



# 2006 Review of the UK HPC Integration Market

C Kitchen, M Guest, M Deegan, I Kozin

April 2007

**©2007 Science and Technology Facilities Council**

Enquiries about copyright, reproduction and requests for additional copies of this report should be addressed to:

Chadwick Library  
Science and Technology Facilities Council  
Daresbury Laboratory  
Daresbury Science and Innovation Campus  
Warrington  
WA4 4AD

Tel: +44(0)1925 603397  
Fax: +44(0)1925 603779  
email: [library@dl.ac.uk](mailto:library@dl.ac.uk)

Science and Technology Facilities Council reports are available online at: <http://epubs.cclrc.ac.uk/>

**ISSN 1362-0207**

Neither the Council nor the Laboratory accept any responsibility for loss or damage arising from the use of information contained in any of their reports or in any communication about their tests or investigations.

# 2006 Review of the UK HPC Integration Market

December 2006

Version 2.1<sup>1</sup>

---

## *Abstract*

In this, the 2<sup>nd</sup> revision of the “UK HPC Integration Market”, we consider the multitude of issues faced by an organisation when deciding how best to procure, maintain and maximise the usage of a given HPC resource. In over-viewing the current HPC landscape, we concentrate on the potential role of HPC integrators in any partnership that looks to maximise this entire process, and whether such organisations based primarily in the UK have the ability to provide the necessary level of expertise required in all phases of the process, from procurement, through installation onto ongoing support of the resource throughout its life cycle.

We consider how current HPC technology roadmaps might impinge on the role of integrators in responding to the undoubted challenges that lie ahead. Crucial issues when considering potential integrator involvement include both size of the hardware solution i.e., number of cores, number of nodes, and the ongoing robustness of open source software solutions that might be deployed on these platforms. Informed by developments over the past 12 months associated with the deployment of systems funded under SRIF, the Scientific Research Investment Fund, we provide an in-depth analysis of the current status and capability of a number of the leading HPC Integrators within the UK. Our primary attention is given to the four major companies who supply the academic community and hence are well known to us – Streamline Computing, ClusterVision, OCF and Compusys. Six other integrators are also considered, albeit with less rigour.

It should be emphasised that all the information in Section 3 is provided directly by the companies and do not necessarily reflect the views of STFC. With the exception of Cambridge On-line systems and Silicon Mechanics, the information in Section 4 has been obtained from the Cluster Integrator’s official websites. The document concludes with a 10-point summary of important considerations when procuring HPC clusters (mid-to-high-end compute clusters).

If the reader has questions regarding issues raised in the document, or requires further feedback they are encouraged to contact us directly – details are available on our website and in Appendix C.

---

## *Contents*

1.	Introduction.....	2
2.	The HPC Landscape .....	3
2.1	Moore’s law and hardware development .....	4
2.2	Operational Challenges with Large Systems.....	5
2.3	Capacity Through Beowulf Clusters .....	6
2.4	Developments in the UK – SRIF3.....	6
2.5	Proprietary Vendors and System Integrators.....	7
2.6	HPC System Integrator Criteria .....	8
3.	State of the UK HPC Integration Marketplace .....	9
3.1	Brief Summary of the Clusters Deployed in 2006. ....	9

---

<sup>1</sup> Revised April 2007

3.2	ClusterVision.....	11
3.3	Compusys.....	14
3.4	OCF Plc.....	18
3.5	StreamLine Computing .....	23
4.	UK HPC Integrators II.....	29
4.1	Cambridge Online Systems Ltd .....	29
4.2	Silicon Mechanics .....	30
4.3	Linux Networx .....	34
4.4	Western Scientific .....	36
4.5	SCC .....	36
4.6	Workstations UK.....	38
5.	Summary and Conclusions .....	38
6.	Bibliography .....	40
7.	APPENDIX 1: Integrator Questionnaire .....	41
8.	APPENDIX 2: Company Contact Details .....	42
8.1	ClusterVision.....	42
8.2	Compusys.....	42
8.3	OCF plc.....	42
8.4	Streamline .....	43
8.5	Cambridge On-Line .....	43
8.6	Linux Networx.....	43
8.7	Silicon Mechanics .....	43
8.8	SCC Plc.....	43
8.9	Western Scientific .....	44
9.	APPENDIX C: Daresbury Contact Details.....	44

---

## 1. Introduction

This paper considers the current HPC landscape and the multitude of issues faced by an organisation when deciding how best to procure, maintain and maximise the usage of any associated HPC resource. Historically many organisations have responded to this challenge by forming partnerships with the associated vendor of the hardware in looking to maximise this entire process. Our specific interest here lies in the potential role of HPC integrators in such a partnership, and whether such organisations in the UK have the ability to provide the necessary level of expertise required in all phases of the process, from procurement, through installation onto ongoing support of the resource throughout its life cycle.

The paper is structured as follows. In section 2 we provide a somewhat extended overview of both current and future HPC landscapes, and how the associated roadmaps might impinge on the role of integrators in responding to the undoubted challenges that lie ahead. We will see from the outset that crucial issues involve both size of the hardware solution i.e., number of nodes, and the ongoing robustness of open source software solutions that might be deployed on these platforms. In considering this landscape, we review briefly developments over the past 12 months associated with the deployment of systems funded under SRIF, the Strategic Research Investment Fund. Sections 3 and 4 present the real detail of this paper - the current status and capability of a number of the leading HPC Integrators within the UK. In section 3 we provide an in-depth analysis of the four major integrators who supply the academic community and hence are well known to us - Streamline Computing, ClusterVision, OCF and Compusys – this analysis has been conducted by a variety of mechanisms that are detailed at the appropriate point; section 4 summarises data from the smaller, less-known organisations.

Finally, we provide a 10-point summary of the findings of this paper in section 5.

## 2. The HPC Landscape

Today's supercomputing choices are the product of commodity market competition, technology evolution, historical hardware and software legacies, and leadership choices within industry. Computer vendors, driven by developments such as DOE's Accelerated Strategic Computing Initiative (ASCI [1]), have aggressively pushed the performance levels of parallel supercomputers higher and higher. Given the economies of scale, it is clear that the immediate future of supercomputing will be dominated by parallel supercomputers built with commodity compute servers, such as those used as web servers, tied together by a high-performance communications fabric. Although currently on this plateau in the evolution of parallel supercomputer architectures, this will not last long. New architectures are already on the drawing boards that will be capable of a quadrillion arithmetic operations per second (PFlops). Such computers cannot be built using the same technology in today's TFlops computers – they would require too much space and consume too much power

**Custom and Commodity Clusters:** The evolution of cluster technologies has proceeded on two fronts: in the first development, beginning in the late nineties, IBM, Compaq and SGI - among others - began creating *proprietary clusters* using their shared-memory servers and custom-designed or semi-commodity networks. The IBM approach has been to use their own custom multi-level switch fabrics to interconnect shared-memory nodes based on their Power workstation processors. Such nodes have been deployed on the UK's HPCx service – initially 32-way Regatta p690 and p690+ nodes comprising power4 and power4+ processors, and more recently power5 p5-575 nodes comprising 16 processors.

At the same time, true *commodity clusters* were being built and deployed based on uni-processor or dual-processor nodes utilizing Compaq Alpha or Intel X-86 processors. These clusters used mainly semi-commodity Myrinet [2] interconnects from Myricom; but smaller examples were sometimes based on gigabit (or slower) Ethernet switch fabrics. In all cases, the system software was built around the Linux open-source operating system. A number of Terascale commodity clusters have now been installed, including Intel x86-based clusters at Los Alamos (LANL) and Lawrence Livermore (LLNL) National Laboratories, and AMD Opteron systems at CINECA (Italy), Geoscience (Texas, USA), SSC Karlsruhe (Germany) utilising a variety of commodity (gigabit Ethernet) and semi-proprietary interconnects, e.g. Myrinet, Infiniband. The most notable systems here are:

- The 53 TFlops (65 TFlops peak) *Thunderbird* cluster at NNSA / Sandia National Laboratory [8] featuring 9024 PowerEdge 1850 Intel Xeon (3.6GHz) processors with Infiniband Interconnect (Number 6 in the Top 500).
- The 47 TFlops (83 TFlops peak) *TSUBAME* grid cluster at GSK Centre (Tokyo Institute of technology) [9] featuring 11088 SunFire x4600 AMD Opteron (2.4/2.6GHz) processors with ClearSpeed accelerators cards and Infiniband Interconnect (Number 9).
- The 42 TFlops (62 TFlops peak) *Jaws* Cluster at Maui HPC Centre (MHPCC) [11], featuring 5200 Dell PowerEdge 1955 Intel Xeon (3.0GHz) processors with Infiniband Interconnect (Number 11).
- The 41 TFlops (55 TFlops peak) *Lonestar* Cluster at the Texas Advance Computing Centre (TACC) [12], featuring 5200 Dell PowerEdge 1955 Intel Xeon (2.66GHz) processors with Infiniband Interconnect (Number 12).
- The 18 TFlops (28 TFlops peak) *Darwin* Cluster at Cambridge University (SRIF3 funded) [13], featuring 2340 Dell PowerEdge 1955 (3.0GHz) processors with InfiniPath Interconnect (Number 20).
- The 11.2 TFlops *Lightning* Cluster at LANL [4] featuring 2,816 of AMD's 64-bit Opteron chip (Number 89), a real competitor to Intel's costly Itanium product line<sup>2</sup>.

---

<sup>2</sup> Note that Lightning was installed by the US HPC Integrator, Linux Networx (see section 4.2)

Overall in the Top 100 clusters, 20 systems are based on the commodity AMD Opteron processor (6 of these systems are in the guise of the Cray XT3/4 product line), and 18 are using Intel's Xeon processor. It is noticeable that the Lightning Cluster which was at position 32 in 2005's Top 500 list has slipped 49 places to number 83.

*It is perhaps worth stressing that commodity clusters based on open source software (i.e., Beowulf clusters) are 2-10X more cost effective than clusters based on proprietary solutions.*

None of these clusters - custom or commodity - have system balance between computation and communications that is competitive with that found historically on true MPPs such as the Cray T3E. Nevertheless, for many important classes of applications, they are capable of achieving high parallel efficiency on a thousand processors or more. They also, in general, lack full-system Reliability, Availability and Serviceability (RAS) features. For truly large systems, this has caused difficulties in running large jobs with long execution times. In addition, few of the large clusters deployed have truly scalable system software i.e. can provide service to dozens of simultaneous users and have fast, scalable system boot-up, and executable loading capabilities.

**Open Source Solutions:** Open Source developments have yielded significant enhancements to the state-of-the-art in operating systems (c.f., Linux OS) and tools (c.f., Apache). Multiple accretions of open software mass have resulted in profitable enterprises that combine these tools into single offerings with support. In addition, there exist many efforts to build clustering tools (c.f., LLNL Chaos, LANL Clustermatic, Sandia CPLANT, NPACI Rocks, OSCAR and others) that extend the desktop environment to medium and high-performance computing. The potential deployment of open source solutions in satisfying requirements of High-end, Terascale Technical Computing is more an open question. What is clear is that current trends in developing and implementing ultra-scale computers fall well below the requirements capable of addressing many scientific challenges. Based on a number of factors, we believe that ultra-scale platforms will be developed using - among other things - commodity-based hardware components and open source software. It is clear, however, that a major investment is required in developing these software components, technologies, and enhancements to attain the 100X to 1,000X in computational capability and capacity that is vitally needed to address scientific requirements.

In a balanced computational environment, hardware and software improvements have contributed equally to the increase in applications performance. What matters most is the shortest time to solution, with code development efforts definitely on the critical path to success. The delivery of fully integrated heterogeneous, multi-vendor parallel computing environments pose considerable challenges from system administration to application development to resource management. If these challenges are beyond current Tier-1 vendors, assessing the potential of HPC integrators to respond to the same challenge raises a number of issues that need to be addressed.

## **2.1 MOORE'S LAW AND HARDWARE DEVELOPMENT**

Chip performance is expected to continue along the Moore's law trend line for the next 5 to 10 years, with the transistor density at least continuing to grow during the foreseeable future. The processor core on IBM's Power4 processor occupies approximately 1/4 of the transistors (and area) of the processor using today's CMOS technology. This implies that in two processor generations from now, i.e., in about three years, the core will only occupy 1/16 of all available transistors. Development economics and high demands on time-to-market will drive the use of common processor cores across many applications using a common base instruction set and a common micro-architecture. It is therefore likely that the extra space enabled by the Moore's law development of transistor density will be used to obtain customised chips, i.e., chips for different purposes which are based on a common core. For example, a low end chip could be obtained by adding fast I/O and L2 cache using the extra space for transistors; a game chip by adding high bandwidth memory interface, high

bandwidth graphics interface and special arithmetic processors; a network processor chip by adding external cache controller, packet processor, protocol accelerator, and Ethernet interface; or a dense server chip by placing several cores plus external cache controller and memory control on a single chip.

We would, however, note the major shift in the semiconductor industry strategy that has become readily apparent during the past two years, whereby increased processor throughput is accomplished by providing more computing elements (“cores”), rather than by increasing the operating frequency. Indeed, the 5-year Semiconductor Industry Roadmap showed a factor of two increase in operating frequency from 3.2 GHz (2004) to 6.4 GHz (2009), whereas Intel, AMD, and others are migrating to multiple (2x, 4x, 8x ...) cores within a single “processor.” Experience has already shown that balance on multi-core processors can be a significant issue; this change in processor design has important implications for Petascale computing (see 2.2).

Finally we note that an alternative trend aimed at improving performance / Watt and total cost of operation will be to use less powerful processors that use much less power, and constrained memory and I/O. Such a trend will mean that many more processors may be required to reach a desired aggregate performance, further increasing the burden on programmers and users of HPC systems.

## 2.2 OPERATIONAL CHALLENGES WITH LARGE SYSTEMS

When delivering and supporting the world’s largest systems (e.g., the ASCI systems - Red, Blue Mountain, Blue-Pacific, White, and Q - and HPCx etc), the associated sites encountered numerous problems with the software environment. These “defects” are often difficult to isolate down to root-cause. In addition, many of the vendor partners providing these platforms do not have systems even within 1/10 the size of these platforms for debugging and patch verification. If a Tier-1 vendor does not have these capabilities, there is little or no chance that an integrator could claim to be able address this shortcoming. All these factors contribute to the process of integrating large systems, resulting in inordinately long periods of “debugging” time (over a year). Several lessons learned in these efforts suggest that high-performance computing goals might best be served by open source development efforts where the end-sites and their associated communities directly participate in the development and/or testing and/or integration of the technology on platforms at scale.

