LBNL-XXXXX

ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

Case Study: JouleX Energy Management (JEM) Solution at Lawrence Berkeley National Laboratory (LBNL)

Dr. Magnus K. Herrlin, President, ANCIS Incorportated
Environmental Energy Technologies Division

July 2012

Acknowledgement

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231

Case Study: JouleX Energy Management (JEM) Solution at Lawrence Berkeley National Laboratory (LBNL)

Dr. Magnus K. Herrlin, President, ANCIS Incorporated

Case Summary

Data centers are complex environments that contain a large number of energy-consuming IT devices. Saving energy at the IT equipment level has a ripple effect on a number of energy-consuming support systems, so it can be a particularly effective energy-reduction strategy. For example, reducing heat produced by IT equipment translates into lower energy consumption for the cooling system, which can reduce capital and operating expenses, as well as greenhouse gas emissions.

This case study sought to demonstrate the capabilities of the JouleX Energy Manager (JEM), a promising automated IT assessment and management tool that can be useful in helping to manage the data center environment. The JEM tool enables energy savings by monitoring, analyzing, and managing the power state of network-connected IT equipment and support systems. The tool was tested in a data center at Lawrence Berkeley National Laboratory (LBNL).

Another focus for the study was to evaluate how well the JEM could interface with the DC Pro tool suite, developed by the U.S. Department of Energy (DOE). The DC Pro tool suite provides data center assessment tools that enable users to baseline energy use and identify energy efficiency opportunities in their data centers. To help identify opportunities related to IT equipment, an IT assessment tool is currently being developed in collaboration with the Green Grid organization. However, if commercially available energy assessment and management tools such as JEM can collect IT equipment data and communicate directly with the DC Pro tools, they could be used by data center energy practitioners or other data center stakeholders to help streamline the assessment process and identify better recommendations for efficiency improvements.

Several factors have sparked DOE's interest in such energy-management tools. First, DOE requires that each of its data centers utilize the DC Pro tools. Second, it is considering the use of IT assessment and management tools for DOE data centers. Third, the DOE Data Center Energy Practitioner (DCEP) program, which qualifies individuals to be data center energy assessors, is in the process of adding IT equipment assessment capability. Helping DCEP candidates learn about the capabilities of such automated tools will improve their ability to recommend efficiency measures related to IT equipment.

This study demonstrated the advantage of monitoring energy usage at the device level, so that operators and energy managers can understand where and how much energy is being used. This will lead to an understanding of where efficiency improvements can be made. The JEM can support this process by collecting equipment data automatically and managing the IT equipment for efficient use of energy. Although this study demonstrated the JEM on a limited scale, it provided an excellent picture of the energy used by IT devices, and this information revealed opportunities to save energy and reduce carbon emissions.

However, effective use of JEM at LBNL was impeded by three data access and security issues: (1) limited connectivity, (2) multiple private subnets, and (3) no remote monitoring access. As a result, the originally planned power monitoring scope was significantly curtailed, and power data had to be manually exported to JouleX for analysis.

In a heavily firewalled and segmented environment such as at LBNL, dynamic power monitoring (accessing the operating system) may be difficult when strict security policies are in place. Although JEM has numerous methods to acquire power data that do not require any privileged accounts, the additional requirements need to be planned well in advance with the security teams. For this class of tools, a well thought-out deployment plan will be needed to address scope and security needs.

Since the reduction of energy consumption and carbon emissions is a major goal for most data centers, tools like JEM could prove to be invaluable. A recommended next step is to install such tools in other DOE data centers, to gain better control over those environments and demonstrate a larger-scale implementation. A broader trial could also compare and contrast the capabilities of various systems. Knowledge and use of these tools would help DCEPs identify energy-savings opportunities in those data centers by more effectively gathering the necessary data for the DC Pro assessment process.

Introduction

Implementing green business initiatives has become a top priority for many organizations, and there is a growing awareness that a sustainable business is important to the financial bottom line. In many companies and institutions, reducing data center energy usage presents the largest opportunity to reduce its carbon footprint and save money.

