

## Japan's Next-Generation Supercomputer Project

May 15, 2007

Focus Point: Technical Computing

### KEY POINTS

- Japan's next-generation supercomputer project was announced in July 2005 with a plan to develop a 10 Petaflops general purpose computer by 2011. The government agency overseeing the venture—the Ministry of Education, Culture, Sports, Science and Technology (MEXT)—is likely to allocate a budget of around JPY 110 billion (around US\$ 920 million) to the project.
- MEXT selected RIKEN, a government-backed physics and chemical research institute, to develop the supercomputer. It's also becoming clear that all three major high-performance computing (HPC) vendors, NEC, Fujitsu and Hitachi, will be included in the project as suppliers.
- For RIKEN, this project is a key opportunity to expand its project portfolio and strengthen its international visibility. RIKEN also has the ability to grow its budget at a time when government-funded research organizations are legally required to decrease base funding.
- After much discussion, RIKEN announced on March 28, 2007 that the supercomputer would be developed and housed in Kobe.
- We expect the system's architecture to be publicly released by MEXT in late July or early August. We suggest that the design will feature a new multi-core chip design with up to 256 cores per CPU. The overall system will be made up of 5000 nodes with 9 CPUs per node. Automated parallelization techniques will be used to deal with such a large number of cores in a single node.
- The Keisoku System will feature a proprietary, extremely high bandwidth interconnect. There is some indication that the system may not use an external storage approach, but will rather rely on a number of local disks attached to each node.
- To date, despite project leaders' recommendations for more emphasis on software, the supercomputer project has been mainly positioned as a hardware-centric project. This raises concerns that the project will lack software innovation.

### INTRODUCTION

After almost two years of debate regarding project location and system architecture, the details of Japan's next-generation supercomputer project are finally taking shape. On March 28, 2007, RIKEN, The Institute of Physical and Chemical Research, which is responsible for the development, construction and operation of the supercomputer, announced that Kobe will be the site for the supercomputer project. And while the architecture of the system has not been publicly announced yet, it is becoming clear that all three major Japanese vendors — NEC, Fujitsu, and Hitachi — will be involved in the project as suppliers. Although this is a key win for these vendors, the supercomputer project could fail to adequately boost their business since the participation of three

project contributors may cancel the benefits for any single provider and only prolong the lives of mediocre platforms and technologies that can't stand on their own. In what follows, Springboard Research discusses the present status of the supercomputer project, its impact upon technical computing and potential relevance to fields beyond the narrow scope of technical computing.

## **EXECUTIVE SUMMARY**

Japan launched its next-generation supercomputer project nearly two years ago when the Ministry of Education, Culture, Sports, Science and Technology (MEXT) informally challenged select research centers and universities to build a follow-up system to Japan's Earth Simulator. The Earth Simulator, until 2004 the world's fastest computer, is a large vector meteorological modeling supercomputer developed by organizations including the Japanese Atomic Energy Agency (JAEA) and NEC. Determined to once again lay claim to the development of the world's fastest computer, in 2005 MEXT quickly settled on a next-generation supercomputer design submitted by scientists at RIKEN, a government-backed physics and chemical research institute. The accepted proposal called for a 10 Petaflops ( $10^{15}$  calculations per second) general purpose supercomputer, to be completed by 2011. Also called the KEISOKU Project, from the Japanese number "KEI," which denotes  $10^{16}$  and "SOKU" for "speed," the supercomputer's development has since been held up as it's taken well over a year for MEXT and RIKEN to work out a more detailed plan for the system. Various political hurdles including the selection of a new facility to house the supercomputer and key uncertainties regarding the supercomputer's system architecture have further postponed the project. Recent developments including the site selection of Kobe and the probable inclusion of all three major Japanese computer vendors in the project suggest KEISOKU may now be on track despite significant delays, but concerns remain. Springboard Research questions the positive impact KEISOKU will have on Japan's IT industry, and notably the technical computing businesses of the Japanese vendors, since there is every indication that the project has failed to select a single architecture, and a single vendor, but rather is likely to support all three Japanese vendors. The implication is that there will simply be not enough resources and funding for any of the three vendors to develop an innovative architecture that can make a broad impact not just in Japan but also abroad. Further, we do not see yet what innovative research on applications will eventually come out of this project. Yet, the entire technical computing business — even in areas like weather and climate modeling — is now all application-driven.

## **BACKGROUND: FROM THE EARTH SIMULATOR TO KEISOKU**

*The KEISOKU project is the latest in a series of large-scale HPC projects and was conceived as a follow-on project to the Earth Simulator, which was completed in 2001. As such, understanding the history and dynamics of this project helps to provide a foundation for the context, significance and possible ramifications of KEISOKU.*

Completed in 2001, the Japanese Earth Simulator system led the list of the world's fastest computer systems for two years. While it has by now been overtaken by several other systems, its unique architecture earns the Earth Simulator a reputation for the highest-

performing system for many applications, notably weather and climate modeling, which do not perform well on large, microprocessor-based parallel systems.

