

## KISTI Selects Sun for Next Generation Supercomputer

May 4, 2007

Focus Point: Technical Computing

### INTRODUCTION

The Korean Institute for Science and Technology Information (KISTI), a public research organization located inside the KAIST campus in the high-tech metropolis Daejeon that provides information and computing services to the Korean research community, has selected Sun Microsystems as the main vendor for its next generation cluster supercomputer. While contract details remains to be worked out, and no contract has awarded at this time, it is likely that Sun will provide KISTI with the hardware, system software, and integration services for a computer system that could be as large as 300 Teraflops once completed next year. KISTI will also purchase a smaller SMP-system with a target performance of 40 Teraflops from IBM.

### KEY OBSERVATIONS

Following the surprise win of the Tokyo Institute of Technology procurement, TSUBAME, at the end of 2005, followed by the successful bid for a similar system presently being installed at the University of Texas, Sun has aggressively pursued a number of high-profile opportunities in the technical computing market. While a number of questions remain open, the KISTI procurement now demonstrates that Sun's intends to build a long-term presence in the high-end computing market. While prizing appears to have been aggressive, it was less price than high-level management attention and commitment by Sun to building a joint research center with KISTI that appears to have been the crucial factor to select Sun over IBM and HP, the two other vendors selected to make a proposal.

#### • KISTI's Future Generation Supercomputer Project

The history of KISTI's supercomputing center goes back to the 1960s and the electronic computing office. The present Supercomputing Center was established in 2001 as the product of a merger of two offices within the organization. KISTI had in the past purchase systems from Cray, IBM, and Compaq but it was with the procurement of a large IBM p690 SMP server cluster in 2001 that KISTI established itself as a leading technical computing center. Installation of the US\$ 27 million 3.7 Tflop/s peak performance system was completed in 2003. For the present procurement, KISTI had eventually selected three vendors—IBM, Sun, and HP— to make a proposal.

Installation of KISTI's next generation system, which had been postponed many times, will be completed in two phases: 20% of the 300TF system will be installed this year while the remainder of 80% will only be installed in 2008. We expect contract to be awarded to Sun and IBM in late May; contracts will be for the entire system. KISTI also plans to purchase a high-speed storage system with a performance target of 0.5 GB/s in I/O performance for each Teraflops of compute performance from Sun Microsystems. This would yield a gigantic 100GB/s storage system—by far the highest performing storage system in Asia.

Aggressive pricing has been clearly important in winning this deal. While KISTI had asked for proposals for a 200 Teraflops system, Sun eventually bit a system with a theoretic peak performance of 300 Teraflops, But, while the economics of Sun systems for TITECH (a computer system within the TOP10 for only US\$ 20 million) were certainly an added attraction, a pledge by Sun to invest up to US\$ 10 million into a joint research center with KISTI appears to have been an important additional factor in the decision to select Sun as preferred vendor.

#### • **TSUBAME: Buying into the Global HPC Market**

The announcement, at Supercomputing 2005, that Sun Microsystems had won the procurement for the new supercomputer at the Global Science Information and Computing Center (GSIC) at the Tokyo Institute of Technology (TITECH) did surprise many HPC insiders, both in Japan and abroad. The Sun-built "grid cluster" TSUBAME was installed successfully by the end of March 2006 and, at a total performance of 50.4 Teraflops, TSUBAME ranked as Number 8 in the latest Top500 list. The GSIC procurement was the brainchild of Prof. Satoshi Matsuoka, a parallel computing and cluster specialist and professor of computer science at TITECH. Dr. Matsuoka's lab had build a number of large clusters using AMD processors, including the Presto III cluster with 78 AMD Athol™ processors, the first AMD systems to be listed in the Top500 list in July 2001. Dr. Matsuoka was critically involved in managing the overall procurement process, drafting the specifications for the system, and eventually also in bringing in the right partners for this procurement—notably AMD and the accelerator-board vendor ClearSpeed.

The Tsubame system consists of 639 Sun Fire X6400 compute servers, each with 8 sockets of 800 Series Dual-Core 2.4 GHz Opterons and 32GB of memory and an additional 16 systems with 8 2.6 GHz Opterons (800 Series) and 64 GB of memory each. The storage systems is based on 42 "Thumper" storage servers, each with 48 500GB SATA drives and two double-core 2.4 GHz Opterons. The entire system is linked together by an Infiniband Fat Tree network with 8 Voltaire ISR 9288 Infiniband switches. In addition to the Thumper systems, 96TB of NEC iStorage are used for mission-critical data (such as all user-related data, etc.).

While initially TITECH had hoped for a total performance of around 85TF, once equipped with all 600 ClearSpeed accelerator cards, maximum performance of the system running Linpack is presently only 47,38 Teraflops (TF). In an initial benchmark in May 2006 the system had reached 38,18TF. TSUBAME ranked as the number 9 system in the November 2006 Top500 list. According to TITECH engineers, it is still difficult to obtain full performance from the installed ClearSpeed CSX600 PCI-x boards. In order to complete its Linpack benchmark, GSIC had to considerably lower the clock speed of the boards, thus also lowering total performance. GSIC engineers believe that further improvements are possible over the next few months, but it seems unlikely that TSUBAME will ever reach the 85TF Libpack target. TITECH also appears to have experienced problems with the storage subsystem.

With Sun and ClearSpeed, TITECH was certainly taking a risk, especially given Sun's limited experience with large x86 cluster systems and its relatively weak HPC team in Japan. Overall, AMD's support appears to have been crucial for this procurement. TITECH was significant both for the Japanese Asian HPC market since.