These conclusions are applicable to the broader high-end technical computing (HEC) community for the following reasons:

- HEC customers require the ability to have operating system and tools source code for root cause analysis and debugging. In many cases – at least in the USA – the HEC customer sites have the requisite in-house expertise not only to do root-cause fault isolation, but also to formulate and implement bug fixes (that may possibly include re-architecting portions of the solution). With Open Source community development efforts, these fixes can be fed back into the community for the benefit of all.
- Taken as a group, vendors of proprietary solutions indicate that they cannot make a profit providing solutions that span the entire HEC space. Thus HEC sites are left with two options: live with mediocre solutions that fulfil only part of their requirements, or implement major portions of the solution in-house on top of the vendor provided proprietary foundations. Neither of these two options is advantageous for the HEC or vendor community in the long run. With the Open Source community development model, HEC sites can contribute their solutions back to the community because it is based on a common software base and can be contributed back without fear that the development would benefit only one particular vendor offering (and thus inhibit competition in the future and lock the HEC site into one vendor proprietary foundation). With this wide HEC scope, the diversity of the debugging environment and developed solutions benefits the entire community.

- Change has been a dramatic feature of the HPC landscape for many years. However, recent introductions of many disruptive technologies (e.g., killer micros, IA-32 commodity hardware, open source software) have radically increased the rate of change in the industry. As a result, the time span for vendor support of provided hardware and software platforms has been reduced to shorter than the full lifecycle of implemented solutions at HEC sites. In addition, timescales under which support is withdrawn (from announcement to withdrawal of support) by vendors (measured in quarters) is typically much shorter than the governmental planning period for replacements (measured in years). This mismatch in timescales has led HEC sites to require Open Source based products as a hedge against “change of support” status.

*Processors vs. Processing Elements.* As noted above, a major shift in the semiconductor industry strategy will lead to increased processor throughput being accomplished through more computing elements (“cores”), thereby increasing the burden of parallelization. Future baseline architectures for 100 TFlops and 1 PFlop might thus appear as below:

- 100 TFlops:
  - 5 GFlops each “processing element”
  - 20,000 processing elements
- 1 PFlop
  - 10 GFlops each processing element
  - 100,000 processing elements

This increased reliance on parallelization and hence system size will merely act to emphasise the operational challenges associated with Large Systems, challenges that stretch the resources of proprietary vendors to the limit and are realistically beyond the reach of the HPC integrators central to this paper.

### **2.3 CAPACITY THROUGH BEOWULF CLUSTERS**

The success ASCI and international HEC efforts to deliver complex modelling capability to the community has led to an overwhelming demand for “capacity at the capability level.” Now that large scale predictive scientific simulation has been adopted as a critical component of scientific endeavour, there is great demand to run parallel applications for parameter studies, convergence studies and the like at small to medium scale of parallelism. These simulations run at a scale of parallelization that was a “capability run” for previous generations of ASCI platforms. However, rather than one or even a few capability runs, literally hundreds of delivered TFlops of computation are required for these production calculations.

There is now widespread debate on how best to meet these crushing capacity demands. Although the strategy - and indeed a detailed analysis of the requirements - is still ongoing, one aspect is certain: the solution will be dependent on commodity clusters that are based on open source software (i.e., Beowulf clusters). Again, this is being driven by the fact that Beowulf clusters are 2-10x more cost effective than clusters based on proprietary solutions. It would seem clear from the outset that *the role of system integrators in providing this scale of resource is far more credible than is the case for their involvement at the very high-end i.e. in the capability regime.*

### **2.4 DEVELOPMENTS IN THE UK – SRIF3**

The mid-range HEI computing landscape in the UK had changed dramatically over the past 12 months thanks in large part to the major investments coming from SRIF, the Science Research Investment Fund. SRIF is a joint initiative by the Office of Science and Technology (OST) and the Department of Education and Skills (DfES) [14]. Its purpose is to contribute to higher education institutions’ (HEI) long-term sustainable research strategies and address past



under-investment in research infrastructure. SRIF also takes into account the need for institutions to make their expertise and facilities more open to access by business and encourage collaboration between HEIs, industry, charitable bodies, Government and other partners. Once grants are assigned to the universities, it is the universities decision on how to allocate resources to the various departments.

#### **2.4.1 SRIF-3: 2006-08**

Capital grants made available to institutions from April 2006 must be spent by March 2008. Funding is distributed by formula as a conditional allocation. £903 million has been allocated for research capital to English Higher Education Institutions, of which £38M has been assigned to computer procurements. Heriot-Watt University is leading the administrative effort in coordinating these procurements. The idea behind centralising the procurement process is to deliver maximum return on investment (ROI), while ensuring that the smaller procurements are not overshadowed by some of the more lucrative £1M+ tenders.

The University HEC sector is now the major provider of HPC resources to the UK research base. As a result of the SRIF3 investment, the capability of this sector has increased 100 fold increase in during the 2005-2008 timeframe. The University HEC sector now has the potential to provide four times the capability of the UK National HPC service provider tier (HPCx and HECToR) over the 2006-2008 timeframe, with a number of 2000+ core systems scheduled to come on line during 2007-2008, resulting from the three procurement Tranches undertaken during 2006/7. The procurement and subsequent installation of major systems at the Universities of Birmingham, Bristol, Cambridge, Cardiff, Edinburgh, Warwick and UCL has led to major changes within the Integrator marketplace, and the evolving relationships between these organisations and their Tier-1 counterparts which is developed further below.

### **2.5 PROPRIETARY VENDORS AND SYSTEM INTEGRATORS**

Market dynamics will not support substantial unique developments for ultra large systems. A guiding principle for many of the proprietary vendors when building HPC systems is to leverage standard technology wherever possible, and to judiciously add custom technology only when this is required. Further, any new technology must ultimately be applicable to mainstream commercial applications and systems for it to be viable from a business perspective.

The vendor is becoming more and more of a system integrator. In this context it should be mentioned that commodity-based Beowulf systems are rapidly becoming the systems of choice, at least within the academic community. This growth is fostered in the UK by the emergence of a number of companies who do provide crucial added-value services around clusters that address existing shortcomings in the standard, open source based environments, while keeping costs to a reasonable level (see section 3). Procuring systems from such integrators may be more prudent than relying on the more traditional Tier-1 companies such as IBM or HP given the existing cost differential in the associated products, although caution should be applied and reference sites always sought for feedback regarding the previous installations (preferably of similar sized systems). While some way removed from providing credible high-end alternatives for e.g. HECToR (Cray), CoH; the emergence of enabling technology infrastructures provided by toolkits such as OSCAR [6] and ROCKS [7] has taken much of the hassle away from supporting such systems, at least up to 128 CPUs. The emergence of essential features such as check point restarting, concurrent file systems etc. needs to be closely monitored in judging the "fit-for-purpose" nature of commodity systems at the high-end.

## 2.6 HPC SYSTEM INTEGRATOR CRITERIA

Continuing the theme of section 2.5, we need to consider an appropriate set of performance criteria that might enable us to differentiate between the various cluster integrators in the market. Based on an objective assessment of their standing against these criteria, it should be possible to identify potential HEC candidates - the more performance target-compliant a vendor is, the more likely they are to provide a viable alternative to the more traditional Tier-1 companies.

We would consider from the outset that the following issues need to be assessed when judging the viability of a given integrator:

1. *In-house technical expertise:*

The 'added-value' that the vendor brings to the procurement through software stacks, finely tuned OS etc.

Number of developers - the vendor's ability to develop, support and maintain the software so as to sustain 'cutting-edge' technology and assist with various code porting and optimization tasks.

2. *Size of the company:*

Is the company sufficiently staffed to be able to support large scale compute clusters over multiple installation sites?

Turn-over - is the company a practical long-term prospect with the potential to be 'self-sufficient' with respect to large clusters (> 1000 processors) within the next few years?

3. *Current install-base:*

An important factor is the integrator's current success in the small- to mid-range computer cluster market and the actual install base (whether solely to the UK or if they have a presence overseas, particularly in Europe or the USA).

This again provides information on the potential longevity of an integrator as well as feedback from the community regarding their overall performance.

4. *Support infrastructure:*

The number of technical and software engineers in place to support the cluster throughout its lifetime and importantly whether this is in-house or out-sourced. If out-sourced, whether there are any plans to change this in the near future.

These points formed the basis of discussions with the key cluster integrators at the outset of this analysis exercise – a more detailed account of these points can be found at Appendix 1.

In addition to the above, the possibility of an integrator solution rather than one from the more traditional Tier-1 vendor is very much dependent on the nature of the HPC resource(s) under consideration. Some obvious examples would include:

- The *size of the system* in question – is this targeting less than 1000 processing elements, a domain in which most of the integrators have experience, or does the system in question exceed, say, 10+ TFlop. If the latter, it is worth mentioning that the current national HECToR procurement rejected the use of integrators at a fairly early stage in the proceedings having considered their capabilities through a series of presentations at SC'2003 in Dallas. While US Integrators certainly have experience in the 1000+ CPU domain, this was not in general the case for their UK counterparts, certainly at the outset of SRIF3. The roll out of several 2000+-core systems is certainly changing this landscape.
- The expected *usage pattern and environment* around the resource – is this being driven by *Capability* or *Capacity* requirements? We would expect, based on some of the considerations above, that integrators would be capable of providing the latter requirement far more effectively than the former.

- The level of *RAS features* expected of the HPC solution. Demanding levels of RAS (say 95+ %) around truly large systems are exceptionally difficult to sustain, particularly in a *Capability* regime when running large jobs with long execution times. Few of the large clusters deployed have truly scalable system software i.e. can provide service to dozens of simultaneous users and have fast, scalable system boot-up, and executable loading capabilities. Assuming such features appear in any contract around the services to be provided, it is extremely unlikely that any integrator would be in the position to accept the risk involved in committing to high levels.

### 3. State of the UK HPC Integration Marketplace

Consideration is given in the remainder of this document to the current status and capability of a number of the leading HPC Integrators within the UK. We provide an in-depth analysis of the four major integrators who supply the academic community and hence are well known to us - Streamline Computing, ClusterVision, OCF and Compusys – plus a less rigorous analysis of some of the other players. Our analysis has been conducted by a variety of mechanisms; given that we know each well, we asked for a response to a set of discussion points following a phone conversation. These points are reproduced in Appendix 1. In this section we provide an overview of each of the four major UK HPC Integrators together with their response to each of the eight discussion points, while Section 4 summarises data from the smaller, less-known organisations.

As mentioned before, there is no doubt that the vendor is becoming more and more of a system integrator, driven by commodity-based Beowulf systems rapidly becoming the systems of choice, at least within the academic community, more information on the variety of solutions sold by the main integrators in this document can be found in section 3.1.

Section 5 provides a summary of the headline issues to be considered between procuring systems from the integrators central to this paper rather than from the traditional Tier-1 companies such as IBM or HP given the undoubted cost differential in the associated products. Important considerations such as where is the dividing line between simply rolling out technology infrastructures provided by toolkits such as OSCAR [6] and ROCKS [7], against providing the level of technical expertise to address the typical RAS requirements associated with high-end solutions for e.g. HECToR?

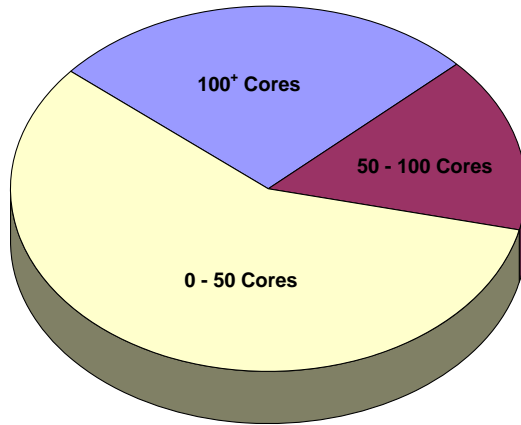
Note that much of the detailed data provided by each integrator makes up the company overviews below. To reiterate these are entirely the companies opinions and do not necessarily reflect STFC's views.

Additional company contact information is included in Appendix 2.

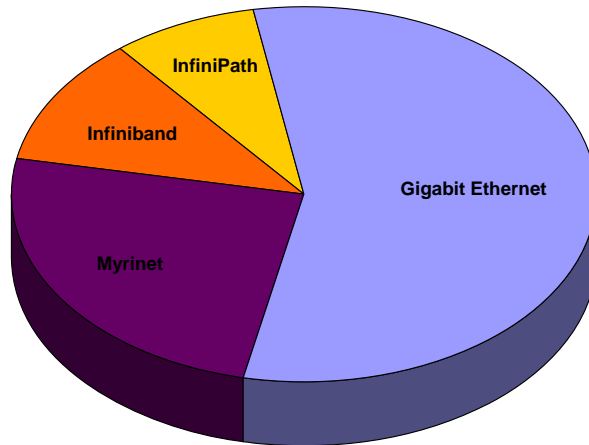
#### 3.1 BRIEF SUMMARY OF THE CLUSTERS DEPLOYED IN 2006.

Commodity-based linux clusters are rapidly becoming the systems of choice for academia. To try to illustrate this, sales information provided by the principle Integrators in section 3 for the 2006 period has been analysed up to give the reader an indication of:

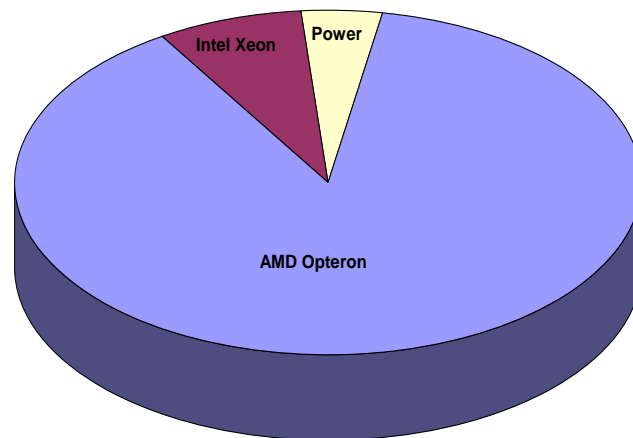
- The Size of the systems being procured;




- Interconnects being deployed;



- Processor Uptake;



- Types of systems being built (rack-based vs blade centres)



From these sales figures, AMD clearly dominated the early part of the 2006 market place, and it wasn't until the release of the Intel Xeon 5100 series (codenamed 'woodcrest'), in particular the 5150/5160 in Q3 that Intel started to make a major impact on sales. Standard rack mount clusters accounted for the majority of the procurements (particularly in the academic sector) with gigabit ethernet taking the lead on the interconnect, but the High performance interconnects (low latency, high bandwidth) accounting for a reasonable percentage of the sales (especially on some of the larger clusters).