The surprising success of the Earth Simulator as a bold hardware development effort has contributed to funding initiatives in the U.S. and to a lesser extent, in Europe, to build similar systems. Interestingly, in Japan, the Earth Simulator had been initially highly controversial, since it was seen by many scientists as a one-off project that drained resources from other projects, while not being well integrated into the Japanese technical computing community. For example, it took several years for the Earth Simulator to connect to the Japanese academic backbone network, SuperSINET, a fact that put considerable limitations on external usage of the system. Early limitations of the Earth Simulator's storage environment also meant that the system could not be used in an optimal fashion. Finally, a strict initial limitation to climate research and insufficient efforts to develop other software platforms to run on the system have limited the Earth Simulator's output considerably. Furthermore, an internal policy that required users to be physically present at the Earth Simulator Center in Yokohama had angered users in Japan and abroad. Political restrictions have, both internally and internationally, meant that only a small number of users had access to the system.

To address these issues, MEXT initiated consultations with major HPC users to plan a follow-up system to the Earth Simulator. Following an internal competition in 2005 that came down to a decision between two competing proposals from the Earth Simulator Research Center (positioned as a follow-up project for the ES system) and RIKEN (for building a new hybrid architecture that would link special-purpose computer blades and general-purpose computer blades in a common, modular network infrastructure), MEXT selected the RIKEN proposal, which was also supported and developed by NEC.

## **PROJECT OVERVIEW**

In 2006, MEXT proposed a budget of JPY 110 billion (around US\$ 950 million) for their multi-year effort to build a 10 PFlops computer system that groups together four distinct efforts supported by the Japanese education ministry:

- To build a computer system that reaches a Linpack performance of at least 10 Petaflops.
- To develop software applications for “Grand Challenge” scientific applications in nanoscience and biology.
- To further develop the Japanese all-optical research network SuperSINET and build a “layered” supercomputer infrastructure that links all computer systems at national universities and research centers.
- To create a national Center of Excellence (CEO) for high-performance technical computing that supports hardware and infrastructure software development, as well as software application development and that also functions as a training ground for future generations of computational scientists in Japan.

### **• System Architecture**

The decision regarding system architecture has been complicated by the fact that the high-performance computing business of all three Japanese vendors participating in the project is presently at a turning point. NEC has continued to develop its SX

architecture for vector systems, but average annual sales over the past few years have slipped, making continued development of the product line unfeasible. Fujitsu and Hitachi don't presently offer competitive HPC products and product development activities have wound down considerably for both manufacturers. Thus, for all three companies, the KEISOKU project has broader implications for their technical computing business and HPC-related activities. While NEC will continue the development of its vector-based SX architecture for the supercomputer, Hitachi has suggested a Blue-Genie-like architecture and Fujitsu is likely to build a large parallel system with multi-core SMP nodes that makes extensive use of automated parallelization.

While the system architecture has not publicly announced and, at this stage, remains a tightly guarded secret, Springboard has learned from sources close to the project that the main system will feature a architecture with a proprietary interconnect and a novel multi-core chip design. While the possibility of a parallel vector architecture had been discussed early on in the project, we believe that Vector architectures could not meet the power consumption limits imposed by RIKEN. We expect the system to have around 5000 node with a 9 processors SMP-like cluster node. The choice of 9 processors per core is based upon two considerations: (1) With 9 processors, 8 processors can be used for calculation while one processors does I/O, which yields a maximum 8 to 1 ration of I/O "compute processors" to " I/O processors". (2) Alternatively, the 9 CPU design reflects the fact that many simulations are 3D based, thus giving the programmed to allocated 3 CPUs to each dimension. Since the goal is to reach at least 10 PFlops on a standard Linpack benchmark the CPUs for the system will feature a design with as many as 128 and up to 256 cores. In other words, each node of the system will be constituted by a large cluster with probably over thousand cores. We believe that this novel multi-core chip will be designed by Fujitsu. In order for programmers to be able to deal with such a huge number of cores within a single node, we expect Fujitsu to invest considerably in the development of a special software stack that will take automated parallelization techniques to a new level.

One of the most impressive features of the Earth Simulator was an extremely high performance proprietary network. While many large systems in the US use standard Infiniband interconnects (or else interconnects that use at least the physical Infiniband layer but add special high-bandwidth I/O processors, such as in the Cray Red Storm/XT3/XT4 design), we predict a highly proprietary network design for the Keisoku systems.

#### • Project Site

As previously mentioned, RIKEN announced on March 28, 2007 that it had selected Kobe as the site for the facilities of the next-generation system and for a new center of excellence for high-end computing and scientific applications. Kobe is already the location of the RIKEN Center for Developmental Biology (CDB) and the new supercomputer facility will be located adjacent to the CDB in an artificial island that houses a number of research centers, companies, and hospital facilities. RIKEN's site decision comes after the release of a report on March 23 by a special committee for next-generation supercomputer site selection. While the committee initially considered a total of 15 applications, the final competition included 5 sites, several of which are

also existing RIKEN sites: Wako near Tokyo, Yokohama, Sendai, Kobe, and Osaka. Among the 5 locations, Wako, Sendai, Kobe, and Osaka were deemed as appropriate sites for the new facility, based on an extensive evaluation of the costs and benefits of each site. Based on this analysis, RIKEN then selected Kobe as the site for the project.

