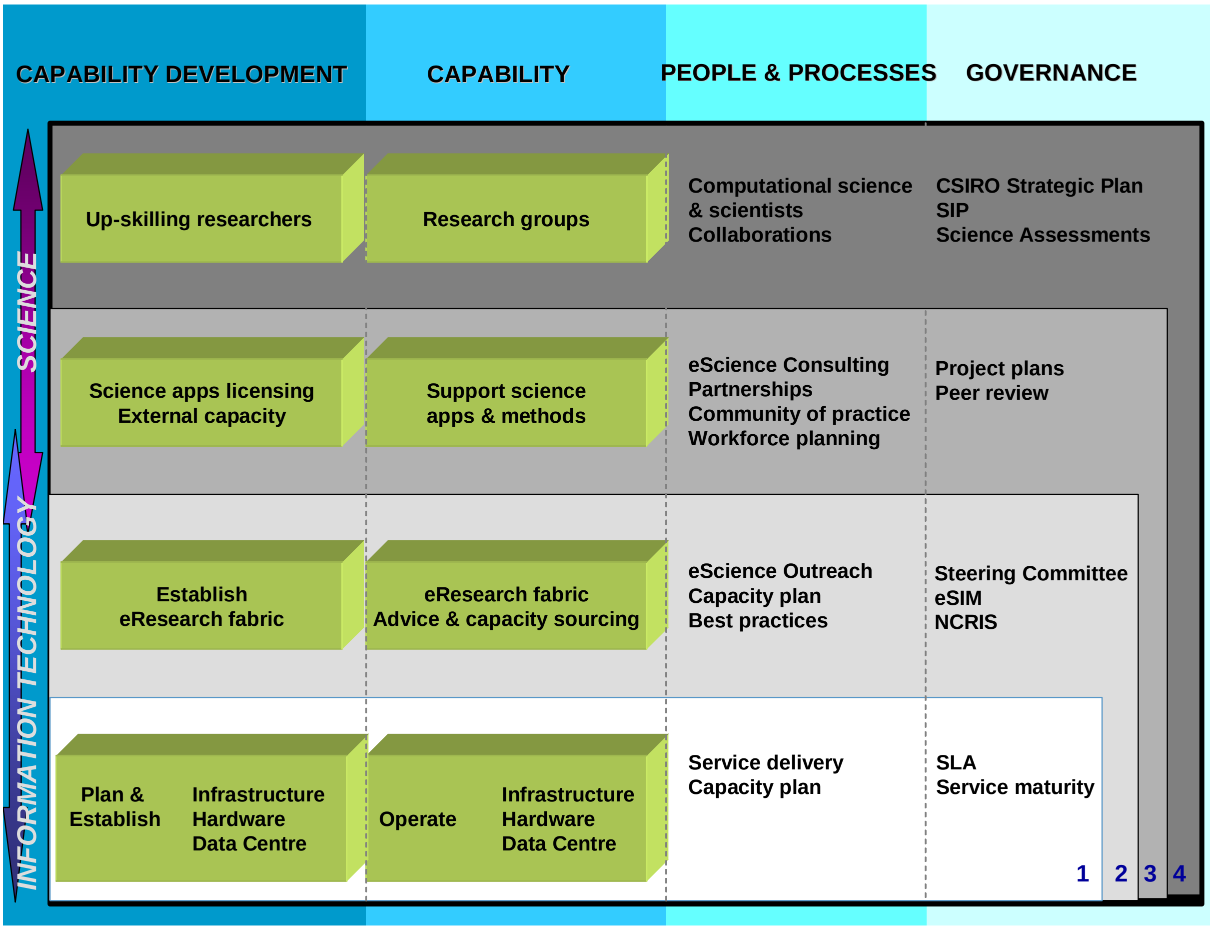

Figure 6. Proposed Structure for CSIRO's ASC capability.

Specific recommendations regarding the Structure and Governance, Provision of Services, Community of Practice and Financial Basis of this capability are listed below.

Structure and Governance

Recommendation 1.

Advanced Scientific Computing needs to be recognised and promoted as a CSIRO enterprise capability, which enables computational science to be a large, growing and hugely valuable portion of our research portfolio.

Recommendation 2.

The base layer of this capability in CSIRO can reside primarily within an expanded IM&T team based on the present HPSC and e-Science groups. The interim label of "ASC Team" is used here to identify this function.

Recommendation 3.

A Steering Committee should be created to oversee operation of the ASC Team and assess the optimum investment balance to match ongoing user requirements. Ideally the committee would consist of the Director of Information Services & Property, the Executive Managers for e-Science and Service Delivery of IM&T, together with the Group Executives or appropriate delegates from each CSIRO Group, plus at least one external representative. This function might be served by a larger scale committee that also encompasses e.g. the eSIM and PfC projects.

Recommendation 4.

All decisions by the ASC Team associated with scientific computing should be made on the basis of maximising the net science impact, per unit cost to CSIRO at enterprise level. For major decisions, this will draw on the expertise of the Steering Committee.

Provision of Services

Advanced Scientific Computing requires a comprehensive and integrated set of services, focusing on compute servers, storage, software, development tools, systems management and science support, incorporated into an e-Research and information management framework.

Recommendation 5.

CSIRO needs to maximise expertise and leverage in Advanced Scientific Computing services by continually identifying and pursuing suitable partnership opportunities, at the local, regional and national level. The aim is to combine the collaborative and economic benefits of shared partner facilities with the provision of a diversity of systems and services that are technologically and geographically diverse.

Recommendation 6.

CSIRO should renew the HPCCC partnership agreement with the Bureau of Meteorology as the basis for provision of a replacement supercomputing system. CSIRO will contribute 25% to the facility, or approximately $2.1M p.a. over 4 years. While CSIRO's proportional share remains constant, the overall investment is a significant increase over previous years. This reflects both the growing importance of climate modelling to the national interest – as focused through the formation of CAWCR – and the evolution of HPC architectures and CSIRO user requirements, notably in simulation science, such that non-CMAR use of this facility will be significantly greater than for the current NEC system. The linked Bureau and CSIRO requirements will lead to a system that is likely to be significantly differentiated from other national facilities. The facility will be available to all CSIRO researchers, focusing on those for whom this differentiation is most valuable.

Recommendation 7.

CSIRO should invest in the National Computing Infrastructure initiative to meet the bulk of projected growth in magnitude and diversity of user demand. The initial level of investment is proposed to be approximately $1.0M p.a. in the national peak facility, plus $1.0M p.a. in regional or "shoulder" partner facilities, to provide hardware, software, storage and support services.

Recommendation 8.

CSIRO requires ongoing in-house provision of at least one shared Linux-based cluster system of significant size, to provide maximum flexibility of access, configuration, prioritisation and security. The architecture of such a system should evolve towards maximum ease of use to new users and ease of migration from desktop environments.

Recommendation 9.

CSIRO needs to implement a production Windows shared cluster facility. The facility can be provided using partitioning and/or virtualisation of a larger system as appropriate.

Recommendation 10.

CSIRO should continue to refine processes for procuring and managing local cluster systems where required by specific research groups or sites. Diverse science requirements for such systems necessitate access to a broad range of system architectures and specifications, influencing choice of vendors.

Recommendation 11.

CSIRO should seek staged migration of existing cluster applications and establishment of new users onto shared systems, for cases where these systems can meet the research needs. These needs – resource, access, risk etc – would be discussed by IM&T and the researchers at appropriate times, notably initiation of new users/projects and renewal of legacy systems.

Recommendation 12.

CSIRO IM&T should continue to grow and mature support for fast single user workstations to meet needs for small scale scientific computing requirements as a commodity service, using standardised hardware and operating system configurations where appropriate.

Recommendation 13.

CSIRO e-Science should expand implementation of the Condor cycle harvesting system and incorporate it into the suite of facilities available to researchers across CSIRO. The facility should target one or more key science applications to be deployed on PCs across the whole of CSIRO. This would create true One-CSIRO research activity, where all staff have a notional level of involvement.

Recommendation 14.

CSIRO must expand data storage capacity, through a mixture of managed disk storage for midsize datasets requiring full on-line access, and HSM for large science data. A second HSM facility should be established as soon as practicable, to provide a mirror for offsite near-line backup to the current HPSC facility. This system should be collocated with a significant compute server at a CSIRO Data Centre or NCRIS partner site.

Recommendation 15.

CSIRO ASC needs an increased total and proportional investment in specialist staff, focusing on support for new users, scientific software implementation and science application areas. These additional staffing needs can be met by a combination of partner services and in-house staff. The former can add significant value especially through applications specialists, while the latter are particularly important for supporting the development of novel techniques, algorithms and software IP. Such staff should be readily accessible across all CSIRO sites, disciplines and facilities as needed. , and provide users with comprehensive information on their options for meeting their computational needs, so they can make the most suitable choice for where to run their computations.

Developing a Community of Practice

CSIRO currently does not have an established computational science community. These research groups are largely focused on their own areas, and often see each other as competitors. CSIRO IM&T can play a valuable role in strengthening links with, and between, such groups, by

- providing effective collaborative IT platforms and tools
- implementing policies that minimise sources of internal competition and conflict
- participating in and facilitating network initiatives such as Computational & Simulation Science
- helping to identify overlaps in skills, expertise and interests across CSIRO and with partners.

Recommendation 16.

A significant proportion of the enlarged CSIRO ASC staff should be assigned to an enlarged outreach program. The outreach group would be more proactive in raising general awareness regarding ASC opportunities and challenges, and provide road shows, seminars, courses and training material, which are suitable for a broad range of potential users. Where appropriate these staff can be temporarily embedded within sites to maximise contact with users.

Recommendation 17.

CSIRO IM&T should seek opportunities to play a more proactive partnership and/or facilitator role where needed by specific research themes with advanced computing requirements and targets.

Recommendation 18.

CSIRO should maintain a register of scientific computing resources, including a skills database and a library of scientific software and tools that have been developed in-house in CSIRO. Such a register will form a significant component in CSIRO's e-Science Information Management assets.

Recommendation 19.

CSIRO IM&T should facilitate an Advanced Scientific Computing culture across CSIRO. A key aim is to highlight the science value of scalability as a concept, by encouraging increased adoption of ASC tools, workflows, parallel programming and related technologies that maximise the throughput and hence the value of CSIRO's investment in scientific staff and computing infrastructure.

Securing a Financial Basis for Growth

The basic principle governing the provision of funds to Advanced Scientific Computing resources is to maximise the cost-effectiveness to CSIRO as a whole, by providing Business Units and their researchers with incentives to choose the most cost-effective options.

This can be done by defining two primary funding mechanisms, namely (i) enterprise level funding, and (ii) funding attributed to research themes, e.g. obtained or linked through the Science Investment Process or from independent project sources.

An enterprise funding model is appropriate for managed *shared* resources, i.e. *those that are accessible to any CSIRO researcher*. Such an approach is consistent with the One-CSIRO Foundations component of the CSIRO Strategic Plan. Note that ASC is unique as a science tool or instrument, in that it can be of substantial value to any scientist in CSIRO, regardless of discipline, application, location or seniority. As such, the ASC service can be a legitimate component of the consolidated overhead spread across all staff. It also maximises incentive and minimises entry barriers to those parts of CSIRO that have not yet been able to embrace fully the transformational capabilities of computational science.

Theme-based funding is appropriate for resources that are *dedicated* to that particular theme. Typically, dedicated resources are of most value in applications that have associated fund streams (e.g. external clients, instruments). Annual expenditure through this mechanism will be governed by demand for such resources.

Table IV. Proposed fund sources for scientific computing services.

System Type	Hardware	Software	Storage	Operational	Systems Admin	Applications Support
Dedicated standalone local						
Dedicated custom co-located						
Dedicated standard co-located	THEMES					
Shared standard co-located	IM&T					
Shared cycle harvesting	NO COST, (Already Funded)					

Table IV shows how an optimum split can be defined by classifying different resource allocations into separate service components.

Recommendation 20.

CSIRO's Advanced Scientific Computing capability should be grown over the period 2007-2011. The proposed optimum funding mix for services is depicted by the dividing line in

Table IV.

Recommendation 21.

CSIRO needs to refine and clarify the procedure for research groups to specify dedicated large scale computing requirements for funding through the Science Investment Process. A key issue is the blurred distinction between capital and operating expenditure for both hardware and software, noting that the short optimum operating life of high-performance hardware (and of many project requirements) increases the attractiveness of leasing and partnering options.

Recommendation 22.

The ASC Team needs to implement annual rolling capacity and team workforce plans, taking into account current usage patterns, new science initiatives, partnership arrangements, external trends and other relevant information. The capacity and workforce planning need to be closely linked due to the strong coupling of requirements. Users of enterprise-funded resources will be requested annually to provide brief summaries of their achievements and future requirements as part of their access conditions, with mutual benefits in capacity planning, investment justification and publicity.

Recommendation 23.

The ASC Team needs to maintain an ongoing detailed model for evaluating costs of Advanced Scientific Computing services. This costing model is to be used to apportion costs for these facilities and provide cost estimates where required for contract and co-investment research with external partners.

Recommendation 24.

CSIRO's allocated enterprise funding for Advanced Scientific Computing needs to grow to approximately double its present value over the next 3 years. Such staged growth is required to:

- Meet existing and future resource demands
- Enable CSIRO's capability to catch up with external peers after a decade of static investment
- Increase staffing to enhance planning, applications support and outreach.

The additional expenditure would be offset by increased efficiency and reduced demand for other infrastructure items (e.g. laboratory scientific instruments), reflecting the lower cost of computational relative to most experimental science.

Recommendation 25.

CSIRO should consider the option of significantly greater long-term expansion, to go beyond being merely competitive to being a world leader in Advanced Scientific Computing. It is inevitable that computational science will continue to grow in importance over the next 10 years at least, such that towards the end of the next decade it will constitute the majority of CSIRO's research portfolio. Hence in any assessment of priority investment areas to enable sustained growth of CSIRO's future scope and impact, ASC applications and skills constitute an attractive and low-risk option.

APPENDICES

Appendix 1: Glossary

Advanced Cluster	Specialised computing system, typically with higher memory bandwidth and faster networking than a commodity cluster. *See also Cluster, SMP.*
Advanced Computing	Use of computing and information technology systems that exceed the capabilities of standard desktop PCs.
Advanced Scientific Computing (ASC)	Use of computing and information technology systems to perform innovative components of scientific and engineering research, where these systems exceed the capabilities of standard desktop PCs.
Applications Support	Provision of advanced services relating to user applications and associated software. Includes software installation, maintenance, upgrades, porting, optimisation and debugging.
Australian Community Climate and Earth System Simulator (ACCESS)	Coupled terrestrial-ocean-atmosphere simulation suite for climate prediction, being developed by CMAR and the Bureau of Meteorology. http://www.bom.gov.au/bmrc/basic/wksp17/papers/Puri.pdf
Australian National Data Services (ANDS)	One of the three major components of NCRIS initiative 5.16, Platforms for Collaboration. ANDS is the mechanism by which scientific data can be curated, preserved and discovered for future re-use. See *eSIM*.
Australian National University Supercomputing Facility (ANUSF)	Host of the APAC National Facility.
Australian Partnership for Advanced Computing (APAC)	National partnership for developing nationwide advanced computing infrastructure and associated programs in research, education and technology diffusion. APAC is being replaced by NCRIS PfC.
APAC National Facility (APAC-NF)	Provider of the Peak computing facility to researchers across Australia. http://nf.apac.edu.au/
Australian Research Collaboration Service (ARCS)	Joint venture to enhance collaborative e-Research capabilities across Australian research institutions. The delivery mechanism for the Interoperability and Collaboration Infrastructure (ICI) component of NCRIS PfC and related requirements. *See also ICI.* http://www.arcs.org.au

Battelle Institute	Global science and technology enterprise that develops and commercialises technology and manages laboratories for customers. Headquartered in Columbus, Ohio, it manages or co-manages national labs with 20,000 staff members and has an annual R&D budget of A$4.4B. http://www.battelle.org/
Blade	Compact self-contained computer system on a printed circuit board, which is then installed into a rack within a chassis. May contain multiple CPUs, memory and local disk, but with common items such as the power-supply removed to the chassis. A higher density of compute power per unit volume and electrical energy is achieved.
Business Unit (BU)	An identifiable organisational part of CSIRO, which has its own budget and accountability processes – typically a Division or Flagship.
Capability computing system	Large computing system capable of running large individual tasks, typically requiring shared memory and high specification processors.
Capacity computing system	Computing system designed to maximise throughput for large numbers of jobs or users, typically a large cluster system with commodity processors.
Centre for Australian Weather and Climate Research (CAWCR)	Joint research organisation currently being established between the Bureau of Meteorology and CSIRO Marine and Atmospheric Research.
Cluster	Computer system composed of an assembly of smaller units. Generally refers to a networked assembly of *Nodes*, where each node is a *Blade* or desktop PC containing a small number of commodity processors.
CCSD	Computing and Computational Sciences Directorate, the IT and computational services division of Oak Ridge National Laboratory, USA (ORNL). http://computing.ornl.gov/
CSC	Finnish IT Centre for science. Provider of ASC facilities for Finland's university and research sectors. http://www.csc.fi/english
Computational Science	Science performed by computer, typically involving evaluation of mathematical models or theories, numerical simulations, and/or manipulation of large datasets. Typically (but not always) refers to the applications, as well as the underlying techniques.
Condor	Free scheduling software used to perform cycle harvesting.

