Data Center Optimization

Energy Efficiency

Implementing an IT energy management plan for real savings

February 2009

Contents

Executive Summary

Energy costs in the data center continue to rise, and will continue to do so if left unchecked. IBM® Systems Director Active Energy Manager™ is the cornerstone of the IBM energy management portfolio, empowering you to get control of energy costs. Whether you just want to measure and monitor energy usage, or leverage advanced functionality to optimize energy usage, Active Energy Manager is the tool to use. This white paper will help you understand the functionality included with Active Energy Manager and assist you in developing a plan to unleash the potential of an energy-optimized data center. Since all functions of Active Energy Manager can be tried for 60 days at no cost, you have nothing to lose. You can gain insight into optimizing the energy usage in your data center, along with the opportunity to achieve significant energy cost savings.

Introduction

Amidst the current economic downturn, the issues surrounding energy use in the data center continue to be a focus item for data center operations. The focus has changed slightly though, from IT capacity growth to reduced energy use and therefore cost savings. With scrutiny on cost savings, it is important to enable key players to see the return on investment (ROI) of energy-related solutions.

The ability to measure energy usage is the basis of developing the ROI for these energy-related solutions. You can use IBM Systems Director Active Energy Manager to measure, monitor, and manage energy consumption on IBM systems. Active Energy Manager can also monitor storage and other data center equipment through external sensors and partner integration. By providing trending views of the data, operators can easily see how the solutions and changes they implement affect power usage in the data center. By correlating the energy use with information on performance, transaction rates, and data access, you can understand not only the amount of energy use, but how the energy use relates to the real business of IT, and ultimately the overall business. Detailed correlations are enabled by Active Energy Manager and its ability to integrate with higher level managers, such as the products in the Tivoli® management suite.

The ability to measure energy usage is the basis of developing return on investment (ROI) for energy-related solutions.

Beyond the ability to measure and monitor energy-related information, Active Energy Manager also provides the capability to manage the energy use of many newer systems that have the ability to adjust power usage. The energy capabilities of systems vary by type and age of system. Active Energy Manager provides the ability to manage the energy attributes from a central point, and dynamically display the results of any changes that were made. Changes to these attributes can be performed manually, accomplished via scripts, or automated based on policies set by the operator. The trending views can provide the necessary information to set appropriate policies and derive optimal monitoring, threshold points, and action scripts.

So whether you are replacing systems with newer, more energy-efficient models, leveraging new technologies such as virtualization, or using more dynamic energy-saving capabilities of systems, Active Energy Manager is the right choice to measure, monitor, and manage energy usage. With Active Energy Manager you can get both baseline and real-time information to access the ROI of the actions you have taken.

Now that we understand the basis for using Active Energy Manager, let's take a more detailed look into the capabilities of the tool. The main tasks associated with energy management software are:

- Measure— Measure and record power and environmental statistics
- Monitor—Create and monitor thresholds on power and thermal values and events
- Manage—Control system power states and capping functions
- Optimize—Leverage the ability to monitor and manage via policy-based automation

Measuring power consumption throughout the data center

Active Energy Manager helps you understand areas of inefficiency by measuring the power of individual pieces of equipment in the data center.

Many companies know how much total power is used by their data center, however, without knowing where the power is going, it's almost impossible to build a plan to solve areas of inefficiency. To begin to understand areas of inefficiency, you must measure the power of the individual pieces of equipment in the data center. Active Energy Manager is designed to accomplish this measurement.

By communicating directly to firmware running on IBM servers, Active Energy Manager can poll systems at a regular time interval (default is five minutes, but can be as often as one minute) to measure how much power they are consuming. This can be done for all of the IBM platforms, including System x®, BladeCenter®, System z®, and the Power Systems™ servers. It's quick and easy to set up monitoring of many servers, because no agent needs to be installed on these IBM systems. Once a server is discovered by IBM Systems Director, Active Energy Manager automatically starts measuring the energy information for that system. Also, since Active Energy Manager uses a direct connection to the firmware, systems can be monitored when they are idle, and when the operating systems are not operational.

In order to truly get a view of the energy consumption of the entire data center, *all* of its IT equipment needs to be measured. Although newer IBM servers can be monitored through a firmware connection, older IBM servers, as well as non-IBM servers, require different approaches. Active Energy Manager can measure these systems through the following methods:

- Measuring the power at the outlet of a rack-mounted or floor-standing power distribution unit (PDU).
- Communicating to an external sensor.
- Communicating to a facility management application, which monitors power consumed through PDUs.