It should be noted that this is only a snap-shot picture of the market place, with many of the larger SRIF3 systems unaccounted for. These trends can alter rapidly, especially regarding processor uptake, and are extremely dependent on developments, forthcoming roadmaps and the timing of chipset releases from both AMD and Intel, particularly in light of the next generation multi-core solutions. The uptake of gigabit Ethernet solutions in the HPC market place is probably on the decline, although the next generation Ethernet solutions (10Gbit, TCP/IP off load engines) and cheaper low-latency higher bandwidth interconnects entering the market place may well come into the ascendancy. Multi-core solutions will certainly make a big impact of the size of the clusters sold, but will also add an additional level of complexity to both the integrator (cluster software scalability) and to the end-user (usability, programmability etc.).

### 3.2 CLUSTERVISION

ClusterVision have been trading across Europe since 2002 and ClusterVision Ltd was introduced as a direct subsidiary responsible for the UK & Ireland HPC market in November 2004; as part of their European Expansion. This now includes offices in Munich, Germany and Paris, France in addition to the Head Office in Netherlands and the offices in London, UK. Their growth and success since the start of trading has been impressive.

ClusterVision specialise in the design, implementation and support of large-scale compute, storage and database clusters and are the only Euro-wide cluster company to focus solely on cluster technology and development. They are independent in terms of component supply working with both Tier1 and white box manufacturers and every ClusterVision cluster is delivered as a fully functional turnkey system with all the hardware and software integrated and configured for immediate deployment.

ClusterVision's sales and technical teams have designed, built and supported some of the "largest and complex computational and storage clusters in the UK, Benelux and Germany". Many of the key staff hold PhD's in Applied Science disciplines and have years of experience working with both traditional and clustered supercomputers for scientific research. It is this experience in applied scientific research combined with practical experience of a wide range of supercomputing technologies that "provides insight and understanding of customer's

requirements” and enables ClusterVision to provide tailor-made solutions to meet these requirements.

### **3.2.1 Install Base**

2006 has seen a huge increase in HPC demand and an equivalent increase in Sales for ClusterVision as they passed the three years trading mark and became firmly established across Europe as a specialist provider of Linux Clusters. They have continued their position at the forefront of new clustering technologies and responded to the increasing demand for high speed interconnects with new variants such as Myri-10G and InfiniPath; accelerator board technology and new operating systems including the Microsoft Windows Cluster software.

Reference sites include the University of Surrey, one of the first European sites installed with Myri-10G alongside the DAS3 grid in the Netherlands; billed as the “fastest grid in the world”; utilising Myricom’s Myri-10G technology. The 5 Linux supercomputer clusters will have an aggregate theoretical peak performance of more than 3.8 teraflops. ClusterVision have also installed the ScotGrid node at the University of Glasgow (140 nodes + storage) and were awarded the contract for the National Grid Service 2 (over 550 dual core processors in dual and quad configuration). The service will provide HPC access to the Academic community and includes in its configuration 4 ClearSpeed technology accelerator cards, capable of 50 teraflops performance with only 25 watts power. This procurement co-ordinated through CCLRC was one of several successes for ClusterVision at that Laboratory including a number of Compute and Storage nodes to increase capacity for the Particle Physicists as the UK LCG Tier 1 site.(over 80 nodes supplied).

Working with the British Antarctic Survey engaged ClusterVision in the delivery not only of a turn key cluster solution but also in the power supply and air conditioning, using their own in house expertise, and working with sub contractors to deliver the compute resource and the environment in which it will be housed (41 nodes). At the University of Cambridge, ClusterVision worked with partner Dell Corp. to install, integrate and provide support on 600 dual processor Xeon Woodcrest nodes; with GE and the new InfiniPath interconnect throughout. The Cluster will run the ClusterVisionOS software. A second delivery to the Department of Chemistry at the University of Cambridge with DDR InfiniBand across 47 nodes, coincided with the installation of the HPC Facility.

The installation of the HPC Facility at Cambridge and the installation at Surrey were awarded under the SRIF [14] Collaborative procurement co-ordinated by Heriot Watt University. ClusterVision have also been awarded contracts from The Scottish Association for Marine Science, Keele University and the largest award currently direct to ClusterVision, at the University of Bristol. This large cluster with two associated clusters at the University of Bristol was the first project where ClusterVision partnered with IBM (it will total over 630 compute nodes) and a fore runner to more success for ClusterVision/IBM in more recent rounds of the SRIF procurement at the University of Edinburgh (128 nodes) and the University of Birmingham (256 nodes). Outside of SRIF but maintaining aspects of hardware and software collaboration ClusterVision has partnered with IBM to install the NEMO cluster at Proudman Oceanographic Laboratory (90 nodes) and are working with Dell to deliver and support 120 nodes at King’s College London.

The current division between Academia and industry installations for ClusterVision is approximately 70 – 30. While there has been a greater degree of interest from the private sector in the last 12 months this has been matched and surpassed by large projects across the Public Sector such as the current SRIF collective procurements in which ClusterVision has been especially successful. Inside and outside of the current Heriot Watt collective procurement there has been a shift towards larger procurements in general either across departments or the entire campus to provide resources for all schools, departments and associated Research Bodies. This pooling of resources, encouraged by full economic costing models had meant that the average cluster size has increased to greater than 35 nodes with the largest individual cluster to date being 600 nodes for the University of Cambridge.

### 3.2.2 Company Details and Size

To deliver these systems ClusterVision have also undergone a rapid growth in key staff areas to allow for their continued concentration on core areas of development such as the ClusterVisionOS™ and to achieve their Customers schedules. The company now totals 28 permanent staff across 4 offices. The Head Office is in Amsterdam, with offices in Germany, France and the UK. Their market stretches across Northern Europe and expansion includes new markets in Ireland, at the University of Ireland: Galway and Africa where success has already been achieved, for example with the University of Limpopo. Any of the projects referred to and many others could be counted as a reference and a contact can be made available on approach.

The head count of permanent staff is broken down as:

Function	Number of FTEs
Management	2
Sales	4
Logistics/Projects	3
Cluster Engineers	7
Software Developers	4
Hardware/Productions Engineers	5
Office Administration	3
<b>Total</b>	<b>28</b>

There are also a number of staff available under part time contract.

### 3.2.3 International Presence

ClusterVision's own expansion in the Netherlands, Germany, France, Switzerland and the UK mirrors the growth in high performance computing within these and the surrounding countries in Northern Europe and also the collaboration between them in developing Grid initiatives, even a single European Grid or Resource. Success for ClusterVision has been achieved already in what is one of the largest worldwide Grid Projects to date; The Large Hadron Collider Computing Grid at CERN, Switzerland; supplying over 425 dual processor nodes to the World's largest Particle Physics Laboratory. (Currently the only Tier 2 Vendor to do so from the UK).

### 3.2.4 Company Expertise

ClusterVision's profile within Europe has risen strongly in the last twelve months and has drawn significant interest; ClusterVision have been approached by, engaged with and been successful with several of the Tier 1 suppliers. This has been due to the Company's expertise in their chosen market of cluster computing; an understanding of the components to identify the most resilient on the market, and the development of their own cluster software stack, the ClusterVisionOS™.

The ClusterVisionOS is a Linux based operating system and software environment specifically developed to ease the administration and use of any Linux cluster. It includes all the software required for the effective utilisation of cluster computers.

The ClusterVisionOS has been a key component of recent awards such as the University of Cambridge HPC Facility with partner Dell allowing an expansion of the role of ClusterVision as a system integrator working with Tier 1 partners.

### ***3.2.5 Marketplace***

The cluster market has developed around commodity based components which, coupled with open source software technology like ClusterVisionOS can match the performance and stability of traditional super computers for a fraction of the cost. However Full Economic Costing looks not only at the initial financial outlay but at the surrounding costs, the power requirements and cooling and this has had a large impact on the market from Manufacturers to Users. Every quotation produced by ClusterVision includes information on the power, heat output and weight of the equipment to allow decisions to be made within the FEC criteria.

ClusterVision have maintained their position at the forefront of Cluster technologies, developing working and collaborative relationships with Manufacturers and Users to provide access to information; and to support technical changes ClusterVision continually run their own and customer benchmarks to identify the optimum configuration, either for a specific application or a range of cross campus projects. They also provide remote access to Customers to trial the ClusterVisionOS and to compare configurations. They focus on designing the best available resource for their customers and have produced a series of “firsts” both in the UK and in Europe: first Opteron cluster in the UK – University of Manchester, first production ready InfiniBand cluster in Europe, University of Utrecht; first Opteron Dual core in the UK – CCLRC RAL, also first with a cluster incorporating the ClearSpeed card in the UK and first to deliver IBM x3455 servers incorporating AMD’s Socket F technology at the Universities of Surrey and Bristol.

The flexibility this early adoption pattern allows ClusterVision to maintain a technical competitive edge which in turn can fund future technology growth and customer support.

### ***3.2.6 Relationship to Tier-1 Organisations***

With the growth in the HPC Market both in terms of user numbers and cluster size there has been recognition of the role of a dedicated “System Integrator” by the Tier 1 suppliers; whereby an opportunity is co-ordinated and sold as a turn key solution with on-going collaboration and contact. As this is an area in which ClusterVision excel the relationship with each Tier 1 has also progressed and brought success. In 2006 ClusterVision have worked successfully as both a prime and a sub contractor whilst maintaining their own independence in the marketplace.

### ***3.2.7 Relationship to Other System Integrators***

ClusterVision maintains links with certain established System Integrators with particular software developments and collaborations in mind to support the technology roadmap for example.

### ***3.2.8 Additional Information***

The ClusterVisionOS has become a prime tool for collaboration to develop key features for large scale clusters, such as software driven power management. There are a number of projects under way which will see ClusterVision answer the long term needs for Customers in both Academic and Industrial circles. Therefore ClusterVision’s’ added value remains in the expert supply of tailor made, turn key solutions that can be deployed upon delivery to the science of the user; that they are a single point of contact for hardware and software, and that the ClusterVisionOS is widely acknowledged as the most useful means of managing a cluster.

## ***3.3 COMPUSYS***

Compusys is a privately owned, UK based Computer Solutions Provider and Systems Integrator, and was formed in 1987. Originally an Olivetti PC Systems reseller, Compusys soon began building and assembling its own brand of PC systems, and by 1989, this had



overtaken branded systems sales. Current platforms include Desktop PCs, Workstations, Servers and Storage, and laptops and portables.

The Compusys organisation has continued to grow, but “has retained it’s strong focus on high levels of product quality and customer service”. By the mid 1990’s, Compusys began to realize that broadening its product and service portfolio was essential to the long term survival of the company.

In 1997, Compusys formed its Networking Division, to provide turn-key networking and integration projects for new and existing customers. This provided clients with “a one-stop shop for their entire desktop, networking and back-office requirements”. Within two years, Compusys’ Networking Division had enjoyed significant success with many turn key campuswide LAN and WAN deployments, for further and higher education establishments, the Police, and commercial customers.

In 1999, Compusys formed its’ HPC division, after recognising the future potential of this commodity-based solution. Compusys HPC has grown and evolved “to become one of the leading HPC Integrators in the UK”, and is the UK’s longest established HPC business unit. Compusys also has an e-Business Consulting Division, providing web, content management and e-business solutions to local government, emergency services and academic institutions, which has several significant and prestigious clients.

### ***3.3.1 Install Base***

Compusys HPC have been building, supplying, installing and supporting HPC Linux Clusters since 1999. In this time, the company has built up “experience and expertise in the design, building, deployment and support of Linux Computational Clusters”.

Initial deployments were based on PA RISC and Alpha Processors, as these were the fastest processors available at that time. In their first two years of HPC, Compusys deployed Alpha Clusters to a significant number of Academic Sites, including the University of Liverpool, the University of Leeds, Cranfield University, and many other leading research sites.

During 2000 and 2001, Intel Architecture performance improved, with PIII Xeon systems catching and eventually overtaking the performance of the Alpha based platforms. Compusys again were “at the forefront of HPC Deployments in the UK, with significant deployments taking place at every leading research University in the UK”. Over the last four years, Compusys have continued to provide cluster solutions to UK, European and Inter-continental customers.

The company’s average run rate for the delivery and deployment of systems over the last four years has been around 50 per annum. The majority of these systems have been below 64 Nodes, with around 15-20% of these being above 128 Nodes. The majority of these installations have been to academic institutions, as the market for HPC has been far more mature than the commercial sector. However, Compusys HPC have continued to expand into commercial markets, with successes in the Automotive, Manufacturing, Bio-Informatics and Financial sectors. The commercial markets currently account for around 15% of Compusys HPC.

### ***3.3.2 Company Details and Size***

Compusys is arguably unique in that it provides all its own service, support, maintenance for all systems sold by the company. This includes all desktops, servers, laptops and HPC Cluster Solutions. The company employs its own field engineering, support, helpdesk and administration staff, and so does not rely on any third party in the provision of any of its service delivery. Compusys continues to provide “a broad range of solutions, to an even broader range of customers”. Compusys now employs over 120 staff.

### *3.3.3 International Presence*

Compusys HPC have been providing and supporting cluster solutions outside of the UK for over 5 Years. Their first significant international installation was a 1000 CPU Alpha Cluster for the Moscow Academy of Science, which was installed in 2001. Since then, Compusys have continued to provide systems and solutions to clients across mainland Europe, with installations in many major University and Research sites in Germany and Austria. Compusys HPC have also shipped a number of HPC Clusters to the United States, where their Clusters are incorporated into sophisticated manufacturing systems in Silicon Valley.

Compusys HPC continue to provide solutions across Europe, and are working on projects in Germany, Austria, France, Spain, Italy, Croatia, Switzerland, and other eastern European countries. The company's international coverage also extends to emerging HPC Markets, such as South Africa, Australia and India, and are exploring new opportunities and forging new partnerships.

### *3.3.4 Company Expertise*

Building stable, manageable, high performance clusters from commodity components is a skilled job, and for this Compusys HPC have "both in-house expertise and a proven track record of over 250 successful cluster deployments. This experience continually feeds back into the cluster solutions that we sell, thereby enabling us to raise our standards even further".

