Sustainable Data Centers: Enabled by Supply and Demand Side Management

Prith Banerjee, Chandrakant D. Patel, Cullen Bash, Parthasarathy Ranganathan
Hewlett-Packard Laboratories, 1501 Page Mill Road, Palo Alto, CA-94304
{firstname.lastname}@hp.com

ABSTRACT

The environmental impact of data centers is significant and is growing rapidly. Servers alone in the US consumed 1.2% of the nation’s energy in 2005, according to the EPA. In the following year, the EPA found that the cost of energy rose by 10%. However, there are many opportunities for greater efficiency through integrated design and management of data center components. To that end, we propose a sustainable data center that replaces conventional resource delivery models with a framework centered around the supply and demand side management of *all* data center resources including IT, power and cooling. We have identified five elements for achieving this vision: data center scale lifecycle design, flexible and configurable building blocks, pervasive cross-layer sensing, knowledge discovery and visualization, and autonomous control. We describe these principles and provide selected results that quantify the potential for savings.

Categories and Subject Descriptors

C.0 [**Computer System Designs**]: General – *system architectures;* K.4 [**Computing milieux**]: General – *computers and society.*

General Terms

Design, Management, Economics

Keywords

Sustainability, Exascale, Data Centers

1. INTRODUCTION

A recent study found that IT is responsible for about 2% of global greenhouse gas emissions [1], about as much as the aviation industry. Furthermore, it projected that that this share would double by the year 2020. Increasing environmental concern and regulatory action will soon force a paradigm shift in how IT solutions are designed and managed across their lifecycles. Data centers are a prominent component of this impact, as well as the fastest growing.

To turn this crisis into an opportunity, we propose the development of a suite of technologies for a *sustainable data center* (SDC). Specifically, we propose developing technologies to reduce the environmental footprint of the data center to such an extent that the services offered from such a facility would be more environmentally friendly than conventional services offered within the physical infrastructures. The SDC would be enabled through supply and demand side resource management:

Supply Side

- Design of physical infrastructure with focus on lifecycle engineering and management, and the available energy required to extract, manufacture, operate and reclaim components;
- Utilization of local resources to minimize destruction of available energy in transmission, and construction of transmission infrastructure.

Demand Side

- Provisioning data center resources based on the needs and service level agreement of the user through use of flexible building blocks, pervasive sensing, knowledge discovery and policy based control

Developing and demonstrating SDC requires the multi-disciplinary collaboration of mechanical engineers, electrical engineers, computer scientists, and others. The *hardware* infrastructure of the data center consists of thousands of servers hosting revenue-generating services, interconnected with each other and the outside world via networking equipment, and relying on storage devices for persistent data. These hardware elements are managed by a datacenter-wide *software* stack that spans the platforms and virtualization layers. The data center also

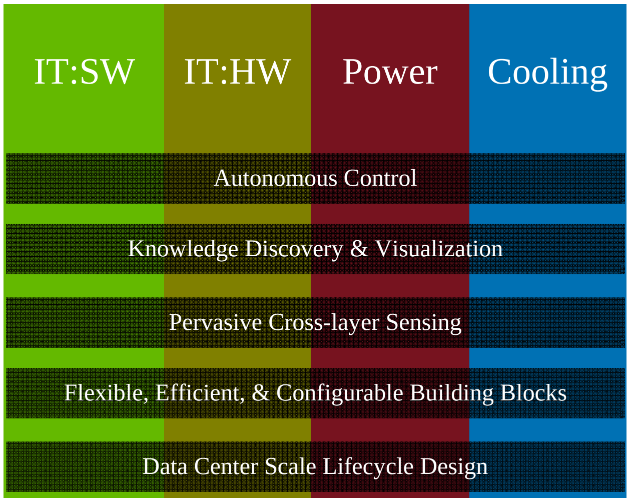

Figure 1: Sustainable Data Center Key Elements

DAC’09, July 26-31, 2009, San Francisco, California, USA

has a *power* infrastructure that feeds electricity to all of the equipment, and a *cooling* infrastructure that removes heat from the equipment. The economic and environmental burden of the latter two infrastructures often equals or exceeds that of the compute infrastructure [2].