Saving energy at the IT equipment level can have a significant impact because it has a ripple effect on a number of energy-consuming support systems. For example, reduced heat from the IT equipment translates into lower energy consumption for the cooling system, including chillers, pumps, and air-moving equipment. Such reductions can reduce capital and operating expenses, as well as carbon emissions. A data center with a "typical" PUE[1] of 1.8 would save almost twice the energy and money overall, compared to the energy saved at the IT load.

This case study sought to demonstrate the capabilities of the JouleX Energy Manager (JEM), an automated IT assessment and management tool. The JEM enables energy savings by monitoring, analyzing, and managing the power state of network-connected IT equipment and infrastructure systems ("devices" in short). It was used to perform power monitoring and an energy and utilization assessment of IT equipment in one of Lawrence Berkeley National Laboratory's (LBNL) data centers.

[1] Power Usage Effectiveness (PUE) is defined as Total Facility Energy divided by IT Equipment Energy. Data centers benchmarked by LBNL averaged a PUE of approximately 1.8.

Another focus for the study was to evaluate whether such tools could interface with the DC Pro tool suite, developed by the U.S. Department of Energy (DOE). The CD Pro tool suite provides data center assessment tools that enable users to baseline energy use and identify energy efficiency opportunities in their data centers. To help identify opportunities related to IT equipment, an IT assessment tool is currently being developed in collaboration with the Green Grid organization. However if commercially available energy assessment and management tools such as JEM are able to collect IT equipment data and communicate directly with the DC Pro tools, they could be used by data center energy practitioners or other data center stakeholders to help streamline the assessment process and identify better recommendations for efficiency improvements. Such tools are available now and can be utilized to supplement the existing DC Pro tools.

Several factors have sparked DOE's interest in such energy-management tools. First, DOE requires each of its data centers to utilize the DC Pro tool suite. Second, it is considering the use of assessment and management tools for DOE data centers. Third, the DOE Data Center Energy Practitioner (DCEP) program, which qualifies individuals to be data center energy assessors, is in the process of adding IT equipment assessment capability. It would be helpful for DCEP assessors to be cognizant of commercially available tools. Lawrence Berkeley National Laboratory studied JEM for possible use within its data centers, as well as to complement the DC Pro tool suite and the DCEP program. LBNL observed and analyzed the installation and operational experience of this tool. Valuable conclusions were drawn on JEM's ease of use, functionality, and its ability to assist in power monitoring and energy assessment in a data center with tight security policies.

The JouleX Energy Manager (JEM)

The JouleX Energy Manager has the capacity to help data center operators and energy managers reduce energy use in data centers by discovering, monitoring, analyzing, and managing servers and network-connected devices with no client-side agents, stubs, or hardware. It provides a global view of energy consumption for devices such as PCs, servers, VoIP devices, printers, network switching devices, IP power switches, and HVAC systems.

The JEM builds its own database of devices so that it is not necessary to create one manually. It uses an automated process to retrieve power and other information from each network-connected device. It will attempt to retrieve active power data if the device is instrumented. This automated process delivers very accurate information. If the device being monitored is not instrumented, JEM tries to use other methods to access power data: the JouleX dynamic method (tapping into the Operating System [OS]) and static method (using look-up tables) (see Figure 1).

