Case Study: Sun Microsystems Energy-Efficient Modular Cooling Systems

Silicon Valley Leadership Group

accenture

High performance. Delivered.

Sun Microsystems hosted the assessment of four commercially available energy-efficient modular cooling systems for data centers. The systems assessed were designed with different goals, different loads, and in some cases different functions (e.g., operation at various chilled water temperatures or control of inlet air temperatures. Due to differences in system capability, each assessment used unique system-specific inlet air temperatures, supply water temperatures, and server loads. Thus the performance metrics cannot be compared against one another directly without understanding the design variations and test conditions. The goal was to compare the efficiency of the modular approach against conventional cooling systems.

The Customer Case

In 2003 the Rocky Mountain Institute hosted an energy summit focused on identifying opportunities in data centers. Lawrence Berkeley National Laboratory (LBNL) also developed a data center "Research Roadmap" for the California Energy Commission. Recommendations from both of these efforts suggested that modular construction could solve observed problems of inefficient operation when centers are not fully occupied, by distributing power more economically or cooling high density IT equipment more efficiently. LBNL initiated the study to investigate energy implications of commercially available modular cooling systems compared to that of traditional data centers.
Data centers require continuous cooling to maintain inlet air environmental conditions within ASHRAE recommended operating ranges for their IT equipment. The shift to IT equipment having higher

Sun microsystems

Berkeley Lab

APC by Schneider Electric

Liebert

modius

POWER ASSURE

Switch to perfection RITTAL

VETTE CORP

IBM

power density has increased the amount of power and cooling required per square foot of floor space. This increase often exceeds what the data centers were originally designed to handle. As a result cooling systems in many data centers are not well suited to handle dense IT equipment. To fix the problem operators try to overcool the space by mixing hot and cold airstreams or over-sizing the cooling systems. The outcome: systems often operate inefficiently.

Electrical and HVAC (power and cooling) systems in most data centers also operate considerably below their design basis for much of their life. Loading conditions are dynamic due to varying computing and storage loads in the short term and incremental build-out in the long term. If excess power and cooling capacity is operating, these systems often operate inefficiently. A new generation of commercially available modular cooling systems offers the potential to satisfy both fully and partially loaded conditions while achieving higher efficiency than legacy systems operating in this dynamic environment.

For this study the four commercially available modular, scalable server rack cooling solutions were evaluated under a range of conditions reflecting different inlet air temperatures and server loads, and, for some systems as appropriate, elevated chilled water temperatures. The systems evaluated were: APC's ACRC100 InRow Cooling Units, Vette/IBM's Rear Door Heat Exchanger, Liebert®'s XD cooling system, and Rittal's LCP+. These units were installed at Sun Microsystems' data center in Santa Clara, California.

Each system's performance was evaluated with a range of server loads (at 25%, 50%, 75%, and 100% for two systems and at 50% and 100% for the other two systems) in the cooled racks, and a range of inlet air temperatures (e.g., at 68°F, 72°F, 76°F, and 80°F). For the in-row and overhead cooling units, the water supply temperature to the modular units was fixed at 45°F (the available chilled water temperature for the campus). The water temperatures for the passive cooling and liquid cooled rack units were elevated to 50°F, 60°F, and 70°F, respectively by connecting an additional coolant distribution unit (CDU) to the cooling module to explore the feasibility of using higher temperature chilled water. The CDU consisted of a water-to-water heat exchanger that provides isolation between the primary chilled water and secondary cool water loops that are connected to the module. The CDU used internal secondary loop pumping and measured primary and secondary water pressures, temperatures and flow rates. (The 68°F inlet air temperature case was not evaluated when the water temperature was at 70°F. The other modules were not evaluated at higher water supply temperatures because their designs precluded connecting the CDU). All units were provided controllability of inlet air temperatures except for the passive cooling system.

Data collection was performed using the Modius OpenData® Data Center infrastructure Manager, a manufacturer-neutral measurement system that collects real-time continuous performance information from site infrastructure equipment in a data center. Over 1500 data points were measured over the study's duration, gathered from the manufacturer's equipment, the Sun servers, and an array of power meters, flow meters, pressure transducers, and RTD temperature sensors. Data was gathered via a variety of network and serial protocols. Similar data points were measured for each rack cooling technology and stored in a shared relational data base at a remote server. Over 40 million individual measurements were taken during the study. Real-time data was available through a web application, allowing LBNL personnel and each manufacturer's engineers to monitor and manage the study remotely. Access controls ensured that each manufacturer could see only its own data while LBNL had access to all data.

To avoid running cable from the central plant, chiller plant data was retrieved from the building management system trend log, and power demand to the data center IT equipment and the chiller plant was recorded manually. In addition the Coefficient of Performance (COP) was computed. COP of a cooling system is the ratio of the heat removed by it to the work supplied to it. The COP is unit less, with a higher value representing higher efficiency for the cooling module.