Centralised Facility / Consolidated Facility	Computing system located and managed in a common area, typically either one or more *Shared* systems, and/or a set of collocated *Dedicated* systems. *See also Local Facility*.
Central Processing Unit (CPU)	Integrated circuit at the heart of a computer which executes program instructions, also referred to simply as "processor". *See Multicore processor*.
CSIRO Data Store	HSM system operated by CSIRO HPSC group and its predecessors since 1991. Has grown in capacity from 0.4TBytes in 1994 to its present size of ~800TBytes.
CSIRO IM&T eScience	Group within CSIRO IM&T focused on scientific computing, including high performance computing (CSIRO HPSC), workstations, grid services and cycle harvesting. http://escience.arrc.csiro.au/twiki/bin/view/Main/WebHome
CSIRO HPSC	Group within CSIRO IM&T that provides access to high performance computing facilities for researchers, including in-house compute systems, the CSIRO Data Store and access to HPCCC and APAC partner facilities. http://www.hpsc.csiro.au
CSIRO IM&T	CSIRO Information Management and Technology function. Provides IM services including library and records, email, service desk, desktop support, IT security and networks, storage and applications.
Cycle Harvesting / Cycle Stealing	Distribution of computing tasks to make use of idle CPU time on desktop computers. *See Condor*.
Data Migration Facility (DMF)	Commercial HSM software used on CSIRO Data Store.
Dedicated Facility	Computing system assigned to a limited range of applications or staff, such as a specific project team. *See also Shared Facility*.
Ensemble	Set of multiple realisations of a physical or model system, typically with perturbed initial or boundary conditions, over which characteristics of the system may be averaged.
e-Repository	Storage infrastructure and environment for managing digital information.
e-Research	Use of information technology to support existing and new forms of research.
e-Science	Use of information technology platforms for conducting scientific research, typically of a computationally intensive, data-intensive or dynamically collaborative nature using distributed networks. *See also CSIRO IM&T eScience*

eScience Information Management (eSIM)	CSIRO IM&T policy framework for curation and stewardship of digital information, including science data, images, software, publications and other knowledge assets, to facilitate collaboration and future research through ready semantic discovery of prior digital objects and data.
Flop/s, Teraflop/s	Floating point Operations Per Second, widely used as a performance measure for scientific computing.
FPGA	Field-Programmable Gate Array. Specialised processor that can be programmed to perform repeated sequences of operations with high efficiency.
Fraunhofer-Gesellschaft (FhG)	German society of 56 institutes for performing applied research in the engineering sciences, with 12,500 staff working at ~40 locations.
Grid Computing	Application of computing resources across organisational boundaries to a single problem at the same time - usually a scientific problem that requires a great number of computer processing cycles or access to large and/or shared datasets.
Grid Services	Software, hardware and protocols that enable users to transparently access grid resources on demand.
Hierarchical File Store (HFS)	Storage system involving multiple layers of media, typically both disk and tape, to maximise the trade-off between cost and capacity of the media.
Hierarchical Storage Management (HSM)	Software product for managing a Hierarchical File Store, i.e. for automating the exchange of files between different layers of storage in a manner that is largely transparent to the user who sees all files as belonging within a single filesystem.
High Performance Workstation	Fast desktop computing system, typically with multicore or other high specification processors, used for numerically or graphically intensive work.
High Performance Computing (HPC)	Generic term referring to specialist hardware for handling large calculations, datasets and/or user bases, for both *Capability* and *Capacity* roles. The HPC abbreviation is now increasingly used to denote *High Productivity Computing.*
High Performance Computing and Communications Centre (HPCCC)	Joint Venture between CSIRO and the Bureau of Meteorology, providing specialised systems and support for operational and research computing. http://www.hpccc.gov.au/
High Performance Scientific Computing (HPSC)	High Performance Computing applied to science and engineering research. *See also CSIRO HPSC.*
Informatics	The science and practice of information processing and management of information systems, for extraction of scientific knowledge from data.

Information Management & Technology (IM&T)	*See CSIRO IM&T above.*
Interoperability and Communications Infrastructure (ICI)	NCRIS PfC initiative to enhance collaborative e-Research capabilities across Australian research institutions, to be delivered through the Australian Collaborative Research Service (ARCS). *See ARCS.*
Information Technology Infrastructure Library (ITIL)	UK Government sponsored initiative providing a set of adaptable processes and techniques for Service Management and Service Delivery in information technology. Used by many large organisations in Australia. Used by CSIRO IM&T to manage its service delivery quality and reliability.
iVEC	Formerly Interactive Virtual Environment Centre - PAC based in Western Australia.
ITWM	Fraunhofer–Gesellschaft Institute for Industrial Mathematics, specialising in the development of mathematical applications for industry, technology and economy.
Linpack	Linear algebra benchmark widely used to measure a computer's performance for numerically intensive calculations. *See TOP500.*
Local Facility	Computer system located close to the users, usually but not always *dedicated* to those users. *See also Shared Facility.*
Model	Simplification or idealisation of reality.
Modelling	Development and application of (mathematical) models. Sometimes used interchangeably with *Simulation.*
Multicore processor	Integrated circuit with more than one CPU (core). Each unit is usually able to run multiple instructions at the same time and each has its own cache. Currently dual-core and quad-core architectures are most common.
MyCluster	Colloquial term for small local cluster system, typically dedicated to one or a few users.
National Center for Computational Sciences (NCCS)	U.S. DoE organisation operated by ORNL, which develops and provides ASC services to national laboratories, universities, and industry. http://nccs.gov/aboutus/index.html
NCF	National Computing Facility (Netherlands). A division of NWO, NCF provides ASC services to Dutch universities and research organisations such as TNO. http://www.nwo.nl/nwohome.nsf/pages/ACPP_4X6R5C_Eng

National Collaboration Network (NCN)	Proposal for high bandwidth network based on CSIRO ICTC technology, incorporated into CSIRO's submission to the NCRIS Platforms for Collaboration capability. *See PfC, ICI, ARCS.*
National Collaborative Research Infrastructure Strategy (NCRIS)	$540M DEST program to boost collaboration among the component institutions of the Australian national innovation system. http://ncris.dest.gov.au
National Computing Infrastructure (NCI)	Component of NCRIS PfC to provide large-scale computing infrastructure for Australian research institutions.
Node	A blade or a commodity computer that serves as the building block in a cluster. Each node may contain multiple CPUs. Each node is managed by a single instance of an operating system, typically Linux or Windows.
Non-Uniform Memory Access (NUMA)	Computer architecture that uses a fast interconnect bus to mimic shared memory. *See Advanced Cluster, SMP.*
NWO	Netherlands Organisation for Scientific Research. Funding administration body. http://www.nwo.nl
ORNL	Oak Ridge National Laboratory. The US Department of Energy's largest science and energy laboratory, operated by Battelle. Staffing 4200, budget $1.2B
PACs	Partnerships for Advanced Computing – shorthand term referring to the state based high-performance computing centres that form APAC.
Parallelism, Coarse-Grained	Programming paradigm whereby multiple instances of the program execute as separate processes, with specified inter-process communication. Well suited to *Cluster* systems.
Parallelism, Fine-Grained	Programming paradigm whereby a single instance of the program generates multiple processes or threads for specified tasks. Well suited to *SMP* or *Multicore* systems.
Parallel Vector Machine (PVM)	Large specialised computer system featuring vector processors and shared memory. *See Vector Processor, Supercomputer.*
Parameter Sweep	Workflow for running multiple instances of a program to test the effects of changing specific parameters, variables or conditions.
Peak Facility	Largest computing system available in a specified community, country or organisation. Typically a *Capability Computing* system. In NCRIS context, the ANU ASC facility.

Peak Performance	Hypothetical maximum performance of a computer system, determined by multiplying the clock speed by the maximum number of possible floating-point operations per clock cycle. *See Top500, Linpack.*
Platforms for Collaboration (PfC)	NCRIS initiative, which provides a computation foundation for other discipline specific initiatives. PfC has three major components being computing infrastructure (NCI), digital information management curation and stewardship (ANDS), and collaborative tools (ICI/ARCS). It also includes a network bandwidth capability (AREN) and a federated access permissions facility (AAF).
QCIF	Queensland Cyber Infrastructure Foundation
SAPAC	South Australian Partnership for Advanced Computing.
Scalability	Ease of adapting or expanding hardware, software or workflow to manage increased demand.
Scalar Processor	A CPU that operates only on single numerical values at a time. *See also Vector Processor.*
Scientific Computing	Use of computing and information technology systems to perform innovative components of scientific and engineering research.
Shared Facility	Computing system or resource assigned to multiple users, groups or organisations, in order to maximise capability or utilisation. *See also Dedicated Facility.*
Simulation	Alternative representation of reality, for example a computer experiment performed on a *model* system to observe a physical process or event.
SIP	CSIRO Science Investment Process.
Supercomputer	Term used to describe a specialised computing system that is among the fastest currently available. *See TOP500.*
Symmetric Multiprocessor (SMP)	Single computer system consisting of multiple processors that share a common operating system and memory.
Theory	Postulated relationship or interaction between observable quantities used to describe natural phenomena. Relationships that are mathematically complex may require computational evaluation.
TPAC	Tasmanian Partnership for Advanced Computing
TNO	Netherlands Organisation for Applied Scientific Research. http://www.tno.nl

TOP500	List of the fastest 500 computers in the world, as measured by the Linpack linear algebra benchmark. The list is updated every 6 months. http://www.top500.org
Thread	Sequence of tasks performed within a process.
User Support	Base level support services to assist users in accessing resources and performing day-to-day tasks.
Vector Processor	A CPU designed to run repeated mathematical operations on a series of data elements within a single clock cycle. *See also scalar processor.*
VPAC	Victorian Partnership for Advanced Computing
Virtualisation	Abstraction of a physical resource to make it appear different logically than it is physically. Enables the resource to be pooled and/or shared and to create a management layer or control point for the virtualised resource. Leads to increased utilisation, greater flexibility and higher availability of IT resources.
VTT	Multidisciplinary Finnish scientific research organisation, employing about 2800 staff.
Water Resources Observation Network (WRON)	Project to develop technical framework and standards required to support water information management, as part of WfHC Flagship. http://wron.net.au/
Wet Science	Generic term for laboratory science, field measurements etc, as opposed to computational science.
Workflow	Repeatable pattern of activity or operating procedure, typically involving a high level of automation of and interaction between component tasks.
Workstation	Single computing, data or visualisation system, typically with one or a few CPUs in a single operating system image and dedicated to a specific user and/or project task.

Appendix 2: History of High Performance Computing in CSIRO

1. Introduction

CSIRO has had a long involvement in computing, starting with CSIRAC, whose construction started before CSIRO was formed from its predecessor CSIR in 1949. This history sketches details of peak systems, organisation, partnerships, and support groups.

2. CSIRAC[20]

In 1947, Maston Beard and Trevor Pearcey led a research group at the Sydney-based Radiophysics Laboratory of the Council for Scientific and Industrial Research (now known as CSIRO), to design and build an electronic computer.

The resources they had available included the vacuum tube or "valve" technology and the pulse techniques developed for radar systems during World War II. Their developments paralleled, but were to a considerable extent independent of computer developments in Europe and the USA.

The CSIR Mk1 ran its first test programs in late 1949, and it was the fifth electronic stored program computer ever developed. It embodied many features novel at the time and was able to operate more than 1000 times faster than the best mechanical calculators. The machine was officially opened in 1951 and used to solve problems both for the Radiophysics Laboratory and outside organisations. It was decommissioned in 1955 and shipped to Melbourne.

On 14 June 1956 the Mk1 was recommissioned and renamed CSIRAC and the new Computation Laboratory at the University of Melbourne was officially opened. CSIRAC was available as a general computing workhorse - from June 1956 to June 1964 over 700 computing projects were processed.

In November 1964, Dr. Frank Hirst switched CSIRAC off for the last time and it was donated to the Museum of Victoria.

CSIRAC is now the world's oldest surviving computer. The machine and its detailed history can be seen on display at the Melbourne Museum[21].

The following quote is from ref. 22.

> In 1955 a committee was established to decide whether development of CSIRAC should continue. The committee consisted of Hartree and Comrie of Cambridge University and Myers from the University of Sydney. Apparently they concluded that Australia should concentrate on primary industry and not waste resources attempting to compete with Britain and the US in computers.

[20] This section is taken from: http://www.csse.unimelb.edu.au/dept/about/csirac/.
[21] http://melbourne.museum.vic.gov.au/exhibitions/exh_science.asp?ID=561786.
[22] J. Deane, A Picture History of CSIRO Radiophysics 1939-1984, CSIRO Division of Radiophysics, 1985.

3. Computing Research Section, Division of Computing Research, CSIRONET

Around 1963 CSIRO formed the Computing Research Section, which later became the Division of Computing Research. The headquarters were at Black Mountain, ACT. A CDC 3600 was acquired and installed there, with smaller CDC 3200 systems in Melbourne, Sydney, and possibly Adelaide and Brisbane.

Jobs were locally prepared on punch cards, and sent by courier for processing at the major centres, with line printer output returned by courier. Innovative work was done on operating systems, an automatic hierarchical storage facility, video displays, and networking. By 1969, CSIRO had a packet-switched network, which was eventually extended to most CSIRO sites, allowing interactive usage, remote job entry, and output printing.

In 1973, CSIRO acquired a CDC Cyber 76, then the premier supercomputer in the world. CSIRO provided services to government departments and others.

The CDC 3600 was de-commissioned in 1977.

Around the late 1970s, CSIRO acquired FACOM systems, and provided further services, mainly in commercial data processing to government departments. It also acquired a Braegan Automated Tape Library, and built a Terabit File Store.

In 1984, CSIRO acquired a CDC Cyber 205, one of the supercomputers of its day.

The CDC Cyber 76 was de-commissioned in 1985.

By the early 1980s, the name CSIRONET was used for the whole organisation.

In the mid-1980s, when the total CSIRONET staff had grown to about 150, the Division of Information Technology was formed from the research staff of CSIRONET, and the services were sold to a private company in two tranches over two years. CSIRO bought some services from CSIRONET, but managed the Cyber 205 separately, with an Advanced Computing Support Group to work with users. CSIRO ran a Merit Allocation Scheme to allow non-CSIRO scientists to have access to the 205, but this led to poor utilisation, because of the difficulty in using both the Cyber 205 and the incompatible front-ends, and because of the charging policy for the front-end usage.

CSIRO ceased using CSIRONET at the end of the life of the Cyber 205 in 1990.

4. Supercomputing Facilities Task Force, JSF, SSG

In 1989, CSIRO formed a Supercomputing Facilities Task Force to work on acquiring new facilities and services to meet CSIRO’s computing needs. The task force was chaired by Mike Coulthard, and included Charles Johnson, Bob Smart and Robert Bell.

A short request for information was issued, and proposals were evaluated. After benchmarking, the Policy Committee on Computing accepted the proposal from Cray Research in conjunction with Leading Edge Technologies for a facility at Port Melbourne. A Cray Y-MP2/216 was acquired by CSIRO, installed at Port Melbourne in March 1990, and run by LET and Cray Research, with resources shared between the parties in the Joint Supercomputing Facility. A Supercomputing Support Group was formed, and based in the Carlton Offices of the Division of Information Technology, with Robert Bell as User Support Manager. A Supercomputing Facility Users Management Committee was formed to advise CSIRO on the running of the facility, which was managed by DIT under John O'Callaghan.

The system was allocated using a development fund structure, to which Divisions were invited to contribute monthly, with the shares on the system set up to be proportional to the contributions. On 14th November 1991, the CRI Data Migration Facility was initiated on the system, to provide hierarchical storage management for users.

In 1992, Leading Edge Technologies went bankrupt.

5. CSF, University of Melbourne, Cray Research, SSG

When the demise of Leading Edge was imminent, CSIRO worked with Cray Research to provide alternative facilities. CSIRO acquired a Cray Y-MP3/64, which was installed at the University of Melbourne, and managed by Cray Research and the University. The University had a 10% share. The arrangements worked very smoothly, with Alan Bell being the Director of the University's Information Technology Services at the time.

In June 1993, a StorageTek Automated Library System was acquired with money from the Development Fund, to provide automatic tape mounting for DMF. In January 1994, a fourth processor was acquired cheaply resulting from access payments by a commercial user (Fluid Thinking).

At the end of its life in September 1997, the system had averaged 98.5% CPU utilisation.

6. HPCCC

In late 1995, CSIRO CEO Malcolm McIntosh proposed the establishment of a new supercomputing facility with the Bureau of Meteorology. This led to the formation of the High Performance Computing and Communications Centre in April 1997. Following a joint tender process, an NEC SX-4 supercomputer was delivered in October 1997, and CSIRO acquired a Cray J90, to manage DMF storage and act as a front-end. The four CSIRO staff of the Supercomputing Support Group moved to the Bureau's Head Office in Melbourne in September 1997.

In 1999, an NEC SX-5 was acquired, and in 2001, a second SX-5 to replace the SX-4.

In 2002-3, the HPCCC tendered again, and contracted with NEC to acquire an SX-6/TX7 system to be situated in Docklands at the Bureau's new Head Office. The initial system was delivered in December 2003, and staff moved to Docklands in August 2004.

7. CSIRO HPSC

In 2002-3, CSIRO reviewed its High Performance Computing services. The review concluded that CSIRO needed to diversify its major computing systems beyond the NEC vector systems, and that more support staff were needed. The review also concluded that CSIRO was under-investing in HPC, but did not present an adequate case for growing the investment.

A new group, CSIRO High Performance Scientific Computing was formed, with Dr Rhys Francis as Director. Staff numbers were increased to 12. An SGI Altix system was acquired to host the Data Store in the new location, and was subsequently expanded to provide a system with large shared memory for parallel processing and large-scale data analysis. IBM Bladecenter cluster nodes were acquired to provide a new range of computing services, with portions of this system being dedicated to specific research groups. The diversification was funded by reducing CSIRO’s commitment to the HPCCC and its vector systems.