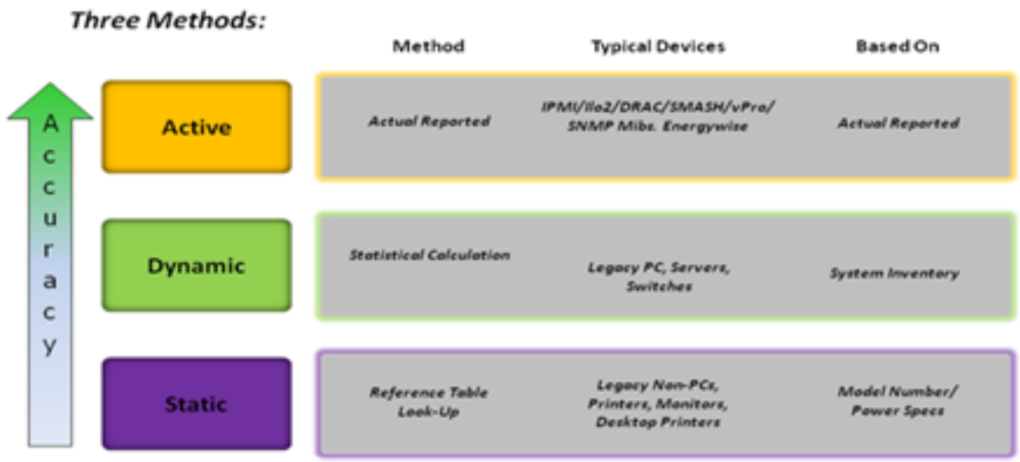

Figure 1: Main Methods to Retrieve Power Data

The JEM pinpoints underutilized and low-density servers (prime candidates for virtualization) and identifies dead and idle servers for retirement. It can also be used to allocate the right amount of power to the devices that need to perform productive work and minimize power when the devices are idling or operating at less than full capacity. The JEM measures system and application utilization and power loads at the server, virtual machine, and application levels while dynamically allocating computing resources as they are needed. These actions can enable energy efficiency and/or demand response measures. The tool's over-time and real-time dashboards help with this process.

LBNL Data Center 50B-1275

LBNL's Data Center 50B-1275 serves both research and business needs at the laboratory. None of the work is classified as "secret," however; this does not mean that anyone can have access to information stored on the computer systems. Indeed, data security is a high priority at the laboratory.

There are two main types of computer systems in the center: research systems and business systems. Linux-based computational cluster systems are owned by different research groups, but all are managed by the laboratory's IT group. None of the clusters are connected to the outside world. Each cluster has a private subnet so that information cannot be accessed by people in other research groups. The clusters are controlled by a scheduling system, which automatically schedules the programs and number of processors to run. Naturally, this makes remote power management (turning down or powering off servers) difficult.

The business system has several subnets managed by different groups, each with its own firewall toward the outside world (the network). Each firewall has a different level of security, depending on the type of work performed. Most of the business systems are older and would have required by-passing the firewalls to access the operating systems. There is only one server for each of the services provided, making it difficult to manage the power remotely without affecting the service level. Therefore, Data Center 50B-1275 is not suitable for remote power management.

For systems that are suitable for remote power management, however, a general concern is what would happen if there were a disruption in the network or the network connection, since the JEM controls devices by using the network. Even without remote power management, however, manual management and the visibility of power consumption at the device level can still add value by providing information on power use.

Data Access and Security at LBNL

During the case study, LBNL's foremost data access and security concern was that no information should leave the data center. One of the laboratory's main goals was to demonstrate how the JEM could gather information without the need for access credentials. Its broad capabilities for monitoring energy consumption proved useful even though there were operational constraints. In terms of security, the dynamic method is the most risky and complex since the operating system must be accessed. However, to address these concerns, the JEM could be deployed and confined within the data center, so that the collected data also stayed local to the collection network.

The computational cluster system is not networked outside of the data center, so no information leaves the laboratory. For the business systems, the only way to provide access is to allow access to all subnets by going around the firewalls. One single system (the JEM) would then have access to the subnets and the operating system of each computer system. If any system were to be compromised, the impact could be quite significant for the laboratory. In this regard, the data center is like a collocation facility, which increases the complexity of the security issues.

Data center security is taken very seriously at most data centers, and this is addressed by one of the features of JEM release 3.0—a new software version released in early 2012. In that release, "collectors" sit in the data center behind the firewalls of segmented networks and report the energy back to a central point. This resolves the issue of traversing multiple firewalls and maintaining the security settings as each customer wishes. Indeed, for the case study, one of the laptops was moved behind the firewalls to alleviate the firewall issue.

The JEM can also monitor plug loads for any device by directly connecting to a power distribution unit (PDU), without accessing the operating system. This could resolve some of the security issues incurred at the laboratory. For a collocation type of data center, this could be an attractive option.

Finally, if the devices are instrumented, no credentials are necessary. However, the JouleX staff did not have remote monitoring access to data because of LBNL's security procedures. All data had to be captured by LBNL and sent manually to JouleX for analysis. A clearer overview of the different methods to retrieve power and energy data, and their pros and cons, would help the overall set-up process.

