CASE STUDY WHITE PAPER

2009 BEST-IN-CLASS WINNER
For Green IT Beyond the Data Center

University of Notre Dame

2009 Green Enterprise IT Award

Uptime Institute™

Grid Heating Clusters: Transforming Cooling Constraints into Thermal Benefits

By Paul R. Brenner, PhD, PE, Douglas Thain, PhD, and Dewitt Latimer, PhD

Publisher's Note: This paper has been included in *The Path Forward v4.0*: *Revolutionizing Data Center Efficiency* because it contains significant information value for the data center professional community which the Uptime Institute serves. *The Path Forward* is the formal compilation of the papers, presentations, and proceedings of the annual Green Enterprise IT Symposium. This paper was submitted as part of the formal awards nomination application and, in part, formed the basis for the evaluation by an independent panel of judges. The Institute does not endorse any opinion or position

Executive Summary

A 2007 EPA report [1] estimates the US spent $4.5 billion on electrical power to operate and cool IT and HPC servers in 2006. The same report forecasts a staggering 2011 national IT electric energy expenditure of 7.4 billion. The computation/information technology industry has responded with innovations in processors, hardware, and systems software to reduce the growth of the oppressive power requirements. In parallel to these economic IT energy concerns, recognition of public demand for environmental stewardship and sustainability has come center stage. The general public and environmental advocacy groups are demanding proactive steps toward conservation and green processes. In response, the technology industry has launched collaboratives such as The Green Grid and multiple corporate initiatives such as Sun Microsystems' ECO Responsibility Initiative, HP's ECO Solutions, and IBM's Project Big Green.

To meet the continued growth in the number of IT/HPC users [2] and capability demands of those users [3], new energy-focused design paradigms are required. We introduce the Grid Heating (GH) paradigm. GH recognizes that, despite evolving low power architectures, demands for increased capability will drive up power consumption toward economic limits on par with capital equipment costs. In contrast to the design of a single facility for centralized compute infrastructure, GH capitalizes on grid and virtualization technologies to distribute compute infrastructure (along with exhaust heat as shown in Fig. 1) in-line with existing municipal and industrial thermal requirements.

Figure 1 – Thermal images from one of Notre Dame's data centers.

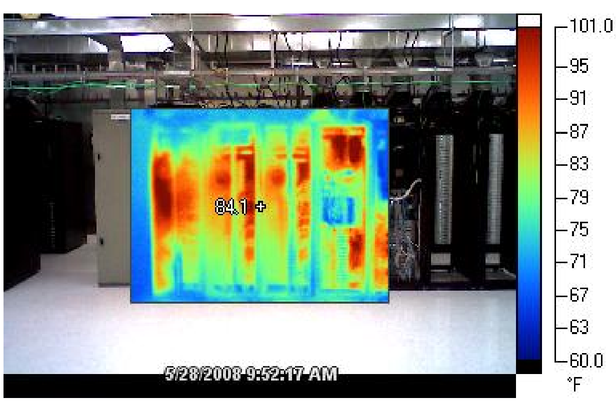

The Principles of Grid Heating

GH's fundamental core is the recognition that computational infrastructure can be strategically grid distributed inline with municipal facilities and industrial processes requiring the thermal byproduct. Frameworks built upon this core reduce or remove cooling requirements and realize cost sharing on primary utility expenditures. The technology builds upon efficiency improvements for the traditional centralized data center made by multiple organizations such as the High-Performance Buildings for High-Tech Industries Team at LBNL [4], the ASHRAE Technical Committee 9.9 for Mission Critical Facilities, Technology Spaces, and Electronic Equipment [5], and the Uptime Institute [6,7].

Individual data centers have reutilized the thermal byproduct to the benefit of their own facility [8]. However, to utilize all of the thermal energy effectively year-round, a grid distributed approach is desirable. GH models recognize that the transformation and/or transportation of the waste heat quickly reduces efficiency, and therefore targets the distribution and scale of each heating grid node to match the geographic and thermal requirements of the target heat sink. GH deployments must also address the physical and practical considerations of operating a computational infrastructure. A sample list of physical considerations must include basic hardware operational requirements of temperature, humidity, and air particulate. For example, practical factors include suitable bandwidth/data locality, security, system administrator access, reliability/redundancy, and acoustics.