Total cooling delivered by the cooling module, in kW, was calculated from the secondary-loop chilled water temperature rise and chilled water flow rate. For this evaluation the portion of chiller pumping power required to deliver the chilled water volume in the primary-loop was ignored for all units. Module System Efficiency (MSE) was calculated as the ratio of total cooling power to cooling transported by the modular system, in kW/ton. A Power Index (PI) was computed as the ratio of power demand for the cooling module to computer load (measured at the PDU supporting the server racks).

Project Outcome

Overall, the modular cooling systems provided effective cooling to maintain acceptable temperature ranges for the servers for all the operating scenarios studied. A comparison of the kilowatt per ton performance is set out below.

Sun designed and operated two electrically powered modular cooling systems to control the temperature over inlet air with chilled water supplied at the fixed temperature of 45°F supplied by Sun's chiller plant. Depending on operating conditions, the time required for achieving steady-state cooling could be 30 minutes or longer. The cooling delivery coefficient of performance (COP, ratio of cooling provided by the module to the power to drive the module) for one system ranged from 19.6 up to 35.5, varying by a factor of 2 depending on the server load and inlet air temperature.

The power index (PI, ratio of power demand of the cooling module to compute load) for this system ranged from .03-.08. The corresponding range for the other system was from 4.7 to 13.1, varying by a factor of 3, and its PI ranged from .1 to .16.

One enclosed electrically powered modular cooling system was designed and operated to control the temperature over inlet air with chilled water supplied at a temperature varying from 50°F up to 70°F. For selected test conditions, the system achieved set points and provided steady-state cooling within 15-30 minutes. Its cooling delivery COP ranged from 4.7 to 13.1 under elevated water temperatures and server loads below its design condition, varying by nearly a factor of 3, depending on server load, inlet air temperature, and water supply temperature. The corresponding PI ranged from .1 to .2.

The passive modular cooling system (which uses no electric power) was designed and operated to use chilled water supplied at a temperature varying from 50°F to 70°F in combination with air-cooling using the surrounding air from the existing data center. The passive system could not control over inlet air temperatures by itself. For the selected test conditions, the time required for the secondary water loop to achieve steady-state cooling was 15-30 minutes. Its COP varied from 64.4 to 229.9, while PI varied from .0004 to .0009. These values indicate dramatically higher energy efficiency than the other systems, as one would expect since it uses only hydraulic power inherited from the chilled water system.

The actively powered systems varied in their efficiency as server loads changed, but in different ways. One system's COP was highest, and its PI was lowest, for higher server loads – it was most efficient for the highest loads. The other system's COPs were highest for lower server loads, but their MSE shows the opposite pattern. The point is that exclusive comparison of the COPs or MSEs results is insufficient. For instance, the efficiency metrics for one system seems lower due to the inability of the test setup to exceed partial load. Comparison must also account for differences in system designs and test conditions.