If you have multiple data center locations, the energy cost calculator could help you determine the most efficient location to run a particular workload.

When you have a rack with a PDU that is capable of outlet level measurement, Active Energy Manager can measure any equipment that plugs into the rack, such as storage and networking equipment. Rack-mounted PDUs are discovered and listed as Power Units in IBM Systems Director. In most cases, the power units are SNMP devices, and Active Energy Manager does its measuring through that SNMP interface. Active Energy Manager also includes a configuration user interface that allows you to associate a device with a PDU outlet. You can simply type in the name of the device on the user interface, or you can view a list of managed devices, and choose the one that is plugged into a specific outlet. Similar functionality is also available when you associate an external sensor with a device. Active Energy Manager Version 4.1 supports IBM DPI PDU+, Eaton PDUs, Raritan PDUs, and SynapSense wireless sensors.

The use of a facility management application is another method that can be used to measure energy consumption. In Active Energy Manager 4.1, a connection to Emerson Liebert's SiteScan application can be used to discover power and cooling equipment in a data center. SiteScan is able to monitor Emerson Liebert equipment, as well as that from third parties. Since PDUs can also be measured with a facility management application, Active Energy Manager has additional sources of data for measuring IT equipment.

With all of the energy data being collected, Active Energy Manager allows you to view and analyze power and temperature information over the past hour, the past 12 hours, and all the way up to the past year. You can view the data in a graphical format, or review individual readings in a table view. By viewing power and temperature data in graphical form, you can pinpoint energy peaks for a certain piece of equipment. For example, you may find out that there is a spike in power consumption when a batch job is running. You may also find times each week when there is very little CPU activity; these intervals represent potential opportunities to save energy. In addition, by viewing trends of the aggregate power of groups of servers, or of the entire data center, you can see how power usage is changing over time.

By tracking these trends, you can help predict when a data center will run out of capacity, and when it is time to add on.

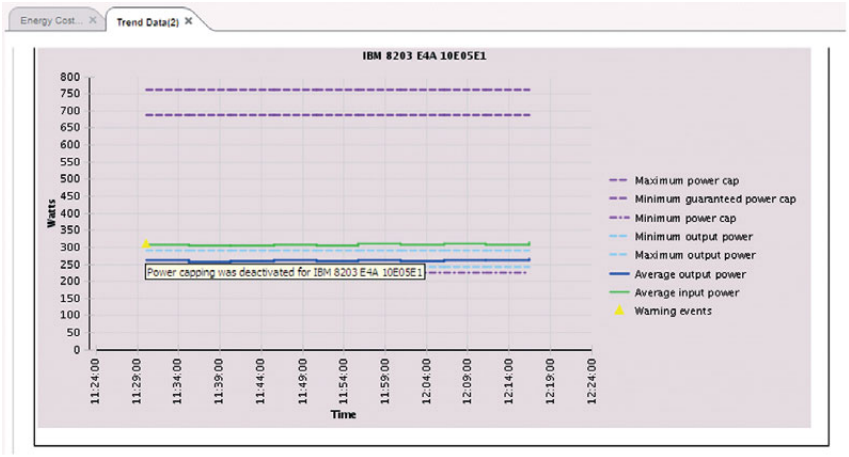

Figure 1: Trend data for approximately 45 minutes in graphical format

Another method of viewing power usage is through the energy cost calculator function in Active Energy Manager. You can enter the rate that you are being charged for energy usage (cost per kilowatt hour) and a cooling rate factor (defaulted to 1.5). By using these factors, as well as the power being monitored for a server, Active Energy Manager calculates how much it has cost to power a server for the specified length of time. If a client has multiple data center locations, and pays a different price per kilowatt hour at the various locations, this information could be useful in determining the most efficient location and server on which to run a particular workload.

