



To:
DCSC Board

Cc.:

From:
DCSC Secretariat and DCSC System Administrators

Subject:
DCSC SC07 Conclusions

Blegdamsvej 17
DK-2100 Copenhagen Ø
Phone +45 3532 5453
Fax +45 3532 5016
info@dcsc.dk
www.dcsc.dk

DCSC system administrator fact-finding-trip to *International Conference for High Performance Computing, Networking, Storage and Analysis (SC07)*, Reno, USA, 10-16 November 2007

5. December 2007
/RB

1. Introduction

The DCSC board decided 1 may 2007 to financially support 50% of a DCSC system administrator fact-finding-trip to *International Conference for High Performance Computing, Networking, Storage and Analysis (SC07)*, Reno, USA, 10-16 November 2007. The regional operating centres under DCSC where asked to finance the remanding 50%.

The aim of the trip was manifold:

- To bring together all DCSC system administrators for a longer period so as to facilitate better day-to-day communication, sharing of technical experiences, and problem solving. This is namely important due to the decentralised nature of DCSC.
- To meet all the major High Performance Computing (HPC) vendors, so as to introduce the whole DCSC operational organisation, to discuss mutual current and future trends, developments and expectations.
- To get hands on experience of present and upcoming technology developments as well vendor strategies, future products and services.
- To report to the DCSC board on all the above issues.

The DCSC delegation consisted of four of the five regional operating centres under DCSC as well as the DCSC secretariat, plus two strategic DCSC partners (Nordic DataGrid Facility, NDGF; a representative from the Danish GRID development environment). In all 10 persons participated as part of the DCSC delegation.

In addition to the SC07 conference, a number of closed vendor meetings had been organised. The delegation conveys its special thanks to DCSC/KU and vendors for using time and effort to plan these very enlightening meetings. The DCSC delegation met with: IBM and also visited there research lab in Almaden; HP; Dell; SUN; SGI and DataDirect Networks.

2. Findings

A number of significant trends within HPC where identified:

- **Power and cooling**

Challenges, technical as well as financial, within power consumption and cooling are not going to go away. Quite to the contrary. The general available

processors (AMD/INTEL) are as low power/cooling as they are going to get. The general expectation of rapid increase in processor speed over time is no longer valid. To mitigate this, i.e. to continue to provide increased performance, vendors are adding more and more cores to the processors to increase performance while keeping power/cooling at the current level. The individual cores in a multi core processor are generally getting slower and slower, so as to save energy.

- **Memory Bandwidth**

As more and more cores are added to the processors, the memory bandwidth per individual core goes down. This is, obviously, a major problem for memory bound codes. This problem is not going to go away.

- **Code Porting**

As all processor vendors are moving towards an expanded multi-core architecture, code porting becomes a necessity for *all* codes not already fully parallel. Moving serial codes to parallel will for at least a subset of codes running in DCSC take many years. Hence, the sooner this is acknowledged at the strategic level, the better. Some code porting will probably be in the 5+ year range. Compilers will not be able aid this for the users in the foreseeable future.

- **InfiniBand**

InfiniBand is becoming standard in clusters. In the near future it will probably move onto the motherboards as standard. Fibre cables are going to make it more manageable.

- **Field-programmable gate array (FPGA)**

FPGA's are already here but are not making much headway. While they appear very interesting for certain codes, the technology as such does not seem "*prime time ready*", mostly due to lack of compiler support.

- **Massive Parallel Processors (MPP)**

MPP's are coming back strong. While MPP is a very interesting architecture, it seems better suited to specific codes, and generally carry a higher price tag - in general and per compute cycle.

- **Constellations**

As cluster nodes with multiple cores in reality act as SMP machines, clusters are quickly turning into constellations. SMP coding directives will also become part of modern cluster programming. Today clusters are offered with nodes of 8 cores, next year probably 12 or 16. The distributed memory architecture of real SMP architectures offsets the memory bandwidth issue for several years, thus leveraging memory intensive compute problems.

- **Storage**

Storage capacity is evolving quite fast as storage requirements rise and rise. Clustered file systems are also getting more and more mature and moving slowly towards the commodity market. Flash based disks are most likely going to become quite standard in the near future. The packaging density is also going up.

- **Hardware offloading**

Hardware offloading of e.g. MPI calls is looking very interesting. It remains to be seen over the next couple of years, if it pays off.

- **Microsoft entering the HPC field**

Microsoft is moving into HPC. While their current technical solution is pretty basic (1 queue, FIFO and so on), they will most likely become a big player over the next couple of years, mainly for uses bound to a windows platform and with limited HPC competences.

- **INTEL**

INTEL's next generation chip, with on-board memory controllers, will most likely become available in Q4 2008. They will most likely perform better than AMD, also on memory intensive codes.

3. Main Conclusions

In terms of technological trends, there is consensus in the delegation, as to what the most significant trend is: That is the fact that one can no longer count on the usual rapid increase in processor speed over time. This will mean that users will have to adapt to multi-cores HPC to a hitherto unseen extent. In turn, this will mean that code porting will be unavoidable for the majority of users, since existing code will no longer run faster.

In terms of DCSC organisational and strategy gains, there is also consensus in the delegation, that the SC07 fact finding trip was a great team-building success, which gave a significant amount of knowledge transfers, both to DCSC and within DCSC. The strengthened personal network within the DCSC delegation is believed to be important for the future coherence within the DCSC organisation, in terms of sharing knowledge and experiences, and solving any future challenges, which might need to involve the whole DCSC organisation.

The vendor market is indeed moving towards standardisation on the cluster architecture. However, it is also evident that specialised architectures have legitimate utilisation with the DCSC user community.

The dedicated vendor meetings were found to be not only informative but also important for future vendor negotiations. The fact that DCSC, for the first time ever, meet vendors as a unified entity is believed to have strengthened the importance of Denmark and DCSC as a current or potential customer.