CSIRO collaborated more closely with APAC and partners, to undertake Computational Tools and Techniques development, and to provide Grid services. CSIRO purchased time on the APAC National Facility, and also made an investment of $0.5M into iVEC, to provide local facilities for WA users.

In July 2006, CSIRO HPSC became a part of CSIRO IM&T.

Appendix 3: Technology Trends

1. Hardware Options

Different computational science applications require different hardware and software platforms to be effective. From a hardware perspective, a modern desktop computer is now powerful enough itself to be a valuable research tool. It can be used to run simulations and perform visualisations that would have required more specialised hardware only a few years earlier. However, a single commodity PC is still limited in its capability to handle large-scale computing problems.

One inexpensive alternative is the use of unused computing time on groups of desktop PCs, known as “cycle harvesting”; alternative labels include cycle stealing, CPU scavenging and volunteer computing. The technique is best known through the SETI@home project[23], but is also used for studying particle physics, protein folding, new malaria and HIV vaccines and other large-scale applications[24]. This approach is only viable for a minority of applications, where the computation can be divided into many independent small-memory tasks over long periods.

A more common option is a conventional cluster of networked commodity PC components. Cluster systems have proven to be very cost effective and many vendors now sell dedicated cluster hardware to meet market demand. Modern cluster systems are no longer simply a network of desktop boxes, but rather are optimised by packing large numbers of systems together, reducing energy usage and heat generation by as much as a factor of six, and allowing for greater reliability, better space utilisation and easier systems maintenance.

On a large scale, “capability” computing systems have hardware specifically designed for high-performance computing. Such machines typically have hundreds or thousands of processors and large shared memory capacity. Some vendors offer specialised processors that can perform specific numerical operations with much greater performance than commodity processors.

2. Software Options

Ultimately, scientific computing is about software; the hardware merely provides a platform for running this software. There are a broad range of complex issues concerning development and manageability of software and in a HPC context, and the relative importance of different aspects of these issues differs from commodity computing contexts.

The cost of developing and maintaining software is high, and ASC constitutes a broad set of niche markets with complex discipline-specific software. Only limited advantage can be drawn from the economies of scale present in more general computing, though these are key for many underlying services, software integration and software standards, such as operating systems (proprietary and Linux), scripting frameworks (python, perl, java, R, …) and communication/service standards and frameworks (web services, MPI).

For specific domains, availability of appropriate software and the closely related issues of *portability* and *portability of performance* are significant. To achieve an appreciable proportion of the aggregate theoretical performance of computing hardware for real-world problems is difficult and requires parallelisation. Software must necessarily also be complex to deal with complex models.

[23] http://setiathome.berkeley.edu/.
[24] http://boinc.berkeley.edu/.

As a result, the cost of developing software is very high and its development/deployment is usually focused on a quite limited range of platforms. Use of libraries (re-using well-tested code) can help, although libraries are not immune to portability, performance and availability issues. The software can constitute intellectual property, becoming a source of competitive edge in research, or a product to be protected or exploited. Community approaches are becoming quite prevalent in the research environment, as individual groups can afford neither to develop their own software/models in isolation, nor pay for commercial development/software. In this context, competitive edge is gained by integration of software and/or services sold related to the software.

Additionally, the cost of migration from one software package to another – or starting with a package from scratch – can also be very high. Licensing regimes can also cause difficulties, as commercial vendors are not consistent with their treatment of cores, CPUs/sockets, machines, clusters and sites. Moving computational science applications between small scale and large scales usually involves encountering these problems.

Both the commercial and non-commercial aspects of software costs can overwhelm the hardware costs, particularly if the cost of support personnel is taken into account.

The framework that the ASC software runs in can help or hinder multidisciplinary activity across different scientific domains. Choosing appropriate standards, module interfaces etc. is often the key to enabling different research groups to work together effectively.

3. Data Growth Rate

The faster computer hardware allows scientists to use higher resolution models and perform ensemble runs. These increased capabilities, coupled with new data collection devices, are leading to a huge increase in the amount of data being generated. Hence, data generation in most disciplines is growing at an exponential rate.

This naturally leads to problems with data storage and the delays in transferring large volumes of data across computer networks. As a result, it is often preferable to move the computation to the data rather than the other way around. Once the data is available, however, processing and visualisation can be equally problematic. These problems can be further exacerbated by unbalanced computer systems, which can leave fast CPUs idle, waiting for data to be retrieved from disk.

4. Storage Choices

The amount (and therefore cost) of storage required by research projects will vary widely. Moreover, the choice for the type of management/backup/failover of the storage is also a key question, which may in turn vary the amount of storage required.

The key questions are:

- How much data is required to be available on-line at any time?
- How is the data to be protected against different hazards, such as disk failure, rogue process, human error or environmental disaster?
- Where does the data need to be located?
- What is the ability and cost to reacquire the data? (In some cases, the data captured may be unique and therefore not able to be recaptured).

Protecting the data: The term storage usually covers a mixture of disk and tape, where the tape is used to at least backup and therefore protect the disk contents. Backup can also be done to more disks or to optical media such as DVD. Both tape and optical media can be automated using robotic enclosures. They are also easily removed off-site for safer storage, e.g. protection against fire. While disk costs are continually falling, on-line disk storage, backed up to tape with off-site storage, is the most expensive data storage option. However, it can be significantly less expensive than the cost of losing some critical data.

- 100% Available data: Disk can either be a RAID array (Redundant Array of Inexpensive Disk) or just separate drives. The most usual RAID is RAID1 where data is mirrored to two separate disk drives, or RAID5 with distributed parity, so in the event of a single disk failure no data is lost and processing continues. However, hardware failure is only a small percentage of the causes of data loss – software and human causes are far more common, and RAID does not guard against these
- Curating the digital data: Curation is the process by which scientific data is procedurally filed away with appropriate metadata for indefinite retention. Subsequent users can then semantically discover and use the curated data to benefit their own work. In turn, they will curate their results for others to use
- Managing large amounts of data online: It is prohibitively expensive to keep vast amounts of data on-line "in case they are needed". Furthermore, the energy costs are no longer negligible. Hierarchical Storage Management is a software product, which manages files on disk. When a file meets defined criteria, typically related to size and age, it is automatically moved to robotic tape, but a "stub" is left in its place so that it looks like the file is on disk. This process is called staging or migration. If a program/user accesses the stub, the HSM software transparently reloads the file to disk (re-staging or recall). This multiplies the effective size of the disk many fold, reducing cost by similar multiples. The data all appears to be on-line. It also provides backup and deep archive at the same time. HPC sites increasingly favour the use of HSM because it requires no manual intervention to migrate the data to/from tape as required, and is the only economic way to do this.

Which to choose? It may be that during a research project, the scientific data need only be protected so that months of work are not lost due to disk failure, systems or human error. An automated backup process to local or remote secondary storage will achieve the result, but will be difficult and costly for large datasets.

HSM managed storage at HPSC (and soon at iVEC) is available via the intranet, and the eSIM project should deliver the framework for curation.

Local and shared clusters may well benefit from local HSM managed storage, since it reduces the requirements for manual intervention, and more importantly, reduces the opportunities for inadvertent human errors. It will also bring the benefits of large-scale storage, close to the clusters.

There are few instances where unprotected storage would be a responsible choice for important data.

5. CPU Trends

In line with Moore's Law[25], chip manufacturers have managed to double the performance of their CPUs approximately every 18 months over the last 30 years. As a result, scientists have been able to

[25] G. E. Moore, *Electronics* **38** (1965) 114.

rely on performance improvements in running their own applications, with little effort on their part other than upgrading to faster hardware.

Unfortunately, due to fundamental physical issues with heat generation, the rate of performance increase of commodity CPUs has now slowed dramatically. This means that serial scientific applications will no longer run faster as they are moved to new systems.

One option for researchers is the use of specialised CPUs such as Field Programmable Gate Arrays (FPGAs) or Digital Signal Processors (DSPs). In the past, specialised architectures (e.g. vector, RISC, DSP) have typically not kept up with commodity processors for cost-effectiveness, because they have not benefited from the same economies of scale. Typically, such architectures require specialist software or programming expertise to fully exploit the hardware potential.

As performance growth of individual commodity CPUs flattens out, chip manufacturers are developing alternatives to increase performance per unit cost. One favoured approach is the use of multicore CPUs, whereby multiple processing units (cores) are built onto a single chip. Two-core and four-core CPUs are now relatively commonplace and the number of cores is likely to grow rapidly over time. Computing clusters based on multicore systems are now becoming common.

Two independent concurrent instances of serial code running on a dual-core CPU generally run no faster than on a single CPU and in many instances will run much more slowly. This is because of contention between the processes for use of shared chip components including shared on-chip memory/cache and data paths for off-chip access. Parallel programming techniques, whereby sections of a program run concurrently on multiple CPUs, are necessary to address many of these problems if the potential of these multicore CPUs is to be realised.

6. Parallel Programming

Unfortunately, parallel programming is difficult for a variety of reasons. It often necessitates a complete redesign of the original program; debugging is hard because different parts of the code are running concurrently, and code optimisation is non-trivial. Furthermore, there are fundamental limits to the scalability of parallel code, as expressed by Amdahl's Law.[26] This leads to an intimate relationship between software performance and the relative performance of different aspects of the underlying hardware, notably latency and bandwidth to access data or memory in different locations.

The hardware architecture has a direct impact on the type of parallel programming model used. Where parallel processors share memory, *fine-grained parallelism* may be used, whereby a program runs as a single instance or process except when it explicitly spawns multiple processes. OpenMP is a widely used standard for fine-grained parallelism.

For distributed memory systems such as clusters, one needs to adopt *coarse-grained parallelism*, where multiple instances of the program execute as separate processes in separate memory domains (e.g. cluster nodes) except where explicit inter-process communications are called. Typically, coarse-grained parallelisation is more difficult than fine-grained parallelisation. The most commonly used standard for coarse-grained parallelism is MPI (Message Passing Interface).

---

[26] G. Amdahl, *Proc. AFIPS* **30** (1967) 483.

The future of HPC will see increasing requirements for large codes to run on systems where small numbers of CPUs share memory, e.g. multicore clusters. This requires a mixture of both fine-grained and coarse-grained parallelism, which further increases the complexity of the programming task.

7. Grid Technologies

A more recent development is the concept of grid computing, whereby users are detached from the physical location of the data or computing resource that they are using.

IBM defines grid computing[27] as "the ability, using a set of open standards and protocols, to gain access to applications and data, processing power, storage capacity and a vast array of other computing resources over the Internet. A grid is a type of parallel and distributed system that enables the sharing, selection, and aggregation of resources distributed across 'multiple' administrative domains based on their (resources) availability, capacity, performance, cost and users' quality-of-service requirements."

Grid environments are based on a software "middleware" layer that can provide:

- access to distributed compute and storage systems
- access to shared and/or geographically dispersed – or *federated* – datasets
- access to shared software resources through portal environments
- authentication protocols for identifying access rights of users while maintaining site security.

By providing standardised environments that span across organisational boundaries, grid technologies are the key to enabling large scale collaborative e-Research.

8. Economies of Utilisation

The demands for CPU usage by individual scientists, or small groups of scientists, usually fluctuate substantially over time. This means that their maximum system requirements are typically rather larger than their average requirements. The more users are aggregated together, the more these fluctuations cancel each other out. For the aggregated population, the average and maximum requirements thus converge. This leads to more efficient utilisation of resources, which generate significant economies, as described below.

[27] http://www-306.ibm.com/software/globalization/terminology/gh.jsp.

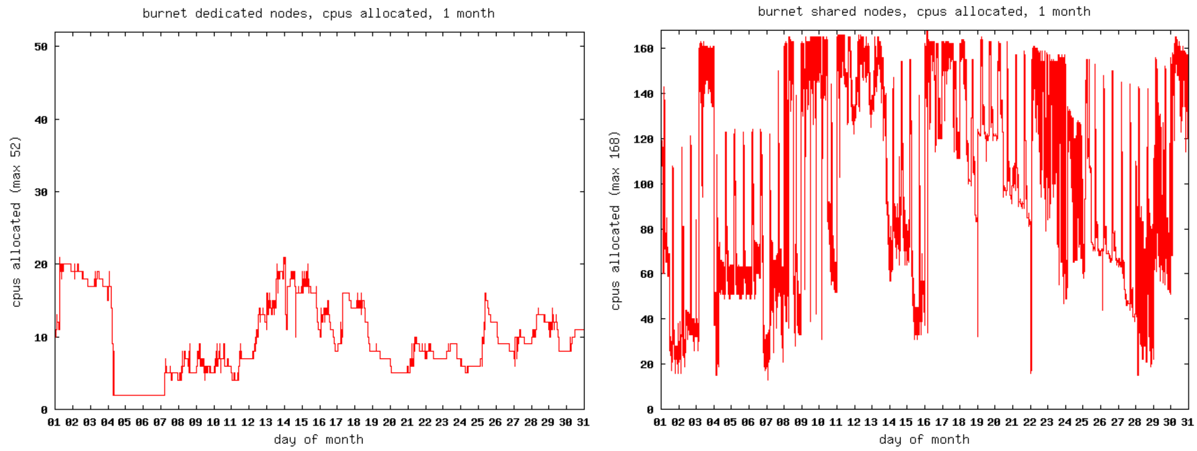

Figure 7. CPU utilisation of project-dedicated and shared access clusters.

The effect is illustrated in Figure 7, which compares typical monthly utilisation for a consolidated shared cluster versus a smaller but equivalent cluster dedicated to users in a specific project area; in each case, the vertical axis represents the maximum capacity of the cluster. Equivalent ongoing data can be seen for HPSC shared and dedicated systems[28] using the Ganglia monitoring tool[29].

In Table V we can see the relative utilisation of the Altix SMP and IBM cluster systems, averaged over all users during 2006-7, using the following measures:

A. Sum of maximum CPU allocation for all users, divided by average allocation for all users

B. Sum of maximum CPU allocation for all users, divided by total size of system

C. Sum of usage-weighted CPU allocation for all users, divided by average allocation for all users

D. Sum of usage-weighted CPU allocation for all users, divided by total size of system.

Measures C and D correspond to the second moments of the usage histogram of each user[30], i.e. *they reflect the number of CPUs allocated to users while they are doing the bulk of their work*.

Table V. Utilisation economies for shared HPSC systems.

System	Type	CPUs	Users	Economy Factor			
				A	B	C	D
SGI Altix	NUMA	120	81	21.5	11.7	9.3	5.1
IBM BladeCenter	Cluster	220	75	17.4	8.3	8.1	3.9

The economy of grouping the users together varies substantially – from a factor of ~4 to >20 – depending on the measure used. All measures reflect that *the cost savings are substantial and worth exploiting where possible*. These measurements are for mid-sized systems; increasing system size will

28 http://intra.hpsc.csiro.au/ganglia/.
29 http://ganglia.sourceforge.net/.
30 Details on calculations and usage data are available on request from hpchelp@csiro.au.

increase this economy. Measure C shows that for typical work patterns for a community of ~80 users the relative value of aggregating these users onto a shared system is a factor of 8-9, i.e. approximately an order of magnitude.

Appendix 4: Cost of Ownership

Summary

The total cost of ownership (TCO) for operating a computing system can be estimated using the following sum:

Total = compute system
+ software
+ storage
+ operating expenses

Indicative costs for the compute and operating components are given in Table VI, while storage costs are given in Table VII. These values are largely derived from a detailed evaluation of HPSC and local cluster systems[31] performed in 2005, and provisional estimates from the IM&T Foundation Project. Annual costs are given, based on an assumed 3 year life cycle.

- Hardware and operating expenses are typically roughly proportional to system size, i.e. number of CPUs (sockets). These costs do not vary greatly over time, but the CPU performance increases through increased clock speed, improved architecture or greater number of cores.
- Operating expenses can be further divided into environmental costs (power, cooling, and floor space) and staffing costs (installation, systems administration, user support, and networking).
- Software costs are highly variable and application-specific, and hence no "typical" amounts are itemised here.
- Storage costs have some economies of scale, but are treated here as being approximately proportional to size of holdings.
- Storage costs per terabyte diminish significantly over time – typically about 40% reduction per annum.

In many cases, some of these costs are hidden (but still real), for example, where computer systems occupy laboratory space or are self-supported by researchers.

Table VI. Compute and Infrastructure Costs.

ASC type	Acquire ($/CPU PA)	per CPU PA		Typical number of CPUs	Processor PA cost (000's)	Environmentals PA (000's)	Floor-space PA (000's)	People Cost PA (000's)	Total cost PA (000's)	Peak Power (Gflops/CPU single core)	Power (Tflops)	Cost TFlop/PA @peak performance (000's)
		Environmentals	Staffing L1-L4									
SMP Altix Cluster	$1,390	$275	$1,320	128	$178	$35	$81	$169	$463	5.2	0.6656	$696
Blade Cluster	$1,000	$200	$1,800	400	$400	$80	$16	$720	$1,216	6.4	2.56	$475
MyCluster	$1,000	$200	$1,800	10	$10	$2	$8	$18	$38	6.4	0.064	$594
Workstation	$1,500	$200	$1,000	2	$3	$0.400	$0	$1	$4	6.4	0.0128	$344
Cycle harvesting (Condor)	n/a	$40 @ 40% utilization	1.5 EFT	2400	n/a	$96	$0	$225	$321	6.4	21.6	$15

[31] https://teams.csiro.au/sites/ascs/References/HPSC%20Costings%202005-6.xls.

Table VII. Storage Costs for 2007.