Monitoring power and thermal measurements with thresholds

Active Energy Manager allows you to monitor the power and environmental measurements in the data center. By integrating Active Energy Manager into the threshold support in IBM Systems Director, you can now be notified whenever any measured metric reaches a user-specified value. You can monitor the following metrics:

- Amperage capacity of a PDU
- Exhaust temperature
- Average input power (AC)
- Average output power (DC)
- Effective CPU speed
- Ambient temperature
- Humidity

For any of these metrics, both a *warning* level and a *critical* level can be specified. For example, if you want to be warned when a server has an input temperature higher than 75 degrees Fahrenheit, and informed of a critical problem when it goes above 80 degrees, you can select *Ambient temperature*, and then fill in 75 and 80 in the appropriate fields on the screen. Monitoring support is key to the operational management of the power and thermal attributes of the data center. Later in this article, you will see how monitoring is used with the management support to optimize energy use in the data center.

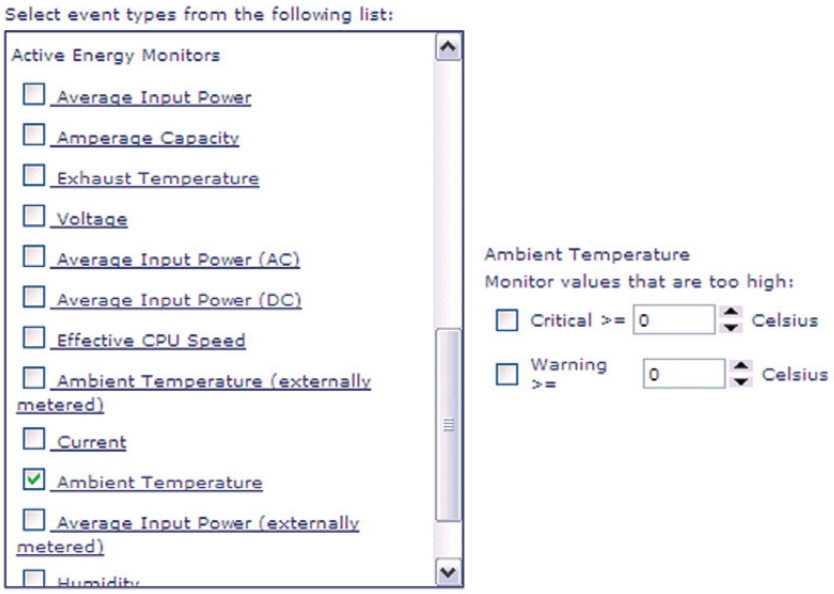

Figure 2: Setting a threshold for ambient temperature

Managing power usage with power modes and power caps

On Power Systems servers, you can save up to 30% of the system's energy usage with ***static low power mode***.

Once you have a good understanding of where the power is being consumed within a data center, the next step is to perform power management. With Active Energy Manager, this means using power savings modes and power capping. These features give you the opportunity to put servers in low power modes when they are not being used as heavily, and to allocate less power for servers when appropriate. To effectively use these power management modes, it is a good idea to monitor the power usage of the servers to get a baseline usage. This baseline will help determine when it's best to use power management functions.

The management function that can save the most energy for Power Systems servers is called static low power mode. With this mode, the system reduces the clock speed and voltage of the processors, and can save up to 30% of the system's energy usage. On a system that is being heavily used, static low power mode could result in reduced performance, so it is recommended that static low power mode be used only during times of lower utilization.

One method to utilize static low power mode is to use the built-in scheduler in IBM Systems Director. For example, you could schedule the system to go into static low power mode at 6:00 p.m. after people have left work for the day, and then have it turned off at 8:00 a.m. before people and workload return. Similarly, this type of power throttling could be exercised on weekends. By using monitoring data from the past 30 days, you can determine the best times to use static low power mode. Another option is to write a script that monitors CPU utilization for a system, and have static low power mode turned on when utilization falls below a specified level (and turned off when utilization rises). An example of such a script is included in the appendix of the IBM EnergyScale™ for POWER6™ Processor-Based Systems white paper.

Another function for saving energy is called dynamic power savings mode, and is also available on Power Systems servers. When dynamic power savings mode is enabled, the firmware of the system continuously monitors the utilization of the system, and adjusts the CPU clock speed and voltage to provide enough power to run the current workload. The less the system is utilized, the more power savings

Many clients allocate power and cooling in their data centers based on nameplate ratings, which can be very inefficient.

are achieved. In addition, you can specify whether you want to favor performance or favor power when enabling dynamic power savings mode. With favor performance, the peak frequency of the processors may be greater than 100%. With favor power, the processors are limited to 95% of nominal frequency under full utilization.