Compusys stands apart from most other HPC specialists in the UK, as they are the only provider whose solutions are integrated using 100% in-house resources. No third parties are used in any part of the solution, as Compusys HPC is "a true cluster solutions provider". Their solutions are built on their own systems hardware, manufactured in their own assembly facilities, adjacently located to their HPC Labs. All of their HPC Clusters are fully built, configured, tested and signed off in the HPC Labs prior to shipment to site. This ensures the systems are fully operational on day 1.

Even after installation, Compusys "continues to provide its own service and support to deal with any HPC issue, hardware or software, itself. Compusys HPC support staff are available to determine the source of the issue and to arrange and take the appropriate action, whether that be scheduling a Compusys field engineer to rectify a hardware fault, or remotely logging into the cluster to fix a software problem".

Compusys have direct relationships with all of the leading vendors of HPC cluster products; both hardware and software. For example, they "collaborate with motherboard and systems hardware manufacturers to help them to design suitable HPC platforms. This collaboration is reciprocated when it comes to support, allowing Compusys to efficiently resolve any operational issues that may occur through direct contact with the designers and engineers of the products sold. The direct relationships also support a preferential pricing and commercial support model – essential when fighting against stiff competition from Tier-1 Vendors".

The company's approach to the deployment of their solutions embodies transparency. "Compusys quotes are clear and explicit, detailing all of the components that are needed for a cluster. This approach allows customers to evaluate the offer effectively, and make like for like comparisons. This open approach is followed through to deployment, where a set of agreed sign-off parameters and formalised procedures are used to ensure that the customer get what they expect. Compusys' implementation specification document (ISD) is a key part of these procedures. It specifies how the cluster will be built, the exact software configuration, and the sign off tests that will be run for system acceptance. This reassures both the client and Compusys that the system will be completed to everyone's satisfaction, within an acceptable timeframe".

All Compusys HPC clusters are hot staged in their HPC Labs facility. After the cluster compute platforms leave their production line, they are built into a cluster in the company's Labs. This includes all Cabinet preparation, with the installation of all required communications cabling, mains wiring and power distribution, and the fitting of external 16

Amp mains connectors. All cables are fully labelled at both ends for easy service and system identification. Once fully built, the systems are passed into the hot staging area. It is here that the complete software environment is built and applied to the cluster, and where internal quality assurance sign off tests are performed. “The testing is designed to identify any components that could fail during the first few months of operation, and to ensure the software environment is configured to the agreed specification. The software environment provides all of the required tools to effectively use, schedule, monitor, manage and control the cluster, and incorporates several in-house developed modules and enhancements, only available with a Compusys Cluster”.

The deployment of a cluster at the customer site is simple, as it is already fully operational, meaning that “it takes a matter of hours to deploy”. All sites are surveyed prior to deployment, to make sure the systems will fit in the space allocated, and that they can be handled through the building to their final position.

### **3.3.5 Marketplace**

Compusys see the HPC market continuing to grow, as the technology continues to mature, and new innovations drive performance ever higher. However, the issue of Full Economic Costing is now playing an increasing part in decision-making process, as the factors that affect FEC are becoming more visible, and sites are now reaching their capacity to provide the space, mains power and air conditioning required to run a super computer cluster.

Compusys are already seeing the cost of ownership benefits of Dual Core technology putting such solutions high on the shopping list of prospective buyers. Compusys own cost benefit analysis shows potential savings of up to 40% per annum in mains powers and air conditioning running costs with a Dual Core Processor cluster deployment. On a cluster with 500 CPU Cores, this can be a significant cost saving, resulting in a reduced Full Economic Cost.

Compusys are partners with AMD, and were launch partners for the Dual Core processors. This gave Compusys early access to the Dual Core parts, for testing and benchmarking ahead of the official launch. Compusys are already working on Dual Core deployments, with the first installation of this technology due to take place shortly. The company are also heavily involved with other leading HPC technology companies, and are looking forward to the launch of many other new technologies (most of which are currently under NDA), which will have a significant impact on the performance and ownership costs of commodity cluster solutions.

Compusys are also leading the drive to “take the mature products and technologies offered to the academic markets, to the commercial sector, with the drive for greater efficiency and value steering users away from large shared memory systems”.

Finally, the company believes that “commodity clusters will be the utility resource that everyone dreams of. Developments are already underway to take standard commercial business applications, and deploy them on commodity clusters, to save cost, increase performance and enable true scale-out computing. In the future, a cluster will just be a resource for processing, with virtual machines running applications for an entire organization. As demands from individual instances increase, resources can be dynamically allocated to those VMs that need it. The days of deploying a new server each time you deploy a new application will be over. With Grid Tools linking each cluster into a pool, most organizations will have the flexibility to cope with all their computing demands, with less hardware, being more flexible, and saving money”.

Compusys aim to play a leading part in bringing the above technologies to market, and is working closely with Global HPC vendors, as they prepare to launch their new products over the next 12 months.

### ***3.3.6 Relationship to Tier-1 Organisations***

Compusys has a neutral stance on the position of Tier-1 vendors in the HPC marketplace, especially in the UK and Europe. Although Compusys have competed head to head with every Tier-1 vendor at some point either currently, or in the past, they believe that their levels of experience and expertise, plus the fact that all their services are provided by in-house staff, ensures that they can match or exceed the levels of service provided by “the so-called big names”.

Compusys is part of a group turning over in excess of \$100 Million a year. Compusys are more than able to finance contracts of well in excess of £2 million and have done so several times in the past. Tier-1 support is only an issue, if a particular Tier-1 vendor has decided to offer products at prices below cost.

It is also a little known fact, that most Tier-1 vendors in the HPC space do not actually manufacture their own HPC products. For example, IBM use MSI to make their Dual AMD platforms, as Compusys have used and deployed this same chassis and motherboard, branded MSI, not IBM. SUN offers a 1U Dual AMD and a 3U Quad AMD server platform for its HPC solutions-these systems are made for SUN by Newisys, and are available from Compusys too.

The company believes that most of the time, Tier-1 vendors bring a low cost price for the computer hardware, and the comfort factor of a known name. However, when “you strip back the offerings from the Tier-1 vendors, and find that the hardware is not truly made from the vendor, and that the integration of the cluster is actually being carried out by a smaller specialist, because the Tier-1 vendor doesn’t have the knowledge or expertise to do it themselves, you have to ask where the value comes from”.

### ***3.3.7 Relationship to Other System Integrators***

Compusys has and maintains contacts with the other HPC specialists in the UK, including Streamline and OCF. Compusys and the other vendors meet at various industry gatherings throughout the year, and obviously follow each others progress. With “the high degree of skill and understanding required to install and support Linux HPC Clusters”, Compusys see themselves as specialist’s contractors to the traditional Systems Integrators, as the steep ramp-up required for a new SI to enter the market would be hugely expensive, and without a track record, difficult to market. Compusys have already done work for Computacenter, and for HPC and Clustering projects.

## ***3.4 OCF PLC***

OCF is a specialist independent reseller of high performance technical compute (HPTC), high performance visualisation (HPV) solutions and Enterprise Computing Infrastructure for Storage and Server Technologies. The company “continues to evolve its skills in accordance with technological advances and breakthroughs in order to remain at the cutting edge of HPTC and HPV developments and Infrastructure provision”. OCF has forged a strong relationship with IBM and the majority of its solutions are now based upon IBM hardware. These solutions range from individual workstations and Servers to large bespoke enterprise systems providing “organisations with maximum compute power and complete data management facilities by adopting a collaborative approach to solving customers IT challenges and working closely with a number of technology partners” in the following areas:

- Development and Deployment of Server Infrastructure
- Complete Data Management Solutions Providing automated ILM (Information Lifecycle Management) allowing organisations to attain the maximum return from all assets.
- High performance visualisation workstations
- High performance immersive group visualisation environments

- The whole range of compute servers, from individual Linux servers, to clustered Linux based solutions, through to high end SMP servers
- High performance computer interconnect technology
- Independent software vendors

### 3.4.1 *Install Base*

The division between public/private installations has historically been 70/30%. This ratio has recently been moving towards a 50:50 split due to two factors:

- The successful deployment of some large private sector contracts
- The delays in the implementation of the public sector SRIF programme.

OCF does keep some private installations confidential for commercial reasons.

A number of customer sites would “be happy to act as a reference” for the UK and the associated market developments – please contact OCF for named contacts.

Perhaps of most current relevance would be the work done by OCF at Southampton University with the IRIDIS cluster. From the initial installation of 330 Opteron cores linked by a Gigabit Ethernet network almost three years ago, the system has been increased to over 1000 Opteron cores (with a section connected by a Myrinet high performance interconnect), over 700GB of RAM, 25TB of storage and a full remote management solution. As far as high performance storage is concerned, a recent GPFS installation by OCF has yielded aggregate performance figures of over 1500MB/s with Infiniband, writing 16GB files with 1MB block sizes, a 33X speedup over the previous “Traditional Cluster Storage” architecture. A recent win by OCF will provide over 4GB/s of throughput to over 100TB of storage.

### 3.4.2 *Company Details and Size*

The overall size of the company is shown in the Table below:

<b>Function</b>	<b>No. of FTEs</b>
Managerial/Supervisory	1
Sales	6
Service	2
Operational/Administrative	3
Sub-contractors	
Consultants Design Experts	
Design Experts	4
Technical (Implementation, Roll out, Support & Maintain)	5
Others:	
<b>Total:</b>	<b>21</b>

OCF’s Technical team comprises a team of nine. This team has “expertise to configure, supply and install large server and storage solution infrastructure”. In addition the company has recruited and developed software engineering skills, giving OCF the ability “to assist it’s customers in such matters as Solution Software design, integration of Management software and Benchmarking Code for HPC Servers”. It is interesting to note that its sales team is composed almost entirely of technical specialists who have transitioned into a sales environment.

In-house warehousing, and engineering workshops at OCF Sheffield Head Quarters enable OCF’s technicians to be able to test all equipment prior to customer shipment, and to evaluate and benchmark new technologies as they are released. The company also provide a customer help desk and support facility (available Monday – Friday 08:00 – 18:00 excluding public holidays). OCF build and support procedures are governed as part of their ISO9000 accreditation and IBM reseller partnership agreement. Strict internal procedures are in place

to ensure that all IT equipment delivered to OCF plc is fully tested prior to shipment to their customers.

OCF provides maintenance for IBM equipment through IBM Global Services, with OCF acting as the first point of contact for complete projects. Maintenance provided can be customised from next day return to base, through to seven day, 24 hour on-site support. Options include service upgrades for in-warranty machines; extended maintenance for post-warranty machines; experienced technicians; and extensive parts distribution.

OCF has a dedicated Support Hotline to which all calls (technical queries, delivery information, returns, repairs etc.) are logged through the Hotline onto an electronic Call Management system. This in turn is set to escalate any calls that exceed the agreed time span. Problems relating to hardware are relayed to the manufacturer within 2 hours of receipt, and the Support Hotline monitors progress.

OCF has been specialising in HPC and HPV perhaps longer than any other UK integrator. This has enabled it to develop an unrivalled ecosystem of HPC and HPV partners, ensuring that technical challenges can be referred to technical specialists outside its own employee base. IBM alone has a team of over 50 people in the UK dedicated to HPC and HPV; by far the greatest commitment to the sector of any vendor. As a Premier Partner, OCF enjoys unrivalled access to this team and further through to the worldwide resources contained within IBM.

### ***3.4.3 Geographical Outreach***

OCF's primary market area is the United Kingdom and Southern Ireland. In the past OCF have performed contracts in Scandinavia and Holland and has acted as a European business partner for a Life Sciences ISV with business currently quoted in Italy, Switzerland and France. OCF does not have any overseas locations and all contracts are serviced with staff based in their Sheffield office.

### ***3.4.4 Company Expertise***

OCF has supplied computational clusters to a wide range of academia and corporate clients over many years, and as a result has implemented a wide range of CPU architectures: x86/x86-64, Itanium, Power/PowerPC and Alpha utilising nodes from a wide range of Tier1 vendors and white box manufacturers. OCF has also integrated clusters across a wide variety of interconnects such as Infiniband, InfiniPath, Myrinet2000, Dolphin Wulfskit and varying Ethernet technologies including Gigabit and FastEthernet, 10G, RDMA and TOEs whilst implementing management tools such as CSM, IBM Director, GridEngine, Tivoli Storage Manager and Scali Manage in order to provide a robust, easy to use operating environment. OCF's engineers are certified to an expert-level in both Linux and AIX and have been early partners and implementers of Microsoft for Windows CCS. Many of OCF's solutions are delivered on IBM hardware and software environments, utilising the best of IBM's market leading technology and support. OCF has implemented more IBM-based technical computing and cluster solutions than any other company (other than IBM itself) in Europe. However, OCF is not simply a 'cluster company' – it specialises in the whole spectrum of HPC and HPV and includes large SMP solutions and other traditional alternatives to clustered solutions in its portfolio.

OCF has a strong dedicated visualisation team and has done extensive work with IBM on DeepView, IBM's advanced parallel and distributed visualisation technology, to provide immersive multi-dimensional environments and deliver high performance collaboration solutions to remote users. OCF is the only company to partner with IBM to demonstrate DeepView and has been involved with this technology since before its release to the market. As a result, OCF is uniquely able to implement and support visualisation solutions for the UK HPC market.

OCF has developed advanced capabilities to enable the delivery of High Performance Storage solution. It provides extremely fast, enterprise class data delivery with low cost commodity components and the IBM GPFS parallel file system. OCF specialise in integrating these storage environments into the customer's existing infrastructure, extending what has traditionally been considered a "cluster-only" technology across the organisation it serves. These storage technologies are backed by IBM's world class support organisation and have been successfully implemented in a number of commercial and academic customers. Additionally OCF is currently investigating the implementation of CLOS networks for large cluster environments.

OCF is full accredited to the ISO 9000 standard and has a clearly defined methodology for the design and delivery of all of its HPC and HPV solutions. Wherever possible this includes the pre-build of the complete solutions at its workshops in Sheffield, with benchmarks and acceptance tests completed to identify any problems prior to installation at the customers premises.

OCF believe that the first step to getting the optimum performance from any compute resource is to optimise the design of the hardware and operating environments, concentrating it's expertise on the infrastructure and operating system framework. For code-level performance tuning, OCF complements its own expanding capabilities with those of its extensive partner ecosystem, including both commercial and academic organisations, to ensure that its customers have access to the widest possible 'best of breed' expertise.

### ***3.4.5 Marketplace***

One of the current hot topics in the HPC world is the desire by funding bodies and corporate customers alike to make investment decisions based upon Full Economic Costing (FEC) matrices. This is becoming more important as the environmental costs of compute equipment are becoming more understood and the costs of power increase. Technologies such as Grid computing, server and storage virtualisation, advanced power and cooling technologies are all a response to this challenge in the world of Commercial computing.