Storage Type	Cost/TB PA	Comments
CSIRO HPSC HSM storage	$1,000	Old / large files migrate to tape
CSIRO Data Centre fibre channel disk*	$16,000	Fast disk
CSIRO Data Centre SATA disk*	$10,000	
CSIRO Data Centre tape*	$9,000	Incl. incremental backups & offsite retention
Unmanaged local disk storage	$300	

* Provisional costing by Storage Project.

Example

The total cost of ownership to CSIRO for a 20-CPU cluster with 5 TB initial SATA and archival tape backup is approximately 20 × $3,800 + 5 × ($10,000 + $9,000) = $76,000 + $95,000 = $171,000 per year for three years, or $513,000 in total, not including any software costs. This is for a nominal hardware acquisition cost of $60,000.

This calculation assumes that the storage requirement will grow at the same rate at which the storage medium gets cheaper. This corresponds to a factor of 1 / (1 – 0.4) = 1.67, or 67% growth per year, i.e. growth from 5 TB to 14 TB over the 3 years.

If the storage requirement is constant at 5 TB over the 3 years then the total cost over 3 years is reduced by $99,000, to $414,000.

If the growing 5-14 TB storage requirement is HSM rather than disk + tape then the total cost over 3 years is $237,000.

Assumptions

Floor space

Assumes a single rack or two racks with $2.3m^2$ footprints and access, each at industry rate of $3,500/$m^2$ per annum.

Altix SMP – $81,333 PA based on HPSC costs incurred in 2005

Blade Cluster – $16,000 PA 2 × ($2.3m^2$ × $3,500)

MyCluster – $8,000 PA ($2.3m^2$ × $3,500)

Environmentals

0.14kWh/CPU Power × 1.3 (+30% for A/C) - $200/CPU PA @ $0.13 per kWh.

Hardware Acquisition

Acquisition - Altix SMP

Depreciation costs ($101k PA) + service costs ($77k PA) over 4 years. Total $712k for 128 CPUs = $5.56k/CPU for 4 years, or $1.39k/CPU PA.

Acquisition - Clusters

Typical configuration options, depreciated over 3 years:

- Dual 3.6GHz Xeon, dual core 2.8GHz Opteron
- 4 GB RAM
- 100 GB scratch disk per dual CPU node
- GigEthernet, Myrinet, Infiniband

People

Staff per CPU per annum (CSOF5)	System administration	$900	Installation, troubleshooting, upgrades, 0.006 EFT per CPU
	Systems integration	$300	Networking, storage, desktop access etc
	User support	$600	Applications support, porting, optimisation, training, documentation, helpline etc

Support costs for workstations

1500 workstations out of 13,000 desktops/computers in CSIRO, ~11% of support staff.

Approximately 100 staff providing general support for all desktops/computers, hence ~10 staff supporting 1500 workstations.

10 EFTs $150,000 PA or $1,500,000 for 1500 workstations = $1,000 per workstation (multi CPUs) PA.

Peak performance

- Altix Cluster single core CPUs ~ 5.2 GFlop/s per CPU
- Blade Cluster quad core CPUs ~6.4 GFlop/s per CPU single core
- MyCluster quad core CPUs ~ 6.4 GFlop/s per CPU single core
- Workstation quad core CPUs ~6.4 GFlop/s per CPU single core[32]
- Cycle Harvesting ~6.4GFlop/s per CPU

Cycle harvesting

[32] e.g. Intel Woodcrest CPUs can perform 4 floating point operations per clock cycle, older Intel Xeons 2 flops per cycle.

2400 CPUs = 40% of the total of 6000 machines available.

Environmental costs \$0.01/hour for an idle PC, \$0.02/hour for a fully utilised PC. Average utilisation of a PC assumed to be 40%.

Additional power cost: 2400 × (\$0.01/h × 16h/day × 220days + \$0.02/h × 24h/day × 145days) × 40%

Equivalent of:

- \$35.20/PC pa for additional 100% utilisation during business hours,
- \$69.60/PC pa for additional 100% utilisation on weekends and public holidays,
- \$125.76/PC over 3 years additional power costs for an average utilisation of 40%

or \$301,824 for 2400 PC for 3 years @40%.

1.5 EFT support.

Peak processing power of a desktop machine ~ 6.4 GFlop/s per CPU or total peak performance 6.4 GFlop/s × 2400.

Appendix 5: Funding and Allocation Models

Management of Advanced Scientific Computing facilities involves several connected processes:

- Funding source
- Management of user allocations
- Management accounting, cost recovery and attribution
- Performance measurement and management.

This Appendix identifies options for these processes and discusses the relative advantages and disadvantages of these options.

History of ASC Funding Mechanisms in CSIRO[33]

During the 1970's and 1980's, CSIRONET shared systems were accessed through a charging mechanism, whereby costs were recovered on a monthly basis from group budgets using a complex utilisation formula. This mechanism had the following drawbacks:

- The computing service had mainly fixed (sunk) costs; hence the cost of running the service was not closely related to the amount of usage
- Usage was inhibited since users could not estimate their usage or costs in advance, leading to unpredictable and significant impact on group budgets
- The charges drove scientists to find alternatives – PCs, minicomputers, collaborators' resources – leading to a vicious circle of ever-increasing costs for the dwindling remaining users
- The net result was that CSIRO had multimillion-dollar systems that were grossly underutilised because scientists could not afford to use them.

In the 1990's a charge-back mechanism was introduced, whereby usage costs were attributed to divisions which then had their appropriation funds increased to match the costs. In addition, a Development Fund system was used from 1990 to 1997, whereby divisions contributed a proportional charge, which governed their access prioritisation using a Fair Share Scheduler.

- This system led to high utilisation (over 98% from 1992 to 1997)
- The Development Fund was used to purchase software and establish the CSIRO Data Store
- The charge-back system was a disincentive to some divisions, since it increased their measured appropriation earnings at a time when they were being assessed by their appropriation/external balance rather than by meeting a total budget.

In 2004 the charge-back system was dropped. Since then HPSC systems have had simple open access, with manual intervention for prioritisation where required to give an equitable balance between users.

- This period has seen a significant increase in the range of users, applications and divisions running on the facilities

---

[33] From B. Arrold *et al*, "HPSC and the HPCCC – Discussion Paper" Appendix C, 2006. http://escience.csiro.au/twiki/pub/Main/StrategyDocuments/CSIRO_HPSC-HPCCC_paper_July_2006_-_Brennan_John_pivotal_paper.pdf.

- System utilisation has still been significantly lower than in the previous decade, as observed worldwide due to the inherent difficulty of efficiently scheduling inhomogeneous jobs on large parallel systems
- The level of intervention required to date has been minimal to nonexistent.

CSIRO users have had access to portions of CSIRO's partner share of the APAC national facility on a first-come first-served basis. After a recent clarification, in 2007 CSIRO users became explicitly eligible for Merit Allocation Scheme (MAS) grants, which provide fixed annual allocations via peer reviewed proposals.

Funding Source Options

The two basic options are (i) corporate level funding, e.g. through CSIRO IM&T, or (ii) user-attributed funding, e.g. through SIP theme bids.

The main arguments for and against corporate level funding are as follows:

+ Historically shown to be most successful and cost-effective for organisation as a whole
+ Minimises barrier to entry – still the biggest ASC issue in CSIRO
+ Encourages widest possible usage of resource, i.e. growth of new areas of computational science
+ Can be used to create incentives for best-for-CSIRO practices, e.g. resource & knowledge sharing, in alignment with the CSIRO Strategic Plan
+ Builds internal collaborative behaviour by minimising direct internal competition
+ Much easier to project requirements collectively rather than individual users/groups
+ Provides users with optimum mechanism to match their priorities to the available resources
+ Encourages scalable codes & workflows
+ Ready access to quality facilities becomes part of CSIRO differentiation & advantage (project governance, employment and usage agreements prevent misuse/waste that might occur for similar policy in academic environment)
+ Optimum environment for attracting young talent
+ Prevents "double hurdles", e.g. external grants, access to NCRIS via merit allocation

– Requires explicit new funding commitment to grow
– Difficult to directly link investment levels to science needs
– Requires additional policy for allocating resources if demand exceeds supply
– Limited incentive to waste resource, e.g. optimised code allows faster results but doesn't save user any money

The main arguments for and against attributed funding are:

+ Clear link between investment levels and science needs
+ Clear link between investment levels and resource allocation
+ Creates new funding pipelines that bypass enterprise cost reduction pressures (e.g. IM&T)
+ Consistent treatment of internal & externally funded activity

+ Provides greater sense of ownership, control and predictability for users
– Subject to funding discontinuities and anomalies
– Project and science investment timelines generally not commensurate with major ASC acquisitions & partnerships
– Disincentive for users to experiment
– Inflexible for new users and projects, changes in direction etc
– Creates barrier to entry for new users
– No incentive for best-for-CSIRO behaviour e.g. resource sharing
– Creates division and hence diminution of shared resources

Resource Allocation Options

Allocation of resources can be made using three basic choices or philosophies:

1. All you can afford
2. All you are allowed
3. All you can use that is available – i.e. no explicit limit

Option 1 is commensurate with explicit attribution of funds to research projects. If the available resources are smaller than the notional investment then allocations and attributions can be scaled accordingly.

Option 2 is widely used in HPC facilities worldwide where these are accessed by users from different institutions, i.e. on a competitive basis. The APAC-NF Merit Allocation Scheme is a typical example. It is typically based on peer (science) assessment, which is an inhibitor for new research activities without established track records.

For shared systems, Option 3 is typically imposed using a scheduler that prioritises jobs based on past usage or other criteria. In principle, this option is most efficient, as it allows maximum flexibility for changes in project directions and priorities, and minimises division of the resources. Compared to the other options, users are not restricted to explicit allocations, nor do they feel obliged to consume excess allocation with low priority work to avoid “waste”.

Appendix 6: Utilisation of Advanced Scientific Computing in CSIRO

Current Computational Science

The information below describes some of the current computational research activity across CSIRO. It is derived largely from Divisional presentations made at the Advanced Scientific Computing Workshop held in August 2007.

Land & Water

WRON

The Water Resources Observation Network (WRON) is CLW's major ASC activity. The WRON systems are mainly used for runoff models, hydrological models, remote sensing temporal processing and uncertainty analysis. CLW are currently building shoulder clusters in other capital cities probably Perth and Adelaide.

Minerals

CFD Group

- Long-standing CFD modelling expertise in minerals processing, chemical processing, power generation, petroleum, food.
- Industrial consulting, model development, strategic / collaborative research

Molecular Simulation

- Predicting material properties and structure.

Online Analysis and Control (OLAC)

- Monte Carlo simulation used by OLAC group to model, design and optimise nuclear instruments. Detailed tracking of neutrons, g-rays, X-rays and electrons through 3D instrument models

Mathematical & Information Sciences

Computational Modelling Group

- Develop and apply particle-based methods (DEM and SPH) for modelling of fluids, solids, particles and bubbles in industrial applications.

Bioinformatics Group

- Algorithmic Analyses
- Genomic Analyses

Mathematics for Mapping & Monitoring Theme

- AGO Land Cover Change Project (operational)
    - Two key algorithms applied to imagery covering entire Australian continent: Spatial-temporal models, Markov Random Fields

- Urban Monitor (prototyping)

Petroleum Resources

Molecular dynamics: e.g. methane hydrates

Seismic processing: stochastic inversion (highly parallel), forward modelling (FD)

Finite difference solutions to partial differential equations: flow in porous media, oil/gas/CO_2, coal

Coupled processes: flow + fracturing, geochemistry, or geomechanics; stratigraphic forward modelling

Plant Industry, Livestock Industries, Entomology

Bioinformatics

- Gene Annotation, Classification, Prediction, Expression
- Microarray Analysis & Data Mining
- Comparative Genomics
- Transcriptomics
- Genotyping
- QTL analysis
- Proteomics
- Phylogenomics

Food Science Australia

Bioinformatics

Sensing, automation, real time process monitoring and control, engineering prototyping

Modelling engineering applications and visualisation

- Modelling food processes and unit operations (Dryers, extruders, baking ovens, High Pressure Processing Vessels)
- Computational Fluid Dynamics (CFD)
- Modelling refrigeration, storage and transport systems

Modelling chromatographic separations

- Image analysis / Cytometry

Microbiology

- Stochastic modelling / multivariate analysis in thermal sterilisation

Data analysis in epidemiological and clinical studies

Experimental design and data analysis from laboratory instruments

- Texture analysis

- Rheology

Exploration & Mining

Predictive Minerals Discovery Theme

- Models the processes that give rise to ore deposition. This will provide a more effective way to identify prospective geologic environments worth testing by drilling.

Sustainable Mining Systems Theme

- Developed software to simulate the mechanical changes that occur in a domain consisting of rock (deformation, stress changes, fracture) and the flow of water and (possibly) methane that occurs in the rock following excavation due to mining.

Marine & Atmospheric Research

BLUElink

- The BLUElink initiative centres on ocean prediction and analysis, and forecasting of day-to-day variations in ocean currents, ocean eddies and temperatures. The forecasts provide information on coastal and ocean currents and eddies, surface and subsurface ocean properties, that impact and are linked to maritime and commercial operations, defence applications, safety-at-sea, ecological sustainability, regional and global climate.

ACCESS

- The Australian Community Climate Earth-System Simulator (ACCESS) is a coupled climate and earth system simulator being developed as a joint initiative of the Bureau of Meteorology and CSIRO in cooperation with the university community in Australia. For CSIRO, it provides the best possible science for use in analysing climate impacts and adaptation, and related fields.

ICT Centre

Contribution here is through the e-Research Theme, whose goal is "*by 2010 develop an integrated e-Research platform that fundamentally changes the paradigm for dispersed teams undertaking research over a distance.*" This theme has the following project portfolio.

- National Collaboration Network – create and maintain the NCN as a national facility.
- Trusted Services – consists of two major components.
    1. the braccetto collaborative platform being developed as a part of the HxI initative between CSIRO ICT Centre, NICTA and DSTO.
    2. applications, tools and services that provide user brokered dynamic, secure resource allocation and deployment on managed infrastructure that facilitate and ensure trusted collaborative environments.
- Self tuning end-to-end Quality of Service – experimentally verify a new resource management protocol applicable to Internet core routers.

Australia Telescope National Facility:

ASKAP

- The Australian Square Kilometre Array Pathfinder is an international collaboration with Australia (CSIRO), Canada (NRC), the Netherlands and South Africa to build an array of dishes capable of high dynamic range imaging and using wide-field-of-view phased array feeds. It will be built in the Mid West region of Western Australia at the Australian SKA candidate site.

Molecular & Health Technologies

Computer aided drug discovery - MolSAR

High-performance computing for screening chemical compounds – structural biology

Quantum chemistry simulation for modelling neurodegenerative diseases

Materials modelling and informatics

Material Science & Engineering

Dynamics of engineering and thermofluids, microfluidics

High speed instrumentation, 3D tomography reconstruction

Theory and modelling of advanced materials, multiscale modelling of corrosion

Energy Technology

Ionic liquids for energy storage technologies

Renewable energy materials

Utilisation of HPSC, iVEC and Condor Facilities 2007

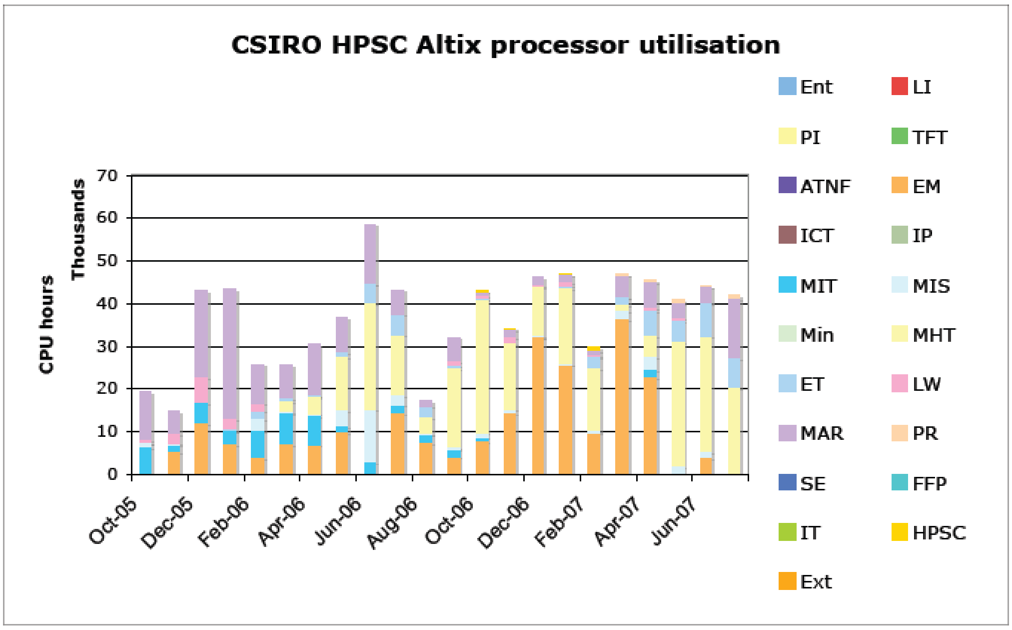

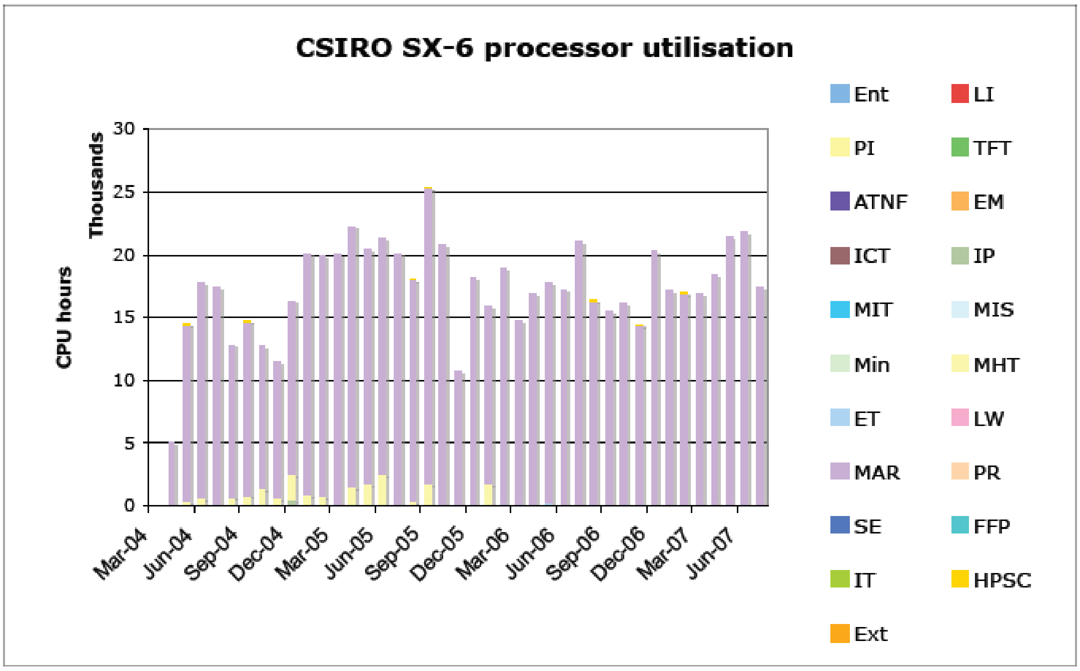
CSIRO SX-6 processor utilisation
CPU hours
Thousands
30
25
20
15
10
5
0
Mar-04
Jun-04
Sep-04
Dec-04
Mar-05
Jun-05
Sep-05
Dec-05
Mar-06
Jun-06
Sep-06
Dec-06
Mar-07
Jun-07
Ent
LI
PI
TFT
ATNF
EM
ICT
IP
MIT
MIS
Min
MHT
ET
LW
MAR
PR
SE
FFP
IT
HPSC
Ext