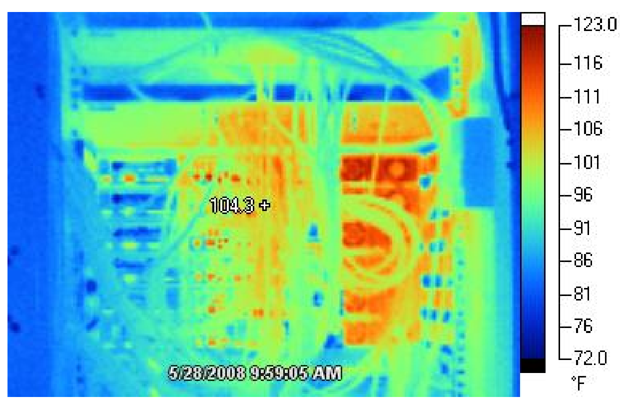

The GH Cluster Framework

This work focuses on a GH Cluster instantiation of a grid heating framework. A GH Cluster framework is based upon a scalable and tightly interconnected multi-node configuration similar to the traditional rack configurations. A full 40-unit rack, partial rack, or multiple racks could be selected to best match the annual heating requirement of a particular facility. The GH Cluster could be a stand-alone enclosure collocated with an industrial heat sink or facility. Waste heat transferred via air or liquid media is directly utilized as shown in Fig. 2. Examples include air heating of a greenhouse, pre-heating hot water for a hospital, or primary heat for a water treatment plant.

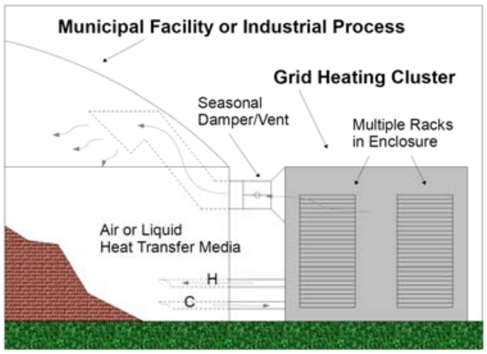

Figure 2 - GH Cluster schematic

For the purposes of this paper, we will introduce and focus on the Sustainable Distributed Data Center (SDDC), a specific GH Cluster prototype. The SDDC deployment is the current phase of Notre Dame's partnership with the City of South Bend and South Bend Botanical Society as presented in the following experimental section.

by the fans, hard drives, etc.); a thermal energy surplus on the order of the input electrical energy will be locally available. We took this basic observation and validated it with a more challenging experiment (*figure 3*). The goal was to provide thermal control of a multiuser workspace through dynamic resource utilization and job scheduling.

Multiple test scenarios were designed and evaluated to capture the cluster's ability to meet and oscillate about temperature targets. In *figure 3*, we present one such test scenario starting with all machines powered off and a room temperature just above 78°F. Once the room had equilibrated to a stable temperature with machines idling, a test script was run to adjust target room temperatures and prompt the grid heating cluster nodes to dynamically accept jobs. As shown in *figure 3*, the target temperature was updated every six hours from 86° to 88° to 90° to 88° and then down to 70°F, which effectively put the machines back in the idle condition. The oscillations in the figure demonstrate the thermal variance in the room as jobs moved from suspend to resume status. Note the graph also reflects the properties of the machines' maximum thermal capacity. Given the successful validation of fine-grained IT exhaust temperature control, we then began the endeavor to build a multi-institution collaboration on larger grid heating prototypes.

