Hot Topic – Cooling Supercomputers

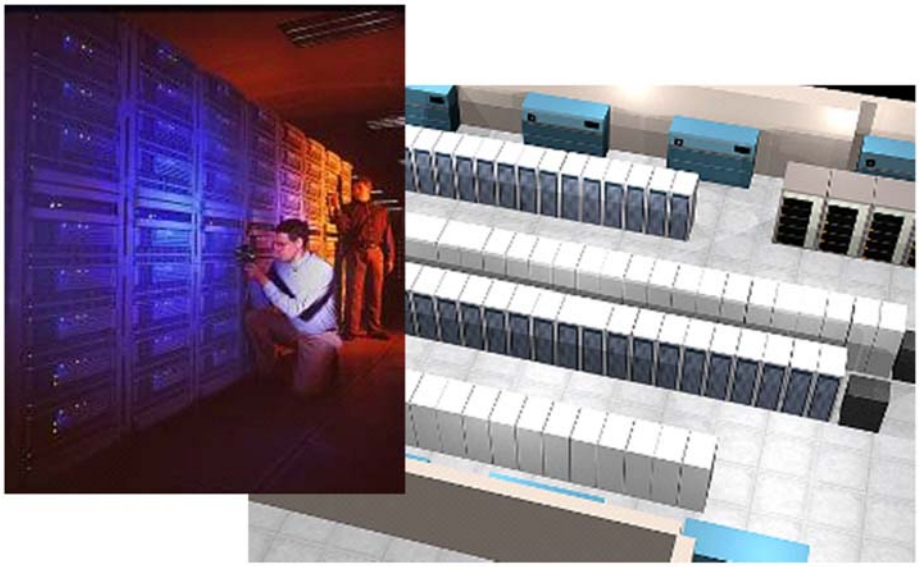

Cooling racks of computers and superconputers is a hot topic

The continued exponential growth in the performance of leadership-class computers is predicated on the ability to perform more computing in less space.

A looming crisis exists in high performance computing. Leadership-class computers are becoming increasingly power hungry and compact which creates problems of economically providing enough power, making the most efficient use possible of the power, and effectively managing the byproduct waste heat.

Few industries have seen the rapid increases in performance that have characterized computing and electronics, but this performance has come with a price. Data center electricity costs are in the range of $3.3 billion annually, as cited in Congressional discussions (H.R. 5646). Compounding the problem is the need to dispose of the massive amounts of heat generated typically using very inefficient convective cooling from airflow.

The continued exponential growth in the performance of leadership-class computers is predicated on the ability to perform more computing in less space. Two key components are 1) the reduction of component size, leading to more powerful chips, and 2) the ability to increase the number of processors, leading to more powerful systems. There is strong pressure to keep the physical size of the system compact, to keep communication

latency manageable, and to minimize new construction costs. The upshot is an increase in power density. There are two efficiency aspects of managing this increased power density. One aspect is to make the system more efficient in terms of the amount of computation performed for the energy expended. The other is to remove the waste heat more efficiently.

Energy Smart Data Center Project

Pacific Northwest National Laboratory's (PNNL's) Energy Smart Data Center project is establishing itself as the national leader in energy efficient and power aware computing by deploying a production system at the Environmental Molecular Science Laboratory (EMSL), evaluating next generation cooling and computing technologies in a testbed, and performing research in power efficient computing.

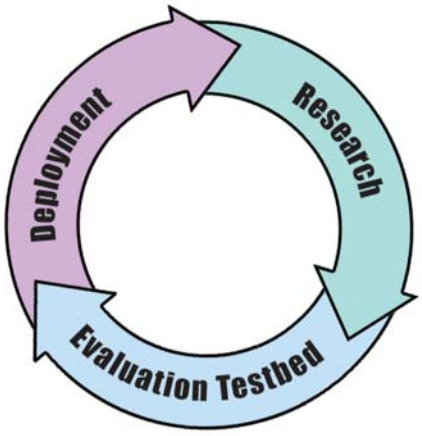

By integrating multiple aspects of heat management, using software, sensors, and hardware along with more efficient implementations involving the physical infrastructure, the Energy Smart Data Center project will be able to drive data center energy efficiency to levels never achieved before. The Energy Smart Data Center project is solidly grounded in efforts that couple research with evaluation testbeds and deployment in production systems.