However, OCF has not yet seen any significant moves within UK academic HPTC market other than in general statements of direction. Whilst public sector customers are increasingly identifying running costs as a factor to be taken into account in their decision making it has not yet become apparent to OCF that such considerations override the traditional drive from the maximum compute power for the minimum cost. Private sector customers tend to be more sophisticated in their capital budgeting processes and take into account FEC as a matter of process.

As regards OCF's own internal investments, they are continuing to increase the peoples skills in OCF's core competency areas i.e. infrastructure and operating systems. Taking the solution up the food chain from there requires the creation of strong mutually beneficial relationships – hence its continued focus upon a partner based approach and evolving relationships with such specialist organisations as Pathscale, Mellanox, Voltaire, IBM, Allinea, Arup, Fluent, Force 10, Level 5, etc. It is convinced that the only sensible model is to establish a strong portfolio of industry leading technology partners and to manage these relationships for the benefit of its customers – this is a true Systems Integration model.

There is a strong body of opinion that future processor development is moving away from satisfying the needs of HPC users and that the HPC market needs to look to technical developments in the areas of co-processors, FPGAs, Cell processors etc. OCF has already supplied Cell processor based systems and has a working relationship with Nallatech to incorporate its FPGA technology into its HPC solution portfolio. It has also developed a relationship with Clearspeed in order to allow its customers to take advantage of the performance acceleration provided by its technology.

The number of HPC users has grown exponentially over the last ten years and it has moved outside the realm of traditional 'expert users'. This has led to a requirement for systems to be far more user friendly and it is OCF's view that supported toolsets become evermore

important, preferably supported by organisations with the credibility and critical mass to provide support and development over the long term. The recent launch by Microsoft of its Compute Cluster software will help at the lower end of the demand scale, whilst such developments as portals for access to large centrally managed systems may well assist the user with higher demands.

Furthermore it is evident that as the need for Petascale computing capability expands, the scale of such systems will likely best be met by the implementation of robust and flexible grid infrastructures. OCF continues to work with its HPC partners to develop solutions to the many problems still facing such complex systems.

### ***3.4.6 Relationship to Tier-1 Organisations***

In responding to the question – “Does your company have the strength and depth to build high-performance, scalable systems which can support Tera & Petascale solutions in the not-too-distant future – Are you able to provide this independently” – OCF believes that the only honest answer is “yes we can build them but not independently”. It does not feel able to address the relative abilities of UK and US integrators.

OCF’s view is that “the only sensible way forward for commodity based Tera & Petascale computing is to leverage the financial strength of a Tier-1 vendor”. OCF have backed that view by developing a very strong relationship with IBM. Risk, reliability and long-term support are vitally important in most areas of business, and they are critical in large scale computing environments. From an integration perspective, the major problem of fault identification and subsequent resolution means that pre-configured supported solutions are very much in vogue with the Tier-1’s, with IBM’s 1350 and 1600 clusters and HP’s XC. The downside of such solutions is that they are significantly more expensive and customers face difficult decisions, especially those funded by the public purse!

By far the majority of systems are purchased to provide ‘capacity’ rather than ‘capability’. Indeed this seems to be the case for all the current SRIF procurements.

In the medium term OCF is not convinced that today’s delivery models will be appropriate. How users obtain their compute cycles may well change radically if there are advances in middleware, bandwidth etc. There has to come a time when Vendors question whether delivering kit is an economic model, or whether the delivery of cycles is more appropriate. IBM currently operate a very crude version of this (On Demand Computing) and the vision of them simply taking servers off the production line and into huge server farms serviced on demand may not be fanciful.

### ***3.4.7 Relationship to Other System Integrators***

This needs to be looked at in two ways. Firstly, with regard to other HPC System Integrators, OCF enjoys a level of healthy competition and differentiates itself by being the only one of the four major Integrators to commit its server and storage solution components to primarily one Vendor – IBM.

With regard to the large traditional System Integrators, OCF believes that the role of these organisations depends to a large degree on the business model being adopted by the provider of the compute cycles. Such organisations are good at billing and on going facilities management. Hence CSC’s tie up with SGI at CESA. They tend to be exceptionally expensive having had a reasonably easy life looking after corporate networks and charging absolute fortunes. OCF believes that a relationship with such organisations will be sensible and currently is in discussions with some of these to act as their HPC expertise provider.

### ***3.4.8 Added Value***

OCF has a long history of proposing innovative and flexible collaboration schemes for public sector procurements. In particular for the current SRIF procurements a detailed package of



measures has been designed, in conjunction with its technology partners, to further evidence OCF's commitment to and support of UK Academic Research. These include access to new and exciting technologies on a no cost basis, funded placements of engineering resource to assist in the management of large systems, access to the research departments of our technology partners (particularly IBM), free quarterly healthcheck visits to ensure continuing efficiency and a whole host of other initiatives.

Whilst OCF believe that the collaborative opportunities that it offers clearly differentiate it from its Integrator peers, it is not yet clear whether such initiatives have a great deal of influence upon procurement decisions when, as with Full Economic Cost model discussed above, the emphasis continues to be the most kit for the least initial cost. This has been one of the more disappointing aspects of the current SRIF3 procurement exercise.

### **3.5 STREAMLINE COMPUTING**

Streamline Computing endeavours to provide state-of-the-art, commodity-based HPC systems to solve its clients' technical and scientific computing problems with the best performance and quality possible.

Streamline Computing is a trading division of Concurrent Thinking Ltd. Alongside its sister division Allinea Software, an HPC tools provider. Concurrent Thinking Ltd was recently included in the 100 fastest growing technology companies in the UK (<http://www.fasttrack.co.uk/home/htm>) and also achieved ISO 9002 accreditation in 2006. Concurrent Thinking Ltd. Has also just closed on a 2<sup>nd</sup> round institutional investment of approximately £1.7M (see <http://www.hpcwire.com/hpc/1032136.html>).

Since spinning out of the HPC facilities of Oxford and Warwick universities, eight of Streamline Computing's Linux supercomputing clusters have been included in the Top500 list at the time of their delivery. Of particular note is a 1024 CPU Sun Opteron cluster to the University of Nottingham (Spring of 2005) together with a multi-site GRID facility in the North West region.

Streamline Computing engineers HPC solutions from commodity components whether sourced internally through Streamline's own supply chain or in partnership with Tier1 and Tier2 vendors, both delivering and supporting these solutions. It has provided over 300 HPC Linux clusters to blue chip companies both nationally and internationally, as well as to nearly all the top-ranking research-led academic institutions in the UK. Streamline delivers the solution stack through the configuration of the front-end node(s), an HPC-optimised software stack, and a customised (and fully customisable) cluster management appliance. The latest minor revision of this management appliance was released in 2006 although a major revision of this will be launched in Q1 2007. The appliance extends the control, management and (dynamic) re-purposing of Linux clusters, treating the provision of compute nodes as canonical. Streamline is currently considering how this solution may be further productised, given the incipient maturity of commoditised server platforms, to alleviate much of the burden of administration and management – a significant contributor to the overall cost of commodity clusters.

Since the last revision of this document, Streamline has continued to execute its strategic plan while increasing its business and customer base. In particular, it has achieved record sales in the current financial year:

1. by enabling and supporting partners in the provision of a Streamline-qualified HPC software stack – this enables Streamline to address large-value procurements with OEM and reseller partners so as to collectively obviate technical and commercial risks;
2. through the growth in its own 'turn-key' Streamline Linux Cluster solutions, especially in the commercial arena;

3. delivering systems based on next generation low-latency interconnects and parallel file systems.

Streamline has a proven track record in tuning operating systems and specialised MPI and scheduling software layers, and has developed expertise in specific applications in the area of engineering, simulation, life science and energy sectors. Streamline-Computing has considerable experience in networking solutions (commodity Ethernet or low latency alternative), storage and file-systems, visualisation and integration of these within the operational environment and policies of the end-user. Streamline maintains a thriving research and development group which actively pursues new technologies and market opportunities. It executes this R&D with internal investment (as witnessed through the development of the Allinea trading division), close involvement with the Score consortium in Japan, and through external funding via the DTI (at a National or Regional level). Of particular note is

1. A Grant for Research and Development which has been investigating ‘Fault Tolerant’ Parallel Computing solutions.
2. The Inter Enterprise Computing project BROADEN (Business Resource Optimisation for Aftermarket and Design on Engineering Networks) aims to build an internal GRID at Rolls Royce PLC as a proving ground for utilising web/grid service technology to fully exploit available IT resources.<sup>3</sup>

**Streamline partners:** Streamline Computing partners, not only with Tier1 vendors, but with a number of technology organisations in the HPC sector to deliver solutions. These range across the complete stack from microprocessor through networking, parallel file-systems, to application development and integration; the company works closely with AMD and Intel and their respective platform providers in order to best deliver these “core” technologies to market. The company is an authorised reseller for most of the software companies providing compiler technology and tools to the HPC market, and if purchased as a component within a Streamline solution, these software products come preinstalled and ready for use and integrated with the Allinea tools for debugging and optimisation.

The majority of clusters delivered by Streamline offer Gigabit Ethernet as the networking solution which meet the demands of the majority of commercial ISV codes as well as meeting the primary requirements of academia in respect of maximising CPU count. In this respect it has worked closely with Nortel in providing multiple stacked configurations for large cluster solutions as well as providing 10 Gigabit Ethernet core capabilities for I/O. In terms of high-performance low-latency networking, Streamline has recently seen an increase in solutions based on Infiniband (and more recently InfiniPath) as well as continued demand for Myrinet.

Streamline clusters (whether through its own OEM channel or Tier1 partners) are shipped with the Streamline Cluster Management Appliance where IPMI or proprietary lights-out management interfaces are available. It is delivered as an appliance based on a 1U form factor commodity server, and ostensibly provides a web-services interface (in version 2) condensing much of the complexity of Linux Clusters to an intuitive set of tools based around management, monitoring and imaging – alleviating much of the administrative burden inherent in Linux clusters. In particular the new features comprise:

- Improved Script and Task Management:
  - Managed repository of preloaded and user-defined scripts
  - Scheduled & manual execution of tasks across a cluster
- Metric Collection & Display:
  - Multiple data collection methods, including ganglia, SNMP ....
  - New Grid/Cluster/rack view to highlight problems at a glance
  - Highly configurable actions on metric threshold breaches

---

<sup>3</sup> See for example <http://www.ngp.org.sg/gridasia/2006/uksingapore/slides/UKSingapore-GrahamHesketh.pdf>

- Hardware configuration & control:
  - Database of hardware, configuration and historical data
  - Support for blade servers and virtual machines
  - Out-of-band management features for compatible nodes
- Image management:
  - Improved Image collection and deployment
  - Provides full audit-trail of deployed images for client nodes
  - Support for RAMdisk and diagnostic boot mode

**Streamline's software stack:** Streamline Computing installs and configures a custom combination of open-source and licensed software on all systems as part of its package. Most systems ship with SuSE Professional or Enterprise Linux or RedHat Enterprise Server Linux (or mixtures thereof) where for example particular support of an ISV package is required. Other Linux distributions such as Scientific Linux, Rocks or older RedHat distributions, are also available with options to support re-imaging of clusters for specific environments as well as the ability to support RAM disk in both disk and diskless clients. All distributions are fully installed and patched up to the prevailing levels, and each system can be configured for automatic operating system updates using Yum or Yast.

For parallel applications using MPI, Streamline has installed most of the numerous choices available, and is now providing these around a "modules" environment. Streamline also has unique access to an open source version of MPI called SCore, developed and widely used in Japan. The SCore parallel computing environment is installed as standard on systems shipped by Streamline. When installed and configured, SCore provides a parallel computing environment, including a custom MPI layer, that can provide significant performance increases for parallel applications compiled to use SCore's drivers and subsystems. SCore provides multiple network support allowing an application to use different networks such as Myrinet and Gigabit Ethernet during execution without any user designation. SCore also provides deadlock detection, fault tolerance with pre-emptive check-pointing, parallel process migration and flexible job distribution including gang and batch scheduling. In addition SCore provides a general parallel shell which provides a powerful system administration tool for cluster-wide operations. Further development of SCore is in process in collaboration with the SCore consortium (see <http://www.pccluster.org/>).

Streamline also has a wide range of expertise in Distributed Resource Management (DRM) and scheduling software and in setting up load distribution systems. Most Streamline clusters ship with SGE, but also support clusters using Torque as well as LSF where appropriate.

### 3.5.1 Install Base

*An idea of the relative number of "small" (32-64) systems compared to larger (128+) Machines*

Whilst a clear majority of systems installed over the last few years fall in the small to medium category, more recently entries into the larger compute market are becoming more commonplace (64 to 128 nodes), with some deployments making the TOP 500 rankings. The largest machine that Streamline has installed include 1000+ processor system for Nottingham University and a middle-eastern Oil company (both with Sun Microsystems). In the industrial space systems tend to be within the 32-128 processor range, primarily constrained by the ISV applications traditionally used, although a 256 processor system has recently been shipped by Streamline to a large engineering company. The split between customers is approx 60% HEI/research and 40% commercial.

### 3.5.2 Company Details and Size

Some of Streamline's company information is given in the table below; of a total of 37 staff within Concurrent Thinking Ltd, 21 reside in the Streamline Division. Details of the skills

base within the organisation are captured above under the sub-heading “Streamline’s Software Stack”.

<b>Function</b>	<b>Streamline Division</b>	<b>Concurrent Thinking Limited</b>
Pre-sales	6	9
Management	2	7
Technical (hardware)	5	7
Technical (software)	8	14
<b>Number of Staff Employed</b>	<b>21</b>	<b>37</b>

### ***3.5.3 European Presence***

Streamline have installed clusters in France, Germany, (as well as the Middle East, USA and Canada), although its primary focus has been within the UK. The approach has been to provide software and support skills to local cluster builders. In servicing Streamline’s large commercial customers, Streamline have worked closely with partners in Europe and in the USA as hardware build partners. Streamline also have provided the same service for Tier 1 partners in the Middle East.

In the Concurrent Thinking business and growth plans, coverage will be increased with offices in the Europe, the US and the Far East.