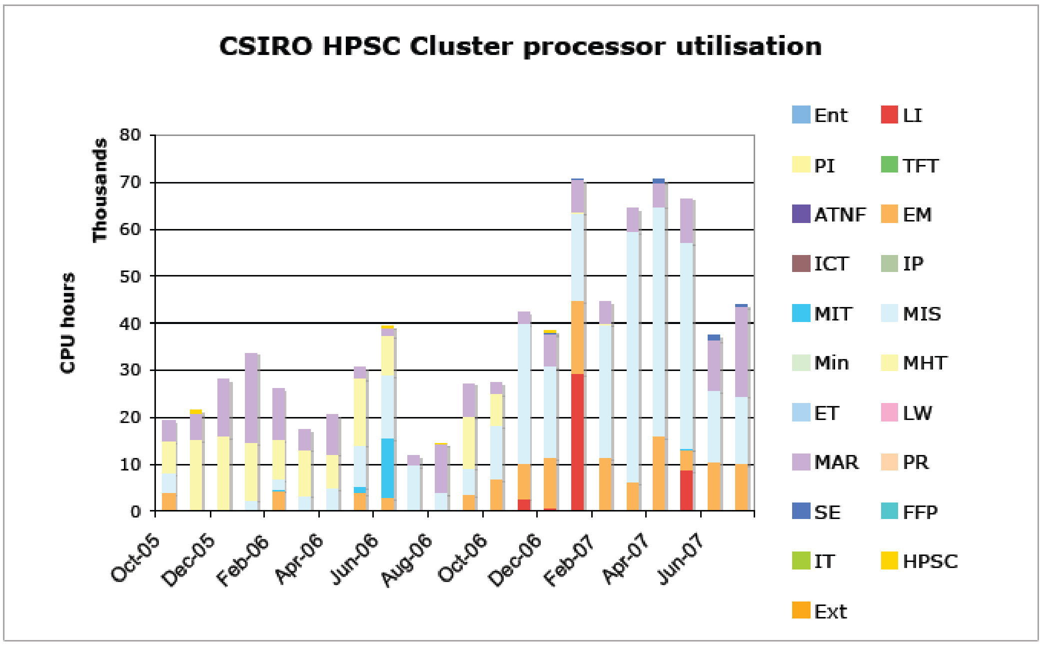
CSIRO HPSC Cluster processor utilisation
CPU hours
Thousands
80
70
60
50
40
30
20
10
0
Oct-05
Dec-05
Feb-06
Apr-06
Jun-06
Aug-06
Oct-06
Dec-06
Feb-07
Apr-07
Jun-07
Ent
LI
PI
TFT
ATNF
EM
ICT
IP
MIT
MIS
Min
MHT
ET
LW
MAR
PR
SE
FFP
IT
HPSC
Ext

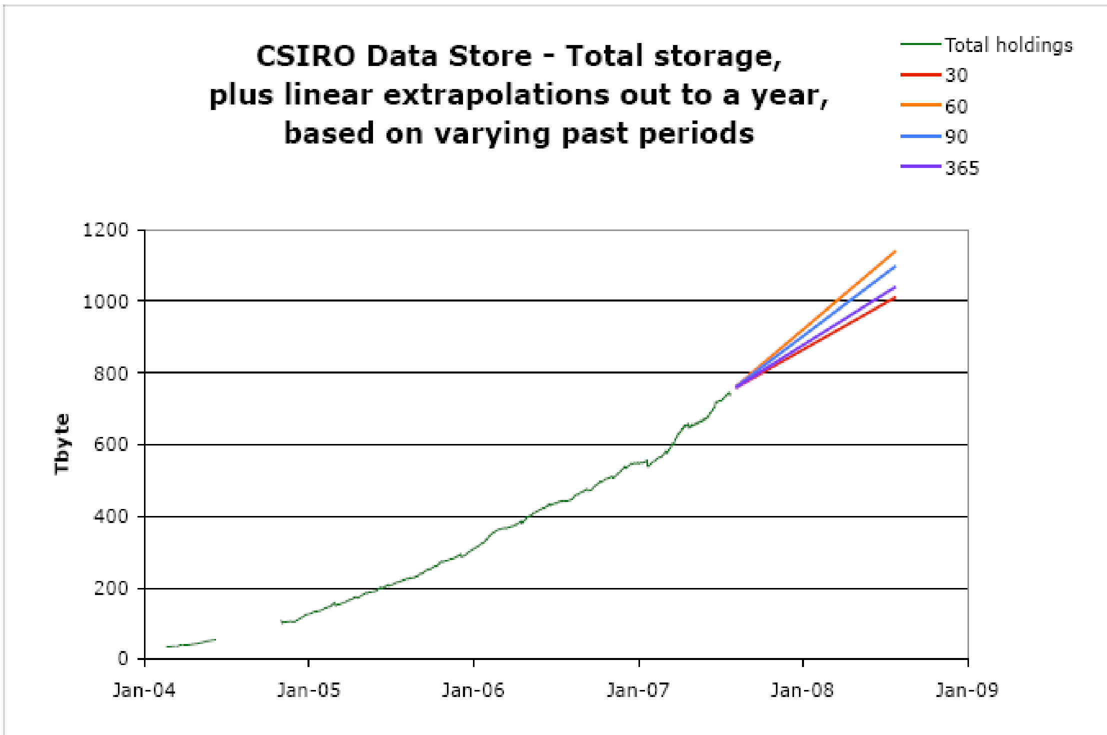

CSIRO Data Store - Total storage, plus linear extrapolations out to a year, based on varying past periods
Total holdings
30
60
90
365
Tbyte
1200
1000
800
600
400
200
0
Jan-04
Jan-05
Jan-06
Jan-07
Jan-08
Jan-09

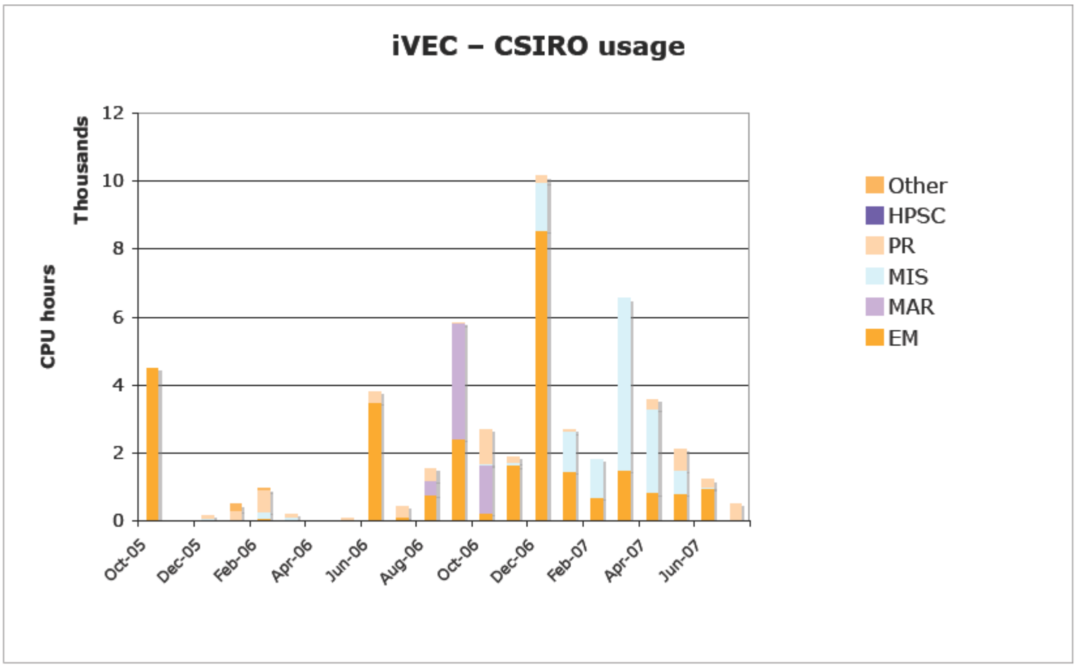

iVEC – CSIRO usage
CPU hours
Thousands
12
10
8
6
4
2
0
Oct-05
Dec-05
Feb-06
Apr-06
Jun-06
Aug-06
Oct-06
Dec-06
Feb-07
Apr-07
Jun-07
Other
HPSC
PR
MIS
MAR
EM

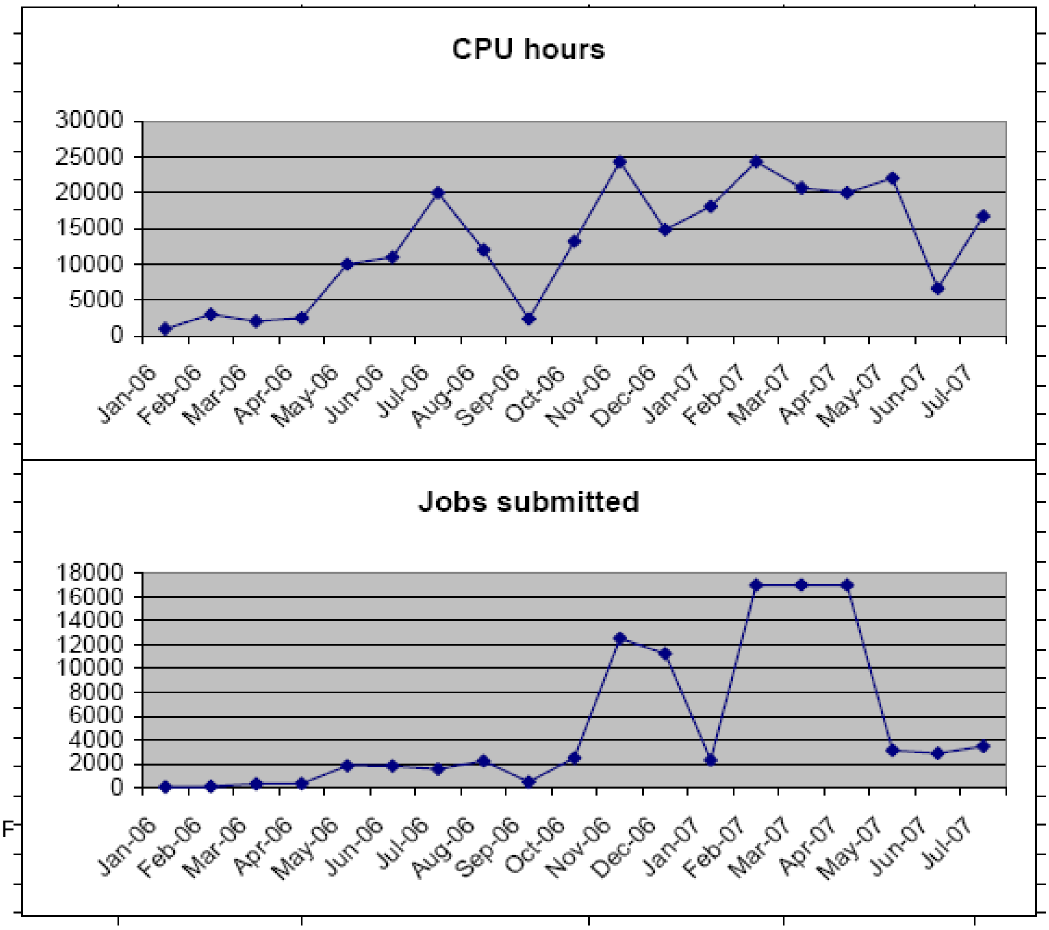
CPU hours
30000
25000
20000
15000
10000
5000
0
Jan-06 Feb-06 Mar-06 Apr-06 May-06 Jun-06 Jul-06 Aug-06 Sep-06 Oct-06 Nov-06 Dec-06 Jan-07 Feb-07 Mar-07 Apr-07 May-07 Jun-07 Jul-07
Jobs submitted
18000
16000
14000
12000
10000
8000
6000
4000
2000
0
Jan-06 Feb-06 Mar-06 Apr-06 May-06 Jun-06 Jul-06 Aug-06 Sep-06 Oct-06 Nov-06 Dec-06 Jan-07 Feb-07 Mar-07 Apr-07 May-07 Jun-07 Jul-07
F

Appendix 7: Survey Results

CSIRO Business Unit submissions

The Advanced Scientific Computing review sought submissions from CSIRO staff in the research, IM&T and administrative domains. David Toll (Director, Information Services and Property) emailed Chiefs and Flagship Directors on 07 June 2007 requesting a nominated contact. These contacts were subsequently emailed an issues paper to stimulate discussion, and a number returned aggregated or individual comments.

These comments have been aggregated and grouped by topic, and are located at https://teams.csiro.au/sites/ascs/Report/Links/ASC%20Collated%20Submissions.pdf .

CSIRO Advanced Scientific Computing Survey

On 23 July 2007 a survey was sent to all nominated Business Unit (BU) representatives. The survey was designed to gather information on current and potential computational science usage, staffing, skills and governance. The survey is reproduced below.

PART ONE – CONTACT DETAILS

PART TWO – GOVERNANCE/PLANNING/FUNDING

Q. 01. Does your Business Unit have a planning or oversight process, to match advanced scientific computing resources to your science needs?

Q 02. If CSIRO were to allocate significant funding ($M's) to increasing our shared capability in Computational Science, where in your opinion would it best be used? (e.g. clusters; vector computers; buying computational services from external providers like APAC, BoM, Amazon; more skilled people?)

PART THREE – SCIENTIFIC COLLABORATION

Q 03. In which of the following ways do you collaborate with researchers from other organizations on Computational Science activities?

- ☐ None
- ☐ Sharing data
- ☐ Sharing code and modules
- ☐ Use external services – e.g. deriving computational results
- ☐ Jointly developing shared systems
- ☐ Other – please provide details below

Q 04. Are you likely to have a requirement for very large scale or geographically distributed "grid" computational resources now or in the near future?
If so, please provide details.

Q 05. Who do you consider to be your peer organisations in Australia or beyond?
(This will provide opportunities to compare resourcing levels for computational activity)

Q.06. What factors would increase your Business Unit's ability to collaborate with other organisations in Computational Science?

PART FOUR – ADVANCED SCIENTIFIC COMPUTING SOFTWARE

Q 07. Scientific Applications are available from a number of sources both internal and external to the organisation.
Please list the major scientific application(s) used within your Business Unit, under the appropriate heading for different types of licenses.

Source	Major Application
CSIRO developed	
Public Domain / Free	
Local License	
Shared License	

Q 08. Who currently supports your Business Unit with assistance in running your locally developed software?

☐ Self
☐ IM&T
☐ Other researchers
Who?____
☐ Vendor
Who?____

Q 09. Would a CSIRO-wide collaborative software development environment (such as Subversion or SourceForge) be of benefit to your Business Unit?
☐ Yes ☐ No
Comments____

Q 10. For software developed by your business unit are you willing to share your applications?

☐ No, this is core Intellectual Property for my research business unit
☐ Yes, but only with CSIRO Researchers
☐ Yes, But only with other Australian Researchers
☐ Yes, Externally as a public domain software tool
☐ Yes, Externally as a Commercially available software tool

Are there any additional Software comments you would like to make?

PART FIVE – SKILLS AND STAFFING
Please note: The CSIRO High Performance Scientific Computing Group (HPSC) are now integrated into CSIRO IM&T

Q 11. How many researchers are currently working in your Business Unit?

Q 12. How many of these researchers currently access Advanced Scientific Computing services at some level?

Q 13 Do you think demand for Advanced Scientific Computing services will grow in your Business Unit?

☐ Yes ☐ No

Comments____

Q.14. What additional resources would help you and your people in getting better computational science outcomes from Advanced Scientific Computing capabilities and infrastructures?