Installation of JEM

The laboratory secured two laptops: one for the cluster systems and one for the business systems. Since the cluster subnets are not routable, a separate laptop was needed. JouleX assisted LBNL in properly installing and configuring the software. Nevertheless, streamlining the installation process would be desirable—especially for the JEM's use in the DCEP program.

Lawrence Berkeley National Laboratory planned to run JEM under different operating conditions with different types of IT equipment, to study its capabilities and performance. It was hoped that the variety of computer systems would demonstrate the tool's monitoring capabilities and provide a good energy picture of both the individual devices and the grouped energy information. The original device monitoring scope was as follows:

- 2 enabled machines with Intelligent Platform Management Interface (IPMI) / Data Center Manageability (DCM)
- 2 non-enabled machines with intelligent power distribution units
- 2 virtual machines host with VMs running.
- 2 legacy machines - 1 Windows and 1 Linux using dynamic modeling
- 2 network switches
- 2 storage class devices
- 1 or more power meters

Monitoring at LBNL

Based on the observations and limitations outlined above, the emphasis was on JEM's power monitoring capabilities at the device level, rather than its power management capabilities.

- Two physical Windows servers and one physical Linux server were monitored by JEM using the dynamic method. These data were collected and manually exported to JouleX for analysis and feedback about the devices. JouleX analyzed the data and produced a report (see next section).

- Simple Network Management Protocol (SNMP) monitoring of one Cisco switch was also arranged. This is an Internet-standard protocol for managing devices on Internet Protocol (IP) networks. It is used mostly in network management systems to monitor network-attached devices. SNMP routes measured power from the device to JEM.

- Due to tight security at the laboratory, there were issues connecting JEM to servers with VMware. Since allowing access to all machines was not acceptable, the options would have been using the dynamic (OS) method on each device individually with local credentials or using the least-attractive static method using look-up tables.

- American Power Conversion (APC) power distribution units (PDUs) were attached to the JEM but were not configured due to an incorrect setup in the JouleX Energy Manager. This issue may have had to do with the SNMP protocol that was used, or the PDUs may not have been enabled to send energy data. JEM can also connect to PDUs through the Cisco EnergyWise protocol, as well as the Data Center Manager from Intel, but this approach was not part of the case study.

- LBNL also tried to connect JEM to a NetApp storage device, but this was not successful since the storage device does not automatically provide information that JEM could use. There is no standard format within storage devices for power analysis. Although JouleX planned to calculate the energy usage using their static method, it was not completed for this case study.

Data Analysis at LBNL

Monitoring energy consumption at the device level enables users to determine quickly where the energy is going. JEM provides a multitude of reports that can report on energy use, power demand, and utilization of the devices. In Table 1, for example, device "sabor.lbl.gov" resulted in the highest energy use and correspondingly the worst carbon dioxide (CO_2) emissions at the power plant. Each device has an associated geographic location, which is tied to the local energy price and how the energy is produced. The cost to operate the device, as well as the carbon footprint, can then be calculated.

Table 1: CO_2 Emissions Associated with Individual Devices

Device	CO2 Emissions [kg]
sabor.lbl.gov (131.243.60.165)	105.01
itprojectdev.lbl.gov (131.243.60.240)	97.53
localhost (127.0.0.1)	3.95
128.3.3.118	0.07
apc-pdu-15u.lbl.gov (131.243.60.128)	0.00
csprotest.lbl.gov (131.243.60.95)	0.00
apc-pdu-15t.lbl.gov (131.243.60.127)	0.00

Note: As noted, PDU data could not be measured, due to incorrect setup.

Another advantage to having energy collected at the device level is that the energy usage can be grouped in many different ways. For example, "pc.linux" consumed more energy than "pc.windows." The data can be selected and divided in any manner the user would like, device-by-device or group-by-group.

