GreenLight Technical Report 2011

Instrument Details

1.A.1. Tom DeFanti, Elazar Harel, Greg Hidley: Design and Implementation of the GreenLight instrument

Project GreenLight has built an instrument to measure the energy usage of computing systems under real-world conditions, with the ultimate goal of getting computer designers and users in the scientific community to re-think the way they do their jobs from an energy efficiency perspective. In Year 1 we completed the construction of the initial GreenLight instrument and preliminary planning for a DC powered server rack in the GreenLight instrument. In Year 2, the infrastructure team:

- Added additional computers to the Instrument
- Supplemented Sun MD sensors with per rack per outlet power sensors using Avocent PM3000 SNMP monitorable PDUs and per server IPMI interfaces to provide power measurements for each device in the instrument.
- Developed and published a data management system to give the research community ready access to the various environmental data generated by the Instrument's sensors.
- Ran baseline power experiments for general IT activities (data transfer, storage, and display)
- Provided access to the instrument to additional researchers
- Added dual mode AC and direct DC powered Intel and Sun Nehalem servers
- Added a cache server to the instrumented environment

- Added dual mode AC/DC powered facility model to the instrument including direct DC powered Intel and Sun Nehalem servers

The instrument consists of fully enclosed computer server facilities within a specially modified and instrumented modular sea storage container (Sun Microsystems Modular Datacenter, or SunMD). UCSD completed the infrastructure design and construction (mechanical, physical, electrical, telecommunications) to support the instrument's structural, power, cooling and data requirements. We installed the Sun MD, specified, ordered, received and configured a variety of networking and server equipment, deployed it in the first instrument and turned it over to applications and systems researchers in late Year 1. We also established a public GreenLight web presence at http://greenlight.calit2.net/ , an internal web based working environment at http://wiki.greenlight.calit2.net/greenlight/ and began the development of a GreenLight data management environment. In Year 2 we implemented the GLIMPSE, http://glimpse.calit2.net, monitoring and data access system, upgraded servers, and designed and installed a DC power distribution system in two of the SunMD racks. The above resources were regularly maintained and updated in Year 3 of the project.

The Sun Modular Datacenter (Sun MD) container houses 1 rack of network and monitoring equipment and 7 racks capable of supporting up to 40 1U servers each. Each rack is filled with various computing and networking equipment to support individual experiments. Custom measurement systems provide energy usage data within the instrument in real-time and include information on power usage and efficiency for the individual electronic devices as well as the server facility infrastructure. These measurements include details of consumption for computation and communication as well as for handling air flow, heat exchanger dissipation, chilled water flow and related thermal management. Energy data, both current and archival, are available in real-time over the Internet to assist energy efficiency research.

In Year 3, as of March 2010, the GreenLight instrument contains the following systems:

Rack 1

- Networking equipment. Optiputer and RCI network availability. Sun related management equipment. 6 10GigE network connections leave the SunMD.

Rack 2

- 24 x Nodes with two Intel Xeon E5440 Processors (2.83GHz) and 8GB memory (Compute nodes)
- 1 nodes with two Intel E5430 Processors (2.66GHz) and 32GB memory (large memory nodes)

Rack 3

- 8 x nodes with two Intel E5430 Processors (2.66GHz) and 32GB memory (large memory nodes)
- 7 x Nodes with two Intel Xeon E5440 Processors (2.83GHz) and 8GB memory (Compute nodes)

Rack 4

- 5 x GPGPU Nodes with dual nVidia 295 cards, two Intel Xeon E5440 Processors (2.83GHz) , and 8GB memory
- 1 nodes with two Intel E5430 Processors (2.66GHz) and 32GB memory (large memory nodes)
- 4 x Convey Nodes (FPGA nodes)

Rack 5

- 1 nodes with two Intel E5430 Processors (2.66GHz) and 32GB memory (large memory nodes)
- 2 x Sunfire X4540 Storage Units, 48TB raw capacity (Thumpers)
- 1 GPGPU Node with dual nVidia 295 graphics cards, two Intel Xeon E5440 Processors (2.83GHz), and 8GB memory.
- 1 HPDMNet server

Rack 6

- 9 x GPGPU Nodes with dual nVidia 295 graphics cards, two Intel Xeon E5440 Processors (2.83GHz), and 8GB memory.

Rack 7

- 5 x Sun Fire X4270 servers with two Intel Xeon X5570 Processors and 24GB memory (Nehalem compute nodes) – DC powered
- 7 x Intel 2600 servers with two Intel Xeon X5550 Processors and 24GB memory (Nehalem compute nodes) – DC powered
- 1 Sunfire X4540 Storage Units, 48TB raw capacity (Thumpers) – DC powered

Rack 8

- DC Power equipment, transformer, rectifier and DC control devices

Switches

- Each rack contains two Extreme Networks Summit X350-48ts switches
- 10Gb Optiputer connectivity is available via a Fujitsu XG2000 switch.