• Awareness-raising workshop	☐ None	☐ Minor	☐ Significant
• Training Workshops	☐ None	☐ Minor	☐ Significant
• Training one-on-one	☐ None	☐ Minor	☐ Significant
• Manuals & self-help	☐ None	☐ Minor	☐ Significant
• CSIRO User group/discussion board	☐ None	☐ Minor	☐ Significant
• Access to computational science specialists	☐ None	☐ Minor	☐ Significant

Q.15. What resources are useful to you in keeping abreast of general developments in Computational Science?

• Publications	☐ None	☐ Minor	☐ Significant
• Conferences	☐ None	☐ Minor	☐ Significant
• Institutional Websites	☐ None	☐ Minor	☐ Significant
• Vendor seminars/websites	☐ None	☐ Minor	☐ Significant
• Training/workshop	☐ None	☐ Minor	☐ Significant

Q 16. How is your Advanced Scientific Computing environment currently being supported?

☐ Supported within the Business Unit or research team, How many Hours/Week? ________(hr/wk)

☐ Supported by another CSIRO Business Unit Who?______________________________(name)

☐ External Service Provider - Who?____(name)

☐ CSIRO IM&T staff - Which team?____ (eg BST/TST, eScience**)**

Q 17. Who provides the (non-research) skills required to support your computational science?

	Self Support	IM&T	External (Who)

- **Procurement** ☐ ☐ ____________
- **System management** ☐ ☐ ____________
- **Hardware support** ☐ ☐ ____________
- **Software support** ☐ ☐ ____________
- **Environmental support** ☐ ☐ ____________
- **Backup and recovery** ☐ ☐ ____________
- **Data management** ☐ ☐ ____________

Q 18. Are you happy with the current support arrangements?

☐ Yes ☐ No
Comments / Suggestions for improvements

PART SIX – INFRASTRUCTURE - SCIENTIFIC COMPUTE/STORAGE USAGE AND GROWTH

Q 19. Please estimate both the computational resources and the annual growth used by your Business Unit (Division, Flagship or Corporate Group)
(We understand that future demand depends on many factors - rough estimates are fine)

Scientific Computing System Type	Service Provider (Name)	Annual usage (CPU-hours)	Annual CPU growth % or +	Storage Usage (terabytes)	Annual storage growth % or +	Science Disciplines (I, P, C, B, M, G, O)	Comments
Desktop Cycle harvesting (e.g. Condor)							
High-end local workstation or server							
Windows PC cluster							
Linux PC cluster							
Advanced cluster (e.g. HPCCC Altix)							
Vector supercomputer (e.g. HPCCC NEC)							
Other							

Q 20. Are any of the following infrastructure items a limiting constraint for computational science activity in your Business Unit?

Computational infrastructure ☐ None ☐ Minor ☐ Significant ☐ Prohibitive
Storage ☐ None ☐ Minor ☐ Significant ☐ Prohibitive

Network	☐ None	☐ Minor	☐ Significant	☐ Prohibitive
Collaboration tools	☐ None	☐ Minor	☐ Significant	☐ Prohibitive
Software Tools	☐ None	☐ Minor	☐ Significant	☐ Prohibitive

If significant/prohibitive, please comment - i.e. what is the constraint?

Outline of Survey Results

In total, 31 responses were received from 13 Business Units. While the original aim was to have each BU to provide a single unified response, several BU's returned multiple responses from individual groups, streams or themes. The results of the survey were collated and presented at the ASC Workshop on 15 August.

A summary of the aggregated survey responses is provided below. Specific information about the original survey content and responses is available on request.

Q1: ASC planning or oversight processes tend to be patchy across Business Units and are mostly at the project level.

Q2: Preferred ASC areas for increased investment:

- more Linux cluster resources
- computationally trained staff
- new Windows cluster(s)

Q3: Majority of ASC collaboration with external researchers is based on sharing data and source code.

Q4: About 50% of respondents predicted a need for grid computing resources.

Q5: A wide range of peer groups was identified.

Q6: Factors which would increase Business Unit's ability to collaborate on ASC with external organisations:

- Faster networks
- Large hierarchical storage
- Collaborative data environment
- MoU's in place for data sharing
- Visitor funding
- Access to more skilled staff
- Guidelines for software IP issues

Q7: Most common source of software for major ASC applications is in-house and public domain.

Q8: Support in running ASC software is mostly self-support and from other scientists.

Q9: Majority of respondents have requirement for a collaborative software development environment?

Q10: Willingness to share in-house developed applications varied widely and was more project-dependent than BU-dependent.

Q11-12: Proportion of EFT staff actively using ASC varied from 3% (FSA) to over 50% (CMAR)

Q13: 100% of respondents predicted an increase in demand for ASC services.

Q14: Road show and training workshops were most frequently identified as resources that would give better outcomes for ASC,

Q15: Principle means for keeping up to date with ASC developments were publications, followed by conferences.

Q16-17: ASC environment support is mostly internal, followed by IM&T.

Q18: The majority of respondents were you satisfied with current support arrangements.

Q19: Predicted CPU and storage requirements identified by survey respondents are plotted below:

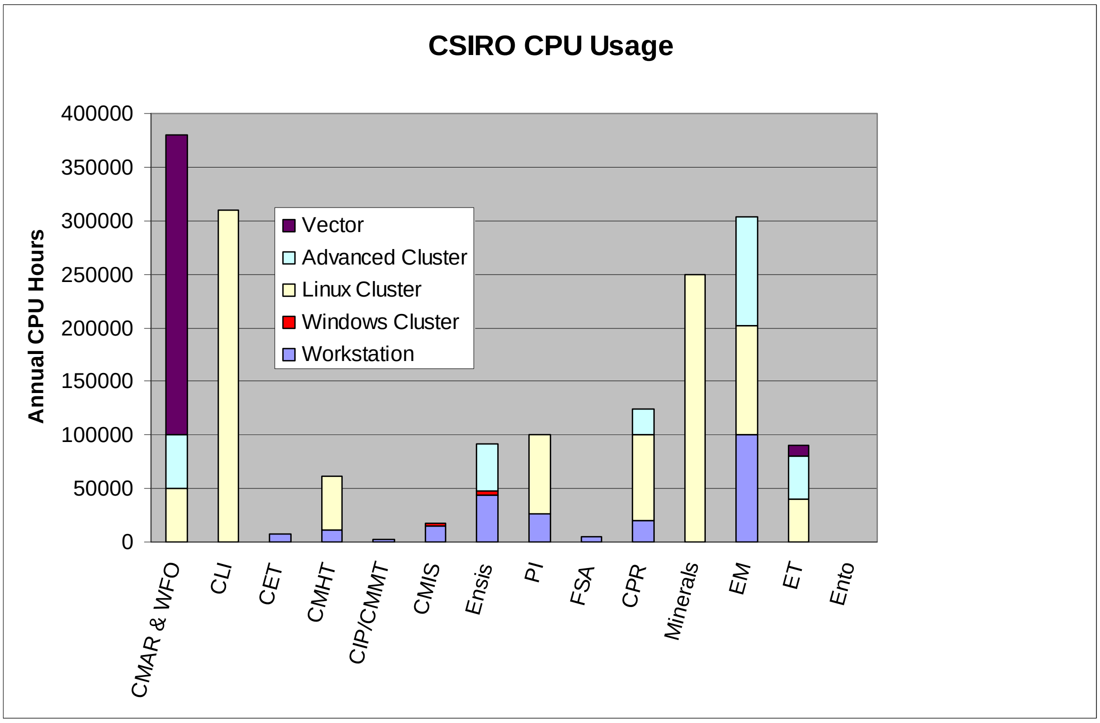

Figure 14. Predicted balance of CPU requirements.

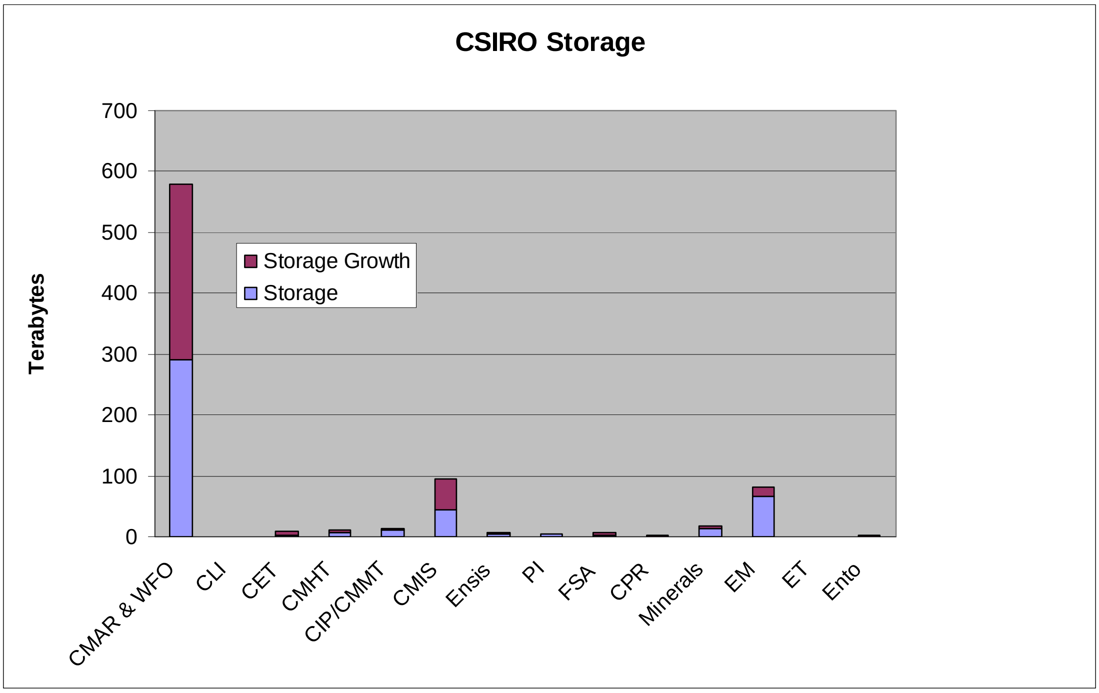

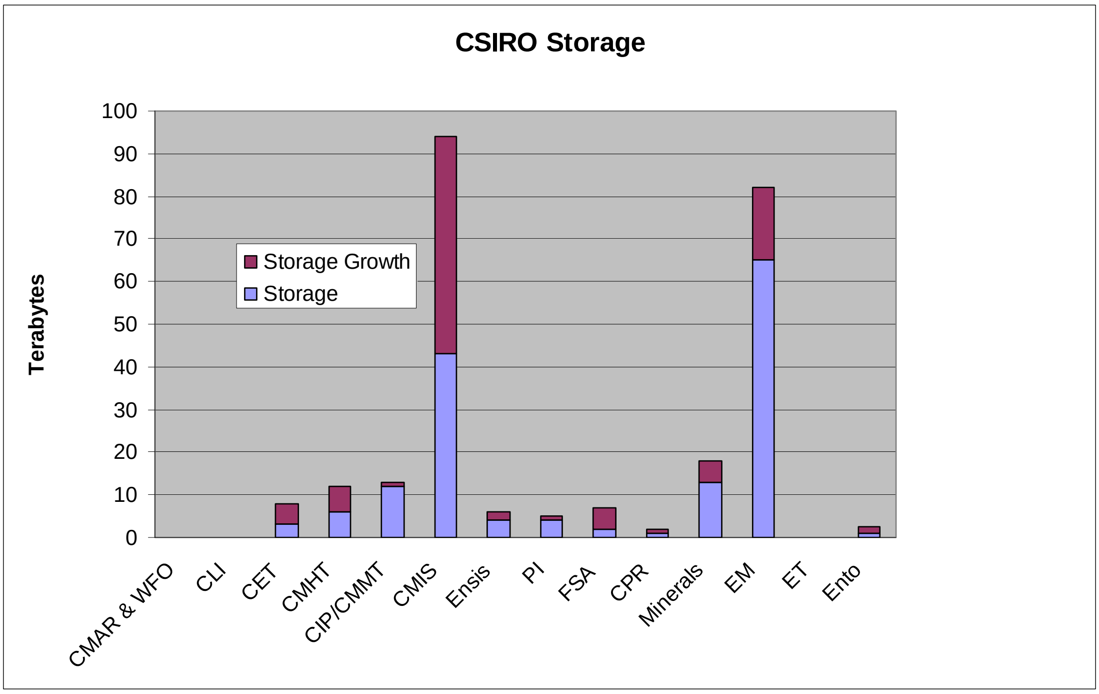

Figure 15. Predicted balance of storage needs and annual growth, with & without CMAR.

Q20: The following were identified as limiting constraints for ASC adoption:

- Computational infrastructure: More CPUs, Windows HPC, Cycle harvesting
- Data storage capacity, backup, archive, accessibility
- Faster networking
- Collaboration tools

- Increased awareness of what software tools are available
- Lack of qualified staff
- Lack of funding for local cluster maintenance

Notes on Survey Responses

Two Business Units with substantial ASC requirements – ATNF and CLW – were not able to respond to the ASC survey.

The major ASC activity in CLW is the Water Resources Observation Network (WRON) project in WfHC. The project operates a cluster with 40 dual core Xeon CPUs, with 50 TB storage and anticipated growth this year of 100 TB. CLW are establishing secondary clusters, most likely in Perth and Adelaide.

ATNF operates significant server and cluster facilities, e.g. associated with the Australia Telescope and Pulsar groups. The major ASC demand for ATNF is the ASKAP project, which will require a 20 TFlop/s system and 60 Petabytes storage per annum from 2009.

In view of the difficulty of capturing full quantitative requirements for all groups in CSIRO, the survey's prime purpose is qualitative assessment of balance of priorities.

Appendix 8: Workshop Summary

An Advanced Scientific Computing Workshop, focusing on CSIRO's current and anticipated future use of computational science, was held in Clayton on 14-15 August 2007.

The workshop provided a forum for some of CSIRO's computational scientists to network with peers, learn about related activities in other locations, and debate the options for expanding scientific computing across CSIRO. The workshop focused on the science application areas and their IT requirements rather than being specifically IT focused.

Following Michael Barber's welcome and overview of simulation science, a series of CSIRO researchers and external experts provided excellent background presentations on a range of topics, including:

- DEST's plans for funding computational infrastructure via the NCRIS PfC initiative
- Strategic and operational insights into the Partnerships for Advanced Computing in Victoria (VPAC), Western Australia (iVEC) and Nationally (APAC-NF)
- Cross-business unit perspectives on advanced scientific computing issues, including modelling, simulation, visualisation, data management, collaboration, infrastructure and support.

A number of these presentations highlighted that rapid growth in the science will soon test the current capacities of CSIRO's computational infrastructure and expertise.

Notes recorded from the presentations and subsequent discussion forum are available from

https://teams.csiro.au/sites/ascs/Report/Links/ASC%20Workshop%20Mindmap.pdf .

The presentations

Tuesday 14th August a.m.

Michael Barber
CSIRO Executive

Keywords: Simulation science; Theory, experiment, AND compute
Someone in CSIRO should at least NEED to use supercomputer from TOP50 even if we don't own one .
CSIRO and Australia do HAVE the scale of problem requiring this?

Bill Appelbe
VPAC

Keywords: VPAC, 100% HPC Service Provider
VPAC has 50-60 people supporting scientists.
HPC investment at $1M pa.
HAVE TO UNDERSTAND PEOPLES' SCIENCE to be effective
Build more general software, rather than just point, for better re-use.

Rob Woodcock
CSIRO Exploration and Mining

Keywords: Understand Uncertainty, Collaboration Ecosystem
Computation=> reduces uncertainty and $s of science
Was 2 models/wk. now workflow => 500 models/wk
Having real impact on exploration industry
Data integration is extreme cost
Poor national standards (e.g. workflow) =obstacles

Rob Bell
CSIRO HPSC

Keywords: CSIRONET&CSIRAC, Data centric model
Computer Cost recovery – you cannot afford to experiment because of cost of making mistakes?
Commodity computing speed has plateau-ed..
Go parallel to go faster – bottlenecks w. data management
Have to collaborate – open networks

The presentations

Tuesday 14th August p.m.

Rhys Francis
CSIRO, AeRIC (DEST)

Keywords: NCRIS, eResearch Challenges, Vital for Australia
Data is central – but no national data community with consensus for future – data-data-data-data – it's all about!!!
Challenge: how to make eScience available for non-IT experts.
AREN => desktop to resource, not just point to point backbone.
Compute – at least $10m/yr – ANU will provide national facility

Ben Evans
ANU

Keywords: APAC National Facility
Peak facility, Alloc^{n}. – Merit (56%), Share and Start-up
NF provides outreach and training services.
DATA MANAGEMENT is MAJOR ISSUE
Better scheduling/suspend/resume eliminates idle time

Chris Mitchell
CSIRO MAR

Keywords: Climate and Weather Xtremes – Global Collaboration
Terrestrial biosphere, Ozone depletion, ice sheets, rainfall – extreme complexity => "**horrendograms**"
4 fold increase in spatial resolution = 100 fold increase in computation requirements.
Global collaboration for eResearch mandated!

Andreas Schiller
Wealth from Oceans

Keywords: 6 themes, Responses to ASC questionnaire. .
Need access across the HPC spectrum – vector => cluster – need HPCCC and APAC collaborations – growth is spread across all sectors, across compute, data and network. IM&T Service Delivery through close collaboration with the science.

The presentations

Tuesday 14th August p.m.

John Taylor
CSIRO MIS

Keywords: Modelling, Bioinformatics
Science constrained by compute power available.
Computational Modelling, Bioinformatics, Urban Monitor
Storage issues, Compute issues, network issues. 100% growth pa.
Many applications scale to parallelism. **Shared data important**
Experience shows: high-value outcomes come from CSIRO developed software.