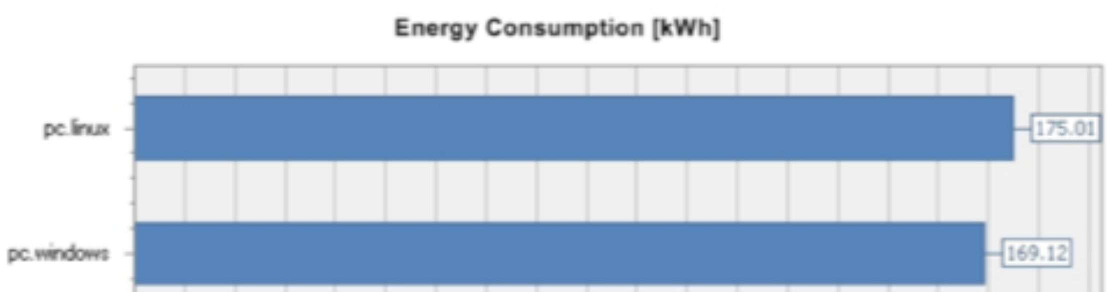

Figure 2: Energy Consumption for Groups

Utilization of each device is also a key performance indicator. Figure 3 shows the utilization of the three devices monitored. Table 1 shows that "sabor.lbl.gov" resulted in the largest carbon emissions and energy consumption, and "itprojectdev.lbl.gov" was a close second. Yet, as shown in Figure 3, the utilization of "itprojectdev.lbl.gov" is four times less than "sabor.lbl.gov." Thus, a significant part of the carbon emissions and energy costs could be cut by either virtualizing "itprojectdev.lbl.gov" or moving the applications to another server.

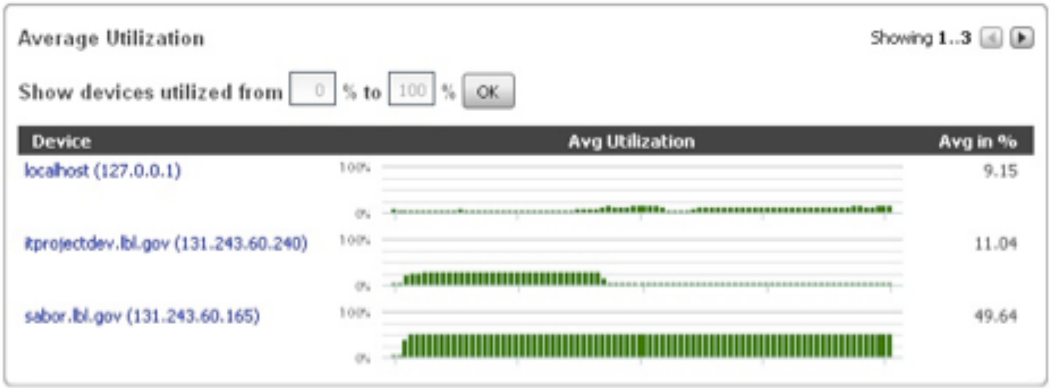

Figure 3: Utilization of Individual Devices

If one were to virtualize, say, 25 servers like "itprojectdev.lbl.com" and put them on a single computer, how much energy and CO_2 emissions could be saved? The chart below which was produced by the JEM tool shows the savings by moving the servers to a single platform.

Table 2: Reductions from Virtualization

Energy Consumption:	Month	Year	3 Years
Dell//PowerEdge 2650	4.64 MWh	56.48 MWh	169.43 MWh
Cisco//N20-B6625-2	522.85 kWh	6.36 MWh	19.08 MWh
» old vs. new	-4.12 MWh	-50.12 MWh	-150.35 MWh
Energy Cost ($0.12/KWh):	**Month**	**Year**	**3 Years**
Dell//PowerEdge 2650	$557.04	$6777.37	$20.33k
Cisco//N20-B6625-2	$62.74	$763.36	$2290.09
» old vs. new	-$494.30	-$6014.01	-$18.04k
CO2 Emissions:	**Month**	**Year**	**3 Years**
Dell//PowerEdge 2650	2.79 tons	33.89 tons	101.66 tons
Cisco//N20-B6625-2	313.71 kg	3.82 tons	11.45 tons
» old vs. new	-2.47 tons	-30.07 tons	-90.21 tons

In a situation where devices can be turned off on nights and/or weekends, JEM offers saving scenarios. The tool lets users run what-if scenarios, to see how much energy and money could be saved annually. For example, turning off the three monitored devices between 7 p.m. and 8 a.m. on weekdays, and completely on weekends, would result in the savings shown in Figure 4.