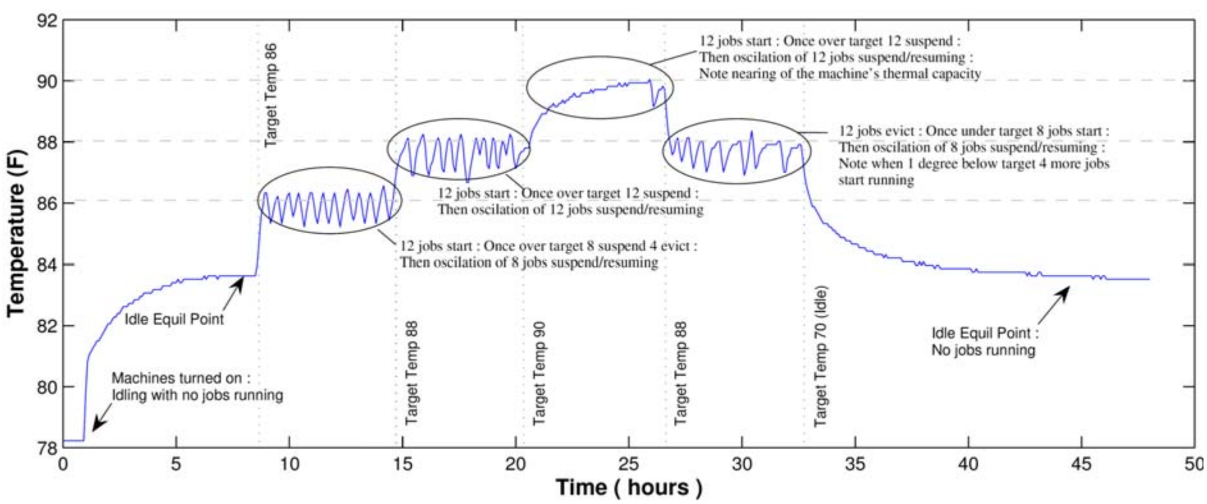

Figure 3 – Dynamic temperature control

3. Experimentation

Temperature Control Validation

Given the basic observation that the electrical energy consumed by modern computational infrastructure is predominantly transformed into thermal energy (we acknowledge minor transformations into mechanical work

Greenhouse Deployment

We proceeded with our first grid heating deployment at the City of South Bend Botanical Conservatories and Greenhouse (SBGH). The SBGH (Fig. 4) is a traditional glass greenhouse facility with multiple interconnected structures of which the most recently constructed was in the

1970s. Municipal funds are not available for capital upgrades, and rising annual natural gas heating costs grew to over $115,000 in 2006. Primary heat is provided via natural gas boilers that provide steam heat to the majority of the building. The target goal is to provide baseline GH Cluster heating capability to the facility in place of fossil fuel heat, using the existing boilers on a limited basis during peak winter weather. The system is consequently designed to utilize free cooling for the majority of the warmer months.

Figure 4 – South Bend Botanical Conservatory and Potawatomi Greenhouse

The University of Notre Dame has partnered with the City of South Bend on a three-phase GH deployment to provide direct thermal benefit to the SBGH and year round operations of our computational equipment. Phase 1 has been completed successfully, and entailed the initial deployment of sufficient computational infrastructure to enable remote submission of scientific simulations from the Notre Dame campus to the SBGH. Primary phase one components included a traditional compute rack (*figure 5*) located within the desert collection dome structure. The rack was configured with traditional 1U compute nodes similar to that used in the thermal control tests. Network interconnectivity to the Notre Dame campus was provided by a local ISP. Successful completion was marked by the delivery of byproduct heat, while facilitating molecular dynamics simulations performed by the LCLS research group to accelerate molecular trajectory generation [9].