PDU/Monitoring

- Each rack contains Avocent PM3024 PDUs with per port power monitoring
- Intermapper Monitoring
- IPMI Monitoring
- DSView3 Monitoring – Avocent Display Tool

The construction of GreenLight is supporting ICT energy efficiency related research activities in a number of areas including computer architecture, software architecture, visualization, media content distribution, energy monitoring and modeling, and domain science applications in Biology, Geology and BioEngineering. It is our belief and project goal that the GreenLight Instrument is enabling these communities of research and application scientists, to come to understand, through this instrumentation, how to measure and then minimize energy consumption, to make use of novel energy/cooling sources coming online at UCSD, and to employ middleware that automates optimal choice of compute/power strategies. This allows domain application researchers to exploit the exponential improvements in silicon technology and provides new quantitative data to support engineering judgments on comparative "computational work per watt" across full-scale applications running on at-scale computing platforms. This is, in turn, helping to re-define fundamentals of systems engineering for green cyberinfrastructure.

Conserving energy, reducing "carbon footprints" and deploying sustainable or renewable energy sources are increasing priorities in public and private institutions, across all sectors of the economy, in the United States and abroad. Project GreenLight is providing the tools to enable the research and understanding of energy conservation in Information and Communications Technology (ICT), one of the world's largest growth industries and leading energy consumer. Given the pervasiveness of ICT in every industry today and its central role in the creation of new industries tomorrow, energy efficiency improvements in ICT will have worldwide and far reaching impact. Project GreenLight is now providing tools and data to help researchers and industry make such improvements and is providing users with information on the energy use of various types of computing clusters.

During Year 3, the project pursued further ways to optimize the work per watt, such as through energy-saving hardware acceleration and data caching, reducing DC/AC/DC conversion losses, increasing the use of renewable energy, and investigating WAN terrestrial transmission options. NLR and CENIC provide 5 x 10GigE WAN and campus network links available to Project GreenLight.

In Year 3, Project GreenLight investigated cache engines to improve streaming data performance. We installed a Cisco cache engine on our switch in support of GreenLight activities, which identifies repeat traffic as a function of the type of media obtained and port used. If remotely accessed data is determined to be static, the local cache hits, significantly reducing network traffic and remote processing. We investigated ways to estimate the resulting power savings and benchmark true performance gains, which visually are quite apparent. In the first tested application which is using remotely served Windows Media Server (WMS) data (the Blue Marble demo with WMS overlays), we saw 70% cache hits and therefore 70% fewer data transfers from the remote server and 70% reduced utilization (bandwidth reduction) of the switching fabric between the remote server and our visualization facility. We will come up with a broader data mix to test in Year 4 to identify the true strengths and weaknesses.

In Year 2 Project GreenLight developed and published web-based tools for energy and environmental monitoring and analysis. These are described in detail in Section 2.A.3. In Year 3, refined and used the tools and practices developed in Year 2 to run wider experiments in areas of multi-media (CineGrid), genomics and mass spectrometry and to help identify best work/watt configurations and maximize PUE at server, rack and facilities levels. We also focused on deployment staging in support of transferring power efficiency lessons learned into semi-production and production environments. As we bring new types of equipment into the instrumented SunMD, we rotate equipment out of the Instrument into semi-production status into the 2nd SunMD, and for production work into the Calit2 1101 server room where we can display to the public best of breed rack configurations.

As proposed and accepted, a No-Cost Extension for approximately 20% of the total project cost involves the in-depth study of visualization, virtual reality, and collaboration systems as extensions to the GreenLight Instrument. Clearly, visualization systems cannot be used within the Sun Modular Data Center itself, since is it not normally occupied by people. Prototypes of Avocent-monitored visualization and collaboration systems will be extended into a surround-stage scale visualization facility at Calit2 in Year 4. This extended instrumented facility, named ***Vroom*** and depicted in Figure 1 below, will foster deep understanding of the energy cost of collaboration over distance.

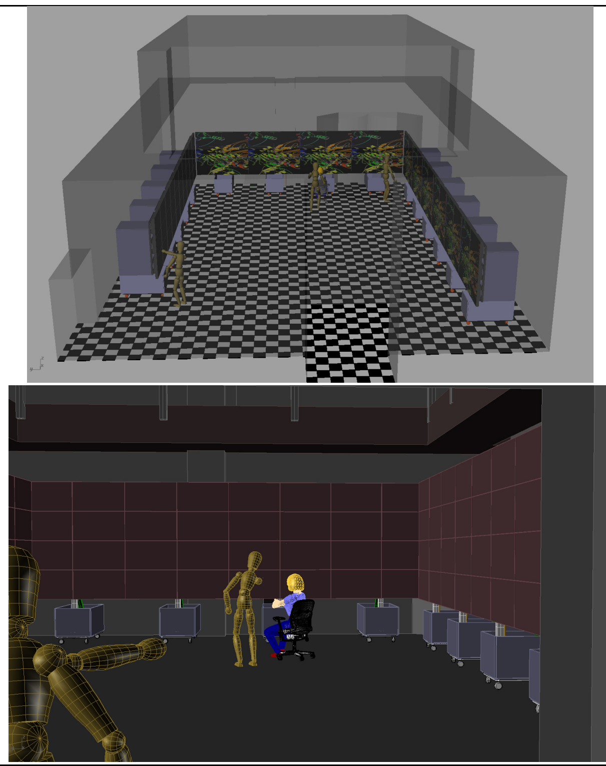

Figure 1: Early depictions or the Vroom facility

The containerized GreenLight instrument has been turned over to the computer science researchers since it has been filled as described above and is in a long-term reasonably stable state. This, of course, was the overall goal; publications are emerging. Long term plans for the GreenLight Instrument involve using it as the core facility for monitoring and analyzing UCSD campus energy consumption.