Michael Barber
for
Alex Zelinsky
CSIRO ICTC

Keywords:ICST Group
ATNF, CMIS and ICTC – all data intensive,
500 scientists, 40 themes, multi-disciplinary
Tensions between delivering to Flagships and deadlines, and developing capability
eResearch rather than eScience – terminology issue
Services are 80%+ of Australian economy, but is <20% of CSIRO involvement – is this a problem?

The presentations

Wednesday 15th August a.m.

Andrew Rohl
iVEC

<u>Keywords: Partnership CSIRO, People cost, People save</u>
$3Mpa = $1.25M infrastructure + **$1.75M people/support**
Cray; SGI Altix; new Cluster – Visualisation Systems.
New Petabyte Datastore. <u>Building/running clusters is expensive</u>
Expert Help always essential. SKA and other astronomy driving growth.

Tracey Hind
CSIRO IM&T

<u>Keywords: eScience Information Management.</u>
Curate and preserve eScience data.
Approved by Executive Team. Parallel with NCRIS' ANDS
Goal 2010 – CSIRO information management environment
Goal for eSI– Find, Share, overlay, analyse, publish, curate it.

Peter Witt
CSIRO Minerals

<u>Keywords: Three groups + Light Metals F/Ship, S/w cost issues</u>
Three groups CFD, Molec Sim, Online Analysis (OLAC).
OLAC - Monte Carlo app. 700k CPU hrs last year via CONDOR.
Molec Sim APAC/VPAC 200k CPU hrs pa -Telepresence microscopy.
Comp Fluid Dynamics – World Leading research
Lots of different clusters used. Planned growth will be a challenge.
Need more machines and cheaper software!

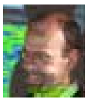

Jonathan Ennis-King
CSIRO Petroleum Res.

<u>Keywords: Less Data = more work. LOBOS</u>
135 people – 15-25 doing heavy computation.
Complex simulations. Open source software.
Run on workstations having tried APAC etc. CSIRO Webpage issues
Central resources can cause lag. Local cluster –less power, more control
Using APAC, iVEC, VPAC, HPSC and lots of local cluster work.
Pros and cons of local clusters.

The presentations

Wednesday 15th August a.m.

Gavin Kennedy
CSIRO Plant Ind.

<u>Keywords: Data Avalanche, Plant Genome sequencing etc</u>
Many many processes for gene/DNA analysis and sequencing
The Challenge – growing collection, annotations, reference databases
CSIRO Bioinformatics Cluster => 66 DELL blades - >300k CPU hrs/yr
NOW => Require Massively //L Sequencing – Workflow (e.g.Taverna) –
Data collaboration – Data Management Model. **NCRIS 5.1 support please!**

Sean McWilliam
CSIRO Livestock Ind.

<u>Keywords: Animal Genome Sequencing, kilocluster needed</u>
All s/w tools used are open source.
e.g. Sheep genome sequ. 2 months processing time! This is a small project – can't afford to make mistakes.
Future, wish list => massive growth in computational, information management and skills capability to remain relevant.

Peter Fitch
CSIRO Land & Water

<u>Keywords: Stakes are high, Data Management and skilled people!</u>
Land & Water usage. Collaborate wide range stakeholders.
Required to manage uncertainties more than before.
Connectedness of eco subsystems means coupled models needed.
Storage storage storage! 0.5 FTE on data management => not enough.
Need expert support to get going. Hosted apps.would be good e.g. Matlab
WRON – a critical project => 60TB store, dell cluster => 100TB needed
Cost is not barrier – access to skilled people is barrier

Jay Sellahewa
CSIRO Food Science Australia

<u>Keywords: Complex Modelling, 3D visualisations, HPC needed</u>
Food supply chain beyond primary production
Some bioinformatics; process control/engineering etc, epidemiology.
Open source and locally developed s/w., CAD, Visualisation some lic. s/w
Desktop modelling. 3D visualisation with help from CSIRO Minerals.
Will need expert support in the future w compute and data modelling.
Saving money and time by computer simulation - very powerful!

The presentations

Wednesday 15th August a.m.

Karl Gordon
CSIRO Entomology

Keywords: Genome sequ., TBytes/project
The of metazoan genomics - >12M species
Honey Bee - international consortium.
Massively //^{L} sequencing.
Eg. ABI's SOLiD system => up to 500Mb per day
Emerging=> Resequencing, Ecogenomics,
Metagenomics = > TB's per project.

Marcus Foster
CSIRO IM&T

Justin Baker
CSIRO IM&T

Keywords: Aggregated results from Survey
and

Brainstorm/Discussion

Alf Uhlherr
CSIRO MHT

WORKSHOP HOST

Appendix 9: External Organisational Benchmarks

Fraunhofer-Gesellschaft

The Fraunhofer-Gesellschaft (FhG) is a society of 56 institutes who perform applied research in the engineering sciences. Their annual budget of $2.0B (€1.2B) consists of approximately 40% core funding, 25% external public sector and 35% private sector contracts. They have 12500 staff working at ~40 locations.

The operations of the FhG institutes are relatively autonomous, which makes it difficult to assess the degree to which these different institutes utilise high performance computing.

ICT research and computational science activity are primarily concentrated in the Information and Communication Technology Group. The ICT Group is the largest European research alliance for information and communication technology. The competences of the 17 member institutes are pooled in strategic networks and marketed in common. This network allows for specific, branch-typical and entire solutions from applied research: especially tailored IT solutions, competent technology consulting, as well as preliminary research for new products and services. Periodic economic summits provide a platform for the appropriate partners from industry and research. With 3,000 employees from 17 institutes, and an annual budget of more than 190 million Euros, the ICT Group is the largest research alliance in Europe.

Some of the major institutes in the ICT group include:

- Industrial Mathematics (ITWM) – 150 staff
  The ITWM has an annual budget of $17.4M (€10.6M) and conducts computational and mathematical research and consulting for industry, primarily in automobile construction, engineering, textiles, microelectronics and computing. ITWM currently operates the 1096 processor "Hercules" Dell PowerEdge cluster. With a peak performance of 10.2 teraflop/s, this is the most powerful system within the FhG and the ninth largest in Germany. The system is used by several Fraunhofer sites to develop and run applications in the areas of multi-physics simulation, molecular dynamics, fluid dynamics, microstructure analysis and seismic imaging. Commercial simulation applications like Abaqus, Nastran, Permas, Fluent, MAGMAsoft and others are available on the machine. The system serves as a general-purpose computer that allows users to utilise individual nodes, run mainstream parallel MPI applications, or use the Fraunhofer Virtual Machine (FVM)
- Algorithms and Scientific Computing (SCAI) – 125 staff
  The SCAI engages in computer simulations in product and process development and has strong industry partnerships. SCAI staff design and optimise industrial applications on HPC systems. The aim is to reduce development times, make experiments less expensive and optimise technical products. SCAI Bioinformatics focuses on solutions for information extraction
- Applied Information Technology (FIT) – 110 staff
  FIT investigates human-centred computing in a process context. The usability and usefulness of information and cooperation systems is optimized in their interplay with human work practice, organisation and process

- Computer Architecture and Software Technology (FIRST)
  FIRST focuses on developing information technologies for intelligent data analysis, embedded and safety-relevant systems and innovative human-computer interaction technologies.

In addition to the ICT Group, the FhG also has two separate "alliances" – groups of FhG institutes - dedicated to Grid Computing and a Numerical Simulations of Products and Processes.

The Ernst Mach Institute (EMI) - with 235 staff and an operating budget of $21.7M (€13.1M) in 2006 - is also a major user of computational science. The EMI deals with physical-technical aspects of high-speed, mechanical, and fluid-dynamic processes. These include experimental and numerical analyses of shock waves in solids, fluids, and gases, flow and combustion processes, impact and penetration processes, the response of structures to shock loads, dynamic material response at high strain and at high strain rates. The EMI have operated a 256-processor IBM cluster with a peak speed of one teraflop/s, for their high-speed mechanics and fluid dynamics calculations.

Battelle

The Battelle Memorial Institute is a global science and technology enterprise that consists of 19000 staff and has an annual R&D budget of $4.4B (US$3.8B). Battelle's major scientific computing facility is at Oak Ridge National Laboratory, which it manages under contract for the U.S. Department of Energy (DoE). This facility is one of the largest in the world. Battelle also manages or co-manages the National Laboratories at Pacific Northwest (PNL), Brookhaven (BNL), Idaho (INL) and the National Renewable Energy Laboratory (NREL), which also operate substantial HPC facilities.

ORNL

Oak Ridge National Laboratory is the DoE's largest science and energy laboratory, and was established in 1943 as a part of the Manhattan Project to pioneer a method for producing and separating plutonium. During the 1950s and 1960s, ORNL became an international centre for the study of nuclear energy and related research in the physical and life sciences. With the creation of DOE in the 1970s, ORNL's mission broadened to include a variety of energy technologies and strategies.

Currently ORNL has a staff of more than 4,200 and annually hosts approximately 3,000 guest researchers. Annual funding exceeds $1.45B ($US1.2 billion). As an international leader in a range of scientific areas that support the DoE's mission, ORNL has six major mission roles: neutron science, energy, high-performance computing, systems biology, materials science at the nanoscale, and national security. ORNL's leadership role in the nation's energy future includes hosting the U.S. project office for the ITER international fusion experiment and the Office of Science sponsored Bioenergy Science Center.

ORNL's advanced scientific computing operations are governed by the Computing and Computational Sciences Directorate (CCSD). The CCSD govern state-of-the-art research and development in computer and computational sciences in support of DoE's missions and programs. They develop and deploy leading edge computing and information technology capabilities to keep computational sciences at a level comparable to experimental sciences in the pursuit of scientific discovery and technical innovation.

CCSD's main computing facility is the National Center for Computational Sciences (NCCS) at Oak Ridge. NCCS was established in 1992, and in 2004 was designated by the Secretary of Energy as the Leadership Computing Facility for the nation, to provide for unclassified research a resource 100 times more powerful than current capabilities. NCCS provide access to a range of HPC systems.

- **Jaguar:** currently ranked the second most powerful supercomputer in the world, this is a combined Cray XT3/XT4 system with a peak performance of 119 trillion floating point operations per second (119 teraflop/s). The system contains 11,708 processors, 11,508 of which are compute processors and the remainder service processors. Jaguar is to be upgraded to a peak performance of 250 teraflops in late 2007. The upgrade will replace the current dual-core processors with quad-core processors.
- **Phoenix:** a Cray X1, with 512 multi-streaming vector processors and 2 TB of globally addressable memory.
- **Ram:** a 256-processor SGI Altix with 2 TB of shared memory. Each processor is a 1.5 GHz Intel Itanium2.
- **Cheetah:** a 27-node IBM p690 system, where each node has thirty-two 1.3 GHz Power4 processors. Most of the nodes have 32 GB of memory, but five of the nodes have 64 GB of memory, and two have 128 GB of memory.
- **Eagle:** a 184-node IBM RS/6000 SP, 176 "thin" nodes have four 375 MHz Power3-II processors and 2GB of memory while 8 "wide" nodes have two 375 MHz Power3-II processors and 2GB of memory.

Oak Ridge also hosts the following computational science divisions:

- The Computational Sciences and Engineering Division (CSE), which develops and applies creative information technology and modelling and simulation research solutions for National Security and National Energy Infrastructure needs
- The Computer Science and Mathematics Division (CSM), which is ORNL's premier source of basic and applied research in high-performance computing, applied mathematics, and intelligent systems. Basic and applied research programs are focused on computational sciences, intelligent systems, and information technologies
- The Joint Institute of Computational Sciences (JICS), which is a partnership with the University of Tennessee. JICS's purpose is to advance scientific discovery and state-of-the-art engineering and to further knowledge of computational modelling and simulation, by utilisation of the NCCS system, and to educate a new generation of scientists and engineers well versed in the application of computational modelling and simulation for solving for the most challenging scientific and engineering problems.

TNO

TNO is a research institute with an annual budget of $940M (€578M) and 4300 staff based at 24 European locations, mostly in the Netherlands. TNO

- performs research for private companies, government bodies and public organisations
- provides contract research and specialist consultancy as well as grants, licenses for patents and specialist software
- tests and certifies products and services, and issues an independent evaluation of quality
- sets up new companies to market innovations.

TNO also performs military research for the Netherlands, which is the only country in the world where the responsibility of the Ministry of Defence.

Research is focused on 5 core areas:

- Quality of Life
- Defence, Security and Safety
- Science and Industry
- Built Environment and Geosciences
- Information and Communication Technology

TNO's high performance computing needs are met by its partnership with NCF.

NCF

The Netherlands National Computing Facility (NCF)[34] is the body that administers and funds high performance computing infrastructure for public scientific research in the Netherlands. Long term funding for the installation, running costs and future upgrades or replacements of the national supercomputer facilities is provided by NWO, which is the Netherlands Organisation for Scientific Research. NWO[35] is the principal body providing research funding for the Dutch Universities, and also operates a range of specialist research institutes.

NCF funds and administers access to computing resources via a peer review mechanism. The machines themselves are currently housed and operated by SARA[36]. SARA facilities are also accessible to commercial users via alternative non-NWO channels.

Over the years, NCF have invested in a variety of HPC systems to serve public scientific research in an efficient way. For HPC, this has meant a low-latency, large shared memory supercomputer to accommodate the most demanding applications on one side (capability computing) and compute clusters for high-throughput applications on the other side (capacity computing). Historically, capability "supercomputer" systems – CDC Cyber205, Cray YMP and C90 vector systems – have been operational from 1983 to 1999. More recent initiatives include the following:

- 2000 installation of TERAS (a 'double' SGI Origin 3800 at SARA) with 1 teraflop/s and 1 TB memory
- 2001 co-financing the purchase of a Cray SV1e at RUG[37]
- 2001 co-financing the purchase of a IBM SP (Power3/Power4) at SARA
- 2001 SGI Altix3700 system with 2.2 TFlop/s peak performance and 1 TB memory.
- 2003 installation of the new national supercomputer, an SGI Altix at SARA
- 2006 participating in the upgrade of the national research network (SURFnet) to SURFnet-3 (2 Mb/s) and SURFnet-4 (34 Mb/s)
- 2006: National Supercomputer (TERAS/ASTER) 3 teraflop/s

[34] http://huygens.supercomputer.nl/NCF.
[35] http://www.nwo.nl.
[36] http://www.sara.nl.
[37] http://www.rug.nl.

- 2007: Installation of the IBM "Huygens" system based on POWER5+ technology at SARA. This temporary system has a theoretical peak performance of 14 teraflop/s
- 2008: Planned upgrade to the new IBM System p supercomputer with 60 teraflop/s peak performance, 15 TB of main memory and online storage capacity of 800TB. This will be the fastest computer in The Netherlands and one of the biggest systems in Europe.

VTT

VTT, the Technical Research Centre of Finland, is the biggest contract research organisation in Northern Europe. It is a non-profit research organisation with 2800 staff and an annual turnover of $350M (€217M). VTT is a part of the Finnish innovation system under the domain of the Ministry of Trade and Industry.

VTT is a participant in the Center for Scientific Computing (CSC), which serves as the Finnish IT Center for Science.

CSC

In Finland, the Center for Scientific Computing (CSC)[38] is the organisation responsible for national supercomputing, networking services and support. It is a non-profit company, whose shareholder is the Ministry of Education (MoE). As well as VTT, the Executive Board of CSC has representatives from academia, the MoE, and the Finish Meteorological Institute (FMI).

The first phase of the Cray XT4 supercomputer (**louhi.csc.fi**) has been installed at the CSC. It consists of six login nodes and 1012 compute nodes with Catamount operating system. Each node has a dual-core 2.6 GHz AMD Opteron processor and 2 GB memory.

For development of user's own software there is a programming environment consisting of Portland, GNU Compiler Collection and Pathscale Fortran, C and C++ compilers, MPI, ACML, and LibSci libraries, TotalView debugger and some profiling tools provided by Cray. The batch processing system of Louhi is PBS Pro. Theoretical peak performance of the first phase of Louhi is 10.5 Tflop/s.

The final configuration will be installed during the second half of 2008. The peak performance will then increase up to 70 Tflop/s, making it one of the most powerful computers in Europe.

Louhi is reserved for jobs, which need plenty of computing power and scale well for large numbers of processors. Parallel jobs can be run with minimum 64 and maximum 512 cores. "Grand challenge" problems may be run using the entire machine (2024 cores).

PAC Service Provision Responses – August 2007

A short survey regarding service provisions was sent to the state based Partnerships for Advanced Computing (PACs). The questions and responses are detailed below.

Respondents:

VPAC: Bill Appelbe

SAPAC: Anthony Williams

[38] http://www.csc.fi/english.

iVEC: Andrew Rohl

TPAC: Nathan Bindoff

QCIF: Bernard Pailthorpe

1. What is the approximate ratio of spending on staff versus infrastructure?

VPAC

A hard question to answer, as we mix commercial work with Member support, but at the very least, I think you should be looking at 1-1 spending. Here are the "back of the envelope" calculations.

We spend about $1M p.a. on HPC, on a 'rolling purchase basis" (always have two systems operational)

We support about 500 users, of which about 75-100 are active at any one time

We employ about 8 EFT in support of Members, in everything from systems support to parallel programming and strategic planning (often fractional)... in all, that probably adds up to about $1M pa, but we have also provided eResearch grants that would push the support $ closer to 1-2

SAPAC

We would average around $1m per annum on infrastructure although as you will appreciate this can be "lumpy" from year to year. We would average around the same for staff costs.

iVEC

1.8 staff to 1 infrastructure

TPAC

~350 k per annum direct spend on infrastructure for our HPC facility,

with ~ 2EFT equating to ~200k. Considerable in-kind is provided by the University for housing, and management of services (e.g. aspects of networks etc) that sit around the HPC and data storage facility.