### ***3.5.4 Company Expertise***

Streamline engineers have skills in a number of areas individually and “combine these skills to solve more difficult system level issues”. The company’s staff has a good solid grounding in Linux but bring skills in DRM and GRID, parallel computing, cluster management, system monitoring and high-end storage. Company staff members have PhDs in Parallel Computing, Computational Physics, Computational Biology, Numerical Linear Algebra, and Computational Fluid Dynamics. Recently it has invested in an ISV partner team with over 20 years technical experience in CFD and CAE codes who will work closely with industrial end-users. It has also recently expanded its support infrastructure with new hires who are responsible for managing support queues, hardware replacements and benchmarking.

### ***3.5.5 Marketplace***

#### **Streamline / Concurrent Thinking Strategy**

As cluster sizes grow but commodity hardware becomes cheaper, the “value” in the market will be the ability to make sure that clusters work efficiently “hitting-the-ground-running” for a broad range of applications and function in a more scalable manner operating seamlessly across subsystems including file-systems and visualization through the lifetime of the resource.

To increase the accessibility to HPC, system management and monitoring capabilities will therefore become more important and proffer a more service based approach to HPC. The ability to provide a high level of support, not only on a system level, but also on an application level, will help differentiate the ‘box-shifters’ from the serious HPC oriented companies, ideally separating out the commodity aspects of the system to a value integration.

Previous investment has been made in building this knowledge base within the company around this strategic view. In particular it is of worth noting

- Collaboration with the Score consortium in Japan and the provision of the MPI Score environment widely used in the UK; this technology offers the best latency/bandwidth profile for commodity Ethernet.
- The spin-out of DDT and OPT within Allinea Software now providing software tools critical to parallel and cluster computing in general Linux cluster solutions; this technology is becoming more increasingly important as the amount of parallelism increased within multi-core and embedded technologies
- Continued research initiatives supported by the DTI for high availability computing, GRIDs and visualization involving major industrial customers;
- The development of cluster appliance products towards the general provisioning of optimized parallel software environments for cluster computing as well as the management and control of these systems.

Streamline continues to invest in such capabilities and is keen to work with partners and customers to ensure our cluster technologies maintain Streamline's position as the leading UK integrator in this field. We also value such collaboration as a means to not only develop the right products but also note that having partnerships can ensure that Streamline can deliver these products in a timely manner to the market – in a sense getting the product right. Primary R&D areas for Streamline concentrate on:

- Distributed and Parallel File Systems
- Optimized MPI environments
- Cluster Installation Management and Provisioning
- Parallel Job Control and GRID environments
- Software Tools and Code Optimization
- Distributed, Remote Visualization of Very Large Computational Models

Following the recent investment round, Concurrent Thinking will be focusing on bringing these technologies to marketplace within its own cluster architectures, as well as those of its partners.

### **Trends in the Marketplace**

A significant proportion of cluster sales have been, and continue to be, based on Gigabit technology, or a mix of Gigabit and High Speed Interconnect reflecting the state of parallel computing. For this reason, Streamline has invested significantly in R&D relating to Score providing maximum efficiency on Gigabit. Streamline observe very few customers who run a large proportion of capability jobs on their clusters; the majority run mildly parallel applications on up to 32 processors. Benchmarks has shown that SCore outperforms MPICH and LAM by over 20% on standard test cases using Computational Chemistry and Computational Fluid Dynamics codes using 32 processors and Gigabit. In certain cases, where significant latency intolerant applications exist performance improvements can be even higher. As mentioned previously, Streamline has extended the performance of Score and now provides around 15us latency as measured on the Intel IMB benchmark and over 200 Mbytes/sec on Send/Recv. We are working closely with Score and Ethernet switch vendors to further improve not only this point-to-point performance but also improve collective operations within a clustered-SMP environment.

Notwithstanding this with the new wave of multi-core solutions, with 4 and 8-socket servers, Streamline is witnessing an increased requirement for clusters of modest SMPs connected by Myrinet 10G, Infiniband/Infinipath or Quadrics interconnects. There has been some interest in 10 Gigabit Ethernet as a cluster interconnect, although this technology remains incipient.

While Linux still accounts for the vast majority of business, Streamline has delivered Solaris based clusters, working closely with Sun Microsystems and is also witnessing strong interest in Microsoft Compute Cluster Edition. Streamline representation has been strong in many Microsoft-led UK events over the past year and is beginning to see embryonic Microsoft CCE clusters appearing in dual-boot configurations.

## **Challenges**

It should be noted that Streamline has had exposure to, and has overcome, some very interesting technical and support challenges relating to large cluster systems (buffer overflows on switches and NICs; processor timing problems; interaction between job schedulers and Linux modules etc.) and it is this level of systematic support that is critical to many operations.

A potential major challenge to SI is whether all public procurement decisions relating to clusters become a question of price, then companies like Streamline will be unable to sustain the skill-base needed to resolve such complex problems. Similarly, records demonstrate that Streamline currently provide a significant level of customer support (all of their support queries are logged in a company database) yet that the delivery of this level of service is clearly not profitable for academic customers on a tight budget. As a growing company, Streamline see this as an investment for the future. However if University procurement procedures do not provide the mechanisms for companies like Streamline to demonstrate their value, then the provision of this level of service will become unviable.

### ***3.5.6 Relationship to Tier-1 Organisations***

Very large systems require the financial strength of a Tier-1 vendor to execute larger contracts but need the specialist skills of companies like Streamline to provide the knowledge and skills to build and support the systems. Relationships at all levels with Tier1 vendors are thus critical as a success factor for future growth.

Streamline have built systems based on the major Tier1 hardware providers. They have also provided support for a number of other vendors, both in UK (for systems shipped from the USA) and in Europe M. East and USA, where Streamline partners build and Streamline support. Experience so far demonstrates that Streamline's technical staff has more than equal skills to any other integrator globally, and can compete effectively if these skills can be channelled profitably.

### ***3.5.7 Relationship to Other System Integrators***

Streamline has developed relationships with large scale System Integrators where commercial and industrial end-users have out-sourced supply to these organisations. This arrangement meets the market demand for delivery of the commodity servers in the most efficient manner but enables Streamline to transfer its value in providing the solution. Streamline is continuing to develop further relationships in this area as part of a definitive strategy aimed primarily at the industrial and commercial HPTC sectors and where appropriate in the academic sector too.

### ***3.5.8 Additional Information***

“With a strong UK based team with deep technical skills” Streamline has much to offer UK and European Customers and as a company “can keep a UK technology flag flying in a global market”. Whilst establishing their own Tier-1 relationships, Streamline is grateful for the support from organisations such as CCLRC and other UK government organisations in allowing Streamline to demonstrate its value and the ability to succeed in this market and ultimately benefit UK PLC with exports and technology from a strong UK base. In particular, Streamline wishes to take this opportunity to thank the DTI and its project partners for their support in the execution of R&D projects and looks forward to continued involvement in similar region, national and international initiatives.

“Ultimately, hardware is the most commodity and in that sense the least important decision relating to the purchase of HPC clusters”. As such, the future strategy of Concurrent Thinking is focussed on the delivery know-how in a manner that is independent of the actual hardware providing user choice, best value-for-money and an appropriate level of risk mitigation.

## 4. UK HPC Integrators II

Having considered the current status and capability of the four leading UK HPC Integrators – Streamline Computing, ClusterVision, OCF and Compusys – in section 3, we now provide a far briefer overview of some of the other, less recognised players. Note that much of the information here has been obtained from the organisations web pages and does not map naturally onto the discussion points raised with each of the integrators of section 3. Companies considered below include Cambridge Online Systems, Linux Networx (USA-based, but with personnel situated in the UK), Western Scientific, SCC, Silicon Mechanics (USA-based, but branching out into European solutions) and Workstations UK. In only two cases, Cambridge Online Systems and Silicon Mechanics, did we have the opportunity to raise and discuss the points noted at Appendix 1.

### 4.1 CAMBRIDGE ONLINE SYSTEMS LTD

#### Response from Cambridge Online Systems Ltd ([www.cambridgeonline.net](http://www.cambridgeonline.net))

##### 4.1.1 Install Base

The customer base is within the UK and made up of Research (40%) and Higher Education establishments (30%) and commercial organisations (30%).

The main architectures installed are HP Proliant (with Intel Xeon and AMD Opteron processors) and HP Integrity (Itanium2) with Linux, and legacy HP (Digital) Alpha with Tru64 Unix. Some customers are trialling Microsoft CCS (Compute Cluster Server) for Windows. The majority of systems are either clusters or compute farms with bladeservers being the most popular form factor. The split between ‘small’ (32-64) systems and larger (128) systems is roughly 75:25.

For integer heavy computations, HP Servers with dual core AMD Opteron CPUs have proven popular, having high performance for a low price. With the introduction of the new HP c-Class BladeSystems combining latest dual and planned multi-core technology from Intel or AMD, together with improved Linux management software, the stage is set for users who previously had to share HPC resource, to afford dedicated systems. For users with advanced visualisation and floating-point calculation requirements requiring full 64bit computing, Itanium2 comes of age with the release of Intel Montecito based HP Integrity systems in September 2006.

Over the past three years, noticeable trends have been:-

- shift from proprietary Unix to Linux
- increasing use of industry-standard, ‘commodity’ computing systems
- heterogeneous, mixed-vendor Linux environments
- requirement for performant file system, such as Lustre/HP Scalable File Share.

##### 4.1.2 Company Details and Size

Cambridge Online was established in 1978. Their Applied Technology Group specialises in the provision, development and support of IT infrastructure, covering computer systems, storage, networking and telecommunications. The company partners with industry leading technology vendors and hold ISO9001:2000 quality accreditation, and with a particular focus upon HPC systems, delivers “a highly technical and consultative approach to meeting solution requirements”.

Total staffing of 60 employees is split by sales - 5, technical support/engineers - 15, software development - 30 and management and administration – 10. More than 250 customers are served.

Relevant customer names include the Wellcome Trust Sanger Institute, European Bioinformatics Institute, University of Cambridge, University of East Anglia, Cranfield University and the Medical Research Council. Reference details can be provided following customer requests

#### ***4.1.3 Company Outreach and Presence***

Cambridge Online's business is largely UK-based with only a small, overseas presence (< 10 customers). They have a number of Biotech, Life Sciences and Research customers based in the Midlands/South East of England.

#### ***4.1.4 Company Expertise***

Cambridge Online has "particular expertise and experience" in High Performance Computing. Their portfolio of products from industry leading vendors is complemented by value-added services which include technical consulting, system and network design, system build and configuration, systems integration, network infrastructure design and installation, technical support and system maintenance.

Linux, clustering and open-source software are key drivers for us, to provide customers with solutions to their computations needs. Increasingly we are consulted on storage and file systems – as the compute capacity increases, traditional methods of storage and file systems can prove inadequate. We have the leading UK expertise in the HP productisation of Lustre, SFS (Scalable File Share). HP SFS can solve the I/P bottleneck typically found on Linux clusters requiring scalable storage. Using low-cost, high-performance disk arrays provides scalability, along with very high bandwidth.

Our HPC lab provides facilities for proof-of-concept demonstrations, benchmarking, training and support.

#### ***4.1.5 Relationship to Tier-1 Organisations***

Cambridge Online state they have the "strength and depth" to provide high performance computing, scalable systems supporting tera/petascale solutions and this experience can be demonstrated.

The company works in close partnership with leading vendors for full support and have accredited relationships with Hewlett-Packard (HP), Intel Corporation, Platform Computing Corporation and RedHat amongst others.

They do feel that UK integrator partners are competitive with those in the USA, *provided always that strong Tier-1 relationships exist.*

## **4.2 SILICON MECHANICS**

### **Response from Silicon Mechancis ([www.siliconmechanics.com](http://www.siliconmechanics.com))**

Silicon Mechanics are different from other companies in the cluster marketplace; they are simply focused on building very reliable "industry-standard" (or "white-box") servers and storage systems. Silicon Mechanics is not a company engaged with the system integration of the cluster layer or other applications, but rather they work with the customer to complete the physical framework: systems, racks, power, cooling, integration into racks, site deployment, etc.

A major aim of the company is they work to make the whole process a simple one for the customer. Silicon Mechanics are very flexible. Most of their customers are highly technical IT staff who prefer a fairly direct, simple experience. This strategy of keeping the interaction simple starts with our web site ([www.siliconmechanics.com](http://www.siliconmechanics.com)) and continues through email or verbal communications with the sales and support team.



#### 4.2.1 Install Base

In the six years since the founding of Silicon Mechanics, their customer base has been represented by both the industrial (commercial) and HEI (public sector, higher education, government research) sectors. The focus is providing high-quality rackmount server and storage systems (along with the rack, power distribution, network switching, and hardware-only infrastructure), and not on providing the software infrastructure for building out a cluster. To date, Silicon Mechanics have been very successful in doing this, primarily in the USA.

Part of the process is to ensure that their products will work with over a dozen Open Source operating systems, plus Windows (though the vast majority of their systems are used with some flavour of Linux).

Silicon Mechanic's primary customer is someone who knows how to build the cluster infrastructure, or who has consulting expertise available to do so. Silicon Mechanics will work closely with their customers to ensure they are planning to use the best hardware platforms for their applications.

#### 4.2.2 Company Details and Size

Silicon Mechanics expects growth in its 2006 fiscal year (Jan-Dec) to be 80% over fiscal year 2005. The company's average annual growth rate since 2001 is 118% percent. The current employee count is 42 employees and continues to grow monthly.

The areas most important to server acquisition and support are staffed as follows:

Function	Number of FTEs
Sales	7
Product Engineering	6
Post-Sales Support	5
Production Personnel	12
<b>Total</b>	<b>30</b>

The company expects to ship 7200+ systems in the 2006 calendar year.

In the past 24 months, Silicon Mechanics has executed server system sales with 600+ customers. The customer base is quite varied, with the following industries the most notable:

- College & Universities (for both academic and research use)
- Federal Government Research Laboratories
- Web Retailers
- On-Line Web Hosting
- Game Developers
- Data Storage and Archiving

Because many of the industrial (commercial) customers not use their names in public documents, below is a list of some of Silicon Mechanics public sector installations (not listed in any particular order):

- Massachusetts Institute of Technology
- Fred Hutchinson Cancer Research Center
- Harvard University
- University of California, Berkeley
- Los Alamos National Lab
- Lawrence Livermore National Lab
- Stanford University
- University of San Diego, Scripps Oceanographic Institute

- University of Washington
- Johns Hopkins University and Medical Institute
- Carnegie Mellon University

For customer references, on a case-by-case basis Silicon Mechanics would be open to setting up communications. Two public references are:

1. Flu Cluster at the Public Health Sciences Division of the Fred Hutchinson Cancer Research Center in Seattle, WA:

'Halloran and Longini champion use of stochastic models, which take into account more real-world unpredictability, as well as many factors about the disease and the affected population... Because the models are very complex, researchers use high-performance computers — like the computing cluster Halloran and Longini recently had installed — to generate the simulations.'