Figure 4: Reductions from Turning Devices Off at Night and On Weekends

DOE Experiences with JEM

The DOE Office of the Chief Information Officer also tested JEM, primarily to document the IT energy environment at their headquarters, identify areas of potential energy improvement, and evaluate whether sufficient energy savings could be realized to offset the cost of an active energy management system (in this case JEM). This was a short-term effort, on the order of a few days to get a high-level understanding of the IT operating energy use and to explore whether a tool such as the JEM tool could directly interface with the DC Pro profiling tool. In contrast to the LBNL case study (in monitoring mode in a data center), the installation was quick, and approximately 1,000 devices were monitored in one day. The general conclusion based on that effort's results was that JEM provided significant value and could interface directly with the DC Pro profiling tool.

It should be noted that security is a policy issue. DOE was able to get the necessary approvals quickly and allow JEM to reach all of their network segments, use their organizational active-directory system, and allow access to VMware consoles and other management systems. They were able to provision a host JouleX server within minutes. This approach may also have been feasible at LBNL had there been more time allocated for a detailed cyber security assessment of the JEM system.

Summary and Recommendations

This study of the JouleX Energy Manager (JEM) in data center 50B-1275 at LBNL demonstrated the advantage of monitoring energy use at the device level, so that data center operators and energy managers can better understand where and how much energy is being used in the facility. The JEM was able to generate many reports on energy use, power demand, and utilization of the devices. While the scale of this study was limited, JEM quantified and characterized the energy use of these devices, and this knowledge revealed opportunities to save energy and reduce associated carbon emissions through modified operation.

However, the effective use of the JEM at LBNL was impeded by three data access and security issues: (1) limited connectivity, (2) multiple private subnets, and (3) no remote monitoring access. Because of these issues, the originally planned power monitoring scope was significantly curtailed, and power data had to be manually exported to JouleX for analysis.

A solution to some of the security issues could include JEM's capability to monitor plug loads for any device by connecting to a PDU without the need to access the operating system. These issues are also addressed by new features of JEM v. 3.0 that have "collectors" that can reside in the data center behind the firewalls and report the energy consumption back to a central point, inside or outside the laboratory.

In a heavily firewalled and segmented environment such as at LBNL, or in collocation facility, dynamic power monitoring (accessing the operating system) may be difficult when strict security policies are in place. In such environments, JEM has numerous methods to acquire power data. Many of these methods do not require privileged accounts, but the additional requirements need to be planned well in advance with the security teams. In most enterprise data centers, dynamic power monitoring would be less of a concern, and many tools, such as patch management, already use this level of access. An alternative would be to use active power monitoring (interfacing with metered devices). The least desirable option would be to use static power monitoring (using look-up tables).

For this class of tools, a well thought-out deployment plan will be needed to address scope and security needs. A clearer overview of the different methods to retrieve power data and their pros and cons would help the overall set-up process. Moreover, streamlining the installation process to take advantage of the power monitoring feature would be desirable for use with the DCEP program. In this program, JEM could be used as an example of commercially available software, including its potential capability to collect equipment data automatically for the planned DC Pro IT tool.

Since the reduction of energy consumption and carbon emissions is a major goal for DOE data centers, tools like JEM could prove to be invaluable. A recommended next step could be to install such tools in other DOE data centers, to gain better control over those environments and to demonstrate a larger-scale implementation. This could also help to compare and contrast capabilities of various systems. Knowledge and use of these tools would help DCEPs identify energy-savings opportunities in those data centers by more effectively gathering the necessary data for the DC Pro assessment process.

Solutions such as the JEM hold a lot of promise and could play an important role in managing data centers, especially if the data transfer can satisfy strict security requirements. Inherently, solutions that require access to the IT equipment's operating systems over the Internet will be scrutinized, so interfacing with security requirements needs to be well planned and coordinated.