- **Long-Term Strategy: Focus on High Visibility**

Following the spectacular TITECH and University of Texas procurements, Sun's win at KISTI demonstrated that the company is re-emerging as a player in the global HPC market, a market largely dominated by IBM and HP. Still, despite these important wins, Sun's long-term HPC strategy remains uncertain, especially after the company failed to win the DARPA HPSS competition, despite a proposal judged as highly innovative by many observers.

The company's recent aggressive approach to the HPC market is linked to the acquisition of Kealia, an Opteron server start-up company managed by Sun co-founder Andy von Bechtolsheim. Sun has exploited, with considerable marketing skills, procurements at TITECH and the University of Texas, which were very aggressively priced. There is no doubt that TSUBAME has helped Sun to re-establish itself as a player in high-end technical systems. More importantly, these highly visible wins also provided a launch pad for Sun's Opteron and Thumper world-wide and investments in HPC are clearly used by Sun as a marketing tool to launch its new line of Opteron servers and storage systems. Also, the company has started to build strong partnerships with a number of software companies, such as Cluster File System Inc., which leads development of the open-source LUSTRE file system—presently the only veritable alternative to the IBM GPFS parallel file system.

Still, there remain a number of questions as to Sun's long-term technical computing strategy and a number of weaknesses for Sun's approach to the HPC market in Asia persist. To start with, Sun has experiencing a number of operational and technical challenges at some of its sites. Delivery delays have plagued many recent installations, including TITECH (where Sun had eventually to mobilize a small army of engineers to install the system in record time to meet acceptance deadlines) and the University of Texas. And, while the Tokyo Institute of Technology was widely advertised as the first installation for Sun's new storage server (code-named "Thumper"), no performance benchmarks for the TITECH system have been released so far, suggesting that the storage system may not be fully functional even one year after installation. Again, while not unusual, this raises questions about Sun's commitment to the HPC market.

Also, in most Asian countries—and including in Japan and Korea—Sun has only very limited resources to service complex installations and the number of sales and pre-sales technical staff devoted to the HPC market is tiny. But, again, these need not be demerits. Its extensive focus on a small number of highly visible procurement has helped the company to win against IBM or HP. As in the case of other recent large systems sold by Sun, direct attention to clients by senior management, and especially the Sun founder Andy von Bechtolsheim, has paid off handsomely. Limited internal staff also means that Sun needs to work extensively with local partners, which can help limit internal cost, but carries its own risks.

- **Expect Further Big Wins for Sun in Asia**

While Sun's HPC strategy—both overall and regionally in Asia—remains somewhat unclear, we nonetheless expect the vendor to pocket further big wins in the near

future, and including in Asia. Recent procurements at TITECH and KISTI have shown that Sun is willing to make significant investment, and to sell systems below cost.

In Japan, the TITECH procurement is presently replicated at three sites—the University of Tokyo, Kyoto University, and the University of Tsukuba—and in three parallel procurement with a total value of over US\$ 130 million and there appears to be good potential for Sun to win at least one of these procurements. Also, a large Sun system installed in January 2003 at the University of Tokyo's Institute for Medical Sciences is up for renewal and we expect Sun to win the procurement for replacing the present system. Interest in Opteron-based hardware from Sun is also increasing in industrial markets—such as EDA or CAE and even risk modeling in the financial sector—and at least in those areas where application expertise is not deemed critical Sun has a strong value proposition. Outside Japan, the company is well entrenched into the academic market in Asia. With universities across the region invest more into computing and build larger facilities, there is considerable potential for Sun to increase its presence in the market for mid-sized and large cluster systems. Finally, with its new cluster and storage offerings, Sun is in a strong position to break into the extremely price-sensitive oil and gas exploration market.

#### • Sun and NEC?

While NEC is also an incumbent vendor at KISTI—which had purchased SX-5 and SX-6 systems in 2001 and 2003, respectively—it appears that Sun's new partner NEC was not involved in the KISTI procurement. Under Sun's alliance with NEC, which was announced after the TITECH procurement, the two companies plan to cooperate on large clustered systems, with Sun providing hardware and NEC contributing software and application expertise. We believe this alliance has limited implications for the Asia-Pacific region outside of Japan, given the weakness of the Japanese vendor in the region. By contrast, in Japan, where NEC's expertise and internal resources would certainly benefit Sun in some areas, it is highly unlikely for NEC as the leading Japanese HPC vendor today to promote hardware build by Sun. In the case of TITECH, NEC was the prime contractor, yet the company was not involved at all in the installation and integration of the system, which was completed by Sun alone. If that was the pattern of cooperation in the initial case, one wonders why it should be different in other cases?

#### CONCLUSION

With KISTI, Sun has won a second large and high visible technical computing procurement in Asia. Beyond its regional impact, the KISTI win is important for Sun since it demonstrates that the company can successfully compete in the IBM/HP dominated worldwide technical computing market. While some questions remain open, it demonstrates the vendor's commitment to the world-wide technical computing market and its willingness to invest in important and highly visible partnerships with research organizations in Asia. But, KISTI is also an important defeat for IBM, the dominant HPC vendor in Asia outside of Japan and the incumbent vendor at KISTI. Finally, the KISTI procurement is also an important win for the LUSTRE file system, the only real competitor to IBM's GPFS, which is increasing its share in the high-end HPC storage market in Asia.