Figure 1 shows the five key elements of SDC along with the four infrastructure verticals with which they interact. The foundation of a sustainable data center is *lifecycle-based design* based on sustainability, service-level agreement (SLA) and total cost of ownership (TCO) goals through the entire data center lifecycle (i.e. build, operation and end-of-life). The data center is designed with *flexible efficient elements* that enable the manipulation of resources within each infrastructure. Examples can include a multiplicity of micro-grids for power generation, variable air-conditioning systems for cooling resource generation and server architectures that are scalable and power proportional. *Pervasive sensing* that cross-cuts the platforms and management layers is necessary in order to gather the information necessary during operation. *Knowledge discovery and data analytics* can then be employed to evaluate resource needs within each infrastructure and ensure the operational requirements are being met. The final element is an *autonomous controller* that holistically optimizes the energy for a given end-user SLA from both the supply and demand side. Finally, aside from cutting across all infrastructures, each element also cuts across multiple academic disciplines and will require close collaboration with historically separate research groups.

2. DATA CENTER LIFECYCLE DESIGN

Sustainable operation must begin with sustainable design. As the first and foundational element of SDC, we propose a Data Center Synthesizer to automate and optimize data center design to facilitate meeting sustainability, SLA and TCO goals. This process begins with a description of the services to be provided and any constraints on the design. At HP Labs, we have developed an SLA decomposition approach that can transform Service Level Objectives (SLOs) for multiple multi-tier applications into the computing, power and cooling resources required to deliver those services, a complex and challenging task since the joint requirements are typically less than the sum of the individual requirements [3-5]. Power distribution and cooling equipment are selected to meet the operating characteristics of the computing equipment, taking into consideration real-time performance management that will be incorporated. The next step is projection of a candidate data center solution based on the selected computing, power and cooling equipment. This step results in a complete model and physical layout of the data center, incorporating all computing, networking and storage devices and the power and cooling infrastructure required to support them. Future enhancements will include provisions to specify anticipated growth rates to enable the equipment selection to be optimized to accommodate that growth, plus additional specifications and operational policies that might affect equipment selection.

The candidate data center is analyzed for sustainability criteria, performance, availability, thermal characteristics and TCO. The sustainability criteria include power consumption for computing and cooling, carbon emissions, heat load and embedded exergy, which accounts for the loss in exergy (usable energy) for all components and manufacturing processes associated with each piece of equipment, from raw materials extraction to finished products. If the design does not meet key customer requirements or if the design can be improved, a re-synthesis process uses the results to modify the specifications and generate another candidate data center. Iteration continues in a fully-automated loop until an optimal solution is reached.

3. FLEXIBLE BUILDING BLOCKS

As an example of innovations in the architectural building blocks driving energy efficiency and sustainability, we next discuss our work on microblades and megaservers [6]. Specifically, our work seeks to understand and design energy-efficient next-generation servers for emerging cloud workloads. We put together a detailed evaluation infrastructure including a new benchmark suite for internet sector workloads, and detailed performance, cost, and power models to quantitatively characterize bottlenecks. Based on the insights from a systematic cost analysis, we propose a new design involving modular cost-effective server blocks ("microblades") to build large powerful computing environments ("megaservers").

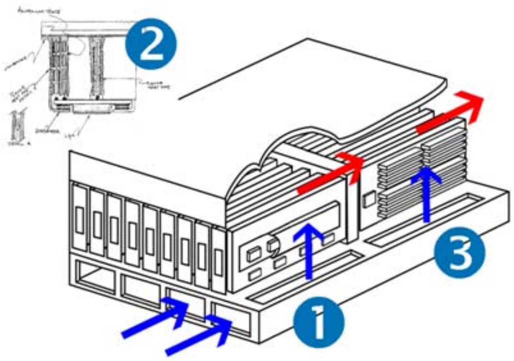

Figure 2: Microblade server.