An important point to note is that with both of the power savings options, static low power mode and dynamic power savings mode, it is possible to save significant energy without noticing any change in workload performance. You will be most successful when you use the monitored data and the historical trends that the data reveals to develop the proper configuration and policy settings that enable these functions.

The other main energy management function in Active Energy Manager is power capping. Power capping allows you to allocate less power to your systems than is typically done by today's means. All systems have a nameplate power rating, which is theoretically the maximum amount of power that could possibly be consumed by that system. However, in reality, most systems won't use more than about half of the power that is stated on the nameplate value because of configuration or usage differences. Many clients will allocate power and cooling in their data centers based on these nameplate ratings, but this can be very inefficient. The result is more power infrastructure and more cooling than what is required to support the systems. Some clients purchase external power meters to measure the actual power usage of systems, and then use those measurements to determine the amount of power and cooling needed for a system. However, this requires buying extra equipment, and it may not accurately determine the worst case scenarios, depending on the conditions in which those measurements were taken.

The Active Energy Manager power capping functionality can accurately allow you to set a cap on the amount of power that a system can consume, and allow you to allocate this smaller amount of power to the server. If the power consumption of a server reaches the user-configured cap, the firmware of the server will reduce the clock speed and voltage of the processors in order to stop the power consumption from exceeding the cap. As an example, let's say you have eight servers in a rack that each has a nameplate power rating of 500 watts. Therefore, you allocated 4000 watts to that rack. Later, you find out via power monitoring that none of these servers are using more than 250 watts each. To give you some breathing room, you decide to set a power cap of 300 watts on each server. Thus, the maximum power that the rack can consume is 2400 watts (8 x 300). Since you already have allocated enough power for 4000 watts, you have 1600 watts to spare, and can now potentially add five more servers (at 300 watts each) into that data center. This example illustrates how power capping can be implemented to add more compute power into a data center without having to add any more power equipment. Not only can power capping help you delay the costs of adding on to your data centers, it can help you grow your business even if the utility company is unable to provide additional power to your data center.

Group capping is a new power capping function that is now available in Active Energy Manager 4.1. Input from some clients has been that they manage their power at a rack level, not at the individual server level. The Active Energy Manager group capping function meets this requirement by allowing a user to create an IBM Systems Director group made up of any combination of servers that support the power capping function. This could include both System x and Power Systems servers. You might choose to group all of the servers in a specific rack, all the

servers in a group of racks, or any set of systems that you wish. You define a cap for the entire group, and then Active Energy Manager sets the caps of the individual servers in order to meet the overall group cap. If the number of servers in the group changes, Active Energy Manager can adjust the individual server caps automatically to compensate.

For a group capping example, you may choose to set the group cap for a specific rack of 10 servers at 3000 watts. If all 10 servers are the same, Active Energy Manager would likely set an initial cap of 300 watts (3000 / 10) for each of the servers. If the 10 servers are a variety of different types, some with a higher valid power cap range, and some with a lower valid power range, Active Energy Manager will determine the best mix of caps to meet that goal. If you attempt to set a group cap to a value that is not possible to meet, based on the valid power cap ranges of the individual systems in that group, then Active Energy Manager will surface an error, and the group cap cannot be set.

Another new capping function, soft power capping, is also available in Active Energy Manager 4.1. Soft power capping allows a power cap to be set to a lower value for a server than the previous power capping function could accomplish. Consequently, a soft power cap should not be used for allocating the power for a server. However, soft power capping can be used to potentially save server power, since the server will attempt to stay below that power cap if at all possible (by slowing the processor clock speed and lowering the voltage). If the soft power cap cannot be maintained, then an alert will be posted, and you can choose to take action if you desire. Soft power capping function is currently only available on POWER6 processor-based blades and the following POWER6 processor-based servers: IBM Power 550 and IBM Power 520 Express.

You can create an automation plan to automatically place a system in power savings mode when it is consuming more than a specified number of watts.

Optimizing energy use through automation

In order to effectively optimize energy use in the data center, IBM Systems Director Active Energy Manager combines measurement, monitoring, and management capabilities to provide manual, event-driven, and automatic optimization to the data center.