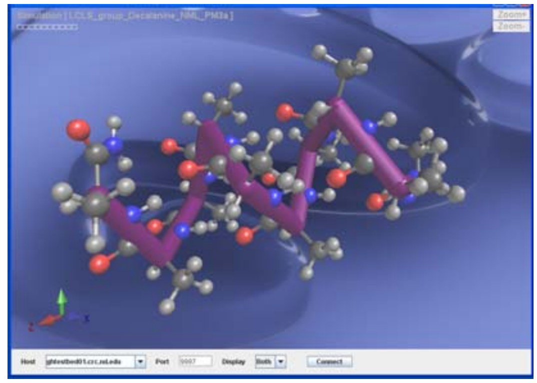

Figure 5 – Phase 1 Grid Heating Cluster prototype and successful remote simulations

We are now in phase two of the deployment with a focus on scaling and securing the compute infrastructure via a Sustainable Distributed Data Center prototype. The SDDC is a 20' wide x 8' long x 8' high containerized data center facility. Using the DOE Data Center Energy Profiler yielded a potential DCiE of 0.81. As of December 2008, SDDC fabrication and necessary site work (utility power, fiber network access, and foundations) at the greenhouse has been completed. The SDDC is slotted for commission in Jan 2009 during the peak winter demand. The prototype is outfitted with a 208 Volt 225 Amp panel providing up to 45kW of power, or an equivalent 150,000 BTU/h of heat. This prototype configuration transforms an annual data center electrical power bill of $35,000 into 150,000 BTU/h of useable heat for the greenhouse.

For year round SDDC operations we determined the suitability of free cooling and supplemental humidification by reviewing respective information from the National Climatic Data Center [10] as shown in *figure 6*. Observe that the average highs do not exceed the maximum ASHRAE allowable inlet face temperature [5]. For the small number of hours annually which are above average, the SDDC is equipped with a modest traditional HVAC apparatus. Heat evacuation to the greenhouse is handled by a ventilation/transmission fan in the SDDC. Proper mixing of inlet and inside air during the cold low humidity winter months will be managed by SDDC outside air louvers. Recent findings on particulate load by LBNL [11] indicate outside air particulate loads are often well below EPA and manufacturer guidelines; however, standard filters and readings relative to the local vicinity are planned. This phase two work also includes a connection to the St. Joe Valley Metronet fiber-optic network shared by South Bend and Notre Dame.

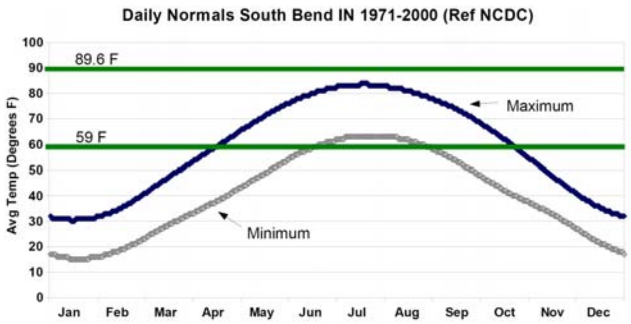

Figure 6 – Daily temperature normals with ASHRAE allowable max/min temperatures

Benefits Derive from Holistic View

It is important to note evolving commercial applications in grid, utility, and cloud computing [12, 13, 14] that would directly benefit from a grid heating configuration. As the computational infrastructure configuration and locality is obscured from the end user, the flexibility to distribute and configure grows, allowing for additional economic and environmental optimizations. Along these lines, the growing acceptance of virtualization [15, 16, and 17] in commercial applications will also allow greater flexibility in the design and deployment of infrastructure.

While these grid frameworks are evolving, a large body of work has studied the problem of managing energy, heat, and load in large centralized data centers. In addition to those works cited earlier in this paper, Schmidt et al [18, 19] provide a good current overview of the mechanical issues of cooling units, heat sinks, fans, and so forth. Given an adequate mechanical infrastructure, several server management techniques can also be applied to reduce energy costs. For example, inactive servers can be shut down, or loads migrated as more energy efficient hardware become idle/available. Chase [20] and Bradley [21] describe techniques for balancing performance, cost, and energy in this situation. To avoid hot spots, it is necessary to map the relation between components and heat [22], and then shape loads so as to evenly distribute the heat [23]. Further, large institutions such as the University of Illinois and NCSA are taking a holistic look at their entire campus utilities infrastructure to efficiently operate their data center. Despite the efficiency benefits of these new techniques, Patel et al. [24] report that a typical data center still consumes about as much energy for cooling as it does for productive work. The advances in efficient traditional infrastructures and a GH framework will serve in tandem to provide greater computational capability while reducing economic and environmental costs.