Our work incorporates volume non-server-class components (CPU, flash) in disaggregated building blocks with ensemble-level provisioning (memory, storage) and novel form-factor, interface, infrastructure optimizations, and novel packaging of power and cooling at the micro-blade level. Our evaluation shows that this approach has the potential for dramatic improvements compared to the state-of-the-art, improving energy efficiency by factors of 4 to 6 in some cases.

As ongoing work, we are currently looking at extending the ideas proposed in this work further to look at co-designed packaging and system architecture, specifically exploring disaggregated blade designs for future dematerialized datacenters.

4. PERVASIVE SENSING

The next element discussed in Figure 1 is pervasive cross-layer sensing and monitoring. In current data centers, the degree of sensing is not well distributed across the infrastructure layers. While the IT layers (particularly the software layer) contains an ubiquitous sensing network, sensing in the power and cooling layers is sparse. We have shown that pervasive sensing in the environmental layers, particularly the cooling layer, is necessary for efficient resource distribution [15][16].

Furthermore, aggregation and coordination of data across the infrastructure layers is required for integrated operation. A key element of this is an underlying architecture for coordinated monitoring and management. To enable this, we have proposed vManage [7], a solution to loosely couple platform and virtualization management and facilitate monitoring and management coordination in data centers. Our solution is comprised of registry and proxy mechanisms that provide unified monitoring and actuation across platform and virtualization domains, and coordinators that provide policy execution for better VM placement and runtime management, including a formal approach to ensure system stability from inefficient management actions. Figure 3 illustrates our approach.

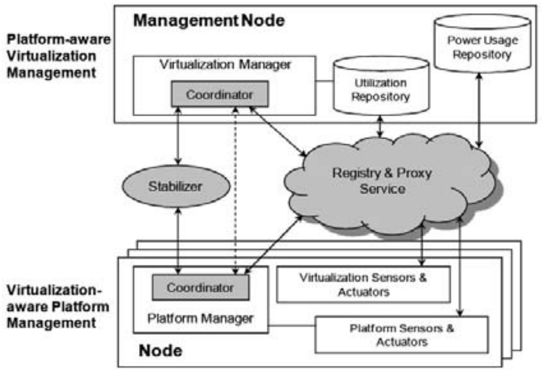

Figure 3: Cross-layer monitoring and management

5. KNOWLEDGE DISCOVERY

Knowledge discovery and data analysis tools are utilized to evaluate the data generated during operation. Knowledge discovery techniques generate insights and knowledge from disparate data sources across the facility. Data center facilities and systems produce huge amounts of data, on the order of megabytes per server per day, including environmental sensor data (e.g., temperature), operational states (e.g., system utilization), and workload information (e.g., user requests). Since the sheer volume of data precludes manual inspection, automated data mining and knowledge discovery techniques are applied to identify trends, patterns, and models for more efficient operation of a data center. In particular, we focus on three key areas: (1) detecting uncorrelated data center events using Principal Component Analysis [8]; (2) visual analytics for thermal state management [9], and (3) temporal data mining of common motifs or patterns for enhancing operational efficiency [10]. Among other benefits, the tools will help optimize resource utilization, predict events, manage growth and improve reliability.

6. AUTONOMOUS CONTROL

The final element of Fig. 1 considers the autonomous control of resources within the infrastructure verticals. Since available energy supply has direct impact on the exchequer, it is sensible to manage the available energy supply based on demand and improve the effectiveness of utilization of available energy for the total IT delivery process at each step. Managing onsite power generation, power delivery, power storage and cooling infrastructure operation to provide critical services to the IT demand layer are critical to sustainable operation.

Figure 4 shows the sustainable data center ecosystem. During runtime, the data center is managed to reduce resource consumption from the IT, power and cooling infrastructures. Service Level Agreements (SLAs) are used to define the operational requirements of the infrastructure based on performance, TCO and sustainability criteria. These criteria are applied to all aspects of data center operation from, for example, workload placement and server consolidation to the distribution of cooling resources based on the provisioned workload. Apart from the distribution of resources like power and cooling, the generation of these resources is also considered in order to take advantage of efficient sources that may have time-varying attributes (like photovoltaics for power generation, or bringing in outside air for cooling during certain times of the day or year).