Cooling Solution Comparison

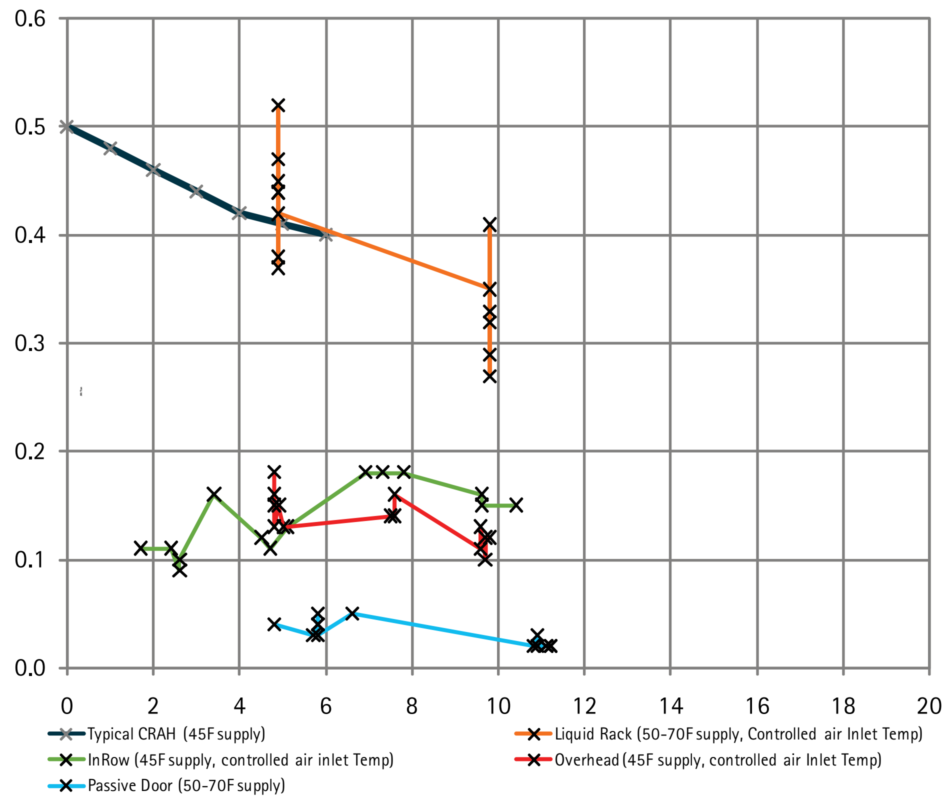

Two liquid cooling systems were tested and demonstrated effective cooling with water temperature elevated above the levels typically supplied by a chiller. One required additional air-cooling; the other did not.

Liquid cooling systems delivered necessary cooling and removed server heat effectively in the racks. All systems in did a good job of removing heat compared to traditional methods. Effective inlet air temperature control and good temperature uniformity was consistently observed in one enclosed modular cooling compared to traditional air cooling.

In evaluating the performance of specific modular systems, assessors took into account each system's unique methodology in temperature and flow controls (air and/or water). While it is difficult to draw direct comparisons of performance metrics among all four cooling systems, opportunities exist for some of the units to improve temperature and flow control, and/or to improve delivery COP by implementing fans and/or pumps with higher efficiency. Examples of potential improvement include using variable-speed-drive motors and optimizing the number of local fans.

Two liquid cooling systems that were tested provided effective cooling with elevated water temperature, in one case without additional air-cooling. This suggests that it is possible to raise the chilled water temperature, which could reduce the use of—or potentially eliminate—chillers in many climates. Free cooling using cooling towers or other evaporative cooling could save significant amounts of energy.

Compared with traditional data center cooling, proper control, integration, and implementation of all the cooling systems studied could lead to greater energy savings.

Companies should assess the overall energy efficiency and economics for their data centers in using liquid cooling at higher temperatures supplied from a cooling tower or a chiller plant operating at a higher set point, rather than at lower temperatures supplied from a chiller.

When considering the application and implementation of modular cooling in data centers, companies need to remember that modular cooling systems in data centers require additional power in addition to the universal chilled water plant for the facility to integrate and optimize them with traditional data center infrastructures.

http://www.sun.com/ecoinnovation

About the Energy Efficient Data Center Demonstration Project

The project's goal is to identify key technology, policy and implementation experts and partners to engage in creating a series of demonstration projects that show emerging technologies and best available energy efficiency technologies and practices associated with operating, equipping and constructing data centers. The project aimed to identify demonstrations for each of the three main categories that impact data center energy utilization:

- operation & capital efficiency
- equipment (server, storage & networking equipment)
- data center design & construction (power distribution & transformation, cooling systems, configuration, and energy sources, etc.).

The project also identified member organizations that have retrofitted existing data centers and/or built new ones where some or all of these practices and technologies are being incorporated into their designs, construction and operations.

About Accenture Technology Labs

Accenture Technology Labs, the dedicated technology research and development (R&D) organization within Accenture, has been turning technology innovation into business results for 20 years. The Labs create a vision of how technology will shape the future and invent the next wave of cutting-edge business solutions. Working closely with Accenture's global network of specialists, Accenture Technology Labs helps clients innovate to achieve high business performance. The Labs are located in San Jose, California; Chicago, Illinois; Sophia Antipolis, France; and Bangalore, India. For more information, please visit our website at www.accenture.com/accenturetechlabs.

Accenture Technology Labs compiled the final report supporting with analytics, background research, and project management.

About The Silicon Valley Leadership Group (SVLG)

The SVLG comprises principal officers and senior managers of member companies who work with local, regional, state, and federal government officials to address major public policy issues affecting the economic health and quality of life in Silicon Valley. The Leadership Group's vision is to ensure the economic health and a high quality of life in Silicon Valley for its entire community by advocating for adequate affordable housing, comprehensive regional transportation, reliable energy, a quality K-12 and higher education system, a prepared workforce, a sustainable environment, affordable and available health care, and business and tax policies that keep California and Silicon Valley competitive.

About Accenture

Accenture is a global management consulting, technology services and outsourcing company. Combining unparalleled experience, comprehensive capabilities across all industries and business functions, and extensive research on the world's most successful companies, Accenture collaborates with clients to help them become high-performance businesses and governments. With more than 180,000 people in 49 countries, the company generated net revenues of US$19.70 billion for the fiscal year ended August 31, 2007. Its home page is www.accenture.com.