

---

To:**DCSC Board**

Copy:

---

From:**DCSC Secretariat and DCSC System Administrators**

Subject:

***DCSC SC08 Conclusions***Blegdamsvej 17  
DK-2100 Copenhagen OE  
DenmarkPhone +45 3532 5453  
Fax +45 3532 5016  
E-mail info@dcsc.dk  
Web www.dcsc.dk  
CVR 29979812  
EAN 579800042249025 November 2008  
/system administrators**DCSC system administrator fact-finding-trip to *International Conference for High Performance Computing, Networking, Storage and Analysis (SC08)*, Austin, USA, 15-21 November 2008.****1. Introduction**

The DCSC board decided the 13<sup>th</sup> of June 2008 to financially support each DCSC regional operating centre by up to DKR 15.000,- in a fact-finding-trip to *International Conference for High Performance Computing, Networking, Storage and Analysis (SC08)*, Austin, USA, 15-21 November 2008. The regional operating centres under DCSC were asked to finance the remanding costs.

The aim of the trip was:

- To bring together all DCSC system administrators for a longer period to facilitate better day-to-day communication, sharing of technical experiences, and problem solving. This is especially important due to the decentralised nature of DCSC.
- To meet all the major High Performance Computing (HPC) vendors, in order 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 all five of the five regional operating centres under DCSC as well as the DCSC secretariat. In all six persons participated as part of the DCSC delegation.

In addition to the SC08 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:

- HP, for a two-day briefing where we visited their research lab and production facilities in Houston.
- IBM, for a one-day briefing at their facility in Austin.
- SUN, for a three-day conference in Austin.

- As well as shorter meetings with Dell, DataDirect Networks and Intel.

## 2. Findings

A number of significant trends within HPC were identified:

### Compute

Intel's next generation chip promises a significant performance gain for a number of reasons:

- Memory bottleneck issues have been resolved by moving the memory controller on die, thus removing the need for the north bridge. However the configuration of the new memory subsystem is crucial. By talking directly to the experts from Intel, the DCSC delegation gained significant insights into the proper configurations, apparently not known by all vendors.
- GP-GPU (general purpose graphics processing unit) is promising energy efficient computations. The learning curve for utilizing the potential is however considered steep.
- In general this year we saw a number of heterogeneous systems, e.g. the most powerful computer at this writing is made up of three different processor types. As with the GP-GPU, these system promises efficient computations, at the price of a highly complex development cycle.

### Storage

- Increased focus on parallel file systems, most notably the Lustre file system.
- Storage vendors move toward selling complete storage systems as appliances, delivering service end-points. A move from selling JBODs (just a bunch of disks).
- SSD (solid state disk) is making its entrance into storage systems. Currently it is mostly being used as a fast buffer for HDD storage. The scale-up in IOPS (I/O operations per second) is measured in orders of magnitude.

### Interconnect

- 10 Gigabit Ethernet is making its entrance into the mainstream market, being a promising technology for interconnects of high bandwidth/high latency for servers and blade enclosures.
- QDR (quadruple data rate) Infiniband is now available, promising a doubling in throughput from today's DDR (double data rate) low latency interconnects.

### Data centre

- Most HPC systems are being run without the use of UPS.
- Server density is still rising and will continue to do so in the foreseeable future.
- Server density has necessitated cooling to become an integral part of the compute rack (e.g. cooling doors) as well as water-cooling of the CPU (IBM Power6 systems).
- There is an added awareness on power saving measures, such as automatic hardware power down in the CPU, as well as on the motherboards. This trend is also present in storage systems.

### 3. Main Conclusions

In terms of DCSC organizational and strategy gains, there is consensus in the delegation, that the SC08 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 organization, in terms of sharing knowledge and experiences, and solving any future challenges, which might need to involve the whole DCSC organization.

The dedicated vendor meetings were found not only to deliver crucial technical insights, but also important for future vendor negotiations. The fact that DCSC met vendors as a unified entity is believed to have strengthened the importance of Denmark and DCSC as a current or potential customer.

The vendor market is indeed moving towards standardization of cluster components and interconnect. It is also evident that specialized architectures have the potential for great improvements in performance; this however comes at an added cost in development/maintenance complexity.

Significant gain in performance is expected, with Intel's new processor as well as SSD buffered storage systems.

Having had the all technical briefings and discussions with vendors as a common DCSC entity has enriched the group with a common technical reference, setting the starting point for further discussions within the group of DCSC administrators.