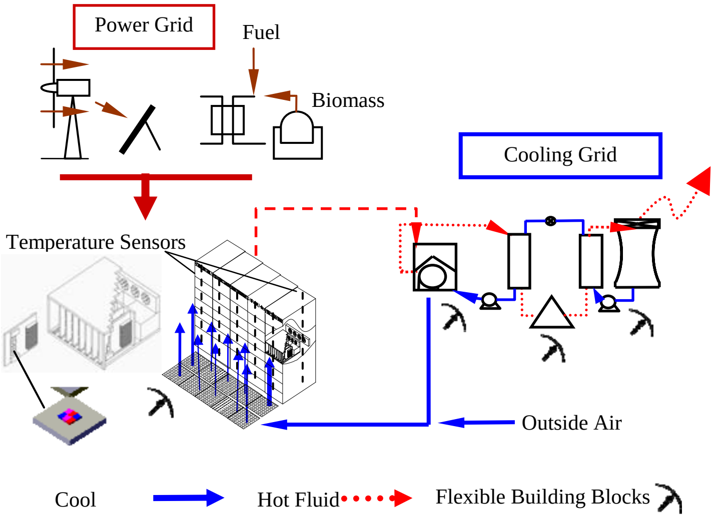

Figure 4: Supply and demand side management

A key characteristic of autonomous control in SDC is that control is integrated across the infrastructure verticals. Since sensing is pervasive and information is shared, resource control decisions are made that consider the entire data center ecosystem rather than on a single portion of that ecosystem. Workload schedulers provide an example. Currently, workload may be scheduled according to performance and availability requirements with supply side attributes (i.e. power and cooling generation and distribution) ignored. However, workload placement location can impact overall data center operational cost. We have developed a method for ranking workload placement according to cooling efficiency in the data center. We've incorporated this ranking into both batch and real-time workload schedulers [11]. We have experimental evidence that quantifies the impact of this integration and shows how the automated migration of a virtualized workload from one set of nodes located in an inefficient thermal zone to a more efficient zone reduced data center energy usage by 27% and improved available thermal capacity by 22% [12].

Furthermore, the flexible building blocks described in Section 2 provide sufficient granularity to provide resources where and when they are needed rather than relying on over-provisioning to meet operational requirements. We recently consolidated 14 laboratory data centers into one large site in Bangalore, India, and applied a dynamic control system that adjusts the utilization of the air conditioning system based on 7,500 temperature sensors deployed throughout the data center. The 40 per cent savings in cooling power consumption achieved in this facility translates to annual savings of approximately $1.2 million, when compared to the conventional approach [13].

Finally, while several past solutions have individually evaluated different techniques to address separate aspects of this problem, in hardware and software, and at local and global levels, there has been not much corresponding work on architectural approaches to coordinating all these solutions. In the absence of such coordination, these solutions are likely to interfere with one another, in unpredictable (and potentially dangerous) ways. We discuss our solution that addresses this problem [14]. We make two key contributions. First, we propose and validate a power management solution (illustrated in Figure 5) that coordinates different individual approaches. Using simulations based on 180 server traces from nine different real-world enterprises, we demonstrate the correctness, stability, and efficiency advantages of our solution. Second, using our unified architecture as the base, we perform a detailed quantitative sensitivity analysis and draw conclusions about the impact of different architectures, implementations, workloads, and system design choices.

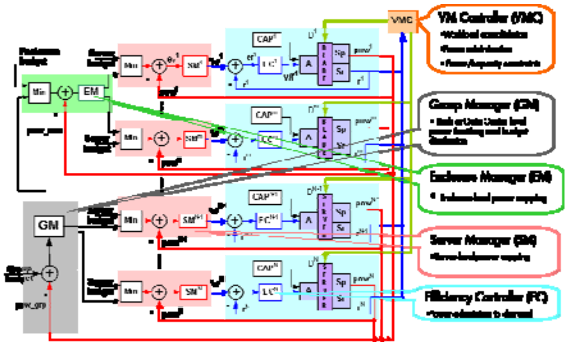

Figure 5: Coordinated policy management across different layers

7. CONCLUSIONS

The environmental impact of IT is a growing concern worldwide. Data centers contribute a large fraction of this impact. This concern, coupled with increasing government interest in regulating data center resource consumption, is resulting in a new approach to management of data centers. This paper describes an integrated life-cycle based approach to the design and management of sustainable data centers enabled by supply and demand side management of resource consumption.