Acknowledgments

This work was partially supported by a Federal NWICG grant. Computational resources were made available by the Center for Research Computing and the Department of Computer Science and Engineering at the University of Notre Dame. We thank our many municipal collaborators at the City of South Bend Botanical Garden and Greenhouse. The authors are appreciative of insightful conversations with Miron Livny, Christopher Sweet, Ed Bensman, Rich Sudlow, and Curt Freeland.

About theUniversity of Notre Dame

The University of Notre Dame, founded in 1842 by Rev. Edward F. Sorin, C.S.C., of the Congregation of Holy Cross, is an independent, national Catholic university located in Notre Dame, Ind., adjacent to the city of South Bend and approximately 90 miles east of Chicago. The University is organized into four colleges—Arts and Letters, Science, Engineering, and the Mendoza College of Business—the School of Architecture, the Law School, the Graduate School, six major research institutes, more than 40 centers and special programs, and the University library system. Building on its strong tradition of undergraduate teaching and graduate studies, Notre Dame is now strengthening its research enterprise with dramatically increased resources and new state–of–the–art facilities. External funding has doubled since 2000, now standing at $83 million annually. This is a rare achievement for a university without a medical school, but the growth trend continues.

About the Author(s)

Paul R. Brenner, PhD, PE is a scientist and engineer for the University of Notre Dame's Center for Research Computing. He specializes in energy efficient data center design for high performance computing and parallel algorithms for scientific applications. Paul is also an engineering officer in the USAF Reserves with DOD data center design experience in Afghanistan and at Wright-Patterson AFB's MSRC. For more information, please e-mail Mr. Brenner at: brenner.6@nd.edu

Douglas Thain, PhD is Assistant Professor of Computer Science and Engineering at the University of Notre Dame. His research interests focus on harnessing large scale computing systems such as clusters, clouds, and grids to attack large problems in science and engineering. Dr. Thain received his Ph.D. from the University of Wisconsin in 2004. For more information, please e-mail Mr. Thain at: ***dthain@nd.edu***

Dewitt Latimer, PhD is the Assistant Vice President, Assistant Provost, Chief Technology Officer (CTO), and Deputy CIO of the University of Notre Dame. In this capacity, he has operational oversight of the IT infrastructure and serves as the University's chief IT strategist. Dr. Latimer has been a leader in higher education's national networking efforts for more than a decade and has played a key role in shaping higher education's position on national information technology policy issues. For more information, please e-mail Mr. Latimer at: *dewitt@nd.edu*

About the Uptime Institute

Uptime Institute is a leading global authority on data centers. Since 1993, it has provided education, consulting, knowledge networks, and expert advisory for data center Facilities and IT organizations interested in maximizing site infrastructure uptime availability. It has pioneered numerous industry innovations, including the Tier Classification System for data center availability, which serves as a de facto industry standard. Site Uptime Network is a private knowledge network with 100 global corporate and government members, mostly at the scale of Fortune 100-sized organizations in North America and EMEA. In 2008, the Institute launched an individual Institute membership program. For the industry as a whole, the Institute certifies data center Tier level and site resiliency, provides site sustainability assessments, and assists data center owners in planning and justifying data center projects. It publishes papers and reports, offers seminars, and produces an annual Green Enterprise IT Symposium, the premier event in the field focused primarily on improving enterprise IT and data center computing energy efficiency. It also sponsors the annual Green Enterprise IT Awards and the Global Green 100 programs. The Institute conducts custom surveys, research and product certifications for industry manufacturers.