Deployment

PNNL's Environmental Molecular Science Laboratory operates MPP2, an HP/Linux Itanium-2 cluster with a peak performance of 11.8 TeraFLOPS. To study the effectiveness of the Isothermal Systems Research, Inc. (ISR) SprayCool™ evaporative cooling technology in a real environment, one of the racks of MPP2 was converted from air cooling to spraycooling. A computational fluid dynamics (CFD) study of a fully spraycooled data center showed that the level of computer room air conditioner (CRAC) supplied cooling could be reduced by 20% (due to reduced airflow need), which could result in a 50% power savings. By rejecting the heat directly to the facility chilled water, or even directly to the cooling tower water, the challenges of facility airflow management are dramatically reduced. While challenges remain, scale-up to a spraycooled center is feasible.

PNNL believes research and evaluation in cooling technologies could result in avoiding $15 million in capital improvements to its own facilities and that similar savings at other large data centers could have an impact on the economy and productivity of multiple market sectors. These cooling technologies could include rear-door heat exchangers and self-contained cooling solutions. Experience with actual implementation costs, performance, production reliability, and energy efficiency could prove vital to implementing a petascale system with minimal changes to floor space and supporting infrastructure.

With the support of the NNSA, PNNL is investigating applying advanced cooling technology developed as part of its Energy Smart Data Center project to a significant fraction of the next Environmental Molecular Science Laboratory supercomputer, expected to be 5-10 times more powerful that then current 11 TF system.

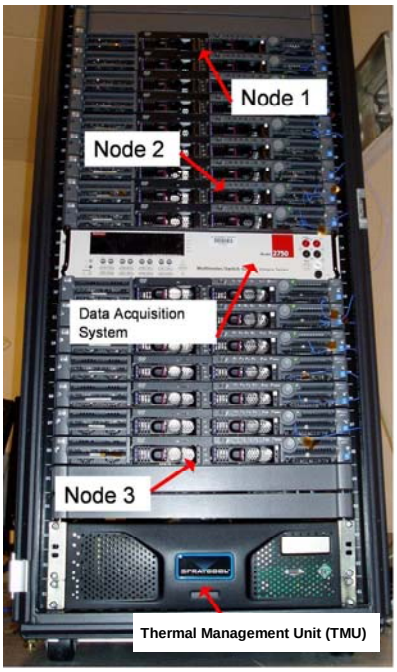

Stack of 16 HP RX1620 servers with SprayCool Thermal Management Unit used as a demonstration vehicle to assess maintainability, acceptance, and adoption by Original Equipment Manufacturers (OEMs).

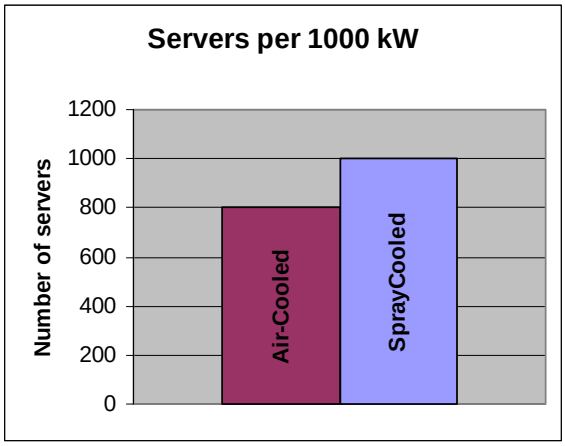

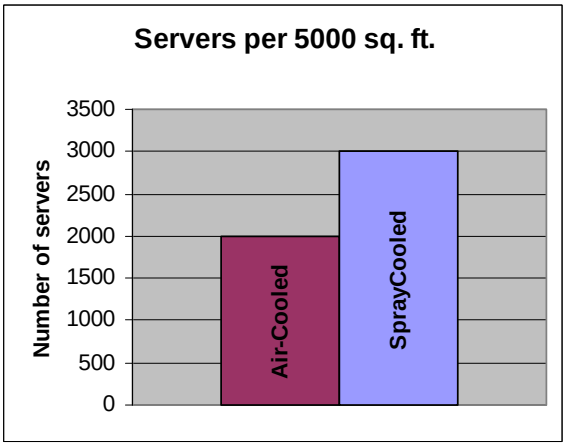

The total cost of ownership (TCO) model developed in Phase I of Energy Smart Data Center and updated in Phase II with operational data from EMSL, predicts that using spraycooling could increase the number of servers by about 20% without increasing power requirements. For space constrained data centers, the model predicts a spraycooled data center could house 50% more servers than an air cooled data center of the same size. Such densification may allow millions of dollars of new construction costs to be deferred. In Phase III of the ESDC project, PNNL will use the testbed facility to validate the TCO model results.

Evaluation

Evaluation at the data center scale requires a fully instrumented facility and a supercomputer of adequate size to provide realistic, credible results. Phase III of the Energy Smart Data Center will build an independent testbed evaluation environment at PNNL. This experimental facility will be operational in early 2007.