8. ACKNOWLEDGEMENTS

The authors would like to collectively acknowledge the members of the Sustainable IT Ecosystem and Exascale Computing Laboratories within Hewlett-Packard Labs, many of whom contributed directly to the research described herein.

9. REFERENCES

[1] Mingay, S., 2008, “IT’s Role in a Low Carbon Economy,” Keynote Address, *Greening the Enterprise 2.0 Conf.*

[2] Belady, C. L., 2007, "In the Data Center, Power and Cooling Costs More Than the IT Equipment It Supports," Electronics Cooling Mag., **13**(1)

[3] G. (John) Janakiraman, Jose Renato Santos, Yoshio Turner, “Automated Multi-Tier System Design for Service Availability,” HP Labs Technical Report, HPL-2003-109.

[4] G. (John) Janakiraman, Jose Renato Santos, Yoshio Turner, “Automated System Design for Availability,” HP Labs Technical Report, HPL-2004-54.

[5] Yuan Chen, et al., “A Systematic and Practical Approach to Generating Policies from Service Level Objectives,” Proceedings of the 11th IFIP/IEEE International Symposium on Integrated Network Management (IM 2009), accepted.

[6] Understanding and Designing New Server Architectures for Emerging Warehouse-Computing Environments, Kevin Lim, Parthasarathy Ranganathan, Jichuan Chang, Chandrakant Patel, Trevor Mudge, Steve Reinhardt, Proceedings of the International Conference on Computer Architecture, 2008

[7] vManage: Loosely Coupled Platform and Virtualization Management in Data Centers, S. Kumar et al, Proceedings of the International Conference on Autonomic Computing, June 2009

[8] Sharma, Marwah, Lugo, Bash, “Application of data analytics to heat transfer phenomena for optimal design and operation of complex systems” 2009 ASME summer Heat Transfer Conference, HT2009-88550

[9] Hao, Sharma, Dayal, Patel, Kleim, "Application of Visual Analytics For Thermal State Management in Large Data Centers", submitted to Eurographics/ IEEE-VGTC Symposium on Visualization (2009)

[10] M. Marwah, R. Sharma, L. Bautista, and W. Lugo. Stream mining of sensor data for anomalous behavior detection in datacenters. Technical Report HPL-2008-40, HP Labs, 2008.

[11] Bash, C.E., Forman, G., “Cool Job Allocation: Measuring the Power Savings of Placing Jobs in Cooling-Efficient Locations in the Data Center”, Usenix, San Jose, 2007.

[12] Bash, C.E., Hyser, C., Hoover, C., ”Thermal Policies and Active Workload Migration Within Data Centers”, InterPack, San Francisco, 2009.

[13] R.K. Sharma, R. Shih, C. Bash, C. Patel, P. Varghese, M. Mekanapurath, S. Velayudhan, V. Manu Kumar (2008), “On Building Next Generation Data Centers: Energy Flow in the Information Technology Stack,” Proceedings of 2008 ACM Compute Conference, Bangalore, India.

[14] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, and X. Zhu.”No “power” struggles: Coordinated multi-level power management for the data center” In Proc. of the 13th International ASPLOS, pages 48–59, Seattle, WA, 2008.

[15] C. D. Patel, C. E. Bash, R. Sharma, M. Beitelmam, and R. J. Friedrich. “Smart cooling of data centers” InterPack, pages 129–137, Kauai, Hawaii, July 2003.

[16] Cader, T., Sharma, R., Bash, C.E., Fox, L., Bhatia, V., Mekanapurath, M. “Operational Benefits of real time monitoring and visualization in data centers”, InterPack, San Francisco, CA, 2009.