[http://www.fhcrc.org/about/pubs/center\\_news/2006/jun1/sart1.html](http://www.fhcrc.org/about/pubs/center_news/2006/jun1/sart1.html)

2. Linden Labs / Second Life

'The Second Life "world" resides in a large array of servers that are owned and maintained by Linden Lab, known collectively as "the grid".'

[http://en.wikipedia.org/wiki/Second\\_Life](http://en.wikipedia.org/wiki/Second_Life)

#### **4.2.3 Company Outreach and Presence**

This question in the original context is not strictly relevant for Silicon Mechanics as it was designed to capture the outreach of majority of the UK Integrators polled for this report: Whilst Silicon Mechanics are located in the northwest United States near Seattle, Washington, most of their customers are outside the 'local area'. Many are outside North America, with customers on all the continents. Silicon Mechanics have considerable experience in working with and supporting distant customers and systems.

Washington state	30%
USA (other states)	65%
International	5%

#### **4.2.4 Company Expertise**

Silicon Mechanics' engineering and development staff design, develop, validate, and document one of the most comprehensive rack-optimized product offerings in the industry. Their product development efforts focus on the needs of customers deploying large, rack-optimized server installations and take into consideration density, power, thermal efficiency, performance, and reliability. They have strong, cooperative relationships with the most advanced technology companies in the industry, such as Intel Corporation, AMD, and Supermicro, and leverage their design efforts and research capabilities in their development process. Silicon Mechanics are able to cost-effectively and efficiently develop products specifically for each customer.

The focus is on developing solutions that make it simple and efficient for a customer to acquire, manage, and service a product through its entire lifecycle. This holistic approach sets Silicon Mechanics apart and ensures that they will remain competitive in environments where IT staff are continually asked to do more with limited resources.

Their products are built to order, and the operations and manufacturing are key components in maintaining a high level of quality and customer satisfaction.

Recently, Silicon Mechanics purchased 19,000+ square feet of state-of-the-art manufacturing space, with an additional 23,000+ square feet of contiguous space available when needed. They expect this facility to accommodate system unit shipment growth of over 600%

over the current unit production volume. This is believed to be central to producing integrated commodity server and cluster configurations for their customers.

#### ***4.2.5 Overview of the Marketplace Perspective***

Silicon Mechanics' believe cluster computing hardware has been the beneficiary of a host of commodity technologies, such as the rack-optimized server, general-purpose microprocessor, Linux, networking, and others. To date, success has been dependent on the integration of these technologies with minimal modification or value-adds. While the forces shaping these technologies are much larger than the influence of the commodity cluster market alone, recognition of clustering is steadily increasing among the companies responsible for these technologies. Silicon Mechanics see the further commoditization of clustering features as they apply to these technologies, making it easier for users to cost-effectively source the hardware framework for powerful clustering configurations.

This has been clearly demonstrated in the research space, where the norm used to be for researchers to share limited resources locally or remotely to do their work. Now they can purchase their own small clusters for less than the pay of a part-time assistant. In the last 12 months, Silicon Mechanics are beginning to see RFQs (USA equivalent to ITTs) for hardware-only proposals for large clusters, without requests for software, middleware, and other professional services that were more common previously.

To respond to these changes, Silicon Mechanics are adding more depth in storage and file systems solutions in response to customers' needs.

#### ***4.2.6 Relationship to Tier-1 Organisations***

Currently, Silicon Mechanics has several large-scale server installations at customer sites, with the largest being greater than 1000 nodes. They believe that their infrastructure is capable of handling tera- & petascale solutions presently. Silicon Mechanics work concurrently with Intel Corporation, as a Premier Provider, and Advanced Micro Devices (AMD), as a Platinum Provider, as well as additional manufacturers like Supermicro Computer, Inc., in the proposal process for large-scale installations. They use these resources readily in an effort to increase the efficiency of acquisition and deployment of these installations.

Many of their customers have made a business decision to move away from Tier 1 hardware platforms, realizing that they can obtain more economical and better-customized solutions by working with Silicon Mechanics as the integrator, and that they themselves have the ability to implement a cluster environment.

However, on occasions there are a small number of opportunities that carry a greater risk and liability profile than our infrastructure will support, and for those Silicon Mechanics will seek to utilize Tier 1 relationships. Currently Silicon Mechanics maintain a relationship with Hewlett-Packard (HP) as a Certified Higher Education Partner, and can work with any customer to determine the benefits of using Silicon Mechanics versus a Tier 1, multinational corporation.

#### ***4.2.7 Added Value***

Whilst not the size of a traditional Tier 1 organization whereby Silicon Mechanics could fund a research project or faculty position, they work closely with the technical community. In the United States, for example, they are a primary sponsor of the national LOPSA (League of Professional System Administrators) organization.

Specific to the Seattle area, just this past year Silicon Mechanics have started a series of quarterly presentations to the local university research community in which they have invited technical guest speakers from Intel, AMD, HP, and Supermicro to discuss their future

technologies and products. They plan to sponsor Linuxfest Northwest 2007, a local, yearly gathering of Linux enthusiasts and professionals.

Silicon Mechanics also sponsor a local, Seattle-based chapter of system administrators, a program they intend to expand into other geographies. They engage in strategic donations of hardware to selected Open Source projects, such as MythTV. Finally, their development staff makes various code libraries available to the Open Source community under the GPL.

### **4.3 LINUX NETWORKX**

**Linux Networkx website ([www.linuxnetworkx.com](http://www.linuxnetworkx.com))**

#### **4.3.1 Install Base**

Linux Networkxs has been responsible for building some of the most powerful supercomputers in the world, and as such have an install base of large systems that is far greater than any of their UK counterparts. The 10 fastest Linux Networkx supercomputers are:

- JVN 13.9 TFlops, 1024 Dual Xeon Nodes (ARL)
- Lightning 11.26 t TFlops, 1408 Dual AMD Opteron 2.0GHz Nodes (ACSI – Los Alamos, Lawrence Livermore & Sandia)
- MCR 11.2 TFlops, 1152 Dual Xeon 2.4GHz Nodes (production Oct 2003)
- Pink 10 TFlops, 1025 Dual Xeon 2.4GHz Nodes
- AIST 3.1 TFlops, 278 Dual Xeon 3.06Ghz Nodes (Japan)
- Jazz 1.73 TFlops, 408 Intel Xeon Processors (Argonne National Lab – LCRC)
- Organge 1.63 TFlops, 256 AMD Opteron 1.6GHz (Los Alamos)
- Catalyst 1.5 TFlops, 128 Dual Xeon 3.06GHz (Sandia National Lab)
- Powell 1.5 TFlops, 128 Dual Xeon 3.06Ghz Nodes (Department of Defense)
- Brahms 1.4 TFlops, 128 Dual Xeon 3.06GHz Nodes (Boeing)

Note the information provided on the above URL is not representative of all Linux Networkx installations, merely those from customer sites who have given permission for Linux Networkx to publish details on their website.

Linux Networkx client base covers a range of sectors including Manufacturing, Life sciences, Government & Research, Entertainment and Oil & Gas. From the website it is not possible to provide a breakdown of academic to commercial sales ratios.

#### **4.3.2 Company Details and Size**

Linux Networkx is based in the USA but has offices/subsidiaries worldwide. The main objective of the company is to help customers improve product development and scientific research by delivering high productivity computing systems. Linux Networkxs strives to raise industry standards through new technologies, higher customer satisfaction by delivering proven computing systems to help customers overcome their most difficult computing challenges.

Linux networkxs workforce consists of hardware and software engineers, installation & integration staff & sales representatives, although the distribution of employees in these fields is not provided.

#### **4.3.3 Company Outreach and Presence**

Linux Networkx's outreach covers most of the world, with sales representatives in the United States, Europe, Middle East, Africa and Asia. The company also has a small number of 'approved resellers' to cover the UK, France and Egypt. Unlike many of the other integrators

considered in this section, Linux Networkx can point to a successful UK installation, having recently installed an Evolocivity II Linux cluster computing system at the European Centre for Medium-Range Weather Forecasts (ECMWF). The system will be used to evaluate the suitability of cluster technology for broader deployment within ECMWF's high performance production environment, primarily as a test bed for various aspects of ECMWF's operational workload. Linux Networkx successfully met ECMWF's requirements for acceptance testing on June 18 2005. The cluster is fairly modest in size, consisting of 64 AMD Opteron 2.2 GHz processors, 128 GB of memory and uses InfiniBand high-speed interconnects from Mellanox. The cluster also includes the companies own cluster management tools, Clusterworx and Icebox, to provide total cluster management from one interface.

#### ***4.3.4 Company Expertise***

Linux Networkxs clearly has a proven track record in the HPC market with a range of systems in the TOP100. The company has incorporated an Active cooling technology to the cluster designs, developed their own management software together with storage to provide a unified high performance computing solution.

The company has partnered with many top-tier application vendors to optimize and pre-integrate software onto the clusters. They also work directly with customers to understand their specific software application and integrate and optimize the performance on their Evolocivity clusters.

#### ***4.3.5 Marketplace***

Linux Networkxs covers a diverse market place with a range of compute solutions, from HPC to grid. Many installations appear to be focussed in the USA, however sales are global with one of the most recent being the ECMWF in the UK. The academic sector is not specifically mentioned on the website.

#### ***4.3.6 Relationship to Tier-1 Organisations***

Linux networkx aims to partner with key software and hardware vendors to jointly design and sell turnkey cluster solutions that are optimised for specific applications and markets. Partnerships have been made with both Intel and AMD in order to deliver some of the fastest supercomputers in the world.

Other important technology solution partnerships have been made with high performance interconnect firms (Mellanox, Myricom, Quadrics), compiler developers (Pathscale, Portland) as well as novel architecture manufacturers (Clearspeed).

#### ***4.3.7 Additional Information***

Linux Networkx in addition to providing hardware have their own Total cluster management packages:

- [Clusterworx](#) A comprehensive Linux cluster management software package;
- [Icebox](#). is a hardware management appliance that combines a serial terminal server and a remote controlled power distribution for simplified cluster management. Scalable to support thousands of nodes, each Icebox has a network connection that allows multiple boxes to create a highly scalable IP-based communication network.

Linux Networkx also have partnerships with several leading independent software vendors (ISVs) to offer fully integrated systems. Optimising and integrating mission critical applications with the ISV enables delivery of high productivity compute clusters

#### **4.4 WESTERN SCIENTIFIC**

##### **Obtained from Western Scientific website ([www.wsm.com](http://www.wsm.com))**

Unfortunately very little in the way of information on client base / technical assistance is provided at the above website.

##### **4.4.1 Company Details and Size**

Western Scientific is a global provider of high-performance computing and storage solutions. Founded 26 years ago, Western Scientific supply an extensive line of computing solutions including the latest Beowulf / HPC clusters, RAID and tape storage, high performance workstations & servers and networking solutions for the multi-users Linux, Unix & Windows marketplace.

##### **4.4.2 Company Outreach and Presence**

Western Scientific main Head quarters are located in the United States. The company also targets the European audience and has an office located in the United Kingdom. They did put in an appearance at the 2004 Machine Evaluation Workshop in December 2004, and while promising much, have been invisible since that event.

##### **4.4.3 Company Expertise**

Western Scientific has partnerships with key Tier-1 organisations including AMD, IBM and Intel. It also has key collaborations with major high performance interconnect providers (Mellanox & Cyclades) and main stream Linux OS (RedHat and SuSE)

#### **4.5 SCC**

##### **Obtained from SCC website ([www.scc.com](http://www.scc.com))**

##### **4.5.1 Install Base**

SCC international client base spans both public and private sectors with specific expertise in Banking, Financial and Professional Services, Manufacturing, Pharmaceuticals, Retail, Leisure, Telecommunications, Transport and Utilities, together with Defence and Intelligence, Education, Health and Local and Civil Government.

##### **4.5.2 Company Details and Size**

SSC has a 28 year history of successful growth, whose business has developed from an initial investment in the UK of €3,000 into a €3 billion turnover business with leading positions in seven key European markets and business partners in over 65 countries.

##### **4.5.3 Company Outreach and Presence**

SCC is a strong company within the European market place – no information is provided for activities outside the European region.

SCC has offices in the following European countries:

- United Kingdom
- Belgium
- France
- Germany
- Italy

- Netherlands
- Spain

#### *4.5.4 Company Expertise*

In order to develop and deliver core solution sets and leverage SCC has developed relationships with key technology vendors. The 'Enterprise Solutions' section of the company is formally organised into 6 key technology pillars.

Each pillar operates under its own management with specialist sales, consultancy and vendor relationship management resources. Each management team "expends considerable time and resource with vendors and service providers from both a leading and emerging marketplace position. This ensures a continual flow of new ideas and technology components for solutions architects to design, test and deliver best of breed solutions with rapid time to market. Distilling their experience with bringing solutions composed of leading technology components to a wide variety of customers allows SCC to develop thought leadership and trusted advisor status".

Enterprise Solutions does not operate independently of the traditional customer account manager or sales contact. Customer engagement is on a planned, integrated basis based on business needs which have been intelligently identified. Enterprise Solutions is a rich source of solutions expertise to be introduced and managed SCC account management teams assigned to customers.

At all stages of engagement SCC Enterprise Solutions works to a process that ensures that value goals are clearly defined at the outset, audited during the project lifecycle and measured at its conclusion.

- **Enterprise Computing – process and transact**  
Key vendor relationships are: HP, IBM, Sun, SGI
- **Enterprise Storage – store and retrieve**  
Key vendor relationships are: HP, IBM, Sun, Veritas, Network Appliance and EMC
- **Enterprise Communications – join and protect**  
Key vendor relationships are: Cisco, Nortel, CheckPoint, Nokia, APC and QinetiQ
- **Enterprise Software – collaborate and analyse**  
Key vendor relationships are: IBM, Oracle, Microsoft, Citrix
- **Enterprise Print Solutions – copy and archive**  
Key vendor relationships are: HP and Xerox
- **Enterprise Solutions Architects – design and transform**  
Key vendor relationships are all the above, ISVs and emerging technology partners (e.g. VMWare)

#### *4.5.5 Marketplace*

As outlined above, SCC covers a variety of commercial clients as well as academic institutions. No data is provided on the website regarding the proportion of commercial to academic install bases.