The testbed environment will investigate the entire heat and power flow of a TOP500 data center testbed. The compute resource will be housed in its own 800 square foot facility. The computer and associated cooling and electrical supplies will be specially instrumented to fully monitor pressure, pressure drops, and flow rates of all cooling fluids (including air) throughout the data center as well as all the power used for computing and cooling.

In addition to demonstrating the viability of the spraycooling approach, reliable full-scale data will be generated to establish accurate coefficient of performance (COP) and life cycle cost of ownership. The project will also provide data to determine the accuracy of airflow models using computational fluid dynamics (CFD).

Fully spraycooled 8U Global Chassis capable of delivering 8kW to 16 boards. A 42U rack system with dual socket 2.4GHz Opteron boards enclosed in Global Chassis is capable of delivering 768GFlops. Liquid cooling provides the opportunity to densify components as well: Opus designed boards are equipped with 8GB/socket densified 3D stacked memory (Tessera, Smart Modular).

Research

A number of research areas contribute to energy efficiency – power awareness; power conversion and distribution; as well as cooling technologies. Possible improvements range from the scale of the micro-electronics chip, through servers and racks of servers to the entire data center. The system as a whole needs to be considered, offering the opportunity to explore power, heat extraction, and management efficiencies at multiple scales. These efficiencies, actively promoted by Congress, are increasingly important to data centers deploying capability and capacity computing.

The Energy Smart Data Center project is investigating advanced cooling methods based on highly efficient liquid evaporative cooling and its ability to enable denser systems. PNNL has teamed with Isothermal Systems Research, Inc. to evaluate, mature, and deploy ISR's SprayCool evaporative cooling technology in a high performance computing environment. Results to date have shown that the technology is not only more efficient than solely air cooling but can indeed enable a denser packaging of components than otherwise possible – e.g., memory DIMMS. The project expects to further the commercialization and reliability of the evaporative cooling technology.

SprayCooling under the Energy Smart Data Center project has demonstrated several benefits:

1. Cooler air exhaust from the SprayCool rack into the facility ambient air
2. Little to no air short-circuiting
3. Significantly lower airflow rate requirement for the SprayCool rack
4. Fewer CRACs
5. Little to no limitations on the data center placement of SprayCool racks.

The Energy Smart Data Center project will investigate how to efficiently convert high-voltage AC electric power to low-voltage DC power used by individual components. DC power distribution holds the promise of larger power savings but requires modifications that are not yet well understood. In Phase III, the Energy Smart Data Center project will investigate efficient conversion technologies.

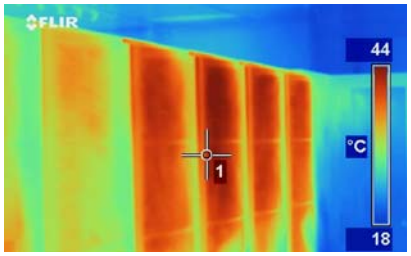

Infrared photography discloses the variability in temperatures in an air cooled data center.

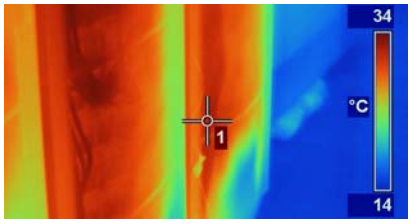

Air-cooled racks in PNNL facility exhibit exhaust temperatures ranging from 44-32C under operation.

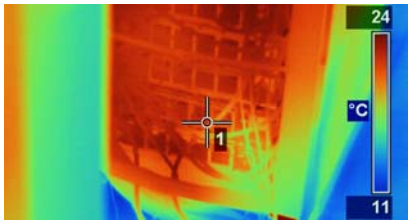

The spraycooled rack under operation runs consistently 10C cooler than the best case air-cooled rack.

Local cooling would provide a more consistent operating temperature for all servers, regardless of their physical location –- making data center design easier. PNNL's testbed facility will provide an environment to study CFD for various local cooling options.

Conclusion

With further support, PNNL will expand the Energy Smart Data Center project to include a more inclusive range of related topics, including power efficiency, power awareness, power conversion, and cooling technologies in a fully integrated program that addresses the critical power and heat management issues facing leadership-class computing.

The combination of research insight, experience deploying in a production facility, and independent testbed evaluation provides a unique capability that will establish the PNNL as the leader in this field.

For more information, contact

Moe Khaleel
Pacific Northwest National Laboratory
P.O. Box 999, K7-36
Richland, WA 99352
509.375.2438
moe.khaleel @pnl.gov

Andrés Márquez
509.372.4457
andres.marquez@pnl.gov

Steve Elbert
509.375.2504
steve.elbert@pnl.gov

Website: esdc.pnl.gov

PNNL-SA-52744