- **Project Impact on Riken**

For RIKEN, Japan's premier science research institute, KEISOKU is a key opportunity to expand its project portfolio and strengthen its international visibility. The supercomputer project is also an opportunity for RIKEN to grow its budget at a time when government-funded research organizations are legally required to gradually decrease base funding and personnel. Still, the project does create problems for RIKEN. Funding for the project is by no means assured and many at RIKEN are worried that, if faced with funding decreases, RIKEN will have little choice but to cut other activities to fund the completion of the KEISOKU system.

The restriction on hiring new staff puts further limitations on RIKEN. As part of the project, RIKEN will build a new center for high-end technical computing that aspires to become a national center of excellence. Yet at the same time the number of people that RIKEN can hire for this project is severely limited. As was the case with other large facilities such as the synchrotron radiation facility SPring-8, it is highly likely that RIKEN will eventually choose to outsource management in order to reduce the headcount of full-time staff employed by the new center.

- **Possible Lack of Software Innovation**

To date, despite project leaders' repeated recommendations for emphasis on software, the KEISOKU project has mainly been positioned as a large hardware development effort and is hardly innovative from a software perspective. Indeed, most project funding appears to be allocated to hardware development and deployment. This raises many concerns. To start with, despite recent efforts by government agencies to fund new software development, Japan's capacity to develop software tools and applications for large technical systems remains limited. In fact, almost all of the infrastructure software for large cluster systems has been developed in the U.S. and the same is true for most dedicated applications. Even the Japanese commercial market for technical computing software — in application areas from life sciences to computer-aided engineering (CAE) and electronic design automation (EDA) — is almost entirely dominated by U.S. or European companies. For automotive CAE in particular, by far the largest market for technical computing in Japan, only a handful of application software tools have been developed by small Japanese vendors (none who've succeeded in building a global market for their products), while the overall market is controlled by increasingly large numbers of U.S. and European companies. Current funding allocations for KEISOKU raise some doubt as to whether the project will change the present situation much and have a lasting impact at all — while Japanese bureaucrats argue that one goal of the project is to help Japanese vendors gain a competitive advantage in the technical market, for the moment there is every indication that KEISOKU will end up as a one-off project. It is true that RIKEN has made extensive efforts to mobilize research groups throughout Japan that are involved in software development, but a number of decisions made over the past two years — including the

selection of senior project managers — suggest that at present, overall project efforts remain hardware-centric.

### **KEISOKU and Japan's Computer Industry: A Domestic Microprocessor Again?**

Much of the strong political support for the US\$ 920 million project has to do with the perception that, over the past decades, Japan's semiconductor and computer industry has largely failed to build its own brand of computer processors for server (or desktop) computers and the implicit political scenario behind the KEISOKU project is that a large public investment will eventually help Japanese vendors to come forth with their own microprocessor architecture. It is true that Japan's large computer vendors, all of which have a long history as suppliers to the telecommunications industry, had failed to foresee the change from proprietary "closed" mainframe systems to "open" RISC-based systems some two decades ago.

Recent trends with multi-core architectures suggest that the computer industry is likely to experience an important shift over the next decade that is perhaps best characterized by the fact that parallel computing techniques (once an exclusive domain of technical computing) are increasingly becoming mainstream, both in standard server architectures and in more specialized applications such as game machines. While, often, this simply means that the number of cores in a CPU is increasing, parallel processing poses far-reaching and difficult software challenges that commercial software vendors are only starting to come to terms with. How likely is the success of CPUs with large numbers of cores? Developed by Sony, IBM, and Toshiba for Sony's Playstation III game machine, the CELL processor is presently the most radical new computer architecture to be sold into a mass market. The large and growing market for high-end game machines represents an increasingly important segment of the overall computer industry (and an extremely important market for IBM's POWER series of microprocessors). Game machines are somewhat less dependent on open standards and developer are probably somewhat more willing to shift to new architectures. It remains yet to see how multi-core architectures with a large number of cores will do in the server market outside of high specialized technical computing applications.

### **CONCLUSION**

After nearly two years of debate, Japan's future-generation supercomputer project is finally taking shape. With RIKEN selected as the organization to implement the project and Kobe designated as the site for the supercomputer, the way is cleared for further project development. More importantly, details of the design of the main system to be developed are starting to become available. We anticipate that both NEC and Fujitsu will be involved in the project, while Hitachi will probably play only a minor role (e.g. limited to component development). For Fujitsu, this will be a key project to re-launch its internal HPC business and we expect the company to start reorganizing its worldwide technical computing business as soon as the project is publicly announced later this year. Of course, both the KEISOKU project and the Japanese vendors have still a long way to go, but the first steps have now been taken!