Because you might like to have some action taken automatically when a threshold is reached, Active Energy Manager now includes an automation wizard that makes it easy to set up automation. On the Active Energy Manager main page, click the Create Automation Plans link in the Automate section to start the wizard. First, the wizard prompts you to select the threshold that you want to set. Then, it asks what action you would like to have taken if that threshold is reached. You can select from various options, including the following:

- Start a program on a system
- Send an e-mail to a mobile phone
- Start a task on a system
- Log to a log file
- Add or remove the event-generating system to/from a group

One possible use of an automation plan could be to automatically place a system in power savings mode when it is detected that the system is consuming more than a specified number of watts. For systems such as the IBM Power 595 and the IBM Power 575 that support power savings, but not power capping, this could be a way of implementing a limited capping function.

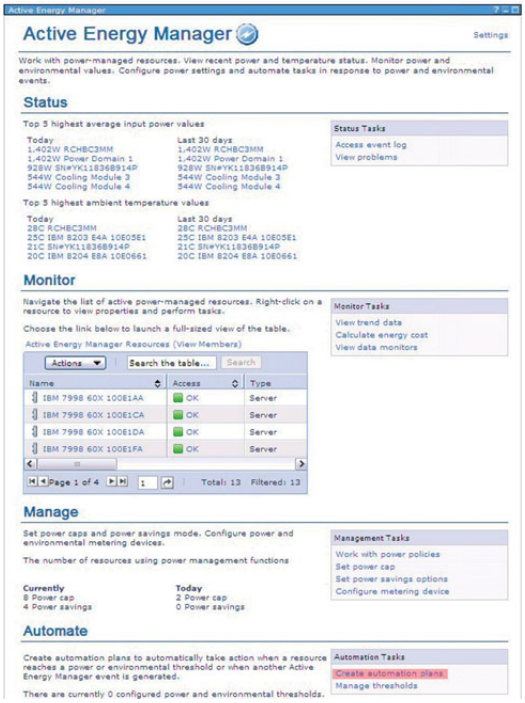

Figure 3: Active Energy Manager main page, with Create automation plans task highlighted

Integrating Active Energy Manager with Tivoli

Active Energy Manager is included as part of the IBM Tivoli Monitoring (ITM) for Energy Management offering, providing broad integration with the Tivoli product set.

Active Energy Manager integrates with IBM Tivoli Monitoring (ITM) for Energy Management to provide energy management at the enterprise scale. Active Energy Manager is included as part of the ITM for Energy Management offering, providing broad integration with the Tivoli product set. ITM for Energy Management interacts with Active Energy Manager to access data and alert information, and allows ITM for Energy Management to control the energy adjustments in the data center. A single instance of ITM for Energy Management can integrate with multiple instances of Active Energy Manager.

ITM for Energy Management provides integration with a range of facilities providers such as Johnson Controls Incorporated (JCI), APC, Eaton, and SynapSense. Through the integration with Active Energy Manager, ITM for Energy Management can also integrate with additional partners such as Emerson Liebert.

Active Energy Manager data can be integrated with a broad range of monitoring information for IT and facilities resources, as well as transaction performance data. These integration capabilities provide system administrators with an unprecedented view of power, temperature, and application performance in their data centers. The power and thermal data can also be placed into the ITM-integrated Tivoli Data Warehouse, enabling historical comparisons, analytics, and in-depth reporting. The data is also available to other Tivoli management products such as Tivoli Usage and Accounting Manager and Tivoli Business Service Manager. Thus, Active Energy Manager and its integration with ITM for Energy Management can serve as a key building block for more comprehensive power-management solutions.

Moving forward with an energy management plan

The following example shows the steps that you can use to evaluate your data center and deploy energy savings solutions.

Install and configure

1. Look at the list of equipment in your data center, and see how it compares to the list of equipment that is supported by Active Energy Manager. See the IBM Systems Director Active Energy Manager V4.1 information center for details.

2. If there is IT equipment in the data center that Active Energy Manager does not support (such as older servers and non-IBM servers), determine if those could be monitored through the use of intelligent PDUs, Synapsense sensors, or the Emerson Liebert SiteScan integration.

3. Once you determine the scope of the equipment that you want to monitor and manage, install IBM Systems Director 6.1 onto one of the servers in your data centers. Then, install Active Energy Manager 4.1 as a plug-in onto that same server.

4. Install any IBM Systems Director and/or Active Energy Manager updates that may be available.

5. Use IBM Systems Director to discover the equipment that you want to monitor and measure for energy use. The IBM Systems Information Center has information on how to discover specific types of equipment.