Uptime Institute, Inc.
2904 Rodeo Park Drive East
Building 100
Santa Fe, NM 87505-6316
Corporate Offices: 505.986.3900
Fax: 505.982.8484
uptimeinstitute.org

References

[1] U.S. Environmental Protection Agency: Report to congress on server and data center energy efficiency public law 109-431. Technical report, U.S. E.P.A (Aug 2007)

[2] Tabor Research: Intersect360(tm) Market Advisory Service: Buyer Views. Technical report (August 2007)

[3] Dongarra, J.J., Meuer, H.W., Strohmaier, E.: 29th top500 Supercomputer Sites. Technical report, Top500.org (June 2007)

[4] Mills, E., Shamshoian, G., Blazek, M., Naughton, P., Seese, R., Tschudi, W., Sartor, D.: The business case for energy management in high-tech industries. Journal of Energy Efficiency **1** (2007) 1–16

[5] 9.9, T.: High Density Data Centers - Case Studies and Best Practices. ASHRAE, Atlanta, GA (2008)

[6] Brill, K.G.: Data center energy efficiency and productivity. Technical report, The Uptime Institute (2007)

[7] Brill, K.G.: Special report: Energy efficiency strategies survey results. Technical report, The Uptime Institute (April 2008)

[8] Intel: Data Center Heat Recovery Helps Intel Create Green Facility. Intel Information Technology Brief (December 2007)

[9] Sweet, C., Petrone, P., Pande, V., Izaguirre, J.: Normal mode partitioning of Langevin dynamics of biomolecules. J. Chem. Phys. (2008)

[10] Center, N.C.D.: Daily normals of temperature, precipitation, and heating and cooling degree days for south bend, in. Climatography of the United States NO.84 (1971-2000)

[11] Tschudi, W.: 24x7 energy efficiency presentation. Technical report, Lawrence Berkley National Laboratory (Feb 2007)

[12] Eilam, T., Appleby, K., Breh, J., Breiter, G., Daur, H., Fakhouri, S.A., Hunt, G.D.H., Lu, T., Miller, S.D., Mummert, L.B., Pershing, J.A.,Wagner, H.: Using a utility computing framework to develop utility systems. IBM Systems Journal **43** (2004) 97–120 16

[13] Raghavan, B., Vishwanath, K., Ramabhadran, S., Yocum, K., Snoeren, A.C.: Cloud control with distributed rate limiting. In: SIGCOMM '07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications, New York, NY, USA, ACM (2007) 337–348

[14] Mitchell, W.: Data center in a BOX. Scientific American **297** (Aug 2007) 90–93

[15] Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., Neugebauer, R., Pratt, I., Warfield, A.: Xen and the art of virtualization. In: SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles, New York, NY, USA, ACM (2003) 164–177

[16] Figueiredo, R., Dinda, P., Fortes, J.: Resource virtualization renaissance. Computer **38** (2005) 28–31 [17] Uhlig, R., Neiger, G., Rodgers, D., Santoni, A., Martins, F., Anderson, A., Bennett, S., Kagi, A., Leung, F., Smith, L.: Intel virtualization technology. Computer **38** (2005) 48–56

[18] Schmidt, R., Cruz, E., Iyengar, M.: Challenges of data center thermal management. IBM Journal of Research and Development **49** (July 2005)

[19] Schmidt, R., Iyengar, M.: Best practices for data center thermal and energy management - review of literature. ASHRAE Transactions **113** (Jan 2007)

[20] Chase, J., Anderson, D., Thakar, P., Vahdat, A., Doyle, R.: Managing energy and server resources in hosting centers. In: Symposium on Operating Systems Principles. (2001)

[21] Bradley, D., Harper, R., Hunter, S.: Workload based power management for parallel computer systems. IBM Journal of Research and Development **47** (2003) 703–718

[22] Moore, J., Chase, J., Ranganathan, P.: Weatherman: Automated, online, and predictive thermal mapping and management for data centers. In: IEEE International Conference on Autonomic Computing. (June 2006)

[23] Moore, J., Chase, J., Ranganathan, P.: Making scheduling "cool": Temperature aware workload placement in data centers. In: USENIX Annual Technical Conference. (April 2005)

[24] Patel, C.D., Bash, C.E., Sharma, R., Beitelmal, M.: Smart cooling of data centers. In: Proceedings of IPACK. (July 2003)