QCIF

Overall, its 1:2 - meaning ~ 1/3 staff, 2/3 infrastructure.

Major infrastructure installation is at UQ, with all other sites having some scale (eg~64-128-256 processors) of installations. As an indication - at JCU, its more like 1:1 .

2. What kind of capacity planning do you do when considering infrastructure upgrades/replacements and staffing levels?

VPAC

We analyse needs, and do that in a semi-quantitative way (look at systems load, requests for support, and have some spreadsheets, etc). It is hard to quantify, as many users can always use up any and all support we provide...

SAPAC

For us infrastructure upgrades are usually funded through ARC LIEF grants or similar and much less frequently State Govt. The process is that we go the research community to find out what thy want/need and then package a research proposal to match. We retire facilities when they are no longer worth paying electricity for and/or when we can no longer afford repairs and/or computer room capacity.

iVEC

Rob Woodcock will tell you at great length why you cannot capacity plan effectively in this area. Our approach is that we are constrained by budget and know we want a staff/infrastructure ratio of about 2:1.

TPAC We choose new HPC facilities based on benchmarks against our key weather and related software codes. We choose our new data storage facilities based on growth of our data holding/digital repositories. Their growth paths differ. We have more staff now, because the services we now provide has widened. In the past we operated a pure HPC facility, now we operate HPC, data storage, tape silo under HSM, and collaboration services, which include grid computing, access grid nodes, etc.

QCIF

QCIF conducts annual, professionally managed Strategic Planning exercises - last in Aug 2006; there is a follow-up operational planning meeting this week.

These involve consultations with relevant state agencies and industry.

Then the HPC Unit (normally in ITS Division) of each member university consults'

with its user community to establish needs. In the past 2 years there have been major compute upgrades at Griffith, QUT and JCU. UQ is currently doing such an assessment, involving major Centres (ARC, etc), Institutes, etc under the auspices of DVC, R.

3. Does your PAC employ a well defined IT management framework (such as ITIL) through which services are provided?

VPAC

Yes, we use an ISO model for project support, and reporting is done in a quantitative way to the Board (cycles used, projects, hours spent in supporting each member projects). We are ISO certified. We have many ITIL aspects to what we do, but are not formally using ITIL (instead embed that in an ISO framework - ISO is required for some commercial clients)

SAPAC

Not at present.

<u>iVEC</u>

No

<u>TPAC</u>

No, we operate in a similar, but smaller scale, like the APAC facility in Canberra.

<u>QCIF</u>

At JCU the general framework is ITIL, and is fully integrated into the board IT service infrastructure of the university, for example storage, directory and backup infrastructure are common. However, we have adapted some additional elements as some of your 'services' also involve development of applications. To this end JCU has TRAC based project management / ticketing environment that users have access to for requesting features, software bugs and so they can be informed of progress.

Griffith Uni is well advanced on Change Management and Service Level Management processes within its ITS Division.

Sorry - I am unsure of the situation at UQ - but see 6 below.

4. Do you have staff that is dedicated to specific research groups ("embedded capability") or do they work on a more remote basis?

<u>VPAC</u>

Yes, we have dedicated staff (embedded capabilities) in certain areas (notably in the AusScope project, working out of Monash; also in staff that are seconded into certain projects or staff that we hire to work on specific projects). We do two types of embedding: physical (the staff is located at a Member site) and logical (the staff are dedicated to a particular project or client on a full-time basis).

<u>SAPAC</u>

We have staff partially embedded but their home office remains SAPAC, although they may spend a day a week at the target research group site. This typically is associated with that research group co-investing in the staff members' salary.

<u>iVEC</u>

We have both kinds of staff, although we have moved more to centering staff at the iVEC facilities (there are three of these), rather than in specific research groups. I suspect we will move back in the other direction as iVEC continues to grow

<u>TPAC</u>

We have 1 staff member who is dedicated to portal development, digital repository portal, and another member who contributes 40% of her time to filling the portal with data sets.

QCIF

At UQ, two APAC-funded staff were deployed to users groups: to Bioinformatics (IMB), and earth system sciences (via ACCESS MNRF - now NCRIS AusScope). That is in transition under NCRIS PfC arrangements, but is likely to continue. Note also that some groups, such as Chem & Phys, have their own HPC support staff employed within research centres.

My understanding is that HPC Unit at QUT uses some staff embedded in subject areas, but I do not have details.

None is embedded at JCU.

5. Are your staff knowledgeable and trained in specific disciplines – such as CFD or bioinformatics – or are they more like general IT specialists?

VPAC

There is a continuous spectrum. We have staff that are general systems administration staff, and staff with specialist skills in certain applications domains (e.g., we have experts in CFD, Health Informatics, Geospatial data processing, Molecular Modelling). We also do some in-house up-skilling/training, on an "apprenticeship" basis (e.g., to train staff up in parallel programming)

SAPAC

Our staff are typically IT knowledgeable but either had on appointment or have since developed a set of discipline specific capabilities. Our goal is to have each staff member able to work closely with two or more discipline groups.

iVEC

We have very few IT specialists - most of our people have domain level knowledge

TPAC

Two staff are system administrators.

One staff member a specialist programmer and grid project leader, and 40% staff member who is a data curator.

Last year I also had a climate systems scientist who specialised in Earth Systems Models, and now looking to replace that person in the near term. I also have a specialist for the specific work on porting and tuning climate codes, like the Unified Model.

QCIF

A few at UQ are - e.g. in bioinformatics, CFD, computational chemistry (as per 4).

Many at UQ are IT specialists.

At JCU, Staff profile is 1x chemist, 1x applied mathematician and 1 IT specialist. All staff are periodically trained in particular tools/disciplines so they can converse more efficiently with user communities.

UQ has just appointed a PhD-level Manager for its HPC Unit, who is actively doing

a skills assessment to guide future recruitment.

6. What kind of tracking system – if any – do you have in place for managing user requests?

VPAC

We use the TR (I think that that is the name) an open-source product...

SAPAC

We are about the turn on Request Tracker for our Helpdesk support and have been planning this for some time. We have been consumed with finishing out new $500k computer room and relocating our new six TFlop quad core SGI system there.

iVEC

Email lists

TPAC

We have a help desk facility.

QCIF

UQ ITS has deployed a commercial product called ITSM for tracking user requests across all if its service areas, including HPC. It does the basics and is being enhanced, in a formal project methodology, to incorporate features such as service level management (e.g. If a SLA breach is imminent, then escalations start occurring automatically)

Such helpdesk systems develop an internal knowledge base to assist with future enquiries.

Users report (via phone, 3365-6000, or email,help@its.uq.edu.au, causing an incident to be created and tracked to resolution.

Griffith uses an in-house developed web app for tracking requests, incidents etc.

At JCU, the use a home-developed environment that involves a phone based CMS and a TRAC-based ticketing system that links with a CVS. This is reported to be very effective. It advantages to a regular wiki can be characterized as:

1. Has a directory-based view so is structured, which gets around our the wiki problem that it can be hard to find pages and easy to get lost in all the content
2. Permissions to view pages can be done on a directory (and associated subdirectories) basis rather than individual pages

3. You can label pages with keywords and search on keywords and text and pull together a virtual directory based on the results
4. It has a simple WYSIWYG editor for pages rather than the making users learn HTML or a wikiML
5.We have developed a Shibboleth interface
6. You can skin it to make nice looking pages aimed at users.

Appendix 10: Sources

Arrold, B. *et al* 2006. CSIRO HPSC and the HPCCC. Discussion Paper.
http://escience.csiro.au/twiki/pub/Main/StrategyDocuments/CSIRO_HPSC-HPCCC_paper_July_2006_-_Brennan_John_pivotal_paper.pdf

Benioff, M. and Lazowska, E. (Chairs), *Computational Science: Ensuring America's Competitiveness*; President's Information Technology Advisory Committee (PITAC) Report, June 2005.
http://www.nitrd.gov/pitac/reports/20050609_computational/computational.pdf

Brill, K.G. 2006. 2005-2010 Heat density trends in data processing equipment. Uptime Institute White Paper.
http://www.uptimeinstitute.org/

CSC. 2006. CSC, The Finnish IT Center For Science. Annual Report 2006.
http://www.csc.fi/english/csc/publications/reports/csc_annual_report_2006

CSIRO, 2007. Submission to NCRIS Platforms for Collaboration (5.16).
http://escience.csiro.au/twiki/pub/Main/StrategyDocuments/CSIRO_NCRIS_Proposal_v_1.4_Final.pdf

CSIRO, 2007. Strategic Plan for 2007-2011.
http://intranet.csiro.au/intranet/sipp/06_Strategy/CSIRO_StrategicPlan_07-11.pdf

CSIRO, 2007. IM&T Operational Plan 2007-2008.
http://intranet.csiro.au/intranet/it/Operational_Plan_2007-2008.pdf

CSIRO IM&T, 2007. CSIRO eScience Information Management Strategy.
http://intranet.csiro.au/intranet/it/CSIRO_eSIM_Strategy.pdf

DEST, 2007. Platforms for Collaboration Investment Plan.
http://www.ncris.dest.gov.au/capabilities/documents/PfC_Investment_Plan_pdf.htm

Elsum, I. 2006. Computational Science: SIP2 Broad Directions Setting, August 2006. Internal CSIRO briefing paper.

Fagerholm, J. *et al*, 2005. CSC Report on Scientific Computing in Finland 2004/2005.
http://www.csc.fi/downloadPublication?uid=265e985845b9c33a08e01136d3ee1032

Fraunhofer-Institut für Techno- und Wirtschaftsmathematik ITWM. Annual Report 2006.
http://www.itwm.fraunhofer.de/zentral/download/jb2006/jb2006_en_72dpi.pdf

Frost R, Francis R. 2005. Costing - Sun cluster support proposal – Final.
https://teams.csiro.au/sites/ascs/References/COSTING%20-%20Sun%20cluster%20support%20proposal%20-%20FINAL.doc

Menuet, R. 2006. Continuous Cooling is required for continuous availability. Uptime Institute White Paper.
http://www.uptimeinstitute.org/

National Energy Research Scientific Computing Center (NERSC).
http://www.nersc.gov/

Netherlands National Computing Facilities foundation.
http://www.nwo.nl/nwohome.nsf/pages/ACPP_4X6R5C_Eng

NSF Cyberinfrastructure Vision for 21st Century Discovery. U.S. National Science Foundation NSF07-28, 2007.
http://www.nsf.gov/pubs/2007/nsf0728/index.jsp

Office of Advanced Scientific Computing Research. U.S. Department of Energy.
http://www.er.doe.gov/ascr/

Powick, A. *et al*, 2006. CSIRO IT Review.
http://intranet.csiro.au/intranet/it/reportsandplans/CSIRO_IT_Review_Intv4.pdf

Risk Solutions, 2006. Hadley Centre (UK Meteorological Office) Review 2006 Final Report.
http://www.defra.gov.uk/environment/climatechange/research/pdf/hadleycentrereview0507.pdf

Scientific Discovery through Advanced Computing (SciDAC). U.S. Department of Energy.
http://www.scidac.gov/

TOP500 Project.
http://www.top500.org

Turner, W.P. 2006. Dollars per kW plus dollars per square foot are a better data centre cost model than dollars per square foot alone. Uptime Institute White Paper.
http://www.uptimeinstitute.org/

Turner, W.P. 2006. Tier Classification defines site infrastructure performance. Uptime Institute White Paper.
http://www.uptimeinstitute.org/

Turner, A. 2007. Supercomputer overhaul is out of this world.
http://www.smh.com.au/news/technology/supercomputer-overhaul-is-out-of-this-world/2007/08/20/1187462175419.html

Uhlherr, A. 2005. Utilisation of High Performance Scientific Computing in CSIRO. Comparative Review.
http://intra.hpsc.csiro.au/user/general_information/HPSC_Comparative_Review.pdf

Verastegui, B. 2005. Computing and Computational Sciences Directorate. Workshop Presentation to ORNL Research Alliance in Math and Science.
http://www.csm.ornl.gov/workshops/RAMSworkshop05/presentations/overview.pdf

Woodcock, R. 2007. Scientific Computing for CSIRO. Presentation to CSIRO scientists.
http://escience.arrc.csiro.au/twiki/pub/Main/Introduction/HPC_roadshow.ppt

Appendix 11: List of Contributors

The contributions of the following individuals towards this review are gratefully acknowledged.

Appelbe, Bill (VPAC)
Arrold, Brennan (CSIRO IM&T, Yarralumla)
Ayers, Greg (CMAR, Aspendale)
Baker, Justin (CSIRO IM&T, Melbourne - HPSC)
Bannon, David (VPAC)
Barber, Michael (Group Exec, North Ryde)
Barris, Wes (LI, St Lucia)
Beavis, James (DET, Clayton)
Bell, Robert (HPSC, Melbourne - HPSC)
Bindoff, Nathan (TPAC)
Breadmore, Andrew (CSIRO IM&T, Clayton)
Bryan, Geoff (CSIRO IM&T, Clayton)
Caris, Con (E&M, QCAT)
Chandry, Scott (FSA, Werribee)
Cornwell, Tim (ATNF, Marsfield)
Courtney, Tim (CSIRO IM&T, Yarralumla)
Craig, Stuart (E&M, QCAT)
Dalrymple, Brian (LI, St Lucia)
Doherty, Gary (ICT Centre, Marsfield)
Dungavell, Ross (CSIRO IM&T, QCAT)
Elsum, Ian (Sci Plan, Canberra)
Ennis-King, Jonathan (CPR, Clayton)
Epa, Vidana (CMHT, Parkville)
Evans, Ben (ANU)
Fagerholm, Juha (CSC, Finland)
Fitch, Peter (CLW, Black Mountain)
Foster, Marcus (CSIRO IM&T, Werribee)
Francis, Rhys (HPSC, Melbourne - HPSC)
Frost, Rozanne (CSIRO IM&T, Yarralumla)
Gladman, Simon (FSA, Werribee)
Gordon, Karl (Entomology, Black Mountain)
Garrett, Geoff (OCE, Canberra)
Griffiths, Cedric (CPR, Kensington)
Gureyev, Tim (CIP & CMMT, Clayton)
Henshall, John (LI, Armidale)
Hill, Rod (Group Exec, Clayton)
Hind, Tracey (CSIRO IM&T, Melbourne - HPSC)
Hipworth, Mark (CSIRO IM&T, Clayton)
Hyatt, Alex (LI, Geelong)
Ignatius, Janne (CSC, Finland)
Kennedy, Gavin (PI, Black Mountain)
Kimber, Will (ICT Centre, ANU - Acton)
Knudson, Rosanna (CSIRO IM&T, Yarralumla)
Koehler, Alan (Ensis, Clayton)

Lambropoulos, Nicholas (DET, Newcastle)
Lovell, David (CMIS, ANU - Acton)
Matuszkiewicz, Tad (CSIRO IM&T, Yarralumla)
McDonough, Aaron (HPSC, Melbourne - HPSC)
McWilliam, Sean (LI, St Lucia)
Michielse, Peter (NCF, Netherlands)
Miller, Shona (CSIRO IM&T, Yarralumla)
Minchin, Stuart (CLW, Black Mountain)
Mitchell, Chris (CMAR, Aspendale)
Morgan, Polly (HPSC, Melbourne - HPSC)
Morrissey, John (CSIRO IM&T, Yarralumla)
Needham, David (CSIRO IM&T, Yarralumla)
Northwood, Colin (CSIRO IM&T, Geelong)
Pailthorpe, Bernard (QCIF)
Perrin, Ewan (CSIRO IM&T, Yarralumla)
Peterson, Nathan (CSIRO IM&T, Clayton)
Poole, Nigel (OEDBusServ, North Ryde)
Rees, Rob (DET, Clayton)
Rohl, Andrew (E&M, Kensington)
Russo, Salvy (Minerals, Clayton)
Schiller, Andreas (CMAR, Hobart)
Scott, Craig (CMIS, Clayton)
Sellahewa, Jay (FSA, North Ryde)
Sheppard, Martin (CSIRO IM&T, Adelaide)
Smart, Bob (HPSC, North Ryde)
Smith, Daniel (HPSC, Melbourne - HPSC)
Stinziani, Antony (CSIRO IM&T, Yarralumla)
Taylor, John (CMIS, ANU - Acton)
Tildesley, Paul (CMAR, Hobart)
Toll, David (Finance, Canberra)
Tribe, Deirdre (CSIRO IM&T, QCAT)
Tse, Winnie (Finance, Canberra)
Tutungi, Vicki (NM F/Ship, Clayton)
Uhlherr, Alf (CMHT, Clayton)
Varghese, Jose (CMHT, Parkville)
Williams, Anthony (SAPAC)
Williams, Gareth (HPSC, Melbourne - HPSC)
Wilson, Bill (CMIS, North Ryde)
Winkler, Dave (CMHT, Clayton)
Witt, Peter (Minerals, Clayton)
Woodcock, Robert (E&M, Kensington)
Youill, Peter (CSIRO IM&T, Yarralumla)
Zacharia, Thomas (ORNL CCSD, USA)
Zelinsky, Alex (ICT Centre, Marsfield)

Contact Us

Phone: 1300 363 400
+61 3 9545 2176

Email: enquiries@csiro.au

Web: www.csiro.au

Your CSIRO

Australia is founding its future on science and innovation. Its national science agency, CSIRO, is a powerhouse of ideas, technologies and skills for building prosperity, growth, health and sustainability. It serves governments, industries, business and communities across the nation.