Proposed Agreement between California Energy Commission and SeaMicro Incorporated

Title: SeaMicro Volume Server Power Reduction Research and Development
Amount: $250,000.00
Term: 24 months
Contact: Paul Roggensack
Committee Meeting: 9/1/2010

Funding

FY	Program	Area	Initiative	Budget	This Project	Remaining Balance	
10	Electric	IAW	ARRA	$750,000	$250,000	$0	0%

Recommendation

Approve this agreement with SeaMicro, Inc. for a 24 month, $250,000.00 cost share grant to supplement the contractor's American Recovery and Reinvestment Act of 2009 (ARRA) award. Under the Energy Efficient Information and Communication Technology DE-FOA-0000107, SeaMicro, Inc., was awarded $9.3 million in U.S. Department of Energy (DOE) ARRA funds and is providing $10.7 million in additional match funding to conduct the Volume Server Power Reduction Research and Development project. The total budget for the project is $20.25 million. Staff recommends placing this item on the discussion agenda of the Commission Business Meeting.

Issue

Electricity consumed by data centers is approximately 3 percent of California's total electrical consumption and is growing rapidly. Energy use by data centers has doubled every five years, and the largest user of energy within data centers is the volume server. Volume servers are the most common server in the data center industry and use about 68 percent of the energy in a data center not counting support equipment. The volume server is the workhorse of the data center, but it is inefficient and poorly packaged.

For today's volume servers, the Central Processing Units (CPU) are not well matched for the workload that has come to dominate the data center. Volume servers use multi-core CPUs that are designed for complex computational challenges and, in terms of computation per dollar, are best-in-class at complex workloads, but are inefficient at simple computational workloads. Volume servers are also called "rack servers" because they are stacked like pizza boxes one on top of another in metal racks. These servers are usually 1.75"-3.5" tall, 18" across, and up to 30" deep. Each server is a discrete unit that is individually managed, controlled, powered, and cooled. These individual servers are linked together with Gigabit Ethernet switches and routers. This results in massive duplication of components as each server has the overhead necessary for individual operation, management, and connectivity, but is never used in isolation.

Background

SeaMicro, Inc. was awarded a $9.3 million grant from DOE to develop the Volume Server Power Reduction Research and Development (Project). The Project will develop a prototype compute appliance that will replace volume servers, use significantly less energy and reduce the footprint and cost.

With the rise of the Internet and the rapid growth of the data center, the CPU has few complex problems. The challenge is how to handle the huge volume of relatively modest computational workloads called web-tier workloads. These workloads are generated by the millions of independent users each wanting to search, view web pages, check email, and read the news, all for free. Multi-core CPUs are particularly inefficient for these simple computational workloads. The mismatch between the CPU in volume servers and the primary workload in the data center is a fundamental cause of the rapid rise in power consumption by volume servers.

The most efficient CPUs for the most common workload in the data center were not intended for servers at all. They were designed for handheld devices and small laptop computers. These CPUs are simpler designs requiring less power, and more efficient at web-tier workloads. They offer more than three times the performance per watt of the large multi-core CPUs for this workload. They provide half the single-thread performance for a sixth of the power draw, a dramatic improvement in computation per unit of power.

With a more efficient CPU for the workload of a data center identified, SeaMicro addressed the inherent inefficiency of server packaging. SeaMicro invented input/output virtualization technology that allows common components to be shared by hundreds of CPUs. This allows removing 90 percent of the components of the server, everything except the CPU and memory, reducing the server to the size of a credit card.

Next, SeaMicro developed an interconnected fabric to link 512 servers in a three-dimensional array inside the compute appliance. The CPUs are managed with load-balancing software for dynamic allocation of workloads. The intelligent allocation of work to CPUs, based on power-use metrics, allows the CPUs to operate in the most efficient range, while allowing CPUs not being utilized to enter deep sleep. The result is a system that uses a quarter of the power in a quarter of the space to do the same work as today's volume server and is less expensive to manufacture.

SeaMicro has obtained venture capital for match funding from Khosla Ventures, Juniper Networks, Crosslink Capital, and Draper Fisher Jurvetson totaling $10.7 million.

Proposed Work

SeaMicro has developed a prototype, innovative compute appliance that is currently in beta testing with customers. The PIER funding, coupled with the ARRA grant and match funding, will be used to:

- Accelerate development engineering and quality assurance testing
- Enhance dynamic power-management software to further drive down power draw
- Accelerate product availability

- Accelerate market adoption by reducing cost and modifying the system to address markets beyond volume servers, such as volume servers working on less-common workloads and mid-range and high-end servers.

The prototype appliance redesigns the volume server with more efficient CPUs and achieves over 75 percent reduction in power consumption per computation for the primary workload in the data center. The compute appliance addresses the following issues concerning the inefficiency of the CPU within volume servers:

- Different and varying workload requirements by the CPU due to the internet.
- Inefficiency in server packaging that left enormous duplication in components.

Using the U.S. EPA's estimated energy use for data centers of 61 billion kilowatt-hours/ year in 2006, the estimated energy savings in California would be 3.7 billion kilowatt-hours/ year, based on 80% of volume servers used for web-tier workloads, 68% of energy in the data center used for powering and cooling data centers, and 15% of national data center floor space located in California.

Justification and Goals

This project "[will develop, and help bring to market] increased energy efficiency in buildings, appliances, lighting, and other applications beyond applicable standards, and that benefit electric utility customers" (Public Resources Code 25620.1.(b)(2)), (Chapter 512, Statues of 2006)); and this project "[will] advance energy science or technologies of value to California citizens..." (Public Resources Code 25620.(c)), and is part of a "full range of research, development, and demonstration activities that . . . are not adequately provided for by competitive and regulated markets (Public Resources Code 25620.1.(a)).

This will be accomplished by:

- Developing a compute appliance that may achieve over 75% power reduction per computation in the data center and occupies less than 25% of the space of traditional volume servers.