#### *4.5.6 Relationship to Tier-1 Organisations*

SCC has a number of key partnerships with Tier-1 organisations including IBM, HP, Sun, SGI. See section 4.4.4 for further details.

## 4.6 WORKSTATIONS UK

### Obtained from Workstations UK ([www.wsm.com](http://www.wsm.com))

Workstations UK Ltd is the European agent for Terrascale Technologies.

TerraGrid is the fastest, most scalable shared storage solution available.

#### 4.6.1 Company Details and Size

Workstations UK is based in Amsersham, Buckinghamshire in the United Kingdom. Inventors of the blade server, Workstations UK has experience of MPI, PVM, SCI, Infiniband interconnects, NAS & SAN storage.

#### 4.6.2 Company Outreach and Presence

Workstations UK operates within EMEA space. The customer base includes:

- Conoco/Phillips
- Sandia National Laboratory
- EBI
- Raytheon
- NNSA
- Defense Intelligence Agency

Workstations UK is currently involved in HPC projects in Norway, Italy, Switzerland (CERN) and Brazil.

#### 4.6.3 Company Expertise

Workstations UK is focussed on the TerraGrid parallel storage platform from TerraScale technologies, and acts as the European agent for TerraGrid.

TerraGrid is used in Geophysical processing, Biotechnology, Digital media, Mechanical and Electrical Engineers and High Performance Computing, where TerraGrid “makes a linux cluster behave like an SMP”.

## 5. Summary and Conclusions

In overviewing the current HPC landscape, this paper has considered the multitude of issues faced by an organisation when deciding how best to procure, maintain and maximise the usage of any associated HPC resource. We have concentrated on the potential role of HPC integrators in any partnership that looks to maximise this entire process, and whether such organisations in the UK have the ability to provide the necessary level of expertise required in all phases of the process, from procurement, through installation onto ongoing support of the resource throughout its life cycle. Our primary conclusions are as follows;

1. Crucial issues when considering potential integrator involvement include both size of the proposed hardware solution i.e., number of nodes, and the ongoing robustness of open source software solutions that might be deployed on these platforms. Specifically;
  - a. The *size of the system* in question – is this targeting less than 1000 processing elements or cores, a domain in which most of the integrators have experience, or does the system in question exceed, say, 10+ TFlop. If the latter, it is worth mentioning that the current national HECToR procurement rejected the use of integrators at an early stage having considered their capabilities through a series of presentations at SC’2003 in Dallas. While US Integrators certainly have extensive experience in the 1000+ CPU domain, this is not in general the case for



their UK counterparts, although the current procurements associated with SRIF3 are changing this landscape. However in fairness to the UK integrator market, tenders of the 1000+ CPU systems are increasing, which has repercussions on companies experience in this market place. It would certainly be in the interest of Europe to give some of the integrators mentioned in this document an opportunity to demonstrate their capabilities in the 10+TFlop arena.

- b. The increased reliance on parallelization and hence system size to accomplish the highest levels of performance will merely act to emphasise the operational challenges associated with extremely large systems, challenges that stretch the resources of proprietary vendors to the limit and are realistically beyond the reach of many of the HPC integrators central to this paper
  - c. The expected *usage pattern and environment* around the resource – is this being driven by *Capability* or *Capacity* requirements? We would again suggest that integrators remain capable of providing the latter requirement far more effectively than the former.
  - d. The level of *RAS features* expected of the HPC solution. Demanding levels of RAS (say 95+%) around truly large systems are exceptionally difficult to sustain, particularly in a *Capability* regime when running large jobs with long execution times. Assuming such features appear in any contract around the services to be provided, it is extremely unlikely that any integrator would be in the position to accept the risk involved in committing to high levels, and would need the partnership framework in partnership with an experienced Tier1 organisation.
2. Our considered view is that while existing UK HPC Integrators certainly do have a valuable role to play in the on-going provision of capacity-based resources i.e. less than 1000 processing elements, the majority are less able to provide added value to high-end capability machines where the focus lies on stringent RAS requirements.
  3. The HPC integration marketplace is growing rapidly. The install base of commodity-based clusters is accelerating at pace in the UK, funded through initiatives such as SRIF, and a large portion of that business is going to the integrators identified in this paper, and not to Tier-1 vendors such as IBM and HP. The reasons for this are easy to understand;
    - a. The focus of Tier-1 activity remains on the larger, proprietary-based machines where the margins are greatest. Companies such as HP and IBM remain primarily focused on their proprietary CPU offerings – e.g. the power series – with much of their pre- and post-sales support targeting such solutions. Interestingly, however, IBM has become far more engaged in the SRIF arena over the past 12 months.
    - b. The margins remain less attractive for Tier-1 organisations when dealing with commodity solutions. We have certainly witnessed Tier-1 vendors discouraging commodity solutions in favour of their own proprietary-based solutions.
  4. Most integrators see the HPC market continuing to grow, as the technology continues to mature, and new innovations drive performance ever higher. One possible caveat here however is the issue of Full Economic Costing. This is now playing an increasing part in the decision-making process as the factors that affect FEC are becoming more visible, and sites are now reaching their capacity to provide the space, mains power and air conditioning required to run a super computer cluster. The days of major injections of capital funding through Universities and SRIF may be drawing to a close, and with it a much needed funding stream for many of the UK integrators.
  5. There is some confusion over the role that Tier-1 vendors actually play in the UK market, a point made by some of the integrators. 80-90% of all HPC Cluster developments in the UK have been carried out by systems integrators, even those ‘sold’ by a Tier-1 vendor. For example, IBM and HP Cluster solutions have been integrated and installed by OCF for a number of years. Dell worked historically with Scali to position their PC offering into a “Dell Cluster”, and more recently with ClusterVision. Streamline recently built

most of the large HPC Clusters sold in the UK by SUN, while Compusys historically were the integration and support specialists for the Cray XD1 supercomputer.

6. This issue of technical competence is seen by all the integrators as the key differentiator, and key to the future of their organisations. As cluster sizes grow but commodity hardware becomes cheaper and cheaper, the “value” in the market will be the ability to make sure that clusters work efficiently for a broad range of applications and function in a more scalable manner and operate seamlessly across sub-systems. System management and monitoring capabilities will become more important, and will help the ability of the integrator and the end-user to support such a system through its lifetime. The ability to provide a high level of support, not only on a system level, but also on an application level, will help differentiate the ‘box-shifters’ from the serious HPC oriented companies.
7. In the commodity-based solutions market, integrators provide cost-effective, technology compelling solutions rivalling those of alternative Tier-1 organisations.
8. One potential engagement model would be to form an Integrator Technology Partnership with those integrators who are deemed appropriate to the task in hand. These organisations typically do not have legacy turf wars that have made previous attempts to structure multi-Tier1 vendor consortiums around high-end HPC solutions extremely difficult (e.g. in the UK’s national HPC procurements - HPC’97 and HPCx), and are far more able to accept such a solution. This would have the obvious advantage of pooling highly competent, but thinly spread, technical expertise.
9. All integrators have existing relationships with Tier-1 vendors, although the nature of these interactions varies considerably. It is clear that deployment of very large systems requires the financial strength of a Tier-1 vendor to execute larger contracts and arguably needs the specialist expertise of key integrators to provide the knowledge and skills to build and support the systems. Productive relationships with Tier-1 vendors are seen as critical success factors for future growth of many of the integrators, and does provide an engagement model for other organisations, assuming that the Tier-1 vendor of choice does not inflict an ineffective integrator (or vice versa).
10. We do not feel that traditional SI houses have a role to play in this arena – they are invisible within the academic space, and would realistically have a steep learning curve to climb to be in a position to deal with the technology issues central to HPC provision. We would suggest that their value is clearly “only perceived at a corporate rather than operational level”.

## 6. Bibliography

- [1] <http://www.sandia.gov/ASCI/>, <http://www.llnl.gov/asci/>
- [2] N. Boden, D. Cohen, R. E. Felderman, A. E. Kulawik, C. L. Seitz, J. N. Seizovic, and W. Su, "Myrinet: A Gigabit-Per-Second Local-Area Network," in *IEEE Micro*, **15**, 1995, pp. 29-36.
- [3] F. Petrini, W.-C. Feng, A. Hoisie, S. Coll, and E. Frachtenberg, "The Quadrics Network: High-Performance Clustering Technology," in *IEEE Micro*, **22**, 2002, pp. 46-57.
- [4] <http://www.fcw.com/fcw/articles/2003/0825/tec-lightning-08-25-03.asp>
- [5] [http://www.brightsurf.com/news/aug\\_03/PNNL\\_news\\_082703.php](http://www.brightsurf.com/news/aug_03/PNNL_news_082703.php)
- [6] OSCAR: <http://www.osl.iu.edu/publications/pubs/2003/oscar:ols03.pdf>, <http://oscar.openclustergroup.org/tiki-index.php>
- [7] ROCKS: <http://www.rocksclusters.org/Rocks/>
- [8] <http://www.bsc.org.es/>
- [9] <http://www.llnl.gov/linux/thunder/>
- [10] [http://news.com.com/2100-1008\\_3-5208220.html?tag=nefd.ledc](http://news.com.com/2100-1008_3-5208220.html?tag=nefd.ledc)

- [11] Maui HPC Centre (MHPCC), <http://www.mhpcc.edu/>
- [12] Texas Advance Computing Centre (TACC), <http://www.tacc.utexas.edu/>
- [13] Cambridge University HPC Service, <http://www.hpc.cam.ac.uk/darwin.html>
- [14] SRIF Scientific Research Investment Fund, [www.hefec.ac.uk/research/srif](http://www.hefec.ac.uk/research/srif)

## 7. APPENDIX 1: Integrator Questionnaire

### Initial Integrator Discussion Points around HPC provision

To best inform the data gathering exercise associated with this document, a set of preliminary questions were devised and discussed with each of the integrators, typically by phone during the 2<sup>nd</sup> and 3<sup>rd</sup> weeks in June. These questions are sketched out below, with the responses of section 3 driven off the following eight points:

1. Understanding the current install base (both in the UK and abroad) and trends in the cluster marketplace. Information on current install base (ideally over the last 4-5 years, providing a picture of changing trends), including where possible the site, architecture, size, procurement date etc. (naturally no financials are expected). An idea of the relative number of "small" (32-64) systems compared to larger (128+) Machines. What is the approximate split between HEI's and industrial installations?
2. Details and Company size/status etc: Company overview - background, status, size etc. In providing this information it would be useful to have a breakdown of the relevant parts of the organisation - sales staff, after-sales support team, technical support (software and hardware if possible) etc., approximate turnover (machines not finances), customers (numbers - names are not necessary although they might prove useful). The details of, say, three customer reference sites would be helpful.
3. Company Outreach and Presence: An extension to the first two points - what presence do you have overseas - in particular in Europe and the USA - and what, if any, is the size of the current install base outside the UK
4. Company areas of Expertise - From the technical perspective, what level of technical expertise do you feel you bring to the system integration market that makes you competitive and a long term prospect in that market place. Please mention any other tie-in's you provide which you feel are relevant.
5. Company perspective of the Marketplace - How do you feel the cluster market place is changing (especially with FEC coming into effect). Any information you can provide as to changes/investments you are making to adapt to this

changing climate, e.g. requirement to increase skills in middleware, compiler, database, files systems etc; need to forge strong links with software companies / interconnect solution; impact of GRID/e-science developments?

6. Relationship to Tier-1 Organisations: Do you feel your company has the strength and depth to build high-performance, scalable systems which can support Tera & Petascale solutions in the not-too-distant future. Are you able to provide this independently, or do you require Tier-1 support in dealing with areas such as risk and liability. Do you feel that the UK integrator providers are competitive with those in the USA?
7. Relationship to Other System Integrators - Do you perceive any role for the more traditional SIs e.g. EDS, CSC in this marketplace. These are pretty invisible to us, but you may have a different perspective over what appears to be a more expensive alternative - at least in the non-academic space?
8. The above pointers are clearly not exhaustive, so if you feel that any other information is important in our trying to understand your position in the marketplace, please feel free to mention it. Can you provide a viable cost-effective alternative to the established blue chip/Tier-1 companies when it comes to procuring mid-range / high-end systems?

## **8. APPENDIX 2: Company Contact Details**

### **8.1 CLUSTERVISION**

<http://www.clustervision.com>

Address: ClusterVision Ltd  
17 Essington House  
Lytton Grove  
London SW15 2ET

Tel: +44 870 080 1990

### **8.2 COMPUSYS**

<http://www.compusys.co.uk>

Address: 1A Bessemer Crescent  
Rabans Lane Industrial Area  
Alesbury  
Buckinghamshire, HP19 8TF

Tel: 0870 745 7575

### **8.3 OCF PLC**

<http://www.ocf.co.uk>

Address: Rotunda Business Centre  
Thorncliffe Road  
Thorncliffe Park  
Sheffield S35 2PG

Tel: +44(0) 1142 572200

#### **8.4 STREAMLINE**

<http://www.streamline-computing.co.uk>

Address: The Innovation Centre  
Warrwick Technology Park  
Gallows Hill  
Warwick CV32 6UW

Tel:+44 (0)1926 623130

#### **8.5 CAMBRIDGE ON-LINE**

<http://www.cosl.co.uk>

Address: Cambridge Online Systems Limited  
163 Cambridge Science Park  
Milton Road  
Cambridge CB4 0GP

Tel: +44(0) 1223 422600

#### **8.6 LINUX NETWORKXS**

Address: Linux Networx GmbH  
Europaallee 10  
67657 Kaiserslautern  
Germany

Tel: +49 631 3031809

#### **8.7 SILICON MECHANICS**

Address: Silicon Mechanics, Inc  
22029 23<sup>rd</sup> Dr SE  
Bothell, WA 98021-441-  
USA

Tel: 001 425-424-000

#### **8.8 SCC PLC**

Address: James House  
Warwick Road  
Birmingham  
B11 2LB

Tel: +44(0) 121 766 7000

## 8.9 WESTERN SCIENTIFIC

Address: Studio 1  
Waterside Park  
Third Avenue Centrum 1000  
Burton on Trent  
Staffs  
England  
DE14 2WQ

Tel: +44(0) 1283 569989

## 9. APPENDIX C: Daresbury Contact Details

Should further information be required regarding any of the issues raised in this document, please contact us using any of the method below:

<http://www.cse.scitech.ac.uk/disco/contact.shtml>

Address: DisCo Group A22  
Christine Kitchen  
STFC Daresbury Laboratory  
Keckwick Lane  
Daresbury  
Warrington  
Cheshire  
England. WA4 4AD

Tel: +44 (0) 1925 603756