6. Launch Active Energy Manager from the IBM Systems Director navigation area by selecting it from the Energy category. Active Energy Manager will then start polling any equipment that is "energy manageable."

7. Adjust the polling intervals if desired. By default, Active Energy Manager polls each device once every five minutes. If you desire more granularity, you can set the polling interval to as often as once per minute. In addition, if you want to poll most servers every five minutes, but poll the most critical servers once every minute, you can configure that, too.

8. If you want to monitor at a group or rack level, instead of at the individual server level, create IBM Systems Director groups.

Measure, monitor, and manage

1. Decide how long you want to monitor the servers in order to get a good sampling of power consumption and temperature over an appropriate sampling period. A one-month sample may be a good start. However, if some months are expected to draw significantly more power than others, you may want to continue monitoring until after one of those heavier months. At the end of this time, export the data so that you can use it as a baseline against which to perform evaluations.

2. Evaluate the power trending curves for representative servers in the data center. Look for trends where there is a lower usage of power, and where spikes in power typically exist. For example, are there times of low usage during nights and weekends that might signal a good time to turn on power savings mode?

3. In the power trending curves, look for places where the power usage is significantly lower than the nameplate power values, and determine if power capping would be useful to allow you to allocate less power to those servers. Look at the valid power capping range for those servers so that you can determine where you could potentially cap those servers.

You can create thresholds and automation so that you can be alerted of an overheating situation immediately.

Optimize

1. If you have concerns about portions of the data center overheating during a cooling system failure, set up some thresholds and automation so that you can be alerted of an overheating situation immediately, and take the appropriate actions.

2. After setting up the power saving, capping, and automations that you think will save you the most energy, continue to run the data center for another month. Then, compare the results with the baseline that you captured previously. You might find it interesting to use the energy cost calculator to see how much money has been saved during the month. You might also find that you've been able to run more workload, but within the same power envelope.

3. Based on the results of the comparison, determine if additional actions might be appropriate. Are more licenses for power capping and power savings required? Could you do a better evaluation if you add more PDUs or sensors to cover a larger portion of the data center? Could you save more energy by buying a new server, and virtualizing all the workload from multiple servers onto that one?

Summary

Active Energy Manager is part of a growing portfolio of IBM products and services to help manage energy use in the data center. Through the use of the monitoring, power saving, and power capping modes described in this paper, you can better optimize your data centers for energy efficiency. Active Energy Manager is capable of meeting your energy management goals at multiple levels—it can monitor individual pieces of IT equipment, groups of equipment, or it can capture a view of your entire data center. It also integrates with Tivoli and partner products for enterprise scale solutions. And, since all functions of Active Energy Manager can be tried for 60 days at no cost, you have nothing to lose, and lots of valuable insight to gain by trying it.

For more information

To learn more about the IBM solutions and offerings discussed in this paper, please contact your IBM marketing representative or IBM Business Partner, or visit the following Web sites:

IBM Systems Director resources

- IBM Systems Director (Website)

IBM Systems Director Active Energy Manager resources

- Active Energy Manager (Website)

IBM Systems Workload Estimator

- IBM Systems Workload Estimator (Website)


Appendix A. Resources

IBM Systems Director resources

- IBM Systems Director (Website)
- IBM Systems Director information center (Website)
- IBM Systems Director Planning, Installation, and Configuration Guide for AIX® (PDF)
- IBM Systems Director Planning, Installation, and Configuration Guide for IBM i (PDF)
- IBM Systems Director Planning, Installation, and Configuration Guide for Linux on Power Systems (PDF)
- IBM Systems Director Planning, Installation, and Configuration Guide for Linux on x86 (PDF)
- IBM Systems Director Planning, Installation, and Configuration Guide for Linux on System z (PDF)
- IBM Systems Director Planning, Installation, and Configuration Guide for Windows (PDF)

IBM Systems Director Active Energy Manager resources

- Active Energy Manager (Website)
- IBM Systems Director Active Energy Manager V4.1 information center (Website)
- IBM EnergyScale for POWER6 Processor-Based Systems (white paper)

IBM Systems Workload Estimator

- IBM Systems Workload Estimator (Website)

Systems and Technology Group
Route 100
Somers, NY 10589

Produced in the United States of America
February 2009

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non IBM products should be addressed with those suppliers.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

OIW03024-USEN-01